

Federal Reserve Bank of New York
Staff Reports

The Role of Start-Ups in Structural Transformation

Robert C. Dent
Fatih Karahan
Benjamin Pugsley
Ayşegül Şahin

Staff Report No. 762
January 2016



This paper presents preliminary findings and is being distributed to economists and other interested readers solely to stimulate discussion and elicit comments. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System. Any errors or omissions are the responsibility of the authors.

The Role of Start-Ups in Structural Transformation

Robert C. Dent, Fatih Karahan, Benjamin Pugsley, and Ayşegül Şahin

Federal Reserve Bank of New York Staff Reports, no. 762

January 2016

JEL classification: E24, J00, J23, L25

Abstract

The U.S. economy has been going through a striking structural transformation—the secular reallocation of employment across sectors—over the past several decades. We propose a decomposition framework to assess the contributions of various margins of firm dynamics to this shift. Using firm-level data, we find that at least 50 percent of the adjustment has been taking place along the entry margin, owing to sectors receiving shares of start-up employment that differ from their overall employment shares. The rest is mostly the result of life cycle differences across sectors. Declining overall entry has a small but growing effect of dampening structural transformation.

Key words: structural transformation, employment dynamics, sectoral reallocation

Dent, Karahan, Pugsley, Şahin: Federal Reserve Bank of New York (e-mails: rob.dent@ny.frb.org, fatih.karahan@ny.frb.org, benjamin.pugsley@ny.frb.org, aysegul.sahin@ny.frb.org). The authors are grateful for many useful comments from Mu-Jeung Yang and excellent research assistance by Harry Wheeler. All results have been reviewed to ensure that no confidential information is disclosed. The views expressed in this paper are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

1 Introduction

The U.S. economy experienced a striking structural transformation—the secular reallocation of employment across sectors—over the past several decades. Most notably, the employment share of manufacturing has declined substantially, matched by an increase in the share of services. Despite a large literature studying the causes and consequences of structural transformation, little is known about the dynamics of reallocation of labor from one sector to the other.¹

There are several margins through which a sector could grow and shrink relative to the rest of the economy. First, differences in growth and survival rates of firms across sectors could cause sectoral reallocation of employment. In addition, differences in sectors’ firm age distribution could affect reallocation since firm age is an important determinant of growth or survival behavior.² We refer to these as life cycle margins. Second, the allocation of employment at the entry stage—which we refer to as the entry margin—could contribute to the gradual shift of employment from one sector to the other.³ Third and finally, because the speed at which differences in entry patterns are reflected in employment shares depends on the aggregate entry rate, changes in the latter could affect the extent of structural transformation.

A thorough understanding of sectoral evolutions requires an understanding of the relative importance of these margins, a task that we undertake in this paper. We build on the framework in Pugsley and Şahin (2014) to dynamically decompose the joint evolution of employment across firm age and sector. While we have analyzed all sectors, we focus on three sectors in this paper: manufacturing, retail trade, and services.

Using data from the Longitudinal Business Database (LBD) and Business Dynamic Statistics (BDS), we find that at least 50 percent of employment reallocation since 1987 has occurred along the entry margin. Interestingly, 85 percent of the decline in manufacturing employment share is predictable from the average life cycle dynamics and the early 1980s distribution of startup employment across sectors. Further changes over time in the distribution of startup employment away from manufacturing, while having a relatively small effect on manufacturing where entry is less important, explain almost one-third of the increase in the services employment share. We find little role for the year-to-year variation in incumbent behavior conditional on firm age in explaining long-term sectoral reallocation. Lastly, we find that a 30-year decline in overall entry—which we refer to as the “startup deficit”—has a small but growing effect of dampening sectoral reallocation through the entry margin.

¹See Rogerson (1987), Duarte and Restuccia (2010) and the references therein.

²See Dunne, Roberts and Samuelson (1988) and Haltiwanger, Jarmin and Miranda (2013), which document the importance of plant and firm age in survival and employment growth.

³The idea of life cycle and entry margins in structural transformation builds on influential earlier work by Gort and Klepper (1982), Klepper (1996) and Agarwal and Gort (1996) studying the patterns of entry, survival and growth in new product markets.

