# The Internal Geography of Firms *

Dominick Bartelme
University of Michigan

Oren Ziv
Michigan State University

July 2019

**Abstract**

We document four facts regarding the geographic location patterns of firms. Firms' establishments are geographically clustered; as they expand, firms become more dispersed but continue to cluster their new establishments; larger, more productive firms are more dispersed even after conditioning on the number of establishments; and establishments that are further away from the geographic center of the firm are smaller. These results are largely driven by the behavior of small and medium-sized firms. These findings are consistent with the hypothesis that smaller firms face significant costs of dispersion which exert a powerful influence on their geography.

*JEL codes*: R32,L25

*Keywords*: economic geography, agglomeration, firm location, firm growth

*We thank Camillo Acosta-Mejia, Andrei Levchenko, Martha Stinson, Junfu Zhang, Jens Wrona and numerous seminar participants for their comments, and C. Luke Watson for excellent research assistance. Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. Contact: dbartelm@umich.edu, orenziv@msu.edu.

# 1 Introduction

Multi-unit firms account for most of the economic activity in the U.S., employing the majority of payroll workers in 2012. An important decision facing such firms is where to locate their establishments, both in absolute terms and in relation to one another. Geographic dispersion allows firm to maximize access to markets and local inputs, but can generate costs for the firm, such as transportation and time costs associated with shipping intermediate inputs between plants or management and coordination costs that increase with distance. An established literature documents some of these costs and their effects on plant performance in specific settings.[1] However, we know little about the general patterns of internal firm geography that are generated by these forces.

This paper uses Census microdata to document four facts regarding the geographic location tendencies of firms within the Continental United States. First, firms' establishments tend to be geographically clustered in the cross section relative to unrelated establishments in the same industries. To document this pattern, we match each real firm to a synthetic firm with identical industrial structure by replacing each establishment in the firm with a randomly chosen establishment in the same (6-digit) industry. We then compare the geographic dispersion of the real and synthetic firms. We find significantly more clustering in the real firms across all sectors, a result that is robust to matching on additional plant characteristics.

Second, growing firms expand outwards from their center, but as they do, they choose new locations that are clustered near their existing establishments. Due to potential selection in growing firms and the fact that cross-sectional clustering may represent location decisions made decades ago, these time-series results are not implied by the cross-section but must be established independently. To do so, we compare the observed growth in each firm's geographic footprint when it adds an establishment to growth we would expect if they chose the new establishment randomly from the distribution of establishments in the same industry. We find that new establishments are on average

[1] Several papers investigate the link between proximity (via trade, travel and communication costs) and plant performance, especially (Giroud, 2013; Kalnins & Lafontaine, 2013; Eichholtz et al., 2015; Alcacer & Delgado, 2016; Charnoz et al., 2018; Atalay et al., 2017). These papers generally finds that distance to management reduces plant performance

further away from the firm's center than existing establishments, but that the observed increase in geographic dispersion is lower than what we would expect if the new establishments simply followed industry location patterns.

Third, larger and more productive firms are more geographically dispersed. We regress firm sales and labor productivity on the firm's dispersion interacted with the number of establishments in the firm and controlling for industry-year effects. While we would expect firms with more establishments to be larger and more productive (and they are), we find that employment and sales per worker are higher for firms whose establishments are more dispersed even after conditioning on the number of establishments. Fourth, establishments that are further away from the firms center tend to be smaller. We regress establishment employment and sales per worker on distance to the firm's center, controlling for establishment firm-industry-year effects, and establishment age. We find that plants that are further away from their firm center have significantly lower employment than comparable plants that are closer, although they have similar sales per worker.

Crucially, our analysis also reveals substantial and systematic heterogeneity in these correlations across the firm size distribution. Each of these patterns is consistently more pronounced for firms with fewer establishments, and attenuates or disappears entirely for larger firms with more establishments. Large firms appear to "defy gravity" in their intra-national geography, suggesting that they either face systematically lower costs or reap systematically larger benefits of geographic dispersion.

These facts are consistent with the hypothesis that there are pervasive costs of geographic dispersion that induce firms to agglomerate in space, although we do not exclude the possibility that other factors such as spatial correlation in firm-level supply or demand shocks play a role as well. The literature has documented the existence of various frictions in moving goods, people or ideas across space that impede coordination between distant plants (Giroud, 2013; Kalnins & Lafontaine, 2013; Eichholtz et al., 2015; Alcacer & Delgado, 2016; Charnoz et al., 2018; Atalay et al., 2017). Our results suggest that these costs exercise a powerful influence on firm geography, albeit quite unevenly across the firm size distribution. To the extent that the costs are responsive to policies such as infrastructure investment, there may be substantial long-run efficiency gains from the spatial reorganization of firms. Our results point to the need for further quantitative research on this topic.

