BU School of Law

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 Raises Returns to Detailed Skills

James Bessen, Erich Denk, and Chen Meng

7/2022

Abstract: Using help-wanted ad data, this paper argues that automation increases demand for detailed skills that are typically unobserved, but which are major determinants of pay. Following automation events, we find that employers request more detailed skills and they substantially increase pay offers (8.7%). Importantly, these increases are not limited to select occupational groups—they apply to both routine and non-routine jobs, to jobs requiring college and those that do not. To explain this phenomenon, we extend the Acemoglu-Restrepo task-based model of automation to consider labor quality, which depends on workers having task-specific skills. We obtain a Remainder Effect: when automation displaces labor on some tasks, it raises the returns to specific skills on the remaining tasks performed by diverse occupational groups. Because not all firms automate, this effect can raise income dispersion within occupations and between firms, including the sorting of skilled workers to high-paying firms. In contrast, labor displacement alone tends to increase between-occupation pay differences.

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 Philippe Aghion, David Autor, Luise Eisfeld, Maarten Goos, Po-Hsuan Hsu, Eric Maskin, Mike Meurer, Abishek Nagaraj, Felix Poege, Ronja Röttger, Anna Salomons, Tim Simcoe, John Turner, and participants in TPRI's seminar for helpful comments.

Introduction

It has long been known that only a part of the rise in income inequality can be explained by worker characteristics that economists typically observe such as occupation, education, and experience. Much of the remaining increase has been attributed to changing returns to unobserved skills (Juhn, Murphy, and Pierce 1993; Acemoglu 2002) as might be explained, for instance, by skill-biased technical change. Because these skills cannot be readily measured, however, analyzing the exact mechanisms behind changes in skill demand has been difficult.[1]

Fortunately, with new data, many of the previously unobserved skills are now partly observable. Employers list detailed skills required for jobs in the title and text of help-wanted advertisements, many of these are highly specific to the tasks and technology of the job (e.g., “third party logistics,” “Java Message Service”), and economists have tied these skill requests to wages and firm performance. Marinescu and Wolthoff (2020) find that advertised job titles explain 90% of the variance in posted salaries. By contrast, six-digit occupation codes only explain a third of the variance. Bana (2021) finds that ad texts account for 87% of the variance in pay including salaries that are not advertised. Deming and Khan (2018) find that select advertised skills are related to pay and also to firm performance. Given that these detailed skills appear to account for much of the variation in pay, might increasing returns to these skills be related to technology?

This paper presents theory and evidence to argue that technological change, specifically automation, increases demand for such detailed skills. Studying over five thousand automation events, we find that adoption of new technology is followed by a significant increase in the number of detailed skills demanded in job advertisements and an 8.7% increase in the pay offered to new hires after controlling for job characteristics.

These findings are important for two reasons: 1) they provide a new explanation of how automation may affect inequality that is different from accounts in the literature, and 2) they help explain the rise in inequality within occupations and between firms/industries. In the literature, the main effect of automation on income inequality is widely seen to occur

[1] One approach has been to use worker fixed effects in wage regressions to capture unobserved skills (Abowd, Kramarz, and Margolis 1999), but this misses dynamic changes.

through the displacement of labor.[2] By definition, automation replaces human workers on some tasks. When automation tends to affect some tasks more than others—for example, routine tasks—demand for occupations that perform those tasks is adversely affected, leading to growing differences in equilibrium wages between occupations. Some researchers see these between-occupation effects as the main source of the growth in economic inequality over the last four decades (Acemoglu and Restrepo 2021).

But the labor displacement story is at some tension with recent findings that most of the rise in wage inequality over the last several decades—as much as 86%—has occurred *within* detailed occupations, not between them (Freeman, Ganguli, and Handel 2020; Hunt and Nunn 2022). Moreover, within-occupation dispersion partly reflects growing pay gaps between firms or industries, and recent studies find that most of the growth in inequality has occurred between firms rather than within firms (Card, Heining, and Kline 2013; Barth, Davis, and Freeman 2018; Song et al. 2019; Lachowska et al. 2020).[3] Our evidence shows that demand for detailed skills increases across educational and occupational groups, for jobs that require a college education as well as jobs that do not, for routine occupations as well as for non-routine ones. Because not all firms automate, this means that automation increases between-firm pay differences and within-occupation pay dispersion. We also show that technology accounts for a substantial part of the sorting of skilled workers to high-paying firms. Thus, while labor displacement might be a significant factor in growing inequality, increased demand for detailed skills provides an additional factor that helps explain important features of inequality growth. Moreover, this new story provides a broader range of conditions where automation positively impacts pay.

We begin with a theoretical model that explains why automation can raise demand for detailed skills. Our model, based on the work of Acemoglu and Restrepo (2018a; 2018b), yields a Remainder Effect (Bessen 2015): automating some tasks can raise the demand for performance quality on the remaining non-automated tasks, in turn raising the desired level of specific skills on those complementary tasks. For example, sales and communication skills

[2] This literature begins with Autor, Levy, and Murnane (2003; see also Brynjolfsson and McAfee 2014; Acemoglu and Restrepo 2018a; 2018b; Benzell et al. 2016; Korinek and Stiglitz 2018; Hémous and Olsen 2022). A related literature looks at displacement and job polarization (Autor, Katz, and Kearney 2008; Goos and Manning 2007; Acemoglu and Autor 2011; Goos, Manning, and Salomons 2014)

[3] Haltiwanger and Spletzer (2020) find the between-firm differences are largely accounted for by inter-industry differences.

became more important for bank tellers when ATMs (automated teller machines) automated many cash handling tasks. More generally, there is evidence that automation changes the nature of tasks performed by different occupations (Autor, Levy, and Murnane 2003; Spitz-Oener 2006; Acemoglu et al. 2020).

Our key insight is that detailed skills matter because the quality of task performance matters. Automation doesn't just alter the relative quantity demanded of different types of labor; it changes the demanded *quality* of labor of each type; quantity does not perfectly substitute for quality. Considerations of product and task quality have been largely missing from models of automation. When production consists of many complementary tasks, the quality of performance on those tasks can be critical. In Kremer's (1993) famous example, the failure of one part doomed the space shuttle *Challenger*. Even when high reliability is not needed, poorly performed tasks can create defects, reducing the value of output, or they can halt production, slowing the rate of output. This may be particularly true when the advantage of automation is not so much reducing costs as it is increasing product quality. Automated machines spin finer yarn than humans, they allow machinists and surgeons to operate at higher precision, and AI systems make more accurate predictions.[4] Yet the quality of task performance often depends critically on the quality of labor, on the ability of labor to perform specific tasks. 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.

Our model extends the Acemoglu-Restrepo model of automation (2018a; 2018b) to include variable task quality as modeled by Kremer and Maskin (Kremer 1993; Kremer and Maskin 1996). When automation increases demand for task quality, firms pay more for workers with task-relevant skills, or they pay to train them in these skills. There are a great many detailed skills important to task performance (16,050 in our data) that are not adequately taught in general education. Many of these skills can only be learned through job experience (e.g., "adhesives industry knowledge") or specialized training, and few have

[4] Researchers find 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; Hirvonen, Stenhammer, and Tuhkuri 2021; Babina et al. 2021).

certifications, prompting firms to screen candidates by requesting these skills in job ads. We argue that greater demand for task-relevant skills means that more of these skills will be listed in job ads.

To validate this model, we study these skills using a difference-in-differences analysis to see how the number of skills requested changes after a major automation event. We study 5,759 automation events identified by major increases in the hiring of software developers.[5] These events mostly involve business process automation technologies such as enterprise resource planning; these types of investments are a major form of automation that is much larger than investment in robots.[6] Consistent with the Remainder Effect, we find significant increases in firm requests for all categories of detailed skills following these automation events, implying greater demand for quality on a wide range of non-automated tasks. Furthermore, we find that these increased skill demands are not limited to specific occupational groups. Skill requests and pay increase both for jobs that require a college education as well as jobs that do not, for routine occupations as well as for non-routine ones. Because these increases occur across most occupations, they can contribute to growing inequality within occupations and between firms. We also find evidence of labor displacement—firms hire fewer workers for routine manual jobs and jobs that don't require college education. Notably, the demands for detailed skills and pay offers increase for these jobs, nevertheless. This implies that the demand for labor quality is orthogonal to the effects of labor displacement, and it may be inadequate to analyze inequality only using occupation-based measures of skill.

