Firm Differences:
Skill Sorting and Software

By James Bessen, Erich Denk, and Chen Meng

April 2021

Abstract:

Recent research shows that much recent rise in wage inequality comes from growing differences between firms, especially sorting of skilled workers to high-paying firms. This paper explores the role of proprietary software in these changes. Using job ad data, we find that proprietary software is strongly associated with firm wage fixed effects and also with firm skills. Software accounts for half or more of skill sorting across firms. Moreover, both skill sorting and firm wage effects are greater for larger firms. The huge growth in proprietary software helps explain the growth in skill sorting that increases wage inequality.

JEL codes: J31, J24, O3

Keywords: skill sorting, income inequality, information technology, firm wage differentials

Thanks for comments from Maarten Goos, Anna Salomons, Emilie Rademakers, Bledi Taska and participants at the TPRI Brown Bag.

Recent research finds that differences between firms play a major role in accounting for rising wage inequality, especially increased sorting of workers by skill across firms.[1] At the same time, differences between firms are affecting industry structure. Productivity differences between firms have grown, large firms have increased their market shares, and they have become more persistently dominant.[2] Are these two trends related and, if so, how?

One commonality is the role of technology. A large literature associates information technology with greater demand for skilled workers (Acemoglu 2002); some papers also relate information technology to plant or firm wages (Doms, Dunne, and Troske 1997; Barth et al. 2020). On the other side, it is well-established that information technology boosts firm productivity and can help explain growing dispersion in productivity across firms (Dunne et al. 2004; Brynjolfsson and Hitt 2003; Bloom, Sadun, and Van Reenen 2012). Moreover, this technology appears to be particularly beneficial to the largest firms. Evidence suggests that increasing use of information technology is responsible for growing industry concentration and is associated with higher labor productivity, markups, and persistent dominance for top firms.[3] If this technology widens the differences in returns to skill between firms, it might very well widen pay differences between firms as well.

This paper explores the extent to which rising between-firm wage differences can be accounted for by firm-specific information technology. We examine both differences in firm fixed effects and differences in sorting of skilled workers to high-paying firms and we look at whether these differences are greater for large firms.

---

[1] Card, Heining, and Kline 2013; Barth, Davis, and Freeman 2018; Song et al. 2019; Lachowska et al. 2020.

[2] (Andrews, Criscuolo, and Gal 2016; Decker et al. 2020; Autor et al. 2020; Bessen et al. 2020)

[3] (Calligaris, Criscuolo, and Marcolin 2018; Crouzet and Eberly 2018; Bessen 2020; Bessen et al. 2020).

A link between proprietary software and wage sorting is potentially very important given that firms have made a huge shift of investment into this type of technology since the turn of the century. In 2019 in the US, private investment in proprietary software—software that firms develop on their own or by contracting others—grew to $234 billion.[4] This figure is about as much as firms' net investment in equipment and does not include associated investments in hardware and human capital. Examples of these systems include Walmart's logistics and inventory management systems or large banks' credit card systems. If these large investments increase the relative demand for more highly skilled workers, they might explain much of the increase in between-firm inequality.

To analyze between-firm wage inequality, we adapt the "canonical" model of skill-biased technical change to cover individual firms in monopsonistic labor markets. The model shows that technology can increase differences in firm pay and skill levels when technology complements skilled workers, when highly skilled workers complement less skilled workers, and when firms differ in their adoption of the technology. To test this model, our main hypothesis is that firm investments in systems using proprietary software increase both firm wages and firm demand for skills, thus increasing sorting and between-firm wage differences. Measuring sorting as the correlations between skills and firm wage fixed effects, proprietary software should account for much of those correlations. Second, noting that firm investments in proprietary software are highly heterogeneous—in particular, firms with more than 1,000 employees are much more likely to hire software developers—we hypothesize that the association between proprietary software and both firm fixed effects

---

[4] Data from BEA, National Income and Product Accounts, Table 9.4U. Software investment and prices, July 31, 2020 revision. This figure excludes purchases of pre-packaged software and excludes development of software for sale.

and skills should be stronger for larger firms.[5] Presumably these firms earn greater returns on software.

The challenge in estimating firm wage effects is to distinguish differences in pay that are attributable strictly to the firms themselves and differences that arise because firms have workers with different qualities, some of which are not observed. Most of the recent studies use the "AKM" method to identify worker fixed effects and firm fixed effects in large databases of linked employee-employer data (Abowd, Kramarz, and Margolis 1999). Presumably, sorting occurs because firms have different labor demands; some firms are willing to pay more and those firms also tend to select more highly skilled workers.

This paper takes a different approach, measuring differences in employer demand directly by using comprehensive data on online help-wanted ads. We measure firm fixed effects from advertised salaries. The job advertisements also provide rich information on the skills that employers demand. While advertised salaries differ from salaries actually paid, on average they are quite similar and they nevertheless provide a clear metric of firms' willingness to pay that is independent of individual worker characteristics.[6] A correlation between AKM firm wage fixed effects and individual skills should be reflected in firm willingness to pay and skills requested in job ads.

