

The Industrial Revolution in Services *

Chang-Tai Hsieh

University of Chicago and NBER

Esteban Rossi-Hansberg

Princeton University and NBER

July 5, 2019

Abstract

The rise in national industry concentration in the US between 1977 and 2013 is driven by a new industrial revolution in three broad non-traded sectors: services, retail, and wholesale. Sectors where national concentration is rising have increased their share of employment, and the expansion is entirely driven by the number of local markets served by firms. Firm employment per market has either increased slightly at the MSA level, or decreased substantially at the county or establishment levels. In industries with increasing concentration, the expansion into more markets is more pronounced for the top 10% firms, but is present for the bottom 90% as well. These trends have not been accompanied by economy-wide concentration. Top U.S. firms are increasingly specialized in sectors with rising industry concentration, but their aggregate employment share has remained roughly stable. We argue that these facts are consistent with the availability of a new set of fixed-cost technologies that enable adopters to produce at lower marginal costs in all markets. We present a simple model of firm size and market entry to describe the menu of new technologies and trace its implications.

*We thank Feng Lin, Harry Li, and Jihoon Sung for extraordinary research assistance. We also thank Rodrigo Adao, Audre Bagnall, Dan Adelman, Pete Klenow, Raghuram Rajan, Richard Rogerson, and Chad Syverson for helpful discussions. The data from the US Census has been reviewed by the U.S. Census Bureau to ensure no confidential information is disclosed.

1. Introduction

Modern production relies on scale: The ability to use a technology to produce the same product or service innumerable times. In manufacturing industries, the steam-engine, electricity, and Ford’s assembly line, together with a number of other inventions, allowed firms to scale production in a single location. For many goods, the cost advantages of the larger scale overwhelmed the cost of transporting the goods to their final consumers. This ability to scale production in a single plant was, however, of little use outside of manufacturing. Producing many cups of coffee, retail or health services in the same location is of no value, since it is impractical to take them to their final consumers. Modern scale production in these sectors had to wait for a different technology, one that allowed firms to replicate the same production process in multiple locations close to consumers.

In this paper we argue that new ICT-based technologies have finally made it possible for firms outside of manufacturing to scale production over a large number of locations. The resulting expansion led to an increase in the national market share of top firms in many industries; a central fact about the US economy in the last three or four decades documented by Autor et al. (2017). This paper argues that this fact, among several others we document, is the result of a new industrial revolution that has taken place in many non-traded service sectors.

Consider Gawande (2012)’s account of how the Cheesecake Factory brought “chain production to complicated sit-down meals.” The Cheesecake Factory has invested in technologies that determine optimal staffing and food purchases for each restaurant and each day. The company also has a well-oiled process via which they introduce new items on their menu. This process starts in a centralized “kitchen” in Calabasas, CA – their R&D facility so to speak – where Cheesecake’s top cooks cull ideas for new dishes and “figure out how to make each recipe reproducible, appealing, and affordable.” The cooks in the R&D

facility then teach the new recipes to the kitchen managers of each restaurant at a bi-annual meeting in California. The kitchen managers then follow a finely honed procedure to teach the new recipes to the cooks in each restaurant. The roll out time, from the time the kitchen managers arrive at Cheesecake's central kitchen in California to when the new dishes are put on the menu in each restaurant, is 7 weeks.

The standardization of production over a large number of establishments that has taken place in sit-down restaurant meals due to companies such as the Cheesecake Factory has taken place in many non-traded sectors. Take hospitals as another example. Four decades ago, about 85% of hospitals were single establishment non-profits. Today, more than 60% of hospitals are owned by for-profit chains or are part of a large network of hospitals owned by an academic institution (such as the University of Chicago Hospitals).¹ As an example of the former, consider the Steward Health Care Group. This company was created by the Cerberus private equity fund in 2010 when it purchased 6 Catholic hospitals in Boston. In Gawande (2012)'s account, Cerberus' goal was to create the "Southwest Airlines of healthcare" by figuring out and codifying best practices and implementing these practices over a large scale. Gawande (2012) describes the scene in Steward's remote intensive care unit (ICU) in a Boston suburb that monitors the ICUs in all of Steward's hospitals:

"Banks of computer screens carried a live feed of cardiac-monitor readings, radiology-imaging scans, and laboratory results from ICU patients throughout Steward's hospitals. Software monitored the stream and produced yellow and red alerts when it detected patterns that raised concerns. Doctors and nurses manned consoles where they could toggle on high-definition video cameras that allowed them to zoom into any ICU room and talk directly to the staff on the scene or to the patients themselves."

Technologies such as the remote ICU has enabled Steward to provide consistent care in all the ICUs in its hospitals. By 2019, Steward had expanded from its 6

¹The employment-weighted share of multi-establishment hospitals in the Longitudinal Business Database increased from 15% in 1977 to 62% in 2013.

original hospitals in Boston to 36 hospitals located in 9 states and Malta.²

We argue that the rise in industry concentration is due to companies similar to the Cheesecake Factory and Steward Healthcare that have adopted technologies that enable them to standardize and scale up the delivery of non-traded *services*. In this sense, what has happened in non-traded services is akin to the industrial revolution unleashed by Henry Ford more than a hundred years ago when Ford introduced mass production to a car industry dominated by independent artisans. We use micro-data from the Longitudinal Business Database from 1977 to 2013 to document the following facts. First, we show that the phenomena of rising concentration documented by Autor et al. (2017) is only seen in three broad sectors – services, wholesale, and retail. As Autor et al. (2017) suggest, top firms have become more efficient over time, but our evidence indicates that this is only true for top firms in these three sectors. In manufacturing, for example, concentration has fallen.

Second, rising concentration in these sectors is entirely driven by an increase the number of local markets served by the top firms. Within a typical market served by a top firm in sectors with increasing concentration, we find that employment of top firms is either constant or falling. Specifically, we find that average employment per establishment of top firms *falls* in sectors with rising concentration. The same is true for employment of top firms in each county they serve.³

Third, we find that *total* employment rises substantially in industries with rising concentration. This is true even when we look at total employment of the smaller firms in these industries. This evidence is consistent with our view that increasing concentration is driven by new ICT-enabled technologies that ultimately raise aggregate industry TFP. It is not consistent with the view that

²Steward's hospitals are in Massachusetts, New York, Ohio, Florida, Arkansas, Louisiana, Texas, Arizona, and New Mexico. Steward also has two hospitals in Malta.

³In a related finding Rossi-Hansberg et al. (2018) show, using the National Establishment Time Series (NETS) data-set, that although sales and employment concentration have increased in most sectors, local concentration has fallen significantly, particularly in Services, Retail and Wholesale.

concentration is due to declining competition or entry barriers, as suggested by Gutierrez and Philippon (2017) and Furman and Orszag (2018), as these forces will result in a *decline* in industry employment.

Fourth, we show that the top firms in the economy as a whole have become increasingly specialized in narrow set of sectors, and these are precisely the non-traded sectors that have undergone an industrial revolution. At the same time, top firms have exited many sectors. The net effect is that there is essentially no change in concentration by the top firms in the economy as a whole. The “super-star” firms of today’s economy are larger in their chosen sectors and have unleashed productivity growth in these sectors, but they are not any larger as a share of the aggregate economy.

In order to make precise the type of technological change that we hypothesize is behind all these secular changes, we first propose a simple theory of firm size and market entry.⁴ Using the theory, we show that a key ingredient of the industrial revolution in services that we document is a new fixed cost technology that lowers the marginal cost in all markets served by the firm. The adoption decision of firms involves a trade-off between a proportional reduction in variable costs and an increase in the fixed cost of the firm. With a large enough fixed cost, only the most efficient firms find it profitable to adopt the new technology, which leads to more concentration in the industry. If firms can decide the extent to which they want to implement the new technological advances, more productive firms will adopt the new technology more fully, also leading to concentration in the industry. Firms that adopt the fixed cost technology serve new markets because the new technology makes it profitable to serve local markets that were previously not viable. Rising input prices due to the expansion of firms that adopt the new technology forces multi-product firms to leave other sectors where the new technology has not occurred, or where their relative productivity is low, so the net effect on total employment

⁴Our theory is reminiscent of Gaubert (2018), but it allows firms to serve multiple local markets, as Ramondo (2014) does in an international context.

of top firms is ambiguous.

The rest of the paper is organized as follows. Section 2 presents our empirical findings organized in five Facts. Section 3 presents the theory and derives the implications of the availability of a menu of new technologies offering combinations of fixed and variable costs. Section 4 discusses the previous literature from the perspective of our empirical findings and their conceptual interpretation and provides some initial computations of the contribution of the industrial revolution in services to aggregate TFP growth. Section 5 concludes.

2. Facts

We use micro-data from the U.S. Census Longitudinal Business Database (LBD). The LBD is based on administrative employment records of every nonfarm private establishment in the U.S. economy. The advantages of the LBD are its broad coverage and quality. The establishment-level variables we use are employment, county, industry (4-digit SIC or 6-digit NAICS), and the ID of the firm that owns the establishment. We restrict to sample to observations from 1977 to 2013 and drop establishments in the public, educational, and mining sectors. We classify each 4-digit SIC and 6-digit NAICS industry into 450 consistently defined industries from 1977 to 2013. Hereafter, when we refer to an industry we mean these 450 industries. We also group counties into metropolitan areas (MSAs) defined consistently over time. We use the firm ID to aggregate employment of establishments to a firm in an industry or in the aggregate economy.

Table 1 shows total employment, number of establishments, and number of firms in the first and last years of our sample. We highlights five facts from the LBD data.

Fact 1. Increase in Industry Concentration

Our first fact, shown in Table 2, is the increase in concentration in the average industry. Table 2 presents the employment share of top 10% firms in an

Table 1: LBD Summary Statistics

	Employment	Establishments	Firms
1977	63.2	4.1	3.4
2013	111.9	6.5	4.9

Note: Employment, establishments, and firms are in millions.