2 Dynamic Decomposition Framework

Let E_t^j denote employment in sector j in year t . E_t^j consists of employment at new firms, or startups, S_t^j , and at incumbents of all ages. Noting that this year's incumbent employment is a function of overall sector employment from the previous year and survival and growth rates, we can express E_t^j as

$$E_t^j = S_t^j + x_t^j(1 + n_t^j)E_{t-1}^j,$$

where x_t^j is the fraction of year $t - 1$ firms that survive into year t , or survival rate, and n_t is the growth in average firm employment, or conditional growth rate. Defining sector j employment share $e_t^j \equiv E_t^j/E_t$ and startup share $s_t^j \equiv S_t^j/S_t$, sector j 's employment share is

$$e_t^j = s_t^j \frac{S_t}{E_t} + e_{t-1}^j \frac{x_t^j(1 + n_t^j)}{E_t/E_{t-1}}. \quad (1)$$

Equation (1) expresses the sector's employment share in terms of the contribution from the aggregate startup employment share S_t/E_t weighted by sector j 's share of startup employment s_t^j and the contribution from sector j 's incumbent growth rate $x_t^j(1 + n_t^j)$ relative to aggregate employment growth E_t/E_{t-1} weighted by the previous year employment share e_{t-1}^j .

Viewed in this way, equation (1) shows the two key channels of structural transformation. First, aggregate startup employment forces reallocation through differences across sectors in startup shares. Second, cross-sector life cycle differences, drive reallocation through sectoral dispersion in incumbent growth relative to aggregate employment growth.

Because both survival and growth vary significantly not only across sector, but also with firm age, a sector's overall incumbent growth rate depends on the distribution of employment across incumbent firm age groups. For current, year t , mature firms, let ω_{mt-1} denote their year $t - 1$ employment share. Letting $G_t^j = x_t^j(1 + n_t^j)$ be the sector's incumbent growth rate, it may be written as an employment-weighted average of young and mature incumbent growth rates

$$G_t^j = \sum_{k \in \{y, m\}} \omega_{kt-1}^j x_{kt}^j (1 + n_{kt}^j). \quad (2)$$

Since these are shares of previous year employment, $\omega_{yt-1} = 1 - \omega_{mt-1}$.

We refer to shifts in a sector's overall incumbent growth rate over time implied by equation (2) as the sector life cycle. This life cycle effect is composed of two parts: (i) expected firm survival x_{kt}^j and growth n_{kt}^j , which vary across firm age group k , (ii) the age composition of incumbents evolves over time. We refer to these parts as the conditional life cycle dynamics and life cycle stage, respectively. Together, they account for the changes in expected incumbent growth over time.

While we do not condition on firm age in equation (1), in our empirical implementation we use age-specific survival and growth rates so that the overall incumbent growth rate varies endogenously

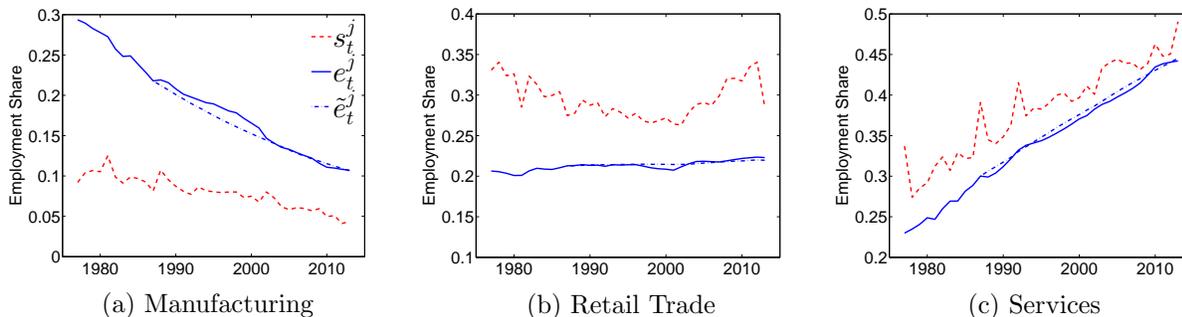


Figure 1: Employment Shares by Industry

Note: U.S. Census Bureau Business Dynamics Statistics and authors' calculations. This figure depicts each sector's share of overall employment (solid lines) and share of startup employment (dotted lines) for 1977 to 2013. The dashed lines show sector employment shares that would have arisen had the age-specific growth and survival rates been flat at their sample averages for 1987-2013.