---

[5] Using the Current Population Survey ASEC files for 2000-17, software developers comprise 3.7% of the workforce for firms with over 1,000 employees, but only 1.3% of the workforce for firms with fewer than 100 employees.

[6] Comparing, the median salary offered (mean if a range was given) for full-time jobs, excluding interns in Burning Glass was 2% more than the median annual earnings for fulltime/full-year workers in the Current Population Survey. Our fixed effects do reflect differences in how much firms are willing to pay for unobserved worker characteristics. That is, some firms will advertise higher salaries knowing that they will be more selective based on these unmeasured characteristics. In this way our fixed effects differ from AKM fixed effects, but they still cleanly identify a basic difference in between-firm pay. Loosely speaking, our fixed effect is equivalent to the AKM fixed effect plus what Song et al. (2019) call "segregation."

We estimate firm wage effects by regressing the log of posted salaries against firm fixed effects with a variety of controls for job characteristics. We derive a measure of firm software investment using the help-wanted ads for each firm (see below). We find a strong positive relationship between the firm fixed effects and software investment in industries where software is not a major part of the firm's product. This relationship holds even after controlling for outsourcing and other sources of rents. Also, we find that the association is substantially stronger for firms with more than 1,000 employees.

The help-wanted ads also provide a rich set of skill characteristics required for different jobs. We identify six skill measures: the total number of specific skills requested, IT skills (for non-IT jobs), "artificial intelligence"/data science skills, "soft" skills, experience required, and education required. Consistent with the prior literature, we find significant correlations between each of these skills and the firm fixed effects, indicating sorting. But we also find that much of the correlation between skills and firm wage fixed effects is accounted for by firm software investment. Moreover, the coefficient on software is substantially larger for large firms. In other words, the sorting of skilled workers to high-paying firms is significantly an artifact of the role of software-enabled systems in raising labor demand for skilled workers, especially at large firms.

While a large literature studies the association between technology and worker skills, relatively few papers look at the links between technology, firm wage effects, and skills independently of individual worker characteristics. Abowd et al. (2007) look at the association between skill measures and firm fixed effects and separately at the association between skill measures and firm computer and software use. Barth et al. (2020) find an association between firm wage fixed effects and software investment and also an association between individual wage fixed effects and software. But neither study shows the role of technology in accounting for the correlation between skills and firm wage fixed effects, that is, the degree of sorting.

This paper makes several contributions to the literature. First, we develop a model to explore possible impacts of technology on firm fixed effects and sorting. Second, we develop a new data source to study firm wage effects. Third, we find that both fixed firm wage effects and sorting by skills are substantially associated with firms' investments in proprietary software. This association accounts for much of worker sorting by skills across firms. Fourth, we find that these effects are stronger for large firms. Rising wage inequality appears to be related to the large shift in corporate investment in proprietary software, especially by large firms.

Technology and Worker Sorting

A substantial literature shows that there are persistent productivity differences between firms that arise for various reasons (Syverson 2011). And productivity differences between firms give rise to wage differences between firms in models of rent-sharing and monopsony (Card et al. 2018). But these productivity differences do not necessarily give rise to increased sorting of workers across firms, as can be seen from a simple model.

Much of the analysis of the effect of technology on wage inequality has been studied with a "canonical" model of skill-biased technical change (Acemoglu 2002). Acemoglu uses a constant elasticity of substitution (CES) aggregate production function in two factors, low skill labor, L, and high skill labor, H. Skill-biased technical change increases the productivity of H relative to the productivity of L. Acemoglu shows that as long as the elasticity of substitution between these two types of labor is greater than one, skill-biased technical change will increase the demand for higher skill workers and their relative wages will rise.

But this is a model for the aggregate economy. The implicit assumption is that all firms adopt the new technology and, indeed, the model is often motivated by pointing to the dramatic drop in the prices of computing that made information technology widely accessible. Yet investments in information technology appear to be highly disparate across firms, accounting for productivity differences (Syverson 2011, 3.3).

We can recast this model to be one for individual firms operating in monopsonistic labor markets. Let each firm produce output according to a constant returns CES production function

$$F(L,H) = aL\left(1+\left(\frac{bH}{L}\right)^{\rho}\right)^{\frac{1}{\rho}}, \qquad \rho < 1$$

where L is low skill labor and H is high skill labor. Two parameters capture aspects of technology/productivity, a, which represents Hicks neutral productivity, and b, the relative productivity of high skill workers (factor augmentation). We assume $b > 1$.

