

Productivity Dispersion and Structural Change in Retail Trade

Dominic Smith
U.S. Bureau of Labor Statistics

G. Jacob Blackwood
Amherst College

Michael D. Giandrea
U.S. Bureau of Labor Statistics

Cheryl Grim
U.S. Census Bureau

Jay Stewart
U.S. Bureau of Labor Statistics

Zoltan Wolf
New Light Technologies

June 28, 2024

Abstract

The retail sector has changed from a sector full of small firms to one dominated by large, national firms. We study how this transformation has impacted productivity levels, growth, and dispersion between 1987 and 2017. We describe this transformation using three overlapping phases: expansion (1980s and 1990s), consolidation (2000s), and stagnation (2010s). We document five findings that help us understand these phases. First, productivity growth was high during the consolidation phase but has fallen more recently. Second, entering establishments drove productivity growth during the expansion phase, but continuing establishments have increased in importance more recently. Third, national chains have more productive establishments than single-unit firms on average, but some single-unit establishments are highly productive. Fourth, productivity dispersion is significant and increasing over time. Finally, more productive firms pay higher wages and grow more quickly. Together, these results suggest that the increasing importance of large national retail firms has been an important driver of productivity and wage growth in the retail sector.

Keywords: retail, reallocation, business cycles, productivity dispersion

JEL Classification: D24, E24, L81

* Corresponding Author: Dominic Smith smith.dominic@bls.gov. Any opinions and conclusions expressed herein are those of the authors and do not represent the views of the U.S. Census Bureau or the Bureau of Labor Statistics. The Census Bureau has ensured appropriate access and use of confidential data and has reviewed these results for disclosure avoidance protection (Project 7526913: CDBRP-FY23-CED006-0016, CDBRP-FY24-CED006-0014). This paper provides a summary of research results. The information is being released for statistical purposes, to inform interested parties, and to encourage discussion of work in progress. The paper does not represent an existing, or a forthcoming new, official BLS statistical data product or production series. Blackwood was also a part-time employee of the U.S. Census Bureau at the time of the writing of this paper. We thank Emek Basker, Cindy Cunningham, Matthew Dey, Lucia Foster, John Haltiwanger, Michael Mandel, Rachel Nesbit, Marshall Reinsdorf, Kirk White, Sabrina Pabilonia, and audience members at the 2023 FSRDC Conference, the November 2023 BLS Technical Advisory Committee Meeting, the 2024 Society for Government Economists, University of Western Ontario, the 2024 Comparative Analysis of Enterprise Data Conference, and 2024 BLS-Census Research Workshop for their helpful comments. This paper is also available as BLS working paper 570.

1. Introduction

The U.S. retail trade sector has transformed since the 1980s, when single-unit firms accounted for the majority of retail sales. This transformation was driven by the expansion of national chains, resulting in a significant decrease in the share of smaller firms (Foster, Haltiwanger, Klimek, Krizan, Ohlmacher, 2016). This shift in structure greatly changed the production process in retail. National chains operate similar stores in many markets, allowing them to reach many consumers and reduce costs through streamlined purchasing and distribution. Additionally, an important byproduct of the rise of national chains was the introduction of new technologies, such as universal product codes, self-checkout machines, and logistics management software. The transition to national chains and shifting production methods have been accompanied by moderate labor productivity growth—2 percent per year on average between 1987 and 2017. However, productivity growth was uneven across industries. Some industries experienced high productivity growth, e.g., nearly 10 percent in Electronic Shopping and Mail-Order Houses (North American Industry Classification System or NAICS 4541 which includes e-commerce establishments), while more traditional industry groups such as Grocery Stores (NAICS 4451) experienced almost no growth despite the adoption of new technologies such as self-checkout machines.

We use data from the Census Bureau’s Census of Retail Trade (CRT) between 1987 and 2017 to explore the implications of the ongoing structural change in the retail sector for measured productivity. We draw out common patterns while also considering heterogeneity in productivity growth across industries, heterogeneity in productivity across establishments, and the contribution of reallocation to productivity growth. In our analysis, we consider particular industries that are influential in the sector while also paying special attention to the role of multi-units of different sizes. We further explore the relationship between productivity, a standard measure of firm performance, and key outcomes like wages and employment growth.

Our discussion is centered around five main findings, which we briefly list here and expand on in the paper.

Finding 1: Productivity growth has been uneven. Growth in productivity has been uneven both across industries and over time. The time variation exhibits a general hump-shaped pattern, with the highest growth rates in the middle of our sample. While this generally holds across industries, there are substantial differences in average growth rates and dynamics across key industries.

Finding 2: The drivers of productivity growth have changed. We explore the role of continuing establishments, entrants, and exits in productivity growth. We further distinguish between the role of expansion of existing firms by opening new establishments and completely new firms. Establishment entry is an important driver of growth, especially the entry of establishments from existing firms. This finding is consistent with work studying the retail sector between 1987 and 1997 (Foster, Haltiwanger, and Krizan, 2006). However, much of the variation in aggregate productivity growth can be accounted for by productivity growth at continuing establishments, which became the most important driver of productivity growth in the last decade of our sample.

Finding 3: Multi-unit firms (chains) are more productive than single-unit firms on average. We describe the distribution of establishments over firm structure and productivity. The distribution of establishments shifted from single-unit and local firms to national firms. These patterns are consistent with national firms expanding by adding new establishments and local and regional firms growing by adding establishments to become national firms. We find that chains are generally more productive across our sample. Furthermore, continuing establishments at national chains experienced high productivity growth from 1997-2007, a period that coincides with the highest productivity growth in the sector. Still, single-unit firms are often very productive; almost 20% of them are in the top quintile of the productivity

distribution.

Finding 4: Productivity dispersion is large. We document general patterns in productivity dispersion within industries and find evidence of substantial dispersion, with significant heterogeneity in *within*-industry dispersion *across* industries. Unweighted within-industry dispersion is generally rising, while hours-weighted dispersion is falling. We further document dispersion in the tails of the productivity distribution and find substantial increases.

Finding 5: Firms with higher productivity pay higher wages and grow more quickly. We examine whether productivity, a standard measure of firm performance, translates into other key labor market outcomes. We find that higher productivity is associated with higher average earnings levels and higher employment growth.

These findings extend previous knowledge of firm and establishment dynamics in the retail sector and expand our understanding of productivity dispersion beyond manufacturing. Although the broad narrative in retail is one of expansion and consolidation of national chains and some potential improvements to production processes, our findings suggest this phenomenon has translated unevenly into productivity growth across industries and establishments. Industry-level heterogeneity suggests care should be taken when specifying mechanisms for the retail sector, and perhaps more focus should be placed on industry-specific mechanisms.

The paper's main contribution is the documentation of productivity dispersion in the retail sector. These findings shed light on how the shift in retail since the 1980s has impacted the distribution of establishments. A priori, it is unclear whether the increasing importance of national retail firms would result in decreased or increased dispersion across establishments within an industry. On the one hand, stores within an industry may have become more similar as a few national chains have come to dominate many retail industries. If stores within a chain have similar productivity, we expect within-industry dispersion of establishment-level

productivity to decrease, while the dispersion of firm-level productivity could increase or decrease. The share of variance accounted for by within-firm variation provides some evidence to this effect. In 2017, 12 percent of the variation in productivity in our sample was due to within-firm variation.¹ On the other hand, entering chains may be much more productive than incumbent stores, which may have survived, despite their low productivity, because of their location or because they operate in niche markets. In this case, we would expect productivity dispersion to increase. Which forces dominate is important for understanding the allocation of inputs and the distribution of productivity increases. Our findings suggest that chains have held back dispersion to some degree, as activity-weighted dispersion does not rise while unweighted dispersion rises. Still, substantial dispersion persists despite the move to chains, and small firms can be found across the productivity distribution.

The remainder of the paper describes how we measure productivity and then documents trends in productivity levels and dispersion.

2. Measuring Establishment Productivity

One of our primary goals is to create statistics that provide insight into productivity in the retail sector that complement the official BLS industry-level labor productivity measures. While industry-level productivity measures provide crucial information about different parts of the economy, they do not tell us anything about what goes on within industries or the mechanisms behind productivity changes. We construct measures of within-industry productivity dispersion for industries in the retail sector using microdata.² There are many reasons why our micro-aggregated estimates might not match the official estimates. So, our

¹ Within-firm productivity dispersion accounted for a larger share of overall productivity dispersion in 2017 (12%) than in 1987 (9%), which is partly mechanical as market share of multi-units rises. The point remains: labor productivity dispersion is largely due to between-firm variation.

