Do Labor Demand Shifts Occur Within Firms or Across Them? Non-Routine Biased Technological Change, 2000-2016

Seth G. Benzell[*], Guillermo Lagarda[†], and Daniel Rock[‡§]

July 16, 2019

Abstract

A large literature has documented occupational shifts in the US away from routine intensive tasks. Theories of skill-biased technological change differ in whether they predict changes in occupational mix within firms, or merely across different firms or industries. Using LinkedIn resume records, BLS OES data, and Compustat employee counts, we estimate occupational employment for publicly traded US firms from 2000 through 2016. We find that faster employment growth among firms that disproportionately employ non-routine workers is the most important cause of SBTC, followed by within firm occupational mix rebalancing. The entry of new firms also plays a role, although firm exit is slightly routine-worker biased. R&D leads firms to have a larger share of routine workers. These results are most consistent with a theory of routine task demand reduction caused by the diffusion of infra-marginally implemented new technologies. We also introduce a new measure of business labor dynamism, capturing the frequency with which firms change their occupational mix. Consistent with trends in productivity and other measures of business and labor market dynamism, this measure has decreased steadily since 2000.

1 Introduction

A large literature has measured wide-scale changes in the composition of US employment. These changes have been biased against occupations that are intensive in routine tasks. Partly for this reason, many have attributed these skill-biased employment shifts to technological change. This paper distinguishes between changes in occupational due to employment changes across and within

[*]MIT Initiative on the Digital Economy; sbenzell@mit.edu

[†]Boston University and Inter-American Development Bank

[‡]MIT Initiative on the Digital Economy and MIT Sloan School of Management

[§]We were provided access to de-identified LinkedIn profile information as participants in LinkedIn's Economic Graph Research program. We thank Brian Lucking, Nicholas Bloom, and John Van Reenen for assisting us with their data. We would like to thank David Autor, Jacqueline Barrett, Erik Brynjolfsson, Pascual Restrepo, Sebastian Steffen, the LinkedIn Economic Graph Research team, and the seminar participants at the MIT Sloan Organizational Economics Lunch workshop for their very helpful comments.

firms, and measures how skill-biased technological change (SBTC) is related to those changes and firm-level investments in R&D.

An occupation, e.g. business analysts or truck drivers, can increase in relative employment share for several reasons. One way this can happen is if firms that have a disproportionate amount of such workers grow faster than other firms. We call this the *across* firm heterogeneous growth effect. Another possibility is that some workers, like data entry clerks, are being fired and replaced with data analysts. We call this the *within* firm rebalancing effect. A final pair of possibilities is that new firms are created that disproportionately employ data analysts or old firms go out of business that disproportionately employed other kinds of workers.

In this paper we measure the relative importance of these components in job creation and destruction. To generate these new granular measures we rely on several data sets. The most important of these is LinkedIn profile records. These records, when combined with information from the Bureau of Labor Statistics, allow us to estimate occupational employment at the firm-year level. Our dataset covers publicly traded US businesses from 2000 through 2016. Over this interval, US public firms added 10.6 million jobs. The firms in our data had 27.0 million employees in 2000, or about 18 percent of the US labor force BLS (2018).

Previous research has found faster employment growth for the highest paid and least routine-intensive occupations (Acemoglu and Autor, 2011). Therefore in our analyses we group occupations by routine-task intensity and initial wage (in 2000). Consistent with this research, we find that the net new employment created over this interval is highly biased against routine labor. 2.95 million more jobs were created in the bottom third of routineness than in the top third. The most important source of this SBTC was was faster growth across firms with high levels of non-routine workers, accounting for 48.9 percent of this bias. The next most important source of this SBTC was within firm rebalancing, which accounts for 30.0 percent of the bias in public firms. New firms also had relatively low levels of routine workers, and accounting for 25.7 percent of the total change. The exit of firms worked slightly against SBTC, explaining -4.6 percent of the trend.

This decomposition is important for better understanding the causes of automation. Leading theories of automation differ in their predictions about whether skill-biased technical change occurs within or across firms. Direct automation (e.g. Benzell et al. (2016)) is the simplest story, i.e. complete replacement of human labor with machines. After acquiring new technologies, such as industrial robots, companies simply fire their automated workers. Other models of interacting industry-level productivity or globalization and aggregate demand suggest that SBTC is caused by a glut of manufactured goods (e.g. Bessen (2018)). Under this theory, routine workers are disproportionately impacted by firings simply because they are concentrated in certain industries. Other theories involving "superstar" firms dominating an industry naturally, such as Autor et al. (2017), lead to a shift toward the occupations and capital varieties most heavily employed in the superstar production function. In these models, within each industry, some new firms or particularly savvy old ones (which are attuned to the new digital economy and have the right mix of non-routine workers to take advantage of it) grow faster and hire more of the workers they always

have. Matching the empirical decomposition we document here is a potential desideratum for future models of automation.

In our framework section, we show that firm TFP-driven changes in aggregate routine employment would tend to produce offsetting within-firm employment mix changes in favor of routine workers. Intuitively, this is because if demand for non-routine workers goes up their wages will go up as well all else constant. All types of firms reduce the share of non-routine workers in employment. Therefore, the fact that within-firm rebalancing is a significant cause of the increase in non-routine employment indicates that production function changes within the firm is a main cause of SBTC.

Which, or to what extent each, theory is true is of much more than academic interest. Some government policies have the explicit goal of preventing technological unemployment or softening the blow of automation on routine workers. Whether these policies will be effective will depend in part on the margins where SBTC occurs. Import restrictions on manufactured products or industrial subsidies for manufactured products will fail if routine workers are increasingly not employed by these firms. Discouraging the exit of firms (either through bailouts or bankruptcy law) or encouraging entrepreneurship will help only insofar as the relevant firms employ routine workers. The Tax Cuts and Jobs Act of 2017, for example, promised to raise the salaries of Rust Belt workers by boosting capital and R&D investment. But if firms that make larger capital investments reduce their demand for routine workers, then this is a dubious solution. Effective policy response is a steep challenge without knowledge of the locus of SBTC or being able to test mechanisms responsible for the shift away from routine workers. This paper provides some description of the different drivers of occupational compositional shifts in the new millennium.

The location of SBTC also has consequences for the long-term future. Theories of directed technological change attempt to understand how economic conditions lead to and are created by innovations. Some theories, such as Peretto and Seater (2013) feature firms which can invest in R&D to change their production functions. In these models, the decision for a firm to develop a skill-biased technology is a function of the prices it faces. The intended consequence of these innovations is to raise demand for workers with low-paid skills. Others, such as Acemoglu (2002) assume that specialized companies create technologies that all other firms implement. In models of that type, technology creators are also sensitive to the relative amounts of workers of each type who might use their tools. The distinction is important because while innovation of the first type will tend to focus on finding new uses for cheap factors, the innovation of the latter type will focus on abundant factors. Under the latter model, higher paid workers may disproportionately benefit from innovation so long as they are sufficiently abundant. We find strong evidence that firms which make additional R&D investments due to tax incentives are more likely to have routine workers.

We also introduce a new measure of within-firm occupational change. Along several dimensions, business dynamism in the United States has been on the decline. From 2000 to 2011, the share of firms less than 5 years old declined about 5 percentage points. Over the same interval, the rates of job creation and destruction decreased from approximately 17 and 14 percent respectively to

13.5 and 12 percent (Decker et al., 2014). We construct a complementary measure, which we call firm employment mix dynamism. This measure tracks how much firms change their occupational mix. We find that firm employment mix dynamism has decreased steadily since 2000. These results indicate that the decrease in the rate of job creation and destruction corresponds in part to a reduction in the rate at which firms' skill demands change, rather than reflecting only churn in employees within a type. Of course, it may also be that firms have better information about the types of labor they will need to hire. The declining within-firm labor dynamism may reflect a more stable business environment.

2 Related Literature and Framework

What is technology's role in shifting labor demand? Early studies of SBTC focused on the education wage premium. The share of workers with a high school or college education increased dramatically in the latter half of the 20th century. Yet, across nations, the wage premium for the educated stayed constant or increased over this interval (Berman et al., 1998). While other contributing factors have been proposed, such as increased globalization (Acemoglu et al., 2016) and decreased unionization Western and Rosenfeld (2011), the consensus view is that a leading cause of this phenomena is technological change (Johnson, 1997).

More recent papers have diagnosed labor demand *polarization* as the cause. From 1980 to 2005, occupations which were highly compensated in 1980 saw disproportionate growth in both wage and employment. The same is true of occupations compensated poorly. Other occupations saw little employment or wage growth. Autor and Dorn (2013) find that areas that specialized historically in industries which use routine tasks intensively (such as manufacturing) saw larger increases in wage and employment polarization. This finding remains after controlling for the offshorability of jobs. They follow Autor et al. (2003) in attributing this to technological advances, in particular in information technology, which tend to substitute for people in routine jobs. Other papers looking at the role of technological change in wage polarization in developed countries are Acemoglu (1999), Goos and Manning (2007), and Goos et al. (2010, 2014).

There is also some industry and small scale firm-level evidence supporting this hypothesis. Bresnahan et al. (2002) and Berman et al. (1994) provide firm and industry-level evidence respectively of IT and R&D having skill-biased impacts on labor demand. Barth et al. (2017) show that manufacturing establishments with a higher proportion of scientists and engineers pay higher wages.

There is also theoretical and empirical evidence of technology playing a role in the rise of the 1%'s share of income. This trend is is driven by labor earnings, not increasing capital income of the wealthy. The increase in top income shares has impacted higher earners of all types Kaplan and Rauh (2013). Rosen (1981) presciently forecasted that economies of scale enabled by new technologies would increase inequality. Innovations in, for example, telecommunications lead more tasks to be winner-take-all where gains might have been more evenly distributed in the past.