Table 2: Employment Share of Top 10% Firms

1977	1987	1997	2007	2013
67.0%	67.4%	69.2%	71.6%	72.4%

Note: Geometric average of employment share of top 10% firms in 450 industries, weighted by industry employment share.

average industry from 1977 to 2013.⁵ The employment share of the top 10% firms in an average industry increased from 68% in 1977 to 73% in 2013, with most of the increase occurring since the late 1980s. Table 2 echoes Autor et al. (2017)'s finding that the share of the top 4, 20, and 40 firms has increased in the average sector over a similar time period.⁶

Table 3 shows the change in the log employment share of the top 10% firms from 1977 to 2013 for each broad industry. The increase in average concentration is primarily driven by three sectors: wholesale, services, and retail, where the average employment share of the top 10% firms increased by about .15 log points between 1977 and 2013. In contrast, concentration in manufacturing *fell* over this time period.

⁵We chose 10% because the smallest number of firms in an industry in our sample is 10.

⁶Autor et al. (2017) use the micro-data from the Economic Censuses.

Table 3: Δ Top Firm Share, 1977-2013

All Industries	.093
Wholesale	.168
Services	.149
Retail	.131
Agriculture	.096
Construction	.071
Finance	.034
Utilities and Transportation	.012
Manufacturing	-.028

Note: Average Δ log employment share from 1977-2013 of top 10% firms in 4-digit industries within each large sector, weighted by average employment share of each industry.

Fact 2. Heterogeneity in Increase in Industry Concentration

Table 4 shows the distribution of the change in the employment share from 1977 to 2013 across our 450 industries. The first row shows the overall distribution. The 90-10 gap in the change in the top 10% share is about .35 log points. The second row shows the dispersion within broad sectors. Given the heterogeneity in the mean change in broad sectors shown in Table 3, it is not surprising that the residual dispersion is smaller than the overall dispersion. However, the residual dispersion is still sizable. The 90-10 gap in the change in residual industry concentration is almost .28 log points.

Fact 3. Industry Employment and Establishment per Firm Rise with Concentration

Table 5 shows total employment of the *industry* as a share of aggregate employment for industries where concentration fell (column 1) vs. industries where concentration increased (column 2). The first thing to see is that industries where concentration *fell* after 1977 accounted for 30% of aggregate employment

Table 4: Distribution of Δ Share of Top Firms, 1977-2013

	10%	25%	50%	75%	90%
Overall	-.057	.008	.078	.182	.290
Within 1-Digit Sector	-.031	.022	.078	.164	.253

Note: Distribution of Δ log employment share of top 10% firms in 450 industries from 1977 to 2013.

Table 5: Employment Share of Industries by Δ Concentration

	Δ Concentration < 0	Δ Concentration > 0
1977	30.2%	69.8%
2013	15.4%	84.6%


Note: Employment share of industries where employment share of top 10% firms fell between 1977 and 2013 (column 1) or increased between 1977 and 2013 (column 2).

in 1977. This reinforces the point that increased concentration by top firms is not universal across industries. More importantly, the second row shows that sectors with more top firm concentration are the ones where total industry employment (as a share of aggregate employment) has also grown. The employment share of industries with increased top firm concentration grew from 70% in 1977 to 85% in 2013.

Figure 1 plots the non-parametric relationship between the change in the employment share of the *industry* against the change in the share of the top firms in the industry (both from 1977 to 2013). There is a clear and robust relationship between industry growth and growth of the top firms in the industry. Table 6 shows the coefficient from an OLS regression of the relationship shown in Figure 1. The elasticity of the change in industry employment to the change

in industry concentration is a precisely estimated 3. The second row in Table 6 regresses the change in log of total employment of the *bottom* 90% of firms in the industry on the change in the employment share of the top 10% firms. The elasticity is smaller but still positive and significantly different from zero. This says that the growth in industry employment with increased concentration is broad based. Firms throughout the size distribution increase employment in sectors with increasing concentration, not only the top 10% firms in the industry, although by definition the increase is larger among the top firms.⁷

Figure 1: Δ Industry Employment vs. Δ Share of Top Firms



Note: Figure shows point estimate and 99% confidence interval of non-parametric regression of Δ log employment share of the industry (relative to aggregate economy) on Δ log employment share of top 10% firms in the industry (relative to industry employment), both from 1977-2013.

Figure 2 plots the non-parametric relationship between the change in the log

⁷The elasticity of the change in total employment of the bottom 90% of firms to the change in top firm concentration is negative (-0.78 ; $s.e.=0.42$) in the non-manufacturing sample. The elasticity of total employment in the industry to the change in top firm concentration is therefore correspondingly smaller (1.12 ; $s.e.=0.41$) in the non-manufacturing industries.

Table 6: Regression of Δ Industry Employment on Δ Concentration

$\Delta \log$ Employment Industry	3.059 (0.315)
$\Delta \log$ Employment Bottom 90%	0.876 (0.309)

Note: Δ Industry employment and concentration from 1977 to 2013. Concentration is log employment share of top 10% firms in industry.


number of establishments per firm (for all firms in the industry) and the change in the log employment share of the top 10% of firms in the industry. The OLS coefficient of this relationship is a precisely estimated 0.745 (shown in the first row in Table 7).⁸ The second row in Table 7 shows the elasticity of the number of establishments per firm for the bottom 90% of firms in the industry. The point estimate is smaller, but is positive and precisely estimated. Hence, all firms in industries with increasing concentration have more establishments, although the increase in the number of establishments per firms is more pronounced for top 10% firms.⁹

Fact 4. Industry Concentration is due to Extensive Margin Growth

We next show that the growth in industry concentration is mostly due to extensive margin growth by the top firms. The change in the employment share of the top firms in an industry can be decomposed into the contribution of growth on the extensive and intensive margins. For example, if we define a

⁸The elasticity of the change number of establishments per firm to the change in top firm concentration is identical (.745; s.e.=.137) in the non-manufacturing industries.

⁹Consistent with our findings, Cao et al. (2019) use data from the Quarterly Census of Employment and Wages between 1990 and 2015 to document an increase in the average number of establishments per firm. They do not relate the change to measures of industry concentration but show that the increase is more pronounced for larger firms and in the service sector.

Figure 2: Δ Establishments/Firm vs Δ Share of Top Firms

Note: Figure shows point estimate and 99% confidence interval of non-parametric regression of $\Delta \log \#$ Establishments/Firm of top 10% firms in the industry relative to all firms in the industry on $\Delta \log$ employment share of top 10% firms in the industry, both from 1977 to 2013.

market as an MSA, the decomposition is:

$$\Delta \log \frac{L_{top}}{L} = \Delta \log \frac{\#MSA_{top}}{\#MSA} + \Delta \log \frac{\frac{L_{top}}{\#MSA_{top}}}{\frac{L}{\#MSA}} \quad (1)$$

The first term in equation 1 is the contribution from growth in the number of MSAs of the top firms and the second term is the contribution from changes in employment per MSA of the top firms (both relative to all firms in the industry).¹⁰ The former we call extensive margin growth and the latter intensive margin growth. Table 8 shows the results of the variance decomposition for the relative number of cities vs employment per city (row 1), relative number

¹⁰ $\#MSA_{top}$ is defined as the sum of the number of MSAs in which the top 10% firms are present. This definition implies that particular MSAs will be counted multiple times if several top 10% firms operate in them. $\#MSA$ is calculated similarly.

Table 7: Regression of Δ Establishment/Firm on Δ Concentration

$\Delta \log$ Establishment/Firm Industry	0.745 (0.083)
$\Delta \log$ Establishment/Firm Bottom 90	0.105 (0.040)

Note: Δ Industry employment and concentration from 1977 to 2013. Concentration is log employment share of top 10% firms in industry.

of counties vs. employment per county (row 2), and relative number of establishments vs. employment per establishment (row 3). The first row shows that 93% of the growth in concentration comes from growth in the number of cities served by top firms, and only 7% comes from increased employment per city. The next two rows show that average employment per county and per establishment of top firms *falls*. So necessarily more than 100% of concentration growth has to come from the increase in the number of counties and establishments served by the top firms.¹¹

Figure 3 plots the non-parametric relationship between our three measures of extensive vs. intensive margin growth of top firms. The top panel shows the point estimate of a non-parametric regression of our three measures of extensive margin growth (# of MSAs, counties or establishments) of the top firms vs. overall growth in the employment share of the top firms. The slopes of all curves are positive, indicating that in industries where top firms have expanded the most, they have done so by expanding geographically into more establishments, counties, and MSAs. The slope increases as we adopt narrower definitions of a market. It is the smallest for MSAs and the largest for establishments. The

¹¹The share of the variation of the change in top firm concentration due to the extensive margin is identical when we consider only the non-manufacturing industries. The share due to extensive margin growth in the non-manufacturing industries are 0.877 (MSAs), 1.112 (counties) and 1.355 (establishments).

Table 8: Δ Concentration Due to Extensive vs. Intensive Margin Growth

	Extensive	Intensive
MSAs	0.932 (0.074)	0.068 (0.074)
Counties	1.161 (0.083)	-0.161 (0.083)
Establishments	1.435 (0.097)	-0.435 (0.097)

Note: Column 1 shows point estimates and standard errors of a regression of $\Delta \log$ # of MSAs/Counties/Establishments of top 10% firms relative to all firms from 1977-2013 on $\Delta \log$ employment share of top 10% firms from 1977-2013. Column 2 shows the same for a regression of $\Delta \log$ employment per MSAs/County/Establishment of top 10% firms relative to all firms from 1977 to 2013 on the same independent variable.

variation in the change in concentration across industries is basically entirely driven by variation across industries in the expansion of top firms into new markets.

