Skill Gaps over the Technology Life Cycle

Ashwini K. Agrawal Daniel Kim Prasanna Tambe*

October 5, 2026

Abstract

Recent technological progress has largely outpaced workers' ability to adapt their skills. Prior work shows that this uneven race between technological change and human capital accumulation generates skill gaps that impede new technology adoption and productivity growth. What is less understood, however, is how this dynamic shapes the skill gaps faced by firms across all stages of the technology life cycle. This paper empirically characterizes the evolution of skill gaps as information technologies diffuse, using firm-level measures constructed by matching job postings to employee resumes. We first document that skill gaps follow a U-shaped pattern: high at the onset of adoption, then declining as firms and workers adjust, but rising again as technologies approach obsolescence. Second, skill gaps reflect shortages in combinations of technical and non-technical skills rather than technical deficits alone. Third, firms with greater organizational capital face systematically smaller gaps. Lastly, firm productivity is inversely related to skill gaps and hump-shaped over the technology life cycle. The findings are consistent with a vintage human capital framework in which skill gaps arise endogenously as a result of firms' and workers' dynamic incentives. The analysis shows how the uneven pace between technological progress and human capital accumulation generates skill gaps and productivity dispersion not only at the technology frontier, but across all stages of the technology life cycle.

*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 Tarek Hassan, Thomaz Teodorovicz, Benjamin Labaschin, Peter Haslag, and seminar audiences at the LSE, McMaster University, Revelio Labs, University of Washington, University of Georgia, the University of Waterloo, Cambridge University, the University of Houston, the SFS Cavalcade, and the 2026 Academy of Management. 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

Recent technological progress has largely outpaced workers' ability to adapt their skills (Goldin and Katz, 2008). This dynamic is particularly salient in light of modern advances in artificial intelligence, data analytics, and automation. Prior research shows that this uneven race between technological change and human capital accumulation generates skill gaps that impede new technology adoption and productivity growth (Adão et al., 2024; Caselli and Coleman, 2006; Acemoglu and Zilibotti, 2001). What remains poorly understood, however, is how this race affects skill gaps across all stages of the technology life cycle—including for firms reliant on older technologies that remain central to production (Cirera et al., 2026).

Characterizing skill gaps across the technology life cycle is theoretically and empirically challenging. Theoretically, it is unclear whether skill gaps decline monotonically with technology age as firms and workers adjust, or whether they reappear as firms and workers face diminishing incentives to invest in skills tied to aging systems. Prior work typically takes skill gaps as exogenous, abstracting from the dynamic incentives of firms and workers to invest in skill formation. This approach limits our ability to distinguish whether observed skill gaps impede technology adoption or reflect a consequence of it.

Empirically, it is challenging to measure skill gaps and identify the dimensions along which firms and workers are misaligned. Many relevant skills are unobservable and difficult to infer from standard data sources. For instance, a job requirement for "Java programming experience" may reflect latent demand for narrow technical expertise, broad problem-solving ability, and social skills needed for collaboration. This problem is compounded by the heterogeneity and inconsistency of skill labels found in practice. Prior approaches mostly rely on one-dimensional proxies that obscure the multiple dimensions of human capital relevant for technology utilization (Deming, 2023; Brunello and Wruuck, 2021).

This paper presents new stylized facts to empirically characterize the evolution of skill gaps and firm productivity over the information technology life cycle. We measure skill gaps

at the firm-level by creating a matched employer–employee dataset that links U.S. firms' job postings with workers' online résumés. Job postings provide detailed information on firm skill requirements, while résumés present the skills supplied by employees hired within close temporal proximity to posted job ads. We quantify the distance between firm requirements and the skills of new hires across multiple dimensions, including managerial, technical, and financial competencies. Job postings also contain detailed information on firms' adoption of specific information technologies, such as data infrastructure systems, enterprise management software, and programming tools. We identify the year of commercial introduction for each technology in our sample and use this variation to characterize the age of firm technologies and production systems over time, following Kalyani et al. (2025) and Tambe (2014).

The key innovation in our construction of skill gaps is the conversion of raw skills data into skill weight vectors (Lazear, 2009) using standard classification methods (Dell, 2025) and large language models (Ludwig et al., 2026). Our approach builds on recent work that uses these tools to extract skill content from unstructured text (Hassan et al., 2026, 2025; Autor et al., 2024; Javadian Sabet et al., 2024). These tools help address two key measurement challenges. First, job postings and résumés contain thousands of heterogeneous and inconsistently defined skill labels, making it difficult to manually construct comparable measures of skill demand and supply without imposing strong assumptions about how observed skill labels map into specific dimensions of human capital. Second, individual skill labels do not readily reflect the underlying capabilities they represent. For example, a requirement for "Java programming" reflects a latent set of technical and analytical competencies that are not directly observable from the label itself.

Large language models (LLM's) are well suited to this problem because they are trained on a vast text corpora to represent words and phrases as vectors (i.e. "embeddings") that capture how firms and workers interpret these skills in practice. For example, LLM's can translate a skill requirement of "Java programming" into a skill vector that quantifies the relative importance of underlying capabilities such as technical expertise, social skills, and

managerial ability. This approach enables us to harmonize raw skill data into a common representation of skill supply and demand that is comparable across firms and workers.

We measure skill gaps as the Euclidean distance between the average skill vector of a firm's new hires and the average skill vector of a firm's corresponding job requirements. Euclidean distance is widely used in the literature (see for example, Bandiera et al. (2025)), because it provides a transparent and intuitive summary statistic of the alignment between worker skills and firm requirements: larger values indicate greater divergence. Because skill weights sum to one, this measure captures the degree to which workers' comparative advantages across skills align with firms' requirements. We use this measure to characterize how skill gaps evolve over the technology life cycle. Specifically, we ask: how do skill gaps evolve as information technologies diffuse, what are the specific skill dimensions and occupations that drive these patterns, how do firm investments in organizational capital impact skill gaps, and how do skill gaps mapin to firm productivity over the technology life cycle?

To interpret our empirical findings, we extend the technology diffusion and vintage human capital framework of Chari and Hopenhayn (1991). Their theory describes how the stock of workers skilled in a given technology rises and falls over its life cycle. We adapt this model to consider how skill gaps endogenously evolve when technological change outpaces human capital accumulation. Although their framework does not explicitly consider skill gaps, we introduce them by microfounding a key assumption in the model—new employee productivity increases with the stock of workers who are skilled in the firm's technology. In our extension, this relationship operates through skill gaps: a larger stock of skilled workers reduces the cost of aligning the job requirements of new employees with the capabilities of available labor, leading to smaller skill gaps and thus higher productivity. For example, workers who become skilled in a technology can perform managerial tasks that are no longer required of new employees. This microfoundation preserves the stationary equilibrium structure of the original model, while providing a tractable framework for interpreting skill gap patterns.

We first document that skill gaps follow a U-shaped progression over the technology life

cycle. Gaps are high when new technologies are adopted, decline as technologies mature, and rise again once technologies approach obsolescence. This process unfolds over decades, consistent with previous evidence on the speed of technology diffusion (Comin and Hobijn, 2010; Kalyani et al., 2025). The pattern is robust to aggregating technologies at the firm-level and to disaggregating across major classes of information technologies, including enterprise systems, data infrastructure tools, and programming languages.

The U-shaped evolution of skill gaps over the technology life cycle arises naturally in our framework. In the model, skill gaps are initially large because workers lack the experience and specific capabilities required by frontier technologies — a consequence of technological progress outpacing worker skill formation. Over time, as technologies diffuse and workers accumulate relevant human capital, firms face lower adjustment costs in aligning job requirements with the capabilities of new hires, leading to smaller skill gaps and higher productivity. However, as newer and more productive technologies emerge, workers' incentives to join firms operating legacy systems begin to diminish, causing skill gaps at these firms to widen again. This general equilibrium mechanism and supporting empirical evidence illustrate how the uneven pace between technological change and human capital accumulation impacts skill gaps not only at the frontier, but across all stages of the technology life cycle.

Our second set of findings decomposes the life-cycle pattern to identify the underlying drivers of skill gaps over the course of technology diffusion. The data challenge the conventional view that IT-driven skill gaps reflect technical deficits alone. Instead, we observe shortages in *combinations* of technical and non-technical skills that jointly enable IT adoption. The largest shortfalls are concentrated in IT-complementary non-technical capabilities such as managerial and coordination skills, alongside technical shortages in areas such as network security and advanced manufacturing. Moreover, the U-shaped pattern appears not only in technical occupations such as engineers and data scientists, but in non-technical occupations such as managers and business analysts who are central to IT-complementary organizational change (Bresnahan et al., 2002; Bartel et al., 2007). Taken together, these

results highlight that the adjustment process between technological progress and human capital accumulation is inherently multidimensional throughout the life cycle.

Our third set of findings identifies a key factor that shapes skill gaps over the life cycle: organizational capital. Firms with greater organizational capital have systematically smaller gaps throughout the technology life cycle. In our framework, organizational capital — which includes management systems, training infrastructure, and workplace design — lowers the cost of aligning job requirements with workers' capabilities, thereby leading to smaller skill gaps and increased worker productivity. The findings extend prior work showing that organizational capacity complements IT adoption and enhances productivity (Brynjolfsson et al., 2021; Bloom et al., 2012; Brynjolfsson et al., 2002; Atkeson and Kehoe, 2005), by suggesting that these investments operate in part by mitigating skill gaps.

Our fourth set of findings maps skill gaps to firm-level productivity over the technology life cycle. We show that productivity moves inversely with skill gaps: total factor productivity (TFP) follows a hump-shaped pattern over the life cycle, and regression estimates indicate that firms with larger skill gaps realize significantly lower TFP. These patterns are consistent with recent work on the J-curve dynamics of productivity for general purpose technologies (Brynjolfsson et al., 2021). Closing skill gaps early in the life cycle requires firms to invest in intangible assets that take time to show measurable productivity effects. At later stages, however, widening skill gaps are associated with lower productivity, as workers have reduced incentives to invest in legacy systems that are less productive than frontier technologies.

We conduct a range of analyses to assess the robustness of our empirical findings. One concern is the sensitivity of the results to the construction of skill gaps. Our baseline measures of skill gaps approximate the alignment between firm skill requirements and worker capabilities using matches based on firm, occupation, and timing. Because we do not observe whether a given hire fills a specific vacancy represented by a job posting, this approach is subject to sampling bias and measurement error——for example, some hires may not correspond to observed job postings, while some postings may remain unfilled (Dalton et al., 2026). To

address these concerns, we construct a variety of alternative skill gap measures. We vary the cross-sectional and time-series criteria used to match workers to job postings, examine both aggregated and disaggregated measures of skill gaps, and reweight observations using industry–occupation data from the Current Population Survey (CPS). Across all specifications, the results are highly stable, indicating that the empirical patterns we document are not driven by the specific manner in which skill gaps are constructed.

Another important concern is the measurement of workers' skill profiles, which are constructed from the skills listed on workers' LinkedIn profiles. The timing for when these skills are reported on the platform is unknown, therefore, these data may include skills acquired after workers begin the jobs matched in our sample (Bloom et al., 2026). This source of measurement error likely attenuates our skill gap measures. Nevertheless, to assess this issue, we re-estimate our results using only observations from the final year of the sample, when the elapsed time between workers' job transitions and observed skills is minimized. The patterns we document persist in this restricted sample, and remain stable as we sequentially include earlier years with longer potential exposure to post-transition skill updates.

We conduct a number of additional tests to further assess the robustness of our findings. Our results hold across alternative definitions of technology vintage, including measures constructed at the firm and firm–occupation levels, as well as across different categories of information technologies. They are also robust to alternative skill taxonomies developed in prior work, including those proposed by Deming and Kahn (2018). In addition, we obtain similar results using alternative distance metrics, such as cosine similarity and Hellinger distance, in place of Euclidean distance. Finally, we replicate our analysis using more advanced large language models, such as ModernBERT, to construct firm and worker skill vectors. The results are even stronger under these more computationally intensive approaches, suggesting that our baseline estimates are conservative.

The main contribution of this paper is an empirical characterization of skill gaps and firm productivity over the information technology life cycle. Interpreted through an extended vin-

tage human capital framework based on Chari and Hopenhayn (1991), the findings show how the uneven race between technological progress and human capital accumulation generates endogenous skill gaps and productivity dispersion not only at the technology frontier, but across all stages of the technology life cycle. Standard frameworks that take skill gaps as exogenous and focus on frontier adoption alone cannot account for our findings.

The findings carry two policy implications. First, workforce development programs should account for the endogenous and multidimensional nature of skill gaps across the entire technology life cycle, not just at the frontier of adoption (Jonas et al., 2026; Dillon et al., 2025; Katz et al., 2022; Card et al., 2018). Policies targeting skill shortages without recognizing these dynamics risk misallocating resources toward gaps that naturally dissipate as technologies mature, while neglecting the shortfalls that accumulate as technologies age into legacy systems. Second, policies promoting new technology adoption may have unintended effects on skill formation incentives across the technology life cycle. Subsidizing frontier technologies such as generative AI may raise productivity in the short run (Agrawal et al., 2019), but may simultaneously reduce workers' incentives to invest in skills specific to established technologies — potentially amplifying, rather than mitigating, productivity dispersion.

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

2 Related Literature

The relationship between technological progress and human capital accumulation is a central determinant of wages, employment, and productivity. Goldin and Katz (2008) study how differences in the pace of technological change and educational attainment over the past century have shaped wage inequality. A large literature on skill-biased technical change

shows that technological progress often raises the relative demand for skilled labor while adjustments in labor supply lag behind (Katz and Murphy, 1992; Autor et al., 1998, 2003; Michaels et al., 2014). More recently, advances in artificial intelligence and automation have sharpened this tension, raising questions about the pace of task displacement versus the speed of new job creation and associated skill acquisition (Acemoglu and Restrepo, 2018; Brynjolfsson and McAfee, 2011, 2014; Autor, 2015; Hassan et al., 2026).

The dynamics of technological change and skill formation also play a central role in technology diffusion and productivity. A growing body of work shows that gaps between the skill requirements of new technologies and the available workforce can impede new technology adoption and suppress productivity (Angelini, 2025; Acemoglu and Restrepo, 2019; Caselli and Coleman, 2006; Nelson and Phelps, 1966). Acemoglu and Zilibotti (2001) show that cross-country productivity differences arise from mismatches between the skill requirements of frontier technologies and the human capital of local labor. Adão et al. (2024) further demonstrate that the speed of new technological adoption depends on the size of skill gaps between workers and new technology requirements. Together with practitioner surveys documenting skill shortages as a primary barrier to adoption (MIT Technology Review Insights, 2023; Brunello and Wruuck, 2021), this literature provides compelling evidence that skill gaps are a key determinant of adoption and productivity dispersion at the technology frontier.

What remains less understood, however, is how these forces affect firms reliant on older technologies that operate away from the frontier. Despite the continued importance of such technologies (Cirera et al., 2026), there is limited theoretical and empirical work characterizing how skill gaps evolve over different stages of the technology life cycle. On the theoretical front, prior work typically treats skill gaps as exogenous and focus on their consequences at the frontier. In practice, however, firms using older technologies such as legacy systems often report hiring difficulties comparable to those of frontier firms (Monaghan and Bass, 2020). Without a framework for understanding the forces that shape skill gaps, it is difficult to interpret whether observed gaps primarily reflect binding constraints on technology

adoption alone or instead reflect the endogenous outcomes of firms' and workers' decisions to invest in technology-specific skills. This ambiguity limits our understanding of how skill gaps contribute to productivity differences across firms, an issue of great importance among managers and policy makers (ManpowerGroup, 2024).

On the empirical front, prior studies of skill gaps largely rely on survey evidence and one-dimensional proxies that obscure the multidimensional nature of human capital (Deming, 2023; Brunello and Wruuck, 2021). As a result, there is limited systematic evidence on whether skill gaps in the context of information technology are driven primarily by technical deficiencies alone or by simultaneous gaps in IT-complementary non-technical capabilities such as managerial and organizational skills (Bresnahan et al., 2002; Brynjolfsson and Hitt, 2000). Moreover, little is known about the specific combinations of skills required in practice (Tambe (2026) is an exception), despite evidence that hiring difficulties often reflect firms' inability to find workers with the appropriate mix of attributes (OECD, 2024).

These theoretical and empirical issues are central to policy debates on workforce development and technology adoption. Understanding the determinants of skill formation is essential for designing effective training policies, particularly in distinguishing between skills that workers would acquire endogenously and those that may require policy intervention (Jonas et al., 2026; Dillon et al., 2025; Katz et al., 2022; Card et al., 2018). Policies aimed at accelerating technology adoption and improving productivity, such as recent initiatives to subsidize AI adoption (Agrawal et al., 2019), also require knowing how such interventions affect worker incentives to invest in technology-specific skills. Ignoring these incentives may lead to policies that inadvertently reduce productivity by widening skill gaps, particularly among firms operating away from the technological frontier (Syverson, 2011).