with the sector life cycle. In the appendix, we extend the framework of Pugsley and Şahin (2014) to correctly adjust the age group employment shares in equation (2) for all sectors consistent with their entry and life cycle dynamics. Starting from an initial distribution of employment, and with exact measures of s_t^j , age-specific survival and growth rates for each sector, and aggregate startup employment S_t , this provides an exact decomposition for the evolution of the joint distribution of employment across sector and firm age.

3 Decomposition Results

3.1 Data

We use data from the Census Longitudinal Business Database and its public use counterpart the Business Dynamics Statistics (BDS) for the years 1977 to 2013, which has near universal coverage of the nonfarm business sector. The method of assigning age ensures that startups (age 0 firms) are composed entirely of new establishments, and is robust to changes in ownership such as mergers and acquisitions.⁴ We define young firms as those ages 1 to 10 and mature firms as those ages 11+. When conditioning on firm age we use data from 1987, which is the first year we can distinguish young and mature firms.

We aggregate by age group and sector and then measure employment, survival and conditional growth rates as in Pugsley and Şahin (2014). With these measures we apply the decomposition to the entire nonfarm business sector economy. For this study, we focus on the manufacturing, retail trade, and services sectors, which account for roughly 75 percent of nonfarm business sector employment.⁵ Figure 1 plots as a solid line each sector's share of overall employment for 1977

⁴See Haltiwanger, Jarmin and Miranda (2013) and Pugsley and Şahin (2014) for details.

⁵In the online appendix we include results for a complete set of sectors.

to 2013. Over these nearly 40 years, the manufacturing sector has steadily declined, the retail trade sector has been roughly flat and the service sector has steadily grown. The dotted line plots each sector’s share of startup employment. Manufacturing’s share of startup employment was 10 percent in 1980 and has since fallen by half to 5 percent. Retail trade’s share of startup employment declined from 30 percent over the 1980s and 1990s before starting to increase in the 2000s. Services’ share of startup employment was already 30 percent in 1980 and has continued to increase. This dispersion in startup shares s_t^j along with the shifts over this period drive reallocation through the entry margin.

3.2 Fluctuations in Life Cycle Dynamics

Changes over time in conditional life cycle dynamics, i.e, survival and conditional growth rates by firm age and sector, in the 1987 to 2013 period explain very little of the employment reallocation. Pugsley and Şahin (2014) show that, while the levels of age-specific growth rates vary across sectors, these relative differences have been stationary over time. This observation suggests little role for fluctuations in growth rates around their age- and sector-specific means in sectoral reallocation of employment. We show that this conjecture is correct by replacing in equation (2) age-specific conditional growth and survival rates with their sample averages for each sector and allowing the firm age distribution to vary endogenously

$$\tilde{G}_t^j = \sum_{k \in \{y, m\}} \tilde{\omega}_{kt-1}^j \bar{x}_k^j (1 + \bar{n}_k^j).$$

Here conditional life cycle dynamics are held fixed; the incumbent growth rate \tilde{G}_t^j varies only because of the shifts in the life cycle stage captured by employment shares $\tilde{\omega}_{kt}^j$. The dashed lines in Figure 1 show that we capture the sectoral evolution of employment well. Although it misses some of the boom and bust, for example the shakeout in retail trade in the late 1990s and early 2000s studied in Foster, Haltiwanger and Krizan (2006), it is the dispersion across sectors in the levels of the survival and growth rates, and not their fluctuations, that primarily account for the reallocation of employment across sectors. In the online appendix we show using the LBD microdata, that the same pattern is true in other nonfarm business sectors, including a high-tech sector. Going forward, we consider only the conditional life cycle dynamics captured by the average age-specific survival and conditional growth rates by sector.

