

Occupational Skill Compositions and the Values of Skills: Panel Estimates from 200 million online Job Postings

Sebastian Steffen

I derive a novel occupation(-industry) level panel of skill demands from the near-universe of tagged online job postings in the US for the last decade (2010-2020). I use this data to study how the skill demands of occupations have changed and how these changes affect the market returns to skills. Low- and medium-wage occupations' skill demands changed more than those of high-wage ones. Thus, lower-wage workers face not only higher risks of direct technological displacement but also increased risks of reskilling in order to stay productive. I show that routine-biased technological change (RBTC) due to automation technologies such as ML can explain some these results, while skill-biased and (endogenously) directed technological change cannot. Technical skills, such as ML, Business, Software, and Data Skills have particularly high implied market values, as do Social Skills and Creativity. These therefore represent lucrative (re-)skill investment opportunities for workers unlike non-cognitive skills. Finally, there is significant heterogeneity by industry with positive returns to the Utilities, Mining and IT Industries and negative ones for the Food and Retail industries, even after controlling for skills.

1 Introduction

Which skills should workers invest in to get jobs and increase their earnings? Skills are intangible investment assets, which have value, face unexpected future returns, and may incur risk of depreciation. After all, skills are also referred to as human *capital* (Becker (1964)). The social sciences have a long history of studying the returns to skills, albeit through proxy variables

such as years of schooling and college attendance (Mincer (1974); Goldin and Katz (2008, 2009); Michaels et al. (2014); Beaudry et al. (2015), wages Autor and Dorn (2013), or ability test scores Hanushek et al. (2015)). Such measures are static, highly aggregated, supply-side based and generally unable to capture dynamic short-term changes in the value of the *actual* skills that workers consider investing time, money, and effort in. With the rapid adoption of novel risk factors, such as artificial intelligence, robotics, and machine learning, uncertainty over which are the 'right' skills has increased.

ML's vast improvements to prediction tasks enables it to substitute, complement, and expand demand for occupations or skills within them (Agrawal et al. (2019)). In particular, ML's ability to learn from large amounts of data make it exceptionally suitable to substitute for routine skills and complement non-routine skills (Goos et al. (2014); Brynjolfsson et al. (2018)) - it is a form of Routine-Biased Technological Change (RBTC). Thus, the widespread adoption of ML may lead to a skill mismatch between the skill demands of jobs and workers' current skills and may lead to technological unemployment (Restrepo (2015)). Many firms plan to adopt ML capabilities to leverage data-driven decision-making but a large number lacks the necessary complements, including skilled workers, to do so (Brynjolfsson and McElheran (2016a)).

While the adoption of ML has naturally led to a rise in the implied market value for ML-related skills, there has also been a depreciation of the value of routine skills. Thus, the relatively higher prevalence of routine skills in medium-wage occupations makes these jobs more susceptible to ML, such that ML may have been an important factor for occupational wage-polarization (Autor and Dorn (2013)). This polarization has led to a considerable rise in the service industry, in particular for workers in health services or last-mile services such as transporting people and packages (Autor (2019)). ML also increased the demand for 'ghost' workers who help to generate the massive labeled datasets necessary to train ML models (Gray and Suri (2019)).

In this paper I study the interactions between the implied market value of skills, changes in occupational skill demands, and ML by leveraging a novel panel dataset of occupation-level skill demands, which I derive from the near-universe of US online job postings of the last decade. Besides estimating returns to skills from this skill demand-side dataset, I also use it to understand what is driving the recent changes in occupational skill demands and which of the three leading theories of technological change can best explain them.

The paper proceeds as follows: section (2) gives an overview of the returns to skills and technological change and automation literature while section (3) explains our data and methodology that allows us to go from annotated job postings to yearly skill shares of occupation-industry cells. In the results section I first show that occupational skill demands are indeed changing and that low and medium-wage occupation-industry cells' skill shares changed more between 2010 and 2018 than those of high-wage ones. Then I present the results from our skill panel regression. As a robustness check I use two different types of skill classifications and show different sets of fixed effects models. Finally, I show the results (4) of our occupational change regression on proxy variables for each of the three leading technological change theories. Section (5) concludes.

2 Literature

2.1 Returns to Skills

Estimating the returns to skills goes back at least 70 years to Jacob Mincer's famous 'Mincer Earnings Regressions' (Mincer (1958)). In his seminal work he estimates:

$$\ln w(s, x) = \alpha_0 + \rho_s s + \beta_0 x + \beta_1 x^2 + \epsilon \quad (1)$$

where $w(s, x)$ is the wage at schooling level s and work experience x ¹. While the Mincer model was a work horse in the past, it has an important shortcoming which is incompatible with the increasing polarization of the wage distribution since the 1980s (Firpo et al. (2011); Lemieux (2006); Autor and Dorn (2013)). In the canonical Mincer model, changes in wages can only occur due to changes in skill prices. Furthermore, *conditional on skills*, wages would be identical across occupations, which is unrealistic.

Another model, by Welch (1969) allows for the skill supply to be bundled, within workers, such that worker i 's wage, w_{it} can be estimated as:

$$w_{it} = \theta_t + \sum_{k=1}^K r_{kt} S_{ik} + u_{it} \quad (2)$$

¹See (Heckman et al. (2003)) for an excellent overview.

, where S_{ik} are the skill components of worker i , and r_{kt} are the desired returns to skills coefficients. However, as discussed in Firpo et al. (2011), this model still cannot account for unequal wages for identical skill bundles in different occupations - it requires that workers' skill bundles could be unbundled, which is not the case (Gibbons et al. (2005)). To allow for wages to differ across occupations for similar skill bundles, Firpo et al. (2011) estimate the following Welch-inspired model for wages of worker i in occupation j :

$$w_{ijt} = \theta_t + \sum_{k=1}^K r_{jkt} S_{ik} + u_{ijt} \quad (3)$$

This model is now general enough to account for technological change as well as offshoring. However, it still does not take industries into account, even though there is strong evidence for industry fixed effects.

Estimating the returns to skills from the skill supply side, i.e. from workers, is very challenging as the skill components of workers are generally latent. Schooling was the earliest proxy variable used by Mincer as well as Becker. However, with the significant rise in college-educated workers, measuring the return to an additional year of school, their proxy for the returns to skills, is no longer relevant. While the definition of the word skill is still hard to pin down, it has certainly moved away from being equivalent to attaining a college degree. Even test scores may only be poor indicators of skill. Thus, the more recent returns to skills estimates, which leverage the Programme for the International Assessment of Adult Competencies (PIAAC) test scores Hanushek et al. (2015) may not be reliable either. Even if the test accurately measured skills around the time of test-taking, it becomes a worse skill measure the longer ago an individual took the test. In addition, there are other confounding hard-to-measure variables, such as grit Duckworth and Gross (2014), or latent, idiosyncratic preferences.

Due to the issues discussed above, I move away from skill supply side estimations and propose a *skill demand side estimation* in a similar style to Firpo et al. (2011). Thus, instead of estimating the returns that a worker gets from his skill bundle, I estimate how the market values skill bundles, by decomposing wages for each occupation into returns to the skills that are demanded for that occupation. By leveraging a large, and novel dataset of nearly all online US job postings of the last decade, I am able to gain additional data granularity and estimate not only an occupation-level panel but also an occupation-industry-level panel, thereby allowing for the theorized

industry-fixed effects.

2.2 Technological Change and Automation

Artificial Intelligence is widely believed to be the next big General Purpose Technology (Brynjolfsson et al., 2018). Automation, more generally, has the capacity to make labor more productive (labor-augmenting automation), to make automation itself more productive (automation at the intensive margin), to introduce new skills into the economy, or to displace a wider range of tasks (automation at the extensive margin)(Acemoglu and Restrepo, 2019). This race between man and machine may lead to a rise of technological unemployment (Acemoglu and Restrepo, 2018). No matter whether automation or (task) innovation 'wins', both forces lead to changes in occupations' underlying skill requirements and force workers to reskill to remain productive. However, given that automation and other IT capital are relative substitutes for workers who perform routine manual and cognitive tasks, but relative complements for workers who perform non-routine cognitive tasks, low-wage workers will face the brunt of the these changes (Autor et al., 2003; Michaels et al., 2014).