3 Conceptual Framework

To interpret our empirical characterization of skill gaps over the technology life cycle, we extend the classic vintage human capital framework of Chari and Hopenhayn (1991), incorporating skill gaps as an endogenous feature of the diffusion process. We also draw on insights from the empirical literature on productivity estimation and general purpose technologies to interpret the observed relationship between skill gaps and TFP (Brynjolfsson et al., 2021). In this section, we outline our framework and its implications for interpreting the data.

3.1 Baseline Model

We first summarize the key aspects of Chari and Hopenhayn (1991). The baseline model is an overlapping-generations economy with an infinite horizon (time t = 0,1,2...). In each time period t, technological progress take the form of a new technology being 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). There are two types of workers, young and old, each of mass 1, who live for two periods. Young workers, who begin life unskilled, choose to work at a firm adopting a particular technology vintage. After one period of experience, young workers acquire technology vintage-specific human capital and become skilled and old. When old, these workers can either join a firm with the same technology and serve as skilled workers, or they can switch to a firm that uses a different technology and serve as unskilled workers.

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. 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 employed in the same firm (i.e. $\frac{\partial f_U}{\partial S_t} > 0$). 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.

Equilibrium requires that workers sort across firms to maximize lifetime utility, firms choose technologies and make hiring decisions to maximize profits, and markets clear. Matching between firms and workers takes place in a frictionless labor market; the model abstracts from search and matching frictions, and unskilled workers are indifferent across firms that adopt technologies of different vintages. 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, then declines as technologies mature into legacy systems, dissipating to zero once the technology becomes obsolete and adoption ceases.

3.2 Extension: Skill Gaps

Our key extension to the framework of Chari and Hopenhayn (1991) is to introduce multidimensional skill gaps. In the original framework, human capital is technology-vintage specific but one-dimensional. We expand this structure by characterizing firms' technology skill requirements and workers' capabilities with skill weight vectors that each sum to one. Technology-vintage specific capital is characterized by a vector of specific skill weights, following the notion of specific human capital introduced by Lazear (2009); human capital is technology-specific insofar as each technology requires a particular combination of skills.

To ground the interpretation of firm skill requirements, we appeal to the related literature and assume that when a firm adopts a technology of a given vintage, it designs a production system with complementary organizational and capital investments built around that technology (Bresnahan et al., 2002; Brynjolfsson and Hitt, 2000; Brynjolfsson et al., 2002). Production involves a set of tasks, and the skills required to perform these tasks define the firm's skill requirements. This mapping from technologies to tasks and, in turn, to skill requirements is based on the task-based approach to human capital (Acemoglu and Autor,

2011).[1] To ground the interpretation of worker skill weights, we assume that workers acquire human capital through education, work experience, training programs, etc. Workers' skill weight vectors capture their comparative advantages across various skills.

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 $\mathcal{N}$). A skill is complementary to IT if its marginal productivity increases with IT usage; a skill is IT-neutral if not.

We denote the *individual* skill gap for skill i between a firm using technology of vintage 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).[2] The skill gap $G^{\tau,s}$

[1] There is a large literature examining how information technologies shape production processes, organizational design, and worker tasks. 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).

[2] We make this assumption for simplicity; this assumption can be relaxed to allow for skill to be defined in terms of imperfect matches between a firm and a worker.

of a firm-worker pairing of this type will thus be *zero*; vintage-specific human capital in the original model is now represented by the specific skill weights desired by the firm (Lazear, 2009). We redefine an unskilled worker as one with 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 is *non-zero*.

We 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 Appendix B, 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 are able to take on managerial roles that reduce 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 intuitive assumption is supported by prior evidence of a negative association between skill gaps and worker productivity (Bandiera et al., 2025; Guvenen et al., 2020; Lise and Postel-Vinay, 2020). Overall, 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, our adaptations microfound the key primitives of Chari and Hopenhayn (1991) while preserving the core mechanics of the original model. As a result, the stationary

equilibrium remains unchanged; for example, the distribution of skilled workers, s_τ, continues to be single-peaked across technology vintages. At the same time, the enriched foundations yield new implications that extend beyond the scope of the original framework, allowing us to interpret the empirical relationship between multidimensional skill gaps and information technology diffusion within a tractable general equilibrium setting. We discuss these implications below and provide a graphical illustration of the model in Figure 1.[3]

3.3 Implications

3.3.1 Skill Gaps over the Technology Life Cycle

The framework predicts that skill gaps between firms and their new hires, $G^{\tau,u}$, follow a U-shaped pattern over the technology life cycle. The reasoning is as follows. We postulate in Section 3.2 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) that is preserved in our framework—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 workers lack the relevant complementary skills. With time, as more workers gain experience working with these technologies, firms face lower adjustment costs in aligning organizational practices and

[3]Our extension maintains the assumption of frictionless labor markets in Chari and Hopenhayn (1991). This modeling choice allows us to abstract from search and matching frictions to maintain tractability, and isolate the role of firms' and workers' skill adjustment incentives over the technology life cycle. An additional modelling choice that maintains consistency with the original framework of Chari and Hopenhayn (1991) is that workers' whose skill vectors do not equal firm requirements along every skill dimension are considered equally unskilled in terms of technology-vintage specific human capital. This assumption allows us to represent multidimensional skill gaps using scalar measures of distance, rather than modeling explicit multidimensional matching between heterogeneous firms and workers. This abstraction is intentional: it enables a transparent and parsimonious mapping between the model and the empirical patterns, while avoiding the additional structure required to model matching frictions in high-dimensional skill spaces. A complementary literature develops richer models of worker–firm sorting and matching across multiple skill dimensions (see, e.g., Fredriksson et al. (2018); Lise and Postel-Vinay (2020); Guvenen et al. (2020); Lindenlaub and Postel-Vinay (2023); Bandiera et al. (2025); Althoff and Reichardt (2026)). Our approach instead focuses on how firms' skill requirements evolve with technology diffusion, taking a reduced-form view of matching in order to isolate the mechanisms of interest.

skill requirements with the capabilities of new hires. However, once technologies become sufficiently mature – what we label as "legacy systems" – and less productive than new technologies, workers have diminished incentives to join firms that continue to rely on these older systems. With fewer skilled workers, firms face higher adjustment costs when reorganizing tasks and skill requirements for new hires, causing skill gaps to rise once again. This general equilibrium mechanism shows how technological progress outpacing worker skill accumulation impacts skill gaps across the entire technology life cycle — not just at the frontier.

3.3.2 Skill Gap Decomposition

Our framework enables us to decompose the skill gap life cycle pattern to identify its individual skill and occupational drivers. If skill gaps are a constraint on IT adoption, we should observe shortfalls in IT-complementary skills rather than gaps in IT-neutral skills. Formally, if firms' demand for IT-complementary capabilities exceeds supply, this will be reflected in positive skill gaps, $G_i^{\tau,u} > 0$, for skills $i \in \mathcal{C}$. Because skill weights sum to one, these shortfalls must be offset by relative excess supply in other dimensions–i.e. IT-neutral skills– implying $G_i^{\tau,u} \leq 0$ for $i \in \mathcal{N}$. Similar reasoning applies to occupations: skill gaps should be more pronounced among occupations central to IT adoption. These occupations may include technical roles such as engineers and programmers, as well as non-technical occupations critical to the organizational complements of IT, such as managers and business analysts (Bresnahan et al., 2002; Bartel et al., 2007; Brynjolfsson and Hitt, 2000).

The framework also clarifies how skill gaps along specific dimensions evolve over the technology life cycle. As the stock of skilled workers, s_τ, expands, firms adjust their skill requirements to better align with the capabilities of new hires. The empirical patterns we document allow us to identify which dimensions of skill demand adjust over time, and how these adjustments differ across technologies and stages of diffusion. In particular, if IT-complementary skills are central to diffusion, the gaps along these dimensions should display a life-cycle pattern that mirrors the aggregate dynamics, reflecting the changing incentives

of firms and workers to invest in IT-complementary skill formation as technologies mature.

3.3.3 Organizational Capital and Skill Gaps

Our framework provides a lens for interpreting firm characteristics that shape skill gaps over the technology life cycle. In our empirical analysis, we focus on the relationship between skill gaps and organizational capital. Prior work shows that organizational capital—encompassing training systems, management practices, and workplace design—facilitates technology adoption and improves productivity (Bresnahan et al., 2002; Bloom et al., 2016; Ichniowski et al., 1997; Atkeson and Kehoe, 2005; Brynjolfsson et al., 2021). We examine whether these capabilities play a role in shaping skill gaps over the technology life cycle.

While the baseline framework of Chari and Hopenhayn (1991) does not explicitly feature organizational capital, our extension allows firms to differ in their ability to align job requirements with the skills of available workers. We interpret organizational capital as a form of technology that shifts this alignment capacity, in the spirit of Bloom et al. (2016). Firms with greater organizational capital face lower costs of adjusting job requirements and equipping workers with the skills needed to perform evolving tasks, enabling closer alignment between task demands and the skill distribution of new hires. In the model, this corresponds to a lower $\frac{\partial G^{\tau,u}}{\partial S_t}$, reflecting an enhanced ability of firms with greater organizational capital to mitigate skill gaps for a given technology vintage and stock of skilled labor. Our framework thus predicts that firms with greater organizational capital should exhibit systematically smaller skill gaps throughout the technology life cycle.

3.3.4 Skill Gaps and Firm Productivity

Lastly, we study the empirical relationship between skill gaps and measured productivity over the technology life cycle. To interpret these patterns, we draw on insights from the literature on productivity measurement and general purpose technologies (Bresnahan and Trajtenberg, 1995; Syverson, 2011). A common view in this literature is that in principle,

similar to Chari and Hopenhayn (1991), newer technologies are more productive than older ones, but these productivity differences may not be immediately visible in measured TFP. A prominent explanation offered by Brynjolfsson et al. (2021) is that new technologies require complementary intangible investments that take time to translate into observable output.

We apply this logic to skill gaps as a proxy for one such intangible asset: human capital. In our setting, skill gaps reflect the extent to which firms have not yet fully aligned workforce capabilities with the requirements of their technology. This interpretation provides a clear mapping from the skill gap life cycle pattern implied by our framework to observed productivity dynamics. Early in the technology life cycle, large skill gaps are associated with lower measured productivity, as firms are investing in worker training and task realignment to help employees adapt to new technologies. As firms and workers accumulate experience and acquire the human capital required by new technologies, skill gaps narrow and productivity rises. In later stages, as technologies mature and worker incentives to invest in legacy systems weaken, skill gaps widen again and are associated with lower productivity. Taken together, our framework thus implies that measured productivity is inversely related to skill gaps over the technology life cycle.

4 Empirical Approach

4.1 Ideal Setting and Challenges

Empirically characterizing skill gaps and their relationship to firm technologies is challenging. Ideally, one would like to observe the specific skills demanded by firms and the skills possessed by workers, and then construct a measure of the "distance" between these attributes for matched firm–worker pairs. Such a measure would allow us to quantify skill gaps and examine how they evolve over the technology life cycle.

In practice, implementing this approach faces several obstacles. First, standard data sources do not provide sufficiently granular detail on workers' human capital. Prior studies

typically rely on survey data or administrative datasets that capture worker demographics, education, and experience (Brunello and Wruuck, 2021). While informative, these sources lack the depth needed to characterize the specific dimensions of human capital that matter for technology adoption. For example, administrative data record education in broad categories; they do not capture specific competencies. Survey data similarly lack technology-specific granularity and are often subject to recall bias.

Second, it is conceptually difficult to identify the specific skills that firms demand, as many relevant capabilities are latent and not directly observable. For example, a job posting that requests "Java programming" may reflect demand not only for coding proficiency, but also for general problem-solving ability and the interpersonal skills required for collaborative software development. More broadly, firms rely on a range of implicit skill requirements—such as communication or adaptability—that are rarely specified explicitly in job postings. Analogous challenges arise when attempting to infer workers' skill portfolios from résumé data, where observed skill listings may only imperfectly capture underlying capabilities.

Prior studies analyzing skill content in job postings often rely on keyword-based classification schemes, mapping observed terms into predefined categories such as "cognitive" or "routine". This approach is poorly suited to our setting, as it requires the econometrician to map observed skill labels to underlying capabilities in a manner that may differ substantially from how firms and workers interpret these skills in practice. This mismatch can generate significant measurement error, distorting estimates of skill gaps and obscuring the underlying relationship between skill demand, skill supply, and technology use.

Third, processing skill data from job postings and worker résumés poses substantial computational challenges. These data contain thousands of heterogeneous and inconsistently defined skill labels that vary widely across firms and occupations. Deming and Kahn (2018) show that skill requirements differ markedly even within narrowly defined occupations, with firm-specific factors accounting for a substantial share of this variation. Similar issues arise on the supply side, where workers use diverse and idiosyncratic language to describe their

skills. This heterogeneity complicates the task of constructing measures of skill gaps that are both computationally tractable and comparable across firms and over time.

4.2 Empirical Methodology

We develop an empirical methodology to overcome the aforementioned challenges that utilizes recent advances in natural language processing and large language models (Hassan et al., 2025; Dell, 2025; Ludwig et al., 2026). We 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 from job postings and worker résumés into a common set of predefined skill groups. This classification step addresses the conceptual challenges noted above—namely, the identification of latent skills and the processing of heterogeneous skill listings—by producing comparable measures of skill attributes that can be used to estimate multidimensional skill gaps. A key advantage of this approach is that BERT leverages contextual information to interpret the meaning of skill descriptions, instead of relying on exact keyword matches. In effect, the model's semantic representations—trained on large text corpora—provide a data-driven mapping from observed skill labels to underlying capabilities that more closely reflects how firms and workers interpret skills in practice. This approach mitigates the measurement error inherent in manual or keyword-based classification schemes and enables consistent comparison of skill attributes across firms, workers, and technologies.

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

[4]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 (Yin et al., 2019). 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 (Kalyani et al., 2025; Tambe, 2014) to characterize firms' information technology vintages. For each technology listed in a job posting, we define its vintage as the difference between the posting year and the year of the technology's commercial introduction. We then aggregate these measures to the firm–year level and to the firm–year–technology-class level (see Figure 5) to characterize the vintage of the IT systems used by firms in a given year. Our main measures weight vintage estimates by the number of workers matched to the corresponding postings in which a technology appears, however, we also compute vintage measures without weights. Overall, these measures provide estimates of the position of firms' IT systems along the technology life cycle.

5 Data

5.1 Data Sources and Sample Construction

We construct a matched employer–employee dataset by combining three primary data sources. First, we use job postings from Lightcast covering U.S. publicly traded firms over the period 2010–2024. These data provide detailed information on firms' skill requirements and are well suited for measuring firm-level skill demand, particularly in the context of technology adoption (Acemoglu et al., 2022; Hershbein and Kahn, 2018; Deming and Kahn, 2018). Dalton et al. (2026) show that Lightcast postings closely track job openings for large establishments during this period, suggesting that postings provide a reliable proxy for hiring demand in our sample. Second, we use worker-level résumé data from Revelio Labs, which aggregates information from online professional profiles (sourced primarily from LinkedIn). These data provide detailed information on workers' self-reported skills and employment histories, including firm affiliations and job start dates. Third, we incorporate firm-level financial and

accounting data from Compustat, and TFP estimates from Imrohoroglu and Tuzel (2014) that are computed using the semiparametric approach developed by Olley and Pakes (1996).

We merge these sources at the firm–occupation–year level to construct measures of skill gaps. The Lightcast and Revelio data form the core inputs: job postings capture firms' skill requirements, while résumé data capture the skills supplied by workers. For each firm, two-digit Standard Occupational Classification (SOC) code, and year, we compare the skill profile implied by job postings with the skill profile of workers who join the firm and hold matching occupations in the subsequent year. To account for timing differences between postings and hiring, we also consider alternative matching windows that include contemporaneous hires and hires up to two years after the posting (discussed further in Section 7).

5.2 Firm Characteristics

Table 1 compares our sample to the broader Compustat population. Our sample comprises 20,858 firm-year observations spanning 2010 to 024. Relative to Compustat firms, sample firms are substantially larger and more established: median sales are $1.4 billion (vs. $127 million), mean employment is 0.023 million (vs. 0.010 million), and median firm age is 23 years (vs. 9 years). Mean and median TFP is also higher for sample firms compared to the broader Compustat population. These differences reflect selection toward firms that are more likely to post jobs online and employ workers with observable online profiles. As a result, our sample primarily captures large, established firms in the U.S. We discuss the implications of this sample composition for the interpretation of our findings in Section 7.

5.3 Information Technology Vintage

A central variable in our analysis is information technology vintage. Figure 2 shows substantial dispersion in technology vintages across firm–year observations. The modal firm operates technologies that are approximately 28–31 years old, and a non-trivial share of firms rely on technologies older than 35 years. These patterns indicate that firms in our sample operate at

markedly different points along the technology life cycle. Moreover, the diffusion patterns we document for information technologies closely mirror the slow and uneven adoption dynamics of general-purpose technologies established in the prior literature (Kalyani et al., 2025; Comin and Hobijn, 2010).