How much does this channel contribute to the overall growth in inequality? Song et al. (2019) find that most of the increase in inequality since 1980 is accounted for by greater sorting of highly skilled workers to high-paying firms. We analyze skill and pay levels for the universe of online help wanted ads, finding a substantial role for information technology in accounting for sorting. We estimate firm fixed effects in pay regressions and find that these

---

[5] Other studies have used the employment of software developers as measures of technology adoption (Tambe and Hitt 2012; Tambe et al. 2019; Bessen 2020; Harrigan, Reshef, and Toubal 2021). A variety of paper have also used 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).

[6] Although robots have featured prominently 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).

are correlated with all our skill measures, indicating sorting. However, we find that information technology accounts for most of these correlations between firm fixed effects and skill measures. In other words, the greater use of information technology and the associated increased demand for detailed skills plausibly account for a significant portion of the growth in skill sorting and hence a significant portion of rising pay inequality.

The key contribution of this paper is that it shows theory and evidence for a largely unrecognized channel by which automation can affect wage inequality. First, we develop a model of automation that includes roles for task quality and detailed skills. The model generates a richer set of outcomes including an impact of automation on inequality within occupational and educational groups and on inequality between firms. Second, we study the micro-level impact of information technology automation to test key aspects of the model, making novel use of job ad data to measure changes in skill demand. We find that automation is followed by significantly increased demand for detailed skills and substantially increased pay across occupational groups at the automating firms, consistent with the Remainder Effect. 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 these correlations are substantially accounted for by this technology, suggesting that the increase in sorting may be closely related to the rise of proprietary information technology.

Other papers find an association between information technology and between-firm wage differences (M. Doms, Dunne, and Troske 1997; Dunne et al. 2004; Barth et al. 2020). But these papers do not connect the adoption of technology with subsequent changes in pay and skill demand. A number of recent papers explore 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] Our paper finds similar effects on pay, but we also show the role of task specific skills. Dillender and Forsythe (2019), in an approach similar to ours, use Burning Glass data to identify firm computer

[7] Several other studies find that advanced technologies increase pay at the adopting firm (Humlum 2019; Rodrigo 2021; Genz et al. 2021; Domini et al. 2022). Graetz and Michaels (2018) find a similar increase at the industry level. 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.

technology adoption; they find greater skill demand and higher pay but their study only covers office and administrative support workers.

Some researchers have explored other links between information technology and notions of skill that extend beyond education. 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 a wide variety of detailed skills and firm pay offers. Lindenlaub (2017) argues that multi-dimensional skills are needed to understand the link between sorting and technology. 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. Cortes et al. (2020) model sorting arising from skill-biased technical change. Another line of research aims to understand unobserved skills using 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). Firpo et al. (2011) explore within-occupation wage dispersion, but they do not observe skills or technology adoption events.

Model

Basic Setup

Tasks and Automation

Our model is a combination of automation models by Acemoglu and Restrepo (2018a; 2018b) 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 be $K=\sum_{i=1}^{I} k_i$ 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),

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

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 or defects increases with the number of tasks prone to failure. Multiple tasks provide multiple opportunities for error. Let $q_i,\ 0 \le q_i \le 1$ 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 written[8]

[8] Where *q* represents a probability of successful completion, then *Y* is expected output and we assume that firms are risk neutral. Note that our model does not require this particular functional form; other functions work as long as it is concave, task quality and quantity are not perfect substitutes, and, for sorting, q_i and q_j have positive cross derivatives.

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

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. 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.

Labor Quality

We posit that task quality on the *i*th task is a continuous, twice differentiable concave function of a number of skills, $q_i = q_i(s_{i1}, s_{i2}, \dots)$, where $s_{ij} \geq 0$ is the level of skill. Then,

$$q_i = \begin{Bmatrix} 1, & i \leq I \\ q_i(s_{i1}, s_{i2}, \dots), & I < i \leq N \end{Bmatrix}. \tag{3}$$

Since we are concerned about advertised skills, we assume that these skills are "general" in the sense that they are valuable to multiple employers, and we assume that they are not covered in typical school curricula (otherwise advertisements would merely stipulate education required). Many of these skills are specific to the technology or the task, and many such specific skills are typically required for any particular job.[10] For several reasons, firms may face significant transaction costs when attempting to hire workers with desired skill levels: the number of skills may be large and the local labor market for any particular skill might be quite thin; when firms need *combinations* of skills, the market will be even thinner; typically, few skills are certifiable; many can only be described in loose language, limiting the ability of firms to accurately screen for them. As a result, employers must assume that even

[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.

[10] The mean number of skills listed per ad in our data is 9.3. The variance is 44 after controlling for 6-digit SOC occupation codes, suggesting that many skills are not listed in advertisements and the total number of skills valuable to job performance is much larger than the mean. The skills in our data are overwhelmingly specific to technology or task.

experienced new hires will require some additional training to become fully productive.[11] But this implies that wage bargaining occurs under asymmetric information. A worker's current employer has better information about her skills than does a potential alternative employer.[12]

We can see these challenges in a model of training that includes learning on the job. First, consider a firm hiring a completely unskilled worker ($s_{ij} = 0, \forall i, j$). Let the cost of training be

$$T_{ik} = \theta_k \sum_j c_{ij} s_{ij}, \quad k = L, H, \quad \theta_L > \theta_H \tag{4}$$

where c_{ij} is the unit cost of training skill *j* on task *i* and θ_k is the "skill group" of worker type *k*—type L workers have higher training costs. We assume that firms can select the applicant skill group, for instance, by screening on education and other requirements.[13] For a given level of quality q, cost minimization of (4) gives us (dropping the task subscript) optimal values $\hat{s}_{ij}$ and a convex minimum cost function

$$\hat{T} = \sum_j c_{ij} \hat{s}_{ij} = \theta_k h(q), \quad h', h'' > 0. \tag{5}$$

We assume further assume that $\lim_{q \to 1} h(q) = \infty$, so that no amount of training can achieve perfect quality.

Next consider hiring an experienced worker. For the moment we will ignore screening for specific skills, returning to that topic below. In this case, the hiring firm does not know the job applicant's actual skill levels, but it does know that on average it will need to spend an additional amount on training of $\hat{T} - T^*$, in order to bring the worker's quality up to the level that maximizes firm profit. T^* can be thought of as the effective level of training of the mean experienced job applicant. If the fully trained worker's marginal productivity is $y(\hat{T})$, then a firm will be willing to offer a wage of $y(\hat{T}) - (\hat{T} - T^*)$ to experienced workers.

---

[11] Nor can firms infer that a worker in a given occupation is fully skilled even if they have prior experience. Firms likely differ in their work organizations or specific technologies so that different combinations of skills are optimal for different firms (Lazear 2009). Our model has a different informational structure than Lazear's.

[12] Other papers on training under asymmetric information include Chang and Wang (1996) and Acemoglu and Pischke (1999).

[13] In our data the number of skills listed is correlated with years of education required with a coefficient of .2615 (.0000).

We can then model the process of training in two stages, a learning stage and a production stage. For an inexperienced worker, the firm offers to subsidize training by paying a wage of w_1 in period 1 and to pay wage w_2 in period 2 conditional on the worker realizing productivity of $y(\hat{T})$ in period 2, for a total wage of $W = w_1 + w_2$. If the worker accepts this offer, she will pay T for training, choosing a level to maximize her second period wage, w_2. We assume both firms and workers are risk neutral, neither are credit constrained, firms are homogenous, and we ignore discounting.

