

Boston University School of Law

Scholarly Commons at Boston University School of Law

Faculty Scholarship

2-24-2022

The Remainder Effect: How Automation Complements Labor Quality

James Bessen

Boston University School of Law

Erich Denk

Boston University School of Law

Chen Meng

Follow this and additional works at: https://scholarship.law.bu.edu/faculty_scholarship



Part of the [Business Law, Public Responsibility, and Ethics Commons](#), [Labor and Employment Law Commons](#), and the [Law and Economics Commons](#)

Recommended Citation

James Bessen, Erich Denk & Chen Meng, *The Remainder Effect: How Automation Complements Labor Quality* (2022).

Available at: https://scholarship.law.bu.edu/faculty_scholarship/1358

This Working Paper is brought to you for free and open access by Scholarly Commons at Boston University School of Law. It has been accepted for inclusion in Faculty Scholarship by an authorized administrator of Scholarly Commons at Boston University School of Law. For more information, please contact lawlessa@bu.edu.



**THE REMAINDER EFFECT: HOW
AUTOMATION COMPLEMENTS LABOR
QUALITY**

Boston University School of Law
Research Paper Series No. 22-3

February 24, 2022

James Bessen

Boston University School of Law

Erich Denk

Boston University School of Law

Chen Meng

Kean University

The Remainder Effect: How Automation Complements Labor Quality

James Bessen, Erich Denk, and Chen Meng

2/24/2022

Abstract:

This paper argues that automation both complements and replaces workers. Extending the Acemoglu-Restrepo model of automation to consider labor quality, we obtain a Remainder Effect: while automation displaces labor on some tasks, it raises the returns to skill on remaining tasks across skill groups. This effect increases between-firm pay inequality while labor displacement affects within-firm inequality. Using job ad data, we find firm adoption of information technologies leads to both greater demand for diverse skills and higher pay across skill groups. This accounts for most of the sorting of skills to high paying firms that is central to rising inequality.

JEL: J31, O33, J23

Keywords: automation, income inequality, skills, information technology, software

Bessen and Denk, TPRI, Boston University School of Law; Meng, Kean University. Thanks to Luise Eisfeld, Maarten Goos, Po-Hsuan Hsu, Eric Maskin, Mike Meurer, Felix Pöge, Ronja Röttger, Anna Salomons and participants in TPRI's seminar for helpful comments.

Introduction

The skill-biased technical change (SBTC) hypothesis holds that technology complements some groups of workers. In contrast, recent economic models of automation posit that automation technologies strictly substitute machines for workers.¹ Labor displacement is seen by some as the main source of growing economic inequality over the last four decades (Acemoglu and Restrepo 2021), leading to calls for redistribution (Korinek and Stiglitz 2018; Benzell et al. 2016) or policies to slow the growth of automation with economic incentives or attempts to influence development engineers (Acemoglu 2021; Brynjolfsson 2021).

Yet some observers have noted that automation may also complement labor in important ways (Autor 2015; Bessen 2015). By definition, automation replaces humans with machines on certain tasks. But automation could, at the same time, complement workers on *other* tasks. This paper presents a model explaining why and how such synergy might occur and empirical evidence that it does occur. The result is a much richer picture of automation that is both cost-reducing and quality-improving, that replaces workers but also increases the demand for diverse skills in a broad range of occupations.

This depiction is important because it helps explain important features about income inequality. Indeed, labor displacement models do not address a key feature of the rise in inequality since 1980, namely that it has largely occurred *between* firms rather than within firms.² Labor displacement affects inequality because it decreases aggregate employment demand for some skill groups relative to others, thus leading to growing wage differences in equilibrium. But in these models, the market wages of different skill groups affect all firms, changing within-firm inequality.³ However, the models can be extended: if automation increases the returns to quality on non-automated tasks, then automating firms might pay more than others and hire higher quality workers, thus increasing sorting.

¹ (Autor, Levy, and Murnane 2003; Acemoglu and Autor 2011; Brynjolfsson and McAfee 2014; Acemoglu and Restrepo 2018a; 2018c; Benzell et al. 2016; Korinek and Stiglitz 2018; Hémous and Olsen 2022).

² (Card, Heining, and Kline 2013; Barth, Davis, and Freeman 2018; Song et al. 2019; Lachowska et al. 2020).

³ Because firms differ in the extent to which they employ different skill groups, these differences might secondarily affect between-firm pay, but only in dilute form. Also, firms might reorganize production in response. For instance, Song et al. (2019) argue that coincident outsourcing of low paying jobs might mask the extent of within-firm changes. We test for this explanation below.

In popular discourse, automation is about reducing costs by cutting labor, not about improving quality. But that is not necessarily the case. Automation includes robots that replace manual labor, but it also includes productivity tools such as spreadsheets (automating the calculation) used by white collar workers or AI tools that automate predictions augmenting humans.⁴ Indeed, researchers have found that advanced technologies are often directed more to improving product quality or creating new products with better quality than they are to saving cost (Brynjolfsson and Hitt 2000; Bresnahan, Brynjolfsson, and Hitt 2002; Bessen et al. 2018; Babina et al. 2020; Hirvonen, Stenhammer, and Tuhkuri 2021). Automated machines can spin finer yarn than humans, they can allow machinists and surgeons to operate at higher precision, and AI systems can make more accurate predictions.

Why is quality important for technology? Quality on complementary tasks is critical to many production processes. In Kremer's (1993) famous example, the failure of one part doomed the space shuttle *Challenger*. Poorly performed tasks can create defects, reducing the value of output, or they can halt production, slowing the rate of output, or they can reduce the reliability of the product. Yet the quality of task performance often depends critically on the quality of labor, on the ability of labor to perform specific tasks. Elon Musk's highly automated Tesla factory fell far short of production quotas because, in his words, "humans are underrated." Clark (1987), comparing workers at highly automated textile mills around the world, found six-fold differences in output per worker, even comparing workers at similar mills using identical equipment and with similar British managers. The differences lay in the varied ability and willingness of these workers to perform non-automated tasks reliably and quickly (see the example below).

This means that there are important dimensions to skill beyond educational or occupational skill groups and, for this reason, analyzing inequality with skill groups alone is inadequate.⁵ In our model, the quality of task output depends on the specific skill or effort of

⁴ And although robots have featured in recent economic papers, US investment in robots was only \$7 billion in 2019, while investment in software, studied here, was over \$400 billion (US Census).

⁵ Lindenlaub (2017) argues that multi-dimensional skills are needed to understand the link between sorting and technology. It is well-recognized that demographic skill groups are at best crude indices of the actual multiple dimensions of skill (Acemoglu 2002, Section 7). Skill groups have other limitations for the analysis of income inequality. For one thing, endogenous selection into skill groups, such as changing access to college education, means that the actual skills of demographic groups change over time. Technology also changes the skills of occupational groups over time (Autor, Levy, and Murnane 2003; Spitz-Oener 2006). Also, there is great and changing variance of wages within skill groups (Hunt and Nunn 2019).

the worker performing the task. But workers differ in their disutility of expending effort on task performance or on learning new skills. Workers with high (low) disutility comprise a low (high) skill group—these are the ones likely to get less (more) education, for instance. Hence, skill groups defined by education or occupation matter, but so, too, do task-specific skill levels. Employers can improve product quality along two margins: by hiring workers from a high skill group and by providing them stronger performance incentives. In this paper, we measure effects on both skill groups and on task-specific skills.

To model automation, we extend the Acemoglu-Restrepo model of automation (2018a; 2018c) to include variable task quality as modeled by Kremer and Maskin (Kremer 1993; Kremer and Maskin 1996). Our model provides a natural explanation for the rise in sorting of skilled workers to high-paying firms. The model generates a Remainder Effect (Bessen 2015): automating some tasks raises the demand for skill and effort on complementary non-automated tasks. The relative importance of between-firm and within-firm wage gaps corresponds to the two margins along which firms manage skill. Between-firm pay differences are driven by differences in labor quality while within-firm pay differences are driven by the relative displacement of different skill groups. This implies that the impact of automation on wage inequality—and the policies needed to counter this inequality—depend on the relative importance of labor displacement and labor quality enhancement. The model also explains why successful implementation of information technologies is linked to management practices that provide stronger incentives for performance (Bloom, Sadun, and Van Reenen 2012).

To test the model predictions, we use rich data on the skills that firms demand and the pay that they offer in online help wanted ads. The skills requested in job ads allow us to identify multiple dimensions of specific skills that employers apparently value as important for achieving quality output, for instance, affecting firm market value (D. Deming and Kahn 2018; Bana 2021). These include education and experience required, measures of cognitive and social skills, other soft skills, information technology skills, and skills related to other technologies and market knowledge.

We test how the demands for these skills are related to firm investments in own-developed software. Firms have been investing heavily in developing information technology systems for their own use, including artificial intelligence applications (AI). Self-developed and custom software grew to \$241 billion in 2020, excluding software developed for use in a

product and much of this investment is in systems that automate business processes such as enterprise resource planning. Following some literature, we measure the adoption of these technologies as the share of software developers in a firm's total hiring (Tambe and Hitt 2012; Tambe et al. 2019; Bessen 2020; Harrigan, Reshef, and Toubal 2021). The entire investment in these technology platforms includes complementary investments in hardware, packaged software, and organizational capabilities.

We first test whether skills demanded and pay increase when firms make major investments in their internal information technology. We identify these episodes as substantial increases in relative hiring of software developers, so-called investment “spikes,” and we analyze them using a difference-in-differences methodology.⁶ This method helps isolate the effects of technology from many possible confounders and we also add controls for outsourcing, labor market conditions, management changes, and productivity and demand shocks. Following a software spike, firms increase their demand for skills across all categories, both for jobs that require a college education as well as jobs that do not. Firms also significantly increase the pay they offer, after controlling for job characteristics, to most skill groups, thus increasing between-firm pay differences among new hires.

Because spiking firms are a select group, this analysis might not reflect the role of information technology in sorting more generally. We also look at the relationship between investment in own-developed software and the demand for skills in the universe of online help wanted ads. We first calculate firm pay fixed effects by regressing salaries offered against job and firm characteristics. These fixed effects are correlated with our various skill measures, indicating sorting. But firm investment in own-developed software also correlates with both the firm fixed effects and skill measures. When software is added to the regression of skill against firm fixed effects, the correlations are substantially reduced. Software accounts for most of the correlations between firm fixed effects and skills. The prominent relationship between information technology and sorting and the relatively recent shift of investment to information technology suggests that much of the rise of sorting can be accounted for by this technology.

⁶ A variety of papers have begun using technology spikes and difference-in-differences or event studies to analyze technology impacts (Bessen and Righi 2019; Bessen et al. 2022; Humlum 2019; Domini et al. 2021; Aghion et al. 2020; Hirvonen, Stenhammer, and Tuhkuri 2021; Rodrigo 2021).

To summarize, this paper makes several contributions. First, we develop a model of automation that includes both cost reduction and quality enhancement, thus generating a richer set of outcomes. The model accounts for automation effects on both within-firm and between-firm pay inequality, reflecting the relative extent to which automation substitutes or complements workers. This difference provides a tool to estimate the relative importance of substitution/complementation and implies different policy choices to counter inequality.

Second, we study the micro-level impact of proprietary information technology to test key aspects of the model. We find that when firms invest in these systems, their demand increases for a diverse set of technical, cognitive, and social skills and these demand increases occur across skill groups—across jobs requiring a college degree and those that do not, across routine jobs as well as nonroutine. Moreover, this increased demand is reflected in higher pay offered, after controlling for job characteristics, contributing to growing between-firm pay differences. While we also find evidence of labor displacement in manual jobs, our overall findings differ from predictions of models of pure labor displacement and from models of skill-biased technical change. In both the labor displacement and skill-biased change stories, technology only affect limited groups of workers; in our story, it affects most groups of workers.

Third, we explore how much of the overall sorting of skilled workers to high-paying firms can be accounted for by proprietary information technology systems by looking at the correlations between firm fixed effects and skills. We find that the majority of these correlations is accounted for by this technology, suggesting that the increase in sorting may be closely related to the rise of proprietary information technology.

Of course, there are important non-technological factors that may contribute to sorting including rent-sharing, firm size (Eeckhout and Kircher 2018), search frictions (Burdett and Mortensen 1998), and monopsony (see for instance Card et al. 2018). Cortes et al. (2020) model sorting arising from skill-biased technical change. In empirical research, technology has been associated with between-firm wage differences (M. Doms, Dunne, and Troske 1997; Dunne et al. 2004; Barth et al. 2020) and the rise in information technology is coincident with the rising importance of sorting to inequality. But little research connects the actual adoption of technology with changes in skill demand. Some research has explored the effects of the adoption of computers or automation technology on firm wages in difference-

in-differences or event studies, generally finding a rise in firm pay following adoption.⁷ However, the increase in firm pay could arise from rent sharing rather than from greater demand for skills. Dillender and Forsythe (2019), in an approach similar to ours, use Burning Glass data to identify firm computer technology adoption; they find greater skill demand and higher pay for office and administrative support workers, suggesting that this technology is labor augmenting. Other papers find that computers or AI change skill demands (Autor, Levy, and Murnane 2003; Spitz-Oener 2006; Acemoglu et al. 2020).

Some studies use worker fixed effects from AKM regressions as a proxy for skill, but these might also reflect rents arising from search frictions (Abowd, Kramarz, and Margolis 1999; Bagger and Lentz 2019). Hakanson et al. (2020) find that worker sorting across firms by ability measured using standardized test scores is related to the rising information technology sector, but they lack firm-level measures of technology. Deming (2017; see also Aghion et al. 2019) finds an association between information technology and soft skills. This paper studies firm-level adoption of technology and both the subsequent firm demand for specific skills and firm pay offers.

A historical example

To fix ideas, it is helpful to look at an example of the Remainder Effect. There is sufficient historical data available for mills weaving coarse cotton cloth during the 19th century in the U.S. to construct an engineering production function that specifies the weavers' main tasks, the frequency of their occurrence, typical times to perform those tasks, and how those tasks were automated over time (this section draws from Bessen 2003; 2012; 2015). Although the process was highly automated, humans still had to perform some critical tasks. When bobbins ran out of thread, humans replenished them; when the edges of the cloth pulled inward as the loom wove, humans straightened them; when threads broke, humans had to stop the machines, unravel the defective cloth, fix the break, and restart the machines.

⁷ Gaggl and Wright (2017) find that computers raise wages in small firms, mostly in managerial, professional, and technical occupations. Bessen et al. (2022) find that automation raises wages in large firms, but wages decline in small firms. Acemoglu, Lelarge, and Restrepo (2020) find that robots raise wages in some regressions. Humlum (2019) and Rodrigo (2021) also find that robots increase pay. Graetz and Michaels (2018) find a similar increase at the industry level.

The productivity of a weaver depended critically on her skill at performing these specific tasks and the effort she applied to them. Many tasks were performed while the loom was stopped, so the weaver's speed affected the rate of output. The reliability of weaver's performance determined whether subsequent defects or failures would occur with greater or lesser frequency. The weaver's attentiveness monitoring the looms affected how quickly faults could be detected and fixed. And of course, the quality of the weaver's performance also affected the occurrence of defects that reduced the value of the cloth.

These skills had to be learned on the job. New hires went through a learning process that quadrupled their output per hour over the course of a year or so. Treating foregone output as a human capital investment (Becker 1993), the human capital of these supposedly "unskilled" weavers was substantial, roughly equivalent to the investments in adult male tradesmen who went through apprenticeships (the weavers were mainly young women). These skills contributed substantially to the rise in productivity over the 19th century. The labor time required per yard of cloth fell by 98%. However, analysis using the engineering production function shows that new inventions cannot account for all this decrease; about a quarter is due, instead, to better quality of labor.

As automation progressed over the century, many of these tasks were automated and no significant new tasks were added in this sector. As weavers performed fewer tasks per yard of cloth, they were assigned more looms to tend, so that their skills on these tasks remained important for productivity. In fact, their skills became even more important. Human capital investments increased (learning curves became steeper) and real pay for weavers rose substantially. Automation substantially increased the returns to skill on the remaining non-automated tasks.

Model

Basic Setup

Tasks and Automation

Our model is a combination of cost saving automation models by Acemoglu and Restrepo (2018a; 2018c) and models of production quality by Kremer and Maskin (Kremer

1993; 1996). We interpret Acemoglu and Restrepo’s model as providing a measure of potential output while Kremer’s model relates actual output to potential output, after accounting for quality-related failures.

We use a simplified version of the Cobb-Douglas instance of Acemoglu and Restrepo’s model (2018a) with constant returns to scale. Let there be N tasks. We keep the number of tasks fixed, ignoring the creation of new tasks, which we discuss further below. Because the production function has constant returns to scale, we allow an indefinite number of firms. Let the tasks be ordered so that the first I tasks are automated and the remaining $N - I$ tasks are performed by labor. The i th automated task uses k_i capital and the j th human task uses l_j labor. Letting the firm’s total capital $K = \sum_{i=1}^I k_i$ and total labor $L = \sum_{i=I+1}^N l_i$, equilibrium potential output can be written, under some assumptions (see Acemoglu and Restrepo 2018a, equation 3),

$$V = A(I)K^\alpha L^{1-\alpha}, \quad \alpha \equiv \frac{I}{N}, \quad \frac{dA}{dI} > 0. \quad (1)$$

where α is capital’s share of output and $A(I)$ is a measure of Hicks-neutral productivity, which we assume to be increasing in the number of automated tasks. We assume that I is exogenously determined by the state of technology. Firms, however, pay a fixed fee to adopt the latest technology so that in some circumstances, only more profitable firms choose to adopt (for a full model of adoption see Bessen et al. 2022 Appendix).

Quality

However, as Kremer (1993) observes, not all potential output is realized if tasks are performed imperfectly. In some production functions, failure of a critical task reduces output to zero (O-ring); in others, imperfect task output reduces the value of output; in yet others, task failures delay production (weaving), reducing the rate of output. The critical assumption here is that quality and quantity are not perfect substitutes. If quality and quantity *were* perfect substitutes, then output could simply be measured in quality-adjusted units and there would be no need to account for quality separately. However, as Kremer and others argue, there are many important instances where this substitution is imperfect, e.g., two mediocre surgeons are not equivalent to one surgeon whose patients have twice the survival rate (Rosen 1981; Kremer 1993).

It is standard in reliability engineering that the probability of failure increases with the number of tasks prone to failure. Multiple tasks provide multiple opportunities to fail. Let q_i , $0 \leq q_i \leq 1$ be the quality of performance of the i th task at a given scale of production. Perfect performance is designated by $q_i = 1$ and complete failure by $q_i = 0$. Then the actual output can be written⁸

$$Y = Q \cdot V, \quad Q \equiv \prod_{i=1}^N q_i. \quad (2)$$

To keep things simple, we assume that machines perform their tasks perfectly, $q_i = 1$, while humans are always at least a bit imperfect.⁹ For the tasks performed by labor, task quality will depend on worker skill or effort. The quality of task production can vary with general skills of the workers performing the task, but in many cases, it will surely depend on task-specific and technology-specific skills. Without loss of significant generality, we assume that workers are assigned to a single task and all workers assigned to a task have the same quality. This way worker skills are task specific. The quality of each task performed by labor is then $q_i = f(e_i)$ where e is effort per worker, either effort expended on the task or effort expended on learning new skills (see below). Then we assign

$$q_i = \begin{cases} 1, & i \leq I \\ f(e_i), & I < i \leq N \end{cases}. \quad (3)$$

We assume that $f(e_i)$ is a monotonically increasing, twice-differentiable continuous function, $f' > 0$, $f'' < 0$, and $\lim_{e \rightarrow \infty} f(e) < 1$ (humans are imperfect).

Labor Quality

Workers deliver a fixed amount of labor—there is no tradeoff with leisure time—but the quality of that labor varies. In a single-period model it is simplest to represent labor quality with a single variable, e , which we might think of as effort, either expended on the

⁸ Where q represents a probability of successful completion, then Y is expected output and we assume that firms are risk neutral.

⁹ A more general model could consider cases where machines have low quality but high efficiency and cases where inefficient machines are adopted because they have higher quality.

task or expended on learning new skills.¹⁰ Given a set of equilibrium prices, each worker's utility can be written as a function of their wage and the effort they exert, $U = U(w, e)$. Let

$$U(w, e) = w - \theta_j g(e), \quad g' > 0, g'' < 0, \theta_j \geq 1$$

where θ_j represents the j th worker's disutility of effort. Workers share the same preferences except for this disutility parameter.

Importantly for our analysis, there are two dimensions to labor quality: the worker's type or skill group, θ_j , and the actual effort/skill exerted at the worker's job, e . We assume that employers (but not third parties) can costlessly observe both θ_j and effort e . Employers can then elicit greater or lesser e by using greater or smaller performance pay incentives. They can also influence the labor quality of their workforces by hiring workers of different types, θ_j , who will be more or less responsive to those incentives.

Employers can elicit a desired level of effort by offering a two-part wage, with fixed amount w_u and performance incentive w_s such that $W = w_u + w_s \cdot e$. Utility maximization for the j th worker occurs when $\frac{w_s}{\theta_j} = g'(e)$, yielding a unique level of effort, $\hat{e}(w_s/\theta_j)$, and the corresponding task quality, $f(\hat{e}(w_s/\theta_j))$.

It is convenient to invert this function, yielding

$$w(q) = w_u + \theta_j \cdot h(q), \quad h', h'' > 0, \quad \lim_{q \rightarrow 1} h'(q) = \infty.$$

We assume competitive labor markets; firms pay the same wages for the same level of effort/skill. Wages differ across firms to the degree that worker effort/skill differs across firms.

Uniform Workers and Firms

We begin the exposition by presenting our model with uniform workers and firms to establish some basic results. Let there be only one type of labor, $\theta_i \equiv \theta$ for all i with otherwise identical firms. We introduce heterogeneity in the next section.

¹⁰ In a multiperiod model, workers may invest time and effort in learning new skills in one period that are used in subsequent periods.

Equilibrium

There is a fixed amount of inelastically supplied labor and capital in the aggregate economy distributed across firms. With uniform labor and firms, firms receive proportional allocations of labor and capital, L and K in equilibrium. Taking output price as numeraire, profits per firm are

$$\begin{aligned} \pi(q_{I+1}, \dots, q_N, k_1, \dots, k_I, l_{I+1}, \dots, l_N; I) \\ = A(I)K^\alpha L^{1-\alpha} \prod_{i=I+1}^N q_i - \sum_{i=1}^I r k_i - \sum_{i=I+1}^N w(q_i) l_i, \end{aligned}$$

where r is the user cost of capital and w is the wage. By the symmetry of the problem, it is straightforward to show that $q_i = q_j$, $k_i = k_j$, and $l_i = l_j$ in the appropriate range in equilibrium. The first order profit maximizing conditions for the three control variables then are

$$\frac{Y}{q_i} - \theta h' l_i = \frac{Y}{N l_i} - w = \frac{Y}{N k_i} - r = 0. \quad (4)$$

A useful result can be obtained by taking the implicit derivative from the first order maximizing condition for q_i (keeping the quality of other tasks fixed),

$$\frac{dq_i}{dA} = \frac{Nw}{\theta A q_i h''(q_i)} > 0. \quad (5)$$

Thus, increases in productivity will increase the equilibrium quality of output. When potential output increases, firms increase incentive pay, workers exert greater effort/skill, and total output increases more than potential output. In other words, an increase in potential output increases the returns to skill/effort.

Remainder Effect

Now consider what happens when the frontier of automated tasks increases from $I - 1$ to I for all firms. Let us assume that the adoption costs of the new technology are negligible so that all firms adopt. Productivity, A , increases and, by implication of the lemma above, this increase should boost labor quality. Aggregate quality also increases because the machine produces with greater quality on task I , that is, $1 > f(e_I)$. Combined, the effect of automation on total output per worker is

$$\Delta \ln \frac{Y}{L} = \Delta \ln A + \Delta \ln Q + \alpha \Delta \ln \frac{K}{L}$$

In this setting, capital and labor will be allocated proportionately across production units in equilibrium, so the last term drops out. Then,

$$\Delta \ln \frac{Y}{L} \approx \Delta \ln A + (N - I) \Delta \ln A \cdot \frac{dq}{dA} \cdot \frac{A}{q} - \ln f(e_I). \quad (6)$$

The second term represents the remainder effect. Automation boosts the returns to quality, increasing equilibrium labor quality. Output increases not only because automation reduces the labor cost of production but also because it increases labor quality. The third term is positive (since $f < 1$, $-\ln f > 0$) and captures the effect of improved quality in the newly automated task.

There is a corresponding change in the wage. Using the first order conditions and $L = (N - I)l_i$, the equilibrium wage is

$$w = \frac{N - I}{N} \cdot \frac{Y}{L}.$$

Following Acemoglu and Restrepo and using (6),

$$\begin{aligned} \Delta \ln w &\approx \frac{d \ln(N - I)}{d I} + \Delta \ln \frac{Y}{L} \\ &\approx -\frac{1}{N - I} + \Delta \ln A + (N - I) \Delta \ln A \cdot \frac{dq}{dA} \cdot \frac{A}{q} - \ln f(e_I) \end{aligned} \quad (7)$$

Acemoglu and Restrepo call the first term the “displacement effect” The second term is an efficiency effect (Acemoglu and Restrepo call it the “productivity effect”). The third term represents the remainder effect and the fourth captures the quality improvement effect. The remainder effect multiplies the base productivity effect, making a positive contribution to wages. Also, the fourth term implies further possible wage increases. In a more general model, this term could possibly be negative—that is, firms might accept inferior quality machines if they deliver a large enough efficiency gain. The sign and magnitude of this term is an empirical matter. However, the addition of the term highlights an important aspect of automation: firms may choose to automate not so much to reduce costs as to provide better quality output. To the extent this is true, the effect on wages will tend to be positive.

Generally, (7) provides reasons beyond Acemoglu and Restrepo why wages might increase.

To keep things simple, we have used single continuous variables for product and labor quality and have kept the number of products and tasks fixed. In a more general

setting, both new tasks and new products might be natural outcomes of a growing demand for greater quality. For example, as the quality of a task becomes more and more valuable with ongoing automation, firms might subdivide that task into two or more new tasks allowing workers to develop more specialized skills. Something like that appears to have happened during the 19th century (Atack, Margo, and Rhode 2019). Similarly, new products might be a form of realizing greater product quality.

Heterogeneity

Now let there be two types of workers: high skill, designated “H,” and low skill, designated “L,” where $\theta_H < \theta_L$. The aggregate supply of each type is fixed.

In general, there are two ways that workers can be assigned to firms: assortative matching, where some firms hire more high skill workers and other firms hire more low skill workers, and cross-matching, where firms hire a mix of high and low skill workers. A theoretical literature identifies a condition under which assortative matching occurs in competitive markets (Becker 1981; Sattinger 1975; 1993; Kremer 1993; Kremer and Maskin 1996), namely a positive cross derivative of output with respect to the qualities of different tasks. Our production function meets this criterion (see also Kremer 1993). In the next section, we consider a stylized model of sorting where firms hire all high skill workers or all low skill workers.

Kremer and Maskin (1996) show that with a slightly different production function, firms will, instead, cross-match under some conditions, hiring both high and low skill workers. This occurs when productivity is more sensitive to some tasks than others. Let us divide tasks into two groups: tasks in the range $I < i \leq J$ are “routine tasks” while tasks in the range $J < i \leq N$ are “nonroutine tasks.” Below we consider an alternative specification that meets the Kremer-Maskin conditions. While real world skill assignments may involve a mix of matching and sorting, these models illustrate in simple form the different effects that automation has on inequality between firms and within firms.¹¹

¹¹ Automation might also affect firms’ choices regarding sorting and cross-matching. Kremer and Maskin (1996) provide a variety of evidence that skill sorting has been increasing and workplaces are becoming more segregated by skill, that is, workers are more likely to work with other workers of similar skill (see also E. Handwerker 2015; E. W. Handwerker, Spletzer, and others 2016). Our model could be extended to address this possibility.