To illustrate the economic content of this measure, Table A1 reports the most frequently cited technologies within each vintage bin. Younger vintages include frontier technologies such as cloud computing platforms (e.g., Amazon Web Services, Microsoft Azure) and modern programming languages (e.g., Python). Older vintages include legacy systems and lower-level languages associated with earlier computing paradigms, such as COBOL, Unix, SQL, and C. The progression across bins closely tracks well-known waves of enterprise technology adoption, providing external validation for our vintage classification.

5.4 Skill Weight Vectors

Figure 3 illustrates the sample distribution of skill weights. To summarize firms' skill requirements, we compute the average entailment scores for each of the 25 skill categories defined by Revelio across all firm job postings in the sample. We compute the corresponding skill weight means for workers in the sample that are matched to job postings. The mean skill gap for each individual skill is the difference between the average firm weight and average worker weight for that skill.

We further group skills into three categories based on their complementarity with information technology deployment.[5] The first category comprises IT-complementary technical skills, including programming languages, database management, networking, and related capabilities. The second comprises IT-complementary non-technical skills — managerial, organizational, and coordination capabilities that are central to the organizational investments firms need to make in order to effectively operate IT systems (Bartel et al., 2007; Brynjolfsson et al., 2002). The third category comprises IT-neutral skills whose productivity does not

[5]Table A2 in the appendix provides the full list of skill groups and describes the categorization process.

increase significantly with information technology, such as public speaking, event planning, and hospitality competencies.

Figure 3 shows that the largest skill gaps are concentrated in IT-complementary non-technical categories: management and leadership (mean gap of 0.0338) and process improvement and program management (0.014). Among IT-complementary technical skills, networking and troubleshooting exhibits the largest gap (0.006), while gaps in programming-intensive categories such as data analysis, engineering, 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 all firm-year observations.

The distribution of skill gaps across categories underscores the relevance of our measure for studying IT adoption. Skill shortages are concentrated in IT-complementary domains—particularly managerial, organizational, and technical capabilities that are essential for implementing new technologies—while IT-neutral skills are comparatively oversupplied. This alignment suggests that our skill gap measure isolates the skill dimensions that are most consequential for information technology adoption, and do not merely reflect arbitrary imbalances between skill demand and supply.

6 Empirical Findings

We begin by graphically depicting the joint evolution of skill gaps and firm productivity over the technology life cycle. Figure 4 plots average skill gaps and average firm TFP against technology vintage, grouping observations into ten equal-sized bins. Skill gaps exhibit a pronounced U-shaped pattern: high for firms adopting new technologies, declining as technologies mature, and rising again as technologies approach obsolescence. Quantitatively, skill gaps fall by roughly 40 percent from their peak among the youngest technologies to a minimum

at mid-life vintages, before increasing again for the oldest technologies to levels comparable to the early stages of the life cycle.

In contrast, firm-level TFP follows a hump-shaped pattern. Productivity is low for firms adopting young technologies, it rises for firms with more mature systems, and then it declines as firms' technologies approach obsolescence. Notably, the decline in TFP coincides (approximately) with the point at which skill gaps begin to widen. Taken together, these patterns reveal a systematic inverse relationship between skill gaps and productivity over the technology life cycle — core empirical facts that we explore in more detail below.

6.1 Skill Gaps over the Technology Life Cycle

We formalize the skill gap pattern depicted in Figure 4 by estimating regressions of skill gaps on technology age. The results are reported in Table 2. Across specifications, the linear term in vintage is negative (ranging between -0.835 and -2.590) and the quadratic term is positive (ranging between 1.301 and 3.923), consistent with a non-monotonic, U-shaped relationship. When we regress skill gaps on indicator variables for technology vintage bins, we observe coefficients that mirror the U-shaped pattern in Figure 4. The coefficient estimates are robust to the inclusion of year and industry fixed effects, indicating that the skill gap pattern holds both across industries and within industries. Specifications with firm fixed effects also yield large and statistically significant vintage coefficient estimates, though estimation relies on relatively limited within-firm variation in technology vintage over the sample period.

We further show that the aggregate patterns are stable across distinct classes of information technologies. Figure 5 depicts a U-shaped pattern in skill gaps for firms deploying enterprise resource systems, database technologies, and software development tools. The consistency of the U-shaped pattern across these categories suggests that it is a general feature of IT diffusion rather than an artifact of a particular technology, industry, or firm. The graphical evidence is reinforced by the regression estimates presented in Table 3, where each column corresponds to a particular technology class. Coefficient estimates for early vintage

bins are negative, while they are positive and significant for later bins—showing that skill gaps follow a U-shaped pattern even after controlling for firm and year fixed effects.

The life cycle pattern that we observe is consistent with our framework, as discussed in Section 3.3. Skill gaps are large early in the life cycle because few workers possess the skills required by new technologies — a *direct* implication of technological change outpacing skill formation. Over time, as adoption expands, workers acquire human capital and firms adjust skill requirements to more closely align with the capabilities of available labor, causing skill gaps to decline. At later stages, however, the arrival of newer and more productive technologies diminishes workers' incentives to match with firms relying on legacy systems, causing skill gaps for these firms to widen — an *indirect*, general equilibrium effect of technological change outstripping skill formation.

6.2 Skill Gap Decomposition

We decompose the life-cycle pattern to identify its underlying drivers over the diffusion process (as discussed in Section 3.3.2). First, we examine individual skill dimensions and show that skill gaps do not reflect deficits in technical skills alone. Instead, they reflect shortages in a broader set of competencies—both technical and non-technical—that are complementary to IT adoption. Second, we analyze variation across occupations and show that the aggregate life-cycle pattern is driven primarily by occupations that play a central role in the implementation and use of new technologies. Taken together, these results highlight that the adjustment process between technological progress and skill formation is inherently multidimensional and complex.

6.2.1 Individual Skill Dimensions

Figure 3 shows systematic differences between firm requirements and worker capabilities across a broad set of skills. The largest shortfalls arise in IT-complementary dimensions, including both technical capabilities (e.g., data analysis, engineering) and non-technical com-

petencies (e.g., management and organizational skills). In contrast, IT-neutral skills—such as public speaking and event planning—tend to be oversupplied. These patterns support the view that constraints on IT-adoption are concentrated in IT-complementary skill shortages.

To further quantify these differences, we group individual skills into three categories: IT-complementary technical, IT-complementary non-technical, and IT-neutral. The classification scheme is described in detail in Table A2. We compute Euclidean distance measures of skill gaps within each category. Figure 6 and Table 4 show how these category-specific gaps evolve with technology vintage. The largest gaps are observed for IT-complementary non-technical skills. Combined with the raw differences in Figure 3, this reflects the fact that firm requirements consistently exceed worker capabilities along these dimensions. In addition, while IT-complementary technical and IT-neutral skill gaps appear similar in magnitude when measured using Euclidean distance, they differ sharply in direction. Technical gaps reflect shortages (requirements exceeding capabilities), whereas IT-neutral gaps largely reflect oversupply. This distinction highlights that aggregate distance measures mask important asymmetries in the underlying skill distributions.

The evolution of these category-specific gaps over the technology life cycle also reveals that the U-shaped pattern in aggregate skill gaps is driven primarily by IT-complementary skills. Adjustments in these dimensions account for both the initial decline and the subsequent increase in overall mismatch. Interpreted through the lens of the model, this pattern is consistent with firms initially reducing gaps in the capabilities most closely tied to productivity gains from IT adoption, followed by a widening of gaps as technologies age and the relative value of associated skills declines.

Taken together, these results show that skill gaps are inherently multidimensional. They do not simply reflect shortages of technical expertise in isolation, but rather deficits in bundles of complementary capabilities—including managerial and organizational skills—that are required to implement and utilize new technologies effectively (Bresnahan et al., 2002; Bartel et al., 2007). This evidence shifts the characterization of skill gaps away from one-

dimensional measures toward a framework in which successful technology adoption depends on the alignment of interdependent technical and non-technical skills. In this framework, shortages in technical capabilities are inseparable from shortfalls in the non-technical skills required to implement, coordinate, and scale information technologies, and the alignment between these capabilities and firm requirements evolves over the diffusion process.

6.2.2 Skill Gaps by Occupation

We next examine heterogeneity in skill gap dynamics across occupations by estimating our baseline specification separately for each 2-digit standard occupational code (SOC). The results, reported in Table 5, reveal that the aggregate U-shaped pattern in skill gaps is driven primarily by occupations that are central to IT implementation and organizational change. In particular, we observe pronounced and statistically significant curvature in management, computer and mathematical, engineering, sales, and production occupations—roles that are closely linked to the adoption and effective use of information technologies. By contrast, occupations with more limited exposure to IT—such as those in social services, education, or personal care—exhibit little systematic relationship between skill gaps and technology vintage. This distinction suggests that the life cycle dynamics documented in the aggregate are not pervasive across the labor market, but instead concentrated in occupations where technology adoption meaningfully reshapes skill requirements.

We also find substantial heterogeneity across occupations in the composition of skill gaps. The specific combinations of technical, managerial, and organizational skills that drive mismatch differ markedly across occupational groups, indicating that human capital constraints are task-specific and depend on the production context. Taken together, these findings reinforce the view that the evolution of skill gaps over the technology life cycle reflects targeted shortages in occupations critical for implementing and adapting to new technologies, rather than a uniform shortage of skills across the workforce.

6.3 Skill Gaps and Organizational Capital

Our third set of findings shows that organizational capital plays a key role in explaining skill gaps throughout the life cycle. Figure 7 illustrates that skill gaps are systematically smaller in firms with higher stocks of organizational capital (measured following Eisfeldt and Papanikolaou (2013)), and similarly, in firms with greater numbers of employees (a proxy for accumulated organizational knowledge, following Atkeson and Kehoe (2005)). Table 6 shows that these patterns remain robust in regression specifications with various controls. For example, Column 3 in Panel A shows that skill gap differences between firms with high and low organizational capital range between 0.026 and 0.045 for both early and later stage vintages. Column B illustrates that vintage effects on skill gaps are even more pronounced when we segment firms based on the number of employees.

Within our framework (discussed in Section 3.3.3), these differences reflect variation in firms' ability to adapt production processes and align skill requirements with worker capabilities as technologies diffuse. Firms with greater organizational capacity——through better management practices, training systems, and internal coordination——face lower costs of adjusting tasks and reallocating responsibilities. As a result, they are better able to reduce misalignment between required and supplied skills, leading to persistently smaller skill gaps over the technology life cycle.

These findings offer a new perspective on the role of organizational capital in technology diffusion. Prior work emphasizes that organizational investments complement IT adoption and enhance productivity (Brynjolfsson et al., 2021, 2002; Bloom et al., 2012; Syverson, 2011; Eisfeldt and Papanikolaou, 2013). Our results point to a specific mechanism underlying this phenomenon: organizational capital mitigates frictions in aligning skill demand with skill supply within the firm. This mechanism is particularly relevant in light of the evidence from Section 3.3.2, which highlights the central role of IT-complementary managerial and organizational skills in driving life cycle patterns in skill gaps. More broadly, the results

suggest that organizational capacity helps firms adjust to skill imbalances that arise as technological change outpaces human capital accumulation.

6.4 Skill Gaps and Firm Productivity

We document a robust, negative empirical relationship between skill gaps and firm productivity over the technology life cycle. As shown in Figure 4, firm productivity (TFP) moves inversely with skill gaps across technology vintages. We examine this relationship more formally by regressing firm TFP on skill gaps while controlling for technology vintage and other firm characteristics. The results, reported in Table 7, indicate that the coefficient on skill gaps is negative across specifications, ranging from -0.339 to -0.374. The consistent negative coefficient on skill gaps, even after controlling for technology vintage, indicates that differences in skill gaps account for meaningful variation in productivity even among firms using technologies of a similar age.

As discussed in Section 3.3.4, we interpret this relationship using insights from the literature on productivity measurement and general purpose technologies (Bresnahan and Trajtenberg, 1995; Brynjolfsson et al., 2021; Syverson, 2011). As firms invest in training and work practices to close skill gaps, they build up human capital and complementary organizational assets. These intangible investments do not generate measurable output immediately, however. Instead, consistent with the J-curve dynamics of productivity (Brynjolfsson et al., 2021), they take time to translate into real productivity gains.

Building on this interpretation, the life cycle pattern in skill gaps maps naturally into observed productivity dynamics. Measured productivity is low when skill gaps are large early in the technology life cycle. It rises as gaps narrow with firm and worker adaptation and the accumulation of valuable intangible capital. At later stages, however, productivity falls as skill gaps widen due to weaker incentives for workers to invest in legacy systems. These patterns illustrate a unique mechanism by which the dynamics of technological progress and human capital accumulation shape productivity dispersion not only at the frontier — but

across all stages of the technology life cycle.

7 Robustness Analysis

We conduct a range of analyses to assess the robustness of our empirical findings to various sources of measurement error and sampling biases.

7.1 Alternative Skill Gap Measures

Our baseline skill gap measures are constructed by comparing the average skill requirements listed in job postings with the average capabilities of newly hired workers, using matches based on firm, two-digit occupation codes, and the time-series restriction that job postings appear in the calendar year prior to the worker's start year. Because we cannot directly observe whether a given hire corresponds to a specific job posting, our skill gap measures are subject to several sources of measurement error and sampling biases.

First, workers and job postings may be subject to occupational misclassification. Second, there are many job postings for vacancies that remain unfilled and many workers who are hired for vacancies that do not correspond to online postings. Third, workers and job postings in our sample may not reflect population data; our sample is skewed towards the types of workers and jobs that are more likely to be matched through online search. Fourth, matches between firms and workers may occur outside of the time windows used in our baseline specification. All of these cases may generate discrepancies between measured job requirements and worker characteristics, leading to inaccurate measures of skill gaps.

We assess the empirical importance of these concerns by implementing a range of alternative matching procedures and presenting our results using different skill gap measures. To address mismatches due to potential misclassification at the occupation level, we relax the requirement that workers and postings share the same occupation code and instead match workers to postings at the firm-level. This approach allows workers to be paired with postings

outside of narrowly defined occupational categories and thus incorporates a broader comparison of firm skill requirements with worker skill sets. The results remain highly stable: Columns 1 to 3 of Table 8 show that across various regression specifications, the U-shaped relationship between skill gaps and technology vintage persists.

To assess the importance of skill vectors being skewed by unfilled vacancies or new hires unrelated to job postings—as well as the importance of the population representativeness of our sample—we reweight our skill gap measures using industry and occupation weights from the Current Population Survey (CPS). This adjustment aligns our sample more closely with the population distribution of jobs and workers, mitigating potential biases arising from differential coverage of occupations and industries in online data. Although this procedure introduces some loss of firm-level variation, the results in Table 9 continue to exhibit a robust U-shaped relationship between skill gaps and technology vintage.

Third, to assess whether our skill gap measures are skewed by the firm-level aggregation of measurement error in occupational matches, sampling rates vary by occupation, we estimate our baseline specifications at the disaggregated firm-occupation-level rather than at the aggregated firm-level. Table 8 (Columns 4–6) shows that the relationship between technology vintage and skill gaps remains U-shaped even at the firm-occupation-level, indicating that our firm-level findings are unlikely to be driven by the aggregation of measurement error due to occupational mismatch.

Fourth, we vary the temporal window used to match workers to job postings to assess the sensitivity of our results to the timing of new hires and job postings used for measuring skill gaps. We first match workers to job postings that are made within the same calendar year—even if jobs postings appear after a new worker is hired. While this matching procedure potentially creates look-ahead bias, this procedure may lead to closer alignment between firm skill requirements and worker skills. Columns 1 to 3 of Table 10 show that across different specifications, the U-shaped relationship between skill gaps and technology vintage remains robust. We also increase the matching window to include two years of job postings that

predate nire hire start dates, and the results remain robust (Columns 4 to 6 of Table 10).

7.2 Measurement Error in Worker Skill Weights

Another important concern is the measurement of workers' skill profiles, which are constructed from the skills listed on workers' LinkedIn profiles. The timing for when these skills are reported on the platform is unknown. Therefore, these data may include skills acquired after workers begin the matched jobs in our sample (Bloom et al., 2026). This timing issue introduces potential measurement error into our estimates of worker skill profiles.

To assess the empirical importance of this concern, we first note that this source of measurement error likely attenuates our measures of skill gaps. If worker skill profiles are posted after we observe firm transitions that comprise our matched sample, then our skill gap measures reflect comparisons between job skill requirements with worker skills after they acquired additional skills that were likely advertised on the postings which they matched. We also re-estimate our results using only observations from the final year of the sample, when the elapsed time between workers' job transitions and the observation of their profiles is shortest. The patterns we document persist in this restricted sample, as shown in Table 11. Moreover, the results remain stable as we sequentially expand the sample to include earlier years with longer potential exposure to post-transition skill updates. These findings indicate that our results are robust to alternative assumptions about the timing of observed worker skills and measurement error in our construction of worker skill weight vectors.