Theories of skill-biased employment demand change, whether due to offshoring, technology,

changing preferences, or some other mechanism, can be broadly divided into a handful of types. They can be caused by changes in the industrial composition of employment, by the heterogeneous growth of firms with unusual employment mixes, by firms changing their mix of employees, or through the entry or exit of firms with unusual employment mixes.

2.1 Framework

To understand how different sources of skill-biased technological change can be distinguished, consider the following model of employment.

The economy makes two commodities: manufactured products (M) and services (S). Each is produced in a perfectly competitive industry. The industries contain I_M and I_S firms respectively. Each firm's production is constant elasticity of substitution (CES) in routine and non-routine labor. Within industries, each firm has an identical production function. We have

$$M = \sum_{i=1}^{I_M} A_M^i (z_{R,M}^{\frac{1}{\sigma}} (L_{R,M})^{\frac{\sigma-1}{\sigma}} + (L_{NR,M})^{\frac{\sigma-1}{\sigma}})^{\frac{\sigma}{\sigma-1}} \tag{1}$$

$$S = \sum_{j=1}^{J_S} A_S^j (z_{R,S}^{\frac{1}{\sigma}} (L_{R,S})^{\frac{\sigma-1}{\sigma}} + (L_{NR,S})^{\frac{\sigma-1}{\sigma}})^{\frac{\sigma}{\sigma-1}} \tag{2}$$

The economy has a fixed total endowment of labor $\overline{L}$. This can be divided into routine labor L_R and non-routine labor L_{NR}, into employment in both sectors E_M and E_S, or by labor and sector (i.e. $L_{R,M}, L_{NR,M}, L_{R,S}, L_{NR,S}$). Total time of the representative household is split between the two possible occupations and industries.

$$\overline{L} = L_R + L_{NR} = E_M + E_S = L_{R,M} + L_{NR,M} + L_{R,S} + L_{NR,S} \tag{3}$$

The model has two periods. In the first period, workers are mobile between firms and jobs such that the wage in all occupations is identical. In the second period, in anticipation of a technological change, workers can move between occupations by paying a re-skilling cost C, which is increasing and convex in the quantity of labor that switches occupations. Welfare for the representative household in the second period is equal to total wage less these re-skilling costs.

$$U = w_R L_R + w_{NR} L_{NR} - C(|L_R - L_R'|) \tag{4}$$

where L_R' is the wage equalizing share of workers in the routine task before the technological change. Workers will move between occupations so as to maximize their total second period wage less re-skilling costs.

$$\frac{\partial C}{\partial L_R} = |w_R - w_{NR}| \tag{5}$$

which is perfectly mobile between occupations and industries.

There is no specification for the form of aggregate demand. Along that dimension, this framework is very general. We now have enough of a framework to begin analyzing how different types of technological change will influence employment across and within firms.

2.2 Total Factor Productivity Changes Across Firms

Suppose that firms in industry M get more productive (i.e. $A_M j$ increases for some or all firms). Total employment of routine workers can be written as

$$L_R = E_M \frac{L_{R,M}}{E_M} + E_S \frac{L_{R,S}}{E_S} \tag{6}$$

where E_M and E_S are total employment in the manufacturing and service industry respectively. Total employment is equal to the number of employees working in each type of firm times the share of workers in each type of firm who are routine.

The total effect of a technological change can be understood as the sum of a change in the share of workers employed *across* different firms, and a change in the mix in the share of workers who are routine *within* different types of firms.

Consider the effect of a technological change which boosts the TFP of some or all manufacturing firms. This could either be a global effect, such as due to a trade shock, or because a new technology raises the total factor productivity of new (or potential new) firms. This change will have an effect on the total share of workers doing routine work through two mechanisms.

The total effect of an increase in the productivity of manufacturing firms can be written as

$$\frac{\partial L_R}{\partial A_M} = \frac{\frac{\partial E_M}{\partial A_M}(\frac{L'_{R,M}}{E'_M} - \frac{L'_{R,S}}{E'_S})}{1 - \frac{\partial \frac{w_R}{w_{NR}}}{\partial L_R}(E'_M \frac{\partial \frac{L_{R,M}}{E_M}}{\partial \frac{w_R}{w_{NR}}} + E'_S \frac{\partial \frac{L_{R,S}}{E_S}}{\partial \frac{w_R}{w_{NR}}})} \tag{7}$$

The numerator of (7) summarizes the across firm effect. When manufacturing firms become more productive, this can increase or decrease the share of workers in that industry. If the products of the two industries are gross complements, the increase in manufacturing firm productivity will decrease the amount of workers needed in those firms (i.e. $0 \geq \frac{\partial E_M}{\partial A_M}$). By assumption manufacturing firms have a higher initial share of workers performing routine tasks ($i.e. \frac{L'_{R,M}}{E'_M} \geq \frac{L'_{R,S}}{E'_S}$. In this case, the across firm effect is to decrease the total amount of routine workers through reducing the share of employment in firms which are initially routine intensive.

The denominator summarizes the within firm effect. A decrease in total relative demand for routine workers will decrease the relative wage of routine workers and vice versa (i.e. $\frac{\partial \frac{w_R}{w_{NR}}}{\partial L_R} \geq 0$). If the relative wage of routine workers increases, then their share of employment in both industries will decrease (i.e. $0 \geq \frac{\partial \frac{L_{R,S}}{E_S}}{\partial \frac{w_R}{w_{NR}}}$ and $0 \geq \frac{\partial \frac{L_{R,M}}{E_M}}{\partial \frac{w_R}{w_{NR}}}$) with the relative importance of both decreases weighed by initial employment in both industries.

Notably, this denominator will always be greater than or equal to one. In other words, any firm level TFP change that leads to a decrease in routine employment will be partially offset by

wage decreases for routine workers. This will lead to within-firm rebalancing which favors routine workers.

If workers are perfectly mobile between occupations (i.e. $C = 0 \; \forall \; L_R$) then the denominator will be 1, because $\frac{\partial \frac{w_R}{w_{NR}}}{\partial L_R} = 0$. In other words, in the special case with perfect labor mobility, (7) reduces to just the numerator. Similarly, if firms are Leontieff in production, and do not change their labor mix as a function of changes in wages (i.e. $\frac{\partial \frac{L_{R,S}}{E_S}}{\partial \frac{w_R}{w_{NR}}} = 0$ and $0 = \frac{\partial \frac{L_{R,M}}{E_M}}{\partial \frac{w_R}{w_{NR}}}$) then the denominator will also equal one. In either case, the effect of the TFP change on total routine employment will be entirely driven between reallocation in employment across firms. In these cases the equation (7) reduces to just the numerator.

Consider the special case where the technological change decreases employment in industry (i.e. $\frac{\partial E_M}{\partial A_M} \leq 0$ and production is Leontieff. Intuitively, in that case the technological change will reduce routine employment so long as manufacturing industry is more routine skill intensive (i.e. $z_M \leq z_S$). Then

$$M = \sum_{i=1}^{I_M} A_M^i [min(z_{R,M} L_{R,M,i}, L_{NR,M,i})] \tag{8}$$

$$S = \sum_{j=1}^{I_S} A_S^j [min(z_{R,S} L_{R,S,j}, L_{NR,S,j})] \tag{9}$$

and the total change in routine employment is only the across firm effect, so the total change in routine employment is

$$\frac{\partial (L_{R,M} + L_{R,S})}{\partial A_M} = \frac{\partial (L_{R,M} + L_{NR,M})}{\partial A_M} (\frac{1}{z_M + 1} - \frac{1}{z_S + 1}) \tag{10}$$

The first effect is through aggregate demand.

Two examples of models with this form of technological change are Bessen (2018) and Autor et al. (2017). In Autor et al. (2017)'s model of superstar firms, increasing industrial concentration explains decreases in labor's share of income. As globalization proceeds, the most productive firms, which tend to have low labor shares, capture an increasing percentage of the market. If this is the correct model of a decrease in labor share, then it could also be a cause of SBTC. If the most productive firms have distinctive employee mixes, then as they disproportionately expand, the occupational composition of the economy will change as well.

In Bessen (2018), reductions in employment for factory workers may be due to TFP changes in the manufacturing industry. That industry disproportionately employs routine manual workers and has seen both increasing import competition and technological progress.

Under either of these hypotheses, the most important source of SBTC will be across-firm growth. Due to the productivity shock, relatively routine-intensive firms will decrease in employment and vice versa. Some of SBTC may also be caused by entry or exit of firms, depending on whether all firms get the TFP shock, or only new ones. However, this form of SBTC is either firm employment

mix neutral (if, for example, within firm production is Leontieff), or will be biased *against* non-routine employment. In general the total employment shift towards non-routine employment will be mitigated by within firm employment mix rebalancing.

In the model, there are three factors leading firms to change their employment mix. Shifts in productivity can directly lead to either more or less hiring depending on price elasticity of demand in the product market. If the two industries have different occupational compositions, this will change the overall composition of employment. The second effect is the direct effect of task specific productivity changes. This leads firms to change their occupation mix. Finally there is an indirect effects through wages. If demand for a certain occupation changes, this can lead to changes in relative wages. Shifts in relative wages causes firms to substitute toward cheaper labor.

2.3 Skill-Biased Technological Changes Within Firms

The other type of technological change possible in this framework is due to a change in Leontieff production functions. This will lead to SBTC within firms in addition to between them.