The bottom panel in Figure 3 shows the non-parametric regression of intensive margin growth (employment per MSA, county and per establishment of the top firms). The figure shows that employment per establishment declines by more in industries where top firm concentration has increased by more. The same is true for employment per county, although the magnitude of the decline with respect to the change in industry concentration is smaller. It is only when we look at employment per MSA that we see more intensive margin growth for the top firms.

Table 9 probes for evidence on the relative size of the markets where the top firms in an industry have entered. Specifically, we measure the size of the local market as total employment (in all industries) in the county or MSA where a given firm in the industry operates an establishment. We then regress $\Delta \log$ size

Table 9: Regression of Δ Market Size on Δ Concentration

$\Delta \log$ Total Employment in County of Top Firms/All Firms	-0.274 (.103)
$\Delta \log$ Total Employment in MSA of Top Firms/All Firms	-0.205 (.061)

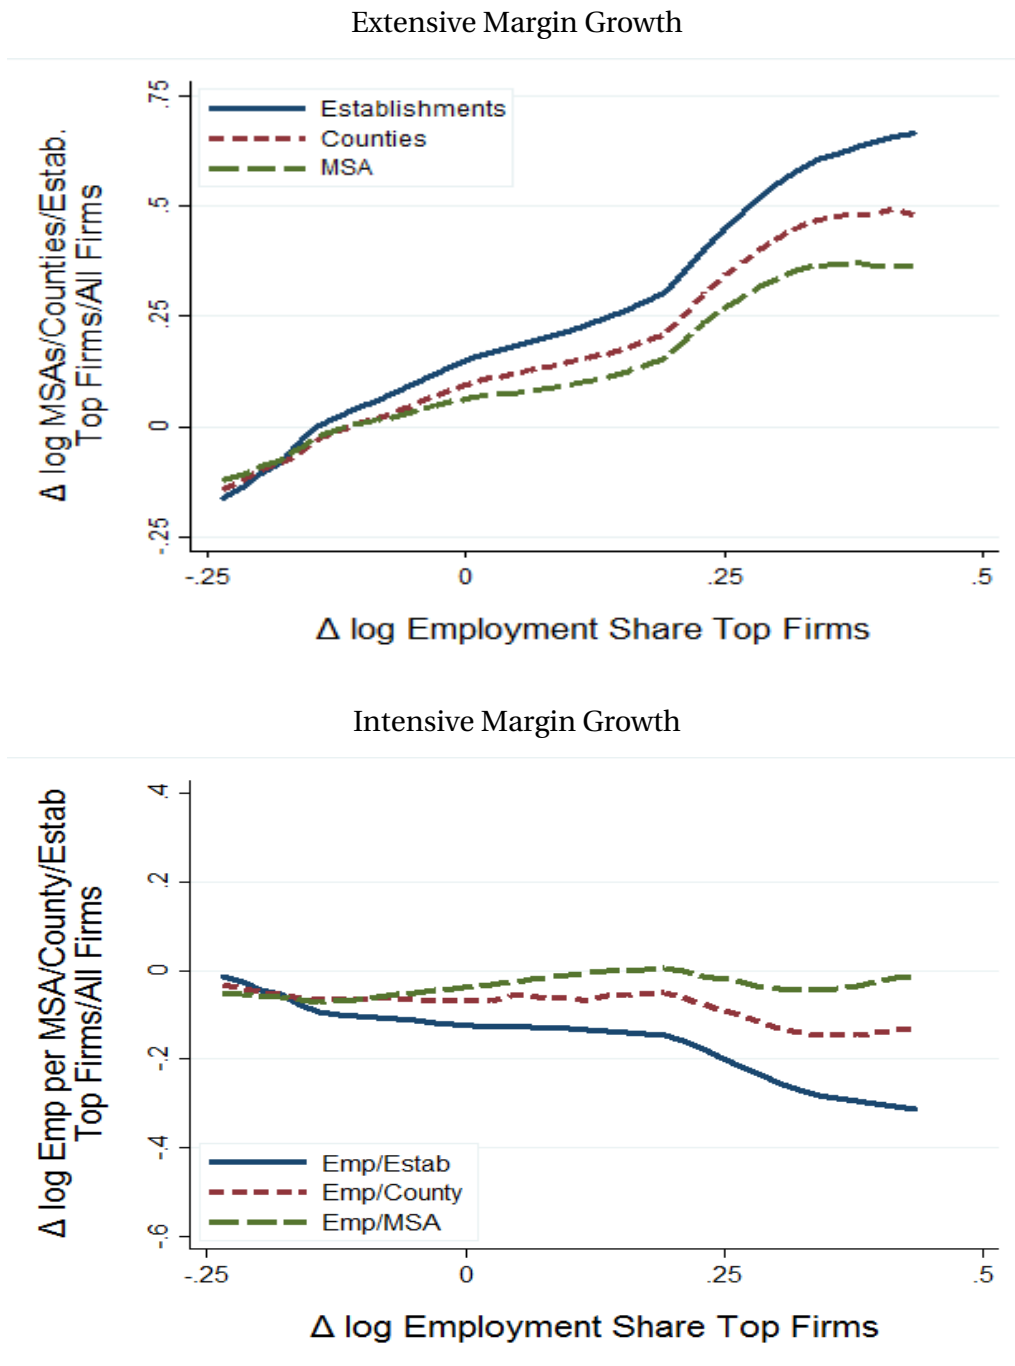
Note: Market size is total employment (in all industries) in the counties or MSAs where top firms operate in the industry relative to total employment in counties where all firms in the industry operate. Concentration is employment share of top 10% firms in industry. $\Delta \log$ market size and concentration from 1977 to 2013.

of the local market of an average top firm relative to all firms in the industry on $\Delta \log$ employment share of the top firm in the industry. Table 9 shows that the elasticity of the change in the relative size of the market of top firms with respect to the change in the market share of top firms is negative and precisely estimated. So top firms on average expand by entering into *smaller* local markets. Of course, the expansion patterns of specific industries might look different. For example, Holmes (2011) shows that Walmart grew by expanding into new local markets that are typically close to its headquarters and *larger* than its existing markets.

Fact 5. No Change in Concentration in Aggregate Economy

All our facts up to now have been about the top firms in an industry. Our last fact is about the top firms in the aggregate economy. The difference between the two is that top firms in the aggregate economy are in multiple industries. Table 10 shows the employment share of the top 0.1%, 0.01% and 0.001% of firms in the aggregate economy. Although the employment share of top firms in an average industry has increased substantially, the employment share of the top firms in the aggregate economy has not. Gutierrez and Philippon (2019) report a similar fact from data on publicly traded firms in Compustat. In our

Figure 3: Extensive vs. Intensive Margin Growth of Top Firms



Note: Top panel shows non-parametric regression of $\Delta \log \#$ of MSAs, Counties or Establishments of top 10% firms relative to all firms on $\Delta \log$ employment share of top 10% firms, both from 1977-2013. Thin solid line is 45° line. Bottom panel shows non-parametric regression of $\Delta \log$ employment per MSA, County or Establishment of top 10% firms relative to all firms on $\Delta \log$ employment share of top 10% firms, both from 1977-2013.

Table 10: Employment Share of Top Firms in Aggregate Economy

	1977	2013	Difference
Top 0.1%	39.3%	40.8%	1.5%
Top 0.01%	21.1%	22.2%	1.1%
Top 0.001%	8.0%	8.4%	0.4%

Note: Table shows the employment share of top 0.1%, 0.01%, and 0.001% firm in the overall economy.

Table 11: # Industries/Firm: Top Firms/All Firms

	1977	2013
Top 0.1%	7.1	4.6
Top 0.01%	21.4	9.4
Top 0.001%	35.0	17.2

Note: Table shows # industries per firm of average top firm (in overall economy) relative to the average firm.

LBD data, the employment share of the top .001% firms increased from 8% in 1977 to only 8.4% in 2013. Over the same period, the employment share of the top 10% firms in an industry increased by 5.4% (Table 2).

Table 11 reconciles the stable top firm share in Table 10 with increasing concentration at the industry level. The table shows the average number of industries per firm among the top firms relative to the number of industries per firm of all firms in the economy. Top firms produce in more industries than the average firm, but less so in 2013 compared to 1977. The number of industries of a top 0.001% firm (relative to the average firm) fell from 35 in 1977 to 17 in 2013.

Table 12: Employment/Industry: Top Firms/All Firms

	1977	2013
Top 0.1%	55.7	88.7
Top 0.01%	98.8	237.0
Top 0.001%	229.5	488.1

Note: Table shows employment per industry of average top firm (in overall economy) relative to the average firm.

The corresponding number for a top 0.01% firm is 21 industries in 1977 and 9 industries in 2013.

Table 11 shows that top firms are more specialized. Table 12 shows that top firms are also larger in the industries they have chosen to specialize in. The table shows average employment per industry of a top firm relative to average employment per industry of all firms. Employment per industry of a top 0.001% firm in 1977 was 230 times larger than that of the average firm. By 2013, the size gap had more than doubled to 488.

Finally, Table 13 shows that the industries that top firms have chosen to specialize in are primarily ones with growing concentration. Specifically, the table shows the employment of top firms in industries with above-median concentration growth as a share of total employment of the top firm. For example, 36% of employment of the top .001% firms in the overall economy in 1977 were in industries with growing concentration. By 2013, almost half of employment of the top .001% firms were in such industries. In summary, top firms are now more specialized, are larger in the chosen industries, and these are precisely the industries that have experienced concentration growth.

Table 13: Share of Industries with Δ Concentration $>$ Median in Employment of Top Firms

	1977	2013
Top 0.1%	43.2	58.8
Top 0.01%	42.7	57.0
Top 0.001%	35.6	48.5

Note: Table shows employment of top firms in industries with above median concentration growth between 1977 and 2013 as a share of the top firm's total employment.