In fact, some of these changes have already manifested themselves. Some argue that the terms routine and non-routine characterize the relationship between tasks/skills and information technology (IT) and find that occupations have shifted towards requiring more analytical and interactive tasks and away from requiring cognitive-routine and manual-routine tasks (Spitz-Oener, 2006). Skills, as a form of task-specific human capital, are an important source of individual wage growth (Gathmann and Schoenberg, 2010). Thus, the relative loss of productivity of routine skills translates to lower wages and an overall more polarized wage and employment share distribution (Autor and Dorn, 2013).

However, medium- and high-wage occupations are not immune to occupational change either. Occupations that heavily rely on IT skills have been shown to change faster due to rapid software innovation (Hershbein and Kahn, 2018). These fast obsolescence rates of specific software skills lead to relatively flatter earnings profiles for STEM workers (Deming and Noray, 2018). Some have argued for a 'great reversal' in demand for cognitive skill and shown that more educated workers have begun to crowd out less educated workers, due to sorting and changes in relative productivity of workers and capital (Beaudry et al., 2015). Automation and IT capital, such as Data-

Driven Decision Making (DDD), have been rapidly adopted and have made plants more productive and efficient, requiring managers and other high-wage occupations to adapt to stay productive (Brynjolfsson and McElheran, 2016b; Bartel et al., 2007). These results suggest that reskilling is both necessary as well as costly, in particular for low-wage workers, and that the dynamics of occupational skill demands are an important phenomenon to study. This is especially true in light of IT enabling faster rates of technological adoption due to low margin costs and scalability.

In this paper I study the effects of skill demands and automation on three occupational outcomes: changes in (i) wages, (ii) employment shares, and (iii) occupational skill demands. In particular, I focus on occupational change and leverage a novel large data set of online job postings between 2010 and 2019 from which I derive a panel of occupational **skill share vectors**.

Our main results show that AI has multiple negative consequences for low and medium-wage occupations: not only do occupations that are highly susceptible to automation correlate with decreased employment shares, they also face much higher occupational change in terms of skills demanded. Thus, the 'lucky' low- and medium-wage workers that do manage to keep their job, still have to reskill relatively more in order to keep it. Notably, an occupation's suitability for machine learning (SML) does not correlate with decreased wages, perhaps due to workers being highly forward-looking and abandoning obsolete skills (Horton and Tambe, 2019).

I also find that certain skill groups correlate particularly strongly with occupational change. These include 'Information Technology' and 'Engineering', consistent with prior findings on the fast obsolescence rates of specific software skills in STEM occupations (Deming and Noray, 2018), as well as 'Maintenance and Repair' and 'Environment' skills, which is consistent with automation and IT displacing routine non-cognitive skills relatively more easily (RBTC) (Goos et al., 2014; Michaels et al., 2014).

3 Data and Methodology

Our methodology relies on several different data sources:

- Burning Glass Technologies (BGT),
- an analytics software company that scrapes and annotates job postings from nearly all online job posting sites and employment search engines,

- The US Bureau of Labor Statistics (BLS) Occupational Employment Survey (OES), which provides historic wages and employment shares at the occupation and occupation-industry level
- Automation-related proxies from several recent papers: Suitability for Machine Learning (SML) from (Brynjolfsson et al., 2018), AI Progress Scores from (Felten et al., 2018), Cognitive Skill Fractions from (Alabdulkreem et al., 2018), and the AI, Software, and Robot Indices created by (Webb, 2019).

I will first describe these data sources in more detail before delving more into the methodology of going from job postings to yearly skill panels at the occupation and occupation-industry levels.

3.1 Millions of Annotated Job Postings

The BGT data covers about 200 million online job vacancy postings posted on over 40,000 distinct online job platforms in the United States between 2010 and 2018 and arguably covers the near-universe of job postings (Figure 1). Each vacancy posting is parsed, deduplicated, and annotated with the posting date, the SOC occupational code, the NAICS industry code, and which skills were demanded among several other variables. The skills data is annotated via BGT’s industry-leading skill parser, which is rule-based and employs string searches as well as disambiguation rules. It maps each job postings’ skills into a detailed skills taxonomy, which consists of 3 levels of granularity.

At the most detailed level, the BGT taxonomy includes $\approx 16,000$ skills - these are nested within 658 skill clusters, which are nested within 28 skill cluster families. For example, *Python* is a skill within the *Scripting Languages* skill cluster, which itself falls into the *Information Technology* skill cluster family. The taxonomy was initially assembled from online resumes and is continuously updated through client feedback, research, and forums. K-Means clustering along with additional qualitative checks were employed to create meaningful skill clusters. Whenever new skills are added to the taxonomy, the labels are refit to the entire history of job postings data. This minimizes potential biases which may arise through BGT’s time-varying ability to capture new skills.

Notably, this taxonomy is significantly more detailed than other skill taxonomies, such as the Bureau of Labor Statistics (BLS)’ O*Net skill taxon-

omy, which contains just 2 levels, with 35 skills mapped into 6 skill groups. Furthermore BGT job postings are scraped daily and are therefore able to capture changes in skill demands at a much higher temporal granularity. O*Net only undergoes yearly updates which generally only cover a subset of occupations.

However, broader taxonomies are still useful as they can be more interpretable, in particular in terms of the routine and non-routine skill distinction of the RBTC theory. Since the BGT taxonomy does not directly map to the O*Net taxonomy, I instead leverage the taxonomy built by Deming and Noray (2018) for additional robustness checks. Their taxonomy maps a subset of the BGT skills into 20 interpretable skill groups, that contain 'social skills', 'cognitive skills', and 'management skills' among others.

The BGT data is ideal for a panel study of occupational skill demands for several reasons. Given that each job can be viewed as a 'bundle of skills' (Restrepo, 2015; Deming and Kahn, 2018), each job posting represents a draw from the 'ground-truth' skill distribution of the job's underlying occupation. Hundreds of thousands of job postings can thus pin down a very tight distribution of skill demands for a given occupation. However, some occupations are notoriously underrepresented in online job platforms, and thus in BGT's data, in particular blue-collar jobs. I therefore remove occupations with very few ($< 1,000$) job postings per year which leaves us with TODO occupations. To ensure that each aggregation cell contains enough job posting data points, I use years, instead of months, as our temporary aggregation variable. This also minimizes potential seasonality effects which may distort skill demands. It furthermore allows us to aggregate to the occupation-industry level as well. Since BGT also annotates job postings with 2-digit industry NAICS codes, I am able to derive a skills demand panel at the SOC6-NAICS2 level with enough job postings in each aggregation cell.

Besides the granularity enabled through its size, this data has another important feature that may be advantageous to study the returns to skills. Unlike most other papers which generally leverage a labor supply side panel with a static skill proxy, I rely on aggregated job postings, which represent the skill demand side. Previous studies face several significant sources of omitted variable biases, such as grit Duckworth and Gross (2014), and idiosyncratic preferences for non-pecuniary benefits Katz and Autor (1999), such as occupational choice, work hours, culture, among others.

3.2 Wages and other Occupational Data

The Bureau of Labor Statistics (BLS) collects and publishes yearly data on employment shares and wages at different levels of aggregation. The BLS classifies occupations according to the Standard Occupational Classification System, SOC codes. The 2010 version defines 840 distinct detailed occupations, which are nested in 461 broad occupations, 97 minor groups, and 23 major groups. It also classifies industries according to the North American Industry Classification System, NAICS codes. The 2012 version defines 20 distinct industry sectors. I use their annual publications of wages and employment shares at the SOC level as well as at the SOC6 x NAICS2 level.