Suppose that a new technology is invented to automate some routine tasks in the manufacturing industry (i.e. z_M increases). We continue to restrict attention to the special case where production within firms is Leontieff. Then the overall impact on routine employment will be

$$\frac{\partial(L_{R,M}+L_{R,S})}{\partial z_M} = -\frac{L_{R,M}+L_{NR,M}}{(z_M+1)^2} + \frac{\partial(L_{R,M}+L_{NR,M})}{\partial z_M}\left(\frac{1}{z_M+1} - \frac{1}{z_S+1}\right) \tag{11}$$

This equation has two terms. The first term is the direct impact of the automation on routine employment in the manufacturing industry. The second term due to any movement between industries. It will be zero if the two industries have the same initial mix of routine and manual workers (i.e. $z_M = z_S$). The second term will also be zero if the productivity effect of easier production of manufactured goods perfectly cancels the price effect (i.e. $\frac{(\partial L_{R,M}+L_{NR,M})}{\partial z_M}$). This would be the case if aggregate demand were Cobb-Douglas.

Examples of models with technological change analogous to this type include Peretto and Seater (2013), Benzell et al. (2016), and Acemoglu and Restrepo (2017). In these models firms can make investments, either in capital or R&D, to reduce their demand for certain types of workers. If this is the correct model of SBTC, we should expect to see occupational changes being driven by firms changing their employment mixes. Depending on the strength of the second term, we would still expect to see a role for the other sources. Additionally, given the key role for firm capital and R&D investment in these models, we would expect to see firms with larger stocks of these investments see more dramatic within firm employment mix rebalancing. While the Leontieff case demonstrates the dynamics under a simple assumption of intense complementarity between labor input types, much of the intuition for the change in labor demand remains relevant under less restrictive assumptions.

3 Data and Methodology

Our dataset is created by merging LinkedIn's database of hundreds of millions of position records with Compustat NA firm financial data, occupation-level BLS OES wage and employment data, and O*NET occupational task data. Employment by firm and occupation is constructed from individual LinkedIn resumes and aggregated to 136 occupational groups. BLS occupations are many to one matched to LinkedIn occupation groups. For LinkedIn occupations nesting several BLS occupations, the weighted (by employment) average characteristics (such as wage) of the occupations are attributed to the LinkedIn occupation.

Individuals employed by some firms and with certain occupations are disproportionately likely to have LinkedIn accounts. We normalize the LinkedIn employment measures to account for both sources of bias.

As in Brynjolfsson et al. (2018b), we use total firm-level employment from Compustat and industry-level occupational composition from BLS. First, we sum firm-level employment in Compustat to measure total public firm employment by (3-digit NAICS) industry.[1]

Starting with the amount of employees we observe within each industry, we use the BLS's estimates of occupational mix by industry to construct the total amount of employees of each occupation in our sample. Aggregating across firms, we then generate an estimate of the total quantities of workers by year in each of the specific occupations in the BLS Occupational Employment Survey. We then estimate LinkedIn coverage by occupation-year. Each firm-occupation-year tuple is deflated (or inflated) for its coverage relative to our Compustat/BLS estimate. Subsequently, we re-adjust our Compustat/BLS measures such that total firm employment is equal to total firm employment in Compustat. This process maps what we observe in the BLS estimates and the LinkedIn employment estimates by year, occupation, and firm to an estimate that adjusts for occupation and firm-specific differences in LinkedIn coverage. This gives us our measure of the occupational labor force within each firm-year.

Firm and year-level measures of firm capital are taken from Compustat and R&D stocks are taken from Peters and Taylor (2017), available via WRDS. R&D stocks are measured as a perpetual inventory using the BEA's industry specific depreciation rate. We extend the Peters and Taylor R&D stocks to 2016, following Brynjolfsson et al. (2018a). We estimate a firm level wage bill by multiplying our measure of each firm's employment by occupation by the BLS mean annual wage.

Our final dataset consists of over 2.5 million firm-year occupation tuples, containing 3662 firms, with an average of 14,687 employees.

We organize occupations and firms into categories for convenient analysis. For occupations, we sort occupations into approximately equal, in terms of year 2000 US publicly traded firm employment, thirds. The first organization is in terms of the routineness of the task. Following Acemoglu and Autor (2011) routineness is measured as the sum of three O*NET questions in 2006 (the ear-

[1] In the case that a firm is missing an accurate *emp* entry, we substitute in a prediction estimate from a regression of *emp* on the number of individuals who have a record of working at the company on LinkedIn, with industry and year fixed effects.

liest year available): (4.C.3.d.3) Pace determined by speed of equipment; (4.A.3.a.3) Controlling machines and processes; (4.C.2.d.1.i) Spend time making repetitive motions. We also organize occupations into equal employment thirds based on year 2000 mean wage.

One way firms are classified is by whether they are in our sample and have positive employees in a given year. If a firm has positive employees in our sample but then leaves the sample before 2016, we say that firm exited. Firms can exit our sample through being bought out, through being de-listed, or through bankruptcy. Firms that have no record of employees, but gained them before 2016 are labeled as entering. We also classify firms into quintiles by estimated wage bill growth over the interval.

4 Net Job Creation Within and Across Firms

Figures 1 and 2 display total net employment growth, and its components, for selected occupations for firms in our sample from 2000 to 2016.

In both figures, the blue bar displays the total net number of jobs created. This is the sum of the next four bars. As figure 1 shows, about one hundred and six thousand data entry clerk jobs were eliminated in our sample over this interval.

The green and red bars in each figure distinguish the share of employment changes driven by heterogeneous growth across firms and within firm employment rebalancing respectively. The green bar (across firm growth) indicates how many net jobs would be added if every firm kept their shares of employment constant between 2016 and 2000.

For occupation group $\hat{j}$, the across firm growth effect is defined as

$$G_{\hat{j}} = \sum_{i=1}^{I} \left(Emp_{i,\hat{j},2016} - \frac{Emp_{i,\hat{j},2016}}{\sum_{j=1}^{J} Emp_{i,j,2016}} \sum_{j=1}^{J} Emp_{i,j,2000} \right) \tag{12}$$

where $Emp_{i,j,t}$ is the number of employees at firm i, of type j in year t. J is the total amount of occupation groups and I the total amount of firms. The green bars would all be of equal heights if every firm had the same employee mix in 2016, or if all firms added employees at the same rate since 2000.

The red bars indicate the share of employment growth due to firms that exist in both 2000 and 2016 changing their employment mixes.

We define the within firm rebalancing effect as

$$W_{\hat{j}} = \sum_{i=1}^{I} \left(Emp_{i,\hat{j},2016} - Emp_{i,\hat{j},2000} \right) - G_{\hat{j}} \tag{13}$$

A firm can only rebalance towards one occupation if it moves away from another, so rebalancing changes sum to zero.[2]

[2] In this draft of the paper within-firm rebalancing does not perfectly add to zero because a handful of occupations that did not exist in 2000 were not assigned routineness-third categories. This will be resolved in future drafts.

The final source of net job growth is firm entry and exit. The teal bar indicates the number of employees of a type in 2000 at firms that leave the sample before 2016. The orange bar indicates the number of employees in 2016 of a type at firms which enter the data after 2000. Exit of firms can only eliminate jobs, and entry can only add them.

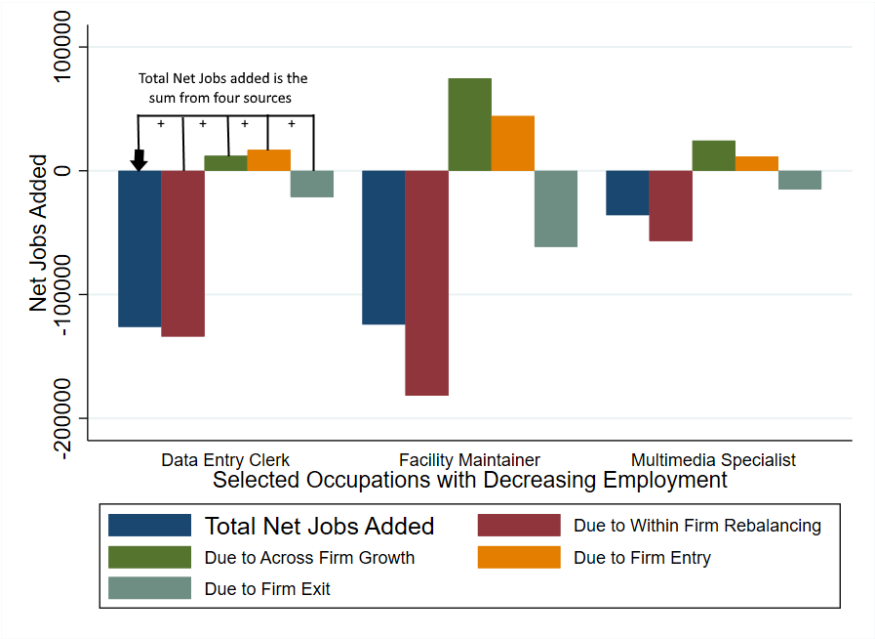

Figure 1: Net job creation by source for US publicly traded corporations from 2000 to 2016. Selected occupations

As shown in figure 1, data entry clerks, facility maintainers, and multimedia specialists lost a net of 126 thousand, 124 thousand, and 36 thousand positions. Data entry clerk is an occupation in the middle third of routineness, and multimedia specialists and facility maintainers are in the top third of routineness. In all of these occupations, the total amount of jobs lost is less than the amount lost due to firings and replacement. For data entry clerks, the amount of jobs lost due to rebalancing is approximately equal to the total amount of jobs lost. This is because firms that employed these types of workers did not grow much, nor did many more firms that hired such workers enter on net.

On the other hand, firms that hire facilities maintainers and multimedia specialists increased their employment somewhat more. Therefore, total job loss for these occupations was more limited. Across all three of these occupations, entry and exit created very few jobs on net. Note that none of these occupations saw net reductions in employment due to firms shrinking in employment.