3. A simple model of firm size and market entry

Our aim in this section is to propose a simple theory of firm production decisions that is rich enough to speak to the facts in the previous section. The main purpose of the theory is to define precisely a form of technological change and trace its implications. This new technology is, we believe, a good abstract description of the innovations that have driven the large secular changes we have documented in the U.S. economy between 1977 and 2013.

3.1. The model

Start by considering a firm i that produces a good j . The firm uses plants to produce in different locations n , out of a continuum of locations with mass N . The price of good j in location n is given by p_{jn} . Assume that the only way to serve market n is to put a plant there. The ability of firms to trade simply determines the size of these markets and therefore the distribution of markets defined below. A firm pays a fixed cost F_j (in units of the numeraire) to produce good j and another fixed cost f_n (in units of the numeraire, but with a magnitude indexed to the local wage) to set up an establishment in market n .

The firm's productivity A_{ij} applies to its establishments in all locations. Labor is the only factor of production, so a firm that hires L_{ijn} units of labor produces $Y_{ijn} = A_{ij}L_{ijn}$ units of output with local revenues given by $R_{ijn} = p_{jn}A_{ij}L_{ijn}$.

We assume that the productivity of potential entrants in the economy is given by Φ with full support on $(0, \infty)$.

Firm i 's profits from the production of good j are given by

$$\Pi_{ij} = \max_{\mathbb{N}_{ij}, L_{ijn}} \int_{\mathbb{N}_{ij}} [p_{jn}A_{ij}L_{ijn} - w_n(L_{ijn} + f_n)] dn - F_j,$$

where \mathbb{N}_{ij} is the set of markets in which firm i enters with product j .

Now suppose that demand is CES and firms compete monopolistically, then $p_{jn} = E_n Y_{ijn}^{-\frac{1}{\sigma}}$, where $\sigma > 1$ is the elasticity of substitution across varieties within an industry and E_n is a function of local expenditure determined in equilibrium.

Then, profits from good j are given by

$$\Pi_{ij} = \max_{\mathbb{N}_{ij}, L_{ijn}} \int_{\mathbb{N}_{ij}} \left[E_n (A_{ij}L_{ijn})^{1-\frac{1}{\sigma}} - w_n(L_{ijn} + f_n) \right] dn - F_j.$$

Conditional on serving market n , profit maximizing employment in the local market is given by:

$$L_{ijn} = A_{ij}^{\sigma-1} \left(\frac{(1 - \frac{1}{\sigma}) E_n}{w_n} \right)^{\sigma}.$$

The firm will serve market n if local profits are positive, namely,

$$\frac{E_n}{w_n} \left(A_{ij}^{\sigma} \left(\frac{(1 - \frac{1}{\sigma}) E_n}{w_n} \right)^{\sigma} \right)^{1-\frac{1}{\sigma}} - \left(A_{ij}^{\sigma-1} \left(\frac{(1 - \frac{1}{\sigma}) E_n}{w_n} \right)^{\sigma} \right) \geq f_n$$

or productivity is above a threshold α defined by

$$A_{ij} \geq \alpha \left(\sigma, f_n, \frac{E_n}{w_n} \right) \equiv \left(\frac{f_n}{\tilde{\sigma} \left(\frac{E_n}{w_n} \right)^{\sigma}} \right)^{\frac{1}{\sigma-1}} \quad (2)$$

where

$$\tilde{\sigma} \equiv \frac{(\sigma - 1)^{\sigma-1}}{\sigma^\sigma}.$$

Hence, the firm is more likely to enter a location where the market entry cost is lower, wages are smaller and total expenditures are larger.

Now consider the decision of the firm to enter industry j . The firm will enter if $\Pi_{ij} > 0$ or


$$\int_{n \text{ s.t. } A_{ij} > \alpha_{jn}} \left[w_n \tilde{\sigma} A_{ij}^{\sigma-1} \left(\frac{E_n}{w_n} \right)^\sigma - w_n f_n \right] dn \geq F_j,$$

where $\alpha_n \equiv \alpha(\sigma, f_n, E_n/w_n)$ is defined in (2). Since the local entry cost f_n is only a function of location, the threshold α is as well.

Suppose that the mass of markets with characteristic $\alpha < \alpha_n$ is given by $\Gamma(\alpha_n)$ with density $\gamma(\alpha_n)$. This distribution $\Gamma(\cdot)$ is determined by parameters, the set of available markets, and the distribution of E_n/w_n which is determined in equilibrium. If all markets are identical and there is free mobility across markets, then the distribution of E_n/w_n would be degenerate and therefore the distribution $\Gamma(\cdot)$ would be degenerate as well. If markets are different in terms of amenities, productivity, housing and other factors, or a variety of other frictions, the distribution of E_n/w_n across markets would not be degenerate, even with free mobility, apart from cases with extreme assumptions on the distribution of the rents of local factors. Limited or frictional mobility would also yield a non-degenerate distribution of E_n/w_n . Here, we stop short of specifying a fundamental model of the distribution $\Gamma(\cdot)$ to gain generality and simplify the exposition.

Figure 4 depicts the markets in which the firm will be active given a distribution $\gamma(\cdot)$ and, to simplify the illustration, a constant f_n and σ that we omit in the notation. Then, the only relevant local characteristic is the ratio E_n/w_n , where a higher ratio means that the market is more profitable. Suppose $\gamma(\cdot)$ is increasing in α and therefore decreasing in E_n/w_n . That is, more desirable markets are more scarce. Firms will choose to operate in all markets for which $A_{it} > \alpha(E_n/w_n)$, namely $E_n/w_n \geq \alpha^{-1}(A_{it})$ since $\alpha(\cdot)$ is decreasing in E_n/w_n .

Figure 4: Determination of Active Markers



We want to determine the total profits of a firm in industry j . Using (2), profits in market j can be written as

$$w_n f_n \left(\left(\frac{A_{ij}}{\alpha_{ij}} \right)^{\sigma-1} - 1 \right) \geq 0 \text{ for } A_{ij} \geq \alpha_{ij}.$$

The mass of markets where a firm with productivity A_{ij} enters is $\Gamma(A_{ij})$, which is the mass of \mathbb{N}_{ij} . Furthermore, if the fixed costs of opening a plant $w_n f_n$ are constant at f , the profits of a firm that enters are given by

$$\Pi \left(\underset{+}{A_{ij}}, \underset{-}{F_j}, f, \Gamma, \sigma \right) = f \int_0^{A_{ij}} \left(\left(\frac{A_{ij}}{\alpha} \right)^{\sigma-1} - 1 \right) \Gamma(d\alpha) - F_j.$$

Denote by $\underline{A}(F_j, f, \Gamma, \sigma)$ the productivity level such that

$$\Pi(\underline{A}, F_j, f, \Gamma, \sigma), F_j, f, \Gamma, \sigma = 0.$$

Since Π is increasing in A_{ij} , there is a unique such value. Therefore, active firms are such that $A_{ij} \geq \underline{A} \left(\underset{+}{F_j}, f, \Gamma, \sigma \right)$.

3.1.1. A parametric example

Suppose that the density of markets with characteristic α is $\gamma(\alpha) = \Omega \alpha^a / f$ for some $a > \sigma - 1$. This distribution implies that there are many markets where it is hard to enter. More so the higher a . $\Omega > 0$ is related to the availability of good markets to enter and is determined in general equilibrium by the level of wages and expenditures. Then

$$\begin{aligned} \Pi \left(\underset{+}{A_{ij}}, \underset{-}{F_j}, f, \Gamma, \sigma \right) &= f \int_0^{A_{ij}} \left(\left(\frac{A_{ij}}{\alpha} \right)^{\sigma-1} - 1 \right) \gamma(\alpha) d\alpha - F_j \\ &= \Omega \left[\int_0^{A_{ij}} A_{ij}^{\sigma-1} \alpha^{1+a-\sigma} d\alpha - \int_0^{A_{ij}} \alpha^a d\alpha \right] - F_j \\ &= \tilde{a} \Omega A_{ij}^{1+a} - F_j \end{aligned}$$

where

$$\tilde{a} = \left[\frac{\sigma - 1}{(2 + a - \sigma)(1 + a)} \right].$$

Note that we defined the level of the distribution $\gamma(\cdot)$ as Ω/f . Other definitions lead to a value of the constant \tilde{a} that depends on f . In the example, $\underline{A} \left(\underset{+}{F_j}, f, a, \underline{\Omega}, \sigma \right)$ is decreasing in Ω since having more profitable markets to enter implies a lower entry productivity threshold.

3.2. A new technology

Now suppose a new technology becomes available. The technology increases the fixed costs of producing a given good in exchange for a reduction in the variable cost (and leaves the fixed cost of creating plants f_n constant). Namely,

adopting the new technology results in an increase in fixed costs to $h^\eta F_j$ and an increase in productivity to hA_{ij} , for $h > 1$ and $\eta > 0$.