3.3 Occupational Automation

Besides these two main data sources I leverage several recent papers' occupation-level automation scores. I include automation as a proxy for the level of routineness of occupations, as routine tasks are more automatable (Autor and Dorn, 2013). Routine tasks are particularly suitable for Machine Learning, which is why our preferred measure is the Suitability for Machine Learning (SML) metric from (Brynjolfsson et al., 2018). I consider additional measures in the appendix. These include the AI scores from (Felten et al., 2018), which are based on the Electronic Frontier Foundation (EFF) AI Progress report, as well as the O*Net-derived cognitive skill fraction as defined in (Alabdulkreem et al., 2018). Finally, I use the recent AI, software, and robot indices created by (Webb, 2019) from patent data.

3.4 From Job Postings to a Panel of Occupational Skill Demands

As previewed above, I aggregate the individual job postings to the occupation-year level, thereby viewing each of the 600 million job ads as a draw from the corresponding occupational skill distribution for occupation i in year t . While the majority of the 840 different occupations listed in the SOC taxonomy are present in our data, I remove military occupations and occupations with fewer than 1,000 postings leaving us with 556 distinct occupations. I also noticed a 'buzzword bias' which led to inflationary usage of certain popular skill words within job postings, thereby misrepresenting actual skill demands. Thus, in our preferred specification I work with the more meaningful skill clusters -

this way mentions of the skills such as 'AWS' and 'Azure' within the same posting are only counted once as the skill cluster 'Cloud Solutions'.

Since the total number of job postings increases from about 80 million in 2010 to 180 million in 2018, I normalize the raw skill cluster counts to derive the *skill share vectors* for each of the 556×9 occupation-year cells.

$$\mathbf{s}_{ikt} = \frac{1}{\sum_{j=1}^S (s_{i,t,j})} (s_{i,t,1}, \dots, s_{i,t,S}) \quad (\text{Skill Shares})$$

represents the share of skill k of all skill demands in occupation i at time t . To illustrate, figures 4 and 5 show the skill demand shares for Data Scientists and Lumberjacks, respectively, based on just the 30 most relevant skill clusters. These top 30 skill clusters cover about 12% and 66% of all skill demands for Data Scientist and Lumberjack occupations, respectively. The changes in skill demands over the past decade are clearly visible and suggest that the implied market skills of skills vary considerably. This confirms that job seekers indeed face considerable risk in terms of how their skills are valued in the market but also that investment in the 'right' skills, i.e. via reskilling, may be very lucrative.

3.5 Measuring Changes in Occupational Skill Demands

To measure more rigorously how large these changes in occupational skill shares are across time, I apply Cosine Distance between the earliest (2010) and latest (2018) skill share vectors:

$$d_{\cos}(\mathbf{s}_{i,2010}, \mathbf{s}_{i,2018}) = 1 - \frac{\mathbf{s}_{i,2010} \cdot \mathbf{s}_{i,2018}}{\|\mathbf{s}_{i,2010}\| \|\mathbf{s}_{i,2018}\|} \quad (\text{Cosine Distance})$$

Notably, I calculate this distance over the entire skill share vectors and not just for the top 30 skills. There are many other suitable distance and similarity metrics to choose from (Cha, 2007), such as the Jensen-Shannon Divergence. However, the Euclidean distance is not one of them as the curse of dimensionality makes it meaningless in high dimensions, which is the case here due to the large number (658) of skills clusters.

Cosine Distance is one of the most widely used distance measures as it is fast to calculate and easy to interpret. It measures the angle between two vectors, with a magnitude of 1 implying perfect alignment, 0 indicating orthogonality, and -1 indicating perfectly opposite alignment. Occupational

skill demands do not change by large magnitudes, such that none of the observed cosine distances falls below 0. For example, the cosine similarity between the 2010 and 2018 skill share vectors of the aforementioned Lumberjacks is 0.37 and is one of the larger changes in our data.

I calculate the Cosine Distance between the 2010 and 2018 skill share vectors for each SOC occupation in our sample. As can be seen in Figure 6, it appears that low wage and medium wage occupations changed significantly more than high wage occupations between 2010 and 2018. This is consistent with routine-biased technological change, as low and medium wage occupations tend to require more routine tasks. I will show the robustness of this relationship with regressions in the results section. Due to the large size of our data I can do the same for each occupation-industry combination, i.e. for each SOC6-NAICS2 cell (Figure 7). I again observe a similar trend. Low and medium-wage occupation-industry combinations changed more than high-wage ones. However, with this additional data granularity there seems to be another point of inflection around the 70th wage percentile.

4 Estimation & Results

4.1 The Implied Market Values of Skills

The returns to skills literature has previously faced challenging data problems:

Firpo et al. (2011) "Ideally, we would like to estimate the skill pricing parameters (r_{jkt}) using repeated cross sections from a large data set containing detailed information on wages, skills, and occupations. We could then look at the contribution of changes in occupational wage setting to the overall changes in the wage structure [...]. **Unfortunately, no such data set exists.**"

While this remains true for skill supply side estimations based on worker panels, I believe that our panel data allows the estimation of skill pricing parameters from the skill demand side. In particular, I estimate the following:

$$w_{(ijt)} = \theta_i + \theta_j + \sum_{k=1}^K r_{ijk} s_{ijkt} + u_{ijt},$$

where w_{ijt} is the wage of occupation i in industry j at time t , the r_{ijk} are the returns to each skill k in occupation i and industry j , the s_{ijkt} are the skill shares of each skill k in occupation i and industry j at time t , and θ_i and θ_j are the occupation and industry-specific fixed effects, respectively. Roughly speaking, by viewing each occupation as a bundle of skills, the wage paid to occupation i in industry j can be decomposed as a weighted average of the values of each skill and how much each skill is demanded for said occupation-industry combination.

Since the skill shares, s_{ijkt} sum to one, the interpretation of the coefficients r_{ijk} requires additional care. There are two issues, the first one being very minor and having a simple fix, the second one being more severe and requiring techniques developed for compositional data. In particular the first problem is due to the skill shares, s_{ijkt} , falling between zero and one. The usual interpretation of coefficient r_{ijk} is that a 1 unit change in s_{ijk} results in a r_{ijk} unit in the outcome, w_{ijt} , *ceteris paribus*, or all else equal. However, 1 unit changes are impossible when each skill share is a proportion that lies between 0 and 1. The easy fix involves considering 0.01 unit changes instead of 1 unit changes. These smaller changes can then be interpreted as $\frac{1}{100}r_{ijk}$ unit changes in Y . Conveniently, a 0.01 unit change is the same as a 1 percentage point change in the skill share.

The second issue is much more severe and is due to the *all else equal* part. Compositional data, such as the skill shares here, always sum to one, such that one cannot hold all other covariates constant when changing one. Geometrically, each point as defined by the K -dimensional skill share vector, s_{ij} is bound to a K -dimensional simplex. In other words, the skill shares are perfectly multicollinear. It is therefore tempting to simply drop one of them and rerun the regression - Stata and other statistical tools do this automatically. However, in the context of compositional covariates, in which only relative instead of absolute amounts matter, this severely biases the regression coefficients and leads to misleading inference statistics ?. Instead, I follow the compositional covariate methodology described by ?. To summarize, I first use the isometric logratio (ILR) transformation to project the K -dimensional skill shares into an orthonormal $K-1$ dimensional space. Even though the bases in this space are uninterpretable, I can still interpret the first base, by aligning it with one of the bases from the K dimensions. This procedure is similar to the Gram-Schmidt orthogonalization technique, in which the first vector only gets normalized, but still points in the same direction as the original.

I can thus iterate through each of the K bases, run the regression on the K-1 bases from the ILR transform and only record the first coefficient.