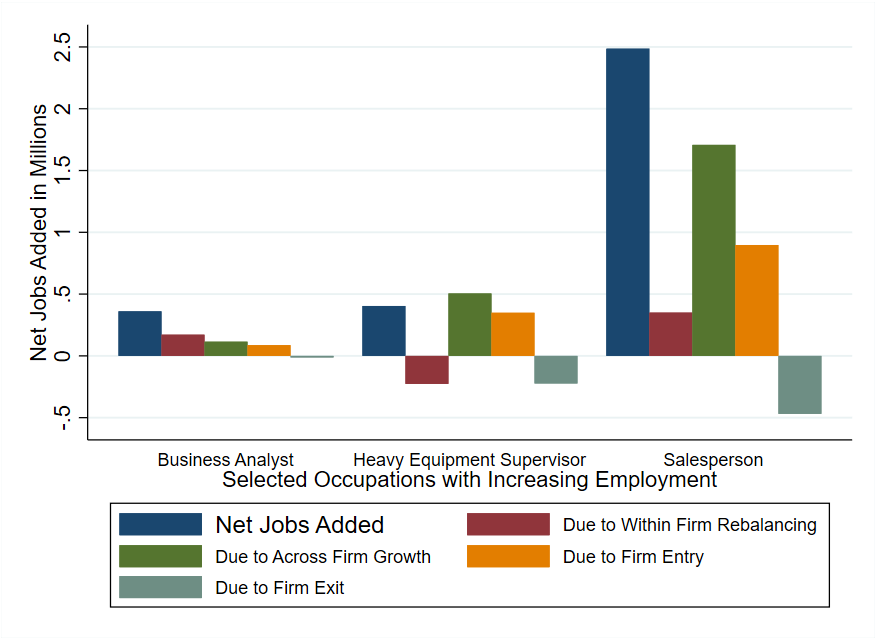

Figure 2: Net job creation by source for US publicly traded corporations from 2000 to 2016. Selected occupations

Figure 2 repeats this decomposition for three of the fastest growing occupations. Salesperson is the most common occupation in the data, and also the fastest growing. The growth in this occupation is almost entirely due to across firm growth. Heavy equipment supervisors also saw growth in employment despite some within firm rebalancing against them. Business analysts also saw employment increase, in roughly equal measure due to across firm growth, within firm rebalancing, and net entry.

Considering the figures as a pair, it is clear that all components of net job creation need not go in the same direction. However, it does seem to be the case that then number of net jobs created lost by shrinking occupations is closer to and more correlated with within firm rebalancing than for growing occupation. 3, reflects this relationship. This figure relates an occupation's employment change due to within firm rebalancing against its total change. As can be seen, the top 12 fastest growing occupations are scattered about far from the 45 degree line. The 12 fastest shrinking occupations are right on top of it. Also notable is the the green triangle in the upper left. This is Heavy Equipment Supervising. This is the only occupation of those selected with positive employment growth despite firms in the sample replacing them on net.

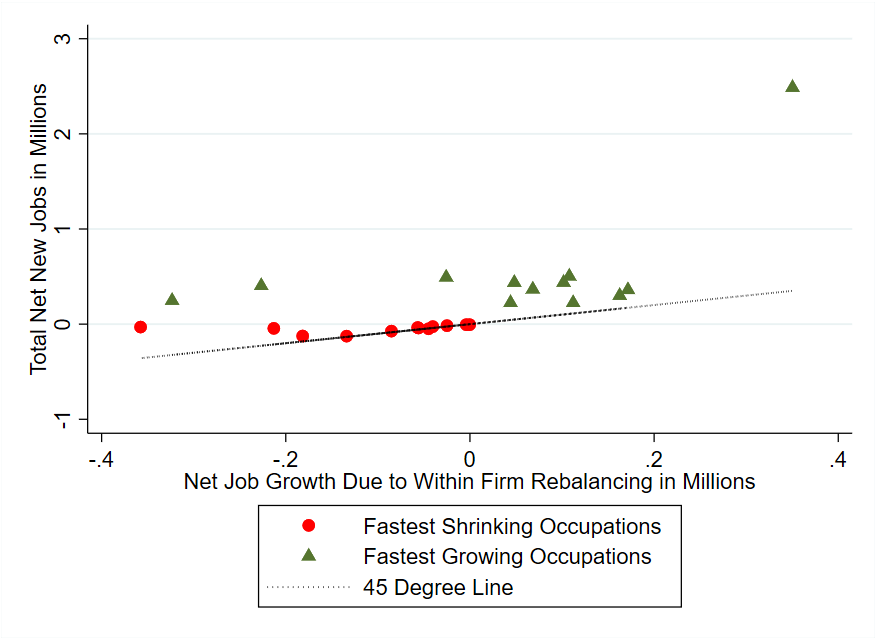

Figure 3: Plot of total net new jobs created and the portion due to within firm rebalancing for the 12 fastest growing and contracting occupations.

5 Skill Biased Technical Change Within and Across Firms

While it is informative to look at individual occupations, the best way to understand SBTC in the economy as a whole is to group occupations into categories. In figure 4 we present the same decomposition of net employment growth into its components with one key difference. Rather than examining individual occupations, this figure sorts all occupations into one of three categories by their routine task intensity. Occupations are sorted into three groups with approximately equal employment in 2000.

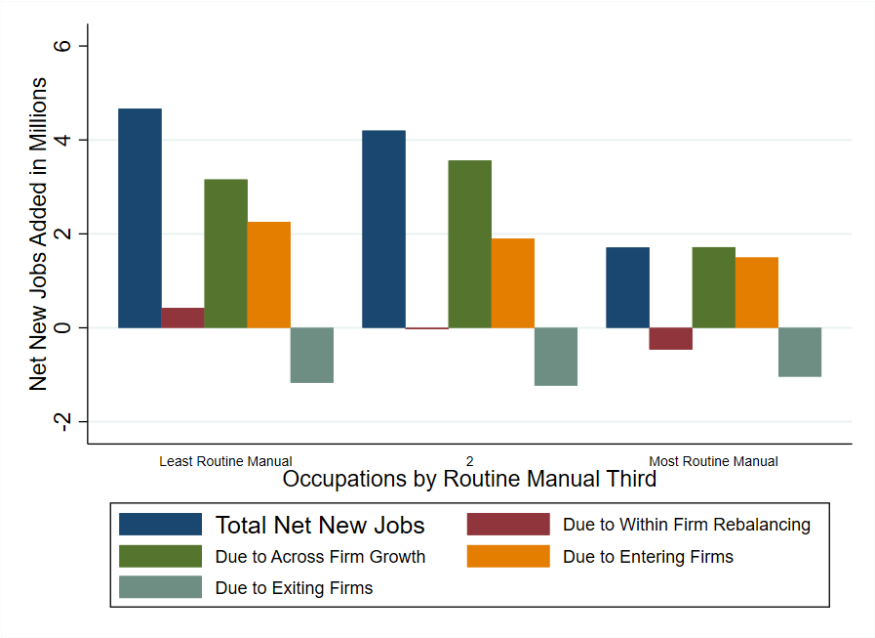

Figure 4: Net job creation by source for US publicly traded corporations from 2000 to 2016. Occupations organized into thirds based on routineness.

SBTC for firms in our sample over this interval was highly biased against manual tasks. As a measure of this, we define the non-routine bias of employment growth as the difference between bottom third routine occupation employment growth less top third routine occupation employment growth. So,

$$GB_R = (E_{R=1,2016} - E_{R=1,2000}) - (E_{R=3,2016} - E_{R=3,2000}) \tag{14}$$

We find that the non-routine bias of employment growth was 2.95 million over this interval. The most important source of this SBTC was was faster growth across firms with high levels of non-routine workers, accounting for 48.9 percent of this bias. The next most important source of this SBTC was within firm rebalancing, which accounts for 30.0 percent of the bias in public firms. New firms also had relatively low levels of routine workers, and accounting for 25.7 percent of the total change. The exit of firms worked slightly against SBTC, explaining -4.6 percent of the trend. These findings are consistent with the hypothesis that non-skill-neutral changes in the production function of firms are critical to to explaining aggregate trends in occupational employment.

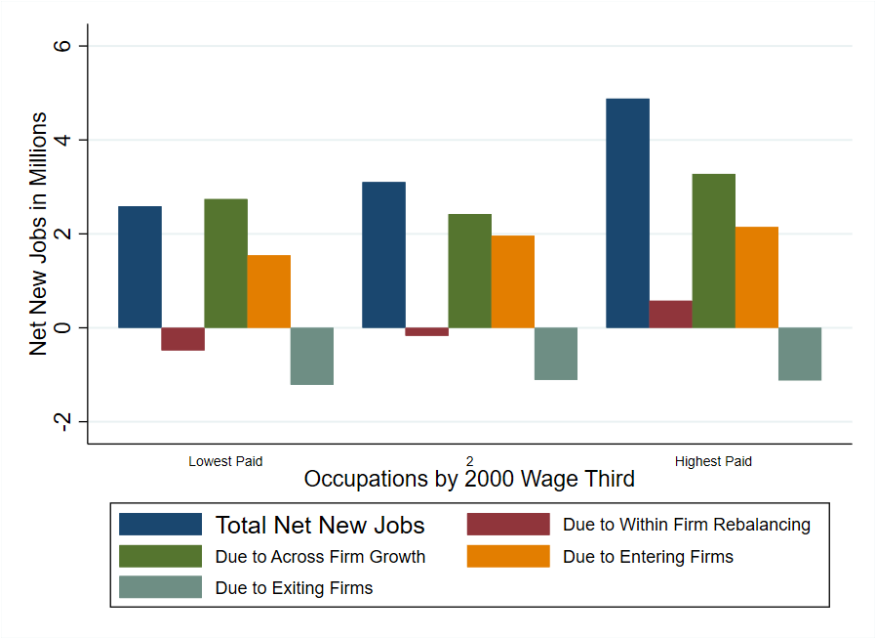

Figure 5: Net job creation by source for US publicly traded corporations from 2000 to 2016. Occupations organized into thirds based on 2000 average wage.