3.3 Additive Decomposition

Next, using only the average conditional life cycle dynamics, we quantify the relative roles of startups and incumbents in driving structural transformation. To do this we consider two experiments. First, using equation (1) we replace s_t^j with the average early 1980s employment share \bar{e}_{1980s}^j , and

Table 1: Percent of employment share change explained

	Incumbent	Entry Margin	
	Life Cycle	\bar{s}_{1980s}	s_t
Manufacturing	32.8	53.8	13.4
Retail Trade	-272.2	509.3	-137.2
Services	48.5	18.8	32.7

U.S. Census Bureau Business Dynamics Statistics and authors' calculations. Fraction of the change in employment share with average life cycle dynamics that is explained by the incumbent and entry margins for different sectors. Incumbent margin is change in employment share with startup shares held fixed at 1980s employment shares. Entry margin is additional change from allowing (i) initial dispersion and then (ii) initial dispersion and changes in startup shares.

using actual aggregate startup employment S_t and life cycle effects in incumbent growth \tilde{G}_t^j for all sectors, we solve for the sequence of sector employment shares. From equation (1), one can see that if $s_t^j = e_t^j$ only dispersion in \tilde{G}_t^j can drive further changes in employment shares. We label the employment reallocation through only incumbent growth, the incumbent life cycle effect. The first column of Table 1 shows the percentage of the overall change in employment share from 1987 to 2013 explained only by incumbent life cycle effects. In manufacturing they account for roughly one-third, whereas in retail trade they predict a decline rather than an increase. The latter is because of the low survival rates in the sector.

In our second experiment, we follow the same procedure, but instead replace s_t^j with its early sample average s_{1980s}^j for all sectors. This introduces further reallocation through the dispersion in startup shares across sectors, but holds those shares fixed instead of allowing them to shift as in Figure 1. Allowing dispersion in startup shares completely offsets the employment share declines in retail trade. It also explains an additional 55 percent of the change in manufacturing employment—in other words, 85 percent of the decline in manufacturing employment is predictable just from the dispersion in startup shares and average life cycle dynamics.

The remaining gap is explained by allowing the startup shares to vary over time as they do in Figure 1. These further changes are relatively small for manufacturing, less than 15 percent, but they are almost one-third of the rise of services. Together, the initial dispersion and evolution of the startup shares account for more than 50 percent of all reallocation in each sector over our sample period. In the online appendix we plot the entire time series of the shares comprising these additive effects.

4 Effects of the Aggregate Startup Deficit

Finally, we quantify what additional employment reallocation there would have been if there had not been a 30-year secular decline in the entry rate, what we refer to as the startup deficit. To do

Table 2: Actual and without startup deficit employment shares

	Overall Change (pp)	
	Actual	No Startup Deficit
Manufacturing	-11.1	-12.0
Retail Trade	0.7	1.3
Services	14.5	15.3

Note: this table presents the actual change in employment shares and the change if the trend in startup employment growth S_t/S_{t-1} were replaced by 2 percent, as in Pugsley and Şahin (2014)

this we repeat the counterfactual experiment in Pugsley and Şahin (2014), which replaces the time-varying drift in aggregate startup employment growth S_t/S_{t-1} with a constant rate of 2 percent, thereby eliminating the startup deficit.

Examining equation (1), increasing the growth rate of aggregate S_t has two effects, first it increases the aggregate startup employment share S_t/E_t and second, over the long run it increases aggregate employment growth E_t/E_{t-1} . Together these place more weight on the dispersion in startup shares, speeding reallocation towards the long run distribution of employment implied by the entry and life cycle dynamics. Reductions in aggregate entry have a dampening effect on reallocation.

As seen in Table 1, quantitatively, the effects are small. Across all sectors, the additional reallocation from eliminating the startup deficit amounts to less than one percentage point within each sector. Although small, the effects of the aggregate decline in entry on reallocation will continue to accumulate over longer horizons through further reductions in S_t/E_t and E_t/E_{t-1} .

5 Conclusions

Our analysis points to the importance of the entry margin in understanding the dynamics of structural transformation in the U.S. economy. Differences in the entry margin, both in terms of dispersion and evolution of startup employment shares, account for more than half of the employment reallocation for each sector. The remaining employment reallocation is due to differences in life cycle margins. Interestingly, in contrast to Agarwal and Gort (1996) analysis of product markets, where life cycle dynamics change over time, we find that holding age-specific survival and conditional growth rates at their averages—so that only endogenous shifts in the age distribution affect incumbent dynamics—captures well the paths of sector employment shares over time. We also identify an additional channel for the aggregate entry rate to affect the pace of structural transformation, finding that the “startup deficit” observed since the 1980s delayed additional employment reallocation.