Two conditions determine the equilibrium wage. First, in the second period, the worker has the option of taking her newly learned skills to another employer at a wage of $y(\hat{T}) - (\hat{T} - T^*)$ as above. This gives rise to a bilateral monopoly. If the worker accepts this wage and switches jobs, the original employer earns nothing; alternatively, if the worker stays, the parties jointly earn $y(T)$. Assuming a Nash bargaining solution, $w_2 = y(T) - \frac{\hat{T}-T^*}{2}$. Then the worker will choose training of $\hat{T}$ that maximizes $y(T) - T$. Second, in period 1 the worker has the choice of forgoing training and taking an unskilled job at wage w_u in period 2. This means that the worker will only accept the firm's offer as along as $W - \hat{T} \geq w_u$. We assume that this constraint binds at equality and there is an interior solution so that

$$W(q) = w_u + \hat{T}(q) = w_u + \theta_k h(q). \tag{6}$$

The second term represents the skill premium over the unskilled wage.

Screening and skill posting

Firms have an alternative to providing additional training to experienced hires: they can screen for the desired skills by requesting those skills in job advertisements. This would eliminate the cost of training (if the screen is accurate), but it also brings a risk—screening reduces the size of the applicant pool (Marinescu and Wolthoff 2020), so there might not be qualified applicants to hire at any given moment, especially if many skills are requested. This problem is compounded by the limitations of verbal descriptions of skills—some candidates with adequate skill might not apply because they mistakenly overestimate the requested skill level. Let p_j be the probability that no applicant applies when the firm requests skill *j* in the ad. Then a firm seeking to hire a trained worker will choose to post skill *j* in the ad if the expected benefit of screening is greater than the benefit of training (dropping the *i* subscript),

$$\left(1-p_j\right)\left(y(\hat{T})-(\hat{T}-T^*)+c_j\left(\hat{s}_j-s_j^*\right)\right)>y(\hat{T})-(\hat{T}-T^*)$$

or

$$c_j\left(\hat{s}_j-s_j^*\right)>\frac{p}{1-p}(w_u+T^*) \tag{7}$$

that is, if savings of training cost from screening exceed the expected losses from the smaller applicant pool. Posted skills will tend to be those with a low probability of hiring failure, p, (thick markets), or the skill will be common in the applicant pool (high s_j^*), or the training costs are high. Labor market tightness should raise p, thus reducing posting. Note that if automation raises the desired skill levels, $\hat{s}_j$, the likelihood a skill will be posted increases, all else equal. For this reason, we expect that if automation increases the demand for labor quality, we should see a rise in the number of detailed skills in jobs ads.

Homogenous Workers and Firms

We begin 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.

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, firm profit is

$$\pi(q_{I+1},\dots,q_N,k_1,\dots k_I,l_{I+1},\dots,l_N;I)$$
$$=A(I)K^{\alpha}L^{1-\alpha}\prod_{i=I+1}^{N}q_i-\sum_{i=1}^{I}rk_i-\sum_{i=I+1}^{N}W(q)l_i,$$

where *r* is the user cost of capital. 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

$$\frac{Y}{q_i}-\theta h' l_i=\frac{Y}{Nl_i}-W=\frac{Y}{Nk_i}-r=0. \tag{8}$$

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),

$$\frac{d\hat{q}_i}{dA} = \frac{NW}{\theta A q_i h''(q_i)} > 0. \tag{9}$$

Thus, increases in productivity will increase the equilibrium quality of output. When potential output increases, firms increase their training subsidies, workers get more training/skill, and total output increases more than potential output. In other words, an increase in potential output increases the returns to task-related skills.

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, $1 > q_I(\hat{s}_{I1}, \hat{s}_{I2}, \dots)$. 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,

$$\Delta \ln \frac{Y}{L} \approx \Delta \ln A + (N - I) \Delta \ln A \cdot \frac{dq}{dA} \cdot \frac{A}{q} - \ln q_I \,. \tag{10}$$

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 $q_I < 1,\ -\ln q_I > 0$) 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 (10),

$$\Delta \ln W \approx \frac{d \ln(N-I)}{d\, I} + \Delta \ln \frac{Y}{L} \tag{11}$$
$$\approx -\frac{1}{N-I} + \Delta \ln A + (N-I)\Delta \ln A \cdot \frac{dq}{dA} \cdot \frac{A}{q} - \ln f(e_I)$$

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. Generally, the quality terms in (11) provide reasons beyond Acemoglu and Restrepo why wages might increase. The extent of these increases is an empirical matter.

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). And this pattern is consistent with the finding of Autor, Salamons, and Seegmiller (2021) that output-augmenting innovations create new tasks. Similarly, new products might be a form of realizing greater product quality.

Heterogeneous Workers and Firms

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. This section explores how these differences in the cost of learning relate to differences in pay between firms and the impact of automation on these pay gaps. There are two overlapping ways that heterogenous labor can give rise to firm heterogeneity: firms may differ in the skill level of the workforces they hire, and they may differ in their adoption of new technology.

In general, there are two ways that workers can be assigned to firms: assortative matching, where some firms hire more high skill workers while 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). Thus, firm heterogeneity emerges naturally from a model with task quality. Below we briefly consider a slightly different production function that gives rise, instead, to cross-matching.

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 (8) then hold separately for each firm type. Combining the first order conditions for quality and labor, for worker/firm type *j*,

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

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

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

High-type firms pay more, and the between-firm wage gap corresponds directly to differences in skill/training. This gap also represents differences within occupations, that is, differences between workers performing the same task in type L and type H firms. Furthermore, it is straightforward to show that productivity is higher in type H firms: $\frac{W_H}{W_L} = \frac{Y_H}{L_H} / \frac{Y_L}{L_L} > 1$.

The difference in productivity is significant because under some common conditions it means that type H firms will be more likely to adopt new technology than type L firms. 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 (10). 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.[14] Of course, there are many other well-known reasons for firms to differ in productivity. From (9), higher productivity firms will have stronger incentives to increase labor quality, hence stronger incentives to hire type H workers. That is, exogenous productivity differences can also give rise to both sorting and to greater adoption of technology by more productive firms. Generally, we might expect technology adoption to be associated with greater skills.

We can ask what happens when type H firms automate but type L firms do not. We calculate the change in ω using an approach like the one used in equation (11). 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, ω increases. Thus, automation can increase pay differences between firms and within occupations. And because it increases the demand for task-related skills at the same time, it increases the magnitude of sorting. These are the key effects we explore in our empirical analysis.

Cross-matching

But firms do not always sort across employee skill groups; they often crossmatch. For example, Acemoglu and Restrepo's model of automation and inequality exogenously assigns high skill workers to nonroutine tasks and low skill workers to routine tasks (2018a; 2018b). Kremer and Maskin (1996) show that with a slightly different production function than the one above, firms will cross-match under some conditions, hiring both high and low

[14] 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.

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."

Then we can specify an alternative production function where

$$q_i = \begin{Bmatrix} 1, & i \leq I \\ 1, & I < i \leq J \\ q_i(s_{i1}, s_{i2}, \dots), & J < i \leq N \end{Bmatrix}$$

and where $I < J < N$. Routine tasks in the range $I < i \leq J$ are not sensitive to the quality of labor 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 high skill workers for nonroutine tasks and low skill workers for routine tasks. Acemoglu and Restrepo argue that routine tasks are more likely to be automated than non-routine tasks and this gives rise to labor displacement. As automation reduces aggregate demand for type L workers relative to type H workers, pay differences between the associated occupational groups grow. In this case, the effect is greater wage inequality within firms, but with complete cross-matching, there is no difference in pay between firms.