² The BLS and the Census Bureau currently jointly release statistics on within-industry productivity dispersion for the manufacturing sector, Dispersion Statistics on Productivity. See Bureau of Labor Statistics, U.S. Department of Labor, and the Census Bureau, U.S. Department of Commerce, 2023.

first step is to construct estimates based on establishment-level data using a similar methodology to the official estimates.³

2.1 Data Description

To measure establishment-level labor productivity, we combine restricted-use establishment-level microdata files from the Census Bureau with public-use industry-level data on prices and hours worked from the BLS. Our establishment-level data come from the Census of Retail Trade (CRT). The CRT is collected every five years in years ending in “2” and “7” as part of the Economic Census. The frame includes all retail establishments of multi-unit firms and a sample of single-unit retail firms, and it collects data about the activity of establishments including their sales and number of employees.⁴ To increase the coverage of single-unit firms, the Census Bureau uses information from administrative records to impute sales data. We link establishments over time and assign industries to these establishments using the Longitudinal Business Database (LBD) (see Chow et al., 2021).

Ideally, we want to construct an output measure from the data that matches the BLS measure as closely as possible. The underlying data source for our output measure, the CRT, is the same as the one used by the BLS. However, the BLS uses the aggregate tabulations of the CRT, which are based on industry codes available in that year. When industry codes change, BLS re-estimates sales by industry to create a consistent series over time. In contrast, our data contain NAICS codes for each establishment each year (Fort and Klimek, 2018). These codes are based on more detail than is available in the aggregate data. These differences lead our estimate of sales to differ more before the transition from Standard Industrial Codes (SIC) to NAICS codes in 2002, and in industries where NAICS codes changed between 2002 and 2017.

³ We describe how the BLS constructs industry-level estimates of productivity growth in the appendix.

⁴ The CRT does not collect information about capital stocks or investment.

2.2 Productivity Measurement

We measure labor input as total hours worked. We calculate total annual hours worked by multiplying total employment by the average annual hours worked per employee in the most detailed NAICS code available in the Current Employment Statistics (CES) data. Thus, all the between-establishment variation in hours worked in a 4-digit industry is due to variation in employment across establishments in the industry and variation in hours across 6-digit industries within the 4-digit industry. We calculate establishment-level (denoted by subscript e) log labor productivity as:

$$LP_e = \ln\left(\frac{\text{real sales}_e}{\text{hours}_i \times \text{employment}_e}\right) \quad (1)$$

where i is the most detailed industry code for which data are available.

There are two concerns with our retail employment measures: (1) employment can be seasonal and CRT employment is measured during the week of March 12, and (2) many employees may be part-time, but the CRT does not differentiate between full- and part-time employees. Because the number of full-time employees in the CRT in March may not reflect each establishment's use of labor inputs over the entire year, we also measure establishment-level log wage productivity as real sales divided by total payroll for the year deflated by the CPI. To illustrate, consider two establishments with equal sales. If one establishment uses one full-time worker and another uses two half-time employees, productivity measured by CRT employment will be much smaller for the establishment with half-time employees. On the other hand, the payroll of the two establishments would be similar, which would imply similar levels of productivity.⁵

⁵ The wage bill is not a perfect measure as it is contaminated by differences in the occupational mix, overtime, and other factors that affect wages such as geography.

3. Five Findings about Retail Productivity

In this section, we present five findings about productivity in the retail trade sector that illustrate essential elements of structural change in retail. Together, these findings provide insight into the three phases of the retail trade transformation. In the 1980s and 1990s, productivity dynamics were determined by the expansion of national chains. In the 2000s, expansion became less important as the expansion of national chains slowed, and those chains consolidated their market position. In the 2010s, we see evidence of stagnation following a recession where low-productivity establishments were more likely to exit.

In addition to looking at the retail sector as a whole, we also highlight five 4-digit NAICS industries that illustrate the heterogeneity within retail trade. Studies have emphasized industry trends, including product (changes in the kinds of output sold) and process (changes in how inputs are used to produce output) innovation, that possibly do not contribute to productivity growth.

We study two industries with major changes in products, two industries with major changes in process, and one with changes in both. For example, in the Electronics and Appliance Stores (4431) industry, which accounts for three percent of the retail sector, stores sell products such as computers and cell phones that have become increasingly important during our study while moving away from appliances and small electronics. General Merchandise Stores, including Warehouse Clubs and Supercenters (4523) is an important and growing industry whose evolution has been extensively studied through the mid-2000s (Basker, 2005; Holmes, 2011). Stores in this industry sell a wide range of products and compete with stores in many other industries (Smith and Ocampo, Forthcoming). These stores began selling groceries during our study and have gradually increased the number of product lines they carry (Basker, Klimek, and Van, 2012).

Turning to changes in process, Grocery Stores (4451) are one of the largest retail

industries and account for 16 to 21 percent of the retail sector during our study. This industry has experienced technological advancements such as universal product codes, self-checkout machines, and computerized inventory control (Basker and Simcoe, 2021; Basker, Becker, Foster, White, and Zawacki, 2019). Clothing Stores (4481) account for about 5 percent of the retail sector and have many large retail firms (U.S. Census Bureau, 2017). While product lines are clearly always evolving in clothing, during this period, the clothing manufacturing process changed substantially. Clothing was increasingly manufactured abroad and imported by either retailers themselves (in the case of large firms) or through intermediaries (Basker and Van, 2010; Smith, 2018). This could have spurred productivity growth at large firms if it increased demand for large firms and allowed them to keep employees busier throughout the workday.

Electronic Shopping and Mail-Order Houses, or E-Commerce, (4541) has changed significantly in both process and products. It is a growing industry that has changed from mail-order to online shopping, which clearly represents innovation in how consumers access retail services.⁶ Adoption of e-commerce has been uneven across products which means that over time the product mix within the industry has changed substantially (Hortacsu and Syverson, 2015).

While there are likely common mechanisms driving productivity dynamics across these five industries, they were also subjected to very different forces between 1987 and 2017.


3.1 Productivity Growth Has Been Uneven

Productivity growth within the retail sector has varied substantially over time and across industries. Figure 1 shows annualized productivity growth rates over the sample period for retail trade. Retail productivity growth increased gradually through 2002 and peaked in the 1997-2007 period with an average annual growth rate of 3 percent over this 10-year period.

⁶ Specifically, 4541 was eliminated and establishments in these industries were reclassified into other industries based on the products they sell. See <https://www.bls.gov/respondents/ars/2022-naics.htm>.

Productivity fell during the Great Recession and recovered somewhat in the 2012-2017 period, although growth was still below the long-term trend.

Figure 1: Average annual productivity growth in retail industries




Notes: Productivity growth is presented as the average compound annual growth rate over each five-year period. The growth rates are calculated using data on sales and employment from the Census of Retail Trade, in addition to BLS data on price deflators and industry group average hours worked. Productivity is calculated for each establishment and then averaged to calculate aggregate productivity. The establishments in the top and bottom one percent of the productivity distribution for each industry group are excluded from the calculation.

In addition to growth being uneven across time, dynamics vary across retail industries.

In Figure 2, we show labor productivity growth over time for the retail sector along with the five industries mentioned above, which highlight some of the factors that may have contributed to the patterns we observe in the retail sector as a whole, while highlighting divergences across industries during the period and relating our findings to industry narratives in the literature.

Figure 2. Average annual productivity growth in selected industries



Notes: Productivity growth is presented as the average compound annual growth rate over each five-year period. Years are presented in order from left to right for each industry. The growth rates are calculated using data on sales and employment from the Census of Retail Trade, in addition to BLS data on price deflators and industry group average hours worked. Productivity is calculated for each establishment and then averaged to calculate industry group productivity or retail sector productivity. The establishments in the top and bottom one percent of the productivity distribution for each industry group are excluded from the calculation. The Retail numbers are identical to those in Figure 1.

Productivity growth was modest for Electronics Retailers between 1987 and 1997, increased sharply in the late 1990s, and remained high (with only a slight dip during the Great Recession) through 2017. This coincides with the reorientation of sales toward computers and cell phones, which may indicate that these products require less labor for each dollar of sales.

On the other hand, Grocery Stores exhibited slow or negative productivity growth for most of the sample period. Grocery Stores have been at the forefront of introducing new technologies to the retail sector, such as universal product codes and self-checkout machines (Basker and Simcoe, 2021; Basker, Becker, Foster, White, and Zawacki, 2019).⁷ Apparently, either these innovations did not affect the average productivity of grocery store workers, or there were countervailing forces that offset these productivity increases. This finding suggests

⁷ Universal product codes began to be introduced before the start of our study, but there was still significant adoption after 1987 (Basker and Simcoe, 2021).

that these technologies may not have been the cause of productivity growth outside of Grocery Stores during this period.

The changes in input sourcing and the development of fast fashion among Clothing Stores have not consistently mapped to productivity growth, which has been uneven over time. Productivity in Clothing Stores peaked in the 1992-2007 period, then reverted to more moderate productivity growth after a spell of negative productivity growth.