Sorting

In a market with complete sorting, some firms, designated by an “H” subscript, hire only high skill workers while other firms hire only low skill workers, designated with an “L” subscript. We assume that both types have the same level of automation initially. The first order profit maximizing conditions (4) then hold separately for each firm type. Combining the first order conditions for quality and labor, for worker/firm type j ,

$$w_j = \frac{Y_j}{Nl_j} = \frac{\theta_j \cdot h'(q_j) \cdot q_j}{N}, \quad j = L, H.$$

In the Appendix we show that in equilibrium, both q_j and the term $\theta_j \cdot h'(q_j) \cdot q_j$ are decreasing in θ_j , all else equal. This means that $w_H > w_L$ and the ratio of between-firm wages is

$$\omega \equiv \frac{w_H}{w_L} = \frac{\theta_H \cdot h'(q_H) \cdot q_H}{\theta_L \cdot h'(q_L) \cdot q_L} > 1.$$

The between-firm wage gap corresponds directly to differences in skill/effort between the firm types. Furthermore, it is straightforward to show that capital intensity and productivity are higher in type H firms:

$$\frac{w_H}{w_L} = \frac{K_H/L_H}{K_L/L_L} = \frac{Y_H/L_H}{Y_L/L_L} > 1.$$

To introduce automation into this setting, note that because type H firms have higher productivity, they also have stronger incentives to adopt new automation technology. The increase in output per worker from automation is $\frac{Y}{L} \Delta \ln A$ and so will be larger for type H firms. This increase will also be greater for the remainder effect term in (6). Suppose that there is a fixed cost per worker needed to adopt an automation technology. Then, in some cases, type H firms will find it profitable to automate while type L firms will not.¹² Given this difference, let us assume that type H firms automate, and type L firms do not. Disparate adoption of automation technologies is, in fact, widely observed and appears in our data as well.

¹² Firms may make temporary profits from automating, yet competition will eventually dissipate these rents. There are other reasons some firms may adopt while other do not: different capabilities of managers and workers or different access to proprietary technologies.

With this assumption, we can calculate ω using an approach like the one used in equation (7). Here, however, we must account for changes in the capital to labor ratios for the two groups. As Y/L increases for H firms, capital also shifts to those firms. In the Appendix we account for this change in the equilibrium solution to derive an approximate lower bound for the change in the between-firm wage ratio:

$$\Delta \ln \omega = \Delta \ln w_H - \Delta \ln w_L \approx > -\frac{1}{N-I} + \frac{N}{N-I-1} \left[\Delta \ln A_H + \Delta \ln Q_H + \frac{1}{I} \right].$$

The first term represents the displacement effect. The expression in brackets captures the productivity and quality effects. Here the displacement effect *decreases* between-firm wage differences while the productivity and remainder effects increase between-firm wage differences. If the productivity and remainder effects are larger than the displacement effect, the between-firm wage gap increases. If, on the other hand, low wage firms tend to automate, contrary to most evidence, then the changes would narrow between-firm differences. And if both types of firms automated, the results are ambiguous. Thus, growing differences in labor quality explain rising between-firm pay gaps if adoption of automation technology is uneven and if the displacement effect is smaller than productivity and remainder effects.

Cross-matching

Kremer and Maskin (1996) show that cross-matching occurs when some tasks are more sensitive to quality than others. The idea is that, under some parameter values, firms will choose to assign high skill workers to sensitive tasks and low skill workers to tasks that are less sensitive.¹³ We can accommodate these notions into our production function by specifying now that

$$q_i = \begin{cases} 1, & i \leq I \\ 1, & I < i \leq J \\ f(e_i), & J < i \leq N \end{cases}$$

where $I < J < N$. Routine tasks in the range $I < i \leq J$ are not sensitive to the quality of labor on those tasks while nonroutine tasks in the range $J < i \leq N$ depend on the skill and effort of workers. With this modification to the production function, firms will prefer to hire

¹³ Acemoglu and Restrepo (2018a; 2018b) exogenously assign high skill workers to nonroutine tasks and low skill workers to routine tasks.

high skill workers for nonroutine tasks and low skill workers for routine tasks. In equilibrium, the fixed stocks of low and high skill workers will be allocated proportionally to firms so that the ratio L_L/L_H of low skill workers to high skill workers will be the same. It is straightforward to show that firms will prefer to assign only high skill workers to nonroutine tasks and only low skill workers to routine tasks. Then first order profit maximizing conditions give us

$$w_H = \frac{Y}{L_H} \frac{(N - J)}{N} = \theta_H h'(q_H) q_H N, \quad w_L = \frac{Y}{L_L} \frac{(J - I)}{N}$$

and the within-firm wage difference ratio is

$$\phi \equiv \frac{w_H}{w_L} = \frac{N - J}{J - I} \cdot \frac{L_L}{L_H}.$$

Note that the within firm wage difference is independent of the quality of the high skill workers. The relative wage within firms depends on the relative supply of workers from different skill groups and the relative demand for routine and nonroutine tasks. While high skill workers will receive higher performance pay, their total pay package is not necessarily greater.

Now consider automation in this setting. In some papers, Acemoglu and Restrepo (2018a; 2018c) study situations where only routine tasks are automated. Then automation can be considered a change in the limit of automation from I to $I + 1$ as above. Then

$$\Delta \ln \phi \approx \frac{d}{dI} \left(\frac{N - J}{J - I} \cdot \frac{L_L}{L_H} \right) = \frac{1}{J - I} > 0.$$

Automation increases within-firm wage differences in this setting. However, automation is not necessarily restricted to routine tasks and then this type of labor displacement might decrease within-firm wage gaps (see Acemoglu and Restrepo 2018b).

But the general point remains that automation influences within-firm wage gaps by way of the displacement effect. In our model as well as the models in the literature, labor displacement directly affects the relative demand for different skill groups within firms and aggregate changes in demand for these groups determines the relative equilibrium wages. Because workers from one skill group are employed at this wage across firms, the effect will be observed as within-firm wage differences. On the other hand, the remainder effect concerns firm-, task-, and technology-specific skills that are not common across different

firms. These affect between-firm wage differences but not differences between skill groups within the firm.

This model provides three testable hypotheses:

1. Automation should increase the demand for task- and technology-specific skills across multiple skill groups;
2. This greater demand should be evident in the firm's greater willingness to pay more for these groups; and,
3. Assuming that automation differentially affects the tasks assigned to different skill groups, it should change the relative employment demand for different skill groups.

The first two hypotheses distinguish this model from pure models of labor displacement: here, automation complements labor. The firm's greater willingness to pay provides an explanation for greater between-firm pay gaps. Our model also differs from the skill-biased technical change hypothesis because the complementary effect of technology is not limited to specific skill groups.

Empirical Analysis

Data

These three hypotheses concern different aspects of firm labor demand: the specific skills demanded, the firm's willingness to pay for different skill groups, and the relative quantities of labor demanded for different skill groups. We measure these aspects of demand using help-wanted advertisements collected by Burning Glass Technologies. Burning Glass scrapes, deduplicates, and cleans the near universe of online job advertisements. A previous analysis of the dataset showed that this it accounts for 60-70% of all job openings and 80-90% of openings requiring a bachelor's degree or more (Carnevale, Jayasundera, and Repnikov 2014). The data include the advertised salary, firm name, industry, occupation, required education and experience, requested skills, and geographic location of the job. Our

sample spans from January 2014 to June 2019.¹⁴ We aggregate the ads by firm and calendar quarter and use this as our unit of observation.

Changes in labor demand should be immediately reflected in help-wanted advertising even though these changes might take longer to appear among the group of employed workers. To the extent that firms demand greater quality on task-specific skills, we should see increases in the specific skills requested in job ads. To the extent that greater demand increases the firm's willingness to pay, we should see higher pay offered for jobs with comparable characteristics. And to the extent that demand changes across skill groups, we should see shifts in the share of job ads directed to different skill groups. We measure these outcomes with the following variables:

Specific skills. Burning Glass collects 16,050 different skills requested in ads as well as experience and education required. We group the specific requests into five mutually exclusive categories: social and cognitive skills as identified by Deming and Khan (2018), other soft skills, information technology and artificial intelligence, and other skills, mainly skills related to other technologies and industry knowledge (see Appendix). We use the mean number of requests per ad for each category and the mean experience and education requested as outcome measures.

Pay offered. Some help wanted ads list a salary offered or a range of salaries. If a range is offered, we take the middle of the range for our salary calculations. The outcome variable is the log Mincer residual in a regression equation including experience, experience squared, education, detailed occupation, state, year, and a measure of labor market tightness. We follow Moscarini and Postel-Vinay (2016) in defining labor market tightness as the ratio between Job Openings and Labor Turnover Survey (JOLTS) statewide openings for the non-farm sector and the state unemployment rate.¹⁵

¹⁴ While Burning Glass provides data prior to 2014, those years used different methods to collect, de-duplicate, and process the data. Because those differences might affect our analysis, we do not use that data. We omit job advertisements that are missing a firm name or salary, are in the public or university sector, are part time, or are internships. To identify ads belonging to the same firm, we cleaned names, removing standard business identifiers (“Inc.,” “Ltd,” “Co.,” etc.) and looking for typos in the most frequently used names in the dataset.

¹⁵ Because most jobs do not list salaries, sample selection bias might affect this measure. Bessen et al. (2020) find that an exogenous change to salary listing does not significantly affect listed salaries, mitigating this concern.

Relative employment. To measure changes in the relative hiring of skill groups, we use the share of job ads for each group. We divide occupations into two sets of skill groups defined by characteristics identified in O*NET, version 17.0. First, we identify whether a bachelor’s degree or higher is required for most jobs in that occupation. Second, we identify occupations as routine cognitive, routine manual, nonroutine cognitive, and nonroutine manual using the indexes for these characteristics developed by Acemoglu and Autor (2011); an occupation is assigned to the job characteristic skill group if its index ranks in the top third.¹⁶

Finally, note that we exclude information technology jobs (SOC 15) from our skill and pay measures to avoid confounding effects.

Implementation

We seek to test the model predictions regarding the adoption of large proprietary information systems. Much of the literature on technology and inequality measures technology as predicted “exposure” to automation, or industry-level investment levels, or proxies such as the share of workers in routine-intensive jobs. To capture impacts on between-firm differences, we thought it important to use firm-level measures of actual technology adoption. These eliminate many potentially confounding correlates.

We measure investment in this technology from the job ad data as the share of jobs going to software developer occupations.¹⁷ This captures investment in firms’ own-developed software and it is correlated with contracted software and other IT measures (Tambe and Hitt 2012; Bessen 2020 fn. 12).

To analyze adoption, we identify “spikes” in developer hiring as events where the share of software developers rose by one percent or more relative to the mean share over the previous four quarters.¹⁸ This approach leverages the finding from the capital investment

¹⁶ These groups are not mutually exclusive.

¹⁷ Occupations in SOC 15 excluding 15-1141, 15-1142, 15-1151, and 15-1152, database, network, and computer administrators and support specialists.

¹⁸ Also, to reduce noise, we eliminate spikes when the firm has fewer than 50 ads in quarter. A variety of robustness checks in the Appendix vary the threshold, finding little effect on results. 19% of firm-quarters are spikes, weighted by the number of job ads. While only about 1% of firms spike, these firms account for 77% of the hiring of software developers.

literature that when uncertain investments are indivisible and irreversible, they will occur in discrete episodes of lumpy investment (Haltiwanger, Cooper, and Power 1999; M. E. Doms and Dunne 1998). We find that investments in own-developed software are also lumpy and persistent (see Appendix Figures A1 and A2), so we use these discrete events in difference-in-differences (DID) regressions and event studies. It is possible that we fail to identify some lumpy investments and incorrectly identify others. For example, firms rely on outside contractors to implement new systems rather than hiring their own developers. To the extent misidentification occurs, our results will be understated.

Do these spike events represent automation? We note generally that most information technology applications involve some degree of automation—they manage information that was formerly managed by humans. This is strictly true for applications that automate business processes such as enterprise resource planning, customer relationship management, and electronic data interchange. In fact, the use of these systems is correlated with bookkeeping measures of automation expenditures (Bessen et al. 2022 Section 2.3). We flag events that specifically include hiring of workers with skills related to these automation applications and find that 81% of our spike event do.¹⁹ Similarly, 31% of the spikes involve firms requesting artificial intelligence skills. Thus our spikes predominately involve applications that automate tasks.

To avoid problems of heterogeneity in our two-way fixed effects regressions, we construct balanced panels around each possible spike quarter and run stacked regressions (Cengiz et al. 2019, Appendix D). Let T_i be the first quarter in which firm i spikes. For each possible spike quarter, p , designating a different cohort, we construct a balanced panel P consisting of observations from $t = p - 5$ to $t = p + 5$ of the treatment group, $T_i = p$, and the control group, $T_i > p + 5$. Because firms that spike are different from firms that do not (see Table A1), we restrict the control group to firms that spike at some point in our data. This means that the treatment and control groups differ only in the timing of their adoption events.²⁰ This gives us a degree of identification by removing fixed or slowly

¹⁹ These are jobs requesting skills with keywords ERP, CRM, EDI, MRP, SAP, Automat*, and Robot*.

²⁰ Bessen et al. (2022, Appendix) provide a model for differential timing. We also duplicate our results for the full sample (Table A4).

changing confounders, such as industry and firm size, and by distinguishing major new investments from maintenance hiring. Our DID specification for outcome variable Y is

$$Y_{ipt} = \delta \cdot \mathbf{1}(t \geq p) + \mu_{ip} + \tau_t + \beta X_{it} + \epsilon_{ipt}. \quad (8)$$

where δ is the average treatment effect, μ_{ip} is the panel x firm fixed effect, τ_t is the time fixed effect, and X_{it} is a vector of control variables.

However, the model is still not fully identified because the timing of adoption is endogenous. While we test for and do not find significant pre-trends in our outcome variables, it is still possible that some other factor is correlated with adoption, occurring simultaneously, and which independently affects outcome variables. We identify and control for four such possible simultaneous confounders:

1. **Labor market tightness.** Tight labor markets might induce firm to automate and might also raise wages and skills demanded (Modestino, Shoag, and Ballance 2019 find tight labor markets *lower* skill requirements). We use the tightness measure described above to control for this confounder.
2. **Outsourcing of low wage jobs.** Perhaps automation facilitates the outsourcing of low wage jobs, mechanically raising the average pay and skill requirements of remaining jobs. We control for the share of “outsourcable” jobs that should track these shifts.²¹
3. **Productivity and demand shocks.** Perhaps firms adopt new technology in response to productivity or demand shocks and these shocks are also passed through to wages. We control for shocks using additional variables obtained from Compustat for the subsample of firms matched to Compustat.²² One variable is the growth in real sales from the quarter before the spike to a year

²¹ The outsourceable occupations are Protective Services (SOC 33), Food and Serving (SOC 35), Building, Grounds, Maintenance (SOC 37), and Transportation and Moving (SOC 53) outside of outsourcing industries, NAICS 484, Truck Transportation, NAICS 561, Administrative and Support Services, NAICS 722, Food Services and Drinking Places, and NAICS 811, Repair and Maintenance.

²² Bledi Taska of Burning Glass provided a preliminary key to match to Compustat, which we supplemented with our own name cleaning algorithm. Further, we used a fuzzy match with distance scores, which was then manually reviewed for those with close distances. The match assigns approximately 63% of the firms in Compustat to a job posting, with 73% of the firm-years being matched to a job posting. The firms that are matched to a posting account for 83% of employment total in Compustat.

earlier. The second control is a third order polynomial in log variable costs and log net capital stock (both deflated).²³

- 4. Management.** Perhaps new managers prefer to adopt technology and also to hire more highly skilled workers. For the entire sample, we add the manager (SOC 11) share of hiring as a control. For the Compustat subsample, we add a binary variable to flag changes of CEO using data obtained from Execucomp.

We find that some of these control variables have weak correlations with the occurrence of spikes (see Table A2), but also, they do not substantively change our results. This gives us a limited form of identification; it is not equivalent to conducting a randomized controlled trial, but our results are identified conditional on the following assumption:

Identification assumption: there are no significant confounders that occur simultaneously with the adoption of these information technology systems other than labor market conditions, outsourcing, productivity and demand shocks, and management changes.

As such, our results are consistent with our model and inconsistent with pure displacement models and with the skill-biased technical change hypothesis. Finally, our spiking results pertain to a select sample of firms. Below we also explore the broader validity of our model to the universe of help-wanted ads.

Findings

Firm Spikes

Table 1 presents stacked difference-in-differences regressions (a balanced panel for each spiking year) where the dependent variables are the number of skills requested in the various categories.²⁴ All of the skill measures show significant increases following the adoption event except for education. The top panel includes all jobs except for IT jobs (SOC 15). We interpret the greater number of skills requested as evidence of greater demand for

²³ In the style of Olley and Pakes (1996) this polynomial is a nonparametric representation of productivity obtained by inverting the demand equation for variable inputs (cost of goods sold).

²⁴ Regressions are weighted by the number of ads and include time and cohort by firm fixed effects as well as controls for labor market tightness, and the shares of management and outsourceable jobs.

specific skills. When firms place greater value on “Teamwork” or on “Adhesives Industry Knowledge,” they will be more likely to specifically request these skills.

Panel B includes the skill measures only for jobs that do not require a college diploma.²⁵ These coefficients tend to be a bit smaller, but as in the larger sample, all are significant and positive except for education. Skill demands appear to rise for both college and non-college jobs, although a bit less for the latter.

Panel C looks at the *share* of skills rather than the number, that is the number of skill requested in each category divided by the total number of skills requested. Following a spike, firms appear to place relatively greater demand on social and soft skills, suggesting organizational changes consistent with Deming (2017). However, these shifts in the composition of skills are small compared to the increases in demand seen in Panel A.²⁶ The overall impact appears to be that firms request more of the kinds of specific skills that they requested before the spike, that is, they demand higher labor quality.

Table 2 examines a broader set of skill groups, namely jobs classified as routine/nonroutine and cognitive/manual as per Acemoglu and Autor (2011). Panel A shows that all groups show significant increases in the mean number of skills requested except for nonroutine manual jobs. These results suggest that the technology complements workers in a wide range of jobs. As we would expect, firms are also willing to pay more to these workers seemingly complemented by software investments—the greater demand for skills does not just reflect the preferences of HR professionals. The dependent variable in Panel B is the log residual wage after controlling for job characteristics. These pay levels rise significantly for all groups except nonroutine manual workers; they rise notably more (9.1%) for nonroutine cognitive jobs.

Table 3 tests the robustness of results to additional controls. Here the sample is limited to firms that are matched to Compustat. Using Compustat and Execucomp data, we add a control (in columns 3 and 6) for the rate of revenue growth, a flag for change of CEO, and a third order polynomial in log capital and log variable costs to capture productivity

²⁵ That is, fewer than half the jobs require a diploma as rated by O*NET.

²⁶ Expressed as percentages, the increases shown in Panel A range from 3% to 13%, much larger than the shifts, which are less than 1%.

nonparametrically. Some of these controls are statistically significant, but they do not meaningfully alter our estimates of the treatment effect.

Our results are also robust to other concerns. Figures 1 and 2 shows event study graphs corresponding to the first column in Table 2.²⁷ The graph shows a significant and persistent increase in the mean number of skills requested and log residual wage following an adoption event. Moreover, there is no evidence of pre-event trends in these outcome variables nor in the other outcome variables used in Table 1, lending support to the parallel trends assumption (see Appendix Table A9). Table A3 tests sensitivity to different spike thresholds and panel lengths; our results are robust to these changes. Table A4 shows regressions using an expanded sample that adds firms that never spike; the results are similar. Table A6 finds that excluding firms in industries that create software products (NAICS 50 and 54) makes little difference to our results. About one third of our spiking firms use artificial intelligence as evidenced by requests for AI skills during the spiking quarter; 81% involve automation technologies. Our main results do not change significantly for these groups of firms (Table A7). We also conduct a placebo test to support the idea that the effects we observe are related to software specifically and not to other technologies or general hiring of higher paid workers. In Table A8, we show results from spikes in the hiring of engineers and technicians constructed in the same way as our software spikes. These personnel may tend to work on technologies that are not so much about automation. Spikes in the hiring of engineering related personnel do not exhibit similar treatment effects, suggesting that it is something specifically about information technology—perhaps automation—that is driving our results.

The increased skill demands and greater pay suggest that proprietary information systems complement labor. Our model suggests that automation can also displace labor. Table 4 shows evidence of displacement. The top panel shows the share of job ads going to each skill group. Following technology investment, relative hiring increases for jobs requiring college degrees and for jobs with cognitive skills, both routine and nonroutine; relative hiring

²⁷ These show the δ_τ coefficients from the following modification of (8):

$$Y_{ipt} = \sum_{\tau=-4}^5 \delta_\tau \cdot \mathbf{1}(\tau = t) + \mu_{ip} + \tau_t + \beta X_{it} + \epsilon_{ipt}.$$

decreases for non-college jobs and manual jobs. Panel B displays the log level of hiring by skill groups. Job ads decrease for occupations that do not require a college degree and for routine manual jobs. Thus, consistent with our model, there is labor displacement that occurs alongside increased demand for skills as seen in the prior tables.

In theory, this displacement contributes to lower equilibrium wages for workers who only have routine manual skills or only a high school diploma. In practice, however, this is difficult to establish empirically because other factors might confound the effect of technology on the pay of different demographic groups (but see Acemoglu and Restrepo 2021). For instance, in 1980 62% of the U.S. workforce had only a high school degree or less; today that figure is 38%. It seems highly likely that expanded access to higher education may have selectively induced some workers—those with lower disutilities of learning—to seek more education. This means that high school educated workers do not comprise a consistent skill group over time and declining pay for this group might reflect declining ability rather than technological effects. Our model provides some insight into the relative importance of labor displacement on wages. Because labor displacement affects market wages, it affects all firms equally; that is, it increases within-firm inequality. The finding that relatively little of the increase in inequality arises within firms—26% according to Song et al. (2019)—suggests that labor displacement is not the dominant driver of rising inequality.

To summarize, given our identification assumption, the evidence on residual wages implies that firm investments in proprietary information technology contributes to between-firm pay differences; the evidence on skills requested implies that these firm pay increases are associated with increased firm skill demands. In other words, these technology investments contribute to sorting of skills to higher paying firms. However, the evidence presented pertains only to a select sample of firms.

Sorting

We can also explore the relationships between firm pay levels, skills requested, and information technology across the entire sample of help wanted ads by looking at sorting of skills to high-paying firms. Studies using linked employee-employer data find that sorting accounts for most of the increase in wage inequality since 1980 (Card, Heining, and Kline 2013; Barth, Davis, and Freeman 2018; Song et al. 2019; Lachowska et al. 2020). Utilizing the AKM method (Abowd, Kramarz, and Margolis 1999), these studies estimate firm pay fixed

effects controlling for observed and unobserved worker heterogeneity with worker fixed effects. The worker effects are positively correlated with the firm effects and this correlation accounts for much of the rise in inequality. Assuming that the worker effects represent worker skills (rather than arising from search frictions or other factors), this correlation represents sorting of skilled workers to high-paying firms.

We alternatively estimate firm pay fixed effects by regressing pay offered in job ads controlling for job characteristics. These pay offers are obviously independent of individual worker heterogeneity. Using log salary as the dependent variable (or the mean of the salary range limits if a range is listed), we calculate firm fixed effects in a regression with controls for detailed occupation, industry, state, year, labor market “tightness,” skills requested, education required, and experience required (see Table A5). The R-squared for this regression is .688. The regression excludes software development occupations to avoid spurious correlation with our key independent variable. This gives us estimates of firm fixed effects for 205,306 firms that posted 85,142,065 help wanted ads, excluding ads for information technology occupations. These firm fixed effects are different from fixed effects derived from the AKM method—our fixed effects reflect differences in pay in hiring, not in the pay of incumbent workers.²⁸ Nevertheless, both methods provide estimates of the firms’ varied willingness to pay for comparable workers. And we can measure sorting by looking at the correlation between these firm fixed effects and actual skill levels demanded in the job ads. These correlations are shown in the top panel of Table 5 which reports regressions of mean skill measures for each firm against firm wage fixed effects. The correlations are all significant, indicating sorting. The standardized coefficients represent the correlation coefficients. These are similar to the correlation of 0.28 between worker fixed effects and firm fixed effects reported by Song et al. (2019) for the period from 2007-13 using the AKM method.²⁹

²⁸ There is a close correspondence between average advertised salaries and average salaries actually paid as observed in the Current Population Survey. Weighting the job ads to match the CPS distribution across occupations, the median log salary range from Burning Glass is from 10.32 to 10.69. The median log CPS salary for new hires is 10.48.

²⁹ Calculated using their figures for $\frac{cov(WFE, FFE)}{\sqrt{var(WFE)var(FFE)}}$.

But it turns out that firm hiring of software developers is correlated with both firm fixed effects and with skill measures.³⁰ The bottom panel adds quadratic terms in the mean share of software developers in hiring. The correlations between worker fixed effects and skill measures drop sharply. The last row shows the magnitude of the decrease in the standardized coefficients as a portion of the correlation coefficient in Panel A. It appears that information technology investments can account for the majority of the sorting of skills to high paying firms in hiring. Given that firm investment in own-developed software has increased more than ten-fold since the 1980s (BEA data), this shift can explain much of the rise in inequality due to sorting.

Conclusion

This paper argues that automation can be both cost-reducing and quality-enhancing; it can replace labor on some tasks while it increases demand for skills on others. Major investments by firms in own-developed information technology are followed by greater demand for specific skills requested in job ads and by higher pay offers. Moreover, demand increases across skill groups, both for jobs requiring college and those that do not, for routine jobs as well as nonroutine jobs. These broad increases contribute to between-firm pay differences and the sorting of skilled workers to high paying firms. Analyzing the universe of help-wanted ads, we find that these information technology investments account for most of the sorting across firms.

This pattern differs from predictions of the skill-biased technical change hypothesis and from theories of labor displacement. Our model provides an explanation: labor quality matters. While automation displaces labor on some tasks, it can also increase the returns to skill on the remaining non-automated tasks. Models that view automation as strictly substituting for labor without also complementing some workers might be incomplete and overly pessimistic. For instance, Acemoglu and Restrepo argue that wages will fall for “so-so innovations” where the productivity gain is small. But if automation raises the demand for

³⁰ See our working paper for a more complete exploration of these relationships (Bessen, Denk, and Meng 2021)

quality on the remaining tasks (remainder effect), wages may rise even with modestly productive innovations.

The matter is ultimately empirical, but here, too, labor quality matters for the analysis. Inequality is frequently measured by differences between occupational or educational groups. But our evidence suggests that skills and inequality change along other dimensions as well. In our model, labor displacement gives rise to greater within-firm inequality, but the evidence suggest that this is a secondary contributor to growing inequality. On the other hand, automation that complements labor can increase between-firm inequality, which appears to be more important.

If so, this suggests a different direction for policy to combat income inequality. Researchers who assume that automation is purely labor displacing have proposed policies to redistribute income, to alter tax incentives to discourage too much automation, and to encourage engineers to not develop automation (Korinek and Stiglitz 2018; Benzell et al. 2016; Acemoglu 2021; Brynjolfsson 2021). But if automation mainly complements workers, giving rise to greater between-firm pay differences, then policy might instead need to focus on reducing differences between firms in the uneven adoption of technology. Indeed, concerns have been raised about slower diffusion of technology (Andrews, Criscuolo, and Gal 2016; Akcigit and Ates 2021). While policy evaluation is beyond the scope of this paper, our analysis highlights that policy should be based on a richer picture of automation, one where technology complements labor as well as substitutes for it, where the quality of labor matters.