Firms will adopt the technology if

$$\Pi(hA_{ij}, h^\eta F_j, f, \Gamma, \sigma) \geq \Pi(A_{ij}, F_j, f, \Gamma, \sigma)$$

This condition can be rewritten as the sum of the profits of the firm in new markets plus the increased profits in old markets being greater or equal than the increase in the firm-product fixed cost relative to local fixed costs, namely,

$$\left[\underbrace{\int_{A_{ij}}^{hA_{ij}} \left(\left(\frac{hA_{ij}}{\alpha} \right)^{\sigma-1} - 1 \right) \Gamma(d\alpha)}_{\text{Value of new markets}} + \underbrace{\int_0^{A_{ij}} \left(\frac{A_{ij}}{\alpha} \right)^{\sigma-1} (h^{\sigma-1} - 1) \Gamma(d\alpha)}_{\text{Increased value of old markets}} \right] \geq (h^\eta - 1) \frac{F_j}{f}. \quad (3)$$

How do the benefits of adopting the new technology, the LHS of condition (3), change with firm productivity? Deriving the LHS of condition (3) with respect to A_{ij} yields,

$$\begin{aligned} \frac{\partial LHS}{\partial A_{ij}} &= \int_0^{hA_{ij}} (\sigma - 1) \left(\frac{h}{\alpha} \right)^{\sigma-1} A_{ij}^{\sigma-2} \Gamma(d\alpha) - \int_0^{A_{ij}} (\sigma - 1) \left(\frac{1}{\alpha} \right)^{\sigma-1} A_{ij}^{\sigma-2} \Gamma(d\alpha) \\ &= (\sigma - 1) A_{ij}^{\sigma-2} \underbrace{\int_{A_{ij}}^{hA_{ij}} \left(\frac{h}{\alpha} \right)^{\sigma-1} \Gamma(d\alpha)}_{>0, \text{ Larger value of new markets for top firms}} \\ &\quad + (\sigma - 1) A_{ij}^{\sigma-2} \underbrace{\int_0^{A_{ij}} \left(\frac{1}{\alpha} \right)^{\sigma-1} (h^{\sigma-1} - 1) \Gamma(d\alpha)}_{>0, \text{ Larger value of existing markets for top firms}} > 0 \end{aligned}$$

since, by Leibniz rule, the derivative of A_{it} in the limit of the integral is equal to zero. Hence, the gains from adopting the new technology increase with a firm's

productivity, while the costs (the RHS of condition 3) are fixed. This implies that there exists a threshold $H(F_j, f, \Gamma, \sigma, h, \eta)$ such that if $A_{ij} \geq H(F_j, f, \Gamma, \sigma, h, \eta)$ then firm i adopts the new technology. The function $H(F_j, f, \Gamma, \sigma, h, \eta)$ is increasing in F_j and η , as simple inspection of (3) indicates.

Some firms will decide not to adopt the new technology as long as

$$\underline{A}(h^\eta F_j, f, \Gamma, \sigma) / h > \underline{A}(F_j, f, \Gamma, \sigma),$$

since the potential entrant productivity distribution Φ has full support in $(0, \infty)$. This is not always the case. In our example with $\gamma(\alpha) = \Omega \alpha^a / f$, it requires $\eta > 1 + a$, since

$$\underline{A}(F_j, f, \Gamma, \sigma) = \left(\frac{F_j}{\tilde{a}\Omega} \right)^{\frac{1}{1+a}}$$

and so

$$\underline{A}(h^\eta F_j, f, \Gamma, \sigma) / h = h^{\frac{\eta}{1+a}-1} \left(\frac{F_j}{\tilde{a}\Omega} \right)^{\frac{1}{1+a}} > \underline{A}(F_j, f, \Gamma, \sigma),$$

if $\eta > 1 + a$, since $h > 1$. Namely, some firms do not adopt if the elasticity of fixed cost to h is larger than one plus the elasticity of the density of tougher markets (higher $1 + \alpha$). More generally, we need the increase in fixed costs in the new technology to be large enough. The arguments above have proven the following proposition:

Proposition 1 *Given the distribution Γ , fixed costs F_j and f , and elasticity of substitution σ , there exists a threshold $H(F_j, f, \Gamma, \sigma, h, \eta) > 0$ such that if $A_{ij} \geq H(F_j, f, \Gamma, \sigma, h, \eta)$ then firm i adopts the new technology. Thus, in equilibrium the highest productivity firms use the new technology (h) and the lowest productivity ones (if active) use the old technology. Firms that adopt the new technology are larger in employment and revenues, enter more markets, and make more profits.*

The model above is consistent with a number of the facts outlined in the previous section. In particular, it is consistent with Fact 1, since it leads to top firms adopting the new technology. It is also consistent with Fact 2 if the new

technology is not available in all sectors, and partially with Fact 3, since adopting firms expand the number of establishments per firm. However, because there is only one technology level h , in the model above firms only have a binary decision: adopt or not. Fact 3 shows us that in industries where concentration has increased, which we interpret as industries where a new technology h is available, the employment of the bottom 90% of firms also increased. It seems that in these industries, firms can choose to adopt the new technology with different degrees of intensity. We introduce this margin in our model next.

3.2.1. A menu of new technologies

Consider the case where firms can choose among a continuum of technologies indexed by $h \geq 1$. If they choose $h = 1$, they maintain the old technology, while if they choose $h > 1$, they choose a technology with the characteristics described above. As before, firm fixed costs are given by $h^\eta F_j$ for $\eta > 0$.

Proposition 1 shows that, given a single h , high productivity firms will adopt while others might not. Here we show that the gains from adopting a new technology h are not only increasing in firm productivity, but that the increase with firm productivity grows with h . That is, the cross-derivative of the LHS of condition (3) is positive. Since the RHS of condition (3) does not depend on h , we conclude that more productive firms choose technologies with higher h . Namely, if we denote by $\phi(A)$ the technology adopted by firms with productivity A , $\phi(A) \geq \phi(A')$ for $A \geq A'$.

The cross-derivative of the LHS of (3) is given by

$$\frac{\partial^2 LHS}{\partial A_{ij} \partial h} = (\sigma - 1) \gamma(h A_{ij}) + (\sigma - 1)^2 A_{ij}^{\sigma-2} \int_0^{h A_{ij}} h^{\sigma-2} \left(\frac{1}{\alpha}\right)^{\sigma-1} \Gamma(d\alpha) > 0.$$

Note also that the RHS of equation (3), $(h^\eta - 1)F_j/f$, is increasing in h for

$\eta > 0$ and the slope grows with η for $h > 1$. Namely,

$$\frac{\partial \frac{(h^\eta - 1)F_j}{f}}{\partial h} = \eta h^{\eta-1} \frac{F_j}{f} > 0,$$

and

$$\frac{\partial \frac{(h^\eta - 1)F_j}{f}}{\partial h \partial \eta} = [\eta h^{\eta-1} + \eta h^{\eta-1} \log h] \frac{F_j}{f} > 0 \text{ for } h > 1.$$

Furthermore, as $\eta \rightarrow 0$, $h^\eta - 1 \rightarrow 0$, and the derivatives above converge to zero. Hence, since the LHS of (3) is strictly positive for $h > 1$, there exists a threshold η_0 such that if $\eta < \eta_0$, low productivity firms also adopt a new technology, although with weakly lower h . In our example this threshold is such that $\eta_0 = 1 + a$, as proven above. Hence we have proven the following proposition:

Proposition 2 *If a firm with productivity A chooses a technology $h = \phi(A)$, then firms in the same sector with technology $A' \leq A$ choose technology $\phi(A') \leq \phi(A)$. That is, $\phi(\cdot)$ is a weakly increasing function. Furthermore, there exists a threshold η_0 such that if $\eta < \eta_0$, $\phi(A) > 1$ for all A .*

The results above allows us to match the fact, described as part of Fact 3 above, that smaller firms in industries that have experienced concentration also expand their employment and number of establishments serving local markets. We now present these results in detail for the case when we keep the distribution of local markets constant.

3.3. Equilibrium implications for a fixed distribution of local markets

Absent general equilibrium effects that determine the distribution Γ , firms that adopt a better technology h , which are the more productive firms, enter more locations. This is simply implied by $\Gamma(hA_{ij})$ increasing in hA_{ij} . Further-

more, when the new technology is available, the difference between the number of markets of productive and unproductive firms increases. Namely,

$$\left. \frac{\partial \Gamma(\phi(A)A)}{\partial A} \right|_{\phi(A)=1} = \gamma(A) [\phi'(A)A + 1] > \gamma(A) = \frac{\partial \Gamma(A)}{\partial A}$$

since $\phi'(A) > 0$ by Proposition 2.

Note also that absent general equilibrium effects, firms that adopt a better technology have larger establishments, e.g. L_{ijn} is increasing in A_{ij} . Note that, even though the new markets where the firm enters are less profitable, the increase in productivity due to the new technology implies that the marginal market has constant employment size. Employment size is given by

$$\begin{aligned} L_{ijn} &= (hA_{ij})^{\sigma-1} \left(1 - \frac{1}{\sigma}\right)^{\sigma} \left(\frac{E_n}{w_n}\right)^{\sigma} \\ &= (\sigma - 1) f \left(\frac{hA_{ij}}{\alpha_n}\right)^{\sigma-1}. \end{aligned} \quad (4)$$


Hence the marginal employment size, when $hA_{ij} = \alpha_n$ is $(\sigma - 1) f$, which does

not depend on h . This is illustrated in Figure 5. Since firms with better technology enter more markets and have more employees per establishment, their employment share necessarily increases when $h > 1$.

Finally, note that in industries where the new technology is very good (η is low) there will be more concentration since top firms will adopt a larger h . In those industries, less productive firms will also adopt a better technology, although with a lower h (see Proposition 2). Hence, in a multi-industry economy with elasticities of substitution across industries greater than one, employment of the whole industry will increase, including employment of the bottom firms.

The arguments above establish the following result for a single industry in partial equilibrium:

Proposition 3 *Given the distribution of markets Γ , the new menu of technologies results in more concentration of employment in more productive firms. Firms,*

Figure 5: The Effect of $h > 1$ on Employment and Market Entry

and more productive firms in particular, enter more markets and are larger in each of them. The effects are more pronounced for small values of η , with no effect if η is very large (since there is no adoption).

Hence, our model can match Facts 1 to 4 if we identify a market n in the theory as a city (MSA) in the data. In that case, national concentration in the industry rises because top firms enter more markets with a larger scale. As Table 8 shows, this is the case for MSA's but not for counties or establishments. Matching Fact 4 for these narrower geographic units, requires us to tweak the effects of the new technology further. We do so in the next section.