Of course, in the real world we see firms both sorting and cross-matching. Interestingly, Kremer and Maskin (1996) provide a variety of evidence that the extent of 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 or occupation (E. Handwerker 2015; E. W. Handwerker, Spletzer, and others 2016). Note that automation might contribute to this trend. To the extent that routine jobs tend to be automated more—that is, jobs where productivity is less sensitive to skill—conditions for sorting equilibria rather than cross-matching tend to increase.

Empirical Analysis

Data

We measure changes in the demand for detailed skills 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 accounts for 60-70% of all job openings and 80-90% of openings requiring a bachelor's degree or more. 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.[15] 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-related skills, we should see increases in the detailed skills requested in job ads. To the extent that greater demand increases the firm's willingness to pay, we should also 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:

Detailed 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 from 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

[15] 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.

between Job Openings and Labor Turnover Survey (JOLTS) statewide openings for the non-farm sector and the state unemployment rate.[16]

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.[17]

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.[18] 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

---

[16] 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.

[17] These groups are not mutually exclusive.

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

previous four quarters.[19] This approach leverages the finding from the capital investment 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, some 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 find that 81% of our spike events involve these specific automation technologies.[20] Similarly, 31% of the spikes involve firms requesting artificial intelligence skills. Thus, our spikes predominately involve applications that automate tasks.

We seek to estimate average treatment effects around these events using fixed effects regressions. A recent literature highlights estimation problems that arise in two-way fixed effects regressions when treatment effects trend over time (de Chaisemartin and D'Haultfœuille 2020; Callaway and Sant'Anna 2020; Goodman-Bacon 2021). To avoid these problems, we follow Cengiz et al. (2019, Appendix D) and construct balanced panels around each possible spike quarter, excluding firms that have previously spiked.[21] We then run stacked regressions as follows (we report alternative estimates in the Appendix). Let T_i be

[19] 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.

[20] These are jobs requesting skills with keywords ERP, CRM, EDI, MRP, SAP, Automat*, and Robot*. See appendix B for the list of skills we identify as related to Artificial Intelligence, which follows Alekseeva et al. 2020.

[21] Our setting differs slightly from the literature in that firms can have multiple treatment events. To the extent that control firms may have spiked prior to our observed sample, our estimates will tend to be understated.

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.[22] This gives us a degree of identification by removing fixed or slowly 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 "outsourceable" jobs that should track these shifts.[23]

---

[22] Bessen et al. (2022, Appendix) provide a model for differential timing. We also duplicate our results for the full sample and for individually estimated cohorts (Tables A4 and A10).

[23] 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.

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 between Burning Glass and Compustat.[24] One variable is the growth in real sales from the quarter before the spike to a year earlier. The second control is a third order polynomial in log variable costs and log net capital stock (both deflated).[25]
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.
5. **Acquisitions.** Perhaps firms change hiring when they acquire other companies. For the Compustat subsample, we use a binary flag if the firm acquired another firm using data from the Thomson Reuters SDC Platinum database.

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: 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, acquisitions, and management changes. 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.

---

[24] 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.

[25] 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).

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.[26] All skill measures show significant increases following the adoption event except for education. The median number of skills requested in our sample is 8, so the total number of skills requested increases 4% over this baseline. The top panel includes all jobs except for IT jobs (SOC 15).

Panel B includes the skill measures only for jobs that do not require a college diploma.[27] 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 skills 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.[28] The overall impact appears to be that firms request more of the kinds of detailed 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.

Thus, automation is followed by significant increases in requests for a wide variety of task-related skills across different occupational and educational groups. Consistent with our

---

[26] Regressions are weighted by the number of ads to reduce heteroscedasticity from sampling variance 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.

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

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

model, we interpret the greater number of skills requested as evidence of greater demand for these skills. When firms place greater value on "Teamwork" or on "Adhesives Industry Knowledge," they will be more likely to specifically request these skills. But could the increase in skill requests reflect something else, instead? Perhaps it reflects just a change in HR practice. There are two reasons to think that the increase in skill requests reflects an increase in the demand for labor quality. First, there is a cost to advertising detailed skills—they reduce the applicant pool, possibly making hiring slower and more costly (Marinescu and Wolthoff 2020). For this reason, firms tend to reduce requested skills when labor markets are tight; the number of skill requests and labor market tightness are negatively correlated in our data (see also Modestino, Shoag, and Ballance 2016).

Second, the increased skill requests are accompanied by a greater willingness to pay for these skills, that is, firms offer to pay more. 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, but the overall increase is also large (8.7%). Because these increases occur at some firms and not others, they contribute to greater between-firm pay gaps.[29] And because these increases affect most occupational and educational groups, they contribute to greater wage dispersion within these groups.

Table 3 tests the robustness of results to additional controls. Here the sample is limited to firms that are matched to Compustat. Using Compustat, Execucomp, and SDC data, we add a control (in columns 3 and 6) for the rate of revenue growth, a flag for change of CEO, one for acquisitions, and a third order polynomial in log capital and log variable costs to capture productivity 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 show event study graphs corresponding to the first column in Table 2.[30] The graphs show significant and

---

[29] This would not be the case if automation were negatively correlated with firm pay levels, but, in fact, automating firms tend to pay higher residual wages, see Appendix Table A1.

[30] These show the δ_τ coefficients from the following modification of (8):

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

persistent increases in the mean number of skills requested and log residual wages 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 little change in our results when we exclude firms in industries that create software products (NAICS 50 and 54). 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 limiting the analysis to 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 to 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. Finally, in Appendix E, we consider the robustness of our results to alternative ways of handling the two-way fixed effects.

This evidence shows that these automation events increase firms' demand for labor quality, for task-related skills, tending to increase pay gaps within occupations and between firms. But this does not rule out changes in the relative demand for the *quantity* of labor between different occupational groups arising from labor displacement. The top panel in Table 4 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 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, we find evidence of shifts in the relative demand for the quantity of labor consistent with the prior literature. However, it is striking that these shifts appear to be independent of the changes in skill demand. For example, the quantity of labor demanded declines in non-college and routine manual occupations, however, at the same time greater skills are demanded in the hiring that occurs in these groups and higher pay is offered. The quantity and quality of labor are apparently not perfect substitutes and, as a result, changing demand for detailed skills represents a distinct channel by which technology affects wage inequality. Labor displacement, by changing aggregate demand for some occupations, affects equilibrium market wages while demand for detailed skills directly affects firm pay offers.

Sorting

But how significant is the contribution of this channel to the actual change in income inequality? The evidence so far only pertains to a select sample of firms. This section looks at the entire universe of firms that advertise job openings to explore one major component of inequality growth, namely, skill sorting. It is well established that some firms pay more for workers with given characteristics; with positive assortative matching, highly skilled workers tend to work at high-paying firms, enhancing wage dispersion. Song et al. (2019) find that the sorting of highly paid workers to high-paying firms accounts for most of the increase in inequality since 1980.

In our model, automation both raises firm pay levels and raises skills demanded, so automation should contribute to sorting. We begin exploring the magnitude of this effect by estimating fixed firm differences in pay. Specifically, we estimate firm pay fixed effects by regressing pay offered in job ads controlling for job characteristics. 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.[31]

---

[31] 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. Nevertheless, 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.

Next, we 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. We regress skill counts s_i for firm *i* against the firm fixed effect, μ_i,

$$s_i = \beta\mu_i + \epsilon_i, \qquad r = \beta\frac{\sigma_\mu}{\sigma_s}$$

where r is the standardized coefficient, which equals the correlation coefficient. The correlations are all significant for every skill measure, indicating sorting. These figures 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.[32]

Consistent with our model, firm hiring of software developers is correlated with both firm fixed effects and with skill measures.[33] The bottom panel adds quadratic terms in the mean share of software developers in hiring,

$$s_i = \beta^*\mu_i + \gamma_1 x_i + \gamma_2 x_i^2 + \epsilon_i, \qquad r^* = \beta^*\frac{\sigma_\mu}{\sigma_s}$$

where x is the share of software developers. 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, $\frac{r-r^*}{r}$. 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 (according to BEA data), this shift can explain much of the rise in inequality due to sorting.