7.3 Technology vintage measures

We evaluate the robustness of our results to different methods for measuring information technology vintages. First, we rerun our analysis with a measure of technology vintage that is based on the average age of all information technologies list in job postings—without weighting the technology ages by the numbers of matched workers in the sample. Second, we study the median technology age in the sample, rather than the mean. Third, we re-

move information technologies such as Microsoft Office products that are commonplace and potentially irrelevant proxies for firm information technologies. These analyses supplement our study of individual classes of technologies and our aggregate IT-vintage measures that comprise our main results. Across all of these alternative measures, our main results continue to hold. Some of these results are discussed in Section 6, while other results are omitted for the sake of brevity and available upon request. We observe a robust U-shape in skill gaps over the life cycle across all of the IT-vintage measures that we construct.

7.4 Choice of language model for classifying skill content

Our main results are based on skill vectors constructed using BERT, a bidirectional transformer language model that has become a standard tool for text classification tasks. A natural concern is whether our findings are sensitive to the choice of language model. Recent advances in transformer architectures — most notably ModernBERT — have yielded models that improve upon BERT's contextual representations, but at substantially greater computational cost: processing our full sample with ModernBERT requires several weeks of compute time, making it infeasible as a primary estimator.

To assess sensitivity of our results to the choice of LLM, we re-estimate our main specifications using both ModernBERT and BERT for one year of data—2020. The results using ModernBERT, presented in Column 3 of Table 12, are slightly stronger than our estimates using BERT in Column 1: the U-shaped relationship between skill gaps and technology vintage is more pronounced when using ModernBERT. This pattern suggests that our baseline estimates are conservative and unrelated to our choice of LLM.

7.5 Alternative Skill Vector Distance Measures

Our main measure of skill gaps is based on the Euclidean distance between the skill vectors that characterize firm requirements and those that characterize worker capabilities. This measure captures the magnitude of reallocation across skill dimensions required to align

worker capabilities with firm demands, placing equal weight on deviations across individual skill categories. To assess whether our findings are sensitive to the choice of distance metric used to estimate skill gaps, we consider two alternative measures commonly used in machine learning: cosine similarity and Hellinger distance.

Cosine similarity captures the alignment in skill vectors by focusing on how closely the relative distribution of skills across categories matches between firms and workers. Cosine similarity is insensitive to the scale of differences across dimensions, as it does not directly capture the extent of reallocation required to close skill gaps. Hellinger distance, by contrast, is designed for comparing probability distributions and therefore respects the compositional nature of our skill vectors, which are nonnegative and sum to one. However, while this property makes it well suited for measuring distributional divergence, its nonlinear transformation complicates interpretation in terms of economic distance or adjustment costs.

Our results are robust to both alternative measures, as shown in Columns 4 and 5 of Table 12. Using cosine similarity, we find that the alignment of worker and firm skill profiles follows a U-shaped pattern over the technology life cycle. Similarly, Hellinger distance yields a U-shaped relationship between skill gaps and technology vintage. Taken together, these findings indicate that our results do not depend on the specific metric used to quantify skill gaps. We nevertheless prefer Euclidean distance as our baseline measure, as it provides a transparent and economically interpretable notion of distance that aligns closely with the adjustment cost microfoundation in our framework.

7.6 Alternative Skill Classifications

We evaluate the robustness of our findings to alternative skill labels. Specifically, we analyze the skill categories proposed by Deming and Kahn (2018), originally devised to study heterogeneity in skill demand across firms. The skill groupings are fewer in number (10 versus 25), and encompass multiple categories that feature in our baseline classification. The advantage of these alternative skill labels is that we may estimate more robust skill weights using fewer

skill labels. The disadvantage of using these labels is that we may ignore important information contained in raw skills data that is not spanned by these categories. For example, skill categories such as "cognitive" and "technical", may not be granular enough to capture differences in skill content across different IT vintages.

The results of this analysis are presented in Figure A1 and Table A3. The U-shaped pattern in skill gaps is more pronounced when we use the alternative skill classification scheme developed by Deming and Kahn (2018). Figure A1 shows that skill gaps for new technologies are approximately 0.381 while skill gaps reach their nadir at approximately 0.260 for technologies of vintages between 32 and 38 years. Table A3 shows that the pattern is robust across regressions that include different controls. These findings indicate that the main empirical patterns that we document are not merely an artifact of the specific skill classification scheme that we deploy in our baseline analysis.

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. Revelio's skill labels can thus help identify skill gaps that might otherwise be obscured by more coarse categories of skills. Second, the labels that we utilize are identified and verified by Revelio Labs to efficiently span the full skill space represented in the underlying résumé data. Thus, the skill labels that we analyze are likely to best capture the breadth and depth of skills that appear in the data.

8 Conclusion

This paper presents new empirical evidence to characterize the evolution of skill gaps over the information technology life cycle. We show that skill gaps — defined as the distance between firm skill requirements and worker skill attributes — follow a systematic U-shaped pattern over the technology diffusion process. Interpreted through an extended vintage human capital

framework based on Chari and Hopenhayn (1991), skill gaps are high at the onset of adoption because workers are slow to adapt to new technologies. With time, skill gaps decline as firms and workers learn. At some point during the diffusion process, however, workers face diminished incentives to work with legacy systems that are less productive than frontier technologies, causing skill gaps to widen again in the later stages of the life cycle.

We further show that skill gaps are driven by shortages in combinations of technical and non-technical skills that complement IT investment, rather than technical skill deficits alone. In addition, firms with greater organizational capital face systematically smaller skill gaps. Lastly, we show that skill gaps are inversely related to firm productivity over the life cycle, consistent with the J-curve dynamics of general purpose technologies (Brynjolfsson et al., 2021): closing skill gaps requires investing in intangible assets that take time to show observable productivity growth.

Taken together, the findings show that the uneven race between technological progress and human capital accumulation leads to skill gaps and productivity dispersion across across the entire technology life cycle. Prior work has mostly focused on the technology frontier, but these dynamics remain critical even for firms reliant on older technologies that remain central to production. As such, our work highlights the importance of firm and worker incentives to invest in skill formation throughout the diffusion process. Failure to account for these dynamics can obscure the interpretation of skill gaps, and may lead to suboptimal policies in the context of workforce development and technology adoption. These issues warrant further inquiry into the additional causes and consequences of skill gaps.

References

Acemoglu, D. and Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. In *Handbook of Labor Economics*, volume 4, Part B, pages 1043–1171. Elsevier.

Acemoglu, D., Autor, D. H., Hazell, J., and Restrepo, P. (2022). Artificial intelligence and jobs: Evidence from online vacancies. *Journal of Labor Economics*, 40(S1):293–340.

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 Restrepo, P. (2019). Artificial intelligence, automation, and work. In Agrawal, A., Gans, J., and Goldfarb, A., editors, *The Economics of Artificial Intelligence: An Agenda*, pages 197–236. University of Chicago Press, Chicago.

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.

Agrawal, A., Gans, J., and Goldfarb, A. (2019). Economic policy for artificial intelligence. *Innovation Policy and the Economy*, 19(1):139–159.

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

Althoff, L. and Reichardt, H. (2026). Task-specific technical change and comparative advantage. Working Paper 35353, National Bureau of Economic Research.

Angelini, D. (2025). Aging population and technology adoption. Technical report, Department of Economics, University of Konstanz. Working Paper Series.

Atkeson, A. and Kehoe, P. J. (2005). Modeling and measuring organization capital. *Journal of Political Economy*, 113(5):1026–1053.

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. (2015). Why are there still so many jobs? the history and future of workplace automation. *Journal of Economic Perspectives*, 29(3):3–30.

Autor, D. H., Katz, L. F., and Krueger, A. B. (1998). Computing inequality: Have computers changed the labor market? *Quarterly Journal of Economics*, 113(4):1169–1214.

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.

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., Moore, G., Simon, L. K., and Wilkie-Rogers, C. (2026). Time travel on professional profiles. Working Paper 35546, National Bureau of Economic Research.

Bloom, N., Sadun, R., and Reenen, J. V. (2012). Americans do IT better: Us multinationals and the productivity miracle. *American Economic Review*, 102(1):167–201.

Bloom, N., Sadun, R., and Van Reenen, J. (2016). Management as a technology? Working Paper 22327, National Bureau of Economic Research. Revised October 2017.

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.

Bresnahan, T. F. and Trajtenberg, M. (1995). General purpose technologies 'engines of growth'? *Journal of Econometrics*, 65(1):83–108.

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

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

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

Brynjolfsson, E. and McAfee, A. (2011). *Race Against the Machine: How the Digital Revolution is Accelerating Innovation, Driving Productivity, and Irreversibly Transforming Employment and the Economy*. Digital Frontier Press, Lexington, MA.

Brynjolfsson, E. and McAfee, A. (2014). *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. W. W. Norton & Company, New York.

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

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.

Caselli, F. and Coleman, Wilbur John, I. (2006). The world technology frontier. *American Economic Review*, 96(3):499–522.

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

Cirera, X., Comin, D., and Cruz, M. (2026). Technology sophistication across establishments. *The Quarterly Journal of Economics*, 141(3):2025–2085.

Comin, D. and Hobijn, B. (2010). An exploration of technology diffusion. *American Economic Review*, 100(5):2031–2059.

Dalton, M., Kahn, L. B., and Mueller, A. I. (2026). Do online job postings capture job vacancies? NBER Working Paper 35697, National Bureau of Economic Research.

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. In *Handbook of the Economics of Education*, volume 7, pages 469–504. Elsevier.

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.

Eisfeldt, A. and Papanikolaou, D. (2013). Organization capital and the cross-section of expected returns. *Journal of Finance*, 68:1365–1406.

Fredriksson, P., Hensvik, L., and Skans, O. N. (2018). Mismatch of talent: Evidence on match quality, entry wages, and job mobility. *American Economic Review*, 108(11):3303–3338.

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.

Hassan, T. A., Hollander, S., Kalyani, A., van Lent, L., Schwedeler, M., and Tahoun, A. (2025). Text as data in economic analysis. *Journal of Economic Perspectives*, 39(3):193–

220.

Hassan, T. A., Kalyani, A., and Restrepo, P. (2026). The skill premium in times of rapid technological change. *NBER Working Paper Series*, (34939).

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

Ichniowski, C., Shaw, K., and Prennushi, G. (1997). The effects of human resource management practices on productivity: A study of steel finishing lines. *American Economic Review*, 87(3):291–313.

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

Javadian Sabet, A., Bana, S. H., Yu, R., and Frank, M. R. (2024). Course-skill atlas: A national longitudinal dataset of skills taught in u.s. higher education curricula. *Scientific Data*, 11:1086.

Jonas, L., Salomons, A., vom Baur, C., and Zierahn-Weilage, U. (2026). Curriculum updates reduce skill mismatch. Technical report, Working Paper.

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. and Murphy, K. (1992). Changes in relative wages, 1963-1987: Supply and demand factors. *Quarterly Journal of Economics*, 107:35–78.

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. and Postel-Vinay, F. (2023). Multidimensional sorting under 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.

Ludwig, J., Mullainathan, S., and Rambachan, A. (2026). Large language models: An applied econometric framework. *Annual Review of Economics*, 18(Volume 18, 2026):283–316.

ManpowerGroup (2024). 2024 global talent shortage: Infographic. Available at: https://go.manpowergroup.com/talent-shortage-2024.

Michaels, G., Natraj, A., and Reenen, J. V. (2014). Has ICT polarized skill demand? Evidence from eleven 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.

Nelson, R. R. and Phelps, E. S. (1966). Investment in humans, technological diffusion, and economic growth. *The American Economic Review*, 56(1/2):69–75.

OECD, editor (2024). *Understanding Skill Gaps in Firms: Results of the PIAAC Employer Module*. OECD Skills Studies. OECD Publishing, Paris. Accessed: October 5, 2026.

Olley, G. S. and Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263–1297.

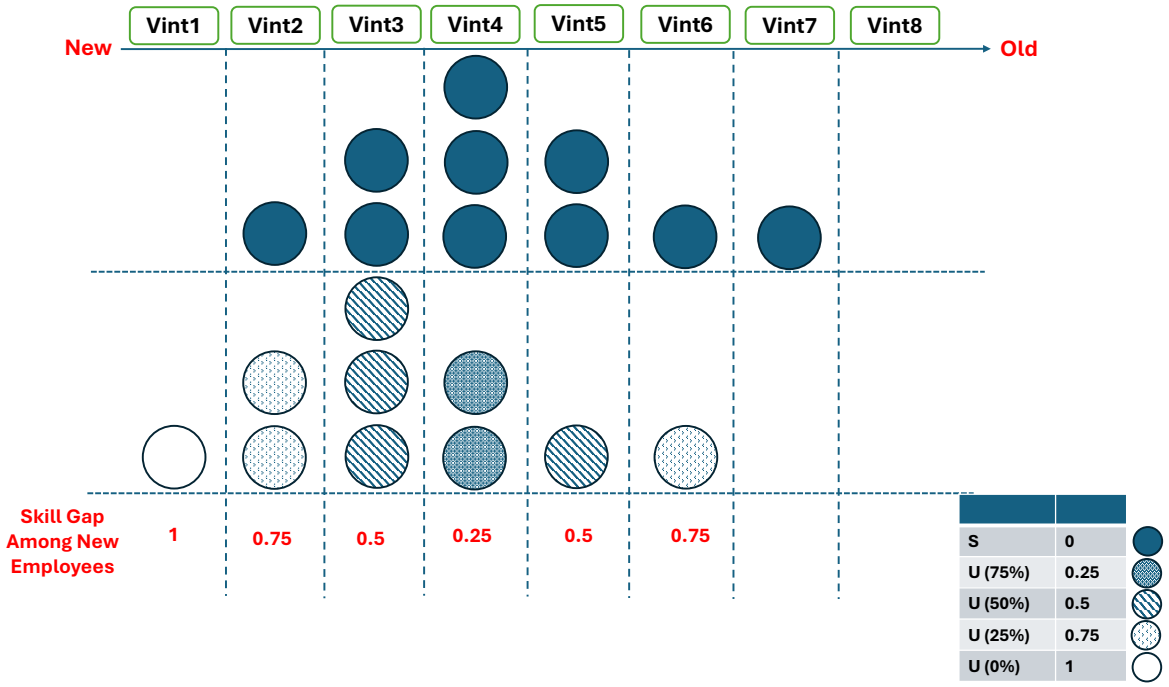
Syverson, C. (2011). What determines productivity? *Journal of Economic Literature*, 49(2):326–65.

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.