The productivity growth of general merchandisers was generally moderate but varied considerably over the sample period. General merchandisers added grocery to their product mix (Arcidiacono, Bayer, Blevins, and Ellickson, 2016; Basker, Klimek, Van, 2012), which may be associated with lower productivity levels, and complicate the understanding of dynamics within this industry. Additionally, the growth of dollar stores in the later part of the sample may be important (Caoui, Hollenbeck, Osborne, 2023). For example, dollar stores may have higher labor productivity due to a reduced focus on the quality of the shopping experience.

The time path of productivity growth in the E-Commerce industry was different from that of the other industries, and it was somewhat surprising.⁸ Productivity growth started out at a relatively high level in the 1987-1992 period, increasing sharply and peaking in the 1992-1997 period, consistent with the idea that new processes were enhancing productivity in the age of the internet. However, productivity then declined steadily over the next 20 years to a level that was below the average productivity growth for the retail trade sector. Changes in the products being sold complicate this analysis. For example, e-commerce establishments initially focused on digital goods such as e-books and music, which may require almost no labor to sell. More recently, establishments in this industry sell essentially all products that may have lower associated labor productivity. Additionally, returns of physical goods may lower productivity.

⁸ E-commerce also includes mail-order houses which would be more important during the initial portion of the sample than towards the end.

Across these industries, we can draw out a few common patterns. Generally speaking, productivity growth rates were highest in the middle of our sample period. Most industries posted annualized growth rates of more than 4 percent in the middle decade, and even grocery stores experienced robust annualized growth of almost 3 percent in the 2002-2007 period. Similarly, during the Great Recession, all industries experienced low productivity growth. All industries except Non-Store Retailers recovered in the final five years of the study, although to levels that were below their peaks in the 2000s. A full account of structural transformation in the retail sector should consider both the common factors and heterogeneity across industries.

4.2 The Drivers of Productivity Growth Have Changed

To examine how the drivers of productivity growth have changed, we decompose industry labor productivity growth over multiple 10-year periods into three sources: continuing establishments, entering establishments, and exiting establishments. Following Foster et al. (2006), labor productivity growth can be decomposed as:

$$\begin{aligned} \Delta LP_{it} = & \sum_{e \in C} s_{e,t-10} \Delta LP_{e,t} + \sum_{e \in C} (LP_{e,t-10} - LP_{i,t-10}) \Delta s_{e,t} + \sum_{e \in C} \Delta LP_{e,t} \Delta s_{e,t} \\ & + \sum_{e \in N} s_{e,t} (LP_{e,t} - LP_{i,t-10}) - \sum_{e \in X} s_{e,t-10} (LP_{e,t-10} - LP_{i,t-10}) \end{aligned} \quad (2)$$

where LP_{it} is log labor productivity in industry i at period t , $s_{e,t}$ is establishment e 's share of output in period t , C is the set of establishments that exist in both periods (continuers), N is the set of establishments that entered between year $t-10$ and t , and X is the set of establishments that exited between year $t-10$ and t .

The first three terms measure the contribution of continuers.⁹ The fourth term is the

⁹ The first term measures the effect of within-establishment productivity growth of continuers, weighted by the establishment's share of 4-digit NAICS industry output in the first period ($s_{e,t-10}$). The second term measures the effect of changing output shares among continuing establishments (reallocation), weighted by the establishment's labor productivity relative to the industry average labor productivity in the first period. The third captures the contribution of the interaction between the within- and between-establishment terms and shows whether labor productivity growth and changes in output shares move together or in opposite directions.


contribution to LP growth due to entering establishments. It is equal to the establishment share of output in year t , multiplied by the labor productivity of that establishment relative to the industry average. The final term is exiting establishments' contribution to LP growth. It is equal to the establishment's share of industry output in year $t-10$ times the establishment's labor productivity relative to the industry average.

We can further decompose the entry and exit components into contributions of continuing *firms* (expansions) and *firm* birth/death. This is an important distinction since the expansion of existing firms through new locations is a distinct phenomenon from *new* establishments of new firms. Likewise, consolidation of existing firms can be distinguished from shuttering of firms altogether.

Figure 3 summarizes the results of this decomposition for the three 10-year intervals in our sample, where the components in the decomposition were regrouped into three components for simplicity. The sum of the within, between-, and cross-components of continuing establishments is labeled as 'Continuers,' the sum of firm entry and entry by establishments of continuing firms is labeled as 'Entrants,' and the sum of firm exit and exit by establishments of continuing firms is labeled as 'Exiters.' Each bar shows the annualized LP growth and the contribution of these three components. Table 1 shows the details of the decomposition, including the firm entry/exit distinction along with within-, between- and cross terms, but we begin with establishment dynamics before returning to the role of firm dynamics.

The first bar in Figure 3 indicates that LP grew by approximately 1.8% yearly between 1987 and 1997. Average annual growth was 3% in the second decade, while it was 1.2% in the third decade. The decomposition for the 1987-1997 period shows the same patterns as Foster et al. (2006), although our results differ from theirs due to differences between the SIC and NAICS definitions of the retail sector.

Figure 3. Average annual productivity growth by establishment status



Notes: Productivity growth is presented as the average compound annual growth rate over each ten-year period. The growth rates are calculated using data on sales and employment from the Census of Retail Trade, in addition to BLS data on price deflators and industry group average hours worked. Productivity is calculated for each establishment and then averaged to calculate retail sector productivity. Following Foster et al. (2006), continuers includes the within, between, and cross components of the productivity growth decomposition calculation. Entrants and exiters are the portions of the productivity growth decomposition due to entry of new establishments and exit of existing establishments, respectively.

The most important determinant of the hump shape in aggregate productivity dynamics is the contribution of continuing establishments. These establishments exhibit a total of about 0.4% annual growth between 1987 and 1997, then 1.25% in the second decade—due primarily to the more than a threefold increase of the within-component (from 0.5 percent to about 1.7 percent; see Table 1 below)—and 0.5% in the last decade. These numbers imply that continuers’ contribution amounted to 20%, 42%, and 43% of overall LP growth in the three decades, respectively. Entry was the primary source of LP growth in the first decade, contributing about one percentage point or 53% of overall LP growth. Its contribution remained stable in the second decade (1.3%), but its share in overall LP growth is only 44% because of the within-establishment dynamics mentioned above. The contribution by this component shrank in the last decade in both absolute (0.2% annually) and relative terms (18% of overall LP growth) because its growth rate decreased below values seen in the first decade, while within-establishment growth was comparable to that in the first decade (see Table 1). The role of exit

in absolute terms is similar across decades (0.5%, 0.4%, 0.5% annually); the relative importance is driven by the dynamics of the other two components (27%, 14%, 39% of total LP growth).

Table 1: Productivity Growth Decomposition, 10-year growth rates

	1987-1997	1997-2007	2007-2017
Within	0.052	0.174	0.065
Between	0.037	0.034	0.044
Cross	-0.052	-0.074	-0.057
Entry	0.099	0.140	0.021
New Firm	0.019	0.014	-0.009
Continuing Firm	0.080	0.126	0.030
Exit	0.050	0.044	0.047
Exiting Firm	0.045	0.036	0.044
Continuing Firm	0.005	0.008	0.003

Notes: Productivity growth is presented as the growth rate over each ten-year period. The growth rates are calculated using data on sales and employment from the Census of Retail Trade, in addition to BLS data on price deflators and industry group average hours worked. Productivity is calculated for each establishment and then averaged to calculate retail sector productivity. Following Foster et al. (2006), continuers includes the within, between, and cross components of the productivity growth decomposition calculation. Entrants and exiters are the portions of the productivity growth decomposition due to entry of new establishments and exit of existing establishments, respectively.

Distinguishing between firm and establishment dynamics illustrates some of the structural changes in the retail sector. In the first period of Figure 3, the largest source of LP growth was establishment entry, but Table 1 shows that this was primarily driven by the expansion of larger firms through the entry of new establishments. Expansion was still important between 1997 and 2007, the period of consolidation. In the third decade, entrants were the least important of the three contributors, reflecting declines in contributions of expanding firms and new firms.¹⁰ Our results suggest that the impact of exiting firms on aggregate productivity growth is about the same in all three periods, while establishment exit of continuing firms has a negligible impact. A possible interpretation of the increase in the contribution of continuing establishments between the 1987-1997 and 1997-2007 periods is that existing firms were able to consolidate their knowledge and experience, which resulted in higher productivity growth at their continuing establishments. The structural changes in the

¹⁰ Somewhat surprisingly, the contribution of new firms turns negative, indicating that new firms are less productive than the typical firm.