Conclusion

This paper argues that automation can increase the demand for a wide range of diverse skills, raising pay even as some jobs are displaced. Moreover, demand increases across skill groups, both for jobs requiring college and those that do not, for routine jobs as

---

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

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

well as nonroutine jobs. These broad increases contribute to wage dispersion within occupations and between firms, and to the sorting of skilled workers to high paying firms.

This analysis provides a richer and more optimistic view about the impact of automation than models that presume only a labor displacing effect. For instance, Acemoglu and Restrepo argue that wages will fall after "so-so innovations" where the productivity gain is small. But if automation raises the demand for quality on the remaining tasks, wages may rise even with modestly productive innovations.

The matter is ultimately empirical, but here, too, a richer view of skills affects the analysis. Inequality is frequently measured by differences between occupational or educational groups. Yet it has long been known that a substantial part of inequality arises from "unobserved" skills within these groups. This paper shows that much of that previously unobserved variation in skills can be observed in help wanted advertising, providing a more complete analysis. Importantly, there is good reason to expect that the thousands of skills requested in job ads behave differently: they face thin markets, they lack credentials, and employers face asymmetric information. While these skills might be correlated with education, educational groups do not make a reliable proxy and they provide at best an incomplete picture of the interaction between technology and pay.

Indeed, evidence shows that most of the rise in inequality occurs within occupations, within educational groups, and between firms. This suggests that labor displacement might not be the dominant driver of growing inequality. If so, different policies might be needed 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 affects inequality via greater demand for detailed skills, 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 and skills, one where technology affects demand for a wide array of skills.

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 Jorn-Steffen Pischke. 1999. "Beyond Becker: Training in Imperfect Labour Markets." *The Economic Journal* 109 (453): 112–42.

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. "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, Lawrence F Katz, and Melissa S Kearney. 2008. "Trends in US Wage Inequality: Revising the Revisionists." *The Review of Economics and Statistics* 90 (2): 300–323.

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.

Autor, David, Anna Salomons, and Bryan Seegmiller. 2021. "New Frontiers: The Origins and Content of New Work, 1940–2018." In *Seminar, Harvard Joint Labor and Public Seminar, Virtual*.

Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. 2021. "Artificial Intelligence, Firm Growth, and Product Innovation." *Firm Growth, and Product Innovation (November 9, 2021)*.

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.

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. 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.

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.

Callaway, Brantly, and Pedro H. C. Sant'Anna. 2020. "Difference-in-Differences with Multiple Time Periods." *Journal of Econometrics*, December. https://doi.org/10.1016/j.jeconom.2020.12.001.

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.

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.

Chaisemartin, Clément de, and Xavier D'Haultfœuille. 2020. "Two-Way Fixed Effects Estimators with Heterogeneous Treatment Effects." *American Economic Review* 110 (9): 2964–96. https://doi.org/10.1257/aer.20181169.

Chang, Chun, and Yijiang Wang. 1996. "Human Capital Investment under Asymmetric Information: The Pigovian Conjecture Revisited." *Journal of Labor Economics* 14 (3): 505–19.

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.

———. 2022. "For Whom the Bell Tolls: The Firm-Level Effects of Automation on Wage and Gender Inequality." *Research Policy* 51 (7): 104533.

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/redy.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.

Firpo, Sergio, Nicole M Fortin, and Thomas Lemieux. 2011. "Occupational Tasks and Changes in the Wage Structure." *Available at SSRN 1778886*.

Freeman, Richard B, Ina Ganguli, and Michael J Handel. 2020. "Within-Occupation Changes Dominate Changes in What Workers Do: A Shift-Share Decomposition, 2005–2015." In *AEA Papers and Proceedings*, 110:394–99.

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.

Genz, Sabrina, Terry Gregory, Markus Janser, Florian Lehmer, and Britta Matthes. 2021. "How Do Workers Adjust When Firms Adopt New Technologies?" *ZEW-Centre for European Economic Research Discussion Paper*, no. 21–073.

Goodman-Bacon, Andrew. 2021. "Difference-in-Differences with Variation in Treatment Timing." *Journal of Econometrics*, Themed Issue: Treatment Effect 1, 225 (2): 254–77. https://doi.org/10.1016/j.jeconom.2021.03.014.

Goos, Maarten, and Alan Manning. 2007. "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain." *The Review of Economics and Statistics* 89 (1): 118–33.

Goos, Maarten, Alan Manning, and Anna Salomons. 2014. "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring." *American Economic Review* 104 (8): 2509–26.

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.

Handwerker, Elizabeth. 2015. "Increased Concentration of Occupations, Outsourcing, and Growing Wage Inequality in the United States." BLS Working Paper. Washington, DC: Bureau of Labor Statistics.

Handwerker, Elizabeth Weber, James R Spletzer, and others. 2016. "The Role of Establishments and the Concentration of Occupations in Wage Inequality." *Research in Labor Economics* 43: 167–93.

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. 2022. "Has US Employment Really Polarized? A Critical Reappraisal." *Labour Economics*, 102117.

Juhn, Chinhui, Kevin M Murphy, and Brooks Pierce. 1993. "Wage Inequality and the Rise in Returns to Skill." *Journal of Political Economy* 101 (3): 410–42.

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.

Lazear, Edward P. 2009. "Firm-Specific Human Capital: A Skill-Weights Approach." *Journal of Political Economy* 117 (5): 914–40.

Lindenlaub, Ilse. 2017. "Sorting Multidimensional Types: Theory and Application." *The Review of Economic Studies* 84 (2): 718–89.

Marinescu, Ioana, and Ronald Wolthoff. 2020. "Opening the Black Box of the Matching Function: The Power of Words." *Journal of Labor Economics* 38 (2): 535–68.

Modestino, Alicia Sasser, Daniel Shoag, and Joshua Ballance. 2016. "Downskilling: Changes in Employer Skill Requirements over the Business Cycle." *Labour Economics*, SOLE/EALE conference issue 2015, 41 (August): 333–47. https://doi.org/10.1016/j.labeco.2016.05.010.

———. 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://rodimirorodrigo.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

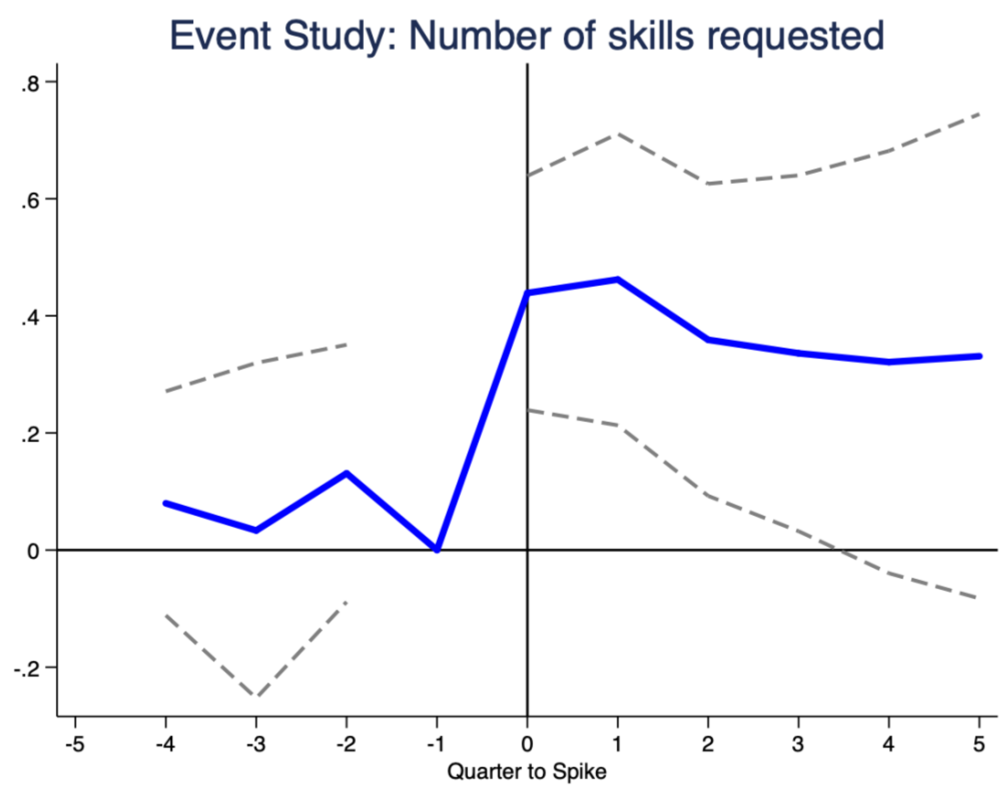

Figure 1. Number of skills requested increases following adoption event.
Note: This figure presents an event study equivalent to Column 1, Panel A, Table 2, reporting the coefficients of quarter dummies for treated firms. The regression is weighted by the number of ads per quarter and it includes fixed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