References

- Agarwal, Rajshree, and Michael Gort.** 1996. “The evolution of markets and entry, exit and survival of firms.” *The review of Economics and Statistics*, 489–498.
- Duarte, Margarida, and Diego Restuccia.** 2010. “The role of the structural transformation in aggregate productivity.” *The Quarterly Journal of Economics*, 125(1): 129–173.
- Dunne, Timothy, Mark J Roberts, and Larry Samuelson.** 1988. “Patterns of firm entry and exit in US manufacturing industries.” *The RAND journal of Economics*, 495–515.
- Foster, Lucia, John Haltiwanger, and Cornell J Krizan.** 2006. “Market selection, reallocation, and restructuring in the US retail trade sector in the 1990s.” *The Review of Economics and Statistics*, 88(4): 748–758.
- Gort, Michael, and Steven Klepper.** 1982. “Time paths in the diffusion of product innovations.” *The economic journal*, 630–653.
- Haltiwanger, John, Ron S Jarmin, and Javier Miranda.** 2013. “Who creates jobs? Small versus large versus young.” *Review of Economics and Statistics*, 95(2): 347–361.
- Klepper, Steven.** 1996. “Entry, Exit, Growth, and Innovation over the Product Life Cycle.” *American Economic Review*, 86(3): 562–83.
- Pugsley, Benjamin W, and Ayşegül Şahin.** 2014. “Grown-up business cycles.” *FRB of New York Staff Report*, , (707).
- Rogerson, Richard.** 1987. “An equilibrium model of sectoral reallocation.” *The Journal of Political Economy*, 824–834.

A Extending the framework from Pugsley and Şahin (2014)

A.1 Conditional distribution of employment by age given sector j

This extends the dynamic decomposition from Pugsley and Şahin (2014). Let \vec{E}_t^j be a 3×1 vector of employment by age group k for sector $j \in \{1, \dots, J\}$. Let S_t^j be a scalar of startup employment in sector j . Let P_t^j be a 3×3 transition matrix across age groups for sector j

$$P_{jt} = \begin{bmatrix} 0 & x_{yt}^j(1+n_{yt}^j) & 0 \\ 0 & q_{yt-1}^j x_{yt}^j(1+n_{yt}^j) & (1-q_{mt-1}^j)x_{mt}^j(1+n_{mt}^j) \\ 0 & 0 & x_{mt}^j(1+n_{mt}^j) \end{bmatrix}$$

where q_{t-1}^j is the fraction of the year $t-1$ young cohort that remains young in year t . This can be easily extended to include additional age groups with an appropriate choice of q_{t-1}^k for each additional age group. The law of motion for sector j is

$$\vec{E}_t^j = P_t^{j'} \vec{E}_{t-1}^j + S_t^j(1, 0, 0)'.$$

A.2 Joint distribution of age and sector

Let \vec{E}_t be a $3 * J$ length vector of employment by age group $k \in \{S, Y, M\}$ and sector $j \in \{1, \dots, J\}$, switching first over age group and then over sector. Let \vec{S}_t be a J length vector of startup employment and define a J length vector of startup employment shares $\vec{s}_t \equiv \vec{S}_t/S_t$ where S_t is aggregate startup employment. Let P_t be a $3 * J \times 3 * J$ block diagonal transition matrix with sector specific transition matrices P_t^j along its diagonal. The law of motion for the joint distribution is

$$\vec{E}_t = P_t' \vec{E}_{t-1} + S_t \vec{s}_t \otimes (1, 0, 0)'.$$

A.3 Aggregating by sector

Let H be an operator that converts the joint distribution over age and sector to a distribution over sector by summing the 3 age groups for each sector

$$H \equiv I_J \otimes (1, 1, 1)$$

Applying to the law of motion

$$\begin{aligned} H \vec{E}_t &= H P_t' \vec{E}_{t-1} + S_t H \vec{s}_t \otimes (1, 0, 0)' \\ &= H P_t' \vec{E}_{t-1} + S_t \vec{s}_t \end{aligned}$$

or in shares

$$H\vec{e}_t = \frac{HP'_t}{1 + g_t} \vec{e}_{t-1} + \frac{S_t}{E_t} \vec{s}_t,$$

where $1 + g_t = E_t/E_{t-1}$

B Additional results

This section presents results for a complete set of sectors. In Figure A.1, we show each sector’s share of overall employment (solid lines) and shares that would have arisen had the age-specific growth and survival rates been flat at their sample averages. Overall, these plots illustrate that changes in conditional life cycle dynamics do not matter much for employment reallocation.