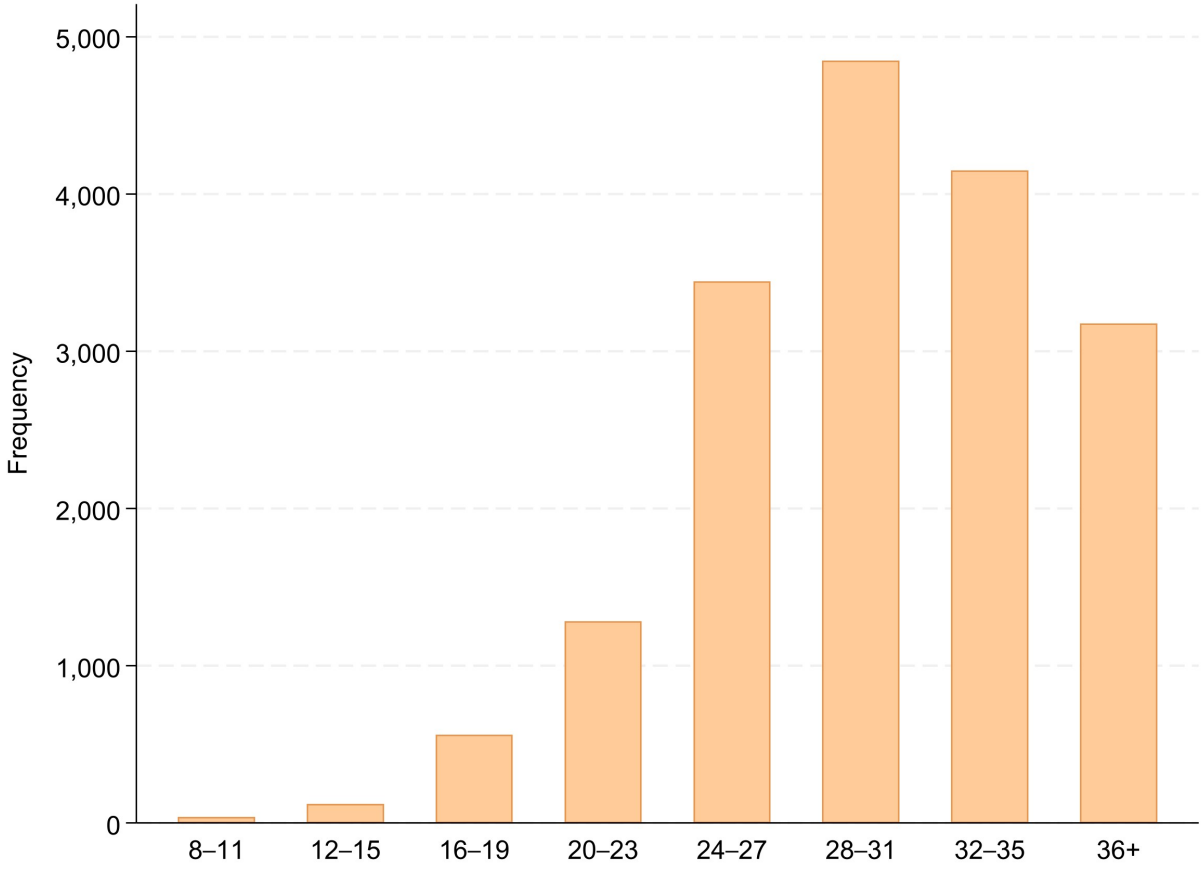
Yin, W., Hay, J., and Roth, D. (2019). Benchmarking zero-shot text classification: Datasets, evaluation and entailment approach. *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*,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

Figure 1: Model Overview



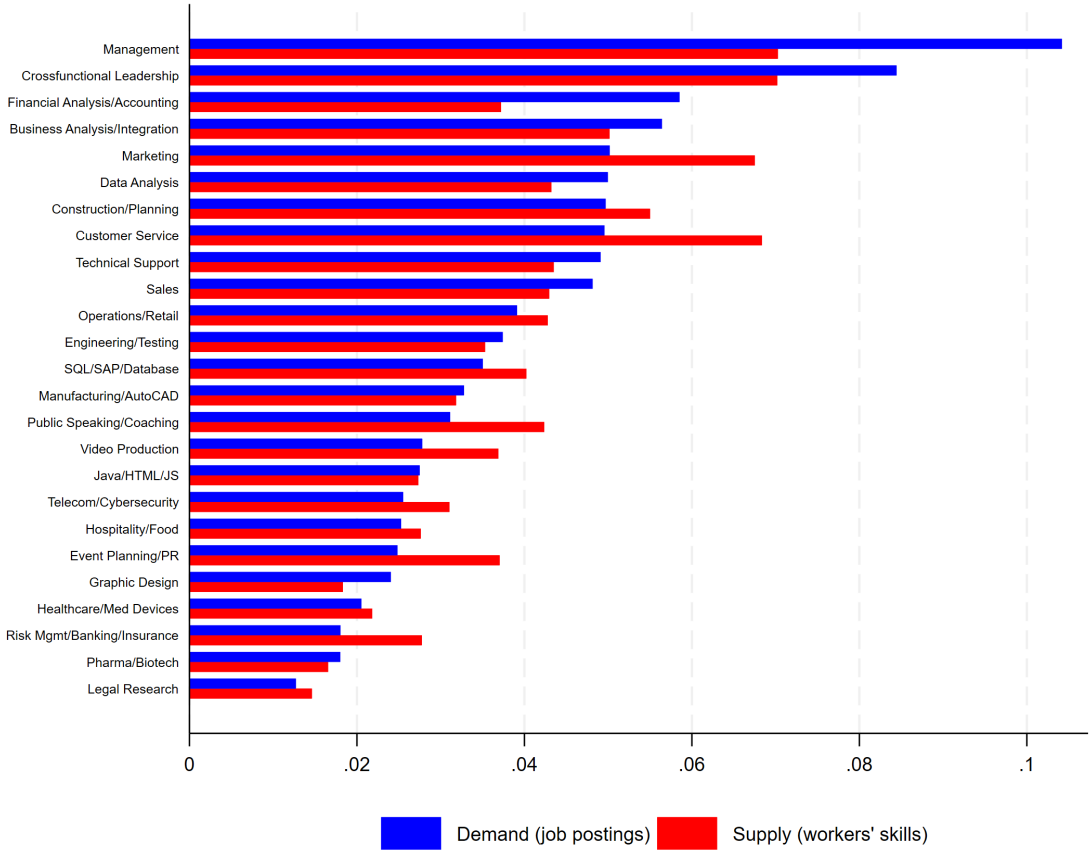
This figure provides a graphic illustration of the model's stationary equilibrium. Each column corresponds to a technology vintage, increasing from the youngest to oldest vintage. In the first period, a single unskilled worker joins a firm adopting the newest technology. Because the firm has no existing skilled employees, the level of skill gap between the new worker and the firm is normalized to 1. After one period of experience, this worker becomes skilled, and two new unskilled workers join the firm. The skill gap for the two new hires is 0.25; the presence of the skilled worker enables the firm to rearrange tasks and skill requirements to more closely align with the capabilities of new hires. The evolution of the skill gap between the firm and its new hires continues each period until the technology reaches middle vintage 4. From this point onward, fewer unskilled workers 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 hires fewer and fewer new employees over time, causing skill gaps between the firm and its new hires to increase. This process continues until the technology becomes obsolete in the last period.

Figure 2: Information Technology Vintage Distribution



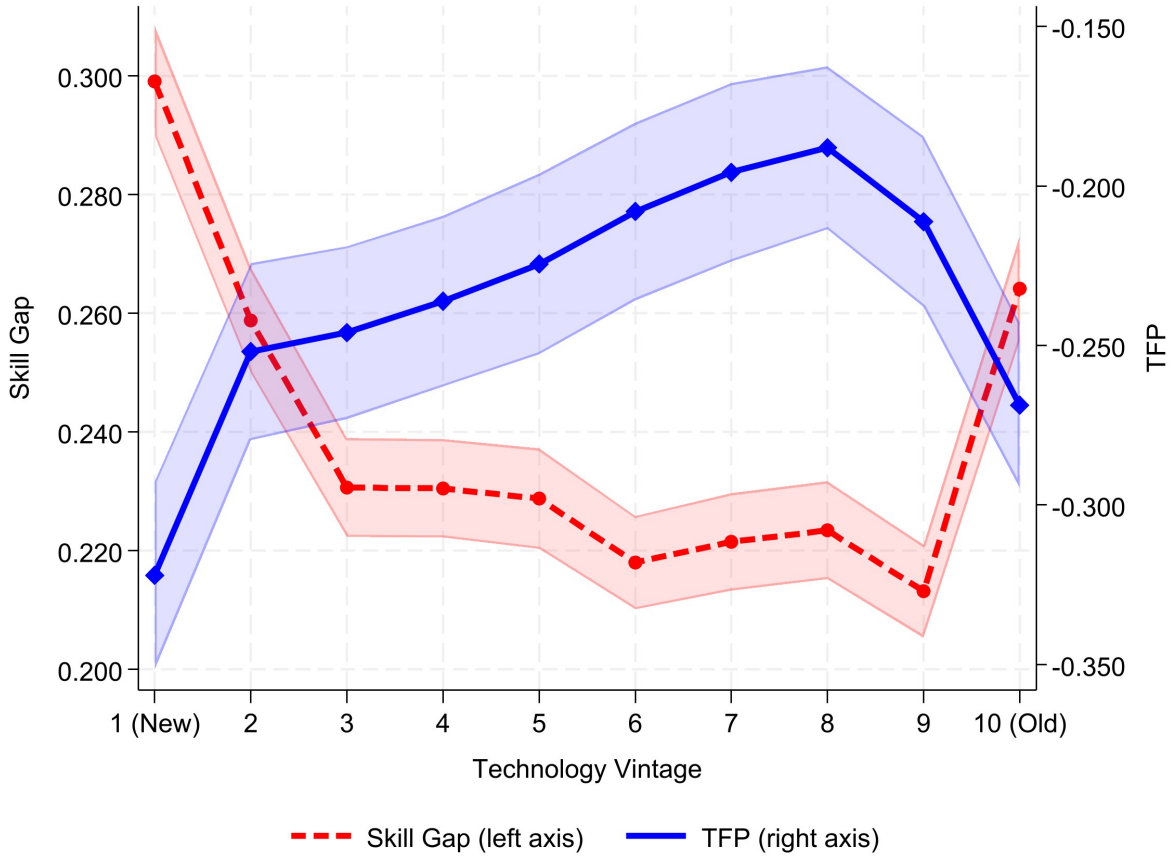
This figure presents a histogram of information technology vintages observed for firms in our sample. Each data point corresponds to a single firm-year observation. The x-axis depicts technology vintage intervals, where technology vintage is computed as the average age of all information technologies listed in the firm's job postings in a given year; age is calculated as the difference between the year of observation and the year of commercial introduction of the technology. The y-axis depicts the frequency (count) of firm-year observations in the sample that belong to each technology vintage interval.

Figure 3: Skill Vector Weights by Skill Category



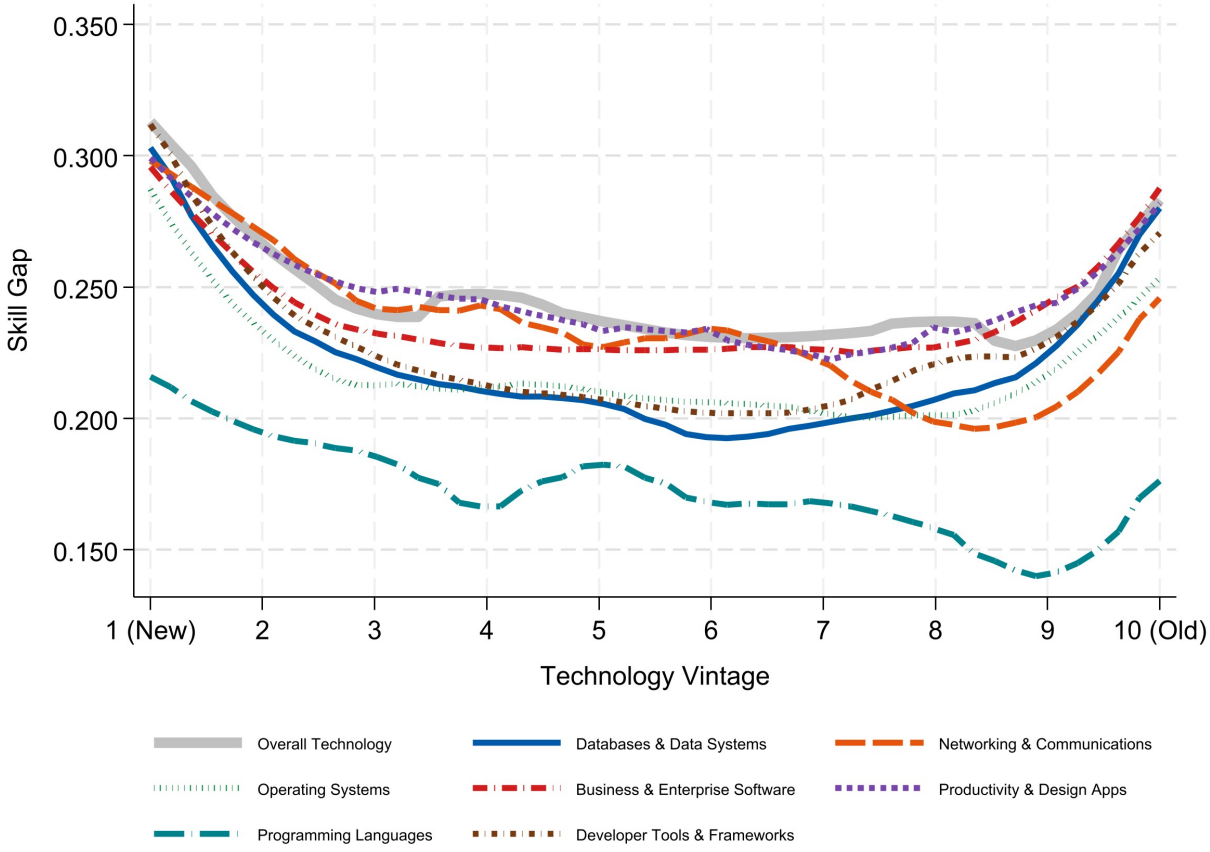
This figure presents the distribution of average skill vector weights for firm skill requirements identified in job postings (*Demand*, in blue) and average skill vector weights of employee skills identified in online résumés (*Supply*, in red) across each of the 25 skill categories used in our analysis (skill definitions provided in Appendix A2 and listed on the Y-axis). The X-axis depicts the sample average weight for each skill category.

Figure 4: Skill Gaps and Firm Productivity over the Technology Life Cycle



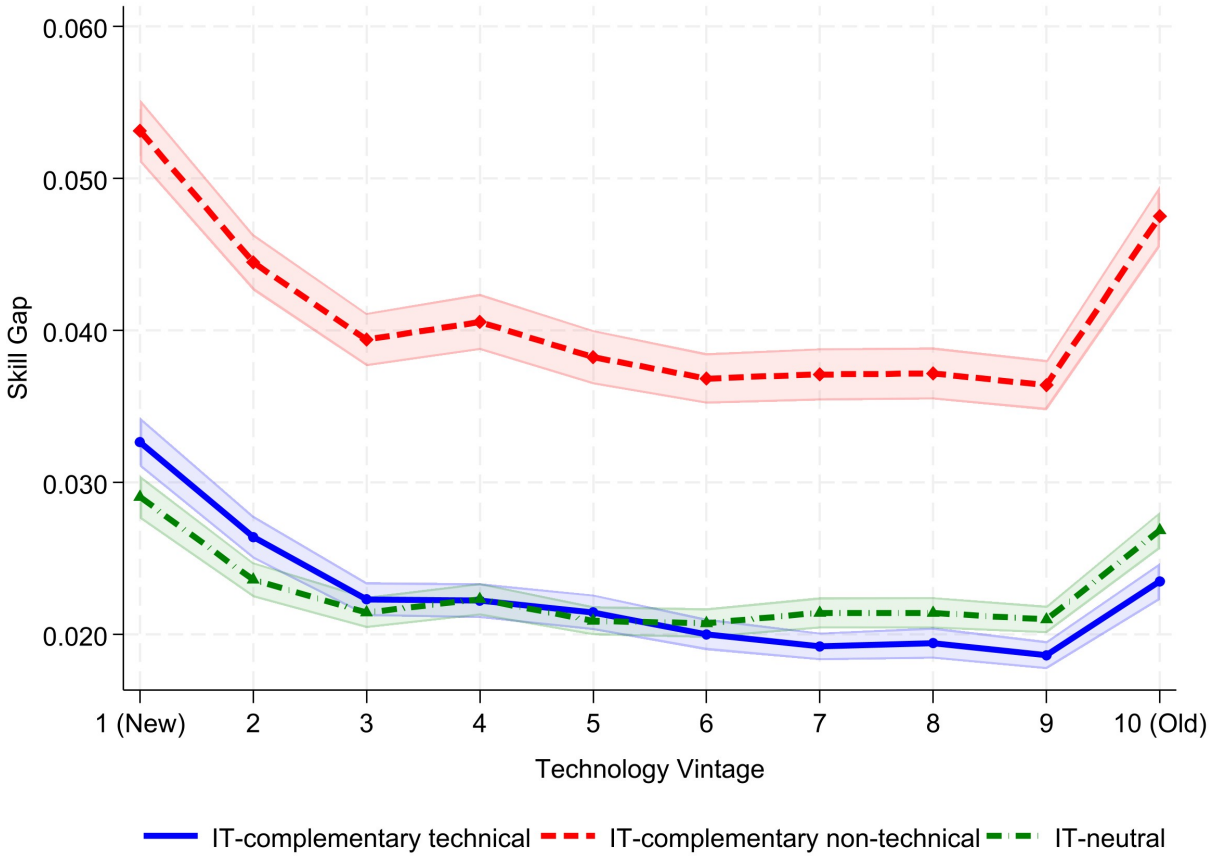
This figure depicts the empirical relationships among skill gaps, firm TFP, and information technology vintage. *Skill gap* is defined as the Euclidean distance between a firm's average skill requirements and the average skill portfolios of its new hires (matched by year and two-digit occupational codes). *TFP* estimates at the firm-year level are taken from Imrohoroglu and Tuzel (2014) and based on the semiparametric methods of Olley and Pakes (1996). *Technology Vintage* is based on the average age of the information technologies used by the firm, weighted by the number of workers matched to job postings at the occupation-firm-year level; age is calculated as the difference between the year of observation and the year of commercial introduction of the technology. Firm-year observations are grouped into ten equal-sized bins based on technology vintage intervals (in years): Bin 1 (0–23.4), Bin 2 (23.4–26.1), Bin 3 (26.1–27.8), Bin 4 (27.8–29.3), Bin 5 (29.3–30.8), Bin 6 (30.8–32.2), Bin 7 (32.2–33.8), Bin 8 (33.8–35.6), Bin 9 (35.6–38.2), and Bin 10 (38.2–91.5). 95% confidence interval bands are presented alongside point estimates.