2007-2017 period reflect the impact of the Great Recession. The subsequent slow economic growth was due to a sharp decline in the contributions of entering establishments, both from entering firms and expanding firms, and continuing establishments. The largest source of labor productivity growth in this decade was the increase in productivity among continuing establishments, while the more efficient ones increased their share of output as well.¹¹

Figure 3 also illustrates the impact of the transformation of retail as firms expanded with new establishments. Entry of new establishments, particularly from existing firms within the industry, helped drive the higher productivity growth rate. Once these establishments, and their parent firms, became established in the industry, the locus of productivity growth shifted to existing establishments. The Great Recession most likely drove the decline in the rate of retail productivity growth overall (to only a 0.6% compound annual growth rate from 2007 to 2012 as shown in Figure 1) and in each of our five industries. Employment shares of four of the five industries increased from 2007 to 2012 (only Electronics experienced decreasing employment share, which may, in part, explain why its productivity growth is by far the highest of the five). Overall, the five industries in Figure 2 accounted for approximately 45 percent of retail employment in the 2010s, an increase from about 38 percent in 1992, with General Merchandise accounting for the entire increase.

4.3 Multi-Unit Firms are More Productive than Single-Unit Firms on Average

To shed light on some of the sources of productivity differences, we explore the relationship between establishment productivity and indicators of establishment status (entrant, exiter, or continuer) and the structure of each establishment's firm. We classify firms into four categories based on whether they are multi-unit and how many states they operate in. Single-

¹¹ An area for future research is the extent to which the fall in the contribution of entry and exit is due to less entry and exit or whether it is due to entrants and exiters looking more like continuing establishments.

unit firms operate one establishment. Multi-unit or chain firms can be one of three types depending on how many states they operate in. A local firm operates multiple establishments in only one state. A regional firm operates establishments in two to five states. Finally, a national firm operates in six or more states. Although we do not use additional information about location or distance, these variables do capture some of the spatial variation that may be relevant in retail trade.

We study these relationships in two ways. First, we regress log productivity on interactions of firm structure and establishment status (with controls for industry). Second, we examine where the establishments of each firm structure category lie in the overall labor productivity distribution. Our results suggest that establishments of chains are more productive than those of single-unit firms on average, but we still find that 36 percent of single-unit firms are in the top 40 percent of establishments in terms of productivity.


First, we present the results of regressing log labor productivity of establishment i in industry s in year t – industry effects removed (captured by λ_s in the regression) – on a set of dummy variables.

$$\begin{aligned} \ln(\text{prod}_{i,s,t}) = & \alpha + \beta_1 \text{cont}_i + \beta_2(\text{cont}_i \times \text{loc}_i) + \beta_3(\text{cont}_i \times \text{reg}_i) + \beta_4(\text{cont}_i \times \text{nat}_i) \\ & + \gamma_1(\text{cont}_i \times \text{EY}_t) + \gamma_2(\text{cont}_i \times \text{loc}_i \times \text{EY}_t) + \gamma_3(\text{cont}_i \times \text{reg}_i \times \text{EY}_t) \\ & + \gamma_4(\text{cont}_i \times \text{nat}_i \times \text{EY}_t) + \delta_1 \text{ent}_i + \delta_2(\text{ent}_i \times \text{loc}_i) + \delta_3(\text{ent}_i \times \text{reg}_i) \\ & + \delta_4(\text{ent}_i \times \text{nat}_i) + \zeta_1 \text{exit}_i + \zeta_2(\text{exit}_i \times \text{loc}_i) + \zeta_3(\text{exit}_i \times \text{reg}_i) \\ & + \zeta_4(\text{exit}_i \times \text{nat}_i) + \lambda_s + \epsilon_{i,t} \end{aligned}$$


We focus on 10-year periods, as we did in the analyses in the previous section. The first group of dummies defines establishment status (entrants, exiters, and continuers). Entrants (ent_i) appear only in the later CRT. Exiters (exit_i) appear only in the first CRT. Continuers (cont_i) are establishments that appear in both CRTs and, therefore, are represented by two observations in each regression. We capture changes in productivity from the initial year to the end year by including the interaction between continuer status and a dummy variable that indicates whether

the observation was in the end-year of a period (EY_t). The second group of dummy variables defines firm structure: single-unit establishments, establishments that belonged to a local (loc_i), regional (reg_i), or national (nat_i) chain. Figure 4 shows the estimated coefficients from these period-specific regressions.¹² In each specification, the reference group is continuing single-unit establishments in the initial year. The four panels of Figure 4 show the coefficient estimates along with 95-percent confidence intervals (full regression results are in the Appendix Table A1). The coefficients represent partial correlations between establishment status and firm type.


Figure 4. Coefficients of regression of log labor productivity on the interactions between establishment entry/exit and firm structure, controlling for industry and year effects.



¹² The full results from this regression are in Table A1 in the appendix.



Panel d. Entrants



Notes: Establishment log labor productivity is regressed on indicators of whether the establishment is a single unit or part of a local, regional, or national firm. Establishments are also identified as continuers if they are in both the first and last year of the period, entrants if they appear only in the last year of the period, and exiters if they appear only in the first year of the period. Continuing single-unit establishments are the excluded category. Coefficient confidence intervals are shown using the red bands on each figure. Robust standard errors are clustered by NAICS-4. Observation counts were rounded according to Census disclosure rules.

The coefficients shown in Panel a of Figure 4 indicate that continuing retailers that belonged to a chain were significantly more productive than single-unit continuers (the reference group) in all three periods. On average, chain-owned continuers' productivity advantage relative to single-unit continuers was 9-20 percent, depending on the period and the type of chain (local/regional/national). The differences across chain types or periods are not statistically significant.

Panel b captures the productivity growth of continuers in each time frame. A positive

coefficient indicates that productivity generally increased among establishments in the group over the decade. The results indicate that among continuers, only single units in the first decade were mildly more productive in the end year than in the first year. However, in the second decade, continuing establishments at national chains became significantly more productive, with productivity growing by an average of 26 percent. Crucially, this was when productivity growth in retail was high, primarily driven by continuing establishments (see Finding 2). Thus, national chains were a significant driver of productivity growth during the consolidation phase.

The coefficients on exiters, shown in Panel c, indicate that exiting single-unit establishments were 29-36 percent less productive than continuing single-units, a difference that is significant both statistically and economically. Exiting stores of local chains were also less productive than the reference group, but to a lesser extent: between 10-15 percent on average. Exiting establishments of regional and national firms were not statistically significantly less productive than continuing single-unit retailers.¹³

Panel d summarizes differences across the groups of entrants. Single-unit entrants' productivity levels do not differ statistically significantly from single-unit continuers. On the other hand, entrants of chains exhibited significantly higher productivity levels except for local chains between 2007-2017. Relating these patterns to the decomposition discussed in the previous section, we can say that the lower contribution of entry in the third decade may be due to either fewer entrants with similar productivity levels or comparably numerous but less productive entrants.¹⁴

But the coefficients do not tell the whole story. The R-squared values for these

¹³ As discussed in section 4.2, exiting establishments among continuing firms contribute little to productivity, which is consistent with the relatively small and statistically insignificant coefficients on exiting establishments of larger chains. An interesting question is how much of the productivity difference of exiting local stores is due to firm exit and exit from continuing firms. We leave this for future research.

¹⁴ These results do not allow us to differentiate between the two explanations because the average productivity level of single-unit continuers changed across decades. For example, the productivity levels of continuing single-unit establishments could have risen.

regressions (see Appendix Table A1), indicate that these characteristics explain a large fraction of the variation in establishment-level productivity. However, the explained variation declined steadily over our sample period, from 47 to 38 to 28 percent. The reason for this decline is not clear, but it warrants further investigation.

Another way to examine the role of firm structure is to look at where the different types of establishments fall in the productivity distribution. Table 2 shows productivity quintiles by firm structure. We construct the table by first expressing each establishment's productivity relative to its industry mean. We then combine the data from these industries and calculate quintiles. In 1987, just over 64 percent of retail establishments were single-unit firms. Approximately 12 percent of retail establishments belong to a local chain, almost 6 percent of establishments belong to a regional chain, and about 18 percent belong to a national chain.

Table 2: Productivity Quintiles by Establishment Characteristic

1987		Chain Type		
Quintile	Single-Unit	Local	Regional	National
1	15	2	1	2
2	14	2	1	3
3	12	3	1	4
4	11	3	1	5
5	12	3	1	4
Total	64	12	6	18

2017		Chain Type		
Quintile	Single-Unit	Local	Regional	National
1	14	1	1	4
2	11	2	1	6
3	10	2	1	7
4	9	2	1	8
5	11	2	1	6
Total	56	8	5	31

Notes: Quintile rows sum to 20; total rows sum to 100. Establishment productivity is calculated using data on sales and employment from the Census of Retail Trade, in addition to BLS data on price deflators and industry group average hours worked. The establishments in the top and bottom one percent of the productivity distribution for each industry group are excluded from the calculation.