Figure 5 repeats this analysis, this time after dividing occupations into thirds based on their 2000 average wage. We find that total employment was not strongly high wage-biased. 2.02 million more jobs were created in the top third of occupations than the bottom. In contrast with Autor and Dorn (2013), we measure the change over this interval as being strictly skill-biased, rather than polarizing. First, by restricting our attention to publicly traded US firms over the interval 2000 to 2016, we are considering a somewhat different subset of the economy than previous studies. Second, the version of LinkedIn's occupation classification system we use can be somewhat coarse across levels of seniority for occupations with similar routine task intensities. This doesn't impact the routineness measure, but does cause us to lump together occupations of dissimilar year 2000 wage in some cases. For this reason throughout the paper we focus on organizing occupations by routineness. Still, the results we do have on the sources of SBTC when occupations are organized by wage are consistent with those on routineness. Namely, rebalancing within firms plays an important role.

6 Firm Employment Mix Dynamism

Given that within firm rebalancing is an important component of SBTC, it is natural to ask how the rate at which firms change their workers has changed across firms and time. We define firm employment mix dynamism as the average absolute value of the change in employment share for occupations in a firm. Or,

$$D_t = \frac{1}{136} \sum_{j=1}^{136} \left| \frac{E_{j,t}}{\overline{E_t}} - \frac{E_{j,t-1}}{\overline{E_{t-1}}} \right| \tag{15}$$

where $E_{j,t}$ is employment at a firm of a given occupation, t is the period (here, years) and $\overline{E_t}$ is total firm employment. Since we have 136 principal employment type groups, this give the average absolute value shift in employment share over all categories.

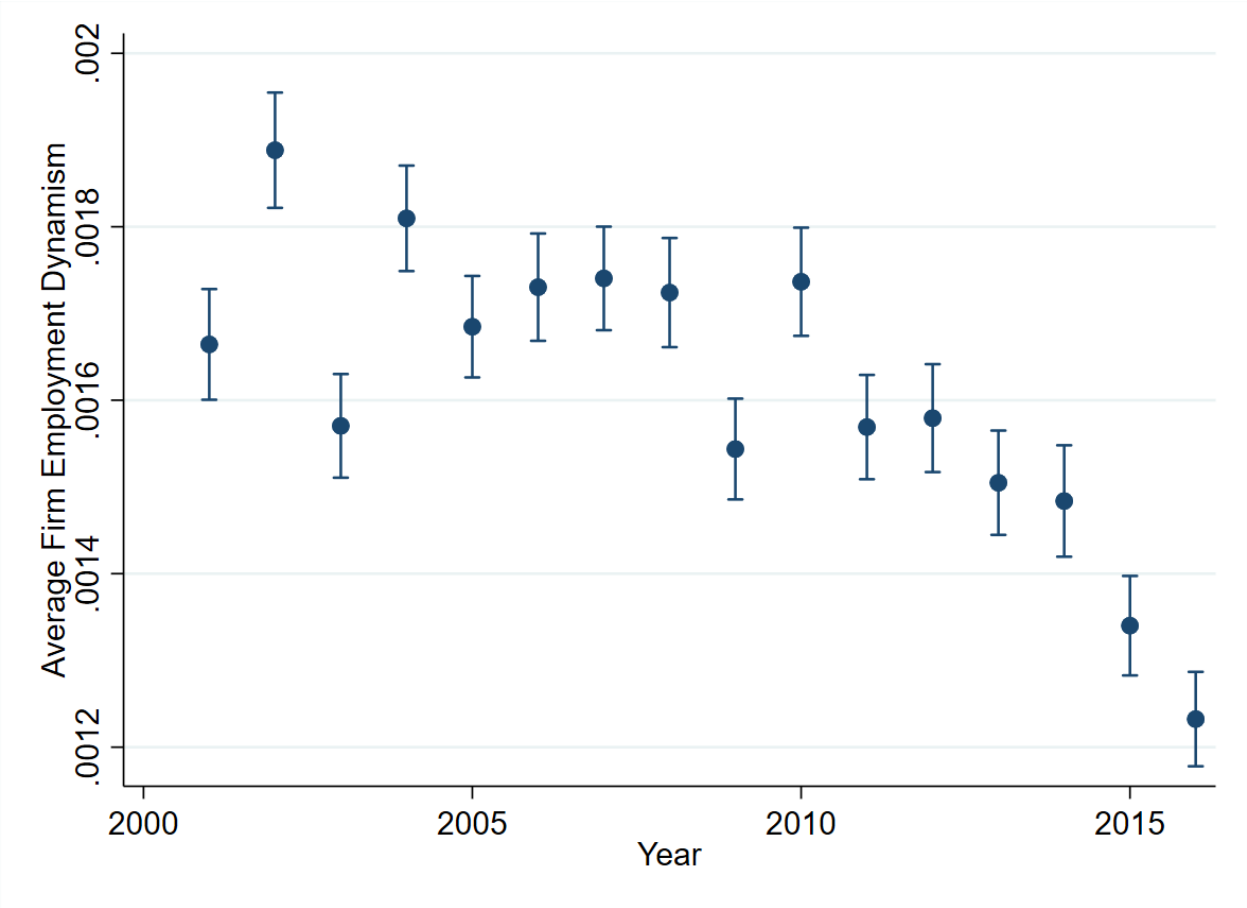

Figure 6: Average employment mix dynamism by year. 95 percent confidence intervals for sample means.

Figures 6 and 7 displays trends in average firm dynamism over time. Figure 6 treats all firms equally, while Figure 7 weighs firms by employment and includes total factor productivity growth for the entire US economy. In both figures there is no year 2000 observation, because both the year and the previous year's firm occupation mix need to be observed. For this reason also, firms must be in the data in consecutive years for an observation to appear. Other measures, summarized by (Decker et al., 2014), had shown that US firm dynamism has been on the decline. Our measure looking at labor mix shifts is consistent with these other measures.

Figure 7: Average employment mix dynamism weighted by firm employment and yearly percentage change in U.S. TFP

We measure a steady decrease in the frequency at which companies change their employees mix. As can also be seen in figure 7, weighing firms by employment when constructing this measure leads to a less dramatic decrease over time. During this period (from 2000 through 2016) the number of firms in the economy overall, as well as publicly traded firms in our sample, declined. The average firm gained more employees. Therefore, it is reasonable to ask whether the decrease in firm employment mix dynamism has been driven or counteracted by an increase in the average size of firms.

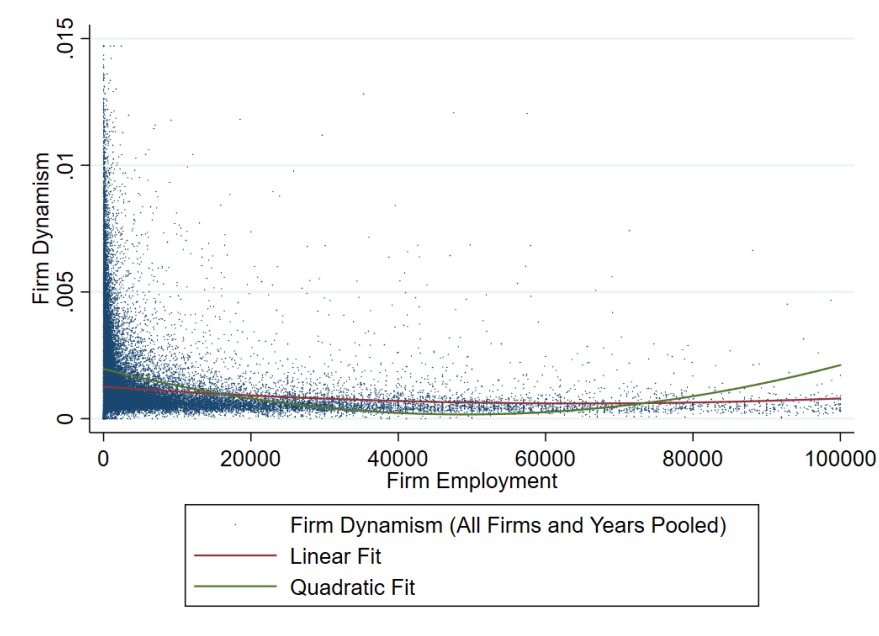

Figure 8: Firm employment mix dynamism and firm employment. All years. Firm-years with less than one hundred thousand employees.

Figure 8 presents the relationship between firm employment mix dynamism and firm employment, pooling all firms and years (with less than one hundred thousand employees). The range of firm variances is highest for very small firms. This is unsurprising because for smaller firms much fewer employees in absolute number need to be replaced (or promoted/demoted to a different occupation in the firm) in order to be measured as having high employment mix dynamism. For these very small firms measurement error also becomes important because each individual LinkedIn observation is weighted to represent several real employees. For large firms, the true and estimated number of employees by occupation is approximately the same. However, for very small firms, discreteness in the number of LinkedIn workers observed becomes important.

Despite this mass of high employment mix dynamism among very small firms, the measure shares no clear relationship overall with firm employment. When weighing by firm employment, a quadratically fit curve is almost perfectly flat. Future work will explore why some firms across all levels of employment have more dynamic labor mixtures.