Each firm faces a rising labor supply curve. We assume a constant elasticity of supply, ψ, such that $w_L(L) = w_L^0 \cdot L^{\frac{1}{\psi}}$ is the wage the firm pays for low skill labor and $w_H(H) = w_H^0 \cdot H^{\frac{1}{\psi}}$ is the wage the firm pays for high skill labor, where the constant $w_H^0 > w_L^0$.[7]

The profit of the firm is $F(L,H) - w_L L - w_H H$. Defining $h \equiv H/L$ and solving the first order conditions,

(1)

[7] One could allow for differentiated elasticities here, though it greatly complicates the model and departs from the literature cited here.

$$\ln \hat{h} = \frac{1}{1-\rho+\frac{1}{\psi}}\left(\rho \ln b - \ln \frac{w_H^0}{w_L^0}\right)$$

(2)

$$\ln \hat{L} = \psi \left[\ln \frac{\psi a}{1+\psi} + \frac{1-\rho}{\rho} \ln\left(1+\left(b\hat{h}\right)^{\rho}\right)\right].$$

From which it follows

(3)

$$\hat{w}_L = w_L^0 \, \hat{L}^{\frac{1}{\psi}}, \quad \hat{w}_H = w_H^0\left(\hat{h}\hat{L}\right)^{\frac{1}{\psi}}, \quad \omega \equiv \frac{\hat{w}_H}{\hat{w}_L} = \frac{w_H^0}{w_L^0} \cdot \hat{h}^{\frac{1}{\psi}}.$$

Finally, we can define the average wage for the firm as,

(4)

$$\bar{w} \equiv \frac{w_L L + w_H H}{L+H} = \frac{a\psi}{1+\psi} \frac{(1+(bh)^{\rho})^{\frac{1}{\rho}}}{1+h}.$$

This model is consistent with the main results of the canonical model of skill-biased technical change. Assuming that the elasticity of substitution $(\sigma = 1/1-\rho)$ between factors is greater than one corresponds to $0 < \rho < 1$. Then an increase in the skill bias, b, increases the relative demand for high skill labor (1), increases the relative wage of high skill labor (3), and increases the average wage (4).[8]

Sorting occurs when firms with higher average wages tend to also employ relatively more high skill workers, that is, when $\bar{w}$ is correlated with $\hat{h}$. The model reveals several aspects of the link between technology and sorting. First, it is straightforward to show that increased productivity dispersion, measured as increased dispersion of a, does not increase

---

[8] It is straightforward to show that the average wage increases if $b > 1$ and $0 < \rho < 1$.

sorting in this model.[9] While difference in a will increase differences in the average wage (4), the relative employment of high skill to low skill labor does not change with a as in (1). So, rising productivity dispersion is not sufficient to generate increased sorting.

Second, if firms have disparate productivity levels, a, but are otherwise similar, then uniform increases in the bias of technology, b, do not increase sorting either. Looking at (1), increasing b will increase the relative demand for high skill workers, but it will do so equally for high productivity firms and low productivity firms. So, skill-biased technical change does not necessarily increase sorting either.

However, if firms realize different levels of skill-biased technical change, the result is different. Firms with higher levels of b will have both greater relative employment of high skill workers and greater average wages. That is, *disparate* increases in b do generate greater sorting by skill. While the canonical model of skill-biased technical change implicitly assumes that all firms equally access new technology, our model suggests that differences in technology adoption might be critical for increased sorting.

Why might some firms adopt new technology to a greater degree than others? One reason might be limited access to the technology because of patents or limited access to employees who have the knowledge to develop and implement new systems. However, this explanation by itself might be hard to reconcile with the sustained growth in sorting over decades. One would expect patents to expire or to be "invented around" and to expect workers to acquire the necessary knowledge over time. Different adoption rates might also arise because firms have different returns to the technology. For example, with economies of

[9] Eeckhout and Kircher (2018) provide models where firms select a single skill and where a dispersion in firm productivity does create sorting.

scale, large firms might find it much more economical to adopt the new technology. Also, investments in technology might represent endogenous sunk costs as in the natural oligopoly models of Shaked and Sutton (Sutton 1991; 2001; Shaked and Sutton 1982; 1983; 1987). In these models, firms invest more or less in technology depending on their position in vertically differentiated markets. This model appears to apply to at least some of the industries that make large investments in proprietary software-enabled systems (Ellickson 2006; 2007).

Our model explains how differences in technology adoption may be related to increased sorting and, generally, how more unequal industry structures are related to greater wage inequality between firms. Assuming that firms employ proprietary information technology differently and that these investments complement skilled workers, the model predicts that these investments should be associated with both higher firm wage fixed effects and higher levels of skills. And assuming that large firms receive greater returns on this technology, these associations should be stronger for large firms. These are the hypotheses we test. The model also indicates other parameters that might affect sorting, for instance, differences in the elasticity of labor supply. Exploring these factors is beyond the scope of this paper.

Data

Our main data are help-wanted advertisements collected by Burning Glass Technologies. Burning Glass is a software company that scrapes, deduplicates, and cleans the near-universe of online job advertisements. A previous analysis of the dataset showed that this it accounts for 60-70% of all job openings and 80-90% of openings requiring a bachelor's degree or more (Carnevale, Jayasundera, and Repnikov 2014). The data include the advertised salary, firm name, industry, occupation, required education and experience, requested skills, and geographic location of the job. Our sample spans from January 2010 to December 2018. 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. We estimate firm wage fixed effects for those ads listing salary information. Using log salary as the dependent variable (or the mean of the salary range limits if a range is listed), we calculated 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 appendix). The R-squared for this regression is .689. The regression excludes software development occupations to avoid spurious correlation with our key independent variable. This gave us estimates of firm fixed effects for 152,004 firms that posted 78,328,283 help wanted ads, excluding ads for information technology occupations.[10] We further matched a subset of the firms advertising in Burning Glass to Compustat to obtain additional variables.[11] Compustat provides measures of firm rents and we use Compustat to identify firms with more than 1,000 employees, “large” firms. While this approach misses large private firms, that simply means our estimates of the large firm effect will tend to be understated.