References

- Abowd, John M, Francis Kramarz, and David N Margolis. 1999. "High Wage Workers and High Wage Firms." *Econometrica* 67 (2): 251–333.
- Acemoglu, Daron. 2002. "Technical Change, Inequality, and the Labor Market." *Journal of Economic Literature* 40 (1): 7–72.
- . 2021. "Harms of AI." National Bureau of Economic Research.
- Acemoglu, Daron, and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." *Handbook of Labor Economics* 4: 1043–1171.
- Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo. 2020. "AI and Jobs: Evidence from Online Vacancies." National Bureau of Economic Research.
- Acemoglu, Daron, Claire Lelarge, and Pascual Restrepo. 2020. "Competing with Robots: Firm-Level Evidence from France." In *AEA Papers and Proceedings*, 110:383–88.
- Acemoglu, Daron, and Pascual Restrepo. 2018a. "Artificial Intelligence, Automation, and Work." In *The Economics of Artificial Intelligence: An Agenda*, 197–236. University of Chicago Press.
- . 2018b. "Low-Skill and High-Skill Automation." *Journal of Human Capital* 12 (2): 204–32.
- . 2018c. "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review* 108 (6): 1488–1542.
- . 2021. "Tasks, Automation, and the Rise in US Wage Inequality." National Bureau of Economic Research.
- Aghion, Philippe, Céline Antonin, Simon Bunel, and Xavier Jaravel. 2020. "What Are the Labor and Product Market Effects of Automation? New Evidence from France."
- Aghion, Philippe, Antonin Bergeaud, Richard W Blundell, and Rachel Griffith. 2019. "The Innovation Premium to Soft Skills in Low-Skilled Occupations." *Available at SSRN* 3489777. <https://ssrn.com/abstract=3489777>.
- Akcigit, Ufuk, and Sina T Ates. 2021. "Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory." *American Economic Journal: Macroeconomics* 13 (1): 257–98.
- Alekseeva, Liudmila, José Azar, Mireia Gine, Sampsa Samila, and Bledi Taska. 2020. "The Demand for AI Skills in the Labor Market."
- Andrews, Dan, Chiara Criscuolo, and Peter N Gal. 2016. "The Best versus the Rest: The Global Productivity Slowdown, Divergence across Firms and the Role of Public Policy."
- Atack, Jeremy, Robert A Margo, and Paul W Rhode. 2019. "'Automation' of Manufacturing in the Late Nineteenth Century: The Hand and Machine Labor Study." *Journal of Economic Perspectives* 33 (2): 51–70.
- Autor, David 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, David H, Frank Levy, and Richard J Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *The Quarterly Journal of Economics* 118 (4): 1279–1333.
- Babina, Tania, Anastassia Fedyk, Alex Xi He, and James Hodson. 2020. "Artificial Intelligence, Firm Growth, and Industry Concentration." *Working Paper*.

- Bagger, Jesper, and Rasmus Lentz. 2019. "An Empirical Model of Wage Dispersion with Sorting." *The Review of Economic Studies* 86 (1): 153–90.
- Bana, Sarah H. 2021. "Using Language Models to Understand Wage Premia," working paper, .
- Barth, Erling, James C Davis, Richard B Freeman, and Kristina McElheran. 2020. "Twisting the Demand Curve: Digitalization and the Older Workforce." *NBER Working Paper*, no. w28094.
- Barth, Erling, James Davis, and Richard B Freeman. 2018. "Augmenting the Human Capital Earnings Equation with Measures of Where People Work." *Journal of Labor Economics* 36 (S1): S71–97.
- Becker, Gary S. 1981. *A Treatise on the Family*. Harvard University Press.
- . 1993. *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*. 3rd ed. University of Chicago press.
- Benzell, Seth G., Laurence J. Kotlikoff, Guillermo LaGarda, and Jeffrey D. Sachs. 2016. "Robots Are Us: Some Economics of Human Replacement." NBER Working Paper 20941. Working Paper Series. National Bureau of Economic Research, Inc. <https://doi.org/10.3386/w20941>.
- Bessen, James. 2003. "Technology and Learning by Factory Workers: The Stretch-out at Lowell, 1842." *The Journal of Economic History* 63 (1): 33–64.
- . 2012. "More Machines, Better Machines... or Better Workers?" *The Journal of Economic History* 72 (1): 44–74.
- . 2015. *Learning by Doing: The Real Connection between Innovation, Wages, and Wealth*. Yale University Press.
- . 2020. "Information Technology and Industry Concentration." *Journal of Law and Economics* 63 (3): 531–55.
- Bessen, James, Erich Denk, and Chen Meng. 2021. "Firm Differences: Skill Sorting and Software." Working Paper.
- Bessen, James, Maarten Goos, Anna Salomons, and Wiljan Van den Berge. 2022. "Automatic Reaction-What Happens to Workers at Firms That Automate?" *Boston Univ. School of Law, Law and Economics Research Paper*.
- Bessen, James, Stephen Michael Impink, Lydia Reichensperger, and Robert Seamans. 2018. "The Business of AI Startups." *Boston Univ. School of Law, Law and Economics Research Paper*, no. 18–28.
- Bessen, James, Chen Meng, and Erich Denk. 2020. "Perpetuating Inequality: What Salary History Bans Reveal about Wages." *Available at SSRN*.
- Bessen, James, and Cesare Righi. 2019. "Information Technology and Firm Employment." *Boston Univ. School of Law, Law and Economics Research Paper*, no. 19–6. <https://ssrn.com/abstract=3371016>.
- Bloom, Nicholas, Raffaella Sadun, and John Van Reenen. 2012. "Americans Do IT Better: US Multinationals and the Productivity Miracle." *American Economic Review* 102 (1): 167–201.
- Bresnahan, Timothy F, Erik Brynjolfsson, and Lorin M Hitt. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence." *The Quarterly Journal of Economics* 117 (1): 339–76.
- Brynjolfsson, Erik. 2021. "Augmentation, Not Automation." *Boston Review*, May 20, 2021. https://bostonreview.net/forum_response/augmentation-not-automation/.

- Brynjolfsson, Erik, and Lorin M Hitt. 2000. "Beyond Computation: Information Technology, Organizational Transformation and Business Performance." *Journal of Economic Perspectives* 14 (4): 23–48.
- Brynjolfsson, Erik, and Andrew McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. WW Norton & Company.
- Burdett, Kenneth, and Dale T Mortensen. 1998. "Wage Differentials, Employer Size, and Unemployment." *International Economic Review*, 257–73.
- Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. 2018. "Firms and Labor Market Inequality: Evidence and Some Theory." *Journal of Labor Economics* 36 (S1): S13–70.
- Card, David, Jörg Heining, and Patrick Kline. 2013. "Workplace Heterogeneity and the Rise of West German Wage Inequality." *The Quarterly Journal of Economics* 128 (3): 967–1015.
- Carnevale, Anthony P, Tamara Jayasundera, and Dmitri Repnikov. 2014. "Understanding Online Job Ads Data." *Georgetown University, Center on Education and the Workforce, Technical Report (April)*.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. "The Effect of Minimum Wages on Low-Wage Jobs*." *The Quarterly Journal of Economics* 134 (3): 1405–54. <https://doi.org/10.1093/qje/qjz014>.
- Clark, Gregory. 1987. "Why Isn't the Whole World Developed? Lessons from the Cotton Mills." *The Journal of Economic History* 47 (1): 141–73.
- Cortes, Guido Matias, Adrian Lerche, Uta Schönberg, and Jeanne Tschopp. 2020. "Technological Change, Firm Heterogeneity and Wage Inequality." Mimeo (https://www.dropbox.com/s/jcpm28binsgqhbp/Draft_v1_0.pdf).
- Deming, David J. 2017. "The Growing Importance of Social Skills in the Labor Market." *The Quarterly Journal of Economics* 132 (4): 1593–1640.
- Deming, David, and Lisa B Kahn. 2018. "Skill Requirements across Firms and Labor Markets: Evidence from Job Postings for Professionals." *Journal of Labor Economics* 36 (S1): S337–69.
- Dillender, Marcus, and Eliza Forsythe. 2019. "Computerization of White Collar Jobs." Upjohn Institute Working Paper 19-310.
- Domini, Giacomo, Marco Grazzi, Daniele Moschella, and Tania Treibich. 2021. "Threats and Opportunities in the Digital Era: Automation Spikes and Employment Dynamics." *Research Policy* 50 (7): 104137.
- Doms, Mark, Timothy Dunne, and Kenneth R Troske. 1997. "Workers, Wages, and Technology." *The Quarterly Journal of Economics* 112 (1): 253–90.
- Doms, Mark E., and Timothy Dunne. 1998. "Capital Adjustment Patterns in Manufacturing Plants." *Review of Economic Dynamics* 1 (2): 409–29. <https://doi.org/10.1006/redo.1998.0009>.
- Dunne, Timothy, Lucia Foster, John Haltiwanger, and Kenneth R Troske. 2004. "Wage and Productivity Dispersion in United States Manufacturing: The Role of Computer Investment." *Journal of Labor Economics* 22 (2): 397–429.
- Eeckhout, Jan, and Philipp Kircher. 2018. "Assortative Matching with Large Firms." *Econometrica* 86 (1): 85–132.
- Gaggl, Paul, and Greg C Wright. 2017. "A Short-Run View of What Computers Do: Evidence from a Uk Tax Incentive." *American Economic Journal: Applied Economics* 9 (3): 262–94.

- Graetz, Georg, and Guy Michaels. 2018. "Robots at Work." *Review of Economics and Statistics* 100 (5): 753–68.
- Hakanson, Christina, Erik Lindqvist, and Jonas Vlachos. 2020. "Firms and Skills: The Evolution of Worker Sorting." *Journal of Human Resources*.
- Haltiwanger, John, Russell Cooper, and Laura Power. 1999. "Machine Replacement and the Business Cycle: Lumps and Bumps." *American Economic Review* 89 (4): 921–46.
- Harrigan, James, Ariell Reshef, and Farid Toubal. 2021. "The March of the Techies: Job Polarization within and between Firms." *Research Policy* 50 (7): 104008.
- Hémous, David, and Morten Olsen. 2022. "The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality." *American Economic Journal: Macroeconomics* 14 (1): 179–223.
- Hirvonen, Johannes, Aapo Stenhammer, and Joonas Tuhkuri. 2021. "New Evidence on the Effect of Technology on Employment and Skill Demand." Working Paper.
- Humlum, Anders. 2019. "Robot Adoption and Labor Market Dynamics." Working Paper.
- Hunt, Jennifer, and Ryan Nunn. 2019. "Is Employment Polarization Informative About Wage Inequality and Is Employment Really Polarizing?" Working paper. National Bureau of Economic Research.
- Khaouja, Imane, Ghita Mezzour, Kathleen M Carley, and Ismail Kassou. 2019. "Building a Soft Skill Taxonomy from Job Openings." *Social Network Analysis and Mining* 9 (1): 43.
- Korinek, Anton, and Joseph E Stiglitz. 2018. "Artificial Intelligence and Its Implications for Income Distribution and Unemployment." In *The Economics of Artificial Intelligence: An Agenda*, 349–90. University of Chicago Press.
- Kremer, Michael. 1993. "The O-Ring Theory of Economic Development." *The Quarterly Journal of Economics* 108 (3): 551–75.
- Kremer, Michael, and Eric Maskin. 1996. "Wage Inequality and Segregation by Skill." National bureau of economic research.
- Lachowska, Marta, Alexandre Mas, Raffaele D Saggio, and Stephen A Woodbury. 2020. "Do Firm Effects Drift? Evidence from Washington Administrative Data." National Bureau of Economic Research.
- Lindenlaub, Ilse. 2017. "Sorting Multidimensional Types: Theory and Application." *The Review of Economic Studies* 84 (2): 718–89.
- Modestino, Alicia Sasser, Daniel Shoag, and Joshua Ballance. 2019. "Upskilling: Do Employers Demand Greater Skill When Workers Are Plentiful?" *Review of Economics and Statistics*, 1–46.
- Moscarini, Giuseppe, and Fabien Postel-Vinay. 2016. "Wage Posting and Business Cycles: A Quantitative Exploration." *Review of Economic Dynamics*, Special Issue in Honor of Dale Mortensen, 19 (January): 135–60. <https://doi.org/10.1016/j.red.2015.11.001>.
- Olley, G Steven, and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64 (6): 263–97.
- Rodrigo, Rodimiro. 2021. "Robot Adoption, Organizational Capital, and the Productivity Paradox." <https://rodimirodrigo.github.io/JMP2021-Rodrigo.pdf>.
- Rosen, Sherwin. 1981. "The Economics of Superstars." *The American Economic Review* 71 (5): 845–58.
- Sattinger, Michael. 1975. "Comparative Advantage and the Distributions of Earnings and Abilities." *Econometrica: Journal of the Econometric Society*, 455–68.
- . 1993. "Assignment Models of the Distribution of Earnings." *Journal of Economic Literature* 31 (2): 831–80.

- Song, Jae, David J Price, Fatih Guvenen, Nicholas Bloom, and Till Von Wachter. 2019. "Firming up Inequality." *The Quarterly Journal of Economics* 134 (1): 1–50.
- Spitz-Oener, Alexandra. 2006. "Technical Change, Job Tasks, and Rising Educational Demands: Looking Outside the Wage Structure." *Journal of Labor Economics* 24 (2): 235–70.
- Tambe, Prasanna, and Lorin M Hitt. 2012. "The Productivity of Information Technology Investments: New Evidence from IT Labor Data." *Information Systems Research* 23 (3-part-1): 599–617.
- Tambe, Prasanna, Lorin M Hitt, Daniel Rock, and Erik Brynjolfsson. 2019. "IT, AI and the Growth of Intangible Capital." *Available at SSRN 3416289*.

Figures

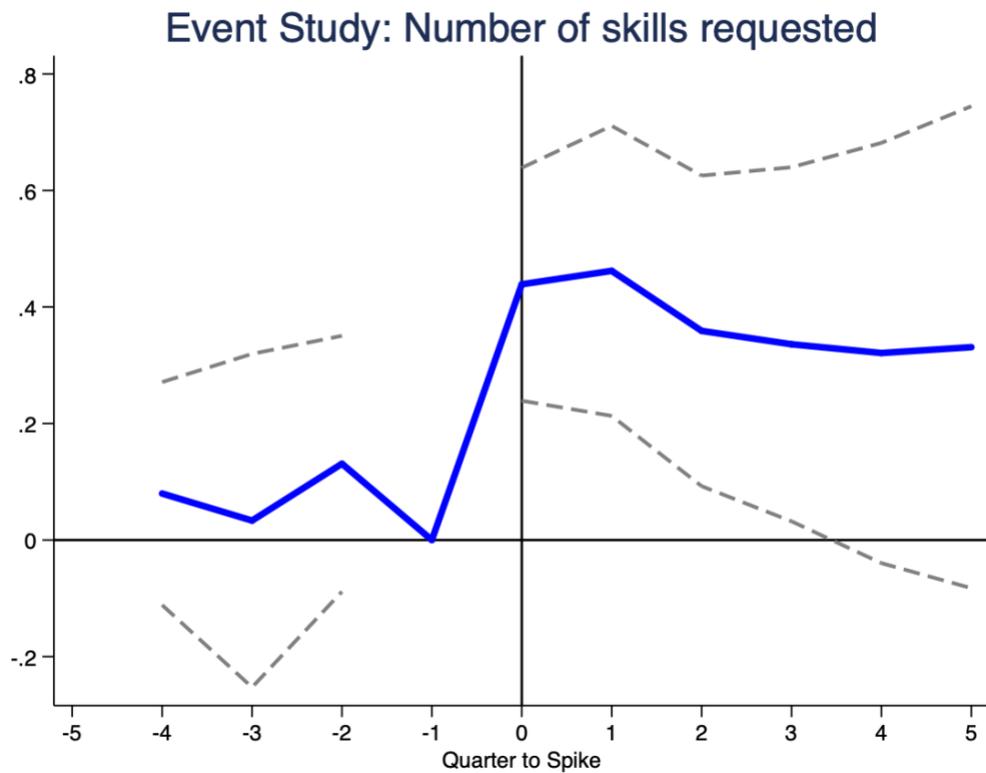


Figure 1. Number of skills requested increases following adoption event.
Note: This figure presents an event study equivalent to Column 1, Panel A, Table 2, reporting the coefficients of quarter dummies for treated firms. The regression is weighted by the number of ads per quarter and it includes fixed effects for quarter and cohort by firm. The dashed lines show the 95% confidence interval with errors clustered by cohort by firm.

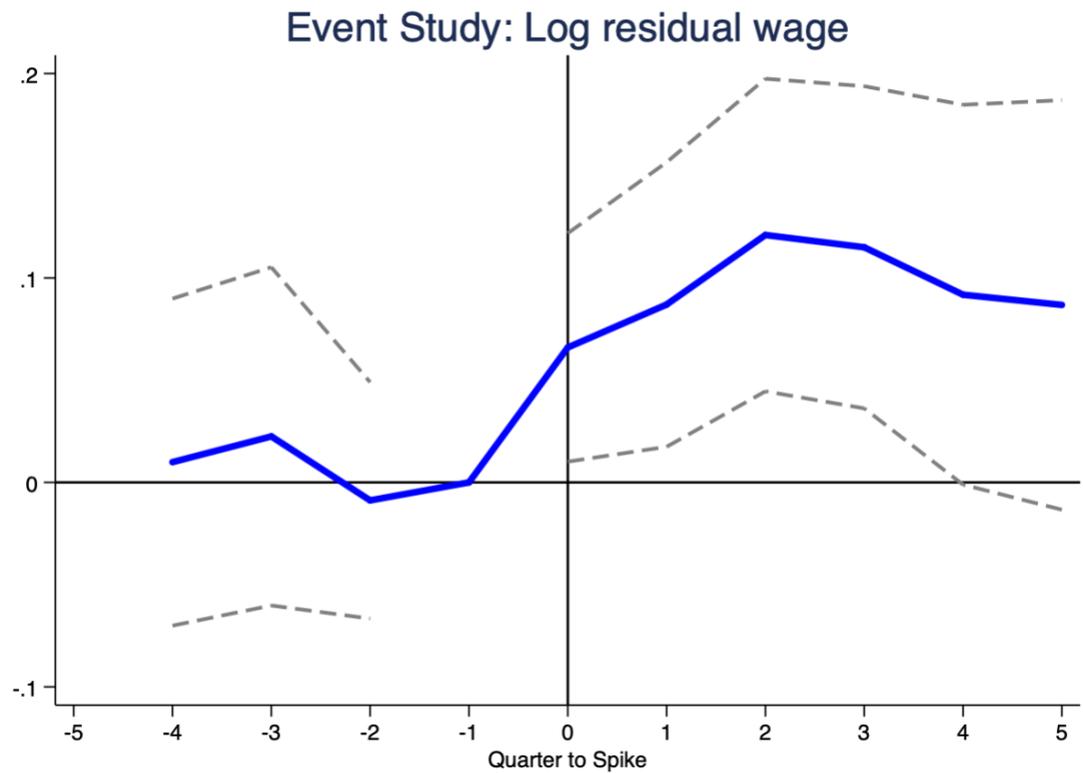


Figure 2. Log residual pay increases following adoption event.
 Note: This figure presents an event study equivalent to Column 1, Panel B, Table 2, reporting the coefficients of quarter dummies for treated firms. The regression is weighted by the number of ads per quarter and it includes fixed effects for quarter and cohort by firm. The dashed lines show the 95% confidence interval with errors clustered by cohort by firm.

Tables

Table 1. Technology Adoption Raises Demands for Specific Skills

Skill measure:	(1) All	(2) IT+AI	(3) Other	(4) Cognitive	(5) Social	(6) Soft	(7) Experience	(8) Education
A. All jobs, number of skills								
Post treatment	0.318*** (0.071)	0.041*** (0.008)	0.173*** (0.059)	0.017*** (0.004)	0.048*** (0.010)	0.038*** (0.010)	0.065*** (0.019)	0.017 (0.020)
Observations	102,086	102,086	102,086	102,086	102,086	102,086	97,045	96,897
R-squared	0.873	0.821	0.868	0.888	0.872	0.868	0.871	0.894
<i>Pre-Spike Means</i>	10.005	0.518	7.437	0.325	0.762	0.962	3.350	14.581
B. Jobs not requiring college diplomas, number of skills								
Post treatment	0.222*** (0.070)	0.031*** (0.009)	0.105* (0.058)	0.010** (0.004)	0.040*** (0.010)	0.037*** (0.011)	0.041* (0.022)	0.035 (0.024)
Observations	95,679	95,679	95,679	95,679	95,679	95,679	87,220	86,775
R-squared	0.840	0.696	0.843	0.833	0.838	0.826	0.808	0.853
C. All Jobs, Share of skills								
Post treatment		0.002* (0.001)	-0.008*** (0.002)	0.001 (0.000)	0.002*** (0.001)	0.004** (0.002)		
Observations		102,086	102,086	102,086	102,086	102,086		
R-squared		0.854	0.847	0.853	0.857	0.755		

Note: these coefficients are from stacked difference-in-differences regressions where a balanced panel (t-5 to t+5) is included for each cohort based on spiking year. The unit of observation is firm by quarter. All firms in the sample spike at some time during the sample period and only observations are included that have not spiked previously. All regressions include controls for labor market tightness, management job share, the outsourceable job share, time and cohort x firm fixed effects and standard errors are clustered by cohort x firm (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). To treat heteroscedasticity arising from sample variance, regressions are weighted by the number of help-wanted ads for each firm-quarter. The top panel includes counts of skills requested on all jobs; the bottom panel counts skills only in occupations where the majority of jobs do not require a college diploma. IT jobs (SOC 15) are excluded from the regressions.

Table 2. Adoption of Technology Raises Skill Demands and Pay Across Skill Groups

Skill group:	(1) All	(2) College not required	(3) Routine Cognitive	(4) Routine Manual	(5) Nonroutine Cognitive	(6) Nonroutine Manual
A. Dependent variable: number of specific skills requested						
Post treatment	0.318*** (0.071)	0.222*** (0.070)	0.398*** (0.087)	0.376*** (0.105)	0.512*** (0.091)	0.153 (0.144)
Observations	102,086	95,679	97,117	69,798	100,449	62,967
R-squared	0.873	0.840	0.803	0.771	0.816	0.732
B. Dependent variable: Log Residual Pay						
Post treatment	0.087*** (0.023)	0.054** (0.024)	0.067** (0.029)	0.067* (0.037)	0.091*** (0.032)	0.023 (0.031)
Observations	29,437	21,073	15,617	10,820	20,092	9,345
R-squared	0.476	0.557	0.543	0.622	0.473	0.627

Note: these coefficients are from stacked difference-in-differences regressions where a balanced panel (t-5 to t+5) is included for each cohort based on spiking year. The unit of observation is firm by quarter. All firms in the sample spike at some time during the sample period and only observations are included that have not spiked previously. All regressions include controls for labor market tightness, management job share, the outsourceable job share, time and cohort x firm fixed effects and standard errors are clustered by cohort x firm (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). To treat heteroscedasticity arising from sample variance, regressions are weighted by the number of help-wanted ads for each firm-quarter. The dependent variable in the top panel is the total number of skills requested per ad; the dependent variable in the bottom panel is the log residual salary offered after controlling for experience, experience squared, education, detailed occupation, state, year, and a measure of labor market tightness. IT jobs (SOC 15) are excluded from the dependent variables.

Table 3. Skill and Pay Treatment Effects are Robust to Controls

Dependent Variable	Number of Skills Requested			Log Residual Pay		
	(1)	(2)	(3)	(4)	(5)	(6)
Post treatment	0.245** (0.117)	0.214* (0.112)	0.213** (0.107)	0.102*** (0.035)	0.102*** (0.035)	0.100*** (0.036)
Labor market tightness		0.284 (1.064)	0.194 (1.097)		-0.849 (0.580)	-0.802 (0.591)
Management jobs		6.653*** (0.676)	6.595*** (0.650)		-0.272 (0.221)	-0.271 (0.201)
Outsourceable jobs		-7.210*** (1.793)	-7.188*** (1.768)		0.050 (0.314)	0.004 (0.316)
Growth Rate of Sales			0.260* (0.154)			0.071 (0.060)
Lag CEO change			-0.966 (1.032)			-0.087* (0.051)
3 rd order productivity polynomial			✓			✓
Polynomial probability value			0.013			0.023
Observations	14,008	14,008	14,008	4,706	4,706	4,706
R-squared	0.873	0.882	0.884	0.461	0.465	0.468

Note: these coefficients are from stacked difference-in-differences regressions where a balanced panel (t-5 to t+5) is included for each cohort based on spiking year. The unit of observation is firm by quarter. All firms in the sample spike at some time during the sample period and only observations are included that have not spiked previously. The sample in this table includes only firms that have been matched to Compustat in order to include additional control variables. All regressions include time and cohort x firm fixed effects and standard errors are clustered by cohort x firm (***) p<0.01, ** p<0.05, * p<0.1). To treat heteroscedasticity arising from sample variance, regressions are weighted by the number of help-wanted ads for each firm-quarter. The dependent variable in the first three columns is the total number of skills requested per ad; the dependent variable in columns 4-6 is the log residual salary offered after controlling for experience, experience squared, education, detailed occupation, state, year, and a measure of labor market tightness. The polynomial used in columns 3 and 6 includes log real cost of goods sold and log real beginning-of-quarter capital. The probability value reported is for the F-test of the null hypothesis that polynomial coefficients are jointly zero. IT jobs (SOC 15) are excluded from the dependent variables.

Table 4: Technology Adoption and Changes in Hiring

Skill Group:	(1) College required	(2) College not required	(3) Routine Cognitive	(4) Routine Manual	(5) Nonroutine Cognitive	(6) Nonroutine Manual
A. Share of Hiring						
Post treatment	0.017*** (0.002)	-0.017*** (0.002)	0.007*** (0.003)	-0.008*** (0.002)	0.021*** (0.002)	-0.006*** (0.002)
Observations	103,547	103,547	103,594	103,594	103,594	103,594
R-squared	0.963	0.963	0.910	0.964	0.957	0.970
B. Log level of Hiring						
Post treatment	0.018 (0.030)	-0.083** (0.033)	0.035 (0.032)	-0.107*** (0.040)	0.029 (0.030)	-0.026 (0.048)
Observations	103,404	103,413	97,567	71,018	100,747	64,290
R-squared	0.920	0.927	0.925	0.925	0.923	0.923

Note: these coefficients are from stacked difference-in-differences regressions where a balanced panel (t-5 to t+5) is included for each cohort based on spiking year. The unit of observation is firm by quarter. All firms in the sample spike at some time during the sample period and only observations are included that have not spiked previously. All regressions include controls for labor market tightness, management job share, the outsourceable job share, time and cohort x firm fixed effects and standard errors are clustered by cohort x firm (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). To treat heteroscedasticity arising from sample variance, regressions are weighted by the number of help-wanted ads for each firm-quarter. The columns designate different skill groups. The dependent variable in the top panel is the group's share of job ads; the dependent variable in the bottom panel is the log of the number of job ads. IT jobs (SOC 15) are excluded from the dependent variables.