7 Skill-Biased Technological Change Within and Across Firms by Industry

To deepen our analysis of within-firm rebalancing, across-firm growth, entry and exit as sources of skill-biased technological change, we repeat our analysis at the industry level. One hypothesis is that industry-level total factor productivity shifts are the most important cause of SBTC. If that is the case we should expect to see non-routine biased employment growth be the most dramatic

in industries that already had a relatively low number of routine workers. This would imply that industrial compositional changes are driving SBTC for the aggregate economy.

Figure 9 displays the initial non-routine bias, initial employment level, and non-routine bias of new jobs created for the 17 largest industries in our data. The initial non-routine bias of an industry is the number of employees in the bottom third of routineness less those in the top third of routineness in 2000.

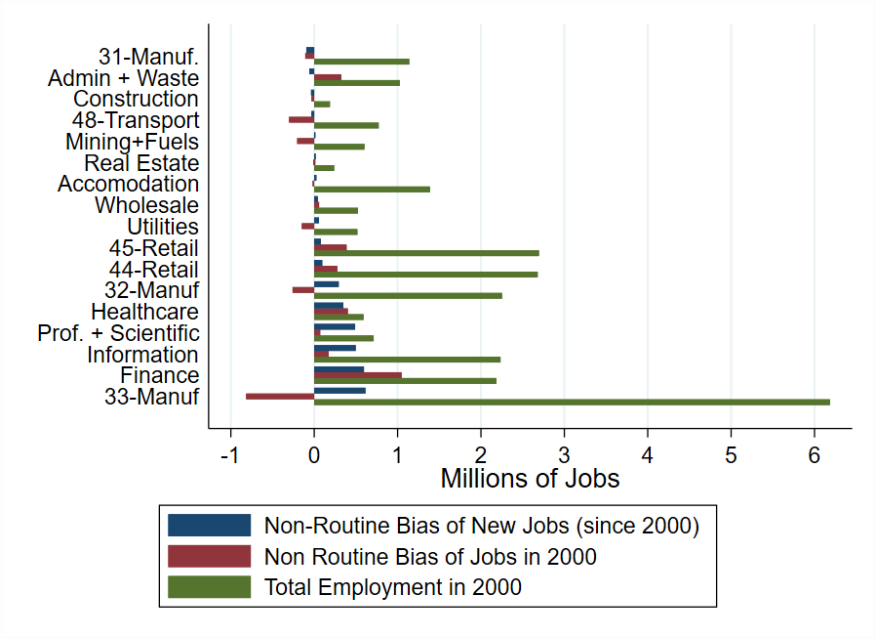

Figure 9: Total employment in 2000, Non-routine bias of employment in 2000, and non routine bias of net new jobs by industry. Largest 17 two-digit industries, sorted by non-routine bias of new job growth. The non-routine bias of employment in a given year is the difference between bottom third routine and top third routine occupation employment (i.e. $EB_R = (E_{R=1,t} - E_{R=3,t})$). The non-routine bias of employment growth is as defined above.

The industry with the most non-routine-biased employment growth is category 33 manufacturing firms, despite being strongly biased against non-routine labor in 2000. Net employment growth in that industry has been sufficient to switch the industry from employing more bottom-third than top-third routine employees to the opposite. The other industries with strong movement towards non-routine workers are mixed between those that were already strongly non-routine-biased to begin with (including Finance and Healthcare), those who were roughly neutral (including Information, Retail, and Professional and Scientific firms), and those that were biased against (32-Manufacturing). No industry got significantly more routine-biased as a result of new hirings.

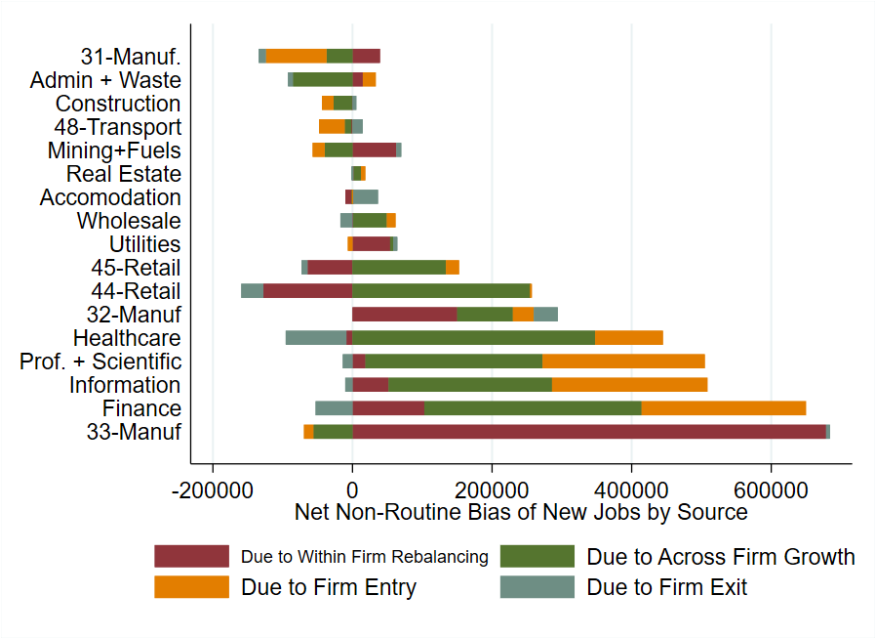

Figure 10: Non-routine bias of new jobs by source. Summing the negative sources to the positive sources yields the blue bar in figure 9, corresponding to the total net non-routine bias of new employment in the industry.

Figure 10 decomposes the non-routine bias of employment growth into its components by industry. For the manufacturing industries, within-firm rebalancing is by far the most important cause of non-routine-biased employment change. At first this might seem mechanical, given that Figure 9 suggests that employment change was very non-routine-biased, and that manufacturing firms have not had large overall employment growth.[3] Yet even in a shrinking industry, non-routine biased employment growth can be caused by firms with a relatively high amount of routine workers contracting faster than others. This was not the case. Notably, the across firm growth effect is actually slightly negative for 33-Manufacturing firms. Of all industries, only Healthcare seems perfectly consistent with the hypothesis that an industry-wide boom is leading to skill-biased technological change. For industries like Finance, Retail, and Professional and Scientific Services, across firm growth is the most important cause of non-routine- biased employment growth. This is more consistent with the Superstar firm hypothesis. The industry-level dynamics vary considerably. Industry compositional shifts are therefore a primary component of SBTC.

It is also interesting to note the great variety in the contributions of firm entry and exit to non-routine-biased employment growth across industries. For Finance, Information, and Professional and Scientific services, three industries with many successful 'unicorn' entrants in recent years, entry has been an important contributor. In these industries exit has been slightly routine-biased, largely because all firms in the industry are non-routine intensive. The non-routine bias of exit closely

[3]For the country as a whole, employment in all manufacturing industries declined from 17.3 million in 2000 to 12.4 million in 2016 BLS (2018). For publicly traded US firms in industry 33-Manufacturing, employment increased by 1.4 million.

inversely follows the routine-bias of firms in the industry in 2000.

A final question we can ask is how employment mix dynamism varies by industry. Figure 11 shows the average, firm employment weighted, firm employment mix dynamism for total economy and selected industries. All industries experience an overall downtrend in dynamism. In this figure an industry's employment mix dynamism is unrelated to industry's routine-biased employment growth due to within-firm-rebalancing. It is unclear whether this is because some industries are inherently more dynamic, or a statistical relic of the fact that some industries have a more diverse array of occupations. The range for the measure is quite narrow, and remarkably stable over the observation period.

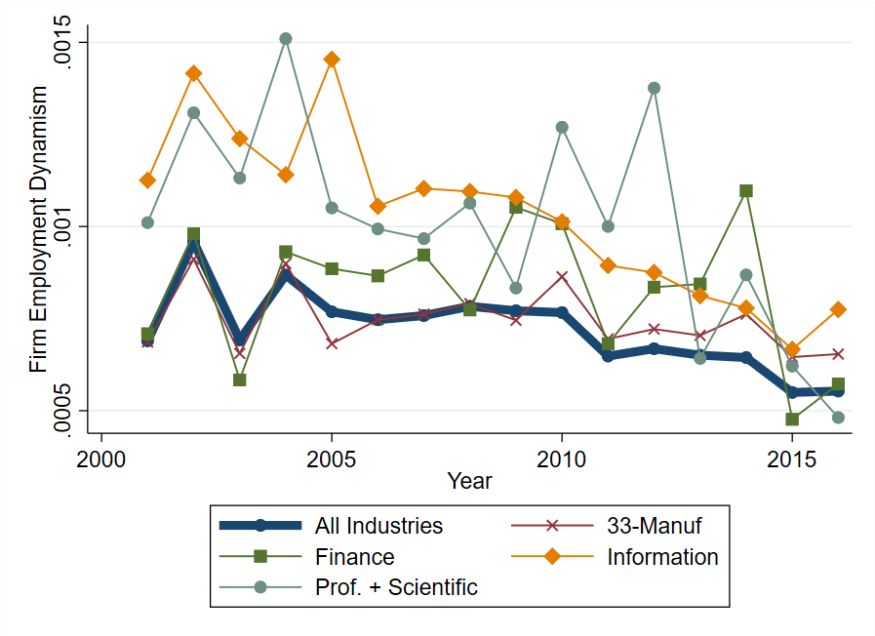

Figure 11: Average, firm employment weighted, firm employment mix dynamism for total economy and selected industries.

8 SBTC and Firm Investment

As documented above, within-firm rebalancing has played an important role in the relative decline of routine-task intensive occupations. As a final exercise, we document the relationship between within-firm rebalancing and other firm characteristics. This is important because to the extent these relationships are causal, they can drive policies meant to help or hurt the automated. We find that R&D investment tends to increase firm level demand for routine occupations, suggesting that R&D subsidies may be an effective mechanism for combating occupational displacement.