[10] These are 75% of the total ads, the remainder being IT jobs and ads with missing data.

[11] 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.

We measure software investment for each firm each year as the share of help-wanted ads in computer and mathematical occupations (SOC 15), excluding administrative and support jobs. A number of papers measure firm investment in developing software as the investment in software developers or, relative to the size of the firm, as the share of software developers in the workforce (Tambe and Hitt 2012; Bessen 2020; Bessen and Righi 2019). This does not capture that portion of proprietary software that firms contract for with third parties nor does it capture complementary investments in human and organizational capital. However, these are likely correlated with the investment in own-developed software.[12] We use the software share of help-wanted ads rather than the share of employees, but these, too, are highly correlated.[13]

The Burning Glass data tabulate specific skills requested in the job ads. In addition to experience and education required, we measure skill with the total number of specific skills requested and the number of skills required in the following categories: IT skills, AI skills, and soft skills. The definitions of these categories are provided in the Appendix.

Modestino et al. (2019) find that the skills listed in job ads change with the business cycle. To control for these effects, we add a measure of labor market tightness by state-month. We follow Moscarini and Postel-Vinay (2016) in defining labor market tightness as the ratio between Job Openings and Labor Turnover Survey (JOLTS) statewide openings for the non-farm sector and the state unemployment rate.

Summary statistics for the main variables are in the Appendix.

---

[12] The employment share of IT workers is highly correlated with BEA measures of software investment that do include contracted software. See Bessen (2020, 537–38).

[13] Comparing measures for NAICS 3-digit industries between employment totals by industry from the American Community Survey of the Census and our measure derived from Burning Glass data, these two measures have a correlation of .945.

Findings

Firm Fixed Effects

Table 1 and Figure 1 explore the basic association between firm wage fixed effects and the software share of firms' workforces. Figure 1 shows a binned scatterplot between firm fixed effects and the software share of the workforce. The relationship appears distinctly positive and concave. Column 1 of Table 1 shows a linear specification and Column 2 adds a quadratic term. Both regressions have highly significant coefficients, with the second providing distinctly better fit.

One possible confounding factor is the extent to which firms outsource certain jobs. If low wage jobs are outsourced, firm mean pay will be higher, all else equal. We identify "outsourceable" jobs, those lower wage occupations that are subject to outsourcing.[14] Column 3 adds a control for the share of each firm's workforce that is in outsourceable occupations. The idea is that firms with a larger share of outsourceable jobs are less likely to have outsourced jobs, all else equal, and hence should have a lower firm wage. We see this to be true and also that the addition of this control has only a modest effect on the software share coefficient relative to the baseline specification.

Another potential concern is that the effect might not be specific to software development jobs. Perhaps technology is raising the demand for all white-collar jobs. To test this, Column 4 includes a control for the managerial share of the workforce (SOC = 11) for a firm. The coefficient on the manager share is small but statistically significant. While the

[14] 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.

manager share is correlated with the software share of hiring, the manager share has little effect on the software share coefficient, suggesting that software is where the real action is.

Column 5 explores differences in these associations by firm size. Our hypothesis is that larger firms get greater returns from their investments in proprietary software systems and hence the association between software share and wage effects should be greater for large firms. The regression confirms this hypothesis and a t-test finds the difference in coefficients to be highly significant (P = .000).

Another concern is that these estimates might be confounded by other sources of rents. That is, rent sharing from other sources might increase firm wages and perhaps software investment is correlated with these other rents. Table 2 adds other measures of firm rents using the subset of our help-wanted database that is matched to Compustat. The first three columns show regressions of firm fixed effects on three measures of rents averaged over the years 2010-18: the log of revenue per employee, net operating margin (earnings before taxes and depreciation over revenues), and Tobin's Q. All coefficients are positive, the first and third significant. The highly significant coefficient in Column 1, .081, corresponds well with similar wage elasticities of rents after controlling for individual worker characteristics in Card et al. (2018).[15] Column 4 shows the regression using just the average software share of help wanted ads for these firms and columns 5-7 add the three rent measures. Note that we should expect both sets of coefficients to be attenuated because measures of firm rents include rents from software investments. We find only modest

[15] For example, Card et al. regress firm wage fixed effects from Portuguese data against log value-added per worker and find estimates of .107 with control for city and industry. Without control the estimate is .137.

attenuation of the coefficients for the software share of the workforce, suggesting that this effect does not come from a spurious correlation with other sources of rents.