Table 5. Information Technology Accounts for Most of the Correlation Between Firm Fixed Effects and Skills

Skill measure:	(1) All	(2) IT+AI	(3) Other	(4) Cognitive	(5) Social	(6) Soft	(7) Experience	(8) Education
Panel A, simple correlation								
Firm FE	2.619*** (0.358)	0.722*** (0.069)	1.338*** (0.234)	0.212*** (0.024)	0.203*** (0.065)	0.144** (0.071)	1.500*** (0.133)	2.912*** (0.353)
Standardized coefficient	<i>0.176</i>	<i>0.246</i>	<i>0.125</i>	<i>0.207</i>	<i>0.111</i>	<i>0.067</i>	<i>0.284</i>	<i>0.203</i>
Observations	205,306	205,306	205,306	205,306	205,306	205,306	205,306	205,306
R-squared	0.031	0.060	0.016	0.043	0.012	0.004	0.081	0.041
Panel B, software controls								
Firm FE	0.613** (0.295)	0.075*** (0.021)	0.557** (0.222)	0.074*** (0.019)	-0.016 (0.059)	-0.078 (0.064)	0.545*** (0.100)	1.330*** (0.323)
Standardized coefficient	<i>0.041</i>	<i>0.026</i>	<i>0.052</i>	<i>0.072</i>	<i>-0.009</i>	<i>-0.036</i>	<i>0.103</i>	<i>0.093</i>
Software share	36.738*** (1.436)	5.601*** (0.170)	20.033*** (1.084)	2.683*** (0.092)	4.105*** (0.223)	4.315*** (0.261)	13.673*** (0.446)	31.272*** (1.203)
Software share ²	-54.94*** (2.254)	-0.035 (0.296)	-37.63*** (1.713)	-4.21*** (0.146)	-6.26*** (0.365)	-6.79*** (0.428)	-15.34*** (0.736)	-49.85*** (1.887)
Observations	205,306	205,306	205,306	205,306	205,306	205,306	205,306	205,306
R-squared	0.255	0.760	0.125	0.278	0.192	0.143	0.454	0.202
SW share of sorting	77%	89%	58%	65%	108%	154%	64%	54%

Note: This table regresses firm mean levels of skill counts, experience and education required against firm wage fixed effects. The unit of observation is the firm. Firm fixed effects are calculated by regressing log salary offered against detailed occupation, industry, state, year, labor market tightness, skills requested, education required, experience required, and firm fixed effects. IT jobs are excluded for the estimates. The regressions are weighted by the number of job ads and errors are robust to heteroscedasticity. The bottom panel ads controls for the share of software developers in firm hiring. The standardized coefficients reflect the correlations between the dependent variables and firm fixed effects. Adding controls for software developers substantially reduces these correlations. The bottom row displays the magnitude of that decrease as one minus the standardized coefficient in Panel B over the standardized coefficient in Panel A.

Appendix

A. Model

Sorting equilibrium

We can write the first order condition for q_i , holding the quality of other tasks, q_j , constant as

$$q_j^{N-I-1}V - \theta h'(q_i)l_i = 0.$$

Taking the implicit derivative,

$$\frac{dq_i}{d\theta} = -\frac{h'(q_i)}{\theta h''(q_i)} < 0.$$

The equilibrium value of q decreases with θ . From this it follows that

$$\frac{d \theta h'(q_i)q_i}{d\theta} = h'(q_i) - \theta(h''(q_i)q_i + h'(q_i))\frac{dq_i}{d\theta} = -\frac{(h'(q_i))^2}{h''(q_i)} < 0.$$

Since, as in the text, $w_j = \theta_j h'(q_i)q_i$, the fact that $\theta_H < \theta_L$ implies that $w_H > w_L$ in equilibrium.

Change in between-firm wage ratio

It is convenient to express output in intensive form,

$$y \equiv \frac{Y}{L} = A \cdot Q \cdot k^\alpha, \quad k \equiv \frac{K}{L}$$

so that the first order profit maximizing condition for labor and capital can be written

$$w = (1 - \alpha)y, \quad k = \frac{\alpha}{r}y.$$

Using these, we have³¹

$$\Delta \ln \omega = \Delta \ln(1 - \alpha_H) + \Delta \ln \frac{y_H}{y_L} \approx -\frac{1}{N-I} + \Delta \ln \frac{y_H}{y_L}.$$

Further,

³¹ α_H increases from $\frac{I-1}{N}$ to $\frac{I}{N}$.

$$\Delta \ln \frac{y_H}{y_L} > \Delta \ln A_H + \Delta \ln Q_H + \alpha_L \Delta \ln \frac{k_H}{k_L}.$$

The last term, which did not appear in the case of uniform workers and firms, captures the shift in capital from low type firms to high type firms as the productivity of the high type firms rises, raising the returns for capital per worker. The expression is an inequality because it ignores the increase in α for high type firms. Also, using the first order condition for capital,

$$\Delta \ln \frac{k_H}{k_L} = \Delta \ln \alpha_H + \Delta \ln \frac{y_H}{y_L} \approx \frac{1}{I} + \Delta \ln \frac{y_H}{y_L}.$$

Substituting this into the previous expression,

$$\Delta \ln \frac{y_H}{y_L} > \frac{1}{1 - \alpha_L} \left[\Delta \ln A_H + \Delta \ln Q_H + \frac{1}{I} \right]$$

and

$$\Delta \ln \omega > -\frac{1}{N - I} + \frac{1}{1 - \alpha_L} \left[\Delta \ln A_H + \Delta \ln Q_H + \frac{1}{I} \right].$$

B. Skill measures

Burning Glass standardizes specific skills requested into 16,050 skills. For our analysis, we constructed 6 mutually exclusive skill categories: IT, AI, cognitive, social, other soft skills, and an additional “other” category. We begin with the definition of social and cognitive skills used by Deming and Khan (2018). We then assign IT, AI, and other soft skills using lists of skill terms not included in the Deming and Khan categories. This last category is the largest and contains many skills related to non-IT technologies and to industry knowledge. For our main analysis, we combine the AI and IT categories, but separate analysis indicates that spikes at firms that hire AI personnel perform much like firms that apparently use non-AI software methods (see Table A7 below). The frequencies with which ads request skills in each category are

<u>Category</u>	<u>Percent of job ads</u>
Other	68.56
IT	13.08
Other soft	8.18
social	6.92
cognitive	3.18
AI	0.08

Cognitive Skills (D. Deming and Kahn 2018)

These skills include the keywords Problem Solving, Research, Analytical, Critical Thinking, Math, and Statistics.

Social Skills (D. Deming and Kahn 2018)

These skills include the keywords Communication, Teamwork, Collaboration, Negotiation, and Presentation.

Other Soft Skills* Keywords (adapted from Khaouja et al. (2019) taxonomy):

Accountability	Ethic	Social skills
Active listening	Flexibility	Speaking
Adaptive	Goal	Strategic thinking
Argumentation	Hospitality	Time management
Coaching	Impartiality	Trustworthy
Commitment	Influence	Verbal communication
Conceptual	Initiative	Writing
Conflict management	Integrity	Written communication
Coordination	Interpersonal communication	
Creativity	Kindness	
Curiosity	Leadership	
Decision	Mentoring	
Decision making	Motivated	
Detail	Optimism	
Diverse	Passion	
Eagerness	Persuasion	
Emotional intelligence	Self-confidence	
Enthusiasm	Self-organized	

* These skills also have synonyms, which were also flagged. For full list of synonyms, please refer to Table 13 in Khaouja et al 2019. To further augment this list, the following commonly requested Burning Glass skills not already identified as a social skill were also flagged as soft skills: Planning, Detail-Oriented, Building Effective Relationships, Energetic, Positive Disposition, Listening, Team Building, Creative Problem Solving, Self-Motivation, Overcoming Obstacles, Multi-Tasking, People Management, Thought Leadership, Team Management. This list excludes skills already identified as social or cognitive skills above.

Other Skills

Skills that do not belong to one of the other five groups are designated as “other”. These skills tend to be industry-specific or firm specific. A majority of skills fit in this category. Examples include 5G Wireless, ACL Surgery, Adhesives Industry Knowledge, and APA Style Guide.

AI Skills (Following Alekseeva et al. (2020))

AI ChatBot	Latent Semantic Analysis	OpenNLP
AI KIBIT	Lexalytics	Pattern Recognition
ANTLR	Lexical Acquisition	Pybrain
Apertium	Lexical Semantics	Random Forests
Artificial Intelligence	Libsvm	Recommender Systems
Automatic Speech Recognition (ASR)	Machine Learning	Semantic Driven Subtractive Clustering Method (SDSCM)
Caffe Deep Learning Framework	Machine Translation (MT)	Semi-Supervised Learning
Chatbot	Machine Vision	Sentiment Analysis / Opinion Mining
Computational Linguistics	Madlib	Sentiment Classification
Computer Vision	Mahout	Speech Recognition
Decision Trees	Microsoft Cognitive Toolkit	Supervised Learning (Machine Learning)
Deep Learning	MLPACK (C++ library)	Support Vector Machines (SVM)
Deeplearning4j	Mlpy	TensorFlow
Distinguo	Modular Audio Recognition Framework (MARF)	Text Mining
Google Cloud Machine Learning Platform	MoSes	Text to Speech (TTS)
Gradient boosting	MXNet	Tokenization
H2O (software)	Natural Language Processing	Torch (Machine Learning)
IBM Watson	Natural Language Toolkit (NLTK)	Unsupervised Learning
Image Processing	ND4J (software)	Virtual Agents
Image Recognition	Nearest Neighbor Algorithm	Vowpal
IPSoft Amelia	Neural Networks	Wabbit
Ithink	Object Recognition	Word2Vec
Keras	Object Tracking	
Latent Dirichlet Allocation	OpenCV	

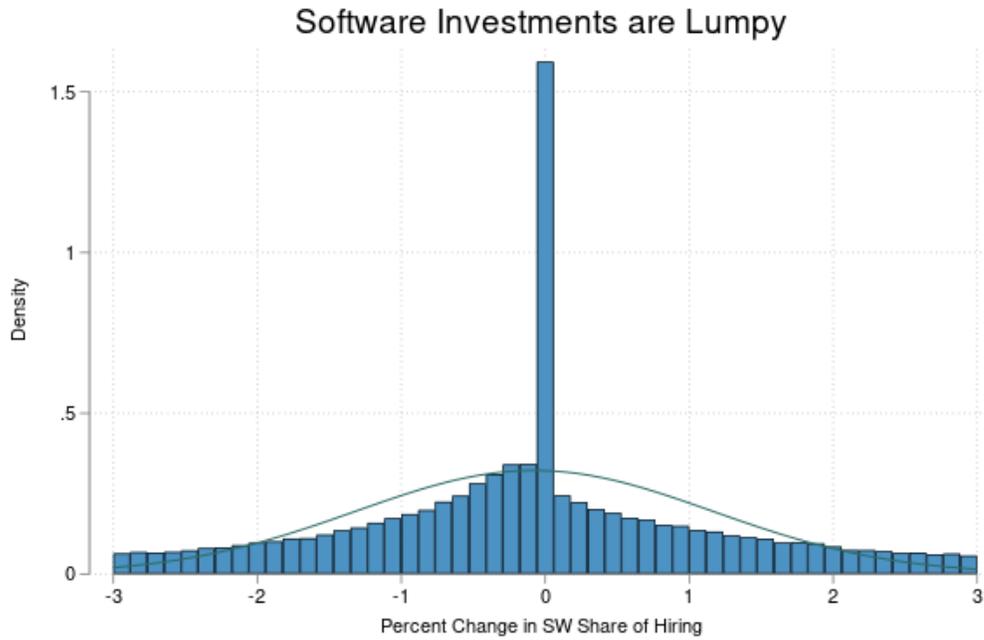
IT Skills (Following Burning Glass Technologies Skill Cluster Families)

Microsoft Development Tools	Enterprise Content Management (ECM)	Productivity Software
Document Management Systems	Internet of Things (IoT)	File Transfer Software
General Networking	Enterprise Management Software	Project Management Software
Software Quality Assurance	Database Administration	Virtual Private Networks
Artificial Intelligence	Android Development	Internet Standards
Operating Systems	Mobile Development	Remote Desktop Software
JavaScript and jQuery	IT Automation	Data Wrangling
Distributed Computing	Configuration Management	Programming Principles
Application Programming Interface (API)	Anti-Malware Software	Network File System (NFS)
Systems Administration	Middleware	Integrated Development Environments (IDEs)
Web Development	Scripting	Disk Imaging
Scripting Languages	Java	Microsoft Office and Productivity Tools
Cloud Solutions	Database Management Systems	Content Management Systems
Cloud Computing	Web Servers	Firewall Software
Software Development Tools	Version Control	Firmware
Data Storage	iOS Stack	Graph Databases
Virtual Machines (VM)	Basic Computer Knowledge	Identity Management
Big Data	Application Development	Partitioning Software
Network Security	Network Protocols	Video Conferencing Software
Data Warehousing	Technical Support	Computer Hardware
Enterprise Messaging	Application Security	Internet Services
Cloud Storage	Typesetting Software	Internet Security
XML Markup Languages	Geographic Information System (GIS) Software	Help Desk Support
Extraction, Transformation, and Loading (ETL)	Data Compression	Management Information System (MIS)
System Design and Implementation	Assembly Languages	Intelligent Maintenance Systems
Network Configuration	Test Automation	Query Languages
Data Synchronization	Telecommunications	Load Balancing
Other Programming Languages	Compiling Tools	Location-based Software
Data Management	Enterprise Resource Planning (ERP)	Video Compression Standards
Web Content	Backup Software	Microsoft SQL Extensions
SAP	Web Design	Advanced Microsoft Excel
Archiving Software	Rule Engines	SQL Databases and Programming
Cybersecurity	Internet Protocols	Device Management
NoSQL Databases	Extensible Languages	Microsoft Windows
Software Development Principles	C and C++	Augmented Reality / Virtual Reality (AR / VR)
IT Management	Desktop and Service Management	Enterprise Information Management
Software Development Methodologies	Mainframe Technologies	Oracle
Content Delivery Network (CDN)	Parallel Computing	Servers
Networking Hardware	Cache (computing)	Data Collection
Information Security	PHP Web	Wiki

Note: There are 1,687 unique skills that Burning Glass identifies as Information Technology skills. From there, they sort these skills into broader categories, which are listed in the table below. Within the category “Microsoft Development Tools” is the Microsoft Office suite, which we omit as an IT skill. We exclude skills flagged as social, cognitive or AI skills. These specific skills include Communications Protocols, Data Communications, Global System for Mobile Communications, Joint Worldwide Intelligence Communications System, Machine-To-Machine (M2M) Communications, Oracle Fusion Middleware Collaboration Suite, and Voice Communications.

C. Lumpy Investment

Figure A1. Lumpiness of Firm Investments



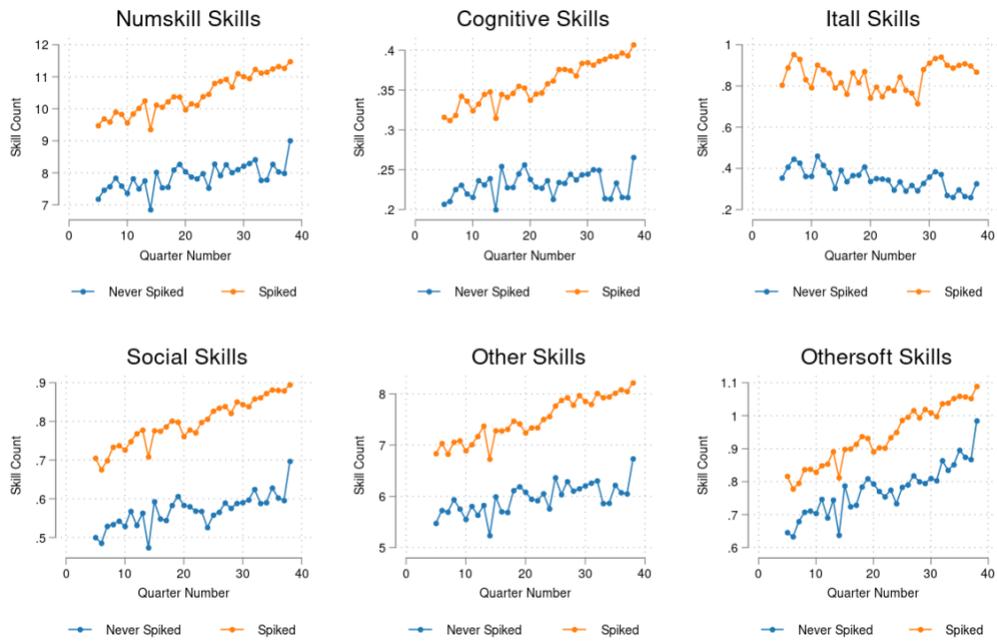
Note: This figure plots changes in software developer share of job advertisements from the average of the previous 4 quarters. The line shows a normal density distribution with the same mean and standard deviation. The distribution is clearly leptokurtic with a peak at zero and fat, “lumpy” tails.

Figure A2. Software Hiring Increases Persist After Spikes



Note: This figure plots an event study of the share of software hiring around hiring spikes. There appears to be a slight anticipation effect, a distinct spike (the threshold is .01), and sustained hiring of software developers at a slightly lower level after the spike.

Figure A3. Skill request trends over time



Note: This figure shows raw trends in skill requests for both spiking (orange) and non-spiking (blue) firms over time. Spiking firms have higher levels of skill requests throughout the sample.

D. Descriptive Statistics and Robustness Checks

Table A1 Summary Statistics

	(1)	(2)	(3)
Sample:	Full sample	Never-Spikers	Spikers
Weighted			
Management Job Share	0.126 (0.190)	0.120 (0.219)	0.139 (0.0960)
Outsourceable Job Share	0.071 (0.183)	0.078 (0.209)	0.056 (0.101)
Labor Market Tightness	0.795 (0.319)	0.837 (0.364)	0.700 (0.139)
IT Share	0.095 (0.160)	0.074 (0.166)	0.108 (0.155)
Residual Wage	0.012 (0.291)	-0.002 (0.336)	0.023 (0.250)
College Required	0.433 (0.279)	0.416 (0.303)	0.471 (0.213)
Routine Cognitive	0.298 (0.284)	0.294 (0.325)	0.307 (0.157)
Routine Manual	0.207 (0.304)	0.224 (0.339)	0.170 (0.201)
Non-Routine Cognitive	0.444 (0.343)	0.423 (0.377)	0.490 (0.243)
Non-Routine Manual	0.158 (0.285)	0.177 (0.320)	0.115 (0.177)
Number of Skills	8.230 (4.895)	7.385 (5.210)	10.062 (3.484)
Unweighted			
Number of Ads/Quarter	85.380 (86.637)	5.980 (1.028)	164.780 (47.623)
Total Firms	2,147,578	2,131,972	15,606

Note: Means given with Standard Deviation in parentheses. Weighted estimates use analytical weights by number of job advertisements.

Table A2. Correlations of Software Spikes and Possibly Correlated Variables

	Lagged Independent Variables				
	(1)	(2)	(3)	(4)	(5)
Panel A. All Firms					
Log Job Ads	0.034*** (0.001)				0.035*** (0.001)
Software share		-0.011 (0.009)			0.036*** (0.008)
Outsourceable jobs			-0.042*** (0.013)		-0.062*** (0.013)
Management jobs				0.034*** (0.011)	0.056*** (0.010)
Observations	89,928	89,928	89,928	89,928	89,928
R-squared	0.023	0.000	0.000	0.000	0.025
Panel B. Compustat					
Labor Productivity	0.006* (0.003)				0.016*** (0.004)
Log COGS		0.014*** (0.002)			
Log Capital			0.008*** (0.002)		0.014*** (0.002)
Sales Growth				0.017 (0.012)	0.028** (0.012)
Observations	14,122	14,122	14,122	14,122	14,122
R-squared	0.001	0.006	0.003	0.000	0.007

Note: This table presents simple OLS regressions between a spike and lagged key variables from both Burning Glass and Compustat. All standard errors are clustered at the firm level. (***) $p < 0.01$, (**) $p < 0.05$, (*) $p < 0.1$)

Table A3 Sensitivity Table

	Panel Size			Spike Threshold		
	(1)	(2)	(3)	(4)	(5)	(6)
	t ± 4	t ± 5	t ± 6	.005	.01	.015
A. Dependent variable: number of specific skills requested						
Post treatment	0.283*** (0.076)	0.318*** (0.071)	0.470*** (0.087)	0.356*** (0.068)	0.318*** (0.071)	0.253*** (0.074)
Labor market tightness	-0.229 (0.263)	-0.176 (0.346)	-0.197 (0.444)	0.264 (0.358)	-0.176 (0.346)	-0.173 (0.363)
Management jobs	4.074*** (0.329)	5.488*** (0.386)	5.516*** (0.576)	4.821*** (0.574)	5.488*** (0.386)	4.916*** (0.334)
Outsourceable jobs	-5.701*** (0.951)	-7.183*** (1.310)	-7.799*** (1.655)	-6.707*** (1.287)	-7.183*** (1.310)	-6.324*** (1.026)
Observations	162,924	102,086	61,377	102,520	102,086	98,609
R-squared	0.892	0.873	0.870	0.888	0.873	0.879
B: Dependent variable: Log Residual Pay						
Post treatment	0.078*** (0.022)	0.087*** (0.023)	0.072*** (0.024)	0.074*** (0.024)	0.087*** (0.023)	0.253*** (0.074)
Labor market tightness	-0.133 (0.128)	-0.279* (0.147)	-0.179 (0.166)	-0.326** (0.130)	-0.279* (0.147)	-0.173 (0.363)
Management jobs	-0.091 (0.091)	0.016 (0.101)	0.198 (0.151)	0.055 (0.107)	0.016 (0.101)	4.916*** (0.334)
Outsourceable jobs	-0.098 (0.126)	0.026 (0.129)	0.270 (0.290)	-0.105 (0.149)	0.026 (0.129)	-6.324*** (1.026)
Observations	42,387	29,437	19,395	28,724	29,437	28,924
R-squared	0.522	0.476	0.411	0.462	0.476	0.450

Note: This table shows how estimates change from changing the size of the balanced panel or threshold for defining a spike. Columns (2) and (5) correspond to estimates in Table 2 Column (1). Construction of panels and additional controls follow those described in Table 2. The unit of observation is firm by quarter. All firms in the sample spike at some time during the sample period and only observations are included that have not spiked previously. All regressions include time and cohort x firm fixed effects and standard errors are clustered by cohort x firm (***) p<0.01, ** p<0.05, * p<0.1).

Table A4 Results for Full Sample And Results for Sample Restricted to Later-Spiking Firms

Sample	(1)	(2)	(3)	(4)
	Number of Skills Requested		Log Residual Wage	
	Later-spiking	Full Sample	Later-spiking	Full Sample
Post treatment	0.318*** (0.071)	0.211*** (0.068)	0.087*** (0.023)	0.074*** (0.020)
Labor market tightness	-0.176 (0.346)	-0.105 (0.115)	-0.279* (0.147)	-0.149*** (0.043)
Management jobs	5.488*** (0.386)	3.249*** (0.106)	0.016 (0.101)	-0.159*** (0.036)
Outsourceable jobs	-7.183*** (1.310)	-2.967*** (0.270)	0.026 (0.129)	0.010 (0.045)
Observations	102,086	1,789,706	29,437	387,844
R-squared	0.873	0.890	0.476	0.513

Note: Our main analysis uses panels with control firms that spike subsequently (“later-spiking”). This table compares this sample with a sample that also includes control firms that never spike. Columns (1) and (3) correspond to Column (1) in Table 2, estimating stacked difference-in-differences regressions where a balanced panel (t-5 to t+5) is included for each cohort based on spiking year. The unit of observation is firm by quarter. All regressions include time and cohort x firm fixed effects and standard errors are clustered by cohort x firm (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). In Columns (1) and (3) firms in the sample spike at some time during the sample period and only observations are included that have not spiked previously. In Columns (2) and (4) we remove this restriction, consequently broadening our sample size. The estimates are similar, but we prefer the estimates provided in the main text.

Table A5 Firm Fixed Effects

	(1) Log of Avg Salary
Other Skill Count	0.003*** (0.000)
Cognitive Count	0.006*** (0.000)
Social Count	0.007*** (0.000)
AI Count	0.035*** (0.002)
IT Count	0.012*** (0.000)
Other Soft Count	0.005*** (0.000)
Minimum of the required experience range in years	0.098*** (0.000)
Experience Required Squared	-0.005*** (0.000)
V/U Labor Market Tightness	-0.001 (0.001)
Observations	4,075,295
R-squared	0.688

Note: This table presents the coefficients used to estimate firm fixed effects. All regressions include occupation, education level, year, and state fixed effects and standard errors are heteroskedastic robust (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). Observations are weighted by occupation share in the Current Population Survey.

Table A6 Non-IT Producing Firms

Sample	(1)	(2)	(3)	(4)
	Number of Skills Requested		Log Residual Wage	
	Full	Non-IT	Full	Non-IT
Post treatment	0.318*** (0.071)	0.358*** (0.078)	0.087*** (0.023)	0.091*** (0.025)
Labor market tightness	-0.176 (0.346)	-0.094 (0.363)	-0.279* (0.147)	-0.294* (0.151)
Management jobs	5.488*** (0.386)	5.900*** (0.430)	0.016 (0.101)	-0.005 (0.107)
Outsourceable jobs	-7.183*** (1.310)	-7.132*** (1.392)	0.026 (0.129)	0.025 (0.135)
Observations	102,086	84,261	29,437	25,597
R-squared	0.873	0.879	0.476	0.480

Note: This table compares the outcomes from Table 2 Column (1) to the same specification excluding IT-producing industries. We defined IT-producing industries as 2-digit NAICS codes 51 and 54. To determine a firm's industry from Burning Glass, we assigned the modal 2-digit industry listed in a firm-year. Columns (1) and (3) correspond to Column (1) in Table 2, estimating stacked difference-in-differences regressions where a balanced panel (t-5 to t+5) is included for each cohort based on spiking year. The unit of observation is firm by quarter. All regressions include time and cohort x firm fixed effects and standard errors are clustered by cohort x firm (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table A7. Firms Using AI and Automation Behave Similarly

VARIABLES	(1) Number of skills requested	(2) Log Residual Wage	(3) Number of skills requested	(4) Log Residual Wage
Non-AI x post treatment	0.330*** (0.089)	0.074*** (0.024)		
AI x post treatment	0.304*** (0.081)	0.096*** (0.027)		
Non-automation x post treatment			0.040 (0.086)	0.100*** (0.037)
Automation x post treatment			0.360*** (0.074)	0.086*** (0.023)
Labor market tightness	-0.177 (0.346)	-0.277* (0.146)	-0.174 (0.345)	-0.280* (0.147)
Management jobs	5.492*** (0.387)	0.012 (0.102)	5.467*** (0.386)	0.017 (0.101)
Outsourceable jobs	-7.179*** (1.308)	0.026 (0.129)	-7.199*** (1.306)	0.026 (0.129)
Observations	102,086	29,437	102,086	29,437
R-squared	0.873	0.476	0.873	0.476

Note: these coefficients are from stacked difference-in-differences regressions where a balanced panel (t-5 to t+5) is included for each cohort based on spiking year. The unit of observation is firm by quarter. All firms in the sample spike at some time during the sample period and only observations are included that have not spiked previously. All regressions include time and cohort x firm fixed effects and standard errors are clustered by cohort x firm (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). To treat heteroscedasticity arising from sample variance, regressions are weighted by the number of help-wanted ads for each firm-quarter. IT jobs (SOC 15) are excluded from the regressions. AI and automation are identified by keywords for skills requested.

Table A8. Placebo: Spikes of engineers and technicians do not display similar effects. Spikes defined for engineers (SOC 17) and technicians (SOC 19) excluding electrical engineers (SOC 172071)

VARIABLES	(1) Number of skills requested	(2) Log Residual Wage
Post treatment	0.094 (0.065)	0.032 (0.033)
Labor market tightness	-0.262 (0.321)	-0.129 (0.127)
Management jobs	5.136*** (0.393)	-0.245 (0.163)
Outsourceable jobs	-6.039*** (0.655)	-0.299*** (0.100)
Observations	97,526	28,920
R-squared	0.884	0.464

Note: these coefficients are from stacked difference-in-differences regressions where a balanced panel (t-5 to t+5) is included for each cohort based on spiking year. The unit of observation is firm by quarter. All firms in the sample spike at some time during the sample period and only observations are included that have not spiked previously. All regressions include time and cohort x firm fixed effects and standard errors are clustered by cohort x firm (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). To treat heteroscedasticity arising from sample variance, regressions are weighted by the number of help-wanted ads for each firm-quarter. IT jobs (SOC 15) are excluded from the regressions.

Table A9. Tests of Pre-trends

F tests of the null hypothesis that event study coefficients are jointly zero prior to the spike, $\delta_{t-2} = \delta_{t-3} = \delta_{t-4} = 0$.