8.1 R&D and Routine Task Intensity

All theories of SBTC entail some connection between innovation and demand for different tasks. In "factor eliminating" models of technological change, such as in Peretto and Seater (2013) firms can make an investment in R&D to change the factor intensity of their production functions. Over the last 35 years, real interest rates have steadily declined and general automation technologies advanced (Benzell and Brynjolfsson, 2018). Therefore, it should be increasingly profitable for firms to develop new products and processes that substitute cheap capital for expensive laborers. Under this hypothesis, we would expect to see firms making greater R&D expenditures experiencing more SBTC. In models of directable technological change at the firm level, such as Acemoglu and Restrepo (2017), firms can choose between either inventing new tasks for low cost labor, or automating old tasks. If routine tasks are both more easily automated and lower wage, then these models are ambiguous in their prediction of the effect of R&D expenditure on routineness and average firm wage. However, if all tasks are equal in their ability to be automated, then R&D should lead to a lower average firm wage.

However, an OLS regression of R&D on firm routineness or average wage is likely to be misleading. This is because many omitted variables exist that could drive both the mix of workers employed by a firm and R&D decisions. Therefore we require an instrument which can predict R&D investment without being otherwise correlated with routineness.

The instruments for R&D we use are state and federal tax laws as codified in Bloom et al. (2013) and updated through 2015 in Lucking et al. (2018). State tax laws vary the attractiveness of R&D through tax-credits, depreciation allowances, and corporate taxes. These tax incentives differentially apply based on where a firms' R&D is located. To determine which state tax laws a firm is exposed to, Bloom et al. (2013) use the 10 year moving average location of a firms' patent holders. Federal tax incentives for R&D vary with time and are a function of a firms' historical R&D and revenues. The specific instruments used are the log of the tax price component of R&D user cost.

The first stage results of the state and federal tax component of R&D cost on a firms' R&D stock is reported in table 1. The results are broadly consistent with those reported in Lucking et al. (2018). In each specification forms of tax cost are negatively correlated with R&D stocks. However, we focus on specifications with firm and time fixed effects, as state R&D credits is likely correlated to a states' latent attractiveness for research, and federal R&D likely vary over time as a function of the macroeconomic environment.

Table 2 reports the effect of R&D, after instrumentation with tax policy, on firm routineness. We measure routineness as the the average routine-task intensity of occupations at a firm, normalized by year. Across specifications with firm fixed effects, we find a significant positive relationship between a firm's R&D stock and the average routineness of its workers.

This effect is quite large. The most complete specification indicates that a firm that doubles its R&D stock will see the routineness of its average occupation increase .355 standard deviations. To give a sense of how large an effect this is, consider that the average information firms' routineness

Table 1: Regressions of Firm Average Occupational Routineness on R&D: First Stage

	(1) Log(R&D Stock)	(2) Log(R&D Stock)	(3) Log(R&D Stock)	(4) Log(R&D Stock)
Federal Cost of RD Invest	-9.361***	-1.621***	-1.290***	-1.198***
	(0.000)	(0.000)	(0.000)	(0.000)
State Cost of RD Invest	-0.695***	-0.0550	-0.0699	0.00280
	(0.000)	(0.576)	(0.252)	(0.924)
Log(Employment)			0.409***	0.148***
			(0.000)	(0.000)
Log(Net PP&E)				0.120**
				(0.003)
Log(Market Value)				0.217***
				(0.000)
Firm FE		X	X	X
Year FE	X	X	X	X
Industry FE	X	X	X	X
N	14262	14247	14236	14105

p-values in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

in 2016 is -.048. In the same year, the construction industry's average routineness is .088. In other words then, a firm increasing its R&D stock by 39% moves it from having the average routineness of an information industry firm to the average routineness of a construction industry firm. It is possible that firms receiving a tax credit to make additional investments in R&D contemporaneously increase the quantity of routine workers hired to implement existing productive processes. It might also be the case that automation activities are infra-marginal, but investment in generating new tasks for routine workers is cross-subsidized by the R&D policy changes. The long-run effects of the R&D policy shifts on routineness are not captured by this analysis. Even if the local average treatment effect for suggests R&D stock increases cause higher average firm routineness, in the long-run firms might devote more resources to SBTC. However, were this result to generalize, it would suggest R&D tax credits as a possible channel to mitigate the short-term impact of automation technologies on labor. Financial constraints may be another possible confound. If firms respond to R&D subsidies by hiring more flexible sources of labor because of anticipated future difficulties in retaining workers, we would expect to see more routineness in the companies that take up the subsidy. Lucking (2019) provides evidence in support of job creation from innovation for these firms, and fails to find support in favor of financial constraints driving the results.

This result, relating firm R&D stock to routineness, is not driven by manufacturing firms alone. Dividing industries into manufacturing (NAICS codes of 31, 32, or 33) or non-manufacturing, (as in table 3, both show a significant negative relationship. However, the F-statistic for non-manufacturing firms is low, indicating these results are potentially biased by a weak instrument.

In addition to these results, R&D stocks, instrumented with tax credits, are strongly positively

Table 2: Regressions of Firm Average Occupational Routineness on R&D: IV Results

	(1)	(2)	(3)	(4)
	Norm_Firm_Routineness	Norm_Firm_Routineness	Norm_Firm_Routineness	Norm_Firm_Routineness
Log(R&D Stock)	-0.0658***	0.378***	0.438***	0.355**
	(0.000)	(0.000)	(0.000)	(0.001)
Log(Employment)			-0.0525	0.104
			(0.223)	(0.075)
Log(Net PP&E)				-0.0602***
				(0.000)
Log(Market Value)				-0.102***
				(0.000)
Firm FE		X	X	X
Year FE	X	X	X	X
Industry FE	X	X	X	X
N	14262	14247	14236	14105
F-Stat	180.4	37.36	22.08	32.43

p-values in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 3: Regressions of Firm Average Occupational Routineness on R&D: IV Results

	(1)	(2)	(3)	(4)
	Norm_Firm_Routineness	Norm_Firm_Routineness	Norm_Firm_Routineness	Norm_Firm_Routineness
Log(R&D Stock)	0.531***	0.435***	0.385*	0.493*
	(0.000)	(0.000)	(0.034)	(0.013)
Log(Employment)		0.0894		0.0887
		(0.085)		(0.569)
Log(Net PP&E)		-0.0533**		-0.138*
		(0.003)		(0.022)
Log(Market Value)		-0.122***		-0.122
		(0.000)		(0.118)
Firm FE	X	X	X	X
Year FE	X	X	X	X
Manuf or non-Manuf?	Manuf	Manuf	non-Manuf	non-Manuf
N	11377	11247	2870	2858
F-Stat	10762.7	19976.9	3.819	3.362

p-values in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 4: Firm Characteristic Relationships with Firm and Year FEs

	(1) Norm_Firm_Routineness	(2) Avg_firm_wage	(3) Log(Employment)	(4) lnMarketValueInTotLiab	(5) Log(Net PP&E)	(6) Log(R&D Stock)
Avg_firm_wage	-0.00000291**		0.000000999***	0.00000220***	-0.000000391	0.00000212**
	(0.001)		(0.000)	(0.000)	(0.128)	(0.001)
Log(Employment)	0.163***	12963.8***		0.387***	0.496***	0.164***
	(0.000)	(0.001)		(0.000)	(0.000)	(0.000)
Log(Market Value)	0.0137	21471.8***	0.291***		0.439***	0.331***
	(0.257)	(0.000)	(0.000)		(0.000)	(0.000)
Log(Net PP&E)	-0.0128	-2115.9	0.206***	0.243***		0.0649
	(0.249)	(0.073)	(0.000)	(0.000)		(0.055)
Log(R&D Stock)	0.00756	10719.3***	0.0640**	0.171***	0.0606**	
	(0.434)	(0.000)	(0.003)	(0.000)	(0.006)	
Anticipated Wage Incr. 1yr	0.000000376	0.110***	0.000000387**	-0.000000554***	-7.50e-08	-0.000000908**
	(0.125)	(0.000)	(0.004)	(0.000)	(0.524)	(0.001)
Anticipated Wage Incr. 2yr	0.000000811*	0.358***	0.000000771***	-0.00000103***	0.000000911**	-0.00000103*
	(0.015)	(0.000)	(0.001)	(0.001)	(0.001)	(0.046)
IT Consultant Share	-1.246*	83154.9**	-1.641	0.400	0.510	0.158
	(0.034)	(0.004)	(0.055)	(0.078)	(0.060)	(0.735)
Norm_Firm_Routineness		-23398.2***	0.101***	0.0113	-0.0191	0.0121
		(0.000)	(0.000)	(0.231)	(0.355)	(0.431)
Firm FE	X	X	X	X	X	X
Year FE	X	X	X	X	X	X
N	23493	23493	23493	23493	23493	23493

p-values in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

associated with firm size, and not strongly related to the firms' average occupational wage.

8.2 Correlations Between Wage Changes, Investment, and Routineness

There are many extant theories that would connect anticipated wage increases, firm investment, and routineness. For example, if investments – either in physical capital or R&D have the potential to change a firm's optimal employment mix, then it may make increased sense for a firm to make these investments if it anticipates wages will increase for it's current employee mix. For this reason and more, we provide a sample of the rich relationships in the data in table 4.

8.3 Firm Occupational Rebalancing and Adjustment Costs

One other factor that mediates the relationship between, wages, occupational mix and investment is adjustment costs. It stands to reason that if firms face fixed costs in hiring and firing workers, then they are more likely to change their occupational mix when growing or contracting. The extent of within firm rebalancing may also be related to growth through occupation-specific productivity changes at the firm level. Technological shocks of this sort will also change overall firm productivity, leading the firm to expand or contract in employment as a function of the demand curve for the firms' products.