Skill Sorting

Table 3 explores the correlation of various skill measures with firm fixed effects. The top panel shows simple regressions of each skill measure on firm fixed effects. All have highly significant correlations, implying the presence of sorting. The second panel adds controls for the software share of the workforce interacted with firm size. The software share variables are significantly correlated with all skill measures and they are substantially larger for large firms. T-tests of the difference between the large firm and small firm software coefficients are all highly significant except for AI skills. The third panel, C, includes the software share with a quadratic term. The associated coefficients are all highly significant and the regression R-squares are larger than in the previous panels.

Note that these regressions exclude IT occupations. In fact, the association seems to hold for most occupational groups. Interacting the software share variable with an indicator for each occupational group (see Appendix) shows all groups have positive coefficients for at least some of the regressions except for transportation occupations; managerial and professional jobs typically have larger coefficients. The implication is that skills complement software in most occupations.

The sorting of worker skills across firms is captured by the correlation between skills and firm wage fixed effects. Panel D presents standardized (beta) coefficients for the firm fixed effect variables in each of the first three panels of the table. These are regression coefficients where each variable has been divided by its standard deviation. For a regression with a single variable, the standardized coefficient is equal to Pearson's correlation coefficient. Ignoring the adjustment for labor market tightness, the first row of panel D shows that the correlation coefficients on firm fixed effects are all economically significant except for AI skills.

How much of these correlation coefficients can be accounted for by proprietary software? Adding software variables in Panels B and C reduces the standardized coefficient on firm fixed effects; this reduction represents the contribution of software to the total correlation. Using the reduction of the coefficient between Panel A and Panel C, most of the correlation between firm fixed effects and IT and soft skill requests is accounted for by software; software accounts for about half of the sorting correlation for the other skill variables. Thus, firm investments in proprietary software account for a substantial portion of the sorting of worker skills to higher paying firms.

Conclusion

Recent research finds that between-firm differences account for much of the rise in wage inequality, especially increased sorting of skilled workers to high-paying firms. This paper finds that firm investments in proprietary software as measured by their demand for software developers are substantially related to between-firm pay differences, both differences in the pay that firms offer for comparable jobs and in the skills of the workers they hire. Firms differ dramatically in the extent to which they invest in proprietary software and related systems and business models, presumably because some firms earn greater returns on the technology than others. If the technology complements skills, then differences in returns imply differences in the marginal productivity of workers by skill across firms, hence differences in pay and skills by firm. This is our finding.

The role of proprietary software and related investments is important for two reasons. First, these investments have grown dramatically over the last 20 years. This suggests that the rise in worker sorting is related to the shift in investment.

Second, this analysis provides a reason why technology is intensifying wage differences between firms. The skill-biased technical change hypothesis assumes that all workers have uniform access to the skill-biased technology. The argument was that low-cost computing made new technology accessible to firms of all types and sizes. But firm investments in proprietary software are highly heterogeneous. In particular, large firms invest much more relative to their sizes, presumably because they earn greater returns. We find that the links between proprietary software and skill sorting are stronger for large firms. This suggests that growing differences in how firms use technology, especially large firms, are closely related to the growing importance of firm differences in understanding wage inequality. The rise in skill sorting is in this way related to rising industry concentration and rising persistence of dominant firms. Firm differences in technology substantially affect firm differences in pay.

References

Abowd, John M, John Haltiwanger, Julia Lane, Kevin L McKinney, and Kristin Sandusky. 2007. "Technology and the Demand for Skill: An Analysis of within and between Firm Differences." National Bureau of Economic Research.

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.

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."

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen. 2020. "The Fall of the Labor Share and the Rise of Superstar Firms." *The Quarterly Journal of Economics* 135 (2): 645–709.

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.

Bessen, James. 2020. "Information Technology and Industry Concentration." *Journal of Law and Economics* 63 (3): 531–55.

Bessen, James, Erich Denk, Joowon Kim, and Cesare Righi. 2020. "Declining Industrial Disruption." *Boston Univ. School of Law, Law and Economics Research Paper*. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3682745.

Bessen, James, and Cesare Righi. 2019. "Information Technology and Firm Employment." *Boston Univ. School of Law, Law and Economics Research Paper*, no. 19–6. https://ssrn.com/abstract=3371016.

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen. 2012. "Americans Do IT Better: US Multinationals and the Productivity Miracle." *American Economic Review* 102 (1): 167–201.

Brynjolfsson, Erik, and Lorin M Hitt. 2003. "Computing Productivity: Firm-Level Evidence." *Review of Economics and Statistics* 85 (4): 793–808.

Calligaris, Sara, Chiara Criscuolo, and Luca Marcolin. 2018. "Mark-Ups in the Digital Era." OECD.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. 2018. "Firms and Labor Market Inequality: Evidence and Some Theory." *Journal of Labor Economics* 36 (S1): S13–70.

Card, David, Jörg Heining, and Patrick Kline. 2013. "Workplace Heterogeneity and the Rise of West German Wage Inequality." *The Quarterly Journal of Economics* 128 (3): 967–1015.