3.4. A technology that reduces local fixed costs too

The model above implies that the advent of the new technology brings increases in an adopter firm's average employment in a market. This prediction is consistent with the evidence if we interpret a market as a city (MSA). However, it is


counterfactual if we interpret a market as a county or a single establishment. To generate declines in the average employment size of adopters we need to allow the new technology to reduce local fixed costs as well.

Suppose that the new menu of technologies is as before but, in addition, local fixed costs are now given by $fh^{-\varphi}$. The exponent $\varphi > 0$ determines the extent to which fixed costs decline with the new technology. The exponent should depend on the definition of a market. For a large geographic area we might think that the cost did not change much beyond the overall firm fixed costs, and so $\varphi = 0$. For a smaller area, like a county or a single establishment, $\varphi > 0$, due to the ease in replicating standardized establishments (as exemplified by companies like Starbucks).

Note from the first line in equation (4) that, given E_n/w_n , the fixed cost does not affect establishment sizes directly. It does, however, determine entry into marginal markets and the size of the smallest establishment of the firm, which is given by $(\sigma - 1)fh^{-\varphi}$. For firms that choose $h > 1$, this implies that the smallest establishment size of the firm falls. Hence, with the new technology, average establishment sizes of existing firms fall for φ large enough. This is illustrated in Figure 6.

Finally it is useful to realize that if η and φ are high enough, average firm size necessarily falls, since firms will choose a small h and so the establishments in the best markets will only increase marginally in size, while (for φ large) the firm will add many new markets with smaller establishments. This is consistent with Fact 4 in the previous section, if we interpret a market as a county or the area served by a single establishment. We summarize these results in the following proposition.

Proposition 4 *Given the distribution of markets Γ , if the new technology also reduces fixed cost to $fh^{-\varphi}$, the minimum employment size of the firm's establishments falls, and average establishment size falls if η and φ are large enough.*

Figure 6: The Effect of $h > 1$ on Employment and Market Entry when $\phi > 0$ 

3.5. Multi-product firms and general equilibrium

Consider a firm that owns a collection of J_i technologies in a variety of industries j with productivity A_{ij} . Each variety requires the firm to pay a fixed cost F_j . When the firm obtains access to the new technology in each sector (with potentially different parameters η_j and φ_j in each sector), it will make a set of upgrading decision $\phi_j(A_{ij})$. In some of its industries the firm might decide to set $\phi_j(A_{ij}) = 1$ (e.g. if η_j is very high or its productivity is too low), or exit.

Denote total industry employment by L_j , and economy wide employment, which we assume fixed, by $\bar{L} = \sum_{j=1}^J L_j$. Then the labor market clearing condition is given by

$$\sum_{j=1}^J \int_{\underline{A}(F_j, f, \Gamma, \sigma)}^{\infty} \int_0^A \left((\sigma - 1) f \left(\frac{\phi_j(A) A}{\alpha} \right)^{\sigma-1} \right) \Gamma_j(d\alpha) \Phi(dA) = \bar{L}, \quad (5)$$

where

$$\int_0^{A_{ij}} \left((\sigma - 1) f \left(\frac{\phi_j(A_{ij}) A_{ij}}{\alpha} \right)^{\sigma-1} \right) \Gamma_j(d\alpha) = L_{ij}$$

denotes firm's i employment in industry j ,

$$\int_{\underline{A}(F_j, f, \Gamma, \sigma)}^{\infty} \int_0^A \left((\sigma - 1) f \left(\frac{\phi_j(A) A}{\alpha} \right)^{\sigma-1} \right) \Gamma_j(d\alpha) \Phi(dA) = L_j$$

denotes total employment in industry j , and

$$\sum_{j=1}^{J_i} \int_0^{A_{ij}} \left((\sigma - 1) f \left(\frac{\phi_j(A_{ij}) A_{ij}}{\alpha} \right)^{\sigma-1} \right) \Gamma_j(d\alpha) = \sum_{j=1}^{J_i} L_{ij}$$

denotes total firm employment. Note that we have recognized above that the distribution Γ_j varies by industry since it depends on expenditure in the industry, E_{jn} .

Overall, as long as (η_j, φ_j) is such that when the new technology is available some firms adopt, $\phi_j(A_{ij}) > 1$, in some industries, those firms will hire more employees for a given distribution of locations Γ_j , as described in Proposition 3. Hence, in equilibrium, the distribution Γ_j has to shift out to satisfy the labor market equilibrium condition, as apparent from condition (5). The shift happens through an increase in the level of wages, which in the specific example of Γ_j used above is represented by Ω . Hence, Ω decreases which selects some firms out of particular industries. As discussed above, $\underline{A}(F_j, f, \Gamma_j, \Omega, \sigma)$, is decreasing in Ω if it denotes the availability of good markets, as in the example.¹²

The implication is that if the bottom firms in the industry do not adopt, namely $\eta_j > \eta_0$, then some firms will exit. If, on the contrary, everyone adopts,

¹²This argument assumes that the distribution Γ_j is invariant to the new technology apart from its level Ω . This will be the case with free mobility (or heterogenous preferences) and constant proportional differences in amenities or productivity. Constant factor shares and constant differences in local factor availability will preserve the distribution under free mobility as well. As stated before, we stop short of developing the fully parameterized fundamental model that gives rise to the distribution Γ_j to gain generality and avoid some distracting notational details. A model with autarkic markets where workers cannot move at all would not be consistent with this assumption.

then all outcomes are possible and depend on the total employment expansion of industries, which in turn depends on the elasticity of substitution in consumption between industry aggregates. In what follows we assume that the bottom firms do not adopt (or adopt only marginally), namely, $\eta_j > \eta_0$.

As long as agents have CES preferences with elasticity of substitution in consumption greater than one across industry aggregates, these implications translate into overall industry employment. Hence, firms gain employment in industries where they are productive and the new technology is good (e.g. low η_j), and lose employment or exit industries in which they are not as productive and the technology change is also large (low η). The latter implication obtains since the top firm in that industry then upgrades more significantly. In industries with high η , nothing happens directly since there is no adoption. The general equilibrium effect makes employment and establishments smaller, makes firms exit local markets, and makes some firms exit the industry.

In sum, the advent of the new technology implies that firms specialize by leaving some markets and investing in others. The total employment size of the firm grows because of the new investments in its main industries, but declines due to tougher competition in its marginal industries and the decline in J_i . The overall implication on employment size is ambiguous. We summarize the findings in the following proposition:

Proposition 5 *The availability of the new technology makes firms specialize, invest, and grow their employment in industries where they are most productive, and reduce employment or exit industries where they are less productive (if $\eta_j > \eta_0$). The overall effect on firm employment is ambiguous.*

Consider a distribution of new technologies η_j across sectors J . Assume that in a group of sectors η_j is large enough so that as described in Proposition 3 firms do not adopt the new technology given the distribution Γ_j . Suppose also, that in other sectors, η is low enough such that some firms adopt. Because the marginal cost of these firms declines ($w_n/(hA_{ij})$), the price $p_{ijn} =$

$(\sigma/(\sigma+1))(w_n/(hA_{ij}))$ declines as well since with CES preferences and monopolistic competition markups are constant. Hence, the industry ideal CES price index, P_{jn} , in those industries falls relative to the price index of industries with higher η_j where firms choose a smaller $\phi_j(A_{ij})$. The result is that consumption expands in industries with low η_j and contracts in industries with high η_j . Note that if the elasticity of substitution in consumption across sectors is greater than one, this implies that output and employment shares in industries with low η_j expand relative to industries with high η_j , as evident from equation (5) and the fact that Γ_j is a decreasing function of expenditure in the industry, E_{jn} . We summarize the result in the following proposition:

Proposition 6 *If the technological innovation η_j is sufficiently heterogeneous across industries so that $\phi_j(\cdot)$ is not identical for all $j \in J$, and the elasticity of substitution of consumption across industries is greater than one, industries with low η_j increase their employment share, while those with high η_j contract.*

The results above show that in general equilibrium, our model is consistent with all the facts outlined in the previous section. We do not show that the employment of top firms does not grow, but the presence of effects in opposite directions implies that the total effect on firm size should be smaller than the effect on overall industrial concentration. Namely, firms get larger in their main sectors, but specialize and drop more marginal ones. The last proposition also illustrates the key parameter leading to the heterogeneity across industries discussed in Fact 2, namely, η_j .

4. Implications

Our interpretation of the secular changes the U.S. economy has experienced for the last four decades is that they amount to a new industrialization process. One that allows firms to expand geographically and deliver its goods and

services to costumers locally. We have argued that this evolution was the result of an underlying technological change that led to reductions in variable costs (and establishment-level fixed costs) in exchange for larger firm-level fixed costs. Section 3 presented a basic model of firm entry across markets to argue that this form of new technological possibilities can explain the array of industry and firm level facts we, and the previous literature, have documented.

Previous work has identified elements of the technological changes we underscore here. Sutton (1991) argues for the presence of new sunk cost technologies and describes their effect on market concentration, although he does not emphasize the increasing geographic scope of firms, nor their resulting specialization. Holmes (2011) focuses on a single firm (Walmart) and studies its geographic expansion and the resulting efficiency gains from investing in its distribution network and inventory system. These technological changes are certainly an example of the general evolution we have in mind. Hortaçsu and Syverson (2015) also provide a consistent description of the evolution of concentration, scale, and geographic expansion of the retail industry. In a recent paper, Ganapati (2018) studies the wholesale industry and argues also that the patterns of industry concentration and expansion of the warehouses and international input use of the top firms, are consistent with the use of technologies with higher fixed and lower marginal costs.