I again run these regressions for both the BGT-provided skill cluster families, as well as for the skill groups provided by Deming (2018). The results can be seen in tables (1) and (4). The corresponding industry-fixed effects can be seen in tables (3) and (6). For the corresponding occupation-fixed effects I only report the top 5 and bottom 5 in tables (2), and (5), as there are over 600 occupations.

The implied market values for the BGT skill cluster families 'Economic Policy', 'Analysis', and 'Design' were the three highest with values of over \$100 for a marginal percentage point. While there is no direct correspondence with the Deming skill groups, the BGT skill cluster families seem to roughly correspond with the highest-valued Deming skill groups. Those were 'ML, AI' as well as 'Business Systems', 'General Software' and 'Creativity'.

4.2 What is driving Occupational Change?

Before looking at occupational changes in skill shares, I present evidence for the (overall) demands for skills (Figure (2)) and, specifically for IT skills (Figure (3)). I can immediately see that there is considerable variation in the demands for different skills. For example, I can see the tremendous rise in demand for 'Big Data' and 'Artificial Intelligence' in high-wage occupations since about 2012, and 2016, respectively. These figures are on a log scale, so the demand for these increased by a magnitude of over 1000. Other important IT skills for high-paid jobs include Java, Javascript, and SQL which each made up almost 1 of every 1000 IT skills demanded in job postings. Considering that IT skills made about 5%, i.e. 1 in every 20 skills demanded in high-wage job postings, these skills appear to be particularly good investments. Another interesting observation, among others, is the shift of health skills from medium-wage occupation to low-wage ones.

Going one level deeper into the aggregation, I can study more detailed skill demands at the occupation level. Figures (4) and Figure (5) show the yearly skill share distribution for data scientists and lumberjacks, respectively. In each year column I can see the share that each skill made up of all skills demanded in job postings associated with these occupations in that year. Again, I can see the rise of 'Big Data' (dark blue) and 'Data Science' (blue) starting in about 2012 for Data Scientists. Naturally, the occupation of lumberjacks had very different skill demands, that include landscaping and,

increasingly, agronomy and farming. It is apparent that the skill distribution for this occupation seems to have changed between 2010 and 2018 and I can measure this change with the aforementioned cosine distance.

I do not purport to know why, specifically, certain occupational skill demands changed or appear to be high. A priori, I would not have expected the occupation of lumberjacks to undergo such changes. Notably, these changes cannot be due to the BGT skills parser classifying skills inconsistently across time. If it indeed misclassified skill words, it would do so *consistently* across time. Thus, sudden jumps in skill demands are due to actual changes in the language of the underlying job postings, for example due to the increased popularity of one word over its synonym. While these changes in language may not represent changes in actual skill demands, they are still important for workers as they can help them better match and be accepted for a job posting. Moreover, by only counting skills within the same skill cluster once instead of multiple times, I can at least partially account for some of these worries, as, for example similar skills such as 'AWS' and 'Azure' are both contained with the skill cluster 'Cloud' skills.

In Figure (6), I plot the cosine distance between 2010 and 2018 skill share vectors over the 2010 wage percentile (each point represents a different occupation). In the literature, the wage percentile is also referred to as the skill percentile, assuming that higher wage proxies higher skill Autor and Dorn (2013). If this assumption holds, the figure supports the 'skill'-biased technological change theory. Lower-wage and medium-wage occupations incur higher changes to their skill share distribution than higher wage ones. These trends remain when the analyses are repeated at the occupation-industry level. A similar positive association between wage and skill share change can be see in Figure (7), which plots the same variables as Figure (6) but each point represents an occupation-industry combination rather than an occupation.

The above patterns support one leading theory of technological change (SBTC). To also test for the implications of and discern multiple leading theories of technical change, I turn to regression analyses. I regress the cosine distance on proxy variables for each of the three leading technological change theories: (1) SBTC, (2) RBTC, and (3) Directed Technological Change:

$$d_{cos}(s_{i,2010}, s_{i,2018}) = \beta_0 + \beta_1 SML_i + \beta_2 w_{i,2010} + \beta_3 w_{i,2010}^2 + \beta_4 wagebill_{i,2010} + \epsilon_i$$

(1) SBTC: To account for SBTC I follow the standard approach used for example by Autor, Dorn (2013) and use real wages from the earliest period,

in this case 2010. Given that wages are polarized, I include squared wages in two of the columns of table (1). In the other columns I use wage tercile dummies to allow for more flexibility in the model.

(2) RBTC: To account for RBTC, I use the SML scores from Brynjolfsson, Mitchell, Rock (2018). The suitability for Machine Learning of an occupation is a good measure for this, because routine tasks are exactly the tasks that ML is particularly suitable for. I prefer this measure to the Routine Task Intensity (RTI) scores from Autor, Dorn (2013) since the latter relies on 5 non-maintained measures from the former Dictionary of Occupational Titles (now O*Net) ².

(3) Directed Technological Change: According to Acemoglu (1998), technological change happens endogenously and its direction is determined by the size of the market for different inventions. Thus, if technological change is assumed to happen at the occupation level, it should correspond with the total wage bill of occupations, i.e. the product of wage and total employment of that occupation.

The results are shown in table (1). Notably, the SML score, our measure of RBTC, is positively correlated with occupational change, i.e. more routine occupations are associated with more changes in skill demands, *ceteris paribus*. SBTC is not significant. It is somewhat puzzling that the wage bill is significant but with a sign opposite of what Acemoglu (1998) would have predicted. Perhaps, an additional factor for the direction of technological change, besides the profitability of innovation, is the level of difficulty of innovation. This would be plausible if ideas are indeed getting harder to find, the more progress has been made ?.

The fixed-effects models in columns (2) and (3) of table (1) show that the implied market values of skills differ considerably. The coefficients can be interpreted as the effective Dollar-amount associated with one additional percentage point in the skill share distribution of an occupation-industry cell. For example, an additional percentage point of 'Economics, Policy' skills is associated with an increase in wage of more than \$400. The fact that some skills have negative coefficients implies that these skills are not a good investment as the market does not value them. However, this does not mean that they are useless - given the bundling of skills into occupations, they may

²Specifically, manual tasks are defined as the DOT score for 'eye-hand-foot coordination', routine tasks are defined as the average of the 'Set limits, tolerances, and standards' and 'finger dexterity' DOT scores, and abstract tasks are defined as the average of 'direction control and planning' and the 'GED Math' score.

still be required to perform those jobs.

I present alternative estimates relying on the Deming (2018) definition of skills in table (4). The interpretation of these results is identical, except that the skill categories differ. I can immediately see the high implied market values for technical skills such as 'Machine Learning', 'AI', and 'General Software' skills. Conversely, non-cognitive skills are associated with negative values. This is in line with the theory of RBTC, as these types of skills are more routine and thus better relative substitutes for automation.

5 Conclusion

Technological change is essential to human progress. However, it also bears risks for some who may be left behind and displaced. Thus, GPTs like Machine Learning, which have the capability to drastically alter society and bring about immense progress, also induce the largest risks. I have shown that what is demanded of workers in terms of skills has changed over the past decade. The fact that low and medium-wage occupations' skill demands changed more than high-wage ones means that besides getting paid less, workers in the former also have to reskill more in order stay productive in and attempt to keep their jobs. However, reskilling may also offer opportunities for social mobility: some skills are highly valued by the market and may well be worth the time and effort investment - in fact, Machine Learning is one of them among other technical skills as well as creative and social skills. Not everyone will be able to acquire or benefit from these skills. In fact, it remains an open question which skills complement each other best in terms of productivity as well as learning.

There have been several theories on how technological innovation progresses - routine-biased (RBTC), wage-biased (SBTC), and endogenously-directed technological change. As automation is becoming more ubiquitous, the routineness of tasks seems so far to be one of the better indicators for where technological progress will progress fastest and therefore where corresponding changes in occupational skill demands will occur.