Carnevale, Anthony P, Tamara Jayasundera, and Dmitri Repnikov. 2014. "Understanding Online Job Ads Data." *Georgetown University, Center on Education and the Workforce, Technical Report (April).*

Crouzet, Nicolas, and Janice Eberly. 2018. "Intangibles, Investment, and Efficiency." In *AEA Papers and Proceedings*, 108:426–31.
Decker, Ryan A, John C Haltiwanger, Ron S Jarmin, and Javier Miranda. 2020. "Changing Business Dynamism and Productivity: Shocks vs. Responsiveness." *American Economic Review* 110 (12): 3952–2990.
Doms, Mark, Timothy Dunne, and Kenneth R Troske. 1997. "Workers, Wages, and Technology." *The Quarterly Journal of Economics* 112 (1): 253–90.
Dunne, Timothy, Lucia Foster, John Haltiwanger, and Kenneth R Troske. 2004. "Wage and Productivity Dispersion in United States Manufacturing: The Role of Computer Investment." *Journal of Labor Economics* 22 (2): 397–429.
Eeckhout, Jan, and Philipp Kircher. 2018. "Assortative Matching with Large Firms." *Econometrica* 86 (1): 85–132.
Ellickson, Paul B. 2006. "Quality Competition in Retailing: A Structural Analysis." *International Journal of Industrial Organization* 24 (3): 521–40.
———. 2007. "Does Sutton Apply to Supermarkets?" *The RAND Journal of Economics* 38 (1): 43–59.
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.
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.
Modestino, Alicia Sasser, Daniel Shoag, and Joshua Ballance. 2019. "Upskilling: Do Employers Demand Greater Skill When Workers Are Plentiful?" *Review of Economics and Statistics*, 1–46.
Moscarini, Giuseppe, and Fabien Postel-Vinay. 2016. "Wage Posting and Business Cycles: A Quantitative Exploration." *Review of Economic Dynamics*, Special Issue in Honor of Dale Mortensen, 19 (January): 135–60. https://doi.org/10.1016/j.red.2015.11.001.
Shaked, Avner, and John Sutton. 1982. "Relaxing Price Competition through Product Differentiation." *The Review of Economic Studies*, 3–13.
———. 1983. "Natural Oligopolies." *Econometrica: Journal of the Econometric Society*, 1469–83.
———. 1987. "Product Differentiation and Industrial Structure." *The Journal of Industrial Economics*, 131–46.
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.
Sutton, John. 1991. *Sunk Costs and Market Structure: Price Competition, Advertising, and the Evolution of Concentration*. MIT press.
———. 2001. *Technology and Market Structure: Theory and History*. MIT press.
Syverson, Chad. 2011. "What Determines Productivity?" *Journal of Economic Literature* 49 (2): 326–65.
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.

Tables and Figures

Table 1. Firm Wage Fixed Effects and Software Share of Hiring
Dependent variable: Firm fixed effect

	(1)	(2)	(3)	(4)	(5)
Software Share	0.560***	0.901***	0.531***	0.563***	
	(0.008)	(0.020)	(0.008)	(0.008)	
Software Share squared		-0.716***			
		(0.038)			
Outsourceable Share			-0.181***		-0.181***
			(0.006)		(0.006)
Manager Share				0.075***	
				(0.004)	
Small Firm x SW Share					0.525***
					(0.008)
Big Firm x SW Share					1.011***
					(0.051)
Observations	144,419	144,419	144,419	144,419	144,419
R-squared	0.034	0.036	0.040	0.036	0.041

Sample excludes IT-producing industries. Observations are firms. Weighted to match the distribution of CPS occupations; standard errors clustered by firm in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 2. Firm Wage Fixed Effects and Firm Rents
Dependent variable: Firm fixed effect

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Log labor productivity	0.081***				0.045**		
	(0.021)				(0.017)		
Operating margins		0.001				0.001	
		(0.001)				(0.001)	
Mean Tobin's Q			0.029*				0.015
			(0.016)				(0.010)
Mean SW share of workforce				0.687***	0.599***	0.687***	0.696***
				(0.184)	(0.182)	(0.184)	(0.201)
Observations	1,630	1,680	832	1,680	1,630	1,680	832
R-squared	0.050	0.000	0.018	0.114	0.127	0.114	0.148

Sample is Compustat firms matched to Burning Glass; variables are firm averages over 2000-2018.
Standard errors clustered by firm in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Table 3. Skills, Firm Fixed Effects, and Software Share for Non-IT Jobs