In 1987, 24 percent of single-unit establishments were in the lowest productivity quintile, and 19 percent belonged to the top quintile. The remaining 57 percent of single-unit stores are in the middle of the productivity distribution. In contrast, establishments that are part of national firms are less likely to be observed in the lowest productivity quintile and more likely to be in the top quintiles. Local and regional chain establishments are similarly more likely to be observed in the top two quintiles. The conclusion from these results is that even though single-unit establishments tend to be less productive than establishments that are part of multi-unit firms, a large fraction of these establishments are highly productive.

Turning to 2017, we see several shifts. The distribution of establishments over firm structure is still dominated by single-unit firms, but national chains have become more important. Specifically, the overall share of single-unit firms fell to 56 percent, while the share of establishments in national chains nearly doubled to 31 percent, with the increase in the share

of national-chain establishments coming disproportionately from local and regional chains. Over this period, the total number of retail establishments increased, but the total number of firms decreased (Smith and Ocampo, Forthcoming). Due to this shift, establishments that are single-unit firms accounted for a smaller fraction of each quintile in 2017 compared with 1987. The shares of establishments in local and regional chains in each quintile declined but by smaller amounts.

Looking at the distribution of establishments across productivity quintiles by type, we find that, except for establishments in national firms, there was remarkably little change between 1987 and 2017. Approximately 45 percent of single-unit establishments were in the first two quintiles in both years, while 36 and 37 percent were in the top two quintiles in 2017. The distributions were similarly stable for establishments in local and regional firms. Somewhat surprisingly, the fraction of national-chain establishments in the bottom two quintiles increased from 27 to 32 percent, while the fraction in the top two quintiles fell from 48 percent to 44 percent, with all the decline coming from the top quintile.

Although single-unit firms tend to be less productive, a large fraction of these establishments are in the top two quintiles. Despite the decline in overall share between 1987 and 2017, single-unit establishments still make up over half of the top productivity quintile. It is possible these single-unit establishments have established successful niche markets in their local economies where margins may be higher. An alternative explanation is that these single-unit firms will grow and become national firms, as was documented by Foster et al. (2016). Expanding the analysis to control for differences across industries seems warranted and could provide insight into these patterns.

4.4 Productivity Dispersion is Significant and Varies Across Industries and Over Time

The results in the previous section shed light on the nature of the productivity

distribution across establishments in the retail sector and how establishment characteristics map onto the productivity distribution. However, they tell us nothing about the size of the differences in productivity across establishments or the amount of productivity variation that exists within narrowly defined industries. Such variation in productivity could indicate that there are potential gains from reallocation (Restuccia and Rogerson, 2008; Hsieh and Klenow, 2009). For example, in the manufacturing sector, the substantial level and secular increase in productivity dispersion since the 1980s is well documented and indicates large potential gains from reallocation (Syverson 2004, Decker et al. 2020, Cunningham et al. 2023). Nonetheless, the market conditions facing retailers were likely far different from those faced by manufacturing plants during this period, and the characteristics of the productivity distribution in the retail sector are not well-documented. We break new ground in this paper by quantifying the within-industry dispersion of labor productivity in the retail sector.

Our primary measure of dispersion is the interquartile range (IQR) of establishment-level log productivity (net of industry and year effects), which is calculated for the 27 4-digit NAICS industries each year. The IQR is defined as the difference in productivity between the 75th and 25th percentiles of the log-labor-productivity distribution in an industry-year cell. We express this difference as a productivity multiple. For example, an IQR value of 100 indicates that an establishment at the 75th percentile is 100 percent more productive than one at the 25th percentile. We also calculate the 99-90 and 10-1 ranges, which are defined analogously. The IQR tells us about dispersion at the center of the distribution, while the 10-1 and 99-90 ranges tell us about dispersion among the least and most productive establishments.

We start by returning to the five industries we highlighted in Section 4.1. Figures 5a-5e show unweighted (each establishment weighted equally) dispersion statistics (IQR, 99-90, and 10-1) for the five industries. We see large productivity dispersion in all five industries but with differences in both levels and trends. Specifically, the IQRs for Grocery Stores, General

Merchandise Stores, and Clothing Stores have similar levels and are flat or increase slightly. For example, the IQR multiple of Grocery Stores increased slightly from 108 in 1987 to 120 in 2017 (Figure 5a). In other words, the establishment at the 75th percentile of the productivity distribution in Grocery Stores was 2.08 times as productive as the establishment at the 25th percentile of the productivity distribution in 1987 and 2.02 times as productive in 2017. The IQRs are somewhat lower for General Merchandise Stores (Figure 5b) and somewhat higher for Clothing Stores (Figure 5c). The remaining two industries—Electronics and Appliance Stores (Figure 5d) and E-commerce (Figure 5e)—are of interest because they recorded the highest and second-highest productivity growth over our sample period. As we can see in the figures, the IQRs for these industries follow different patterns over time. The IQR for Electronics and Appliance Stores started out very high at just over 420 (the establishment at the 75th percentile of the distribution was 5.2 times as productive as the establishment at the 25th percentile of the distribution), decreased steadily until 2002, then remained flat through 2017. Dispersion in the E-commerce industry increased substantially over the 30-year period such that by 2017, an establishment at the 75th percentile was over 4.4 times as productive as an establishment at the 25th percentile.

Figure 5a: Industry Productivity Dispersion, Grocery (4451)




Figure 5b: Industry Productivity Dispersion, General Merchandise (4523)

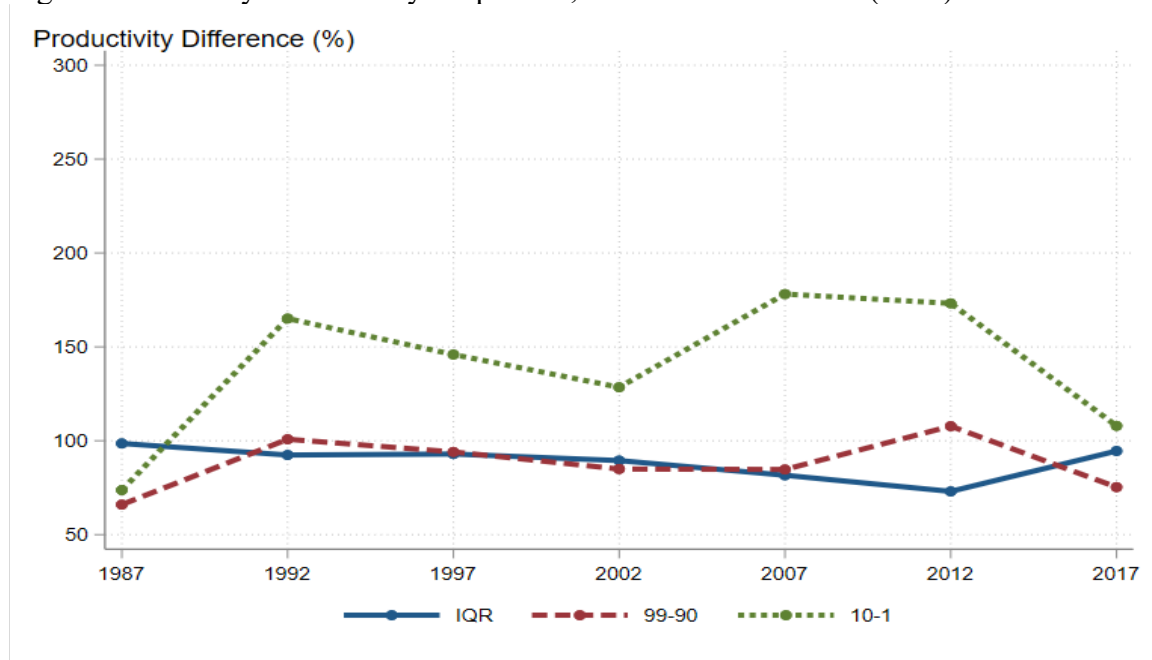


Figure 5c: Industry Productivity Dispersion, Clothing (4481)




Figure 5d: Industry Productivity Dispersion, Electronics and Appliances (4431)





Figure 5e: Industry Productivity Dispersion, Electronic Shopping (4541)



Notes: Establishment productivity is calculated using data on sales and employment from the Census of Retail Trade, in addition to BLS data on price deflators and industry group average hours worked. Productivity is calculated for each establishment and the establishments in the top and bottom one percent of the productivity distribution for each industry group are excluded from the calculation. The y-axis plots the productivity difference of log productivity. The interquartile range is the difference between the 75th percentile establishment and the 25th percentile establishment. Similarly, the 99-90 range is the difference between the 99th and 90th percentiles and the 10-1 range is the difference between the 10th and 1st percentiles.


Dispersion in the tails of the distribution is also large and varies across industries and over time. In some industries, productivity dispersion in the tails is as large or larger than it is in the center of the within-industry distribution. In Grocery Stores in 2017, for example, establishments at the 99th percentile were nearly twice as productive as those in the 90th percentile. At the other end of the distribution, the establishment at the 10th percentile of the productivity distribution is 2.7 times as productive as the establishment at the 1st percentile of the distribution. In Grocery Stores, there is more dispersion in productivity among the lowest productivity establishments than in the middle of the distribution. This is also the case for Clothing Stores and General Merchandise Stores in most years.