Figure 5: Skill Gaps by Information Technology Class



This figure depicts skill gaps over the life cycle for individual classes of information technologies. *Skill gap* is defined as the Euclidean distance between a firm's average skill requirements and the average skill portfolios of its new hires (matched by year and two-digit occupational codes). *Technology Vintage* measures the relative age of a specific information technology class within the firm, weighted by the number of workers matched to job postings that contain the technology class at the occupation-firm-year level. Firm-year observations are grouped into ten equal-sized bins based on technology vintage (in years): Bin 1 (0–23.4), Bin 2 (23.4–26.1), Bin 3 (26.1–27.8), Bin 4 (27.8–29.3), Bin 5 (29.3–30.8), Bin 6 (30.8–32.2), Bin 7 (32.2–33.8), Bin 8 (33.8–35.6), Bin 9 (35.6–38.2), and Bin 10 (38.2–91.5).

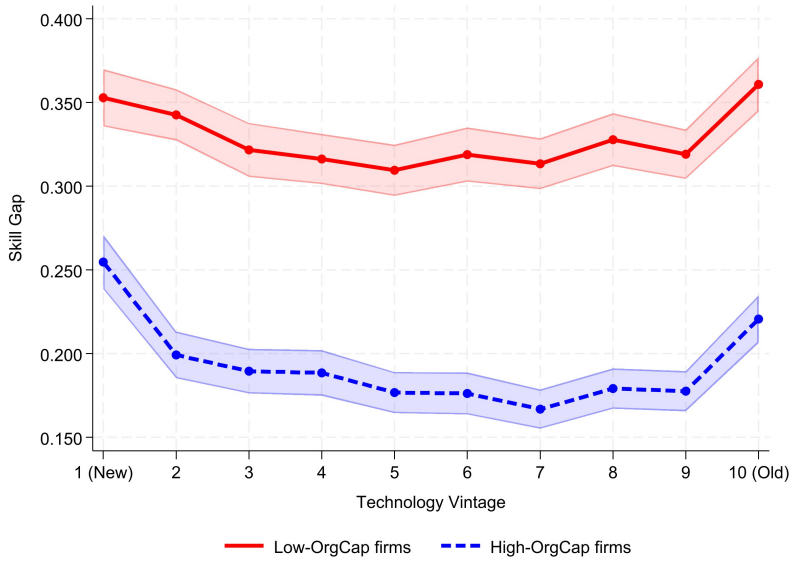
Figure 6: Skill Gaps by IT-Complementarity of Skills



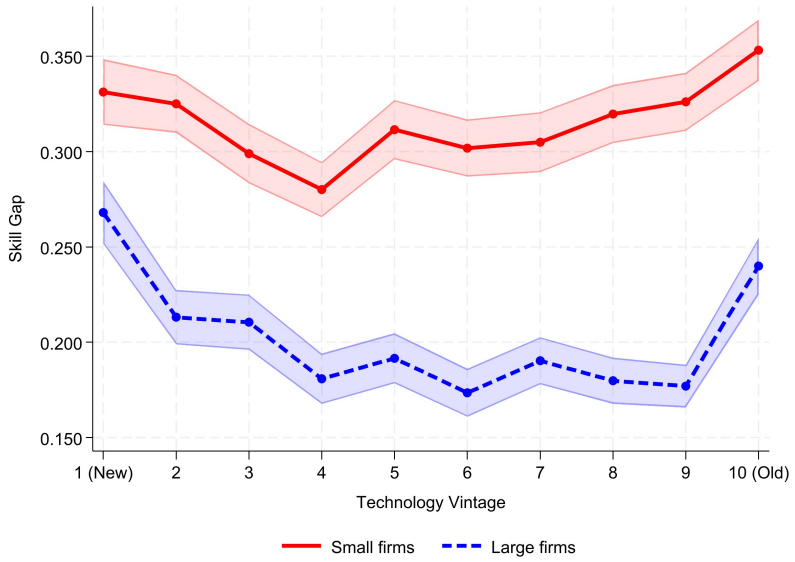
This figure presents the skill gap between firm skill requirements and worker skill attributes across three categories of aggregated skills: IT-complementary technical, IT-complementary non-technical, and IT-neutral skills. Skill definitions, groupings, and a description of the classification process are provided in Appendix A2. The Y-axis depicts the average skill gap for skills belonging to each category. *Skill gap* is defined as the Euclidean distance between a firm's average skill requirements and the average skill portfolios of its new hires (matched by year and two-digit occupational codes)—segmented by skill group. The X-axis represents information technology vintage bins as defined in Figure 4. 95% confidence intervals are presented alongside point estimates.

Figure 7: Skill Gaps by Firms Characteristics

(a) Organizational Capital



(b) Firm Size



This figure depicts average skill gaps (Y-axis) across technology vintages, for sample firms split by median firm stock of organizational capital as per Eisfeldt and Papanikolaou (2013) (Panel a) and median firm size based on number of employees (Panel b). *Skill gap* is the Euclidean distance between a firm's average skill requirements and the average skill portfolios of its new hires (matched by year and two-digit occ. codes). *Technology Vintage* on the X-axis is defined as the average age of information technologies used by the firm, weighted by the number of workers matched to job postings at the occupation-firm-year level. Firm-year observations are grouped into ten equal-sized bins based on technology vintage (in years): Bin 1 (0–23.4), Bin 2 (23.4–26.1), Bin 3 (26.1–27.8), Bin 4 (27.8–29.3), Bin 5 (29.3–30.8), Bin 6 (30.8–32.2), Bin 7 (32.2–33.8), Bin 8 (33.8–35.6), Bin 9 (35.6–38.2), and Bin 10 (38.2–91.5). 95% confidence intervals are presented alongside point estimates.

Table 1: Summary Statistics for Sample vs. Population

Variable	Observations	Mean	Std Dev	25th pct.	Median	75th pct.
Panel A: Sample						
Skill gap	20,858	0.271	0.173	0.121	0.225	0.380
Technology vintage	17,772	30.823	6.219	26.956	30.727	34.596
Sales (MM)	20,060	6,598.21	16,544.89	431.57	1,425.17	4,745.10
Firm age	20,162	22.899	11.223	14.000	23.000	36.000
Total asset (MM)	20,062	8,936.99	23,378.56	507.19	1,725.52	6,109.22
Market-to-book	18,404	2.123	1.662	1.159	1.576	2.392
Employment (MM)	19,725	0.023	0.071	0.001	0.005	0.017
Profit	20,162	0.105	0.113	0.066	0.111	0.159
TFP	14,614	-0.242	0.505	-0.485	-0.245	0.021
Panel B: Compustat Population						
Sales (MM)	124,905	2,630.27	8,451.41	6.37	126.94	1,088.52
Firm age	175,342	12.866	11.352	3.000	9.000	20.000
Total asset (MM)	125,471	7,435.26	28,560.08	35.36	375.96	2,475.32
Market-to-book	108,032	7.594	36.613	1.006	1.382	2.574
Employment (MM)	104,029	0.010	0.047	0.000	0.001	0.004
Profit	175,342	-0.238	1.347	-0.003	0.000	0.080
TFP	33,953	-0.317	0.576	-0.580	-0.296	-0.011

Notes: This table presents summary statistics that describe the firms in our sample, and for comparison, the broader Compustat population. The sample period runs from year 2010 through year 2024; *Observations* refers to the number of firm-year observations. *Skill gap* is defined as the Euclidean distance between a firm's average skill requirements and the average skill portfolios of its new hires (matched by year and two-digit occupational codes). *Technology vintage* is the average age of information technologies that appear in all firm-year job postings, weighted by the number of new hires matched to the job postings' firm-year-occupation codes. *TFP* estimates at the firm-year level are taken from Imrohoroglu and Tuzel (2014) and based on the semiparametric methods of Olley and Pakes (1996). Panel A shows statistics for firm-year observations 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 levels. Variables are winsorized at the 1st and 99th percentiles.

Table 2: Skill Gaps over the Technology Life Cycle

	(1)	(2)	(3)	(4)	(5)	(6)
Vintage	-2.590***		-2.570***		-0.835***	
	(0.174)		(0.169)		(0.150)	
Vintage^2	3.923***		3.897***		1.301***	
	(0.269)		(0.254)		(0.225)	
Vintage Bin:						
2		-0.048***		-0.050***		-0.021***
		(0.007)		(0.006)		(0.005)
3		-0.072***		-0.073***		-0.029***
		(0.007)		(0.007)		(0.006)
4		-0.065***		-0.066***		-0.023***
		(0.007)		(0.007)		(0.006)
5		-0.075***		-0.075***		-0.033***
		(0.007)		(0.007)		(0.006)
6		-0.083***		-0.082***		-0.034***
		(0.007)		(0.007)		(0.006)
7		-0.082***		-0.080***		-0.027***
		(0.007)		(0.008)		(0.006)
8		-0.076***		-0.075***		-0.030***
		(0.007)		(0.008)		(0.006)
9		-0.082***		-0.082***		-0.029***
		(0.008)		(0.008)		(0.007)
10		-0.031***		-0.034***		-0.015**
		(0.008)		(0.009)		(0.007)
Constant	0.663***	0.314***	0.659***	0.314***	0.380***	0.275***
	(0.028)	(0.006)	(0.028)	(0.006)	(0.0248)	(0.005)
Observations	17,772	17,772	17,752	17,752	17,589	17,589
R-squared	0.027	0.025	0.082	0.080	0.550	0.551
Year FE	No	No	Yes	Yes	Yes	Yes
Industry/Firm FE	No	No	Industry	Industry	Firm	Firm

Notes: This table presents regression coefficient estimates of skill gaps as a function of firm technology vintage. *Skill gap* is defined as the Euclidean distance between a firm's average skill requirements and the average skill portfolios of its new hires (matched by year and two-digit occupational codes). *Vintage* is the average age of firm information technologies at the firm-year level (weighted by the number of workers matched to job postings based on 2-digit occupational codes), and 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 4). Fixed effects for the year of observation, industry (2-digit SIC code), and firm are included where specified. Standard errors are heteroskedasticity robust and clustered by firm.

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

Table 3: Skill Gaps by Information Technology Class

Vintage Bin	Database (1)	Networks (2)	Oper Sys (3)	ERP (4)	Graphics (5)	Prog. (6)	Software (7)
2	-0.021***	-0.029***	-0.020***	-0.014***	-0.002	0.004	-0.030***
	(0.004)	(0.005)	(0.004)	(0.005)	(0.005)	(0.004)	(0.005)
3	-0.014***	-0.032***	-0.008*	-0.019***	-0.004	0.008*	-0.033***
	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)	(0.005)
4	-0.002	-0.025***	0.003	-0.031***	-0.011**	0.023***	-0.027***
	(0.005)	(0.005)	(0.005)	(0.004)	(0.005)	(0.006)	(0.005)
5	0.004	-0.017***	0.010*	-0.030***	-0.007	0.033***	-0.014***
	(0.005)	(0.005)	(0.005)	(0.004)	(0.005)	(0.006)	(0.005)
6	0.015***	0.004	0.019***	-0.013***	-0.022***	0.038***	0.001
	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)	(0.006)	(0.005)
7	0.028***	0.012**	0.026***	-0.001	-0.020***	0.049***	0.012**
	(0.005)	(0.005)	(0.006)	(0.005)	(0.005)	(0.006)	(0.006)
8	0.033***	0.025***	0.029***	0.014***	-0.004	0.038***	0.022***
	(0.006)	(0.005)	(0.006)	(0.005)	(0.005)	(0.006)	(0.006)
9	0.039***	0.030***	0.054***	0.019***	0.021***	0.044***	0.027***
	(0.006)	(0.006)	(0.007)	(0.005)	(0.005)	(0.006)	(0.006)
10	0.063***	0.051***	0.066***	0.025***	0.029***	0.045***	0.055***
	(0.007)	(0.006)	(0.007)	(0.006)	(0.006)	(0.007)	(0.007)
Constant	0.237***	0.249***	0.234***	0.256***	0.253***	0.223***	0.250***
	(0.004)	(0.004)	(0.004)	(0.003)	(0.004)	(0.004)	(0.004)
Observations	17,589	17,589	17,589	17,589	17,589	17,589	17,589
R-squared	0.559	0.561	0.559	0.556	0.554	0.553	0.561
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table presents regression coefficient estimates of skill gaps as a function of firm technology vintage, across the following individual classes of information technologies: (1): Database, Storage & Data Engineering Systems (2): Computer Networking & Communications Software (3): Operating Systems & System Software (4): Enterprise & Business Management Systems (5): Productivity, Graphics & User Applications (6): Programming Languages & Execution Runtimes (7): Software Frameworks, Tools & Utilities. *Skill gap* is defined as the Euclidean distance between a firm's average skill requirements and the average skill portfolios of its new hires (matched by year and two-digit occupational codes) – for positions where the job posting contains the information technology class. *Vintage bins* correspond to deciles of firm information technology vintages by individual classes of technologies. Firm and year fixed effects are included. Standard errors are heteroskedasticity robust and clustered by firm.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Skill Gaps by IT-complementarity of skills

	Technical skills		Non-technical skills		Neutral skills	
	(1)	(2)	(3)	(4)	(5)	(6)
Vintage	-0.099***		-0.198***		-0.070**	
	(0.027)		(0.035)		(0.027)	
Vintage^2	0.158***		0.319***		0.100**	
	(0.041)		(0.054)		(0.041)	
Vintage Bin:						
2		-0.002*		-0.004***		-0.003***
		(0.001)		(0.001)		(0.001)
3		-0.003***		-0.005***		-0.004***
		(0.001)		(0.001)		(0.001)
4		-0.003***		-0.004***		-0.003**
		(0.001)		(0.001)		(0.001)
5		-0.002**		-0.007***		-0.005***
		(0.001)		(0.001)		(0.001)
6		-0.003**		-0.007***		-0.005***
		(0.001)		(0.001)		(0.001)
7		-0.002		-0.005***		-0.004***
		(0.001)		(0.001)		(0.001)
8		-0.002		-0.006***		-0.004***
		(0.001)		(0.001)		(0.001)
9		-0.002		-0.005***		-0.004***
		(0.001)		(0.002)		(0.001)
10		0.000		-0.002		-0.003**
		(0.001)		(0.002)		(0.001)
Constant	0.036***	0.023***	0.068***	0.043***	0.035***	0.027***
	(0.004)	(0.001)	(0.006)	(0.001)	(0.005)	(0.001)
Observations	17,589	17,589	17,589	17,589	17,589	17,589
R-squared	0.366	0.366	0.470	0.470	0.413	0.414
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table presents regression coefficient estimates of skill gaps across specific subgroups of individual skills, as a function of technology vintage. Columns (1) and (2) depict regression coefficients for IT-complementary technical skill gaps. Columns (3) and (4) depict regression coefficients for IT-complementary non-technical skill gaps. Columns (5) and (6) depict regression coefficients for IT-neutral skill gaps. *Skill gap* for each subgroup of skills is defined as the Euclidean distance between a firm's average skill requirements and the average skill portfolios of its new hires (matched by year and two-digit occupational codes) for all skills within a particular subgroup. Skill classifications into subgroups are provided in Appendix A2. *Vintage* is the average age of firm information technologies at the firm-year level (weighted by the number of workers matched to job postings based on 2-digit occupational codes), and 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 4). Year and firm fixed effects are included where specified. Standard errors are heteroskedasticity robust and clustered by firm.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Skill Gaps over the Technology Life Cycle: Occupation-Level Analysis