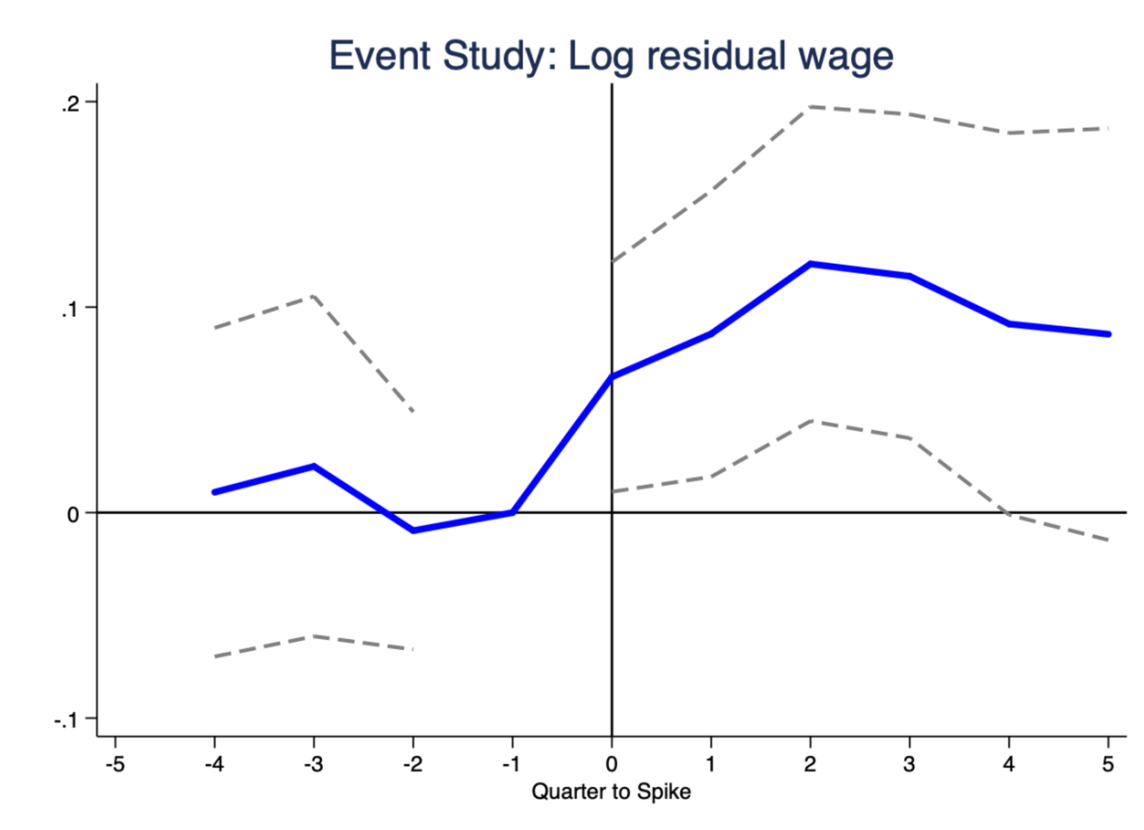

Figure 2. Log residual pay increases following adoption event.
Note: This figure presents an event study equivalent to Column 1, Panel B, Table 2, reporting the coefficients of quarter dummies for treated firms. The regression is weighted by the number of ads per quarter and it includes fixed effects 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 Detailed Skills

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
			Detailed skills					
Skill measure:	All	IT+AI	Other	Cognitive	Social	Soft	Experience	Education
A. All jobs, number of skills								
Post automation event	0.318***	0.041***	0.173***	0.017***	0.048***	0.038***	0.065***	0.017
	(0.071)	(0.008)	(0.059)	(0.004)	(0.010)	(0.010)	(0.019)	(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
Pre-Spike Means	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 automation event	0.222***	0.031***	0.105*	0.010**	0.040***	0.037***	0.041*	0.035
	(0.070)	(0.009)	(0.058)	(0.004)	(0.010)	(0.011)	(0.022)	(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 automation event		0.002*	-0.008***	0.001	0.002***	0.004**		
		(0.001)	(0.002)	(0.000)	(0.001)	(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 detailed skills requested						
Post automation event	0.318***	0.222***	0.398***	0.376***	0.512***	0.153
	(0.071)	(0.070)	(0.087)	(0.105)	(0.091)	(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 automation event	0.087***	0.054**	0.067**	0.067*	0.091***	0.023
	(0.023)	(0.024)	(0.029)	(0.037)	(0.032)	(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

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable	Number of Skills Requested			Log Residual Pay		
Post automation event	0.245**	0.214*	0.219**	0.102***	0.102***	0.094***
	(0.117)	(0.112)	(0.108)	(0.035)	(0.035)	(0.036)
Labor market tightness		0.284	0.111		-0.849	-0.635
		(1.064)	(1.113)		(0.580)	(0.518)
Management jobs		6.653***	6.557***		-0.272	-0.237
		(0.676)	(0.652)		(0.221)	(0.194)
Outsourceable jobs		-7.210***	-7.151***		0.050	-0.013
		(1.793)	(1.719)		(0.314)	(0.328)
Growth Rate of Sales			0.261*			0.069
			(0.155)			(0.058)
Lag CEO change			-0.972			-0.080
			(1.032)			(0.053)
Acquisition			-0.154			0.127***
			(0.189)			(0.037)
3^{rd} order productivity polynomial			✓			✓
Polynomial probability value			0.018			0.030
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 automation event	0.017***	-0.017***	0.007***	-0.008***	0.021***	-0.006***
	(0.002)	(0.002)	(0.003)	(0.002)	(0.002)	(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 automation event	0.018	-0.083**	0.035	-0.107***	0.029	-0.026
	(0.030)	(0.033)	(0.032)	(0.040)	(0.030)	(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

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
			Detailed Skills					
Skill measure:	All	IT+AI	Other	Cognitive	Social	Soft	Experience	Education
Panel A, simple correlation								
Firm FE	2.619***	0.722***	1.338***	0.212***	0.203***	0.144**	1.500***	2.912***
	(0.358)	(0.069)	(0.234)	(0.024)	(0.065)	(0.071)	(0.133)	(0.353)
Standardized coefficient	*0.176*	*0.246*	*0.125*	*0.207*	*0.111*	*0.067*	*0.284*	*0.203*
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**	0.075***	0.557**	0.074***	-0.016	-0.078	0.545***	1.330***
	(0.295)	(0.021)	(0.222)	(0.019)	(0.059)	(0.064)	(0.100)	(0.323)
Standardized coefficient	*0.041*	*0.026*	*0.052*	*0.072*	*-0.009*	*-0.036*	*0.103*	*0.093*
Software share	36.738***	5.601***	20.033***	2.683***	4.105***	4.315***	13.673***	31.272***
	(1.436)	(0.170)	(1.084)	(0.092)	(0.223)	(0.261)	(0.446)	(1.203)
Software share^2	-54.94***	-0.035	-37.63***	-4.21***	-6.26***	-6.79***	-15.34***	-49.85***
	(2.254)	(0.296)	(1.713)	(0.146)	(0.365)	(0.428)	(0.736)	(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

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

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,

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

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

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

Using these, we have[34]

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

Further,

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

[34] α_H increases from $\frac{I-1}{N}$ to $\frac{I}{N}$.