To complete the picture of the Retail Sector, Figures 6 through 9 summarize the industry-level dispersion statistics. Figures 6 and 7 show the establishment-based means

(unweighted and industry weighted) of the three dispersion statistics, while Figures 8 and 9 show analogous mean values for firm-based statistics.


In Figure 6, the mean IQR multiple is relatively flat at around 150 over our sample period. This means that, in industries at the mean IQR, the establishment at the 75th percentile is about 2.5 times as productive as the establishment at the 25th percentile. The mean 10-1 range ratcheted up between 1987 and 1992 and between 2002 and 2007 but was relatively flat otherwise. The 99-90 range increased between 1987 and 1992 and was relatively flat for the rest of the sample period. Notably, the mean 10-1 range was about the same as the mean IQR. This is remarkable, given that there are five times as many establishments between the 75th and 25th percentiles as there are between the 10th and 1st percentiles. The mean 99-90 range was lower than the other two, indicating that productivity dispersion among the most productive establishments is smaller than among those in the middle half and bottom decile of the within-industry distribution on average. This is in contrast to what was found in the manufacturing sector, where dispersion in the right tail is comparable to the IQR (Cunningham et al.,2023).

Figure 6: Average Dispersion for Establishments (Unweighted)




Notes: Lines represent the mean across industries in the within-industry statistics as indicated by the legend. NAICS-4 industries are given equal weight. Establishments within each industry are given equal weight. Calculations use data from the CRT, LBD, and BLS industry-productivity statistics.

Figure 7: Average Dispersion for Establishments (Industry Weighted)




Notes: Lines represent the mean across industries in the within-industry statistics as indicated by the legend. NAICS-4 industries are weighted according to their share of total hours within a year. Establishments within each industry are given equal weight. Calculations use data from the CRT, LBD, and BLS industry-productivity statistics.

Figure 8: Average Dispersion for Firms (Unweighted)



Notes: Lines represent the mean across industries in the within-industry statistics as indicated by the legend. NAICS-4 industries are given equal weight. Firms within each industry are given equal weight. Calculations use data from the CRT, LBD, and BLS industry-productivity statistics.

Figure 9: Average Dispersion for Firms (Industry Weighted)



Notes: Lines represent the mean across industries in the within-industry statistics as indicated by the legend. NAICS-4 industries are weighted according to their share of total hours within a year. Firms within each industry are given equal weight. Calculations use data from the CRT, LBD, and BLS industry-productivity statistics.

A comparison of Figures 7, 8, and 9 to Figure 6 helps illustrate two features of the retail trade sector that deserve special attention. The first is that activity is heavily concentrated in a few industries (Smith and Ocampo, Forthcoming). The two largest 4-digit industries in retail are Grocery Stores and General Merchandise Stores. These two industries accounted for approximately one-fifth of retail sales in 2017. This suggests that industry-weighted dispersion statistics could be different if large industries' dispersion patterns are different. We can see this by comparing industry-weighted dispersion statistics in Figure 7 to the unweighted statistics in Figure 6.¹⁵ The average IQR is lower and exhibits a slight upward trend when we weight by

¹⁵ We weight industries by their share of total annual hours in the retail sector when calculating these cross-industry

industry, indicating that there is less dispersion in larger industries but that high-dispersion industries have become relatively larger. The mean 99-90 and 10-1 ranges look remarkably similar. They are at about the same level in both the unweighted and industry-weighted graphs, although there is slightly more year-to-year variation in the 99-90 range when weighted by industry. Thus, industry size matters more in the middle of the distribution than in the tails.

The second feature is the period of firm expansion (1987 – 1997), which was followed by the consolidation period (1997 – 2007) and saw large national and regional firms opening new establishments. These large firms often have decades of experience in opening and operating stores that look very similar across the regions where they operate. The effects of consolidation and the rise of national chains can be seen in the dispersion statistics by comparing unweighted firm-level (Figure 8) to the unweighted establishment-level statistics (Figure 6). The graphs for the 99-90 and 10-1 ranges again look remarkably similar, with the main difference showing up in the middle of the distribution. The graphs for the tails are at about the same level and have very similar trends. The firm-based IQR is higher, which we would expect given that establishments in the same firm likely have very similar levels of productivity. Again, the main differences between firm-based and establishment-based measures are in the middle of the distribution. The story is almost exactly the same when we weight by industry. That is, industry weighting does not alter our comparison of establishment-based and firm-based dispersion statistics.

Comparing unweighted and industry-weighted firm-based dispersion statistics (Figures 8 and 9) sheds additional light on the second feature. Like our other comparisons, dispersion in the tails of the two distributions looks very similar in both level and trend. Again, the differences are in the middle of the distribution. The industry-weighted and unweighted graphs

statistics, but we do not weight establishments when constructing industry IQRs. In the next draft, we will consider weighting by revenue shares.

end up at about the same level, but they have sharply different trends. The industry-weighted graph starts at a much lower level compared with the unweighted graph (136 vs. 161) and increases faster, with both ending up around 175. Thus, dispersion in the middle increased faster among firms in relatively large industries, and high-dispersion (firm-based) industries became relatively larger.

Our comparisons of firm-based and establishment-based dispersion measures are consistent with our finding that at least 90 percent of the total variation across establishments is due to differences across firms and with the expansion of national chains, which operate many nearly identical establishments. And in all of these comparisons—establishment vs. firm and unweighted vs. industry-weighted—the differences showed up in the middle of the distribution (the IQRs).

4.5 Firms with higher productivity pay higher wages and grow more quickly

Productivity is a primary indicator of firm and establishment performance, and we would expect it to be correlated with employment growth and earnings. We examine the relationships between productivity and these firm-level outcomes in Table 3.¹⁶

Column 1 of Table 3 shows results from regressing the natural log of average earnings (calculated as the ratio of payroll to employment) on the natural log of sales per hour, interacted with the firm structure dummies described in Section 4.3. The first coefficient in column 1 suggests that a 1 percent increase in single-unit firm productivity is associated with higher

¹⁶ Ideally, we would use total factor productivity for this purpose. But as noted in footnote 4, the CRT only has information on sales, employment, and payroll. Therefore, we use labor productivity as defined above as a measure of performance. It is worth noting, by using labor productivity as our performance measure in regressions where the dependent variables are labor related, we cannot make causal inferences. For example, in the wage regressions there may be division bias in the coefficient of LP. In the growth regressions, coefficients may be positively biased (LP higher because denominator is lower, which also implies growth rate is higher).

average earnings of approximately 0.5 percent. This association is even stronger for regional and national firms—an additional 0.08 and 0.07, although the difference between regional and national chains is not statistically significant.

Table 3: The relationship between productivity and other outcomes interacted with firm structure indicators

	Log-earnings- per-worker	Earnings-per- worker, (5-year change in %)	Employment (5- year change in %)
ln(output-per- hour) [†]	0.504*** (0.0245)	0.520*** (0.0192)	0.222*** (0.0098)
Firm structure*Productivity			
Local	0.0307 (0.0192)	0.0207* (0.0102)	-0.0197 (0.0120)
Regional	0.0844*** (0.0205)	0.00802 (0.0198)	0.0105 (0.0134)
National	0.0673* (0.038)	0.0106 (0.0525)	0.177*** (0.0487)
Constant	2.616*** (0.0002)	0.131*** (0.00172)	-0.0437*** (0.0012)
Observations	4454000	2033000	2033000
R-squared	0.514	0.363	0.078

Notes: The regressions in columns 2 and 3 are based on using the first available lag of log-labor-productivity. The reference group is single-unit firms. Robust standard errors clustered by NAICS-4 and are in parentheses. Values were rounded according to Census Bureau disclosure rules. ***, **, * denote statistical significance at 1%, 5%, 10% levels, respectively.

Column 2 shows results from regressing the percentage change in average earnings on productivity, again interacted with firm type. We find that a one percentage point increase in the natural log of real sales per hour of single-establishment firms leads to a 0.52 percentage point increase in average earnings within the firm. The association is slightly larger for local firms with an estimated coefficient of 0.02 but is significant only at the 10 percent level. The association between productivity and earnings growth for regional and national firms is not statistically different from single-establishment firms. Thus, although wages are higher in more productive regional and national firms compared with single-establishment firms, wages do not

grow faster in these firms. The last column of Table 3 shows the effect of productivity on employment growth: a 1 percent more productive firm grows about 0.22 percent faster, with 0.19 percent additional growth at national chains.

The findings in this section are generally in line with results from the firm dynamics literature, which has focused on the manufacturing sector (due to data availability). We find the expected (positive) relationships between earnings and productivity and between earnings and employment growth and productivity. We also find firm structure is important in these relationships.