It is perhaps hard to set apart a number of concurrent technological changes, all of which are naturally intertwined. Information and communication technology (ICT) started in the 60's with the systematic use of corporate databases, then continued with the invention and expansion of personal computers, electronic communication technologies and the internet, and the invention and subsequent explosion in the use of smartphones (Hobijn and Jovanovic, 2001). A number of papers have studied the way in which ICT has changed the organization of production (Caroli and Van Reenen, 2001), the decentralization of decision making (Bresnahan et al., 2002), the span of control of managers (Rajan and Wulf, 2006; Garicano and Rossi-Hansberg, 2006), and the distribution of

firm sizes (Garicano and Rossi-Hansberg, 2004). More recently, Aghion et al. (2019b) study the growth implications of the ability of firms to manage more establishments due to improvements in ICT.

The form of technological change we emphasize here was certainly enabled by ICT, at least partly, which explains its timing. The examples of fixed-cost based technologies described in the introduction and the literature all have a component that was facilitated either by better data collection and analysis or by better communication and diffusion of information. It is undoubtedly the case that new business processes that reduce the cost of managing many different establishments requires easy communication, as well as cheap data gathering and processing. Managing many hospitals and exploiting the synergies between them would be impractical without the heavy use of ICT-based systems. Thus, ICT is an essential part of the industrialization of services. It is the general purpose technology, as defined by Rosenberg and Trajtenberg (2004), that has enabled the geographic expansion of firms (particularly in retail, services, and wholesale) by allowing them to replicate and control establishments dispersed across space.

Another phenomenon closely related to the new industrial revolution in services is the rise in intangible capital. As Haskel and Westlake (2017) document, intangible investments became increasingly important during the period of our analysis. As they argue, intangible investments in marketing, technology, information, or training, all facilitate scale and replication and as such amount to the use of new technologies with higher fixed (or sunk) costs. Hence, the rapid expansion of intangibles is a consequence of the type of technological change we uncover.

Finally, there is a large recent literature that has interpreted the increase in industry concentration as an indication of the augmented market power of top firms, perhaps facilitated by entry barriers or regulatory capture. This view has been supported by evidence that points to increasing profits and markups (Gutierrez and Philippon, 2017; De Loecker et al., 2018) and a decrease in mar-

ket dynamism (Decker et al., 2017). Together with a number of other papers in the literature (Hopenhayn et al., 2018; Syverson, 2019; Edmond et al., 2019), we argue that the industrialization of services that we document is technological, not institutional. Nevertheless, although we chose to model this process in a world with CES preferences and, therefore, fixed markups, in a model with variable markups these same technological changes could generate increases in markups. We chose not to focus on this dimension of the industrial revolution of services partly because the magnitude of the overall trend in markups is still controversial (see Traina, 2018, or Karabarbounis and Neiman, 2018) and partly because the geographic expansion of top firms leads to declines in local concentration (Rossi-Hansberg et al., 2018) that could enhance competition.

We end this section with a back-of-the-envelope calculation of the implication of the technological revolution in the service sector for aggregate TFP growth. With CES preferences across industries with elasticity of substitution ρ , aggregate TFP is given by

$$TFP_t = \left(\sum_{i=1}^N TFP_{jt}^{\rho-1} \right)^{\frac{1}{\rho-1}},$$

where TFP_{jt} is aggregate TFP of industry j in year t . In our model, aggregate TFP in an industry is determined by the set of firms active in each market. Hence, the entry of new top firms in a city increases industry TFP in the city, and by extension, aggregate industry TFP in the country as a whole.¹³

We proceed in three steps. First, the share of the change in aggregate TFP due to TFP growth of industry j is given by

$$\frac{TFP_{jt+1}^{\rho-1} - TFP_{jt}^{\rho-1}}{TFP_{t+1}^{\rho-1} - TFP_t^{\rho-1}}.$$

¹³The aggregate measure of TFP in an industry is defined as $TFP_{jt} = Y_{jt}/L_{jt} = R_{jt}/(P_{jt}L_{jt})$. In the data, L_{jt} and R_{jt} can be easily measured, but measuring P_{jt} is complicated since it requires the prices per unit of quality of goods and services sold in each market, plus a methodology to aggregate them across locations. These complications are particularly salient for the service industries, where quality adjusted prices are notoriously hard to measure.

Since the employment share of an industry is given by $L_{jt} = \left(\frac{TFP_{jt}}{TFP_t}\right)^{\rho-1}$, the contribution of TFP growth in an industry to aggregate TFP growth can be written as

$$\frac{L_{jt+1} \cdot \tilde{g}_t - L_{jt}}{\tilde{g}_t - 1},$$

where $\tilde{g}_t \equiv (\rho-1) \Delta \log TFP_t$. Note that we only need to know \tilde{g} and the employment shares of the industry for this calculation – we do not need to know the change in industry TFP. Intuitively, the change in the employment share of an industry captures the effect of the change in industry TFP relative to aggregate TFP.

Second, we need to parse out the contribution of the technological changes we described to industry employment. We have shown that the incidence of the industrial revolution shows up in the market share of the top firms in the industry. We can then calculate the share of aggregate TFP growth due to the industrial revolution by the contribution of TFP growth in industries where top firm concentration has increased:

$$\sum_{j \in \Delta Top > 0} \frac{L_{jt+1} \cdot \tilde{g}_t - L_{jt}}{\tilde{g}_t - 1}. \quad (6)$$

So given an estimate of \tilde{g}_t , we can calculate the contribution of the revolution we document on aggregate TFP growth with only data on industry employment and top firm shares.

The third step is to calculate \tilde{g} . One option is to take the BLS' official numbers of aggregate TFP growth. However, the industries where the employment share of top firms have grown, such as sit-down restaurants and hospitals, are precisely the sectors where measurement of quality is very difficult. Instead of using the BLS' numbers for aggregate TFP, we adopt the following approach. The growth rate of aggregate TFP is given by

$$\Delta \log TFP_t = \sum_{j=1}^N \alpha_{jt} \Delta \log TFP_{jt} \quad (7)$$

Table 14: Aggregate Growth due to “Service Revolution,” 1977-2013

	$\rho = 2$	$\rho = 3$	$\rho = 4$
Aggregate TFP Growth Rate	3.7%	2.4%	1.9%
Due to “Service Revolution”	90%	88%	87%

Note: The first row shows the weighted average of industry TFP from 1977 to 2013 imputed from the growth rate of the employment share of the industry relative to the employment share in manufacturing and BLS’ estimate of TFP growth in manufacturing. The second row uses equation 6 to impute the share of the service sector revolution in aggregate TFP growth from the change in employment shares of sectors where top firm concentration increased between 1977 and 2013.

where α_{jt} is the Sato-Vartia weighted average of the employment shares of the industry in year t and $t + 1$.¹⁴ Suppose there is an industry j (say manufacturing) for which we have a reliable number for industry TFP growth. Then the growth rate of TFP for another industry k is given by

$$\Delta \log TFP_{kt} = \frac{1}{\rho - 1} \Delta \log \left(\frac{L_{kt}}{L_{jt}} \right) + \Delta \log TFP_{jt} \quad (8)$$

So we make an assumption about the value of ρ to calculate the growth rate of industry TFP from equation 8. Then we use equation 7 to calculate the growth rate of aggregate TFP. After we do this, we then use equation 6 to calculate the contribution of the technological revolution to aggregate TFP growth.¹⁵

Table 14 shows the results from this calculation using the growth rate of manufacturing TFP from 1977 to 2013 from the BLS as the benchmark (1% per year). Using this number and $\rho = 2$, we get that aggregate TFP grew by 3.7% per

¹⁴The Sato-Vartia weight is defined as $\alpha_{jt} \equiv \frac{\frac{L_{jt+1} - L_{jt}}{\log L_{jt+1} - \log L_{jt}}}{\sum_{k=1}^N \frac{L_{kt+1} - L_{kt}}{\log L_{kt+1} - \log L_{kt}}}$.

¹⁵This methodology assumes homothetic CES preferences and, therefore, does not consider potential changes in employment shares as the level of income changes. Such effects, that would arise if consumers have non-homothetic preferences as in Buera and Kaboski (2012), would yield different results.

Table 15: Aggregate Growth due to “Service Revolution” by Period

	1977-1987	1987-1997	1997-2013
Aggregate TFP Growth Rate (Official)	0.5%	0.6%	1.9%
Aggregate TFP Growth Rate (Imputed)	2.9%	1.8%	2.7%
Due to “Service Revolution”	76%	86%	83%

Note: Row 1 shows the growth rate of aggregate TFP from 1977-1987, 1987-1997, and 1997-2013 as reported by the BLS. Row 2 shows aggregate TFP imputed from the change in employment share of the sector, the BLS’ official numbers of growth rate of TFP in manufacturing, and assuming $\rho = 3$. The third row shows the share of the service sector revolution in aggregate TFP growth calculated from equation 6.

year from 1977 to 2013. The imputed growth rate of aggregate TFP falls as we consider larger values of ρ . With $\rho = 4$ the imputed growth rate of aggregate TFP is 1.9%. All these numbers are substantially larger than the BLS’ official number for aggregate TFP growth from 1977 to 2013, which is 0.8%. The second row in Table 14 shows the share of aggregate TFP growth due to the “service sector revolution” computed from equation 6. The share of growth due to the “service revolution” is remarkably stable around 88% for all three values of ρ .

Table 15 shows the same calculation for three time periods: 1977-1987, 1987-1997, and 1997-2013. Here we assume $\rho = 3$ and that the BLS’ official estimates of TFP growth for manufacturing are accurate. The first row shows the BLS’ numbers for aggregate TFP growth in each period. The second row shows our imputed TFP growth rate. The gap between the official BLS’ numbers and our imputed numbers are the largest for 1977-1987, and smallest for 1997-2013. The contribution of “service revolution” is also the smallest in the 1977-1987 period.