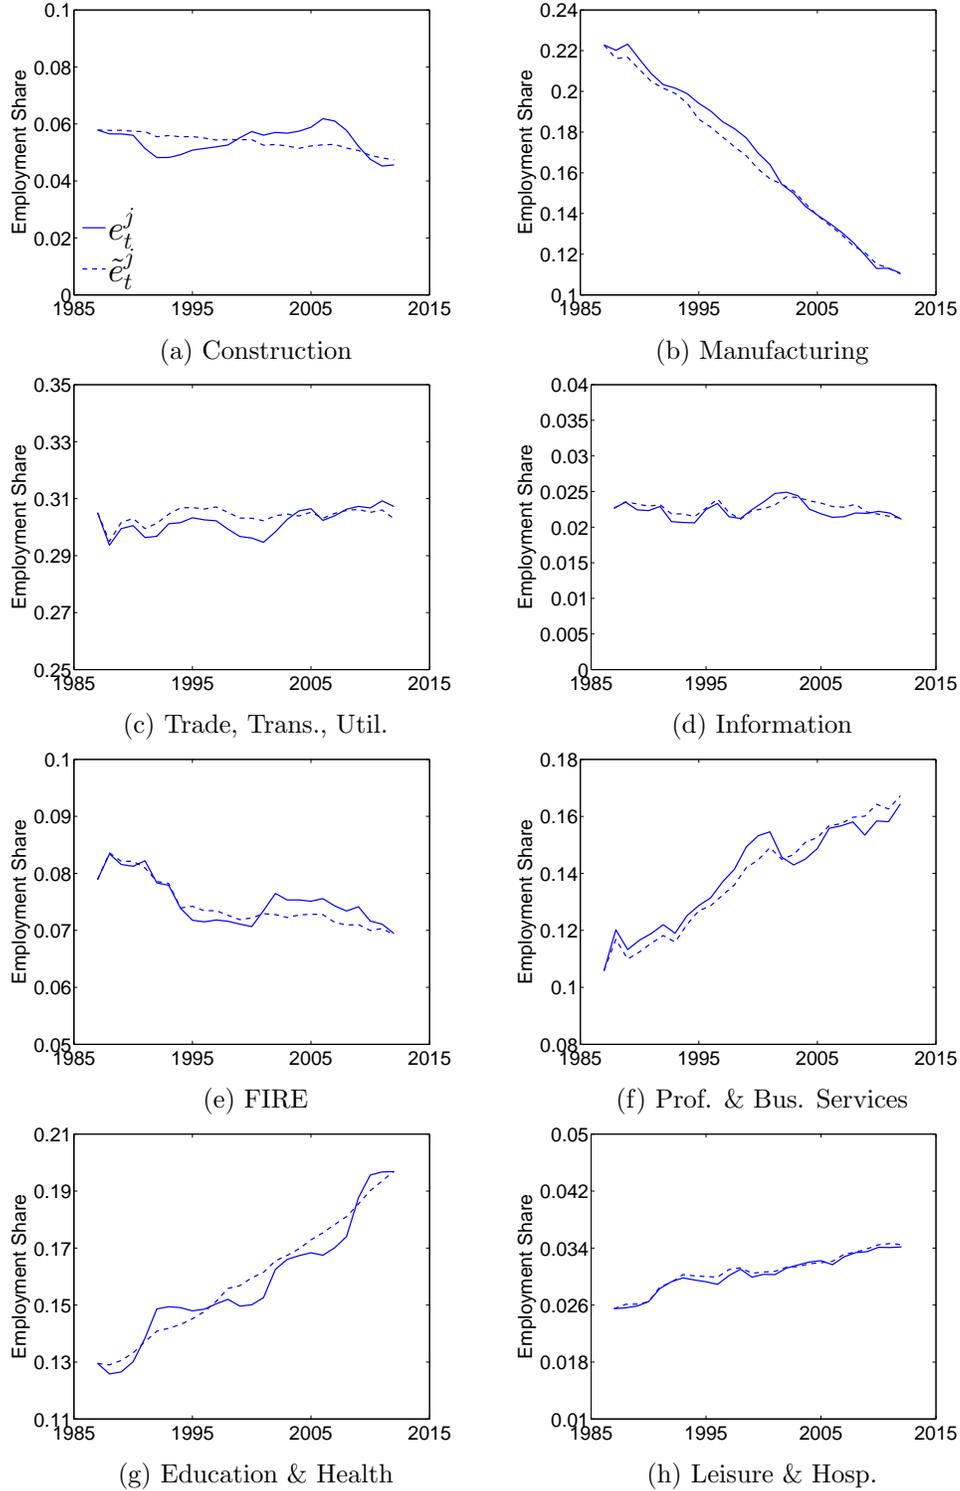
Next, we look at finer categories of sectors and compute, for each 4-digit NAICS sector, the maximum difference between that sector’s employment and the employment share that would have arisen had the age-specific growth and survival rates remained flat at their sample averages. The left panel of Figure A.2 shows the histograms of these share differences, whereas the right panel shows the same for log differences. These plots further confirm the result that changes in conditional life cycle dynamics matter very little for employment reallocation.