A Tables

Table 1: Panel Wage Regression - BGT Skills

	<i>Dependent variable:</i>		
	Annual Wage		
	(1)	(2)	(3)
Economics Policy	-586.46 (379.98)	129.18** (55.69)	437.15*** (49.96)
Analysis	1,145.07*** (352.77)	329.34*** (31.89)	182.87*** (28.60)
Design	-631.36* (347.01)	-85.12*** (21.98)	136.06*** (19.92)
Marketing, PR	-120.86 (340.70)	-56.53*** (13.58)	66.50*** (12.29)
Manufacturing	89.30 (341.08)	42.02*** (11.46)	62.99*** (10.27)
Engineering	-83.58 (346.34)	-10.66 (19.57)	37.61** (17.54)
Business	181.00 (340.48)	56.20*** (11.01)	31.11*** (9.88)
Health Care	183.93 (339.47)	-44.04*** (9.55)	19.88** (8.58)
Public Safety, Security	-25.33 (343.28)	-86.96*** (19.65)	14.06 (17.65)
Architecture, Construction	104.79 (340.95)	-67.77*** (12.63)	8.57 (11.33)
HR	462.34 (340.93)	73.10*** (11.93)	-2.18 (10.72)
Environment	-9.29 (344.34)	-139.32*** (21.09)	-2.28 (18.92)
IT	135.16 (340.98)	-72.05*** (11.43)	-4.02 (10.26)
Agriculture	-200.10 (350.44)	-44.59** (17.63)	-6.54 (15.80)
Energy, Utilities	10.42 (345.05)	-57.59** (25.65)	-6.84 (22.98)
Media, Writing	-174.72 (349.68)	-169.68*** (23.19)	-7.24 (20.85)
Finance	-211.76 (341.37)	-90.02*** (11.59)	-8.07 (10.41)
Customer/Client Support	218.25 (340.01)	-43.08*** (8.68)	-12.75 (7.79)
Supply Chain Logistics	311.72 (340.32)	-44.34*** (8.87)	-30.62*** (7.94)
Sales	-6.32 (340.65)	-97.50*** (11.00)	-36.49*** (9.87)
Administration	342.65 (339.95)	3.02 (9.47)	-39.95*** (8.49)
Maintenance, Repair	116.96 (340.79)	-73.31*** (9.42)	-46.06*** (8.44)
Personal Care	146.16 (340.16)	-98.06*** (9.22)	-56.67*** (8.28)
Industry Knowledge	90.68 (341.06)	-114.13*** (12.98)	-81.75*** (11.65)
Education, Training	286.93 (343.29)	-186.47*** (14.07)	-96.82*** (12.64)
Science, Research	-373.19 (352.09)	-169.87*** (23.56)	-120.87*** (21.11)
Legal	-43.74 (344.72)	-222.45*** (20.71)	-199.83*** (18.55)
Occupation-Fixed Effects	<i>No</i>	<i>Yes</i>	<i>Yes</i>
Industry-Fixed Effects	<i>No</i>	<i>Yes</i>	<i>Yes</i>
Time-Fixed Effects	<i>No</i>	<i>No</i>	<i>Yes</i>
Observations	6,658	58,678	58,678
Adjusted R ²	0.99	0.97	0.98

Note:

*p<0.1; **p<0.05; ***p<0.01

Table 2: Top and Bottom 5 Occupation Fixed Effects from BGT Skills Wage Panel Model 3.

SOC Title	SOC6 Code	FE Value
Anesthesiologists	291061	249400.56
Surgeons	291067	243323.18
Oral and maxillofacial surgeons	291022	228877.01
Obstetricians and gynecologists	291064	224534.88
Orthodontists	291023	212590.82
Tapers	472082	22616.72
Ushers, lobby attendants, and ticket takers	393031	22206.52
Food Processing Workers, All Other	513099	21755.97
Veterinary assistants and laboratory animal caretakers	319096	19933.28
Motion picture projectionists	393021	19456.70

Table 3: Industry Fixed Effects from BGT Skills Wage Panel Model 3.

Industry Title	NAICS2 Code	FE Value
Utilities	22	13483.27
Mining	21	10903.34
Management of Companies and Enterprises	55	9880.43
Professional, Scientific, and Technical Services	54	8935.77
Information	51	8508.38
Finance and Insurance	52	7714.41
Transportation and Warehousing	48-49	4391.52
Manufacturing	31-33	4250.13
Wholesale Trade	42	4181.24
Construction	23	3294.98
Educational Services	61	0.00
Real Estate and Rental and Leasing	53	-318.61
Health Care and Social Assistance	62	-832.50
Admin, Support, Waste Management, Remediation Services	56	-919.16
Agriculture, Forestry, Fishing, and Hunting	11	-1016.43
Other Services (except Public Administration)	81	-2015.06
Arts, Entertainment, and Recreation	71	-2549.52
Retail Trade	44-45	-3711.92
Accomodation and Food Service	72	-4325.95

Table 4: Panel Wage Regression - Deming Skills

	<i>Dependent variable:</i>		
	Annual Wage		
	(1)	(2)	(3)
ML, AI	4,048.60*** (258.86)	4,187.47*** (123.14)	2,760.06*** (111.45)
Business Systems	-2,017.18*** (182.91)	1,101.32*** (50.38)	1,295.10*** (45.31)
General Software	2,957.46*** (301.68)	1,049.75*** (90.69)	923.64*** (81.59)
Creativity	-526.78*** (115.18)	25.85 (40.48)	214.09*** (36.50)
Data Analysis	1,121.32*** (303.66)	622.56*** (88.50)	104.70 (79.67)
Admin, Support	55.19 (48.07)	164.37*** (14.51)	97.81*** (13.06)
Project Management	21.09 (72.82)	100.66*** (23.61)	79.81*** (21.23)
Social	88.11*** (25.19)	114.33*** (10.02)	69.39*** (9.05)
Customer Service	-0.92 (24.80)	11.48 (7.74)	36.03*** (6.96)
Engineering	-1,000.81*** (117.25)	-214.32*** (36.12)	22.31 (32.53)
Cognitive	309.49*** (50.36)	39.49** (15.72)	-18.50 (14.16)
Finance	-990.31*** (72.72)	-108.17*** (17.36)	-20.06 (15.63)
Database	-355.12*** (126.67)	-157.31*** (40.17)	-79.58** (36.11)
Tech Support	-1,016.07*** (104.87)	-436.37*** (35.78)	-89.77*** (32.31)
Computer	-291.76*** (45.65)	-205.69*** (14.22)	-90.39*** (12.82)
Product Marketing	-43.47 (40.19)	-117.47*** (15.14)	-114.12*** (13.61)
Management	-378.05*** (70.64)	-214.38*** (26.33)	-120.48*** (23.68)
Non Cognitive	525.45*** (35.44)	59.06*** (12.03)	-135.84*** (10.95)
Writing	-477.47*** (83.78)	-329.63*** (27.82)	-188.97*** (25.03)
Occupation-Fixed Effects	<i>No</i>	<i>Yes</i>	<i>Yes</i>
Industry-Fixed Effects	<i>No</i>	<i>Yes</i>	<i>Yes</i>
Time-Fixed Effects	<i>No</i>	<i>No</i>	<i>Yes</i>
Observations	6,662	58,873	58,873
Adjusted R ²	0.99	0.97	0.98

Note:

*p<0.1; **p<0.05; ***p<0.01

Table 5: Top and Bottom 5 Occupation Fixed Effects from Deming Skills Wage Panel Model 3.

SOC Title	SOC6 Code	FE Value
Anesthesiologists	291061	249296.67
Surgeons	291067	242321.03
Oral and maxillofacial surgeons	291022	229119.17
Obstetricians and gynecologists	291064	223188.11
Orthodontists	291023	212841.38
Sewing machine operators	516031	21772.52
Veterinary assistants and laboratory animal caretakers	319096	20673.60
Ushers, lobby attendants, and ticket takers	393031	20399.14
Motion picture projectionists	393021	18711.47
Food Processing Workers, All Other	513099	17230.07

Table 6: Industry Fixed Effects from Deming Skills Wage Panel Model 3.