5. Conclusion

The retail sector has transformed from a sector dominated by single-unit firms to one filled with multi-unit firms, many of which have a large geographic footprint. The early part of this transformation was accompanied by moderate productivity growth, driven by the expansion of new establishments of existing firms. The middle part, when productivity was the highest, was driven by existing establishments consolidating their position and becoming more productive. More recently, growth has stagnated for reasons that are not completely clear.

Throughout the period we find establishments of large firms are more productive than those of single-unit firms. Even the exiting establishments of large firms are as productive as continuing single-unit firms on average, although there is significant heterogeneity in terms of the productivity levels of individual establishments of each firm type.

A major contribution of this paper is the analysis of productivity dispersion in retail trade. We find that there is significant dispersion within detailed industries. Much of what we see in the dispersion statistics accords well with the narrative that we typically hear about retail trade.

When we compare industry-weighted statistics with unweighted statistics, we see lower dispersion in the industry-weighted IQR indicating that there is less dispersion in larger industries, such as Grocery Stores. Comparing establishment-based and firm-based dispersion statistics we see greater dispersion in the firm-based IQRs, which is consistent with large chains expanding via the opening of nearly identical establishments. Interestingly, these differences show up only in the middle of the distribution—dispersion in the tails is invariant to industry weighting and unit of observation. And in contrast to earlier findings for manufacturing, dispersion in the upper tail of the distribution is lower than in the middle, which is consistent with the higher productivity of establishment in chains. In the future, we plan to release these dispersion statistics under a joint agreement between the BLS and Census.

References

- Arcidiacono, Peter, Patrick Bayer, Jason R. Blevins, and Paul B. Ellickson. 2016. "Estimation of Dynamic Discrete Choice Models in Continuous Time with an Application to Retail Competition." *Review of Economic Studies*, 83(3): 889-931.
<https://doi.org/10.1093/restud/rdw012>
- Basker, Emek. 2006. "Job Creation or Destruction? Labor-Market Effects of Wal-Mart Expansion," *Review of Economics and Statistics* 87(1): 174-183.
- Basker, Emek, Shawn Klimek, and Pham Hoang Van. 2012. "Supersize It: The Growth of Retail Chains and the Rise of the "Big Box" Store," *Journal of Economics and Management Strategy*, 21(3): 541-582.
- Basker, Emek and Pham Hoang Van. 2010. "Imports "R" Us: Retail Chains as Platforms for Developing-Country Imports," *American Economic Review: Papers & Proceedings*, 100: 414-418.
- Bureau of Labor Statistics, U.S. Department of Labor, and the Census Bureau, U.S. Department of Commerce. 2023. "Dispersion Statistics on Productivity," 1987–2020.
<https://www.bls.gov/productivity/articles-and-research/dispersion-statistics-on-productivity/home.htm> and <https://www.census.gov/disp>.
- Caoui, El Hadi, Brett Hollenbeck, and Matthew Osborne. 2023. "The Impact of Dollar Store Expansion on Local Market Structure and Food Access."
https://bretthollenbeck.com/wp-content/uploads/2023/07/caoui_hollenbeck_osborne_dollarstores_draft_v5.pdf
- Chow, Melissa C., Teresa C. Fort, Christopher Goetz, Nathan Goldschlag, James Lawrence, Elisabeth Ruth Perlman, Martha Stinson, and T. Kirk White. 2021. "Redesigning the Longitudinal Business Database." NBER Working Paper 28839.
<https://doi.org/10.3386/w28839>.
- Cunningham, Cindy, Lucia Foster, Cheryl Grim, John Haltiwanger, Sabrina Wulff Pabilonia, Jay Stewart, and Zoltan Wolf. 2023. "Dispersion in Dispersion: Measuring Establishment Level Differences in Productivity." *Review of Income and Wealth*, 69(4): 999-1032. DOI: [10.1111/roiw.12616](https://doi.org/10.1111/roiw.12616).
- Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2020. "Changing Business Dynamism and Productivity: Shocks versus Responsiveness." *American Economic Review*, 110(12): 3952–90. DOI: [10.1257/aer.20190680](https://doi.org/10.1257/aer.20190680)
- Elvery, Joel, Lucia Foster, C.J. Krizan, and David Talan. 2006. "Preliminary Micro Data Results from the Business List Comparison Project." In 2006 Proceedings of the American Statistical Associations, Business and Economics Statistics Section. Alexandria, VA: American Statistical Association.
- Fort, Teresa and Klimek, Shawn. 2018. The Effects of Industry Classification Changes on US Employment Composition, Working Papers, U.S. Census Bureau, Center for Economic Studies, <https://EconPapers.repec.org/RePEc:cen:wpaper:18-28>.

- Foster, Lucia, John Haltiwanger, Shawn Klimek, Cornell J. Krizan, and Scott Ohlmacher. 2016. "The Evolution of National Retail Chains: How We Got Here." *Handbook on the Economics of Retailing and Distribution*, edited by Emek Basker pp. 7-37, ch 1, <https://dx.doi.org/10.4337/9781783477388>.
- Foster, Lucia, John Haltiwanger, and Cornell J. Krizan. 2006. "Market Selection, Reallocation, and Restructuring in the U.S. Retail Trade Sector in the 1990s." *The Review of Economics and Statistics*, 88(4): 748-58. <https://doi.org/10.1162/rest.88.4.748>
- Holmes, Thomas J., 2011. "The Diffusion of Wal-Mart and Economies of Density," *Econometrica*, 79(1): 253-302.
- Hortacsu, Ali and Chad Syverson. 2015. "The Ongoing Evolution of US Retail: A Format Tug-of-War." *Journal of Economic Perspectives*, 29(4): 89-112. <https://dx.doi.org/10.1257/jep.29.4.89>
- Hsieh, Chang-Tai and Peter J. Klenow. 2009. "Misallocation and Manufacturing TFP in China and India." *Quarterly Journal of Economics*, 124(4): 1403–1448. <https://doi.org/10.1162/qjec.2009.124.4.1403>
- Restuccia, Diego and Richard Rogerson. 2008. "Policy Distortions and Aggregate Productivity with Heterogeneous Establishments." *Review of Economic Dynamics*, 11(4): 707-720. <https://doi.org/10.1016/j.red.2008.05.002>
- Smith, Dominic A. 2018. "Concentration and Foreign Sourcing in the U.S. Retail Sector." https://dominic-smith.com/smith_jmp.pdf
- Smith, Dominic A. and Sergio Ocampo. Forthcoming. "The Evolution of US Retail Concentration." *American Economic Journal: Macroeconomics*. <https://www.aeaweb.org/articles?id=10.1257/mac.20220249&from=f>
- Syverson, Chad. 2004. "Product Substitutability and Productivity Dispersion." *Review of Economics and Statistics*. 86(2): 534–550. doi: <https://doi.org/10.1162/003465304323031094Appendix>
- U.S. Census Bureau. 2017. "Selected Sectors: Single Unit and Multiunit Firms for the U.S.: 2017." *Economic Census, ECN Core Statistics Economic Census: Establishment and Firm Size Statistics for the U.S., Table EC1700SIZESUMU*, 2017, [https://data.census.gov/table/ECNSIZE2017.EC1700SIZESUMU?q=retail trade&y=2017&n=448](https://data.census.gov/table/ECNSIZE2017.EC1700SIZESUMU?q=retail%20trade&y=2017&n=448). Accessed on June 16, 2024.

Appendix

A. BLS Industry-level Productivity

BLS publishes annual measures of real sectoral output, employment levels, hours worked, and labor productivity growth for 27 4-digit North American Industry Classification System (NAICS) retail industries.¹⁷ BLS also makes available a dataset that includes values for selected 5-digit and 6-digit NAICS Retail industries for which the underlying data are of high quality. For each industry, labor productivity (LP) growth is measured as the change in the ratio of indexes of sectoral output and labor. These LP growth rates are chained to construct productivity indexes.

BLS estimates real sectoral *output* using several data sources. Three surveys from the Census Bureau provide nominal output estimates: the Economic Census (conducted every five years), the Annual Retail Trade Survey (ARTS), and the Monthly Retail Trade Survey (MRTS). The ARTS collects data on total annual sales, e-commerce sales, sales taxes, end-of-year inventories, purchases, total operating expenses, and gross margins. The ARTS includes employer businesses classified in the retail sector. Firms without paid employees are included in the BLS estimates through imputation or administrative data provided by other federal agencies. Annual values from the ARTS are adjusted based on the Merchandise Line Sales from the Economic Census (or Product Line Sales in later years). BLS uses the most detailed ARTS sales data available for deflating and then aggregates it to the 4-digit NAICS industry level. BLS does not adjust for resales or changes in inventories.¹⁸

¹⁷ Real sectoral output is output that is sold to entities outside of the industry. For retail trade industries, real sectoral output is virtually the same as real gross output (gross sales). Note, BLS also publishes these measures for two 2-digit industries and 13 3-digit industries.