SOC (Occupation)	Constant	Constant SE	Vintage	Vintage SE	Vintage^2	Vintage^2 SE	Obs.	R-squared
Panel A: Occupations with Positive and Significant Skill Mismatch Curvature								
11 (Management)	0.442***	0.044	-1.417***	0.270	2.127***	0.412	15,221	0.454
13 (Bus. & Finance)	0.461***	0.049	-1.434***	0.302	2.229***	0.464	14,403	0.456
15 (Computer & Math)	0.362***	0.046	-0.876***	0.279	1.271***	0.426	13,752	0.487
17 (Engineering)	0.420***	0.065	-0.897**	0.393	1.177**	0.593	10,397	0.512
31 (Healthcare Support)	0.640***	0.151	-1.669*	0.872	2.438*	1.283	1,597	0.565
37 (Building Maint.)	0.673***	0.122	-1.474**	0.741	2.070*	1.132	2,491	0.501
41 (Sales)	0.370***	0.047	-0.797***	0.282	1.198***	0.423	12,184	0.505
43 (Admin Support)	0.411***	0.048	-0.992***	0.291	1.512***	0.435	13,788	0.461
51 (Production)	0.425***	0.062	-0.780**	0.362	1.008*	0.533	9,651	0.511
Panel B: Occupations with Negative or Insignificant Skill Mismatch Curvature								
19 (Life Science)	0.399***	0.084	-0.368	0.490	0.357	0.717	7,506	0.463
21 (Social Service)	0.435***	0.119	-0.490	0.706	1.003	1.054	3,212	0.499
23 (Legal)	0.535***	0.119	-0.903	0.716	1.447	1.074	5,152	0.482
25 (Education)	0.341**	0.141	0.121	0.851	0.003	1.271	4,410	0.423
27 (Arts & Media)	0.471***	0.081	-0.804	0.495	1.011	0.754	9,699	0.453
29 (Healthcare)	0.331***	0.090	-0.016	0.536	0.061	0.801	5,209	0.553
33 (Protective Service)	0.405***	0.111	-0.373	0.642	0.658	0.924	3,359	0.512
35 (Food Service)	0.498***	0.124	-0.709	0.688	1.159	0.945	3,658	0.492
39 (Personal Care)	0.356***	0.117	-0.045	0.654	0.151	0.913	2,376	0.492
45 (Farming & Forestry)	0.100	0.326	1.326	1.799	-1.074	2.492	527	0.521
47 (Construction)	0.167	0.110	1.153*	0.645	-1.892**	0.950	4,792	0.474
49 (Installation)	0.299***	0.072	-0.045	0.413	0.141	0.593	8,154	0.528
53 (Transportation)	0.328***	0.066	-0.241	0.383	0.332	0.559	8,462	0.539

Notes: This table presents regression coefficient estimates of skill gaps as a quadratic function of technology vintage for each 2-digit SOC occupational code. *Skill gap* is defined as the Euclidean distance between a firm's average skill requirements and the average skill portfolios of its new hires (matched by year and two-digit occupational codes). *Vintage* is the average age of firm information technologies at the firm-year level (weighted by the number of workers matched to job postings based on 2-digit occupational codes), and normalized by 100; *Vintage*^{2} is also included. Panel A (B) reports occupations where the quadratic term is positive and statistically significant (negative and/or insignificant). Standard errors (SE) are clustered by firm.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table 6: Skill Gaps over the Technology Life Cycle: Firm Heterogeneity

Vintage Bin	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Organizational Capital						
	Low	High	Low-High	Low	High	Low-High
2	-0.010	-0.056***	0.045**	-0.011	-0.050***	0.039**
	(0.013)	(0.013)	(0.018)	(0.012)	(0.012)	(0.017)
3	-0.031**	-0.065***	0.034*	-0.033**	-0.060***	0.026
	(0.014)	(0.013)	(0.019)	(0.013)	(0.012)	(0.018)
4	-0.037***	-0.066***	0.030	-0.041***	-0.060***	0.019
	(0.014)	(0.013)	(0.019)	(0.013)	(0.013)	(0.018)
5	-0.043***	-0.078***	0.035*	-0.053***	-0.076***	0.022
	(0.014)	(0.013)	(0.019)	(0.014)	(0.013)	(0.019)
6	-0.034**	-0.078***	0.045**	-0.046***	-0.066***	0.020
	(0.014)	(0.014)	(0.020)	(0.014)	(0.014)	(0.020)
7	-0.039***	-0.088***	0.048**	-0.054***	-0.078***	0.025
	(0.014)	(0.014)	(0.019)	(0.014)	(0.014)	(0.020)
8	-0.025*	-0.076***	0.051**	-0.037**	-0.066***	0.029
	(0.014)	(0.014)	(0.020)	(0.015)	(0.015)	(0.021)
9	-0.034**	-0.077***	0.043**	-0.052***	-0.064***	0.012
	(0.014)	(0.014)	(0.020)	(0.015)	(0.016)	(0.022)
10	0.008	-0.034**	0.042*	-0.011	-0.021	0.010
	(0.016)	(0.015)	(0.022)	(0.016)	(0.017)	(0.023)
Constant	0.353***	0.255***	0.098***	0.362***	0.247***	-
	(0.011)	(0.012)	(0.016)	(0.011)	(0.012)	
Observations	4,386	4,386	8,772	4,385	4,386	8,771
R-squared	0.010	0.031	0.173	0.108	0.109	0.249
Fixed effects	No	No	No	Yes	Yes	Yes
Panel B: Firm size						
	Small	Large	Small-Large	Small	Large	Small-Large
2	-0.006	-0.055***	0.049***	-0.004	-0.045***	0.041**
	(0.013)	(0.014)	(0.019)	(0.012)	(0.013)	(0.018)
3	-0.032**	-0.058***	0.025	-0.034***	-0.051***	0.018
	(0.014)	(0.014)	(0.019)	(0.013)	(0.013)	(0.018)
4	-0.051***	-0.087***	0.036*	-0.050***	-0.081***	0.031
	(0.014)	(0.015)	(0.020)	(0.013)	(0.014)	(0.019)
5	-0.020	-0.076***	0.057***	-0.029**	-0.070***	0.040**
	(0.014)	(0.015)	(0.020)	(0.014)	(0.015)	(0.020)
6	-0.029**	-0.094***	0.065***	-0.036***	-0.085***	0.049**
	(0.014)	(0.015)	(0.020)	(0.013)	(0.015)	(0.020)
7	-0.026*	-0.078***	0.051**	-0.038***	-0.078***	0.040*
	(0.014)	(0.015)	(0.021)	(0.014)	(0.016)	(0.021)
8	-0.012	-0.088***	0.077***	-0.027**	-0.082***	0.054**
	(0.014)	(0.015)	(0.020)	(0.014)	(0.016)	(0.021)

Table 6: Skill Gaps over the Technology Life Cycle: Firm Heterogeneity

Vintage Bin	(1)	(2)	(3)	(4)	(5)	(6)
9	-0.005	-0.091***	0.086***	-0.024*	-0.086***	0.061***
	(0.014)	(0.015)	(0.021)	(0.015)	(0.017)	(0.022)
10	0.022	-0.028*	0.050**	-0.002	-0.031*	0.028
	(0.015)	(0.017)	(0.022)	(0.015)	(0.018)	(0.024)
Constant	0.331***	0.268***	0.063***	0.340***	0.263***	-
	(0.011)	(0.014)	(0.017)	(0.011)	(0.014)	
Observations	4,380	4,380	8,760	4,380	4,380	8,760
R-squared	0.013	0.039	0.135	0.104	0.139	0.219
Fixed effects	No	No	No	Yes	Yes	Yes

Notes: This table presents regression coefficient estimates of skill gaps as a function of technology vintage bins, across firms bifurcated by median organizational capital stock (Panel A) and median size (Panel B). Firm size is defined in terms of employees and is taken from Compustat; organizational capital stocks are measured using Compustat data following Eisfeldt and Papanikolaou (2013). Indicators of technology vintage bins are described in Figure 4. Columns (1) and (2) and Columns (4) and (5) in each panel depict regression coefficients for sample firms that are below vs. above the sample median for the relevant firm characteristic. Columns (3) and (6) of each panel depict the difference in coefficients between the preceding two columns. Year and 2-digit industry fixed effects are included where specified. Standard errors are heteroskedasticity robust and clustered by firm.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Firm Total Factor Productivity (TFP) over the Technology Life Cycle

	(1)	(2)	(3)	(4)	(5)	(6)
Skill Gap	-0.343***	-0.341***	-0.339***	-0.368***	-0.374***	-0.371***
	(0.0471)	(0.0525)	(0.0526)	(0.0469)	(0.0522)	(0.0523)
Vintage		2.453***			2.486***	
		(0.663)			(0.687)	
Vintage^2		-3.513***			-3.978***	
		(0.996)			(1.024)	
Vintage Bin:						
2			0.0611**			0.0655**
			(0.0268)			(0.0267)
3			0.0594**			0.0599**
			(0.0281)			(0.0280)
4			0.0677**			0.0661**
			(0.0281)			(0.0284)
5			0.0787***			0.0751**
			(0.0301)			(0.0307)
6			0.100***			0.0900***
			(0.0302)			(0.0308)
7			0.102***			0.0832***
			(0.0302)			(0.0313)
8			0.112***			0.0810**
			(0.0293)			(0.0322)
9			0.0722**			0.0270
			(0.0312)			(0.0344)
10			0.0502			-0.00726
			(0.0321)			(0.0352)
Constant	-0.154***	-0.564***	-0.226***	-0.148***	-0.520***	-0.203***
	(0.0164)	(0.112)	(0.0302)	(0.0162)	(0.117)	(0.0306)
Observations	14,614	13,084	13,084	14,613	13,083	13,083
R-squared	0.013	0.016	0.017	0.048	0.050	0.052
Year FE	No	No	No	Yes	Yes	Yes
Industry/Firm FE	No	No	No	Industry	Industry	Industry

Notes: This table presents regression coefficient estimates of firm total factor productivity (TFP) as a function of skill gaps and technology vintage. *TFP* estimates at the firm-year level are taken from Imrohoroglu and Tuzel (2014) and based on the semiparametric methods of Olley and Pakes (1996). *Skill gap* is defined as the Euclidean distance between a firm's average skill requirements and the average skill portfolios of its new hires (matched by year and two-digit occupational codes). *Vintage* is the average age of firm information technologies at the firm-year level (weighted by the number of workers matched to job postings based on 2-digit occupational codes), and normalized by 100, with a quadratic term included. Indicators of technology vintage bins are described in Figure 4. Year and industry fixed effects are included where specified. Standard errors are heteroskedasticity robust and clustered by firm.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8: Skill Gaps over the Technology Life Cycle: Alternate Firm-Worker Matching

	Across occupation			Within occupation		
	(1)	(2)	(3)	(4)	(5)	(6)
Vintage	-2.540***			-1.477***		
	(0.166)			(0.071)		
Vintage^2	3.663***			2.068***		
	(0.263)			(0.100)		
Vintage Bin:						
2		-0.044***	-0.012**		-0.043***	-0.014***
		(0.006)	(0.006)		(0.003)	(0.003)
3		-0.062***	-0.025***		-0.053***	-0.013***
		(0.006)	(0.006)		(0.004)	(0.003)
4		-0.062***	-0.022***		-0.071***	-0.021***
		(0.006)	(0.006)		(0.004)	(0.003)
5		-0.065***	-0.024***		-0.066***	-0.019***
		(0.007)	(0.006)		(0.004)	(0.003)
6		-0.065***	-0.021***		-0.061***	-0.016***
		(0.007)	(0.006)		(0.004)	(0.003)
7		-0.065***	-0.026***		-0.078***	-0.021***
		(0.007)	(0.006)		(0.004)	(0.003)
8		-0.073***	-0.025***		-0.066***	-0.018***
		(0.007)	(0.007)		(0.004)	(0.003)
9		-0.066***	-0.021***		-0.067***	-0.016***
		(0.008)	(0.007)		(0.005)	(0.003)
10		-0.033***	-0.013*		-0.027***	-0.006*
		(0.008)	(0.007)		(0.005)	(0.003)
Constant	0.559***	0.192***	0.155***	0.493***	0.298***	0.253***
	(0.026)	(0.006)	(0.005)	(0.012)	(0.003)	(0.002)
Observations	17,998	17,978	17,800	119,756	119,643	115,584
R-squared	0.034	0.056	0.439	0.012	0.029	0.515
Year FE	No	Yes	Yes	No	Yes	Yes
Industry/Firm FE	No	Industry	Firm	No	Industry	Firm X Occ

Notes: This table presents regression coefficient estimates of skill gaps as a function of firm technology vintage, using alternative methods of matching firms and workers to construct skill gaps. Columns (1) to (3) use skill gap measures constructed by matching firm job postings with new employees across all occupations at the *firm-year* level. Columns (4) to (6) use skill gap measures constructed by matching firm job postings with new employees at the *firm-occupation-year* level; skill gaps are at the disaggregated occupation-firm-year level (i.e. skill gaps are not aggregated to the firm-year level). Vintage is the average age of firm information technologies at the firm-year level (weighted by the number of workers matched to job postings based on 2-digit occupational codes), and normalized by 100, with a quadratic term included in Columns (1) and (4). The remaining columns depict piece-wise linear regression coefficients for indicators of technology vintage bins (described in Figure 4). Fixed effects for the year of observation, industry (2-digit SIC code), firm, and firm by occupation pairs, are included where specified. Standard errors are heteroskedasticity robust and clustered by firm.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 9: Skill Gaps over the Technology Life Cycle: CPS-weighted

	(1)	(2)	(3)	(4)	(5)	(6)
Vintage	-1.474***		-1.506***		-0.403***	
	(0.147)		(0.144)		(0.119)	
Vintage^2	1.930***		2.016***		0.593***	
	(0.223)		(0.212)		(0.173)	
Vintage Bin:						
2		-0.028***		-0.030***		-0.005
		(0.006)		(0.006)		(0.005)
3		-0.048***		-0.049***		-0.016***
		(0.007)		(0.007)		(0.005)
4		-0.045***		-0.045***		-0.010*
		(0.007)		(0.007)		(0.006)
5		-0.055***		-0.054***		-0.013**
		(0.007)		(0.007)		(0.006)
6		-0.060***		-0.057***		-0.015***
		(0.007)		(0.007)		(0.006)
7		-0.070***		-0.068***		-0.018***
		(0.007)		(0.007)		(0.006)
8		-0.071***		-0.071***		-0.017***
		(0.007)		(0.008)		(0.006)
9		-0.077***		-0.078***		-0.016**
		(0.007)		(0.008)		(0.006)
10		-0.053***		-0.052***		-0.010
		(0.008)		(0.008)		(0.006)
Constant	0.549***	0.336***	0.550***	0.335***	0.349***	0.296***
	(0.024)	(0.006)	(0.024)	(0.006)	(0.020)	(0.004)
Observations	17,923	17,923	17,903	17,903	17,752	17,752
R-squared	0.017	0.018	0.082	0.082	0.545	0.545
Year FE	No	No	Yes	Yes	Yes	Yes
Industry/Firm FE	No	No	Industry	Industry	Firm	Firm

Notes: This table presents regression coefficient estimates of skill gaps as a function of firm technology vintage, where the number of new hires at a given firm is weighted by Current Population Survey (CPS) industry and occupation population weights. *Skill gap* is defined as the Euclidean distance between a firm's average skill requirements and the average skill portfolios of its CPS-weighted new hires (matched by year and two-digit occupational codes). *Vintage* is the average age of firm information technologies at the firm-year level (weighted by the number of workers matched to job postings based on 2-digit occupational codes), and 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 4). Fixed effects for the year of observation, industry (2-digit SIC code), and firm are included where specified. Standard errors are heteroskedasticity robust and clustered by firm.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 10: Skill Gaps over the Technology Life Cycle: Alternative Temporal Matching

	Within same year			Within two years		
	(1)	(2)	(3)	(4)	(5)	(6)
Vintage	-2.603***			-2.444***		
	(0.165)			(0.193)		
Vintage^2	4.002***			3.614***		
	(0.259)			(0.296)		
Vintage Bin:						
2		-0.061***	-0.030***		-0.046***	-0.013**
		(0.006)	(0.005)		(0.007)	(0.006)
3		-0.078***	-0.037***		-0.052***	-0.018***
		(0.007)	(0.005)		(0.007)	(0.006)
4		-0.079***	-0.035***		-0.057***	-0.021***
		(0.007)	(0.005)		(0.007)	(0.006)
5		-0.080***	-0.035***		-0.067***	-0.029***
		(0.007)	(0.006)		(0.008)	(0.006)
6		-0.086***	-0.037***		-0.075***	-0.031***
		(0.007)	(0.006)		(0.008)	(0.006)
7		-0.091***	-0.035***		-0.068***	-0.025***
		(0.007)	(0.006)		(0.008)	(0.007)
8		-0.085***	-0.034***		-0.072***	-0.025***
		(0.008)	(0.006)		(0.008)	(0.007)
9		-0.084***	-0.030***		-0.070***	-0.023***
		(0.008)	(0.006)		(0.009)	(0.007)
10		-0.036***	-0.018**		-0.031***	-0.018**
		(0.009)	(0.007)		(0.009)	(0.008)
Constant	0.660***	0.323***	0.283***	0.654***	0.310***	0.275***
	(0.026)	(0.006)	(0.005)	(0.031)	(0.007)	(0.005)
Observations	19,136	19,114	19,037	15,968	15,952	15,800
R-squared	0.030	0.087	0.577	0.022	0.076	0.546
Year FE	No	No	Yes	No	Yes	Yes
Industry/Firm FE	No	No	Firm	No	Industry	Firm