Figure 12 reports growth in occupations by routineness for firms sorted by their employment growth. The far left figure shows firms in the bottom quintile of total employment growth, and the far right shows firms in the top quintile. Attention is restricted to firms that operate in both

2000 and 2016, and each category of firms had an equal amount of employees in 2000. As can be seen, within firm rebalancing is more common for firms with shrinking employment than those with constant or increasing employment. Over this interval, the growing firms saw about 90 thousand employees moving between occupational routineness categories, constant employment firms saw about 101 thousand, and shrinking firms saw about 258 thousand. This result suggests that the fastest growing firms in our sample maintain relatively constant employment mixtures.

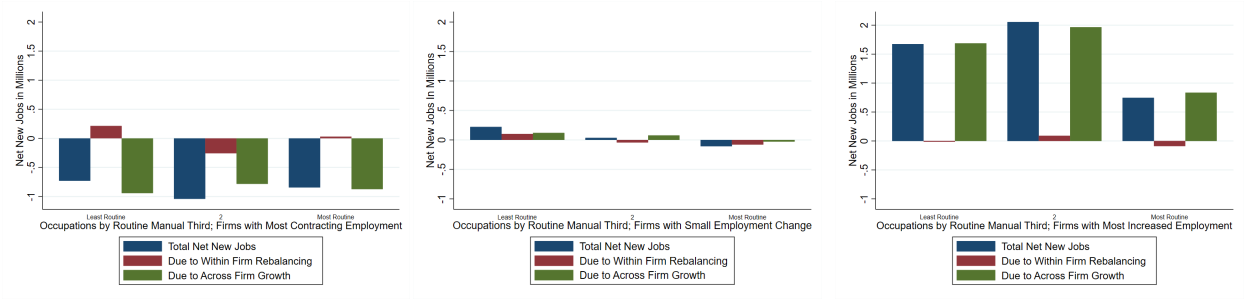

(a) Firms in Bottom Quintile of Employment Growth

(b) Firms in Second Quintile of Employment Growth

(c) Firms in Top Quintile of Employment Growth

Figure 12: Change in employment by non-routineness of occupation in thirds and firm estimate wage bill growth in quintiles. 2000 through 2016. Net new job growth decomposed into sources. Only firms which do not enter or exit included.

This indicates that both causes of SBTC are important. Part of SBTC is driven by the fastest growing firms being routine biased, and part by the fastest contracting disproportionately laying of routine workers.

9 Conclusion

In this paper we decompose, for the first time, aggregate skill-biased technological change into its sources within and across firms. There are several caveats to the conclusions in this study. First, much work remains in exploring alternative specifications and measures for the relationship between investment, growth and firm-level skill-biased employment changes. Second, our data source, while unique and powerful, is constructed using self-reported resumes. To deal with this, we introduce a system of reweighing observations by firm, industry and occupation. What is clear from our analysis is that the shifting of occupational types within firms is an important component of SBTC. Future work will explore the mechanisms behind these shifts in greater detail.

We find that within-firm rebalancing is the second most important source, quantitatively, of non-routine biased employment growth. It is especially important for the increasing non-routineness of labor for firms in initially routine-intensive industries, such as manufacturing. Rebalancing within firms in these industries is the most important sources of skill-biased technological change. Within firm rebalancing explains almost all the decline in employment for the fastest contracting occupations. These observations run counter to theories of SBTC that emphasize firm production

function neutral technological changes alone.

Still, we find that growth across firms is still important in industries that are already relatively intensive in non-routine tasks. Net firm entry also a plays a larger role for growing industries and industries that are initially non-routine intensive.

Contrasting the non-routine bias of within-firm rebalancing across firms with different investment and growth, we find a weak correlation between capital intensity and firms becoming more non-routine intensive. In our regression analysis, we find that firms which do more R&D as a result of a tax break contemporaneously increase their relative share of routine workers. All these trends are consistent with a theory of SBTC driven by technological changes which firms implement infra-marginally.

Consistent both with adjustment costs or with firm-level factor intensity changing technological change, we find that growing and shrinking firms experience more non-routine-biased within-firm rebalancing than firms that do not change in size. In a contribution to the business dynamism literature, we introduce a measure of how frequently firms change their employment mix. We find that this has steadily decreased from 2000 to the present. It is unclear whether a decrease in this measure will be associated with reduced skill-biased tech change in the future however, as there was no clear relationship between industry level dynamism and non-routine biased within-rebalancing.

References

Acemoglu, Daron, "Changes in unemployment and wage inequality: An alternative theory and some evidence," *American Economic Review*, 1999, *89* (5), 1259–1278.

_ , "Directed Technical Change," *Review of Economic Studies*, 2002, *69* (4), 781–809.

_ **and David Autor**, "Skills, tasks and technologies: Implications for employment and earnings," in "Handbook of labor economics," Vol. 4, Elsevier, 2011, pp. 1043–1171.

_ **and Pascual Restrepo**, "The Race Between Man and Machine: Implications of Technology for Growth, Factor Shares and Employment," 2017.

_ **, David Autor, David Dorn, Gordon H Hanson, and Brendan Price**, "Import competition and the great US employment sag of the 2000s," *Journal of Labor Economics*, 2016, *34* (S1), S141–S198.

Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, John Van Reenen et al., *The fall of the labor share and the rise of superstar firms*, National Bureau of Economic Research, 2017.

Autor, David H. and David Dorn, "The growth of low-skill service job and the polarization of the US labor market," *The American Economic Review*, 2013, *103* (5), 1553–1597.

_ , **Frank Levy, and Richard J. Murnane**, "The Skill Content of Recent Technological Change: An Empirical Exploration," *The Quarterly Journal of Economics*, nov 2003, *118* (4), 1279–1333.

Barth, Erling, James C. Davis, Richard B. Freeman, and Andrew J. Wang, "The Effects of Scientists and Engineers on Productivity and Earnings at the Establishment Where They Work," Working Paper 23484, National Bureau of Economic Research June 2017.

Benzell, Seth G. and Erik Brynjolfsson, "Digital Abundance and Scarce Genius: Implications for Interest Rates and Wages," Technical Report, MIT Mimeo 2018.

Benzell, Seth G, Laurence J Kotlikoff, Guillermo Lagarda, and Jeffrey D Sachs, "Robots Are Us: Some Economics of Human Replacement," *NBER Working Paper Series*, 2016, (Working Paper 20941).

Berman, Eli, John Bound, and Stephen Machin, "Implications of Skill-Biased Technological Change: International Evidence," *Source: The Quarterly Journal of Economics*, 1998, *113* (49), 1245–127961.

_ , _ , **and Zvi Griliches**, "Changes in the Demand for Skilled Labor within U.S. Manufacturing: Evidence from the Annual Survey of Manufacturers," *Quarterly Journal of Economics*, 1994, (May).

Bessen, James, "AI and Jobs: the role of demand," in "Economics of Artificial Intelligence," University of Chicago Press, 2018.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen, "Identifying technology spillovers and product market rivalry," *Econometrica*, 2013, *81* (4), 1347–1393.

BLS, "Labor Force Statistics from the Current Population Survey," 2018.

Bresnahan, T. F., E. Brynjolfsson, and L. M. Hitt, "Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence," *The Quarterly Journal of Economics*, feb 2002, *117* (1), 339–376.

Brynjolfsson, Erik, Daniel Rock, and Chad Syverson, "The Productivity J-Curve: How Intangibles Complement General Purpose Technologies," *mimeo*, 2018.

_ , _ , **Prasanna Tambe, and Jacqueline Barrett**, "People Are Our Greatest Asset: Measuring the Market Value of Human Capital," in "in" 2018.

Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda, "The role of entrepreneurship in US job creation and economic dynamism," *Journal of Economic Perspectives*, 2014, *28* (3), 3–24.

Goos, Maarten, Alan Manning, and Anna Salomons, "Explaining Job Polarization in Europe: The Roles of Technology, Globalization and Institutions," *CEP Discussion Papers*, 2010.

_ , _ , _ , **David Autor, Larry Katz, Thomas Lemieux, Stephen Machin, Guy Michaels, John Van Reenen, and Ulrich Zierahn**, "Explaining Job Polarization: Routine-Biased Technological Change and Offshoring," 2014.

_ **and** _ , "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain," *The Review of Economics and Statistics*, 2007, *89*, 118–133.

Johnson, George E, "American Economic Association Changes in Earnings Inequality: The Role of Demand Shifts," *Source: The Journal of Economic Perspectives*, 1997, *11* (2), 41–54.

Kaplan, Steven N. and Joshua Rauh, "It's the Market: The Broad-Based Rise in the Return to Top Talent," *Journal of Economic Perspectives*, 2013, *27* (3), 35–56.

Lucking, Brian, "Do R&D Tax Credits Create Jobs?," *Unpublished Working Paper*, 2019.

_ **, Nicholas Bloom, and John Van Reenen**, "Have R&D Spillovers Changed?," 2018.

Peretto, Pietro F and John J Seater, "Factor-eliminating technical change," *Journal of Monetary Economics*, 2013, *60* (4), 459–473.

Peters, Ryan H and Lucian A Taylor, "Intangible capital and the investment-q relation," *Journal of Financial Economics*, 2017, *123* (2), 251–272.

Rosen, Sherwin, "The Economics of Superstars," *The American Economic Review*, 1981, *71* (5), 845–858.

Western, Bruce and Jake Rosenfeld, "Unions, norms, and the rise in US wage inequality," *American Sociological Review*, 2011, *76* (4), 513–537.