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

Category	Percent of job ads
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 technology-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

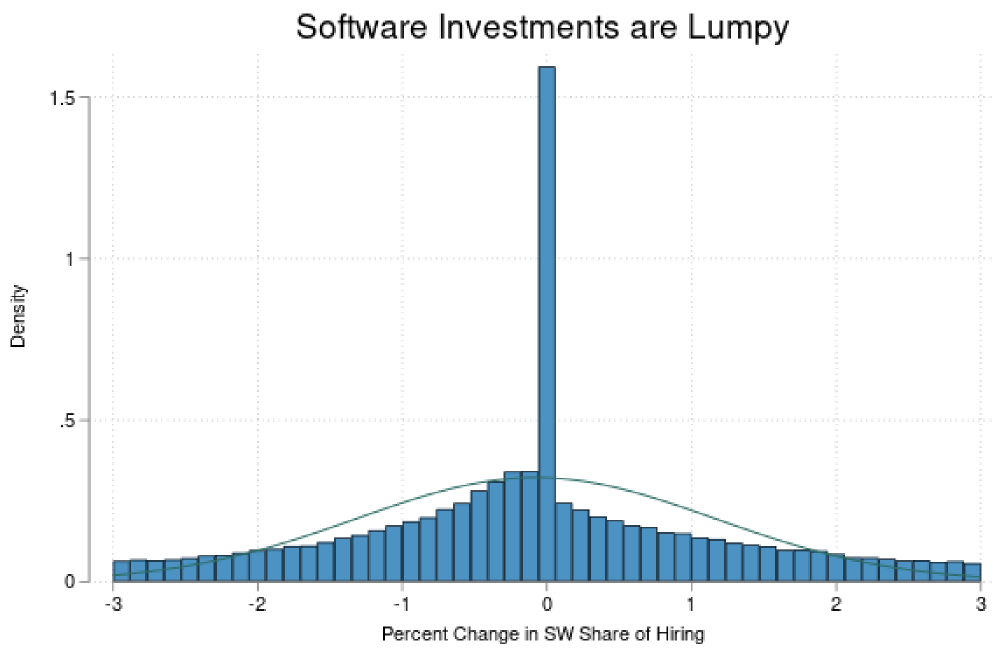

Note: This figure plots changes in software developer share of job advertisements from the average of the previous 4 quarters. The line shows a normal density distribution with the same mean and standard deviation. The distribution is clearly leptokurtic with a peak at zero and fat, “lumpy” tails.

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

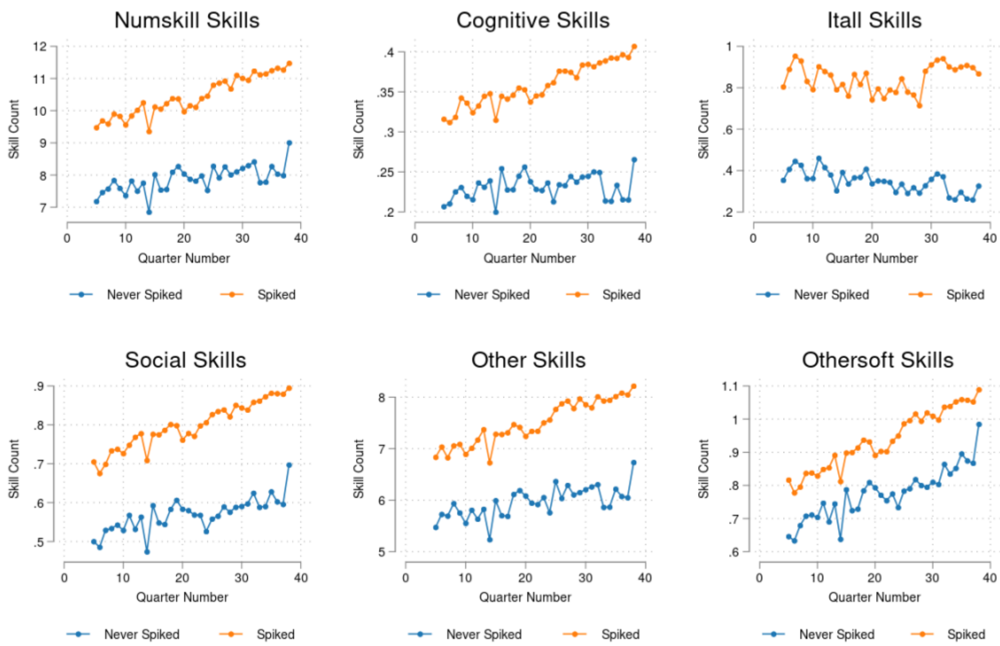

Note: This figure shows raw trends in skill requests for both spiking (orange) and non-spiking (blue) firms over time. Spiking firms have higher levels of skill requests throughout the sample.

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	0.120	0.139
	(0.190)	(0.219)	(0.0960)
Outsourceable Job Share	0.071	0.078	0.056
	(0.183)	(0.209)	(0.101)
Labor Market Tightness	0.795	0.837	0.700
	(0.319)	(0.364)	(0.139)
IT Share	0.095	0.074	0.108
	(0.160)	(0.166)	(0.155)
Residual Wage	0.012	-0.002	0.023
	(0.291)	(0.336)	(0.250)
College Required	0.433	0.416	0.471
	(0.279)	(0.303)	(0.213)
Routine Cognitive	0.298	0.294	0.307
	(0.284)	(0.325)	(0.157)
Routine Manual	0.207	0.224	0.170
	(0.304)	(0.339)	(0.201)
Non-Routine Cognitive	0.444	0.423	0.490
	(0.343)	(0.377)	(0.243)
Non-Routine Manual	0.158	0.177	0.115
	(0.285)	(0.320)	(0.177)
Number of Skills	8.230	7.385	10.062
	(4.895)	(5.210)	(3.484)
Unweighted			
Number of Ads/Quarter	85.380	5.980	164.780
	(86.637)	(1.028)	(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***				0.035***
	(0.001)				(0.001)
Software share		-0.011			0.036***
		(0.009)			(0.008)
Outsourceable jobs			-0.042***		-0.062***
			(0.013)		(0.013)
Management jobs				0.034***	0.056***
				(0.011)	(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*				0.016***
	(0.003)				(0.004)
Log COGS		0.014***			
		(0.002)			
Log Capital			0.008***		0.014***
			(0.002)		(0.002)
Sales Growth				0.017	0.028**
				(0.012)	(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 detailed skills requested						
Post automation event	0.283***	0.318***	0.470***	0.356***	0.318***	0.253***
	(0.076)	(0.071)	(0.087)	(0.068)	(0.071)	(0.074)
Labor market tightness	-0.229	-0.176	-0.197	0.264	-0.176	-0.173
	(0.263)	(0.346)	(0.444)	(0.358)	(0.346)	(0.363)
Management jobs	4.074***	5.488***	5.516***	4.821***	5.488***	4.916***
	(0.329)	(0.386)	(0.576)	(0.574)	(0.386)	(0.334)
Outsourceable jobs	-5.701***	-7.183***	-7.799***	-6.707***	-7.183***	-6.324***
	(0.951)	(1.310)	(1.655)	(1.287)	(1.310)	(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 automation event	0.078***	0.087***	0.072***	0.074***	0.087***	0.253***
	(0.022)	(0.023)	(0.024)	(0.024)	(0.023)	(0.074)
Labor market tightness	-0.133	-0.279*	-0.179	-0.326**	-0.279*	-0.173
	(0.128)	(0.147)	(0.166)	(0.130)	(0.147)	(0.363)
Management jobs	-0.091	0.016	0.198	0.055	0.016	4.916***
	(0.091)	(0.101)	(0.151)	(0.107)	(0.101)	(0.334)
Outsourceable jobs	-0.098	0.026	0.270	-0.105	0.026	-6.324***
	(0.126)	(0.129)	(0.290)	(0.149)	(0.129)	(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