Lastly, we quantify the relative roles of startups and incumbents in driving structural transformation for 8 broad sectors. The results are presented in Table A.1.

Table A.1: Actual and without startup deficit employment shares

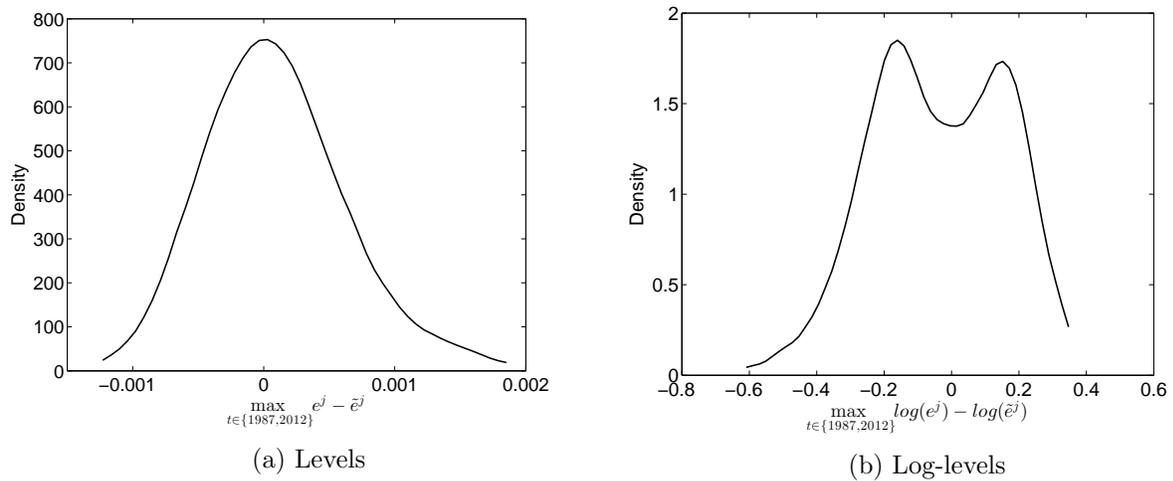
	Incumbent Lifecycle	Entry Margin		Overall Change (pp)	
		Initial	Time-varying	Actual	No Startup Deficit
Mining	-178.2	-37.5	315.7	-0.1	-0.1
Construction	128.4	-87.0	58.6	-1.8	-1.7
Manufacturing	33.6	53.2	13.3	-11.1	-12.0
Trans., Comm., Util.	-218.3	204.2	114.1	-0.3	-0.6
Wholesale Trade	-12.0	42.2	69.8	-1.1	-1.3
Retail Trade	-289.8	526.4	-136.6	0.7	1.3
FIRE	3.9	134.6	-38.5	-0.7	-0.9
Services	47.4	19.7	32.9	14.5	15.3

Notes: This table presents the fraction of the change in employment share that is explained by the incumbent and entry margins for different sectors as well as the actual change in employment shares and the change in the absence of the startup deficit. Source: U.S. Census Bureau Business Dynamics Statistics and authors’ calculations.