Industry Title	NAICS2 Code	FE Value
Utilities	22	13431.18
Mining	21	11152.31
Management of Companies and Enterprises	55	10096.24
Professional, Scientific, and Technical Services	54	8017.92
Information	51	7357.04
Finance and Insurance	52	7200.85
Manufacturing	31-33	5315.15
Transportation and Warehousing	48-49	5044.71
Wholesale Trade	42	4121.35
Construction	23	3364.98
Educational Services	61	0.00
Real Estate and Rental and Leasing	53	-99.76
Health Care and Social Assistance	62	-954.71
Agriculture, Forestry, Fishing, and Hunting	11	-990.07
Admin, Support, Waste Management, Remediation Services	56	-1055.74
Other Services (except Public Administration)	81	-1727.06
Arts, Entertainment, and Recreation	71	-2352.31
Retail Trade	44-45	-2873.29
Accomodation and Food Service	72	-4629.06

Table 7: Occupational Change

	<i>Dependent variable:</i>			
	Occupational (2010-2018) Skill Change			
	(1)	(2)	(3)	(4)
SML Score	0.08*** (0.02)	0.10*** (0.03)	0.06*** (0.02)	0.08*** (0.03)
Log Wage	0.32 (0.23)	0.13 (0.30)		
Log Wage ²	-0.01 (0.01)	-0.01 (0.01)		
Medium Wage Tercile			0.02** (0.01)	0.01 (0.01)
High Wage Tercile			0.01 (0.01)	-0.01 (0.01)
Wage Bill	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)	-0.00*** (0.00)
Constant	-1.96 (1.27)	0.21 (1.75)	-0.10* (0.05)	0.96 (0.64)
Skill Fixed Effects	<i>No</i>	<i>Yes</i>	<i>No</i>	<i>Yes</i>
Observations	650	618	650	618
R ²	0.11	0.20	0.10	0.20
Adjusted R ²	0.10	0.16	0.10	0.16

Note:

*p<0.1; **p<0.05; ***p<0.01

B Figures




Figure 1: Yearly Number of Job Postings scraped by BGT (2010-2018).

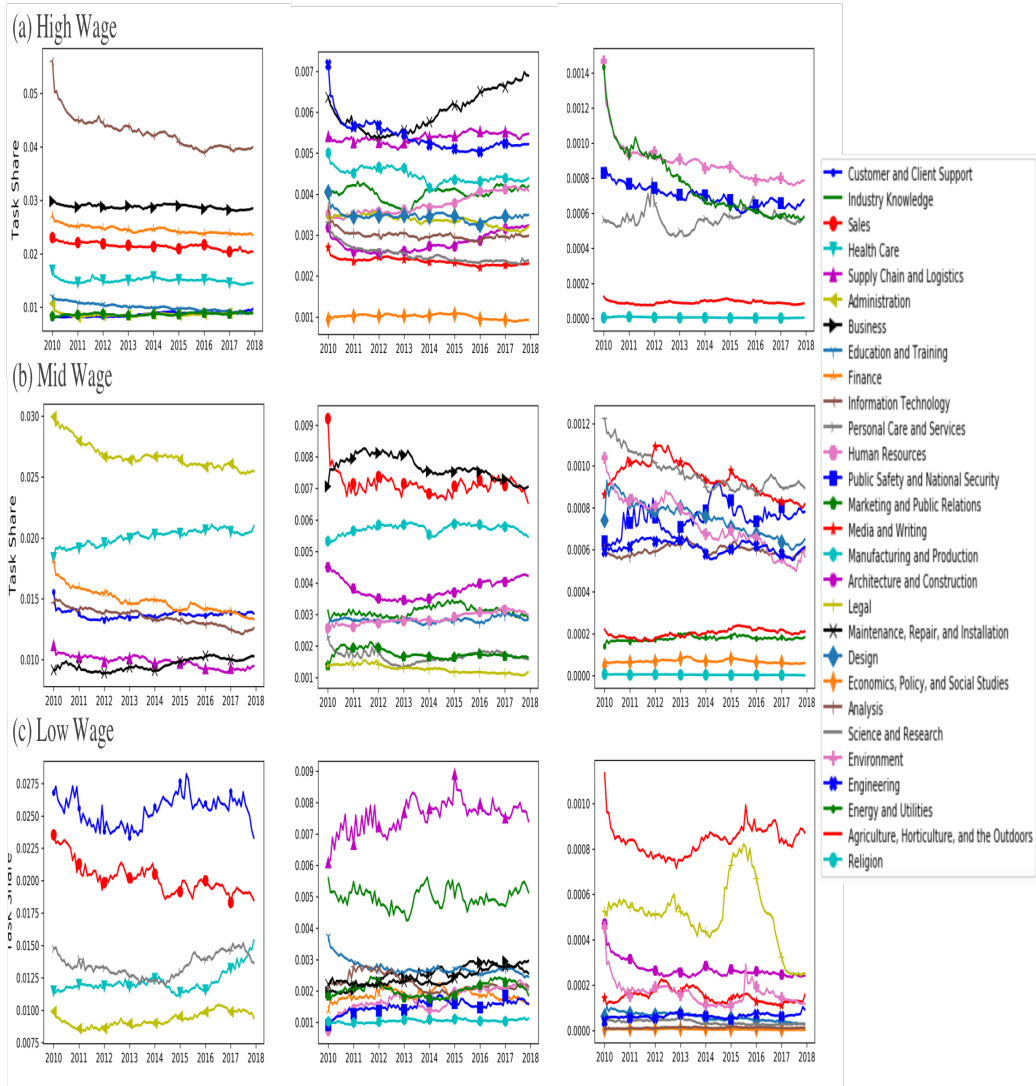


Figure 2: Logged Overall Skill Shares by 2010 Wage Tercile, using the 28 BGT Skill Cluster Families.