Outcome variable	Probability value
Log residual wage	0.723
<u>Skill measures</u>	
All	0.633
IT+AI	0.371
Other	0.553
Cognitive	0.359
Social	0.196
Soft	0.941
Experience	0.972
Education	0.709

Note: These event study regressions are weighted by the number of ads per quarter and they include fixed effects for quarter and cohort by firm.

The Remainder Effect: How Automation Complements Labor Quality

James Bessen, Erich Denk, and Chen Meng

2/24/2022

Abstract:

This paper argues that automation both complements and replaces workers. Extending the Acemoglu-Restrepo model of automation to consider labor quality, we obtain a Remainder Effect: while automation displaces labor on some tasks, it raises the returns to skill on remaining tasks across skill groups. This effect increases between-firm pay inequality while labor displacement affects within-firm inequality. Using job ad data, we find firm adoption of information technologies leads to both greater demand for diverse skills and higher pay across skill groups. This accounts for most of the sorting of skills to high paying firms that is central to rising inequality.

JEL: J31, O33, J23

Keywords: automation, income inequality, skills, information technology, software

Bessen and Denk, TPRI, Boston University School of Law; Meng, Kean University. Thanks to Luise Eisfeld, Maarten Goos, Po-Hsuan Hsu, Eric Maskin, Mike Meurer, Felix Pöge, Ronja Röttger, Anna Salomons and participants in TPRI's seminar for helpful comments.

Introduction

The skill-biased technical change (SBTC) hypothesis holds that technology complements some groups of workers. In contrast, recent economic models of automation posit that automation technologies strictly substitute machines for workers.¹ Labor displacement is seen by some as the main source of growing economic inequality over the last four decades (Acemoglu and Restrepo 2021), leading to calls for redistribution (Korinek and Stiglitz 2018; Benzell et al. 2016) or policies to slow the growth of automation with economic incentives or attempts to influence development engineers (Acemoglu 2021; Brynjolfsson 2021).

Yet some observers have noted that automation may also complement labor in important ways (Autor 2015; Bessen 2015). By definition, automation replaces humans with machines on certain tasks. But automation could, at the same time, complement workers on *other* tasks. This paper presents a model explaining why and how such synergy might occur and empirical evidence that it does occur. The result is a much richer picture of automation that is both cost-reducing and quality-improving, that replaces workers but also increases the demand for diverse skills in a broad range of occupations.

This depiction is important because it helps explain important features about income inequality. Indeed, labor displacement models do not address a key feature of the rise in inequality since 1980, namely that it has largely occurred *between* firms rather than within firms.² Labor displacement affects inequality because it decreases aggregate employment demand for some skill groups relative to others, thus leading to growing wage differences in equilibrium. But in these models, the market wages of different skill groups affect all firms, changing within-firm inequality.³ However, the models can be extended: if automation increases the returns to quality on non-automated tasks, then automating firms might pay more than others and hire higher quality workers, thus increasing sorting.

¹ (Autor, Levy, and Murnane 2003; Acemoglu and Autor 2011; Brynjolfsson and McAfee 2014; Acemoglu and Restrepo 2018a; 2018c; Benzell et al. 2016; Korinek and Stiglitz 2018; Hémous and Olsen 2022).

² (Card, Heining, and Kline 2013; Barth, Davis, and Freeman 2018; Song et al. 2019; Lachowska et al. 2020).

³ Because firms differ in the extent to which they employ different skill groups, these differences might secondarily affect between-firm pay, but only in dilute form. Also, firms might reorganize production in response. For instance, Song et al. (2019) argue that coincident outsourcing of low paying jobs might mask the extent of within-firm changes. We test for this explanation below.

In popular discourse, automation is about reducing costs by cutting labor, not about improving quality. But that is not necessarily the case. Automation includes robots that replace manual labor, but it also includes productivity tools such as spreadsheets (automating the calculation) used by white collar workers or AI tools that automate predictions augmenting humans.⁴ Indeed, researchers have found that advanced technologies are often directed more to improving product quality or creating new products with better quality than they are to saving cost (Brynjolfsson and Hitt 2000; Bresnahan, Brynjolfsson, and Hitt 2002; Bessen et al. 2018; Babina et al. 2020; Hirvonen, Stenhammer, and Tuhkuri 2021). Automated machines can spin finer yarn than humans, they can allow machinists and surgeons to operate at higher precision, and AI systems can make more accurate predictions.

Why is quality important for technology? Quality on complementary tasks is critical to many production processes. In Kremer's (1993) famous example, the failure of one part doomed the space shuttle *Challenger*. Poorly performed tasks can create defects, reducing the value of output, or they can halt production, slowing the rate of output, or they can reduce the reliability of the product. Yet the quality of task performance often depends critically on the quality of labor, on the ability of labor to perform specific tasks. Elon Musk's highly automated Tesla factory fell far short of production quotas because, in his words, "humans are underrated." Clark (1987), comparing workers at highly automated textile mills around the world, found six-fold differences in output per worker, even comparing workers at similar mills using identical equipment and with similar British managers. The differences lay in the varied ability and willingness of these workers to perform non-automated tasks reliably and quickly (see the example below).

This means that there are important dimensions to skill beyond educational or occupational skill groups and, for this reason, analyzing inequality with skill groups alone is inadequate.⁵ In our model, the quality of task output depends on the specific skill or effort of

⁴ And although robots have featured in recent economic papers, US investment in robots was only \$7 billion in 2019, while investment in software, studied here, was over \$400 billion (US Census).

⁵ Lindenlaub (2017) argues that multi-dimensional skills are needed to understand the link between sorting and technology. It is well-recognized that demographic skill groups are at best crude indices of the actual multiple dimensions of skill (Acemoglu 2002, Section 7). Skill groups have other limitations for the analysis of income inequality. For one thing, endogenous selection into skill groups, such as changing access to college education, means that the actual skills of demographic groups change over time. Technology also changes the skills of occupational groups over time (Autor, Levy, and Murnane 2003; Spitz-Oener 2006). Also, there is great and changing variance of wages within skill groups (Hunt and Nunn 2019).

the worker performing the task. But workers differ in their disutility of expending effort on task performance or on learning new skills. Workers with high (low) disutility comprise a low (high) skill group—these are the ones likely to get less (more) education, for instance. Hence, skill groups defined by education or occupation matter, but so, too, do task-specific skill levels. Employers can improve product quality along two margins: by hiring workers from a high skill group and by providing them stronger performance incentives. In this paper, we measure effects on both skill groups and on task-specific skills.

To model automation, we extend the Acemoglu-Restrepo model of automation (2018a; 2018c) to include variable task quality as modeled by Kremer and Maskin (Kremer 1993; Kremer and Maskin 1996). Our model provides a natural explanation for the rise in sorting of skilled workers to high-paying firms. The model generates a Remainder Effect (Bessen 2015): automating some tasks raises the demand for skill and effort on complementary non-automated tasks. The relative importance of between-firm and within-firm wage gaps corresponds to the two margins along which firms manage skill. Between-firm pay differences are driven by differences in labor quality while within-firm pay differences are driven by the relative displacement of different skill groups. This implies that the impact of automation on wage inequality—and the policies needed to counter this inequality—depend on the relative importance of labor displacement and labor quality enhancement. The model also explains why successful implementation of information technologies is linked to management practices that provide stronger incentives for performance (Bloom, Sadun, and Van Reenen 2012).

To test the model predictions, we use rich data on the skills that firms demand and the pay that they offer in online help wanted ads. The skills requested in job ads allow us to identify multiple dimensions of specific skills that employers apparently value as important for achieving quality output, for instance, affecting firm market value (D. Deming and Kahn 2018; Bana 2021). These include education and experience required, measures of cognitive and social skills, other soft skills, information technology skills, and skills related to other technologies and market knowledge.

We test how the demands for these skills are related to firm investments in own-developed software. Firms have been investing heavily in developing information technology systems for their own use, including artificial intelligence applications (AI). Self-developed and custom software grew to \$241 billion in 2020, excluding software developed for use in a

product and much of this investment is in systems that automate business processes such as enterprise resource planning. Following some literature, we measure the adoption of these technologies as the share of software developers in a firm's total hiring (Tambe and Hitt 2012; Tambe et al. 2019; Bessen 2020; Harrigan, Reshef, and Toubal 2021). The entire investment in these technology platforms includes complementary investments in hardware, packaged software, and organizational capabilities.

We first test whether skills demanded and pay increase when firms make major investments in their internal information technology. We identify these episodes as substantial increases in relative hiring of software developers, so-called investment “spikes,” and we analyze them using a difference-in-differences methodology.⁶ This method helps isolate the effects of technology from many possible confounders and we also add controls for outsourcing, labor market conditions, management changes, and productivity and demand shocks. Following a software spike, firms increase their demand for skills across all categories, both for jobs that require a college education as well as jobs that do not. Firms also significantly increase the pay they offer, after controlling for job characteristics, to most skill groups, thus increasing between-firm pay differences among new hires.

Because spiking firms are a select group, this analysis might not reflect the role of information technology in sorting more generally. We also look at the relationship between investment in own-developed software and the demand for skills in the universe of online help wanted ads. We first calculate firm pay fixed effects by regressing salaries offered against job and firm characteristics. These fixed effects are correlated with our various skill measures, indicating sorting. But firm investment in own-developed software also correlates with both the firm fixed effects and skill measures. When software is added to the regression of skill against firm fixed effects, the correlations are substantially reduced. Software accounts for most of the correlations between firm fixed effects and skills. The prominent relationship between information technology and sorting and the relatively recent shift of investment to information technology suggests that much of the rise of sorting can be accounted for by this technology.

⁶ A variety of papers have begun using technology spikes and difference-in-differences or event studies to analyze technology impacts (Bessen and Righi 2019; Bessen et al. 2022; Humlum 2019; Domini et al. 2021; Aghion et al. 2020; Hirvonen, Stenhammer, and Tuhkuri 2021; Rodrigo 2021).

To summarize, this paper makes several contributions. First, we develop a model of automation that includes both cost reduction and quality enhancement, thus generating a richer set of outcomes. The model accounts for automation effects on both within-firm and between-firm pay inequality, reflecting the relative extent to which automation substitutes or complements workers. This difference provides a tool to estimate the relative importance of substitution/complementation and implies different policy choices to counter inequality.

Second, we study the micro-level impact of proprietary information technology to test key aspects of the model. We find that when firms invest in these systems, their demand increases for a diverse set of technical, cognitive, and social skills and these demand increases occur across skill groups—across jobs requiring a college degree and those that do not, across routine jobs as well as nonroutine. Moreover, this increased demand is reflected in higher pay offered, after controlling for job characteristics, contributing to growing between-firm pay differences. While we also find evidence of labor displacement in manual jobs, our overall findings differ from predictions of models of pure labor displacement and from models of skill-biased technical change. In both the labor displacement and skill-biased change stories, technology only affect limited groups of workers; in our story, it affects most groups of workers.

Third, we explore how much of the overall sorting of skilled workers to high-paying firms can be accounted for by proprietary information technology systems by looking at the correlations between firm fixed effects and skills. We find that the majority of these correlations is accounted for by this technology, suggesting that the increase in sorting may be closely related to the rise of proprietary information technology.

Of course, there are important non-technological factors that may contribute to sorting including rent-sharing, firm size (Eeckhout and Kircher 2018), search frictions (Burdett and Mortensen 1998), and monopsony (see for instance Card et al. 2018). Cortes et al. (2020) model sorting arising from skill-biased technical change. In empirical research, technology has been associated with between-firm wage differences (M. Doms, Dunne, and Troske 1997; Dunne et al. 2004; Barth et al. 2020) and the rise in information technology is coincident with the rising importance of sorting to inequality. But little research connects the actual adoption of technology with changes in skill demand. Some research has explored the effects of the adoption of computers or automation technology on firm wages in difference-

in-differences or event studies, generally finding a rise in firm pay following adoption.⁷ However, the increase in firm pay could arise from rent sharing rather than from greater demand for skills. Dillender and Forsythe (2019), in an approach similar to ours, use Burning Glass data to identify firm computer technology adoption; they find greater skill demand and higher pay for office and administrative support workers, suggesting that this technology is labor augmenting. Other papers find that computers or AI change skill demands (Autor, Levy, and Murnane 2003; Spitz-Oener 2006; Acemoglu et al. 2020).

Some studies use worker fixed effects from AKM regressions as a proxy for skill, but these might also reflect rents arising from search frictions (Abowd, Kramarz, and Margolis 1999; Bagger and Lentz 2019). Hakanson et al. (2020) find that worker sorting across firms by ability measured using standardized test scores is related to the rising information technology sector, but they lack firm-level measures of technology. Deming (2017; see also Aghion et al. 2019) finds an association between information technology and soft skills. This paper studies firm-level adoption of technology and both the subsequent firm demand for specific skills and firm pay offers.

A historical example

To fix ideas, it is helpful to look at an example of the Remainder Effect. There is sufficient historical data available for mills weaving coarse cotton cloth during the 19th century in the U.S. to construct an engineering production function that specifies the weavers' main tasks, the frequency of their occurrence, typical times to perform those tasks, and how those tasks were automated over time (this section draws from Bessen 2003; 2012; 2015). Although the process was highly automated, humans still had to perform some critical tasks. When bobbins ran out of thread, humans replenished them; when the edges of the cloth pulled inward as the loom wove, humans straightened them; when threads broke, humans had to stop the machines, unravel the defective cloth, fix the break, and restart the machines.

⁷ Gaggl and Wright (2017) find that computers raise wages in small firms, mostly in managerial, professional, and technical occupations. Bessen et al. (2022) find that automation raises wages in large firms, but wages decline in small firms. Acemoglu, Lelarge, and Restrepo (2020) find that robots raise wages in some regressions. Humlum (2019) and Rodrigo (2021) also find that robots increase pay. Graetz and Michaels (2018) find a similar increase at the industry level.

The productivity of a weaver depended critically on her skill at performing these specific tasks and the effort she applied to them. Many tasks were performed while the loom was stopped, so the weaver's speed affected the rate of output. The reliability of weaver's performance determined whether subsequent defects or failures would occur with greater or lesser frequency. The weaver's attentiveness monitoring the looms affected how quickly faults could be detected and fixed. And of course, the quality of the weaver's performance also affected the occurrence of defects that reduced the value of the cloth.

These skills had to be learned on the job. New hires went through a learning process that quadrupled their output per hour over the course of a year or so. Treating foregone output as a human capital investment (Becker 1993), the human capital of these supposedly "unskilled" weavers was substantial, roughly equivalent to the investments in adult male tradesmen who went through apprenticeships (the weavers were mainly young women). These skills contributed substantially to the rise in productivity over the 19th century. The labor time required per yard of cloth fell by 98%. However, analysis using the engineering production function shows that new inventions cannot account for all this decrease; about a quarter is due, instead, to better quality of labor.

As automation progressed over the century, many of these tasks were automated and no significant new tasks were added in this sector. As weavers performed fewer tasks per yard of cloth, they were assigned more looms to tend, so that their skills on these tasks remained important for productivity. In fact, their skills became even more important. Human capital investments increased (learning curves became steeper) and real pay for weavers rose substantially. Automation substantially increased the returns to skill on the remaining non-automated tasks.

Model

Basic Setup

Tasks and Automation

Our model is a combination of cost saving automation models by Acemoglu and Restrepo (2018a; 2018c) and models of production quality by Kremer and Maskin (Kremer

1993; 1996). We interpret Acemoglu and Restrepo’s model as providing a measure of potential output while Kremer’s model relates actual output to potential output, after accounting for quality-related failures.

We use a simplified version of the Cobb-Douglas instance of Acemoglu and Restrepo’s model (2018a) with constant returns to scale. Let there be N tasks. We keep the number of tasks fixed, ignoring the creation of new tasks, which we discuss further below. Because the production function has constant returns to scale, we allow an indefinite number of firms. Let the tasks be ordered so that the first I tasks are automated and the remaining $N - I$ tasks are performed by labor. The i th automated task uses k_i capital and the j th human task uses l_j labor. Letting the firm’s total capital $K = \sum_{i=1}^I k_i$ and total labor $L = \sum_{i=I+1}^N l_i$, equilibrium potential output can be written, under some assumptions (see Acemoglu and Restrepo 2018a, equation 3),

$$V = A(I)K^\alpha L^{1-\alpha}, \quad \alpha \equiv \frac{I}{N}, \quad \frac{dA}{dI} > 0. \quad (1)$$

where α is capital’s share of output and $A(I)$ is a measure of Hicks-neutral productivity, which we assume to be increasing in the number of automated tasks. We assume that I is exogenously determined by the state of technology. Firms, however, pay a fixed fee to adopt the latest technology so that in some circumstances, only more profitable firms choose to adopt (for a full model of adoption see Bessen et al. 2022 Appendix).

Quality

However, as Kremer (1993) observes, not all potential output is realized if tasks are performed imperfectly. In some production functions, failure of a critical task reduces output to zero (O-ring); in others, imperfect task output reduces the value of output; in yet others, task failures delay production (weaving), reducing the rate of output. The critical assumption here is that quality and quantity are not perfect substitutes. If quality and quantity *were* perfect substitutes, then output could simply be measured in quality-adjusted units and there would be no need to account for quality separately. However, as Kremer and others argue, there are many important instances where this substitution is imperfect, e.g., two mediocre surgeons are not equivalent to one surgeon whose patients have twice the survival rate (Rosen 1981; Kremer 1993).

It is standard in reliability engineering that the probability of failure increases with the number of tasks prone to failure. Multiple tasks provide multiple opportunities to fail. Let q_i , $0 \leq q_i \leq 1$ be the quality of performance of the i th task at a given scale of production. Perfect performance is designated by $q_i = 1$ and complete failure by $q_i = 0$. Then the actual output can be written⁸

$$Y = Q \cdot V, \quad Q \equiv \prod_{i=1}^N q_i. \quad (2)$$

To keep things simple, we assume that machines perform their tasks perfectly, $q_i = 1$, while humans are always at least a bit imperfect.⁹ For the tasks performed by labor, task quality will depend on worker skill or effort. The quality of task production can vary with general skills of the workers performing the task, but in many cases, it will surely depend on task-specific and technology-specific skills. Without loss of significant generality, we assume that workers are assigned to a single task and all workers assigned to a task have the same quality. This way worker skills are task specific. The quality of each task performed by labor is then $q_i = f(e_i)$ where e is effort per worker, either effort expended on the task or effort expended on learning new skills (see below). Then we assign

$$q_i = \begin{cases} 1, & i \leq I \\ f(e_i), & I < i \leq N \end{cases}. \quad (3)$$

We assume that $f(e_i)$ is a monotonically increasing, twice-differentiable continuous function, $f' > 0$, $f'' < 0$, and $\lim_{e \rightarrow \infty} f(e) < 1$ (humans are imperfect).

Labor Quality

Workers deliver a fixed amount of labor—there is no tradeoff with leisure time—but the quality of that labor varies. In a single-period model it is simplest to represent labor quality with a single variable, e , which we might think of as effort, either expended on the

⁸ Where q represents a probability of successful completion, then Y is expected output and we assume that firms are risk neutral.

⁹ A more general model could consider cases where machines have low quality but high efficiency and cases where inefficient machines are adopted because they have higher quality.

task or expended on learning new skills.¹⁰ Given a set of equilibrium prices, each worker's utility can be written as a function of their wage and the effort they exert, $U = U(w, e)$. Let

$$U(w, e) = w - \theta_j g(e), \quad g' > 0, g'' < 0, \theta_j \geq 1$$

where θ_j represents the j th worker's disutility of effort. Workers share the same preferences except for this disutility parameter.

Importantly for our analysis, there are two dimensions to labor quality: the worker's type or skill group, θ_j , and the actual effort/skill exerted at the worker's job, e . We assume that employers (but not third parties) can costlessly observe both θ_j and effort e . Employers can then elicit greater or lesser e by using greater or smaller performance pay incentives. They can also influence the labor quality of their workforces by hiring workers of different types, θ_j , who will be more or less responsive to those incentives.

Employers can elicit a desired level of effort by offering a two-part wage, with fixed amount w_u and performance incentive w_s such that $W = w_u + w_s \cdot e$. Utility maximization for the j th worker occurs when $\frac{w_s}{\theta_j} = g'(e)$, yielding a unique level of effort, $\hat{e}(w_s/\theta_j)$, and the corresponding task quality, $f(\hat{e}(w_s/\theta_j))$.

It is convenient to invert this function, yielding

$$w(q) = w_u + \theta_j \cdot h(q), \quad h', h'' > 0, \quad \lim_{q \rightarrow 1} h'(q) = \infty.$$

We assume competitive labor markets; firms pay the same wages for the same level of effort/skill. Wages differ across firms to the degree that worker effort/skill differs across firms.

Uniform Workers and Firms

We begin the exposition by presenting our model with uniform workers and firms to establish some basic results. Let there be only one type of labor, $\theta_i \equiv \theta$ for all i with otherwise identical firms. We introduce heterogeneity in the next section.

¹⁰ In a multiperiod model, workers may invest time and effort in learning new skills in one period that are used in subsequent periods.

Equilibrium

There is a fixed amount of inelastically supplied labor and capital in the aggregate economy distributed across firms. With uniform labor and firms, firms receive proportional allocations of labor and capital, L and K in equilibrium. Taking output price as numeraire, profits per firm are

$$\begin{aligned} \pi(q_{I+1}, \dots, q_N, k_1, \dots, k_I, l_{I+1}, \dots, l_N; I) \\ = A(I)K^\alpha L^{1-\alpha} \prod_{i=I+1}^N q_i - \sum_{i=1}^I r k_i - \sum_{i=I+1}^N w(q_i) l_i, \end{aligned}$$

where r is the user cost of capital and w is the wage. By the symmetry of the problem, it is straightforward to show that $q_i = q_j$, $k_i = k_j$, and $l_i = l_j$ in the appropriate range in equilibrium. The first order profit maximizing conditions for the three control variables then are

$$\frac{Y}{q_i} - \theta h' l_i = \frac{Y}{N l_i} - w = \frac{Y}{N k_i} - r = 0. \quad (4)$$

A useful result can be obtained by taking the implicit derivative from the first order maximizing condition for q_i (keeping the quality of other tasks fixed),

$$\frac{dq_i}{dA} = \frac{Nw}{\theta A q_i h''(q_i)} > 0. \quad (5)$$

Thus, increases in productivity will increase the equilibrium quality of output. When potential output increases, firms increase incentive pay, workers exert greater effort/skill, and total output increases more than potential output. In other words, an increase in potential output increases the returns to skill/effort.

Remainder Effect

Now consider what happens when the frontier of automated tasks increases from $I - 1$ to I for all firms. Let us assume that the adoption costs of the new technology are negligible so that all firms adopt. Productivity, A , increases and, by implication of the lemma above, this increase should boost labor quality. Aggregate quality also increases because the machine produces with greater quality on task I , that is, $1 > f(e_I)$. Combined, the effect of automation on total output per worker is

$$\Delta \ln \frac{Y}{L} = \Delta \ln A + \Delta \ln Q + \alpha \Delta \ln \frac{K}{L}$$

In this setting, capital and labor will be allocated proportionately across production units in equilibrium, so the last term drops out. Then,

$$\Delta \ln \frac{Y}{L} \approx \Delta \ln A + (N - I) \Delta \ln A \cdot \frac{dq}{dA} \cdot \frac{A}{q} - \ln f(e_I). \quad (6)$$

The second term represents the remainder effect. Automation boosts the returns to quality, increasing equilibrium labor quality. Output increases not only because automation reduces the labor cost of production but also because it increases labor quality. The third term is positive (since $f < 1$, $-\ln f > 0$) and captures the effect of improved quality in the newly automated task.

There is a corresponding change in the wage. Using the first order conditions and $L = (N - I)l_i$, the equilibrium wage is

$$w = \frac{N - I}{N} \cdot \frac{Y}{L}.$$

Following Acemoglu and Restrepo and using (6),

$$\begin{aligned} \Delta \ln w &\approx \frac{d \ln(N - I)}{d I} + \Delta \ln \frac{Y}{L} \\ &\approx -\frac{1}{N - I} + \Delta \ln A + (N - I) \Delta \ln A \cdot \frac{dq}{dA} \cdot \frac{A}{q} - \ln f(e_I) \end{aligned} \quad (7)$$

Acemoglu and Restrepo call the first term the “displacement effect” The second term is an efficiency effect (Acemoglu and Restrepo call it the “productivity effect”). The third term represents the remainder effect and the fourth captures the quality improvement effect. The remainder effect multiplies the base productivity effect, making a positive contribution to wages. Also, the fourth term implies further possible wage increases. In a more general model, this term could possibly be negative—that is, firms might accept inferior quality machines if they deliver a large enough efficiency gain. The sign and magnitude of this term is an empirical matter. However, the addition of the term highlights an important aspect of automation: firms may choose to automate not so much to reduce costs as to provide better quality output. To the extent this is true, the effect on wages will tend to be positive.

Generally, (7) provides reasons beyond Acemoglu and Restrepo why wages might increase.

To keep things simple, we have used single continuous variables for product and labor quality and have kept the number of products and tasks fixed. In a more general

setting, both new tasks and new products might be natural outcomes of a growing demand for greater quality. For example, as the quality of a task becomes more and more valuable with ongoing automation, firms might subdivide that task into two or more new tasks allowing workers to develop more specialized skills. Something like that appears to have happened during the 19th century (Atack, Margo, and Rhode 2019). Similarly, new products might be a form of realizing greater product quality.

Heterogeneity

Now let there be two types of workers: high skill, designated “H,” and low skill, designated “L,” where $\theta_H < \theta_L$. The aggregate supply of each type is fixed.

In general, there are two ways that workers can be assigned to firms: assortative matching, where some firms hire more high skill workers and other firms hire more low skill workers, and cross-matching, where firms hire a mix of high and low skill workers. A theoretical literature identifies a condition under which assortative matching occurs in competitive markets (Becker 1981; Sattinger 1975; 1993; Kremer 1993; Kremer and Maskin 1996), namely a positive cross derivative of output with respect to the qualities of different tasks. Our production function meets this criterion (see also Kremer 1993). In the next section, we consider a stylized model of sorting where firms hire all high skill workers or all low skill workers.

Kremer and Maskin (1996) show that with a slightly different production function, firms will, instead, cross-match under some conditions, hiring both high and low skill workers. This occurs when productivity is more sensitive to some tasks than others. Let us divide tasks into two groups: tasks in the range $I < i \leq J$ are “routine tasks” while tasks in the range $J < i \leq N$ are “nonroutine tasks.” Below we consider an alternative specification that meets the Kremer-Maskin conditions. While real world skill assignments may involve a mix of matching and sorting, these models illustrate in simple form the different effects that automation has on inequality between firms and within firms.¹¹

¹¹ Automation might also affect firms’ choices regarding sorting and cross-matching. Kremer and Maskin (1996) provide a variety of evidence that skill sorting has been increasing and workplaces are becoming more segregated by skill, that is, workers are more likely to work with other workers of similar skill (see also E. Handwerker 2015; E. W. Handwerker, Spletzer, and others 2016). Our model could be extended to address this possibility.