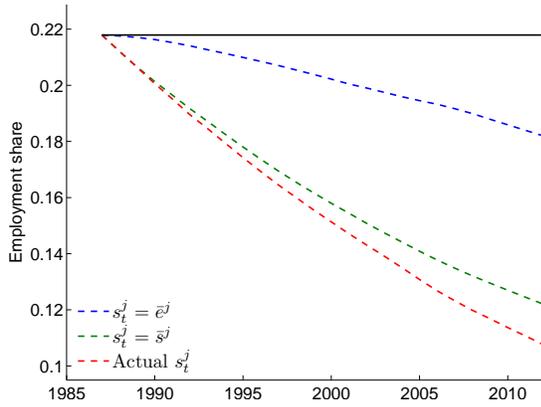
Notes: This figure depicts each sector's share of overall employment (solid lines) for 1987 to 2012. The dashed lines show sector employment shares that would have arisen had the age-specific growth and survival rates been flat at their sample averages. Source: U.S. Census Bureau Longitudinal Business Database and authors' calculations.

Figure A.1: Employment Shares by Industry (LBD)

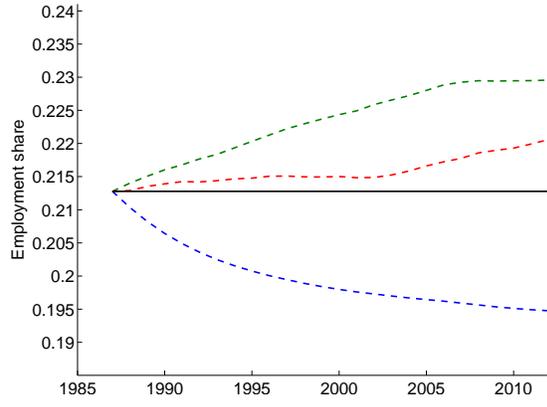


Notes: This figure depicts the maximum difference between NAICS4-level employment shares and the employment share that would have arisen had the age-specific growth and survival rates been flat at their sample averages. The left panel shows changes in the employment share while the right panel shows log changes. Source: U.S. Census Bureau Longitudinal Business Database and authors' calculations.

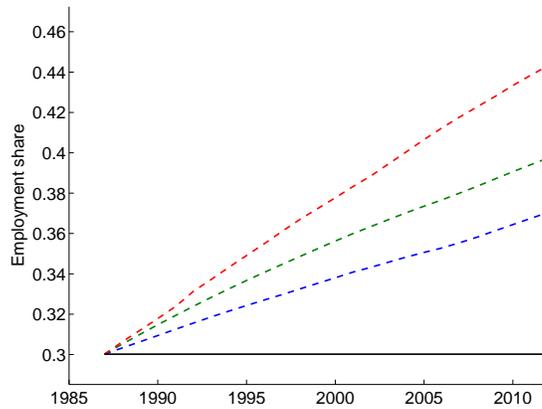
Figure A.2: Employment Shares by Industry (LBD)



(a) Manufacturing



(b) Retail Trade



(c) Services

Notes: This figure shows the effects of the entry and incumbent margins over the period 1987-2013. The solid black lines show the evolution of sector employment shares under the assumption that conditional life cycle dynamics are the same across sectors and the startup employment is allocated according to sectors' overall employment shares. The blue dashed lines highlight the contribution evolution of sector shares by allowing the conditional life cycle dynamics to vary across sectors. The green dashed lines show the evolution of sector shares if, in addition to incumbent life cycle effects, the startups were allocated according to the empirical allocation at the beginning of the sample period. Finally, the red dashed lines allow for the startup shares to vary over time as they do in the data. Source: U.S. Census Bureau Longitudinal Business Database and authors' calculations.

Figure A.3: Decomposition of sectoral reallocation