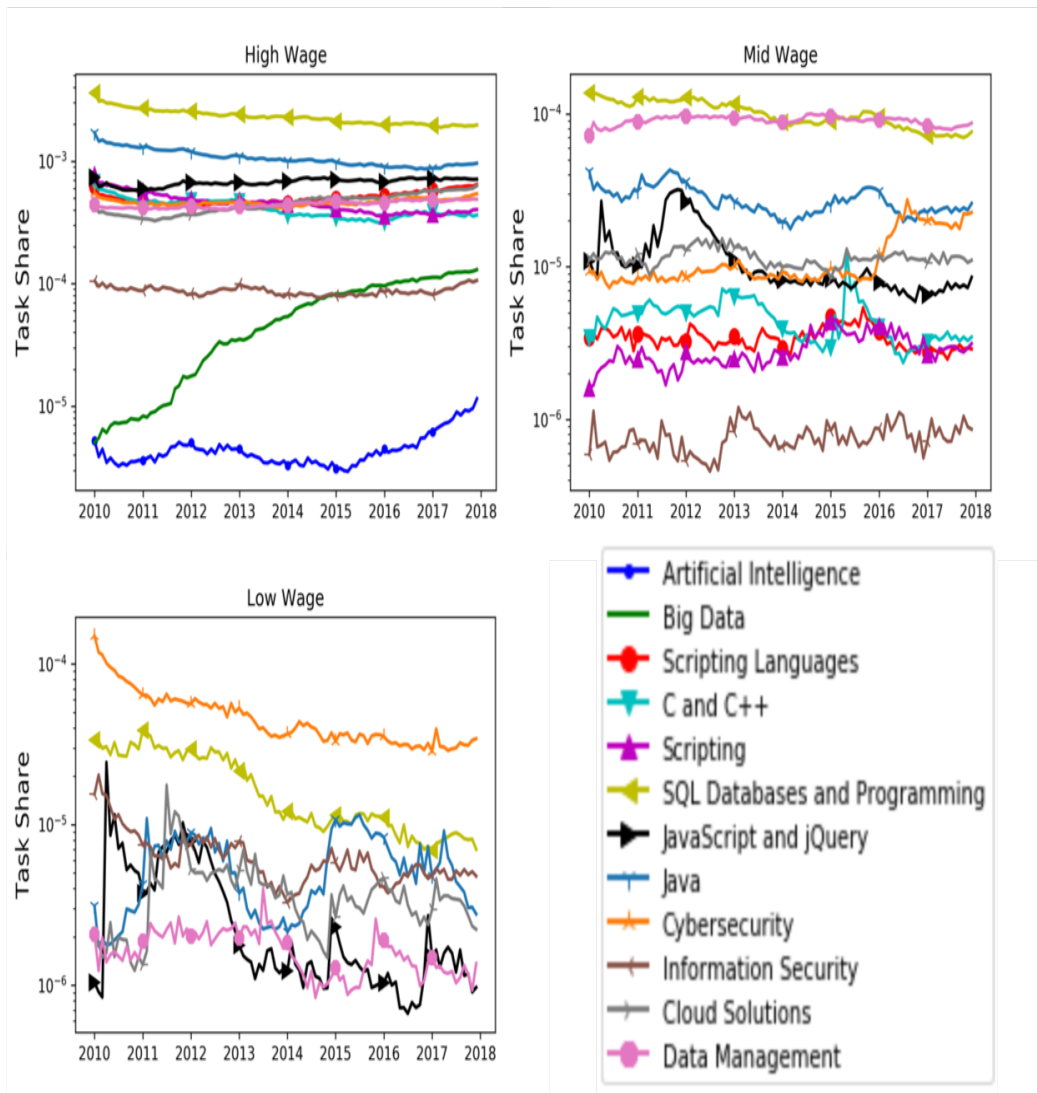


Figure 3: Logged Overall Skill Shares for Skill Clusters within the IT Skill Cluster Family by 2010 Wage Tercile.

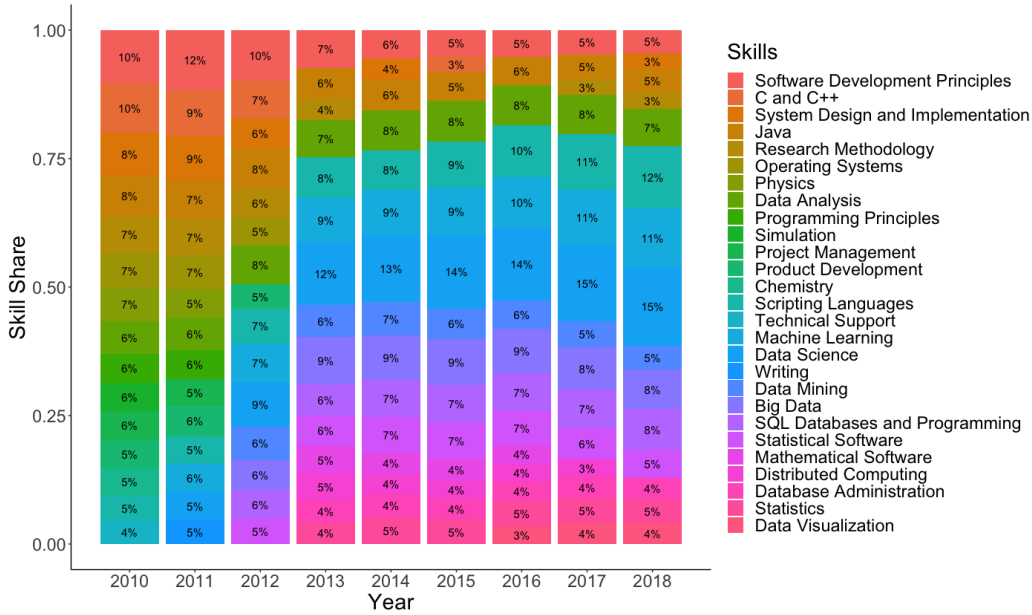


Figure 4: The top 30 skill demand shares of Data Scientists (2010-2018).




Figure 5: The top 30 skill demand shares of Lumberjacks (2010-2018).

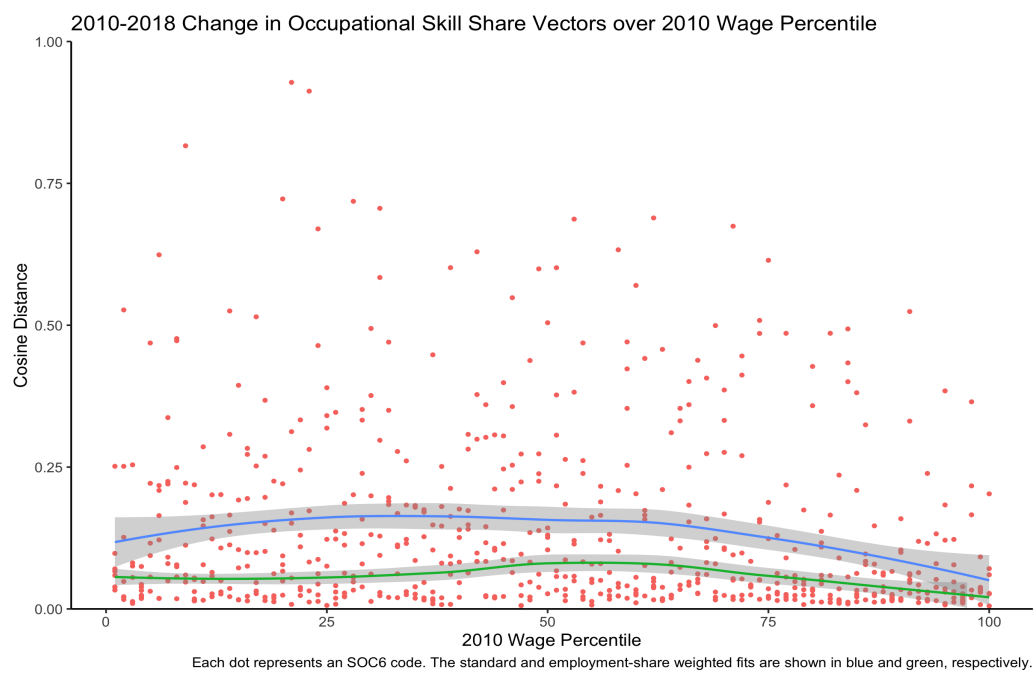


Figure 6: Occupational Cosine Similarity between 2010 and 2018 skill share vectors over 2010 wage percentile.