Dependent variable:	(1) No. skills	(2) IT skills	(3) AI skills	(4) Soft skills	(5) Experience	(6) Education
A. Firm Fixed Effects						
Firm fixed effects	3.672***	0.205***	0.002***	0.099***	1.626***	2.129***
	(0.329)	(0.016)	(0.000)	(0.035)	(0.141)	(0.149)
Observations	78,328,283	78,328,283	78,328,283	78,328,283	38,881,237	49,323,583
R-squared	0.018	0.019	0.001	0.003	0.028	0.163
B. Firm Fixed Effects and Software Share by Firm Size						
Firm fixed effects	2.308***	0.065***	0.001***	0.040	0.876***	1.474***
	(0.300)	(0.009)	(0.000)	(0.034)	(0.127)	(0.135)
Small firm x SW share	8.140***	0.984***	0.006**	0.300***	5.113***	4.052***
	(0.459)	(0.012)	(0.003)	(0.025)	(0.163)	(0.301)
Large firm x SW share	14.108***	1.236***	0.007***	0.696***	6.620***	6.695***
	(1.118)	(0.044)	(0.002)	(0.070)	(0.451)	(0.400)
R-squared	0.053	0.134	0.002	0.018	0.105	0.182
T-test of size coefficients						
Prob[small = large]	0.000	0.000	0.981	0.000	0.000	0.000
C. Firm Fixed Effects and Software Share Quadratic						
Firm fixed effects	1.643***	0.044***	0.001***	0.008	0.705***	1.131***
	(0.279)	(0.008)	(0.000)	(0.033)	(0.127)	(0.138)
SW share	33.122***	1.837***	0.008***	1.584***	11.638***	17.021***
	(1.350)	(0.044)	(0.002)	(0.089)	(0.619)	(0.698)
SW share^2	-43.897***	-1.470***	-0.003**	-2.206***	-11.370***	-23.656***
	(1.866)	(0.062)	(0.001)	(0.128)	(0.948)	(1.053)
R-squared	0.081	0.144	0.002	0.032	0.119	0.199
D. Standardized Coefficients, firm fixed effect						
Panel A	0.134	0.134	0.020	0.053	0.158	0.114
Panel B	0.084	0.043	0.010	0.021	0.085	0.079
Panel C	0.060	0.029	0.009	0.004	0.069	0.061
SW share of sorting (1-Panel C / Panel A)	55%	78%	55%	92%	56%	46%

Sample excludes IT occupations. All regressions include controls for labor market tightness and are weighted to match the distribution of CPS occupations; standard errors clustered by firm in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Figure 1. Binned scatterplot of Firm Wage Fixed Effects and Software Share of Workforce

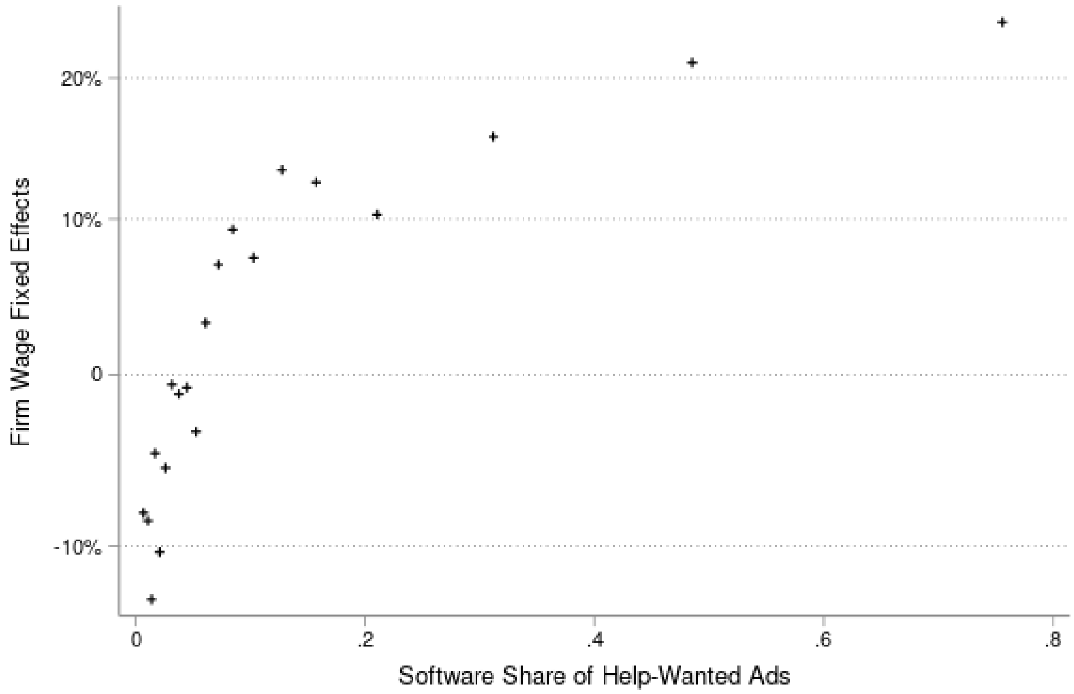

Appendix

Summary Statistics

Table A1.