¹⁸ Resales in retail trade are negligible.

The industry-specific implicit deflators for output (sales) are constructed by dividing the index of current-dollar sales for all establishments in the industry, which are available annually, by the corresponding Tornqvist index of annual constant-dollar sales constructed using product deflators. The Tornqvist index is constructed by combining product-line deflators (available for each industry annually) with product-line-by-industry sales (available for each industry every five years and interpolated for the intervening years). The first step is to match sales for each product line to the appropriate Consumer Price Index (CPI) or Producer Price Index (PPI),¹⁹ and then deflate the individual product-line sales.

BLS measures *labor input* as the total annual hours worked by all persons in an industry. This measure is constructed by combining data from three surveys: the Current Employment Statistics (CES) survey, the Current Population Survey (CPS), and the National Compensation Survey (NCS). BLS uses data on employment and average weekly hours *paid* for production and nonsupervisory employees (henceforth referred to as nonsupervisory workers) from the CES. The NCS data, which include information on paid leave and other types of compensation, are used to estimate an hours-worked-to-hours-paid ratio that adjusts total nonsupervisory hours from an hours-paid to an hours-worked basis by removing paid vacation accrued and sick leave taken. To estimate supervisory employee average weekly hours, BLS uses data from the CPS to calculate a ratio of supervisory to nonsupervisory employee average weekly hours worked, which is then multiplied by the adjusted CES nonsupervisory employee hours (worked). The total hours worked by all employees is the sum

¹⁹ BLS uses the CPI Research Series because the series is more consistent over time. The official CPI is never revised, whereas the CPS Research Series is revised to incorporate the current methodology into the historical data. In cases where there are multiple deflators for a product-line definition, BLS creates a deflator for that product line using relative importance values as weights and then uses that to deflate the single product line. This situation occurs when the deflators are defined at a more detailed level than the product lines.

of the hours worked by nonsupervisory workers and supervisory workers.²⁰ Self-employed and unpaid-family-worker (SE) hours are then added to this total. SE employment and average weekly hours are obtained from the CPS, and multiple data sources (including nonemployer statistics, the Internal Revenue Service, and ARTS, among others) are used to allocate SE hours to industries. The level of aggregation for the inclusion of SE workers is between the 6-digit and 4-digit NAICS industry level.

B. Comparing Micro-Aggregated Data to Published Industry Data


In this section, we compare our micro-aggregated estimates to the official data published by BLS, covering the 1987–2017 period. Based on earlier work comparing similar business data across the two government agencies, we expect that there will be some systematic differences between these measures (Elvery et al., 2006). Even though differences in the levels of the micro-aggregated and published first moments do not directly affect our conclusions about dispersion (because our measures are mean invariant), it is useful to know how far apart the two sets of estimates are. If the first moments are close, then it is more reasonable to think of the micro-based second moments as measuring variation around the published first moments. We start by comparing employment and nominal sales levels and then compare productivity levels. We use the BLS average weekly hours series and deflators directly from BLS industry productivity data, which implies that all level differences come from differences in sales, employment, and sample coverage.

²⁰ BLS recently changed its method for estimating employee hours. The new method uses the CES all-employee hours series as the main data source. It adjusts these data from an hours-paid concept to an hours worked concept using data from the NCS (to account for paid time off) and the CPS (to account for off-the-clock work). In the next draft of this paper, we will use these estimates for our comparison.

B.1 Input and Output Measures

In Figure B1, we compare total nominal retail sales growth across the BLS and Census microdata series. We do not deflate these series since the same deflator would be used for both. The two series follow each other very closely in terms of both changes and levels. Both series start at about \$2.3 trillion in 1987. Sales growth was modest between 1987 and 1992, accelerated between 1992 and 2007, fell between 2007 and 2012, and resumed the accelerated growth between 2012 and 2017. The series diverged starting in 1997, with the BLS series growing at a slightly faster rate.

Figure B1: Retail Sales




Notes: BLS sales is the sales according to the BLS industry statistics, while sample sales is the sales of establishments in the CRT sample.

Next, Figure B2 shows the total number of employees in the retail sector from each series. BLS Retail industry employment in 1987 was just below 14 million workers.

Employment grew slowly between 1987 and 1992 and then accelerated between 1992 and 1997 when it grew by over 10 percent. Employment growth slowed after 1997, with a dip in employment between 2007 and 2012. The aggregated Census microdata follow this general trend but at somewhat lower levels compared with the BLS data. Nearly all the differences between the series come from BLS, including self-employed and unpaid family workers in the employment totals, according to an analysis using unreleased BLS estimates. An implication of this difference is that self-employed and unpaid family workers have become less important, resulting in slower employment growth in the BLS series.

Figure B2: Retail Employment Levels



Notes: BLS employment is the employment according to the BLS industry statistics. Sample employment is the employment of establishments in the CRT sample. The two employment measures differ, in part, due to their treatment of nonemployer establishments.

B.2 Productivity Growth

Figure B3 compares 5-year growth rates of the BLS series and our micro-aggregated

series. The growth rates of the BLS series are calculated from indexes of labor productivity. From the sample microdata, we calculate employment-weighted average labor productivity across all establishments in the sample and then calculate 5-year growth rates. The two series exhibit nearly identical growth rates in three of the six years (1992, 2007, and 2017). In the 1997-2002 and 2002-2007 periods, the microdata show five-year growth rates of around 10 percent, while the BLS estimates are almost 20 percent and over 25 percent.²¹ Between 2002 and 2007 the microdata show productivity declines while the BLS data show slight increases. On average, over the entire period, the BLS series exhibits faster productivity growth than the microdata series, with differences in employment growth being the main culprit.

²¹ The conversion from SIC to NAICS codes seems like a logical place to look for the reasons behind these differences. Also, the BLS and Census business registers do not always assign the same industry codes to the same establishments. In the official BLS productivity statistics, the employment data are based on BLS industry coding while output data are based on Census industry coding. Both employment and output are based on Census industry coding in our micro-aggregated estimates.

Figure B3: 5-Year Productivity Growth Rates by Data Source



Notes: Microdata labor productivity growth plots five-year growth rates using sales and employment from the CRT and hours from BLS industry productivity statistics. Productivity is calculated for each establishment and then averaged to calculate aggregate productivity. The establishments in the top and bottom one percent of the productivity distribution for each industry group are excluded from the calculation. Official BLS Measure is the percent change in index levels.

Table A1: Regression of log labor productivity on the interactions between establishment entry/exit and firm structure, controlling for industry and year effects.

		1987-1997	1997-2007	2007-2017
Continuers				
	Local	0.136*** (0.033)	0.0926*** (0.033)	0.0877*** (0.030)
	Regional	0.156*** (0.054)	0.170*** (0.042)	0.167*** (0.043)
	National	0.191*** (0.053)	0.103* (0.056)	0.196*** (0.066)
Continuers in End-Year				
	Single-Unit	0.0819*** (0.024)	0.0961* (0.053)	-0.0387 (0.034)
	Local	0.0519 (0.032)	0.102* (0.059)	0.00341 (0.028)
	Regional	0.0759 (0.055)	0.0871 (0.052)	0.0186 (0.032)
	National	0.0728 (0.043)	0.264*** (0.064)	-0.00861 (0.046)
Exiters				
	Single-Unit	-0.287*** (0.017)	-0.287*** (0.026)	-0.360*** (0.020)
	Local	-0.116*** (0.025)	-0.107** (0.050)	-0.152*** (0.042)
	Regional	-0.0443 (0.043)	-0.0383 (0.075)	-0.0615 (0.064)
	National	-0.0592 (0.053)	-0.131 (0.080)	-0.104 (0.075)
Entrants				
	Single-Unit	0.0403 (0.033)	0.00287 (0.052)	-0.0808* (0.043)
	Local	0.193*** (0.048)	0.236*** (0.061)	0.0425 (0.058)
	Regional	0.280*** (0.061)	0.332*** (0.072)	0.164** (0.070)
	National	0.216*** (0.067)	0.351*** (0.079)	0.176** (0.075)
	Constant	-2.588*** (0.0210)	-2.387*** (0.0296)	-2.129*** (0.0322)
Observations		2,065,000	2,062,000	2,019,000
R-squared		0.473	0.365	0.279

Notes: The reference group is single-unit continuing establishments. Robust standard errors clustered by NAICS-4 and are in parentheses. Values were rounded according to Census Bureau disclosure rules. Column 1 uses establishments that exist in 1987 and/or 1997, column 2 uses 1997 and/or 2007, and column 3 uses 2007 and/or 2017. Sample: Sample 1 CRT by decade. Asterisks denote the following levels of significance: ***: 1%, **:5%,*:10