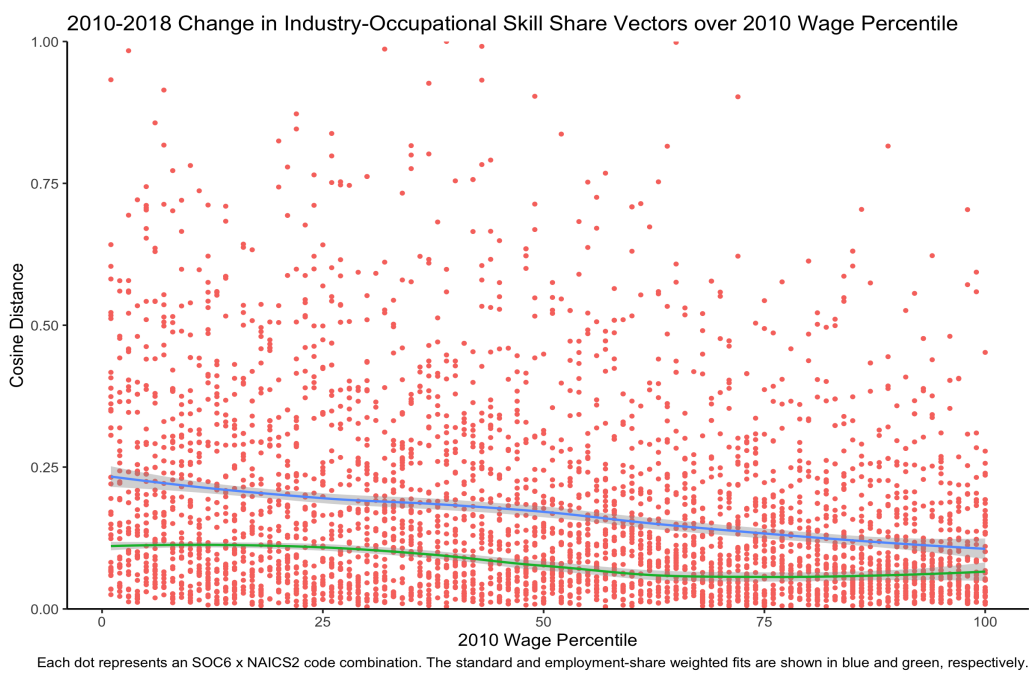


Figure 7: Occupational Cosine Similarity between 2010 and 2018 skill share vectors over 2010 wage percentile.

References

- Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American Economic Review*, 108(6):1488–1542.
- Acemoglu, D. and Restrepo, P. (2019). Automation and New Tasks: How Technology Displaces and Reinstates Labor. *Journal of Economic Perspectives*, 33(2):3–30.
- Agrawal, A., Gans, J. S., and Goldfarb, A. (2019). Artificial Intelligence: The Ambiguous Labor Market Impact of Automating Prediction. *Journal of Economic Perspectives*, 33(2):31–50.
- Alabdulkreem, A., Frank, M. R., Sun, L., Alshebli, B., Hidalgo, C., and Rahwan, I. (2018). Unpacking the polarization of workplace skills. *Science Advances*, 4.
- Autor, D. (2019). Work of the Past, Work of the Future. *National Bureau of Economic Research Working Paper Series*, No. 25588.
- Autor, D. H. and Dorn, D. (2013). The growth of low-skill service jobs and the polarization of the US Labor Market. *American Economic Review*, 103(5):1553–1597.
- Autor, D. H., Levy, F., and Murnane, R. J. (2003). The Skill Content of Recent Technological Change: An Empirical Exploration. *The Quarterly Journal of Economics*, 118(4):1279–1333.
- Bartel, A., Ichniowski, C., and Shaw, K. (2007). How Does Information Technology Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement, and Worker Skills. *Quarterly Journal of Economics*, 122(4):1721–1758.
- Beaudry, P., Green, D. A., and Sand, B. M. (2015). The Great Reversal in the Demand for Skill and Cognitive Tasks. *Journal of Labor Economics*, 34(S1):S199–S247.
- Becker, G. S. (1964). *Human capital; a theoretical and empirical analysis, with special reference to education*. National Bureau of Economic Research.

General series: no. 80. New York, National Bureau of Economic Research; distributed by Columbia University Press, 1964.

- Brynjolfsson, E. and McElheran, K. (2016a). The Rapid Adoption of Data-Driven Decision-Making. *American Economic Review*, 106(5):133–139.
- Brynjolfsson, E. and McElheran, K. (2016b). The Rapid Adoption of Data-Driven Decision-Making. *American Economic Review*, 106(5):133–139.
- Brynjolfsson, E., Mitchell, T., and Rock, D. (2018). What Can Machines Learn and What Does It Mean for Occupations and the Economy? *AEA Papers and Proceedings*, 108:43–47.
- Cha, S.-H. (2007). Comprehensive Survey on Distance/Similarity Measures between Probability Density Functions. *International Journal of Mathematical Models and Methods in Applied Sciences*, 1(4):300–307.
- Deming, D. and Kahn, L. B. (2018). Skill Requirements across Firms and Labor Markets: Evidence from Job Postings for Professionals. *Journal of Labor Economics*, 36(S1):S337–S369.
- Deming, D. J. and Noray, K. (2018). STEM Careers and Technological Change. *NBER Working Paper*.
- Duckworth, A. and Gross, J. J. (2014). Self-Control and Grit. *Current Directions in Psychological Science*, 23(5):319–325.
- Felten, E. W., Raj, M., and Seamans, R. (2018). A Method to Link Advances in Artificial Intelligence to Occupational Abilities. *AEA Papers and Proceedings*, 108:54–57.
- Firpo, S., Fortin, N. M., and Lemieux, T. (2011). Occupational Tasks and Changes in the Wage Structure. *IZA Discussion Paper Series*, (5542).
- Gathmann, C. and Schoenberg, U. (2010). How General Is Human Capital? A Task-Based Approach. *Journal of Labor Economics*, 28(1):1–49.
- Gibbons, R., Katz, L. F., Lemieux, T., and Parent, D. (2005). Comparative Advantage, Learning, and Sectoral Wage Determination. *Journal of Labor Economics*, 23(4):681–723.

- Goldin, C. and Katz, L. F. (2009). The Race between Education and Technology: The Evolution of U.S. Educational Wage Differentials, 1890 to 2005.
- Goldin, C. D. and Katz, L. F. (2008). The Race between Education and Technology: The Evolution of U.S. Educational Wage Differentials, 1890 to 2005.
- Goos, M., Manning, A., and Salomons, A. (2014). Explaining Job Polarization: Routine-Biased Technological Change and Offshoring. *American Economic Review*, 104(8):2509–2526.
- Gray, M. L. and Suri, S. (2019). *Ghost work : how to stop Silicon Valley from building a new global underclass*. Boston : Houghton Mifflin Harcourt, 2019.
- Hanushek, E. A., Schwerdt, G., Wiederhold, S., and Woessmann, L. (2015). Returns to skills around the world: Evidence from PIAAC. *European Economic Review*, 73:103–130.
- Heckman, J. J., Lochner, L. J., and Todd, P. E. (2003). Fifty Years of Mincer Earnings Regressions. *NBER Working Paper*.
- Hershbein, B. and Kahn, L. B. (2018). Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings. *American Economic Review*, 108(7):1737–1772.
- Horton, J. J. and Tambe, P. (2019). The Death of a Technical Skill. *Working Paper*.
- Katz, L. and Autor, D. (1999). Changes in the Wage Structure and Earnings Inequality. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics*, vol. 3A, pages 1463–1555.
- Lemieux, T. (2006). The 'Mincer Equation' Thirty Years after Schooling, Experience, and Earnings. In Grossbard, S., editor, *Jacob Mincer: A Pioneer of Modern Labor Economics*, pages 127–145. New York:, U British Columbia.
- Michaels, G., Natraj, A., and Van Reenen, J. (2014). Has ICT Polarized Skill Demand? Evidence from Eleven Countries over Twenty-Five Years. *Review of Economics and Statistics*, 96(1):60–77.

- Mincer, J. (1958). Investment in Human Capital and Personal Income Distribution. *Journal of Political Economy*, 66(4):281–302.
- Mincer, J. (1974). *Schooling, experience, and earnings*. Human behavior and social institutions: 2. New York, National Bureau of Economic Research; distributed by Columbia University Press, 1974.
- Restrepo, P. (2015). Skill Mismatch and Structural Unemployment. *Working Paper*.
- Spitz-Oener, A. (2006). Technical Change, Job Tasks, and Rising Educational Demands: Looking outside the Wage Structure. *Journal of Labor Economics*, 24(2):235–270.
- Webb, M. (2019). The Impact of Artificial Intelligence on the Labor Market. *NBER Working Paper*.
- Welch, F. (1969). Linear Synthesis of Skill Distribution. *Journal of Human Resources*, 4(3):311–327.