What accounts for the difference between our estimates of aggregate TFP imputed from employment data and the official numbers? In the service sector, the BLS measures the price of real output as the price of a well-defined service in

Table 16: Missing Growth due to Entry of New Establishments, 1977-2013

	MSA	County
Non-Manufacturing	0.8%	0.8%
All Sectors	0.7%	0.6%

Note: The first row shows missing growth in the non-manufacturing sectors from 1977 to 2013. The second row shows missing growth in all sectors from 1977 to 2013. Calculations assume the elasticity of substitution between varieties in an industry is 3.

the same establishment. However, we have shown that the growth of top firms in the service sectors is entirely driven by entry of top firms into new markets. As argued by Aghion et al. (2019a), quality growth due to firm entry into new markets is not measured by the BLS.

Table 16 follows Aghion et al. (2019a) and measures the growth not captured by the BLS due to entry of new establishments. We differ from Aghion et al. (2019a) in that we measure missing growth in each locality, and then aggregate missing growth across all the localities. Specifically, for the set of industries present in a given location at the beginning and ending years (“incumbent industries”), we measure “missing growth” in each industry and location with Aghion et al. (2019a)’s formula. Specifically, for the incumbent industries in a city, “missing growth” due to firm entry can be measured as the weighted average of the product of $1/(\sigma - 1)$ and the change in the log employment share of incumbent establishments in each industry, where σ is the elasticity of substitution between varieties in an industry and the weights are the Sato-Vartia weights of the industry in the city.¹⁶

Second, in each location, primarily smaller cities, there is also entry of brand new service industries. For each new industry in a location, we calculate miss-

¹⁶We assume the elasticity of substitution between varieties σ is 3.

ing growth due to new industries as the product of $1/(\rho - 1)$ and the sum over all the new industries of the log of the inverse of 1 minus the employment share of the new industry in that city, where ρ is the elasticity of substitution across industries.¹⁷ Total missing growth in the given location is the sum of missing growth due to entry into incumbent industries and entry into new industries. Finally, we aggregate missing growth in each city using the employment share of each city in the final year.

The results are shown in Table 16 using two definitions of cities: MSAs and counties. As can be seen, there is more missing growth in the non-manufacturing sectors. Missing growth in non-manufacturing averages 0.8% per year from 1977 to 2013. Missing growth in all sectors, including manufacturing, is lower at .6% to .7% per year from 1977 to 2013. The difference is, obviously, missing growth in manufacturing, which is essentially zero. Thus, “missing growth” due to local market entry, can account for about half of the difference between the BLS estimates of aggregate TFP growth and our results in Table 14. The remaining difference measures the aggregate effect of the service revolution on the productivity of incumbent establishments and new establishments *within* these local markets, particularly those of top firms.

5. Conclusion

We show that new technologies have enabled firms that adopt them to scale production over a large number of establishments dispersed across space. Firms that adopt this technology grow by increasing the number of local markets that they serve, but on average are smaller in the markets that they do serve. Unlike Henry Ford’s revolution in manufacturing more than a hundred years ago when manufacturing firms grew by concentrating production in a given location, the new industrial revolution in non-traded sectors takes the form of horizontal expansion across more locations. At the same time, multi-product firms are forced

¹⁷We assume $\rho = 2$.

to exit industries where their productivity is low or where the new technology has had no effect. Empirically we see that top firms in the overall economy are more focused and have larger market shares in their chosen sectors, but their size as a share of employment in the overall economy has not changed.

The result of this new industrial revolution affecting many non-traded sectors is an increase in concentration and employment in these sectors. We see that employment increases even for the bottom 90% of firms in an industry with increasing concentration, suggesting that the new industrial revolution in these sectors is broad based, but obviously has a larger effect on the top 10% of firms.

We leaves three important questions for further work. First, it is important to say more precisely what this new technology is. The timing of these trends suggests that general purpose innovations in information and communication technologies have probably facilitated these fixed-cost based sectoral innovations. We also give some hints in our narrative in the introduction about the Cheesecake Factory and the Steward Health Care Group, but that only scratches the surface. We believe that a blend of quantitative and narrative accounts of this new industrial revolution, in the style of Chandler (1993)'s seminal work on the history of the industrial revolution in U.S. manufacturing, would be very useful. We hope that others (or perhaps we will) take up this challenge in the future.

Second, our story potentially has implications for the distribution of income and the distribution of employment of workers of different skills across locations. On the latter, the fixed cost technology is likely to be skilled worker intensive and top firms may choose to locate these services in larger and skill intensive cities. On the other hand, the expansion of top firms into smaller local markets improves the quality of local services, which may make these locations more attractive.

Third, we provided a back of the envelope calculation of the implication of increased concentration on aggregate TFP for the economy as a whole. It is possible to do something similar for intangible investment using the assumptions

of the model we laid out. But it should be clear that these calculations are only the beginning, and our hope is that more reliable numbers will be forthcoming in the future.

References

- Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter Klenow, and Huiyu Li, “Missing Growth from Creative Destruction,” *Working Paper*, 2019.
- , —, —, —, and —, “A Theory of Falling Growth and Rising Rents,” *Working Paper*, 2019.
- Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen, “The Fall of the Labor Share and the Rise of Superstar Firms,” *NBER Working Papers No. 23396*, 2017.
- Bresnahan, Timothy F., Erik Brynjolfsson, and Lorin M. Hitt, “Information Technology, Workplace Organization, and the Demand for Skilled Labor: Firm-Level Evidence,” *The Quarterly Journal of Economics*, 2002, 117 (1), 339-376.
- Buera, Francisco J. and Joseph P. Kaboski, “The Rise of the Service Economy,” *American Economic Review*, May 2012, 102 (6), 2540–69.
- Cao, Dan, Henry R. Hyatt, Toshihiko Mukoyama, and Erick Sager, “Firm Growth Through New Establishments,” *Working Paper*, 2019.
- Caroli, Eve and John Van Reenen, “Skill-Biased Organizational Change? Evidence from A Panel of British and French Establishments,” *The Quarterly Journal of Economics*, 2001, 116 (4), 1449-1492.
- Chandler, Alfred D., Jr., *The Visible Hand: The Managerial Revolution in American Business*, Harvard University Press, 1993.
- Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda, “Declining Dynamism, Allocative Efficiency, and the Productivity Slowdown,” *American Economic Review*, 2017, 107 (5), 322–326.
- Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu, “How Costly Are Markups?,” *Working Paper*, 2019.
- Furman, Jason and Peter Orszag, “Slower Productivity and Higher Inequality: Are They Related?,” *Working Paper Series WP18-4*, 2018.

- Ganapati, Sharat, “The Modern Wholesaler: Global Sourcing, Domestic Distribution, and Scale Economies,” *Working Paper*, 2018.
- Garicano, Luis and Esteban Rossi-Hansberg, “Inequality and the Organization of Knowledge,” *American Economic Review*, 2004, 94 (2), 197–202.
- and —, “Organization and Inequality in a Knowledge Economy,” *The Quarterly Journal of Economics*, 2006, 121 (4), 1383–1435.
- Gaubert, Cecile, “Firm Sorting and Agglomeration,” *American Economic Review*, November 2018, 108 (11), 3117–53.
- Gawande, Atul, “Big Med,” *New Yorker*, 2012.
- Gutierrez, German and Thomas Philippon, “Declining Competition and Investment in the U.S.,” *NBER Working Paper No. 23583*, 2017.
- and —, “Fading Stars,” *NBER Working Paper No. 25529*, 2019.
- Haskel, Jonathan and Stian Westlake, *Capitalism without Capital: The Rise of the Intangible Economy*, Princeton University Press, 2017.
- Hobijn, Bart and Boyan Jovanovic, “The Information-Technology Revolution and the Stock Market: Evidence,” *American Economic Review*, 2001, 91 (5), 1203–1220.
- Holmes, Thomas J., “The Diffusion of Wal-Mart and Economies of Density,” *Econometrica*, 2011, 79 (1), 253–302.
- Hopenhayn, Hugo, Julian Neira, and Rish Singhania, “From Population Growth to Firm Demographics: Implications for Concentration, Entrepreneurship and the Labor Share,” *NBER Working Papers No. 25382*, 2018.
- Hortaçsu, Ali and Chad Syverson, “The Ongoing Evolution of US Retail: A Format Tug-of-War,” *Journal of Economic Perspectives*, 2015, 29 (4), 89–112.
- Karabarbounis, Loukas and Brent Neiman, “Accounting for Factorless Income,” *NBER Working Papers No. 24404*, 2018.

- Loecker, Jan De, Jan Eeckhout, and Gabriel Unger, “The Rise of Market Power and the Macroeconomic Implications,” *Working Paper*, 2018.
- Rajan, Raghuram and Julie Wulf, “The Flattening Firm: Evidence from Panel Data on the Changing Nature of Corporate Hierarchies,” *The Review of Economics and Statistics*, 2006, 88 (4), 759–773.
- Ramondo, Natalia, “A quantitative approach to multinational production,” *Journal of International Economics*, 2014, 93 (1), 108 – 122.
- Rosenberg, Nathan and Manuel Trajtenberg, “A General-Purpose Technology at Work: The Corliss Steam Engine in the Late-Nineteenth-Century United States,” *The Journal of Economic History*, 2004, 64 (1), 61–99.
- Rossi-Hansberg, Esteban, Pierre-Daniel Sarte, and Nicholas Trachter, “Diverging Trends in National and Local Concentration,” *NBER Working Papers No. 25418*, 2018.
- Sutton, John, *Sunk Costs and Market Structure: Price Competition, Advertising, and the Evolution of Concentration*, MIT Press, 1991.
- Syverson, Chad, “Macroeconomics and Market Power: Facts, Potential Explanations and Open Questions,” 2019.
- Traina, James, “Is Aggregate Market Power Increasing? Production Trends Using Financial Statements,” *New Working Paper Series No. 17*, 2018.