Variable	Mean	Std. Dev.
No. skills	9.147	7.064
IT skills	0.173	0.379
AI skills	0.0004	0.020
Soft skills	0.617	0.486
Experience	3.211	2.709
Education	12.763	4.791
Software share of workforce	0.077	0.124
Firm fixed effect	0.015	0.266

Skill measures

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

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

* These skills also have synonyms, which were also flagged. For full list of synonyms, please refer to Table 13 in Khaouja et al 2019. In addition, the following commonly requested Burning Glass skills were also flagged as soft skills: Communication Skills, Teamwork / Collaboration, Planning, Detail-Oriented, Building Effective Relationships, Verbal / Oral Communication, Energetic, Positive Disposition, Listening, Team Building, Creative Problem Solving, Self-Motivation, Overcoming Obstacles, Multi-Tasking, People Management, Negotiation Skills, Thought Leadership, Team Management

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: (Based on classification by Burning Glass). There are 1,687 unique skills, which can be sorted into broader categories, listed in the table below. Within the category "Microsoft Development Tools" is the Microsoft Office suite, which we omit as an IT skill.

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

Skill sorting by occupation

Table A2.

Dependent variable	(1) No. skills	(2) IT skills	(3) AI skills	(4) Soft skills	(5) Experience	(6) Education
Firm fixed effects	2.385***	0.072***	0.001***	0.050	0.810***	1.453***
	(0.299)	(0.009)	(0.000)	(0.034)	(0.126)	(0.127)
SW share x						
Manager	14.407***	1.106***	0.010***	0.637***	10.203***	7.167***
	(0.546)	(0.020)	(0.002)	(0.026)	(0.296)	(0.289)
Professional	9.384***	1.252***	0.007***	0.244***	4.572***	4.784***
	(0.391)	(0.014)	(0.002)	(0.022)	(0.180)	(0.242)
Health care	5.371***	0.268***	0.004	-0.346***	-1.626***	7.909***
	(1.613)	(0.047)	(0.003)	(0.095)	(0.481)	(0.869)
Business support	-11.758***	0.289***	0.010***	-0.073	1.474***	-9.323***
	(1.462)	(0.056)	(0.004)	(0.104)	(0.511)	(1.359)
Sales	11.655***	0.698***	0.006*	0.724***	2.508***	3.304***
	(0.971)	(0.053)	(0.004)	(0.057)	(0.738)	(0.614)
Administrative	7.192***	0.975***	0.002**	0.519***	-0.364	-0.675**
	(0.628)	(0.034)	(0.001)	(0.050)	(0.463)	(0.304)
Construction/prod.	2.577***	0.807***	0.002**	-0.277***	3.536***	-4.513***
	(0.613)	(0.034)	(0.001)	(0.042)	(0.224)	(0.393)
Transportation	-22.475***	0.075	0.004*	-1.541***	-1.697***	-12.165***
	(2.568)	(0.057)	(0.002)	(0.188)	(0.558)	(1.310)
Other	-0.612	0.893***	0.003**	0.243***	6.688***	4.702***
	(1.134)	(0.066)	(0.001)	(0.048)	(0.921)	(0.645)
Observations	85,727,849	85,727,849	85,727,849	85,727,849	44,237,189	54,577,715
R-squared	0.063	0.188	0.003	0.021	0.152	0.193

All regressions include controls for labor market tightness and are weighted to match the distribution of CPS occupations; standard errors clustered by firm in parentheses, *** p<0.01, ** p<0.05, * p<0.1. The groups are determined by SOC codes as follows: managers, 11; professionals, 13-27; health, 29, 31, 39; business support, 33-37; sales, 41; Administrative, 43; construction/production, 45-51; and transportation, 53.

Table A3.

Robustness: Managerial Share

	(1)	(2)	(3)	(4)
Manager Share	0.070***	0.052***		0.160***
	(0.005)	(0.005)		(0.013)
Outsourceable Share		-0.217***		-0.385***
		(0.006)		(0.015)
Small Firm x Manager Share			0.050***	
			(0.005)	
Large Firm x Manager Share			0.617***	
			(0.048)	
Small Firm x Outsourceable Share			-0.216***	
			(0.006)	
Large Firm x Outsourceable Share			-0.179***	
			(0.058)	
Manager Share squared				-0.153***
				(0.017)
Outsourceable Share squared				0.265***
				(0.021)
Observations	144,419	144,419	144,419	144,419
R-squared	0.002	0.011	0.013	0.013
Test: large coefficient = small (probability value)				
Manager Share			0	
Outsourceable Share			0.522	

This table repeats table 1, but with manager shares instead of software shares. The firm fixed effects are partially explained by manager share, but these point estimates are much smaller than the coefficients shown in table 1. Additionally, the small and large firm differences are not statistically significant.

Table A4.

Initial Regression to Calculate Firm Fixed Effects

VARIABLES	Log of Avg Salary
Number of Skills Requested	0.004***
	(0.000)
Non-Microsoft Office IT	0.012***
	(0.001)
AI Skill Required	0.034***
	(0.008)
Soft Skill Required	0.006***
	(0.001)
Experience	0.096***
	(0.000)
Experience Squared	-0.004***
	(0.000)
Labor Market Tightness	0.011***
	(0.002)
Fixed Effects	
2-Digit Occupation	Yes
Educational Attainment	Yes
Year	Yes
State	Yes
Firm Name	Yes
Observations	2,918,605
R-squared	0.688