Notes: This table presents regression coefficient estimates of skill gaps as a function of firm technology vintage, using skill gap measures constructed under alternative time-series assumptions for matching firm job postings with new employees. Columns (1) to (3) depict the results for skill gap measures that are constructed by matching firm job postings with new employees of the same 2-digit occupation code that join the firm in the same year of the job posting. Columns (4) to (6) depict the results for skill gap measures that are constructed by matching firm job postings with new employees of the same 2-digit occupation code that join the firm within 2 years of the job posting.*Skill gap* is the Euclidean distance between a firm's average skill requirements and the average skill portfolios of its new hires. *Vintage* is the average age of firm information technologies at the firm-year level (weighted by the number of workers matched to job postings based on 2-digit occupational codes), and 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 4). Fixed effects for the year of observation, industry (2-digit SIC code), and firm are included where specified. Standard errors are heteroskedasticity robust and clustered by firm.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 11: Skill Gaps over the Technology Life Cycle: Year Subsamples

Vintage Bin	2024 (1)	2023 to 2024 (2)	2022 to 2024 (3)	2021 to 2024 (4)	2020 to 2024 (5)
2	-0.033	-0.022*	-0.034***	-0.047***	-0.041***
	(0.022)	(0.013)	(0.011)	(0.009)	(0.008)
3	-0.032	-0.030**	-0.027**	-0.042***	-0.042***
	(0.023)	(0.013)	(0.011)	(0.009)	(0.008)
4	-0.045**	-0.023*	-0.026**	-0.045***	-0.042***
	(0.022)	(0.013)	(0.011)	(0.009)	(0.008)
5	-0.044**	-0.041***	-0.050***	-0.058***	-0.043***
	(0.022)	(0.013)	(0.010)	(0.009)	(0.008)
6	-0.058***	-0.039***	-0.050***	-0.064***	-0.055***
	(0.021)	(0.013)	(0.010)	(0.009)	(0.008)
7	-0.039*	-0.033***	-0.039***	-0.058***	-0.058***
	(0.021)	(0.013)	(0.010)	(0.009)	(0.008)
8	-0.054***	-0.043***	-0.046***	-0.052***	-0.051***
	(0.021)	(0.013)	(0.010)	(0.009)	(0.007)
9	-0.028	-0.028**	-0.041***	-0.058***	-0.047***
	(0.021)	(0.013)	(0.010)	(0.009)	(0.008)
10	0.008	0.022*	0.011	0.003	-0.001
	(0.022)	(0.013)	(0.010)	(0.009)	(0.008)
Constant	0.166***	0.143***	0.149***	0.162***	0.158***
	(0.016)	(0.010)	(0.008)	(0.007)	(0.006)
Observations	1,086	2,329	3,656	5,028	6,445
R-squared	0.018	0.021	0.022	0.028	0.021

Notes: This table presents regression coefficient estimates of skill gaps as a function of firm technology vintage bins for subsamples restricted to observations that appear in specific years. The sample years that correspond to the columns are the following: (1) 2024, (2) 2023-2024, (3) 2022-2024, (4) 2021-2024, (5) 2020-2024. *Skill gap* is defined as the Euclidean distance between a firm's average skill requirements and the average skill portfolios of its new hires. *Vintage Bin* corresponds to piece-wise linear regression coefficients for indicators of technology vintage bins (described in Figure 4). Standard errors are heteroskedasticity robust and clustered by firm.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 12: Skill Gaps over the Technology Life Cycle: Robustness Analyses

Vintage Bin	Benchmark (1)	10-skill class. (2)	ModernBERT (3)	Cosine Dist. (4)	Hellinger Dist. (5)
2	-0.011	0.012	-0.008	0.015	-0.016
	(0.017)	(0.016)	(0.011)	(0.021)	(0.016)
3	-0.025	-0.018	-0.014	0.035*	-0.028*
	(0.016)	(0.016)	(0.011)	(0.021)	(0.016)
4	-0.019	-0.006	-0.012	0.029	-0.024
	(0.017)	(0.016)	(0.011)	(0.021)	(0.016)
5	-0.036**	-0.030*	-0.029***	0.056***	-0.043***
	(0.016)	(0.016)	(0.010)	(0.021)	(0.016)
6	-0.033**	-0.026*	-0.027***	0.049**	-0.042***
	(0.016)	(0.016)	(0.010)	(0.021)	(0.016)
7	-0.031*	-0.021	-0.020*	0.048**	-0.039**
	(0.016)	(0.016)	(0.010)	(0.020)	(0.016)
8	-0.039**	-0.042***	-0.032***	0.060***	-0.051***
	(0.016)	(0.015)	(0.010)	(0.020)	(0.016)
9	-0.028*	-0.014	-0.027***	0.038*	-0.034**
	(0.016)	(0.015)	(0.010)	(0.020)	(0.016)
10	-0.005	-0.001	-0.010	0.014	-0.013
	(0.016)	(0.016)	(0.010)	(0.021)	(0.016)
Constant	0.140***	0.180***	0.275***	0.847***	0.202***
	(0.012)	(0.012)	(0.008)	(0.016)	(0.012)
Observations	1,417	1,417	1,417	1,417	1,417
R-squared	0.009	0.015	0.014	0.013	0.013

Notes: This table presents regression coefficient estimates of skill gaps as a function of technology vintage bins using a piecewise linear specification for sample data from the year 2020. Column (1) presents results using our benchmark specification depicted in Table 2. Column (2) presents results uses skill gap measures constructed using the 10-skill classification scheme defined by Deming and Kahn (2018). Column (3) uses a measure of skill gaps that is estimated using an alternative large language model—ModernBERT—instead of BERT. Column (4) uses a measure of skill gaps that is based on the cosine similarity between firm and worker skill vectors, rather than Euclidean distance (larger coefficients indicate smaller skill gaps). Column (5) presents results using skill gaps measured by the Hellinger distance between firm and worker skill vectors (larger coefficients indicate larger skill gaps). *Skill gap* is defined as the Euclidean distance between a firm's average skill requirements and the average skill portfolios of its new hires. *Vintage Bin* corresponds to piecewise linear regression coefficients for indicators of technology vintage bins (described in Figure 4). Standard errors are heteroskedasticity robust and clustered by firm.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix

A Information Technology Data

Table A1: Most Frequently Listed Information Technologies by Vintage Bin

Rank	Vintage 1: less than 23.4 years	Vintage 2: between (23.4, 26.1) years
1	Microsoft Outlook	JavaScript (Prog. Language)
2	Java (Programming Language)	Microsoft Access
3	Salesforce	HyperText Mark. Lang. (HTML)
4	Amazon Web Services	Windows Servers
5	Microsoft Azure	PostgreSQL
6	C# (Prog. Language)	Visual Basic (Prog. Language)
7	Microsoft SharePoint	SolidWorks (CAD)
Rank	Vintage 3: between (26.1, 27.8) years	Vintage 4: between (27.8, 29.3) years
1	Linux	Python (Prog. Language)
2	Microsoft Visio	Microsoft SQL Servers
3	R (Prog. Language)	IT Infrastructure Library
4	PL/SQL	Learning Mgt. Systems
5	Applicant Tracking Systems	Business Intelligence Tools
6	SAP ERP	Operational Data Store
7	Web Servers	Web Browsers
Rank	Vintage 5: between (29.3, 30.8) years	Vintage 6: between (30.8, 32.2) years
1	Microsoft Office	Microsoft PowerPoint
2	Dashboard	Adobe Illustrator
3	Perl (Prog. Language)	Transact-SQL
4	Firewall	LabVIEW
5	Adobe Photoshop	MicroStation (CAD Software)
6	PeopleSoft Applications	IT Architecture
7	Bash (Scripting Language)	Presentation Software
Rank	Vintage 7: between (32.2, 33.8) years	Vintage 8: between (33.8, 35.6) years
1	Microsoft Excel	Microsoft Word
2	C++ (Prog. Language)	Microsoft Project
3	Microsoft Windows	MATLAB
4	Field-Programmable Gate Array (FPGA)	IBM DB2
5	Customer Relationship Mgt (CRM)	Verilog
6	IBM Maximo	Primavera (Software)
7	Objective-C (Prog. Language)	OSI Models
Rank	Vintage 9: between (35.6, 38.2) years	Vintage 10: greater than 38.2 years
1	AutoCAD	SQL (Prog. Language)
2	TurboTax	SAP Applications
3	Oracle Databases	Operating Systems
4	Mac OS	App. Prog. Interface (API)
5	Bug Tracking And Management	Spreadsheets
6	Oracle Hyperion	Unix
7	QuickBooks (Accounting Software)	C (Prog. Language)

Table A2: Skill Groupings by IT-complementarity

Aggregated skill groups	Individual Baseline Skill categories
IT-Complementary Technical	Java/HTML/JS
	SQL/SAP/Database
	Data Analysis
	Technical Support
	Telecom/Cybersecurity
	Manufacturing/AutoCAD
	Engineering/Testing
	Graphic Design
IT-Complementary Non-technical	Business Analysis/Integration
	Crossfunctional Leadership
	Management
	Construction/Planning
	Marketing
	Financial Analysis/Accounting
	Operations/Retail
	Risk Mgmt/Banking/Insurance
	Sales
IT-Neutral	Customer Service
	Healthcare/Med Devices
	Pharma/Biotech
	Public Speaking/Coaching
	Hospitality/Food
	Video Production
	Event Planning/PR
	Legal Research

Notes: This table presents the 25 skill groups used in our main analysis, and their classification into IT-complementary technical, IT-complementary non-technical, and IT-neutral categories. To classify skills, we use GPT-5.6 (last accessed in September, 2026) to group baseline skill categories according to their complementarity with information technology in production. A skill is defined as IT-complementary if its marginal productivity increases with the adoption or use of IT, and IT-neutral otherwise. We further distinguish between technical and non-technical IT-complementary skills based on whether the task involves direct interaction with IT systems or supports the organizational processes that rely on IT Bresnahan et al. (2002); Bartel et al. (2007). Operationally, we assign each skill category a score based on its intensity of IT use and the extent to which its productivity depends on IT. Skills with high IT-intensity and IT-complementarity scores are classified as IT-complementary, while those with limited reliance and low complementarity scores are classified as IT-neutral.

B 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.

B.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 $W_{i,t} \geq 0$.

B.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 $\widetilde{G}_t$.

B.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.

B.4 Lagrangian

The Lagrangian is:

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

B.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:

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

B.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

$$\lambda = 0.$$

B.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

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

B.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.

B.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.

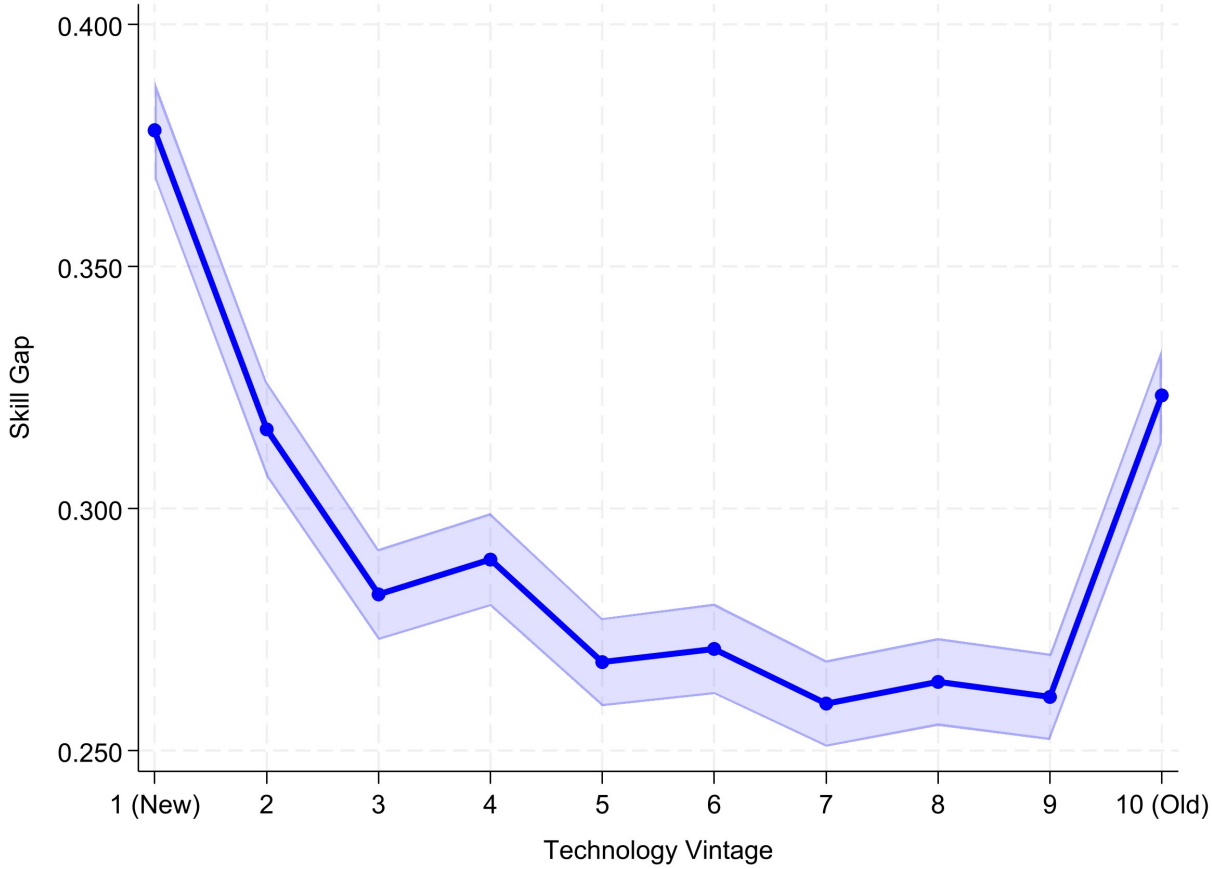
Second derivative.

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

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

C Additional Tables and Figures

Figure A1: Skill Gaps over the Technology Life Cycle: Alternative Skill Classifications



This figure depicts the empirical relationship between the firm's information technology vintage and the average skill gap between the firm and its newly hired employees, using skill vectors based on the 10-dimension skill classifications in Deming and Kahn (2018): Cognitive, Social, Character, Writing, Customer service, Project management, People management, Financial, Computer (general), and Software (specific). *Skill gap* is defined as the Euclidean distance between a firm's skill requirements and the skill portfolios of its new hires. *Technology Vintage* measures the relative age of the information technologies used by the firm, weighted by the number of workers matched to job postings at the occupation-firm-year level. Firm-year observations are grouped into ten equal-sized bins based on technology vintage (in years): Bin 1 (0–23.4), Bin 2 (23.4–26.1), Bin 3 (26.1–27.8), Bin 4 (27.8–29.3), Bin 5 (29.3–30.8), Bin 6 (30.8–32.2), Bin 7 (32.2–33.8), Bin 8 (33.8–35.6), Bin 9 (35.6–38.2), and Bin 10 (38.2–91.5). 95% confidence intervals error bars are presented around point estimates.

Table A3: Skill Gaps Across Technology Vintages - Alternative Skill Classifications

	(1)	(2)	(3)	(4)	(5)	(6)
Vintage	-3.619***		-3.508***		-1.328***	
	(0.219)		(0.208)		(0.184)	
Vintage^2	5.354***		5.299***		2.015***	
	(0.341)		(0.320)		(0.272)	
Vintage Bin #:						
2		-0.067***		-0.068***		-0.032***
		(0.008)		(0.008)		(0.007)
3		-0.097***		-0.096***		-0.044***
		(0.009)		(0.008)		(0.007)
4		-0.090***		-0.088***		-0.038***
		(0.009)		(0.009)		(0.008)
5		-0.110***		-0.105***		-0.055***
		(0.009)		(0.009)		(0.008)
6		-0.111***		-0.103***		-0.045***
		(0.009)		(0.009)		(0.008)
7		-0.121***		-0.111***		-0.042***
		(0.009)		(0.009)		(0.008)
8		-0.115***		-0.103***		-0.047***
		(0.009)		(0.010)		(0.008)
9		-0.119***		-0.105***		-0.039***
		(0.009)		(0.010)		(0.008)
10		-0.058***		-0.046***		-0.022**
		(0.010)		(0.011)		(0.009)
Constant	0.878***	0.381***	0.849***	0.375***	0.500***	0.327***
	(0.035)	(0.008)	(0.034)	(0.008)	(0.031)	(0.006)
Observations	17,777	17,777	17,757	17,757	17,593	17,593
R-squared	0.035	0.031	0.086	0.081	0.528	0.528
Year FE	No	No	Yes	Yes	Yes	Yes
Industry/Firm FE	No	No	Industry	Industry	Firm	Firm

Notes: This table presents regression coefficient estimates of skill gaps as a function of firm technology vintage, where skill gaps are estimated using skill vectors based on the 10-dimension skill classifications in Deming and Kahn (2018): Cognitive, Social, Character, Writing, Customer service, Project management, People management, Financial, Computer (general), and Software (specific). Vintage is the average age of firm information technologies at the firm-year level (weighted by the number of workers matched to job postings based on 2-digit occupational codes), and 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 4). Fixed effects for the year of observation, industry (2-digit SIC code), and firm are included where specified. Standard errors are heteroskedasticity robust and clustered by firm.

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