Sorting

In a market with complete sorting, some firms, designated by an “H” subscript, hire only high skill workers while other firms hire only low skill workers, designated with an “L” subscript. We assume that both types have the same level of automation initially. The first order profit maximizing conditions (4) then hold separately for each firm type. Combining the first order conditions for quality and labor, for worker/firm type j ,

$$w_j = \frac{Y_j}{Nl_j} = \frac{\theta_j \cdot h'(q_j) \cdot q_j}{N}, \quad j = L, H.$$

In the Appendix we show that in equilibrium, both q_j and the term $\theta_j \cdot h'(q_j) \cdot q_j$ are decreasing in θ_j , all else equal. This means that $w_H > w_L$ and the ratio of between-firm wages is

$$\omega \equiv \frac{w_H}{w_L} = \frac{\theta_H \cdot h'(q_H) \cdot q_H}{\theta_L \cdot h'(q_L) \cdot q_L} > 1.$$

The between-firm wage gap corresponds directly to differences in skill/effort between the firm types. Furthermore, it is straightforward to show that capital intensity and productivity are higher in type H firms:

$$\frac{w_H}{w_L} = \frac{K_H/L_H}{K_L/L_L} = \frac{Y_H/L_H}{Y_L/L_L} > 1.$$

To introduce automation into this setting, note that because type H firms have higher productivity, they also have stronger incentives to adopt new automation technology. The increase in output per worker from automation is $\frac{Y}{L} \Delta \ln A$ and so will be larger for type H firms. This increase will also be greater for the remainder effect term in (6). Suppose that there is a fixed cost per worker needed to adopt an automation technology. Then, in some cases, type H firms will find it profitable to automate while type L firms will not.¹² Given this difference, let us assume that type H firms automate, and type L firms do not. Disparate adoption of automation technologies is, in fact, widely observed and appears in our data as well.

¹² Firms may make temporary profits from automating, yet competition will eventually dissipate these rents. There are other reasons some firms may adopt while other do not: different capabilities of managers and workers or different access to proprietary technologies.

With this assumption, we can calculate ω using an approach like the one used in equation (7). Here, however, we must account for changes in the capital to labor ratios for the two groups. As Y/L increases for H firms, capital also shifts to those firms. In the Appendix we account for this change in the equilibrium solution to derive an approximate lower bound for the change in the between-firm wage ratio:

$$\Delta \ln \omega = \Delta \ln w_H - \Delta \ln w_L \approx > -\frac{1}{N-I} + \frac{N}{N-I-1} \left[\Delta \ln A_H + \Delta \ln Q_H + \frac{1}{I} \right].$$

The first term represents the displacement effect. The expression in brackets captures the productivity and quality effects. Here the displacement effect *decreases* between-firm wage differences while the productivity and remainder effects increase between-firm wage differences. If the productivity and remainder effects are larger than the displacement effect, the between-firm wage gap increases. If, on the other hand, low wage firms tend to automate, contrary to most evidence, then the changes would narrow between-firm differences. And if both types of firms automated, the results are ambiguous. Thus, growing differences in labor quality explain rising between-firm pay gaps if adoption of automation technology is uneven and if the displacement effect is smaller than productivity and remainder effects.

Cross-matching

Kremer and Maskin (1996) show that cross-matching occurs when some tasks are more sensitive to quality than others. The idea is that, under some parameter values, firms will choose to assign high skill workers to sensitive tasks and low skill workers to tasks that are less sensitive.¹³ We can accommodate these notions into our production function by specifying now that

$$q_i = \begin{cases} 1, & i \leq I \\ 1, & I < i \leq J \\ f(e_i), & J < i \leq N \end{cases}$$

where $I < J < N$. Routine tasks in the range $I < i \leq J$ are not sensitive to the quality of labor on those tasks while nonroutine tasks in the range $J < i \leq N$ depend on the skill and effort of workers. With this modification to the production function, firms will prefer to hire

¹³ Acemoglu and Restrepo (2018a; 2018b) exogenously assign high skill workers to nonroutine tasks and low skill workers to routine tasks.

high skill workers for nonroutine tasks and low skill workers for routine tasks. In equilibrium, the fixed stocks of low and high skill workers will be allocated proportionally to firms so that the ratio L_L/L_H of low skill workers to high skill workers will be the same. It is straightforward to show that firms will prefer to assign only high skill workers to nonroutine tasks and only low skill workers to routine tasks. Then first order profit maximizing conditions give us

$$w_H = \frac{Y}{L_H} \frac{(N - J)}{N} = \theta_H h'(q_H) q_H N, \quad w_L = \frac{Y}{L_L} \frac{(J - I)}{N}$$

and the within-firm wage difference ratio is

$$\phi \equiv \frac{w_H}{w_L} = \frac{N - J}{J - I} \cdot \frac{L_L}{L_H}.$$

Note that the within firm wage difference is independent of the quality of the high skill workers. The relative wage within firms depends on the relative supply of workers from different skill groups and the relative demand for routine and nonroutine tasks. While high skill workers will receive higher performance pay, their total pay package is not necessarily greater.

Now consider automation in this setting. In some papers, Acemoglu and Restrepo (2018a; 2018c) study situations where only routine tasks are automated. Then automation can be considered a change in the limit of automation from I to $I + 1$ as above. Then

$$\Delta \ln \phi \approx \frac{d}{dI} \left(\frac{N - J}{J - I} \cdot \frac{L_L}{L_H} \right) = \frac{1}{J - I} > 0.$$

Automation increases within-firm wage differences in this setting. However, automation is not necessarily restricted to routine tasks and then this type of labor displacement might decrease within-firm wage gaps (see Acemoglu and Restrepo 2018b).

But the general point remains that automation influences within-firm wage gaps by way of the displacement effect. In our model as well as the models in the literature, labor displacement directly affects the relative demand for different skill groups within firms and aggregate changes in demand for these groups determines the relative equilibrium wages. Because workers from one skill group are employed at this wage across firms, the effect will be observed as within-firm wage differences. On the other hand, the remainder effect concerns firm-, task-, and technology-specific skills that are not common across different

firms. These affect between-firm wage differences but not differences between skill groups within the firm.

This model provides three testable hypotheses:

1. Automation should increase the demand for task- and technology-specific skills across multiple skill groups;
2. This greater demand should be evident in the firm's greater willingness to pay more for these groups; and,
3. Assuming that automation differentially affects the tasks assigned to different skill groups, it should change the relative employment demand for different skill groups.

The first two hypotheses distinguish this model from pure models of labor displacement: here, automation complements labor. The firm's greater willingness to pay provides an explanation for greater between-firm pay gaps. Our model also differs from the skill-biased technical change hypothesis because the complementary effect of technology is not limited to specific skill groups.

Empirical Analysis

Data

These three hypotheses concern different aspects of firm labor demand: the specific skills demanded, the firm's willingness to pay for different skill groups, and the relative quantities of labor demanded for different skill groups. We measure these aspects of demand using help-wanted advertisements collected by Burning Glass Technologies. Burning Glass scrapes, deduplicates, and cleans the near universe of online job advertisements. A previous analysis of the dataset showed that this it accounts for 60-70% of all job openings and 80-90% of openings requiring a bachelor's degree or more (Carnevale, Jayasundera, and Repnikov 2014). The data include the advertised salary, firm name, industry, occupation, required education and experience, requested skills, and geographic location of the job. Our

sample spans from January 2014 to June 2019.¹⁴ We aggregate the ads by firm and calendar quarter and use this as our unit of observation.

Changes in labor demand should be immediately reflected in help-wanted advertising even though these changes might take longer to appear among the group of employed workers. To the extent that firms demand greater quality on task-specific skills, we should see increases in the specific skills requested in job ads. To the extent that greater demand increases the firm's willingness to pay, we should higher pay offered for jobs with comparable characteristics. And to the extent that demand changes across skill groups, we should see shifts in the share of job ads directed to different skill groups. We measure these outcomes with the following variables:

Specific skills. Burning Glass collects 16,050 different skills requested in ads as well as experience and education required. We group the specific requests into five mutually exclusive categories: social and cognitive skills as identified by Deming and Khan (2018), other soft skills, information technology and artificial intelligence, and other skills, mainly skills related to other technologies and industry knowledge (see Appendix). We use the mean number of requests per ad for each category and the mean experience and education requested as outcome measures.

Pay offered. Some help wanted ads list a salary offered or a range of salaries. If a range is offered, we take the middle of the range for our salary calculations. The outcome variable is the log Mincer residual in a regression equation including experience, experience squared, education, detailed occupation, state, year, and a measure of labor market tightness. We follow Moscarini and Postel-Vinay (2016) in defining labor market tightness as the ratio between Job Openings and Labor Turnover Survey (JOLTS) statewide openings for the non-farm sector and the state unemployment rate.¹⁵

¹⁴ While Burning Glass provides data prior to 2014, those years used different methods to collect, de-duplicate, and process the data. Because those differences might affect our analysis, we do not use that data. We omit job advertisements that are missing a firm name or salary, are in the public or university sector, are part time, or are internships. To identify ads belonging to the same firm, we cleaned names, removing standard business identifiers (“Inc.,” “Ltd”, “Co.,” etc.) and looking for typos in the most frequently used names in the dataset.

¹⁵ Because most jobs do not list salaries, sample selection bias might affect this measure. Bessen et al. (2020) find that an exogenous change to salary listing does not significantly affect listed salaries, mitigating this concern.

Relative employment. To measure changes in the relative hiring of skill groups, we use the share of job ads for each group. We divide occupations into two sets of skill groups defined by characteristics identified in O*NET, version 17.0. First, we identify whether a bachelor’s degree or higher is required for most jobs in that occupation. Second, we identify occupations as routine cognitive, routine manual, nonroutine cognitive, and nonroutine manual using the indexes for these characteristics developed by Acemoglu and Autor (2011); an occupation is assigned to the job characteristic skill group if its index ranks in the top third.¹⁶

Finally, note that we exclude information technology jobs (SOC 15) from our skill and pay measures to avoid confounding effects.

Implementation

We seek to test the model predictions regarding the adoption of large proprietary information systems. Much of the literature on technology and inequality measures technology as predicted “exposure” to automation, or industry-level investment levels, or proxies such as the share of workers in routine-intensive jobs. To capture impacts on between-firm differences, we thought it important to use firm-level measures of actual technology adoption. These eliminate many potentially confounding correlates.

We measure investment in this technology from the job ad data as the share of jobs going to software developer occupations.¹⁷ This captures investment in firms’ own-developed software and it is correlated with contracted software and other IT measures (Tambe and Hitt 2012; Bessen 2020 fn. 12).

To analyze adoption, we identify “spikes” in developer hiring as events where the share of software developers rose by one percent or more relative to the mean share over the previous four quarters.¹⁸ This approach leverages the finding from the capital investment

¹⁶ These groups are not mutually exclusive.

¹⁷ Occupations in SOC 15 excluding 15-1141, 15-1142, 15-1151, and 15-1152, database, network, and computer administrators and support specialists.

¹⁸ Also, to reduce noise, we eliminate spikes when the firm has fewer than 50 ads in quarter. A variety of robustness checks in the Appendix vary the threshold, finding little effect on results. 19% of firm-quarters are spikes, weighted by the number of job ads. While only about 1% of firms spike, these firms account for 77% of the hiring of software developers.

literature that when uncertain investments are indivisible and irreversible, they will occur in discrete episodes of lumpy investment (Haltiwanger, Cooper, and Power 1999; M. E. Doms and Dunne 1998). We find that investments in own-developed software are also lumpy and persistent (see Appendix Figures A1 and A2), so we use these discrete events in difference-in-differences (DID) regressions and event studies. It is possible that we fail to identify some lumpy investments and incorrectly identify others. For example, firms rely on outside contractors to implement new systems rather than hiring their own developers. To the extent misidentification occurs, our results will be understated.

Do these spike events represent automation? We note generally that most information technology applications involve some degree of automation—they manage information that was formerly managed by humans. This is strictly true for applications that automate business processes such as enterprise resource planning, customer relationship management, and electronic data interchange. In fact, the use of these systems is correlated with bookkeeping measures of automation expenditures (Bessen et al. 2022 Section 2.3). We flag events that specifically include hiring of workers with skills related to these automation applications and find that 81% of our spike event do.¹⁹ Similarly, 31% of the spikes involve firms requesting artificial intelligence skills. Thus our spikes predominately involve applications that automate tasks.

To avoid problems of heterogeneity in our two-way fixed effects regressions, we construct balanced panels around each possible spike quarter and run stacked regressions (Cengiz et al. 2019, Appendix D). Let T_i be the first quarter in which firm i spikes. For each possible spike quarter, p , designating a different cohort, we construct a balanced panel P consisting of observations from $t = p - 5$ to $t = p + 5$ of the treatment group, $T_i = p$, and the control group, $T_i > p + 5$. Because firms that spike are different from firms that do not (see Table A1), we restrict the control group to firms that spike at some point in our data. This means that the treatment and control groups differ only in the timing of their adoption events.²⁰ This gives us a degree of identification by removing fixed or slowly

¹⁹ These are jobs requesting skills with keywords ERP, CRM, EDI, MRP, SAP, Automat*, and Robot*.

²⁰ Bessen et al. (2022, Appendix) provide a model for differential timing. We also duplicate our results for the full sample (Table A4).

changing confounders, such as industry and firm size, and by distinguishing major new investments from maintenance hiring. Our DID specification for outcome variable Y is

$$Y_{ipt} = \delta \cdot \mathbf{1}(t \geq p) + \mu_{ip} + \tau_t + \beta X_{it} + \epsilon_{ipt}. \quad (8)$$

where δ is the average treatment effect, μ_{ip} is the panel x firm fixed effect, τ_t is the time fixed effect, and X_{it} is a vector of control variables.

However, the model is still not fully identified because the timing of adoption is endogenous. While we test for and do not find significant pre-trends in our outcome variables, it is still possible that some other factor is correlated with adoption, occurring simultaneously, and which independently affects outcome variables. We identify and control for four such possible simultaneous confounders:

1. **Labor market tightness.** Tight labor markets might induce firm to automate and might also raise wages and skills demanded (Modestino, Shoag, and Ballance 2019 find tight labor markets *lower* skill requirements). We use the tightness measure described above to control for this confounder.
2. **Outsourcing of low wage jobs.** Perhaps automation facilitates the outsourcing of low wage jobs, mechanically raising the average pay and skill requirements of remaining jobs. We control for the share of “outsourcable” jobs that should track these shifts.²¹
3. **Productivity and demand shocks.** Perhaps firms adopt new technology in response to productivity or demand shocks and these shocks are also passed through to wages. We control for shocks using additional variables obtained from Compustat for the subsample of firms matched to Compustat.²² One variable is the growth in real sales from the quarter before the spike to a year

²¹ The outsourceable occupations are Protective Services (SOC 33), Food and Serving (SOC 35), Building, Grounds, Maintenance (SOC 37), and Transportation and Moving (SOC 53) outside of outsourcing industries, NAICS 484, Truck Transportation, NAICS 561, Administrative and Support Services, NAICS 722, Food Services and Drinking Places, and NAICS 811, Repair and Maintenance.

²² Bledi Taska of Burning Glass provided a preliminary key to match to Compustat, which we supplemented with our own name cleaning algorithm. Further, we used a fuzzy match with distance scores, which was then manually reviewed for those with close distances. The match assigns approximately 63% of the firms in Compustat to a job posting, with 73% of the firm-years being matched to a job posting. The firms that are matched to a posting account for 83% of employment total in Compustat.

earlier. The second control is a third order polynomial in log variable costs and log net capital stock (both deflated).²³

- 4. Management.** Perhaps new managers prefer to adopt technology and also to hire more highly skilled workers. For the entire sample, we add the manager (SOC 11) share of hiring as a control. For the Compustat subsample, we add a binary variable to flag changes of CEO using data obtained from Execucomp.

We find that some of these control variables have weak correlations with the occurrence of spikes (see Table A2), but also, they do not substantively change our results. This gives us a limited form of identification; it is not equivalent to conducting a randomized controlled trial, but our results are identified conditional on the following assumption:

Identification assumption: there are no significant confounders that occur simultaneously with the adoption of these information technology systems other than labor market conditions, outsourcing, productivity and demand shocks, and management changes.

As such, our results are consistent with our model and inconsistent with pure displacement models and with the skill-biased technical change hypothesis. Finally, our spiking results pertain to a select sample of firms. Below we also explore the broader validity of our model to the universe of help-wanted ads.

Findings

Firm Spikes

Table 1 presents stacked difference-in-differences regressions (a balanced panel for each spiking year) where the dependent variables are the number of skills requested in the various categories.²⁴ All of the skill measures show significant increases following the adoption event except for education. The top panel includes all jobs except for IT jobs (SOC 15). We interpret the greater number of skills requested as evidence of greater demand for

²³ In the style of Olley and Pakes (1996) this polynomial is a nonparametric representation of productivity obtained by inverting the demand equation for variable inputs (cost of goods sold).

²⁴ Regressions are weighted by the number of ads and include time and cohort by firm fixed effects as well as controls for labor market tightness, and the shares of management and outsourceable jobs.

specific skills. When firms place greater value on “Teamwork” or on “Adhesives Industry Knowledge,” they will be more likely to specifically request these skills.

Panel B includes the skill measures only for jobs that do not require a college diploma.²⁵ These coefficients tend to be a bit smaller, but as in the larger sample, all are significant and positive except for education. Skill demands appear to rise for both college and non-college jobs, although a bit less for the latter.

Panel C looks at the *share* of skills rather than the number, that is the number of skill requested in each category divided by the total number of skills requested. Following a spike, firms appear to place relatively greater demand on social and soft skills, suggesting organizational changes consistent with Deming (2017). However, these shifts in the composition of skills are small compared to the increases in demand seen in Panel A.²⁶ The overall impact appears to be that firms request more of the kinds of specific skills that they requested before the spike, that is, they demand higher labor quality.

Table 2 examines a broader set of skill groups, namely jobs classified as routine/nonroutine and cognitive/manual as per Acemoglu and Autor (2011). Panel A shows that all groups show significant increases in the mean number of skills requested except for nonroutine manual jobs. These results suggest that the technology complements workers in a wide range of jobs. As we would expect, firms are also willing to pay more to these workers seemingly complemented by software investments—the greater demand for skills does not just reflect the preferences of HR professionals. The dependent variable in Panel B is the log residual wage after controlling for job characteristics. These pay levels rise significantly for all groups except nonroutine manual workers; they rise notably more (9.1%) for nonroutine cognitive jobs.

Table 3 tests the robustness of results to additional controls. Here the sample is limited to firms that are matched to Compustat. Using Compustat and Execucomp data, we add a control (in columns 3 and 6) for the rate of revenue growth, a flag for change of CEO, and a third order polynomial in log capital and log variable costs to capture productivity

²⁵ That is, fewer than half the jobs require a diploma as rated by O*NET.

²⁶ Expressed as percentages, the increases shown in Panel A range from 3% to 13%, much larger than the shifts, which are less than 1%.

nonparametrically. Some of these controls are statistically significant, but they do not meaningfully alter our estimates of the treatment effect.

Our results are also robust to other concerns. Figures 1 and 2 shows event study graphs corresponding to the first column in Table 2.²⁷ The graph shows a significant and persistent increase in the mean number of skills requested and log residual wage following an adoption event. Moreover, there is no evidence of pre-event trends in these outcome variables nor in the other outcome variables used in Table 1, lending support to the parallel trends assumption (see Appendix Table A9). Table A3 tests sensitivity to different spike thresholds and panel lengths; our results are robust to these changes. Table A4 shows regressions using an expanded sample that adds firms that never spike; the results are similar. Table A6 finds that excluding firms in industries that create software products (NAICS 50 and 54) makes little difference to our results. About one third of our spiking firms use artificial intelligence as evidenced by requests for AI skills during the spiking quarter; 81% involve automation technologies. Our main results do not change significantly for these groups of firms (Table A7). We also conduct a placebo test to support the idea that the effects we observe are related to software specifically and not to other technologies or general hiring of higher paid workers. In Table A8, we show results from spikes in the hiring of engineers and technicians constructed in the same way as our software spikes. These personnel may tend to work on technologies that are not so much about automation. Spikes in the hiring of engineering related personnel do not exhibit similar treatment effects, suggesting that it is something specifically about information technology—perhaps automation—that is driving our results.

The increased skill demands and greater pay suggest that proprietary information systems complement labor. Our model suggests that automation can also displace labor. Table 4 shows evidence of displacement. The top panel shows the share of job ads going to each skill group. Following technology investment, relative hiring increases for jobs requiring college degrees and for jobs with cognitive skills, both routine and nonroutine; relative hiring

²⁷ These show the δ_τ coefficients from the following modification of (8):

$$Y_{ipt} = \sum_{\tau=-4}^5 \delta_\tau \cdot \mathbf{1}(\tau = t) + \mu_{ip} + \tau_t + \beta X_{it} + \epsilon_{ipt}.$$

decreases for non-college jobs and manual jobs. Panel B displays the log level of hiring by skill groups. Job ads decrease for occupations that do not require a college degree and for routine manual jobs. Thus, consistent with our model, there is labor displacement that occurs alongside increased demand for skills as seen in the prior tables.

In theory, this displacement contributes to lower equilibrium wages for workers who only have routine manual skills or only a high school diploma. In practice, however, this is difficult to establish empirically because other factors might confound the effect of technology on the pay of different demographic groups (but see Acemoglu and Restrepo 2021). For instance, in 1980 62% of the U.S. workforce had only a high school degree or less; today that figure is 38%. It seems highly likely that expanded access to higher education may have selectively induced some workers—those with lower disutilities of learning—to seek more education. This means that high school educated workers do not comprise a consistent skill group over time and declining pay for this group might reflect declining ability rather than technological effects. Our model provides some insight into the relative importance of labor displacement on wages. Because labor displacement affects market wages, it affects all firms equally; that is, it increases within-firm inequality. The finding that relatively little of the increase in inequality arises within firms—26% according to Song et al. (2019)—suggests that labor displacement is not the dominant driver of rising inequality.

To summarize, given our identification assumption, the evidence on residual wages implies that firm investments in proprietary information technology contributes to between-firm pay differences; the evidence on skills requested implies that these firm pay increases are associated with increased firm skill demands. In other words, these technology investments contribute to sorting of skills to higher paying firms. However, the evidence presented pertains only to a select sample of firms.

Sorting

We can also explore the relationships between firm pay levels, skills requested, and information technology across the entire sample of help wanted ads by looking at sorting of skills to high-paying firms. Studies using linked employee-employer data find that sorting accounts for most of the increase in wage inequality since 1980 (Card, Heining, and Kline 2013; Barth, Davis, and Freeman 2018; Song et al. 2019; Lachowska et al. 2020). Utilizing the AKM method (Abowd, Kramarz, and Margolis 1999), these studies estimate firm pay fixed

effects controlling for observed and unobserved worker heterogeneity with worker fixed effects. The worker effects are positively correlated with the firm effects and this correlation accounts for much of the rise in inequality. Assuming that the worker effects represent worker skills (rather than arising from search frictions or other factors), this correlation represents sorting of skilled workers to high-paying firms.

We alternatively estimate firm pay fixed effects by regressing pay offered in job ads controlling for job characteristics. These pay offers are obviously independent of individual worker heterogeneity. Using log salary as the dependent variable (or the mean of the salary range limits if a range is listed), we calculate firm fixed effects in a regression with controls for detailed occupation, industry, state, year, labor market “tightness,” skills requested, education required, and experience required (see Table A5). The R-squared for this regression is .688. The regression excludes software development occupations to avoid spurious correlation with our key independent variable. This gives us estimates of firm fixed effects for 205,306 firms that posted 85,142,065 help wanted ads, excluding ads for information technology occupations. These firm fixed effects are different from fixed effects derived from the AKM method—our fixed effects reflect differences in pay in hiring, not in the pay of incumbent workers.²⁸ Nevertheless, both methods provide estimates of the firms’ varied willingness to pay for comparable workers. And we can measure sorting by looking at the correlation between these firm fixed effects and actual skill levels demanded in the job ads. These correlations are shown in the top panel of Table 5 which reports regressions of mean skill measures for each firm against firm wage fixed effects. The correlations are all significant, indicating sorting. The standardized coefficients represent the correlation coefficients. These are similar to the correlation of 0.28 between worker fixed effects and firm fixed effects reported by Song et al. (2019) for the period from 2007-13 using the AKM method.²⁹

²⁸ There is a close correspondence between average advertised salaries and average salaries actually paid as observed in the Current Population Survey. Weighting the job ads to match the CPS distribution across occupations, the median log salary range from Burning Glass is from 10.32 to 10.69. The median log CPS salary for new hires is 10.48.

²⁹ Calculated using their figures for $\frac{cov(WFE, FFE)}{\sqrt{var(WFE)var(FFE)}}$.

But it turns out that firm hiring of software developers is correlated with both firm fixed effects and with skill measures.³⁰ The bottom panel adds quadratic terms in the mean share of software developers in hiring. The correlations between worker fixed effects and skill measures drop sharply. The last row shows the magnitude of the decrease in the standardized coefficients as a portion of the correlation coefficient in Panel A. It appears that information technology investments can account for the majority of the sorting of skills to high paying firms in hiring. Given that firm investment in own-developed software has increased more than ten-fold since the 1980s (BEA data), this shift can explain much of the rise in inequality due to sorting.

Conclusion

This paper argues that automation can be both cost-reducing and quality-enhancing; it can replace labor on some tasks while it increases demand for skills on others. Major investments by firms in own-developed information technology are followed by greater demand for specific skills requested in job ads and by higher pay offers. Moreover, demand increases across skill groups, both for jobs requiring college and those that do not, for routine jobs as well as nonroutine jobs. These broad increases contribute to between-firm pay differences and the sorting of skilled workers to high paying firms. Analyzing the universe of help-wanted ads, we find that these information technology investments account for most of the sorting across firms.

This pattern differs from predictions of the skill-biased technical change hypothesis and from theories of labor displacement. Our model provides an explanation: labor quality matters. While automation displaces labor on some tasks, it can also increase the returns to skill on the remaining non-automated tasks. Models that view automation as strictly substituting for labor without also complementing some workers might be incomplete and overly pessimistic. For instance, Acemoglu and Restrepo argue that wages will fall for “so-so innovations” where the productivity gain is small. But if automation raises the demand for

³⁰ See our working paper for a more complete exploration of these relationships (Bessen, Denk, and Meng 2021)

quality on the remaining tasks (remainder effect), wages may rise even with modestly productive innovations.

The matter is ultimately empirical, but here, too, labor quality matters for the analysis. Inequality is frequently measured by differences between occupational or educational groups. But our evidence suggests that skills and inequality change along other dimensions as well. In our model, labor displacement gives rise to greater within-firm inequality, but the evidence suggest that this is a secondary contributor to growing inequality. On the other hand, automation that complements labor can increase between-firm inequality, which appears to be more important.

If so, this suggests a different direction for policy to combat income inequality. Researchers who assume that automation is purely labor displacing have proposed policies to redistribute income, to alter tax incentives to discourage too much automation, and to encourage engineers to not develop automation (Korinek and Stiglitz 2018; Benzell et al. 2016; Acemoglu 2021; Brynjolfsson 2021). But if automation mainly complements workers, giving rise to greater between-firm pay differences, then policy might instead need to focus on reducing differences between firms in the uneven adoption of technology. Indeed, concerns have been raised about slower diffusion of technology (Andrews, Criscuolo, and Gal 2016; Akcigit and Ates 2021). While policy evaluation is beyond the scope of this paper, our analysis highlights that policy should be based on a richer picture of automation, one where technology complements labor as well as substitutes for it, where the quality of labor matters.