	(1)	(2)	(3)	(4)
	Number of Skills Requested		Log Residual Wage	
Sample	Later-spiking	Full Sample	Later-spiking	Full Sample
Post automation event	0.318***	0.211***	0.087***	0.074***
	(0.071)	(0.068)	(0.023)	(0.020)
Labor market tightness	-0.176	-0.105	-0.279*	-0.149***
	(0.346)	(0.115)	(0.147)	(0.043)
Management jobs	5.488***	3.249***	0.016	-0.159***
	(0.386)	(0.106)	(0.101)	(0.036)
Outsourceable jobs	-7.183***	-2.967***	0.026	0.010
	(1.310)	(0.270)	(0.129)	(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

	(1)	(2)	(3)	(4)
	Number of Skills Requested		Log Residual Wage	
Sample	Full	Non-IT	Full	Non-IT
Post automation event	0.318***	0.358***	0.087***	0.091***
	(0.071)	(0.078)	(0.023)	(0.025)
Labor market tightness	-0.176	-0.094	-0.279*	-0.294*
	(0.346)	(0.363)	(0.147)	(0.151)
Management jobs	5.488***	5.900***	0.016	-0.005
	(0.386)	(0.430)	(0.101)	(0.107)
Outsourceable jobs	-7.183***	-7.132***	0.026	0.025
	(1.310)	(1.392)	(0.129)	(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 automation event	0.330***	0.074***		
	(0.089)	(0.024)		
AI x post automation event	0.304***	0.096***		
	(0.081)	(0.027)		
Non-automation x post automation event			0.040	0.100***
			(0.086)	(0.037)
Automation x post automation event			0.360***	0.086***
			(0.074)	(0.023)
Labor market tightness	-0.177	-0.277*	-0.174	-0.280*
	(0.346)	(0.146)	(0.345)	(0.147)
Management jobs	5.492***	0.012	5.467***	0.017
	(0.387)	(0.102)	(0.386)	(0.101)
Outsourceable jobs	-7.179***	0.026	-7.199***	0.026
	(1.308)	(0.129)	(1.306)	(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 automation event	0.094	0.032
	(0.065)	(0.033)
Labor market tightness	-0.262	-0.129
	(0.321)	(0.127)
Management jobs	5.136***	-0.245
	(0.393)	(0.163)
Outsourceable jobs	-6.039***	-0.299***
	(0.655)	(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
Skill measures	
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.

E. Two-way Fixed Effects

Our main estimation uses the method of Cengiz et al. (2019, Appendix D) to avoid biases that arise in difference-in-differences regressions with heterogeneous treatment effects (de Chaisemartin and D'Haultfœuille 2020; Callaway and Sant'Anna 2020; Goodman-Bacon 2021). We construct balanced panels around each possible spike quarter, including 5 years prior to the spike year and 5 years after, and excluding firms that have previously spiked.

We can explore this issue further. First, Table A10 explores treatment heterogeneity over time by estimating each balanced panel separately for both the restricted (only including firms that spike at some point in the sample) and full samples. While there is noisy year-to-year variation, there is no evidence of trends in the estimation coefficients and the means of the separate coefficients are close to the means from the stacked-panel regressions in Table 2.

This approach has the drawback of using only balanced panels. Firms that are entering or exiting the sample might behave differently. As an alternative, Table A11 shows standard two-way fixed effects regressions where the sample drops observations of treated firms that are more than five quarters after the automation event. The specification for firm i at quarter t with spikes at quarter p is

$$Y_{it} = \delta \cdot \mathbf{1}(t \geq p) + \mu_i + \tau_t + \beta X_{it} + \epsilon_{it}.$$

This approach has the disadvantage of dropping some data, but it limits biases arising from negative cell weights as in de Chaisemartin and D'Haultfœuille (2020). The coefficient estimates are similar to those in Table 2, although a bit smaller with larger standard errors. The last line of the table shows the sum of negative weights.[35] All the weights sum to 1, so the negative weights are quite small, indicating that significant bias from heterogeneous treatment effects is likely small.

Combined, these considerations suggest our results are robust regarding heterogeneous treatment effects.

[35] We use the Stata routine *twowayfeweights* developed by de Chaisemartin, D'Haultfœuille, and Deeb.

Table A10. Estimated Coefficients on the Post-Automation Dummy Variable for Individual Cohorts

Dependent variable	Number of skills requested		Log residual wage	
Sample	Later spikers	Full sample	Later spikers	Full sample
	(1)	(2)	(3)	(4)
Quarter				
2015q2	0.445 (0.328)	0.153 (0.181)	0.012 (0.064)	0.042 (0.063)
2015q3	0.455 (0.143)	0.347 (0.126)	0.089 (0.042)	0.093 (0.048)
2015q4	0.211 (0.175)	0.093 (0.170)	0.108 (0.064)	0.142 (0.063)
2016q1	0.270 (0.202)	0.184 (0.209)	-0.042 (0.040)	0.006 (0.036)
2016q2	0.570 (0.196)	0.264 (0.326)	0.124 (0.058)	0.124 (0.067)
2016q3	0.112 (0.090)	0.062 (0.091)	0.016 (0.046)	-0.009 (0.068)
2016q4	0.640 (0.408)	0.826 (0.507)	0.125 (0.038)	0.028 (0.049)
2017q1	0.511 (0.183)	0.146 (0.168)	0.204 (0.101)	0.152 (0.085)
2017q2	0.060 (0.130)	0.187 (0.129)	0.138 (0.111)	0.091 (0.072)
2017q3	0.303 (0.163)	0.309 (0.149)	-0.060 (0.052)	-0.020 (0.037)
2017q4	0.322 (0.250)	0.525 (0.215)	0.039 (0.050)	0.018 (0.048)
2018q1	0.100 (0.226)	0.109 (0.171)	0.127 (0.058)	0.130 (0.038)
Mean	0.333 (0.065)	0.267 (0.066)	0.073 (0.019)	0.066 (0.017)
Total *N*	102086	1789706	29437	387844

Note: these coefficients are from difference-in-differences regressions for balanced panels (t-5 to t+5) around each spiking year cohort. The unit of observation is firm by quarter. In Columns (1) and (3) firms in the estimation sample spike at some time in our data and only observations are included that have not spiked previously. In Columns (2) and (4) we remove this restriction, consequently broadening our sample size. All regressions include controls for labor market tightness, management job share, the outsourceable job share, and firm fixed effects and standard errors are clustered by firm. 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. The means of the cohort estimates are not significantly different from the stacked difference-in-differences estimates found in Table 2, Column 1.

Table A11. Two-way Fixed Effects OLS Regression

	(1)	(2)	(3)	(4)
	Spiking firms only		Full sample	
VARIABLES	Number of skills requested	Log residual wage	Number of skills requested	Log residual wage
Post automation event	0.201*	0.067**	0.168**	0.048**
	(0.103)	(0.028)	(0.075)	(0.022)
Observations	152,895	42,145	3,401,107	358,388
R-squared	0.854	0.433	0.798	0.484
Sum of negative weights	-0.001	0.000	-0.002	0.000

Note: these coefficients are from standard two-way fixed effects regressions. This sample drops observations of treated firms more than 5 quarters after they have spiked. The unit of observation is firm by quarter. All firms in the first sample (columns 1 and 2) spike at some time during the sample period. 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 bottom line reports the sum of negative weights as described in de Chaisemartin and D'Haultfœuille (2020); the sum of all weights is 1.