References

- Abowd, John M, Francis Kramarz, and David N Margolis. 1999. "High Wage Workers and High Wage Firms." *Econometrica* 67 (2): 251–333.
- Acemoglu, Daron. 2002. "Technical Change, Inequality, and the Labor Market." *Journal of Economic Literature* 40 (1): 7–72.
- . 2021. "Harms of AI." National Bureau of Economic Research.
- Acemoglu, Daron, and David Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." *Handbook of Labor Economics* 4: 1043–1171.
- Acemoglu, Daron, David Autor, Jonathon Hazell, and Pascual Restrepo. 2020. "AI and Jobs: Evidence from Online Vacancies." National Bureau of Economic Research.
- Acemoglu, Daron, Claire Lelarge, and Pascual Restrepo. 2020. "Competing with Robots: Firm-Level Evidence from France." In *AEA Papers and Proceedings*, 110:383–88.
- Acemoglu, Daron, and Pascual Restrepo. 2018a. "Artificial Intelligence, Automation, and Work." In *The Economics of Artificial Intelligence: An Agenda*, 197–236. University of Chicago Press.
- . 2018b. "Low-Skill and High-Skill Automation." *Journal of Human Capital* 12 (2): 204–32.
- . 2018c. "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment." *American Economic Review* 108 (6): 1488–1542.
- . 2021. "Tasks, Automation, and the Rise in US Wage Inequality." National Bureau of Economic Research.
- Aghion, Philippe, Céline Antonin, Simon Bunel, and Xavier Jaravel. 2020. "What Are the Labor and Product Market Effects of Automation? New Evidence from France."
- Aghion, Philippe, Antonin Bergeaud, Richard W Blundell, and Rachel Griffith. 2019. "The Innovation Premium to Soft Skills in Low-Skilled Occupations." *Available at SSRN* 3489777. <https://ssrn.com/abstract=3489777>.
- Akcigit, Ufuk, and Sina T Ates. 2021. "Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory." *American Economic Journal: Macroeconomics* 13 (1): 257–98.
- Alekseeva, Liudmila, José Azar, Mireia Gine, Sampsa Samila, and Bledi Taska. 2020. "The Demand for AI Skills in the Labor Market."
- Andrews, Dan, Chiara Criscuolo, and Peter N Gal. 2016. "The Best versus the Rest: The Global Productivity Slowdown, Divergence across Firms and the Role of Public Policy."
- Atack, Jeremy, Robert A Margo, and Paul W Rhode. 2019. "'Automation' of Manufacturing in the Late Nineteenth Century: The Hand and Machine Labor Study." *Journal of Economic Perspectives* 33 (2): 51–70.
- Autor, David 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, David H, Frank Levy, and Richard J Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *The Quarterly Journal of Economics* 118 (4): 1279–1333.
- Babina, Tania, Anastassia Fedyk, Alex Xi He, and James Hodson. 2020. "Artificial Intelligence, Firm Growth, and Industry Concentration." *Working Paper*.

- Bagger, Jesper, and Rasmus Lentz. 2019. "An Empirical Model of Wage Dispersion with Sorting." *The Review of Economic Studies* 86 (1): 153–90.
- Bana, Sarah H. 2021. "Using Language Models to Understand Wage Premia," working paper, .
- Barth, Erling, James C Davis, Richard B Freeman, and Kristina McElheran. 2020. "Twisting the Demand Curve: Digitalization and the Older Workforce." *NBER Working Paper*, no. w28094.
- Barth, Erling, James Davis, and Richard B Freeman. 2018. "Augmenting the Human Capital Earnings Equation with Measures of Where People Work." *Journal of Labor Economics* 36 (S1): S71–97.
- Becker, Gary S. 1981. *A Treatise on the Family*. Harvard University Press.
- . 1993. *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*. 3rd ed. University of Chicago press.
- Benzell, Seth G., Laurence J. Kotlikoff, Guillermo LaGarda, and Jeffrey D. Sachs. 2016. "Robots Are Us: Some Economics of Human Replacement." NBER Working Paper 20941. Working Paper Series. National Bureau of Economic Research, Inc. <https://doi.org/10.3386/w20941>.
- Bessen, James. 2003. "Technology and Learning by Factory Workers: The Stretch-out at Lowell, 1842." *The Journal of Economic History* 63 (1): 33–64.
- . 2012. "More Machines, Better Machines... or Better Workers?" *The Journal of Economic History* 72 (1): 44–74.
- . 2015. *Learning by Doing: The Real Connection between Innovation, Wages, and Wealth*. Yale University Press.
- . 2020. "Information Technology and Industry Concentration." *Journal of Law and Economics* 63 (3): 531–55.
- Bessen, James, Erich Denk, and Chen Meng. 2021. "Firm Differences: Skill Sorting and Software." Working Paper.
- Bessen, James, Maarten Goos, Anna Salomons, and Wiljan Van den Berge. 2022. "Automatic Reaction-What Happens to Workers at Firms That Automate?" *Boston Univ. School of Law, Law and Economics Research Paper*.
- Bessen, James, Stephen Michael Impink, Lydia Reichensperger, and Robert Seamans. 2018. "The Business of AI Startups." *Boston Univ. School of Law, Law and Economics Research Paper*, no. 18–28.
- Bessen, James, Chen Meng, and Erich Denk. 2020. "Perpetuating Inequality: What Salary History Bans Reveal about Wages." *Available at SSRN*.
- Bessen, James, and Cesare Righi. 2019. "Information Technology and Firm Employment." *Boston Univ. School of Law, Law and Economics Research Paper*, no. 19–6. <https://ssrn.com/abstract=3371016>.
- Bloom, Nicholas, Raffaella Sadun, and John Van Reenen. 2012. "Americans Do IT Better: US Multinationals and the Productivity Miracle." *American Economic Review* 102 (1): 167–201.
- Bresnahan, Timothy F, Erik Brynjolfsson, and Lorin M Hitt. 2002. "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence." *The Quarterly Journal of Economics* 117 (1): 339–76.
- Brynjolfsson, Erik. 2021. "Augmentation, Not Automation." *Boston Review*, May 20, 2021. https://bostonreview.net/forum_response/augmentation-not-automation/.

- Brynjolfsson, Erik, and Lorin M Hitt. 2000. "Beyond Computation: Information Technology, Organizational Transformation and Business Performance." *Journal of Economic Perspectives* 14 (4): 23–48.
- Brynjolfsson, Erik, and Andrew McAfee. 2014. *The Second Machine Age: Work, Progress, and Prosperity in a Time of Brilliant Technologies*. WW Norton & Company.
- Burdett, Kenneth, and Dale T Mortensen. 1998. "Wage Differentials, Employer Size, and Unemployment." *International Economic Review*, 257–73.
- Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. 2018. "Firms and Labor Market Inequality: Evidence and Some Theory." *Journal of Labor Economics* 36 (S1): S13–70.
- Card, David, Jörg Heining, and Patrick Kline. 2013. "Workplace Heterogeneity and the Rise of West German Wage Inequality." *The Quarterly Journal of Economics* 128 (3): 967–1015.
- Carnevale, Anthony P, Tamara Jayasundera, and Dmitri Repnikov. 2014. "Understanding Online Job Ads Data." *Georgetown University, Center on Education and the Workforce, Technical Report (April)*.
- Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer. 2019. "The Effect of Minimum Wages on Low-Wage Jobs*." *The Quarterly Journal of Economics* 134 (3): 1405–54. <https://doi.org/10.1093/qje/qjz014>.
- Clark, Gregory. 1987. "Why Isn't the Whole World Developed? Lessons from the Cotton Mills." *The Journal of Economic History* 47 (1): 141–73.
- Cortes, Guido Matias, Adrian Lerche, Uta Schönberg, and Jeanne Tschopp. 2020. "Technological Change, Firm Heterogeneity and Wage Inequality." Mimeo (https://www.dropbox.com/s/jcpm28binsgqhb/Draft_v1_0.pdf).
- Deming, David J. 2017. "The Growing Importance of Social Skills in the Labor Market." *The Quarterly Journal of Economics* 132 (4): 1593–1640.
- Deming, David, and Lisa B Kahn. 2018. "Skill Requirements across Firms and Labor Markets: Evidence from Job Postings for Professionals." *Journal of Labor Economics* 36 (S1): S337–69.
- Dillender, Marcus, and Eliza Forsythe. 2019. "Computerization of White Collar Jobs." Upjohn Institute Working Paper 19-310.
- Domini, Giacomo, Marco Grazzi, Daniele Moschella, and Tania Treibich. 2021. "Threats and Opportunities in the Digital Era: Automation Spikes and Employment Dynamics." *Research Policy* 50 (7): 104137.
- Doms, Mark, Timothy Dunne, and Kenneth R Troske. 1997. "Workers, Wages, and Technology." *The Quarterly Journal of Economics* 112 (1): 253–90.
- Doms, Mark E., and Timothy Dunne. 1998. "Capital Adjustment Patterns in Manufacturing Plants." *Review of Economic Dynamics* 1 (2): 409–29. <https://doi.org/10.1006/redo.1998.0009>.
- Dunne, Timothy, Lucia Foster, John Haltiwanger, and Kenneth R Troske. 2004. "Wage and Productivity Dispersion in United States Manufacturing: The Role of Computer Investment." *Journal of Labor Economics* 22 (2): 397–429.
- Eeckhout, Jan, and Philipp Kircher. 2018. "Assortative Matching with Large Firms." *Econometrica* 86 (1): 85–132.
- Gaggl, Paul, and Greg C Wright. 2017. "A Short-Run View of What Computers Do: Evidence from a Uk Tax Incentive." *American Economic Journal: Applied Economics* 9 (3): 262–94.

- Graetz, Georg, and Guy Michaels. 2018. "Robots at Work." *Review of Economics and Statistics* 100 (5): 753–68.
- Hakanson, Christina, Erik Lindqvist, and Jonas Vlachos. 2020. "Firms and Skills: The Evolution of Worker Sorting." *Journal of Human Resources*.
- Haltiwanger, John, Russell Cooper, and Laura Power. 1999. "Machine Replacement and the Business Cycle: Lumps and Bumps." *American Economic Review* 89 (4): 921–46.
- Harrigan, James, Ariell Reshef, and Farid Toubal. 2021. "The March of the Techies: Job Polarization within and between Firms." *Research Policy* 50 (7): 104008.
- Hémous, David, and Morten Olsen. 2022. "The Rise of the Machines: Automation, Horizontal Innovation and Income Inequality." *American Economic Journal: Macroeconomics* 14 (1): 179–223.
- Hirvonen, Johannes, Aapo Stenhammer, and Joonas Tuhkuri. 2021. "New Evidence on the Effect of Technology on Employment and Skill Demand." Working Paper.
- Humlum, Anders. 2019. "Robot Adoption and Labor Market Dynamics." Working Paper.
- Hunt, Jennifer, and Ryan Nunn. 2019. "Is Employment Polarization Informative About Wage Inequality and Is Employment Really Polarizing?" Working paper. National Bureau of Economic Research.
- Khaouja, Imane, Ghita Mezzour, Kathleen M Carley, and Ismail Kassou. 2019. "Building a Soft Skill Taxonomy from Job Openings." *Social Network Analysis and Mining* 9 (1): 43.
- Korinek, Anton, and Joseph E Stiglitz. 2018. "Artificial Intelligence and Its Implications for Income Distribution and Unemployment." In *The Economics of Artificial Intelligence: An Agenda*, 349–90. University of Chicago Press.
- Kremer, Michael. 1993. "The O-Ring Theory of Economic Development." *The Quarterly Journal of Economics* 108 (3): 551–75.
- Kremer, Michael, and Eric Maskin. 1996. "Wage Inequality and Segregation by Skill." National bureau of economic research.
- Lachowska, Marta, Alexandre Mas, Raffaele D Saggio, and Stephen A Woodbury. 2020. "Do Firm Effects Drift? Evidence from Washington Administrative Data." National Bureau of Economic Research.
- Lindenlaub, Ilse. 2017. "Sorting Multidimensional Types: Theory and Application." *The Review of Economic Studies* 84 (2): 718–89.
- Modestino, Alicia Sasser, Daniel Shoag, and Joshua Ballance. 2019. "Upskilling: Do Employers Demand Greater Skill When Workers Are Plentiful?" *Review of Economics and Statistics*, 1–46.
- Moscarini, Giuseppe, and Fabien Postel-Vinay. 2016. "Wage Posting and Business Cycles: A Quantitative Exploration." *Review of Economic Dynamics*, Special Issue in Honor of Dale Mortensen, 19 (January): 135–60. <https://doi.org/10.1016/j.red.2015.11.001>.
- Olley, G Steven, and Ariel Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64 (6): 263–97.
- Rodrigo, Rodimiro. 2021. "Robot Adoption, Organizational Capital, and the Productivity Paradox." <https://rodimirodrigo.github.io/JMP2021-Rodrigo.pdf>.
- Rosen, Sherwin. 1981. "The Economics of Superstars." *The American Economic Review* 71 (5): 845–58.
- Sattinger, Michael. 1975. "Comparative Advantage and the Distributions of Earnings and Abilities." *Econometrica: Journal of the Econometric Society*, 455–68.
- . 1993. "Assignment Models of the Distribution of Earnings." *Journal of Economic Literature* 31 (2): 831–80.

- Song, Jae, David J Price, Fatih Guvenen, Nicholas Bloom, and Till Von Wachter. 2019. "Firming up Inequality." *The Quarterly Journal of Economics* 134 (1): 1–50.
- Spitz-Oener, Alexandra. 2006. "Technical Change, Job Tasks, and Rising Educational Demands: Looking Outside the Wage Structure." *Journal of Labor Economics* 24 (2): 235–70.
- Tambe, Prasanna, and Lorin M Hitt. 2012. "The Productivity of Information Technology Investments: New Evidence from IT Labor Data." *Information Systems Research* 23 (3-part-1): 599–617.
- Tambe, Prasanna, Lorin M Hitt, Daniel Rock, and Erik Brynjolfsson. 2019. "IT, AI and the Growth of Intangible Capital." *Available at SSRN 3416289*.

Figures

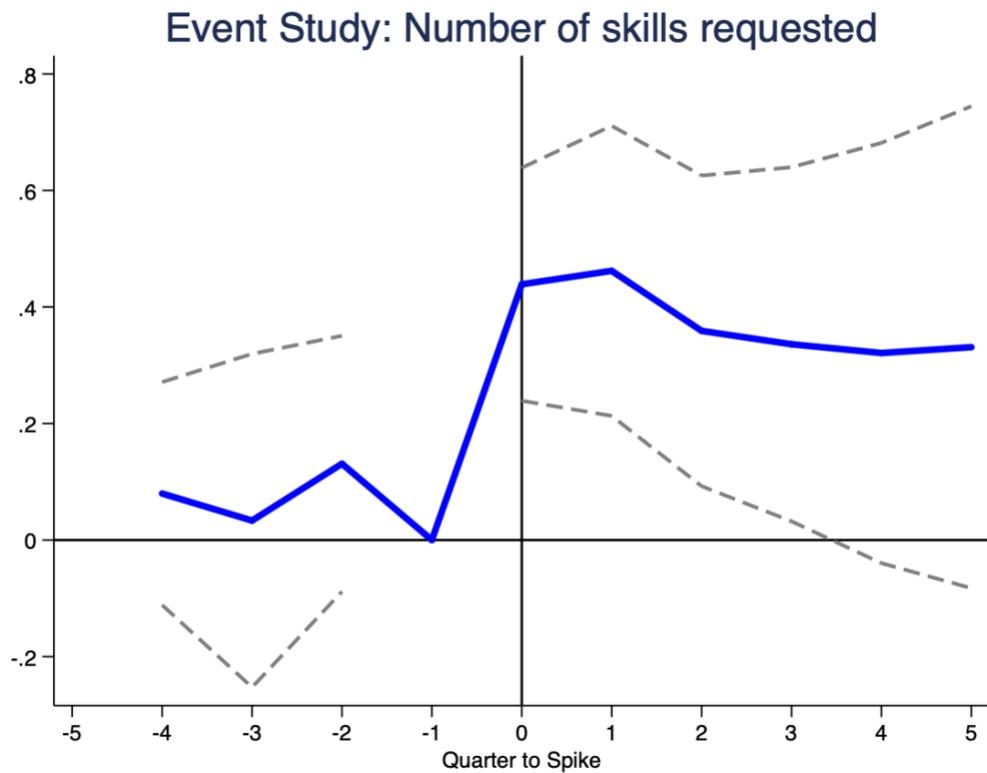


Figure 1. Number of skills requested increases following adoption event.
Note: This figure presents an event study equivalent to Column 1, Panel A, Table 2, reporting the coefficients of quarter dummies for treated firms. The regression is weighted by the number of ads per quarter and it includes fixed effects for quarter and cohort by firm. The dashed lines show the 95% confidence interval with errors clustered by cohort by firm.

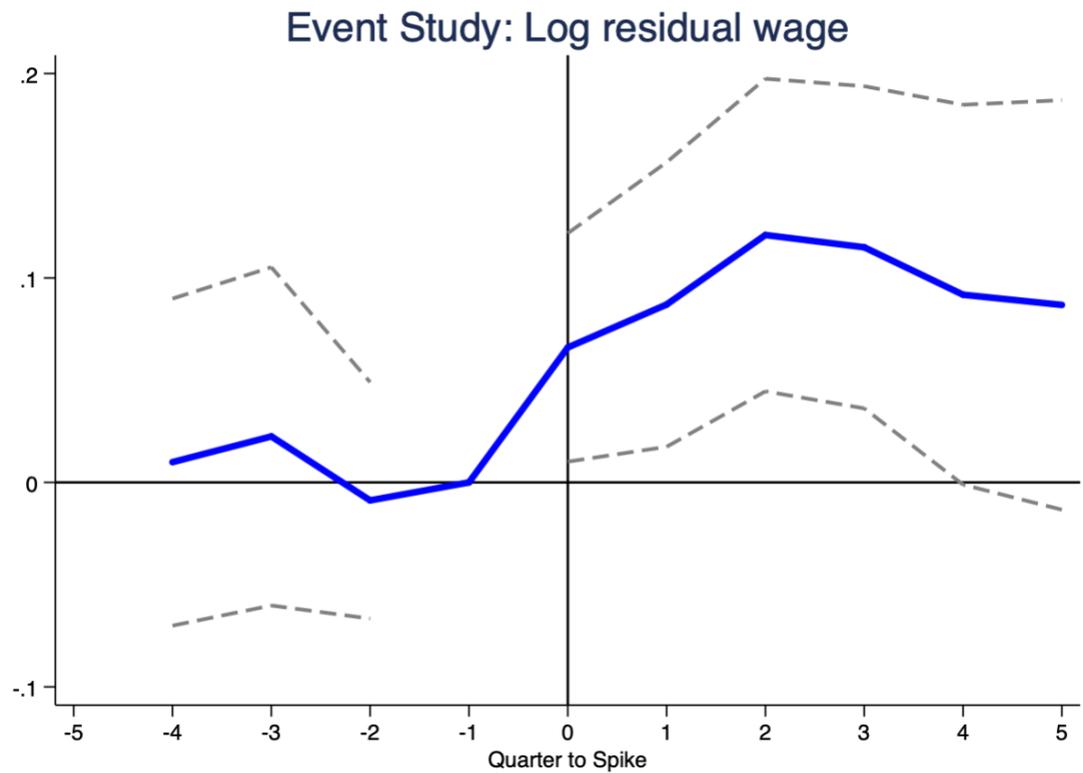


Figure 2. Log residual pay increases following adoption event.
 Note: This figure presents an event study equivalent to Column 1, Panel B, Table 2, reporting the coefficients of quarter dummies for treated firms. The regression is weighted by the number of ads per quarter and it includes fixed effects for quarter and cohort by firm. The dashed lines show the 95% confidence interval with errors clustered by cohort by firm.

Tables

Table 1. Technology Adoption Raises Demands for Specific Skills

Skill measure:	(1) All	(2) IT+AI	(3) Other	(4) Cognitive	(5) Social	(6) Soft	(7) Experience	(8) Education
A. All jobs, number of skills								
Post treatment	0.318*** (0.071)	0.041*** (0.008)	0.173*** (0.059)	0.017*** (0.004)	0.048*** (0.010)	0.038*** (0.010)	0.065*** (0.019)	0.017 (0.020)
Observations	102,086	102,086	102,086	102,086	102,086	102,086	97,045	96,897
R-squared	0.873	0.821	0.868	0.888	0.872	0.868	0.871	0.894
<hr/>								
<i>Pre-Spike Means</i>	10.005	0.518	7.437	0.325	0.762	0.962	3.350	14.581
B. Jobs not requiring college diplomas, number of skills								
Post treatment	0.222*** (0.070)	0.031*** (0.009)	0.105* (0.058)	0.010** (0.004)	0.040*** (0.010)	0.037*** (0.011)	0.041* (0.022)	0.035 (0.024)
Observations	95,679	95,679	95,679	95,679	95,679	95,679	87,220	86,775
R-squared	0.840	0.696	0.843	0.833	0.838	0.826	0.808	0.853
<hr/>								
C. All Jobs, Share of skills								
Post treatment		0.002* (0.001)	-0.008*** (0.002)	0.001 (0.000)	0.002*** (0.001)	0.004** (0.002)		
Observations		102,086	102,086	102,086	102,086	102,086		
R-squared		0.854	0.847	0.853	0.857	0.755		

Note: these coefficients are from stacked difference-in-differences regressions where a balanced panel (t-5 to t+5) is included for each cohort based on spiking year. The unit of observation is firm by quarter. All firms in the sample spike at some time during the sample period and only observations are included that have not spiked previously. All regressions include controls for labor market tightness, management job share, the outsourceable job share, time and cohort x firm fixed effects and standard errors are clustered by cohort x firm (***) p<0.01, ** p<0.05, * p<0.1). To treat heteroscedasticity arising from sample variance, regressions are weighted by the number of help-wanted ads for each firm-quarter. The top panel includes counts of skills requested on all jobs; the bottom panel counts skills only in occupations where the majority of jobs do not require a college diploma. IT jobs (SOC 15) are excluded from the regressions.

Table 2. Adoption of Technology Raises Skill Demands and Pay Across Skill Groups

Skill group:	(1) All	(2) College not required	(3) Routine Cognitive	(4) Routine Manual	(5) Nonroutine Cognitive	(6) Nonroutine Manual
A. Dependent variable: number of specific skills requested						
Post treatment	0.318*** (0.071)	0.222*** (0.070)	0.398*** (0.087)	0.376*** (0.105)	0.512*** (0.091)	0.153 (0.144)
Observations	102,086	95,679	97,117	69,798	100,449	62,967
R-squared	0.873	0.840	0.803	0.771	0.816	0.732
B. Dependent variable: Log Residual Pay						
Post treatment	0.087*** (0.023)	0.054** (0.024)	0.067** (0.029)	0.067* (0.037)	0.091*** (0.032)	0.023 (0.031)
Observations	29,437	21,073	15,617	10,820	20,092	9,345
R-squared	0.476	0.557	0.543	0.622	0.473	0.627

Note: these coefficients are from stacked difference-in-differences regressions where a balanced panel (t-5 to t+5) is included for each cohort based on spiking year. The unit of observation is firm by quarter. All firms in the sample spike at some time during the sample period and only observations are included that have not spiked previously. All regressions include controls for labor market tightness, management job share, the outsourceable job share, time and cohort x firm fixed effects and standard errors are clustered by cohort x firm (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). To treat heteroscedasticity arising from sample variance, regressions are weighted by the number of help-wanted ads for each firm-quarter. The dependent variable in the top panel is the total number of skills requested per ad; the dependent variable in the bottom panel is the log residual salary offered after controlling for experience, experience squared, education, detailed occupation, state, year, and a measure of labor market tightness. IT jobs (SOC 15) are excluded from the dependent variables.

Table 3. Skill and Pay Treatment Effects are Robust to Controls

Dependent Variable	Number of Skills Requested			Log Residual Pay		
	(1)	(2)	(3)	(4)	(5)	(6)
Post treatment	0.245** (0.117)	0.214* (0.112)	0.213** (0.107)	0.102*** (0.035)	0.102*** (0.035)	0.100*** (0.036)
Labor market tightness		0.284 (1.064)	0.194 (1.097)		-0.849 (0.580)	-0.802 (0.591)
Management jobs		6.653*** (0.676)	6.595*** (0.650)		-0.272 (0.221)	-0.271 (0.201)
Outsourceable jobs		-7.210*** (1.793)	-7.188*** (1.768)		0.050 (0.314)	0.004 (0.316)
Growth Rate of Sales			0.260* (0.154)			0.071 (0.060)
Lag CEO change			-0.966 (1.032)			-0.087* (0.051)
3 rd order productivity polynomial			✓			✓
Polynomial probability value			0.013			0.023
Observations	14,008	14,008	14,008	4,706	4,706	4,706
R-squared	0.873	0.882	0.884	0.461	0.465	0.468

Note: these coefficients are from stacked difference-in-differences regressions where a balanced panel (t-5 to t+5) is included for each cohort based on spiking year. The unit of observation is firm by quarter. All firms in the sample spike at some time during the sample period and only observations are included that have not spiked previously. The sample in this table includes only firms that have been matched to Compustat in order to include additional control variables. All regressions include time and cohort x firm fixed effects and standard errors are clustered by cohort x firm (***) p<0.01, ** p<0.05, * p<0.1). To treat heteroscedasticity arising from sample variance, regressions are weighted by the number of help-wanted ads for each firm-quarter. The dependent variable in the first three columns is the total number of skills requested per ad; the dependent variable in columns 4-6 is the log residual salary offered after controlling for experience, experience squared, education, detailed occupation, state, year, and a measure of labor market tightness. The polynomial used in columns 3 and 6 includes log real cost of goods sold and log real beginning-of-quarter capital. The probability value reported is for the F-test of the null hypothesis that polynomial coefficients are jointly zero. IT jobs (SOC 15) are excluded from the dependent variables.

Table 4: Technology Adoption and Changes in Hiring

Skill Group:	(1) College required	(2) College not required	(3) Routine Cognitive	(4) Routine Manual	(5) Nonroutine Cognitive	(6) Nonroutine Manual
A. Share of Hiring						
Post treatment	0.017*** (0.002)	-0.017*** (0.002)	0.007*** (0.003)	-0.008*** (0.002)	0.021*** (0.002)	-0.006*** (0.002)
Observations	103,547	103,547	103,594	103,594	103,594	103,594
R-squared	0.963	0.963	0.910	0.964	0.957	0.970
B. Log level of Hiring						
Post treatment	0.018 (0.030)	-0.083** (0.033)	0.035 (0.032)	-0.107*** (0.040)	0.029 (0.030)	-0.026 (0.048)
Observations	103,404	103,413	97,567	71,018	100,747	64,290
R-squared	0.920	0.927	0.925	0.925	0.923	0.923

Note: these coefficients are from stacked difference-in-differences regressions where a balanced panel (t-5 to t+5) is included for each cohort based on spiking year. The unit of observation is firm by quarter. All firms in the sample spike at some time during the sample period and only observations are included that have not spiked previously. All regressions include controls for labor market tightness, management job share, the outsourceable job share, time and cohort x firm fixed effects and standard errors are clustered by cohort x firm (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). To treat heteroscedasticity arising from sample variance, regressions are weighted by the number of help-wanted ads for each firm-quarter. The columns designate different skill groups. The dependent variable in the top panel is the group's share of job ads; the dependent variable in the bottom panel is the log of the number of job ads. IT jobs (SOC 15) are excluded from the dependent variables.

Table 5. Information Technology Accounts for Most of the Correlation Between Firm Fixed Effects and Skills

Skill measure:	(1) All	(2) IT+AI	(3) Other	(4) Cognitive	(5) Social	(6) Soft	(7) Experience	(8) Education
Panel A, simple correlation								
Firm FE	2.619*** (0.358)	0.722*** (0.069)	1.338*** (0.234)	0.212*** (0.024)	0.203*** (0.065)	0.144** (0.071)	1.500*** (0.133)	2.912*** (0.353)
Standardized coefficient	<i>0.176</i>	<i>0.246</i>	<i>0.125</i>	<i>0.207</i>	<i>0.111</i>	<i>0.067</i>	<i>0.284</i>	<i>0.203</i>
Observations	205,306	205,306	205,306	205,306	205,306	205,306	205,306	205,306
R-squared	0.031	0.060	0.016	0.043	0.012	0.004	0.081	0.041
Panel B, software controls								
Firm FE	0.613** (0.295)	0.075*** (0.021)	0.557** (0.222)	0.074*** (0.019)	-0.016 (0.059)	-0.078 (0.064)	0.545*** (0.100)	1.330*** (0.323)
Standardized coefficient	<i>0.041</i>	<i>0.026</i>	<i>0.052</i>	<i>0.072</i>	<i>-0.009</i>	<i>-0.036</i>	<i>0.103</i>	<i>0.093</i>
Software share	36.738*** (1.436)	5.601*** (0.170)	20.033*** (1.084)	2.683*** (0.092)	4.105*** (0.223)	4.315*** (0.261)	13.673*** (0.446)	31.272*** (1.203)
Software share ²	-54.94*** (2.254)	-0.035 (0.296)	-37.63*** (1.713)	-4.21*** (0.146)	-6.26*** (0.365)	-6.79*** (0.428)	-15.34*** (0.736)	-49.85*** (1.887)
Observations	205,306	205,306	205,306	205,306	205,306	205,306	205,306	205,306
R-squared	0.255	0.760	0.125	0.278	0.192	0.143	0.454	0.202
SW share of sorting	77%	89%	58%	65%	108%	154%	64%	54%

Note: This table regresses firm mean levels of skill counts, experience and education required against firm wage fixed effects. The unit of observation is the firm. Firm fixed effects are calculated by regressing log salary offered against detailed occupation, industry, state, year, labor market tightness, skills requested, education required, experience required, and firm fixed effects. IT jobs are excluded for the estimates. The regressions are weighted by the number of job ads and errors are robust to heteroscedasticity. The bottom panel ads controls for the share of software developers in firm hiring. The standardized coefficients reflect the correlations between the dependent variables and firm fixed effects. Adding controls for software developers substantially reduces these correlations. The bottom row displays the magnitude of that decrease as one minus the standardized coefficient in Panel B over the standardized coefficient in Panel A.

Appendix

A. Model

Sorting equilibrium

We can write the first order condition for q_i , holding the quality of other tasks, q_j , constant as

$$q_j^{N-I-1}V - \theta h'(q_i)l_i = 0.$$

Taking the implicit derivative,

$$\frac{dq_i}{d\theta} = -\frac{h'(q_i)}{\theta h''(q_i)} < 0.$$

The equilibrium value of q decreases with θ . From this it follows that

$$\frac{d \theta h'(q_i)q_i}{d\theta} = h'(q_i) - \theta(h''(q_i)q_i + h'(q_i))\frac{dq_i}{d\theta} = -\frac{(h'(q_i))^2}{h''(q_i)} < 0.$$

Since, as in the text, $w_j = \theta_j h'(q_i)q_i$, the fact that $\theta_H < \theta_L$ implies that $w_H > w_L$ in equilibrium.

Change in between-firm wage ratio

It is convenient to express output in intensive form,

$$y \equiv \frac{Y}{L} = A \cdot Q \cdot k^\alpha, \quad k \equiv \frac{K}{L}$$

so that the first order profit maximizing condition for labor and capital can be written

$$w = (1 - \alpha)y, \quad k = \frac{\alpha}{r}y.$$

Using these, we have³¹

$$\Delta \ln \omega = \Delta \ln(1 - \alpha_H) + \Delta \ln \frac{y_H}{y_L} \approx -\frac{1}{N-I} + \Delta \ln \frac{y_H}{y_L}.$$

Further,

³¹ α_H increases from $\frac{I-1}{N}$ to $\frac{I}{N}$.

$$\Delta \ln \frac{y_H}{y_L} > \Delta \ln A_H + \Delta \ln Q_H + \alpha_L \Delta \ln \frac{k_H}{k_L}.$$

The last term, which did not appear in the case of uniform workers and firms, captures the shift in capital from low type firms to high type firms as the productivity of the high type firms rises, raising the returns for capital per worker. The expression is an inequality because it ignores the increase in α for high type firms. Also, using the first order condition for capital,

$$\Delta \ln \frac{k_H}{k_L} = \Delta \ln \alpha_H + \Delta \ln \frac{y_H}{y_L} \approx \frac{1}{I} + \Delta \ln \frac{y_H}{y_L}.$$

Substituting this into the previous expression,

$$\Delta \ln \frac{y_H}{y_L} > \frac{1}{1 - \alpha_L} \left[\Delta \ln A_H + \Delta \ln Q_H + \frac{1}{I} \right]$$

and

$$\Delta \ln \omega > -\frac{1}{N - I} + \frac{1}{1 - \alpha_L} \left[\Delta \ln A_H + \Delta \ln Q_H + \frac{1}{I} \right].$$

B. Skill measures

Burning Glass standardizes specific skills requested into 16,050 skills. For our analysis, we constructed 6 mutually exclusive skill categories: IT, AI, cognitive, social, other soft skills, and an additional “other” category. We begin with the definition of social and cognitive skills used by Deming and Khan (2018). We then assign IT, AI, and other soft skills using lists of skill terms not included in the Deming and Khan categories. This last category is the largest and contains many skills related to non-IT technologies and to industry knowledge. For our main analysis, we combine the AI and IT categories, but separate analysis indicates that spikes at firms that hire AI personnel perform much like firms that apparently use non-AI software methods (see Table A7 below). The frequencies with which ads request skills in each category are

<u>Category</u>	<u>Percent of job ads</u>
Other	68.56
IT	13.08
Other soft	8.18
social	6.92
cognitive	3.18
AI	0.08

Cognitive Skills (D. Deming and Kahn 2018)

These skills include the keywords Problem Solving, Research, Analytical, Critical Thinking, Math, and Statistics.

Social Skills (D. Deming and Kahn 2018)

These skills include the keywords Communication, Teamwork, Collaboration, Negotiation, and Presentation.

Other Soft Skills* Keywords (adapted from Khaouja et al. (2019) taxonomy):

Accountability	Ethic	Social skills
Active listening	Flexibility	Speaking
Adaptive	Goal	Strategic thinking
Argumentation	Hospitality	Time management
Coaching	Impartiality	Trustworthy
Commitment	Influence	Verbal communication
Conceptual	Initiative	Writing
Conflict management	Integrity	Written communication
Coordination	Interpersonal communication	
Creativity	Kindness	
Curiosity	Leadership	
Decision	Mentoring	
Decision making	Motivated	
Detail	Optimism	
Diverse	Passion	
Eagerness	Persuasion	
Emotional intelligence	Self-confidence	
Enthusiasm	Self-organized	

* These skills also have synonyms, which were also flagged. For full list of synonyms, please refer to Table 13 in Khaouja et al 2019. To further augment this list, the following commonly requested Burning Glass skills not already identified as a social skill were also flagged as soft skills: Planning, Detail-Oriented, Building Effective Relationships, Energetic, Positive Disposition, Listening, Team Building, Creative Problem Solving, Self-Motivation, Overcoming Obstacles, Multi-Tasking, People Management, Thought Leadership, Team Management. This list excludes skills already identified as social or cognitive skills above.

Other Skills

Skills that do not belong to one of the other five groups are designated as “other”. These skills tend to be industry-specific or firm specific. A majority of skills fit in this category. Examples include 5G Wireless, ACL Surgery, Adhesives Industry Knowledge, and APA Style Guide.

AI Skills (Following Alekseeva et al. (2020))

AI ChatBot	Latent Semantic Analysis	OpenNLP
AI KIBIT	Lexalytics	Pattern Recognition
ANTLR	Lexical Acquisition	Pybrain
Apertium	Lexical Semantics	Random Forests
Artificial Intelligence	Libsvm	Recommender Systems
Automatic Speech Recognition (ASR)	Machine Learning	Semantic Driven Subtractive Clustering Method (SDSCM)
Caffe Deep Learning Framework	Machine Translation (MT)	Semi-Supervised Learning
Chatbot	Machine Vision	Sentiment Analysis / Opinion Mining
Computational Linguistics	Madlib	Sentiment Classification
Computer Vision	Mahout	Speech Recognition
Decision Trees	Microsoft Cognitive Toolkit	Supervised Learning (Machine Learning)
Deep Learning	MLPACK (C++ library)	Support Vector Machines (SVM)
Deeplearning4j	Mlpy	TensorFlow
Distinguo	Modular Audio Recognition Framework (MARF)	Text Mining
Google Cloud Machine Learning Platform	MoSes	Text to Speech (TTS)
Gradient boosting	MXNet	Tokenization
H2O (software)	Natural Language Processing	Torch (Machine Learning)
IBM Watson	Natural Language Toolkit (NLTK)	Unsupervised Learning
Image Processing	ND4J (software)	Virtual Agents
Image Recognition	Nearest Neighbor Algorithm	Vowpal
IPSoft Amelia	Neural Networks	Wabbit
Ithink	Object Recognition	Word2Vec
Keras	Object Tracking	
Latent Dirichlet Allocation	OpenCV	

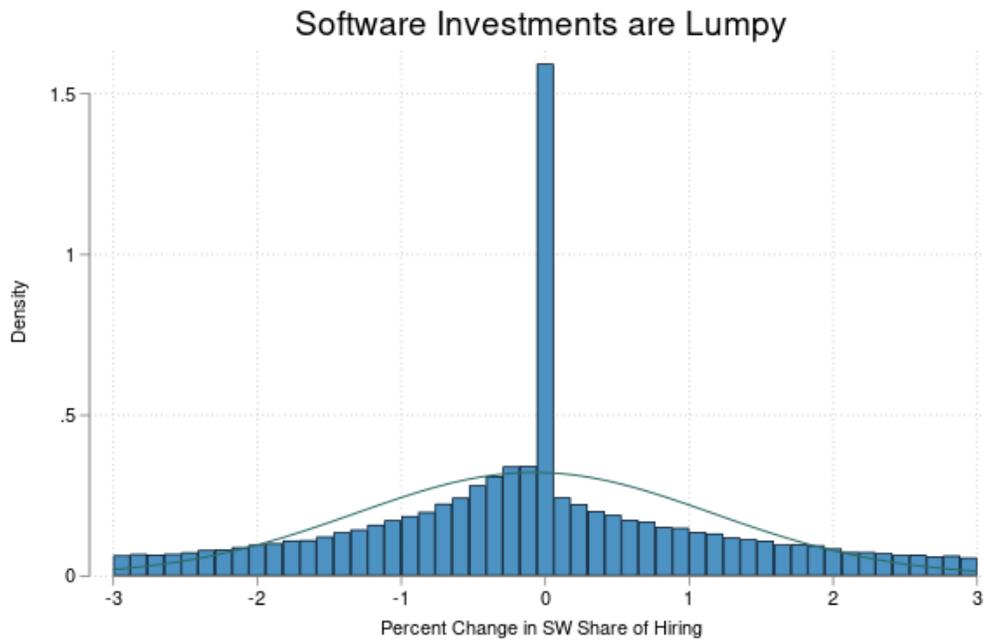
IT Skills (Following Burning Glass Technologies Skill Cluster Families)

Microsoft Development Tools	Enterprise Content Management (ECM)	Productivity Software
Document Management Systems	Internet of Things (IoT)	File Transfer Software
General Networking	Enterprise Management Software	Project Management Software
Software Quality Assurance	Database Administration	Virtual Private Networks
Artificial Intelligence	Android Development	Internet Standards
Operating Systems	Mobile Development	Remote Desktop Software
JavaScript and jQuery	IT Automation	Data Wrangling
Distributed Computing	Configuration Management	Programming Principles
Application Programming Interface (API)	Anti-Malware Software	Network File System (NFS)
Systems Administration	Middleware	Integrated Development Environments (IDEs)
Web Development	Scripting	Disk Imaging
Scripting Languages	Java	Microsoft Office and Productivity Tools
Cloud Solutions	Database Management Systems	Content Management Systems
Cloud Computing	Web Servers	Firewall Software
Software Development Tools	Version Control	Firmware
Data Storage	iOS Stack	Graph Databases
Virtual Machines (VM)	Basic Computer Knowledge	Identity Management
Big Data	Application Development	Partitioning Software
Network Security	Network Protocols	Video Conferencing Software
Data Warehousing	Technical Support	Computer Hardware
Enterprise Messaging	Application Security	Internet Services
Cloud Storage	Typesetting Software	Internet Security
XML Markup Languages	Geographic Information System (GIS) Software	Help Desk Support
Extraction, Transformation, and Loading (ETL)	Data Compression	Management Information System (MIS)
System Design and Implementation	Assembly Languages	Intelligent Maintenance Systems
Network Configuration	Test Automation	Query Languages
Data Synchronization	Telecommunications	Load Balancing
Other Programming Languages	Compiling Tools	Location-based Software
Data Management	Enterprise Resource Planning (ERP)	Video Compression Standards
Web Content	Backup Software	Microsoft SQL Extensions
SAP	Web Design	Advanced Microsoft Excel
Archiving Software	Rule Engines	SQL Databases and Programming
Cybersecurity	Internet Protocols	Device Management
NoSQL Databases	Extensible Languages	Microsoft Windows
Software Development Principles	C and C++	Augmented Reality / Virtual Reality (AR / VR)
IT Management	Desktop and Service Management	Enterprise Information Management
Software Development Methodologies	Mainframe Technologies	Oracle
Content Delivery Network (CDN)	Parallel Computing	Servers
Networking Hardware	Cache (computing)	Data Collection
Information Security	PHP Web	Wiki

Note: There are 1,687 unique skills that Burning Glass identifies as Information Technology skills. From there, they sort these skills into broader categories, which are listed in the table below. Within the category “Microsoft Development Tools” is the Microsoft Office suite, which we omit as an IT skill. We exclude skills flagged as social, cognitive or AI skills. These specific skills include Communications Protocols, Data Communications, Global System for Mobile Communications, Joint Worldwide Intelligence Communications System, Machine-To-Machine (M2M) Communications, Oracle Fusion Middleware Collaboration Suite, and Voice Communications.

C. Lumpy Investment

Figure A1. Lumpiness of Firm Investments



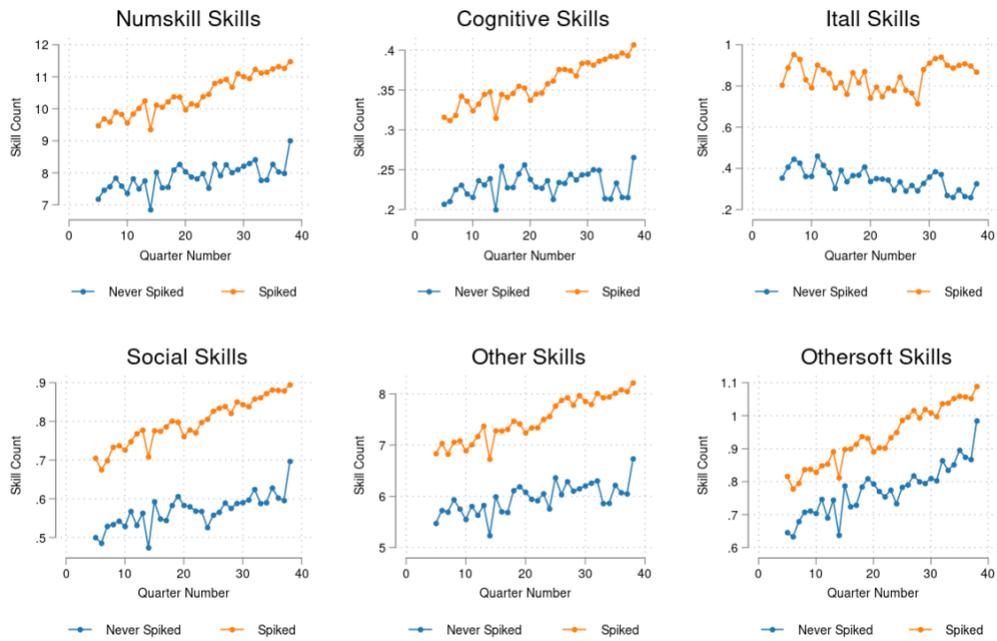
Note: This figure plots changes in software developer share of job advertisements from the average of the previous 4 quarters. The line shows a normal density distribution with the same mean and standard deviation. The distribution is clearly leptokurtic with a peak at zero and fat, “lumpy” tails.

Figure A2. Software Hiring Increases Persist After Spikes



Note: This figure plots an event study of the share of software hiring around hiring spikes. There appears to be a slight anticipation effect, a distinct spike (the threshold is .01), and sustained hiring of software developers at a slightly lower level after the spike.

Figure A3. Skill request trends over time



Note: This figure shows raw trends in skill requests for both spiking (orange) and non-spiking (blue) firms over time. Spiking firms have higher levels of skill requests throughout the sample.

D. Descriptive Statistics and Robustness Checks

Table A1 Summary Statistics

	(1)	(2)	(3)
Sample:	Full sample	Never-Spikers	Spikers
Weighted			
Management Job Share	0.126 (0.190)	0.120 (0.219)	0.139 (0.0960)
Outsourceable Job Share	0.071 (0.183)	0.078 (0.209)	0.056 (0.101)
Labor Market Tightness	0.795 (0.319)	0.837 (0.364)	0.700 (0.139)
IT Share	0.095 (0.160)	0.074 (0.166)	0.108 (0.155)
Residual Wage	0.012 (0.291)	-0.002 (0.336)	0.023 (0.250)
College Required	0.433 (0.279)	0.416 (0.303)	0.471 (0.213)
Routine Cognitive	0.298 (0.284)	0.294 (0.325)	0.307 (0.157)
Routine Manual	0.207 (0.304)	0.224 (0.339)	0.170 (0.201)
Non-Routine Cognitive	0.444 (0.343)	0.423 (0.377)	0.490 (0.243)
Non-Routine Manual	0.158 (0.285)	0.177 (0.320)	0.115 (0.177)
Number of Skills	8.230 (4.895)	7.385 (5.210)	10.062 (3.484)
Unweighted			
Number of Ads/Quarter	85.380 (86.637)	5.980 (1.028)	164.780 (47.623)
Total Firms	2,147,578	2,131,972	15,606

Note: Means given with Standard Deviation in parentheses. Weighted estimates use analytical weights by number of job advertisements.

Table A2. Correlations of Software Spikes and Possibly Correlated Variables

	Lagged Independent Variables				
	(1)	(2)	(3)	(4)	(5)
Panel A. All Firms					
Log Job Ads	0.034*** (0.001)				0.035*** (0.001)
Software share		-0.011 (0.009)			0.036*** (0.008)
Outsourceable jobs			-0.042*** (0.013)		-0.062*** (0.013)
Management jobs				0.034*** (0.011)	0.056*** (0.010)
Observations	89,928	89,928	89,928	89,928	89,928
R-squared	0.023	0.000	0.000	0.000	0.025
Panel B. Compustat					
Labor Productivity	0.006* (0.003)				0.016*** (0.004)
Log COGS		0.014*** (0.002)			
Log Capital			0.008*** (0.002)		0.014*** (0.002)
Sales Growth				0.017 (0.012)	0.028** (0.012)
Observations	14,122	14,122	14,122	14,122	14,122
R-squared	0.001	0.006	0.003	0.000	0.007

Note: This table presents simple OLS regressions between a spike and lagged key variables from both Burning Glass and Compustat. All standard errors are clustered at the firm level. (***) $p < 0.01$, (**) $p < 0.05$, (*) $p < 0.1$)

Table A3 Sensitivity Table

	Panel Size			Spike Threshold		
	(1)	(2)	(3)	(4)	(5)	(6)
	t ± 4	t ± 5	t ± 6	.005	.01	.015
A. Dependent variable: number of specific skills requested						
Post treatment	0.283*** (0.076)	0.318*** (0.071)	0.470*** (0.087)	0.356*** (0.068)	0.318*** (0.071)	0.253*** (0.074)
Labor market tightness	-0.229 (0.263)	-0.176 (0.346)	-0.197 (0.444)	0.264 (0.358)	-0.176 (0.346)	-0.173 (0.363)
Management jobs	4.074*** (0.329)	5.488*** (0.386)	5.516*** (0.576)	4.821*** (0.574)	5.488*** (0.386)	4.916*** (0.334)
Outsourceable jobs	-5.701*** (0.951)	-7.183*** (1.310)	-7.799*** (1.655)	-6.707*** (1.287)	-7.183*** (1.310)	-6.324*** (1.026)
Observations	162,924	102,086	61,377	102,520	102,086	98,609
R-squared	0.892	0.873	0.870	0.888	0.873	0.879
B: Dependent variable: Log Residual Pay						
Post treatment	0.078*** (0.022)	0.087*** (0.023)	0.072*** (0.024)	0.074*** (0.024)	0.087*** (0.023)	0.253*** (0.074)
Labor market tightness	-0.133 (0.128)	-0.279* (0.147)	-0.179 (0.166)	-0.326** (0.130)	-0.279* (0.147)	-0.173 (0.363)
Management jobs	-0.091 (0.091)	0.016 (0.101)	0.198 (0.151)	0.055 (0.107)	0.016 (0.101)	4.916*** (0.334)
Outsourceable jobs	-0.098 (0.126)	0.026 (0.129)	0.270 (0.290)	-0.105 (0.149)	0.026 (0.129)	-6.324*** (1.026)
Observations	42,387	29,437	19,395	28,724	29,437	28,924
R-squared	0.522	0.476	0.411	0.462	0.476	0.450

Note: This table shows how estimates change from changing the size of the balanced panel or threshold for defining a spike. Columns (2) and (5) correspond to estimates in Table 2 Column (1). Construction of panels and additional controls follow those described in Table 2. The unit of observation is firm by quarter. All firms in the sample spike at some time during the sample period and only observations are included that have not spiked previously. All regressions include time and cohort x firm fixed effects and standard errors are clustered by cohort x firm (***) p<0.01, ** p<0.05, * p<0.1).

Table A4 Results for Full Sample And Results for Sample Restricted to Later-Spiking Firms

Sample	(1)	(2)	(3)	(4)
	Number of Skills Requested		Log Residual Wage	
	Later-spiking	Full Sample	Later-spiking	Full Sample
Post treatment	0.318*** (0.071)	0.211*** (0.068)	0.087*** (0.023)	0.074*** (0.020)
Labor market tightness	-0.176 (0.346)	-0.105 (0.115)	-0.279* (0.147)	-0.149*** (0.043)
Management jobs	5.488*** (0.386)	3.249*** (0.106)	0.016 (0.101)	-0.159*** (0.036)
Outsourceable jobs	-7.183*** (1.310)	-2.967*** (0.270)	0.026 (0.129)	0.010 (0.045)
Observations	102,086	1,789,706	29,437	387,844
R-squared	0.873	0.890	0.476	0.513

Note: Our main analysis uses panels with control firms that spike subsequently (“later-spiking”). This table compares this sample with a sample that also includes control firms that never spike. Columns (1) and (3) correspond to Column (1) in Table 2, estimating stacked difference-in-differences regressions where a balanced panel (t-5 to t+5) is included for each cohort based on spiking year. The unit of observation is firm by quarter. All regressions include time and cohort x firm fixed effects and standard errors are clustered by cohort x firm (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). In Columns (1) and (3) firms in the sample spike at some time during the sample period and only observations are included that have not spiked previously. In Columns (2) and (4) we remove this restriction, consequently broadening our sample size. The estimates are similar, but we prefer the estimates provided in the main text.

Table A5 Firm Fixed Effects

	(1) Log of Avg Salary
Other Skill Count	0.003*** (0.000)
Cognitive Count	0.006*** (0.000)
Social Count	0.007*** (0.000)
AI Count	0.035*** (0.002)
IT Count	0.012*** (0.000)
Other Soft Count	0.005*** (0.000)
Minimum of the required experience range in years	0.098*** (0.000)
Experience Required Squared	-0.005*** (0.000)
V/U Labor Market Tightness	-0.001 (0.001)
Observations	4,075,295
R-squared	0.688

Note: This table presents the coefficients used to estimate firm fixed effects. All regressions include occupation, education level, year, and state fixed effects and standard errors are heteroskedastic robust (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). Observations are weighted by occupation share in the Current Population Survey.

Table A6 Non-IT Producing Firms

Sample	(1)	(2)	(3)	(4)
	Number of Skills Requested		Log Residual Wage	
	Full	Non-IT	Full	Non-IT
Post treatment	0.318*** (0.071)	0.358*** (0.078)	0.087*** (0.023)	0.091*** (0.025)
Labor market tightness	-0.176 (0.346)	-0.094 (0.363)	-0.279* (0.147)	-0.294* (0.151)
Management jobs	5.488*** (0.386)	5.900*** (0.430)	0.016 (0.101)	-0.005 (0.107)
Outsourceable jobs	-7.183*** (1.310)	-7.132*** (1.392)	0.026 (0.129)	0.025 (0.135)
Observations	102,086	84,261	29,437	25,597
R-squared	0.873	0.879	0.476	0.480

Note: This table compares the outcomes from Table 2 Column (1) to the same specification excluding IT-producing industries. We defined IT-producing industries as 2-digit NAICS codes 51 and 54. To determine a firm's industry from Burning Glass, we assigned the modal 2-digit industry listed in a firm-year. Columns (1) and (3) correspond to Column (1) in Table 2, estimating stacked difference-in-differences regressions where a balanced panel (t-5 to t+5) is included for each cohort based on spiking year. The unit of observation is firm by quarter. All regressions include time and cohort x firm fixed effects and standard errors are clustered by cohort x firm (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$).

Table A7. Firms Using AI and Automation Behave Similarly

VARIABLES	(1) Number of skills requested	(2) Log Residual Wage	(3) Number of skills requested	(4) Log Residual Wage
Non-AI x post treatment	0.330*** (0.089)	0.074*** (0.024)		
AI x post treatment	0.304*** (0.081)	0.096*** (0.027)		
Non-automation x post treatment			0.040 (0.086)	0.100*** (0.037)
Automation x post treatment			0.360*** (0.074)	0.086*** (0.023)
Labor market tightness	-0.177 (0.346)	-0.277* (0.146)	-0.174 (0.345)	-0.280* (0.147)
Management jobs	5.492*** (0.387)	0.012 (0.102)	5.467*** (0.386)	0.017 (0.101)
Outsourceable jobs	-7.179*** (1.308)	0.026 (0.129)	-7.199*** (1.306)	0.026 (0.129)
Observations	102,086	29,437	102,086	29,437
R-squared	0.873	0.476	0.873	0.476

Note: these coefficients are from stacked difference-in-differences regressions where a balanced panel (t-5 to t+5) is included for each cohort based on spiking year. The unit of observation is firm by quarter. All firms in the sample spike at some time during the sample period and only observations are included that have not spiked previously. All regressions include time and cohort x firm fixed effects and standard errors are clustered by cohort x firm (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). To treat heteroscedasticity arising from sample variance, regressions are weighted by the number of help-wanted ads for each firm-quarter. IT jobs (SOC 15) are excluded from the regressions. AI and automation are identified by keywords for skills requested.

Table A8. Placebo: Spikes of engineers and technicians do not display similar effects. Spikes defined for engineers (SOC 17) and technicians (SOC 19) excluding electrical engineers (SOC 172071)

VARIABLES	(1) Number of skills requested	(2) Log Residual Wage
Post treatment	0.094 (0.065)	0.032 (0.033)
Labor market tightness	-0.262 (0.321)	-0.129 (0.127)
Management jobs	5.136*** (0.393)	-0.245 (0.163)
Outsourceable jobs	-6.039*** (0.655)	-0.299*** (0.100)
Observations	97,526	28,920
R-squared	0.884	0.464

Note: these coefficients are from stacked difference-in-differences regressions where a balanced panel (t-5 to t+5) is included for each cohort based on spiking year. The unit of observation is firm by quarter. All firms in the sample spike at some time during the sample period and only observations are included that have not spiked previously. All regressions include time and cohort x firm fixed effects and standard errors are clustered by cohort x firm (***) $p < 0.01$, ** $p < 0.05$, * $p < 0.1$). To treat heteroscedasticity arising from sample variance, regressions are weighted by the number of help-wanted ads for each firm-quarter. IT jobs (SOC 15) are excluded from the regressions.

Table A9. Tests of Pre-trends

F tests of the null hypothesis that event study coefficients are jointly zero prior to the spike, $\delta_{t-2} = \delta_{t-3} = \delta_{t-4} = 0$.

Outcome variable	Probability value
Log residual wage	0.723
<u>Skill measures</u>	
All	0.633
IT+AI	0.371
Other	0.553
Cognitive	0.359
Social	0.196
Soft	0.941
Experience	0.972
Education	0.709

Note: These event study regressions are weighted by the number of ads per quarter and they include fixed effects for quarter and cohort by firm.