Our paper contributes to the empirical literature on plant location in multi-unit firms. Our finding of geographical clustering in the cross-section has antecedents in work of Henderson & Ono (2008) on headquarter location of manufactures and Behrens & Sharunova (2015) on Canadian manufactures, which finds evidence for clustering in the cross section as well. Holmes (2011) studies the geographic expansion path of Walmart and shows that new stores tended to expand the geographic footprint of the firm. Our contribution is to generalize these findings to all sectors of the economy, explore clustering in the time series, and present some entirely new, related facts (e.g. fact 3). Furthermore, we document for the first time the pervasive and meaningful heterogeneity in these patterns across the firm size distribution.

Our paper is also related to work on the international geography of multinational activity (see Antràs & Yeaple (2014) for a review). Multinational expansion within the firm is also geographically and culturally concentrated (see Alfaro & Chen (2018) for a review), and larger and more productive firms are more likely to become multinationals (Tomiura, 2007). This concentration is plausibly linked to trade and communication costs which reduce affiliate productivity (Keller & Yeaple, 2013; Bahar, 2014; Gumpert, 2018). Our work shows patterns similar to the key finding of Yeaple (2009), that multinationals enter markets in a way that may be responsive to distance, and the finding of Irarrazabal et al. (2013) and Antràs & Yeaple (2014) that within firms, sales fall with distance. However, we show that these relationships attenuate with firm size and that the largest firms “defy gravity.” This heterogeneity has not yet been explored in the literature on multinationals.

Section 2 describes our data and the measures we construct from it. In section 3, we lay out our facts. Section 4 discusses economic implications and concludes.

## 2 Data and Measurement

**Measuring Dispersion** We are interested in the geographic footprint of the firm. Our principle measure of this will be the average number of miles between firms' constituent plants and the geographic centroid of the firm. This provides a computationally tractable measure of these establishments' dispersion in a way that includes multilateral resistance – taking into account each establishment's relation to all others in the firm. Throughout this paper, we refer to this measure as the "dispersion" of the firm and refer to the centroid as the "center of the firm".

Appendix A formally describes this measure and compares it to the mean of the Duranton & Overman (2005) CDF, for which our measure is a close proxy. Appendix Figure 1a plots the relationship between these measures for a set of synthetically generated establishments and reports an adjusted $R^2$ between the measures of 0.94.

Different theoretical foundations for the patterns we report may suggest alternative measures of the firm's footprint. For example, if the reader holds the hypothesis that the distance costs underlying our results are management costs, she may prefer to look at dispersion with respect to firm headquarters. Our results are robust to measuring plants' distance to the firm's oldest establishment, as well as to the reported headquarters of the firm, for the sub-sample for which firm headquarters is reported.

**Sample** Our main sample uses multi-unit firms responding to Economic Censuses which we observe in 5-year intervals between 1992 and 2012. Facts 1 and 3, in Section 3 use this sample and the sub-sample comprised of firms that are modally manufacturers (by sales). Fact 4, which uses establishment-level variation within firms, uses the main sample and a sub-sample of manufacturing establishments in any firm. Fact 2 uses a sub-sample of firms that expanded between Census years. Here we discuss how each sample and sub-sample is created.

Our main sample consists of the universe of establishments that are a part of firms with two or more establishments in a given year in the Longitudinal Business Database (LBD) and responding to Economic Censuses between (and including) 1992 and 2012, where all establishments in the firm have positive sales, and more than one employee, and where the firm exists in at least two zip codes so that an internal firm distance can be measured. To ensure data quality, we remove extreme outliers: any establishment above the top 0.05 percentile in sales, employment, payroll, and (for manufacturing establishments) value added.[2]

Because the patterns in the data vary by sector, we also extract a sub-sample of firms whose modal sales are in the manufacturing sector.[3] All statistics in Section 3 are reported for both the full sample and the manufacturing firm sub-sample. For establishment-level regressions we use a sub-sample of manufacturing establishments.

Fact 2 focuses on the growth patterns of firms. We isolate firms that expanded by comparing the number of establishments in each firms between each Census wave, isolating firm-year observations where the firm moved one size class (defined below) from the preceding Census. We find approximately 44,000 and 4,900 firm-year observations across our full sample and our sub-sample of manufacturing firms, respectively. Because a significant fraction of establishments report zero employment in their first year, we include in this sample establishments in their first year that report fewer than two employees but that do report more than one employee in subsequent Census waves.

Appendix A provides further description and summary statistics for our samples and the measures we employ in our main analysis.

**Firm and establishment measures** To operationalize our measurement of within-firm dispersion on this sample, we use the firm identifiers and establishment zip codes from the LBD to first calculate the geographic center of the firm, then each establishment's distance from that point. Key outcome variables in Section 3 will be employment and sales per worker, which are observed or calculated at the establishment and firm levels.

[2]However, all our results are robust to eliminating all of these cuts as well as to restricting our sample to establishments with greater than five employees.

[3]Firms may have both manufacturing and non-manufacturing operations. To isolate firms where the chief activity is manufacturing, we first follow Fort & Klimek (2016) to assign consistent NAICS codes to establishments before 1997. We then calculate the sector of the firms' modal establishment, weighing establishments by their sales. Manufacturing firms are then defined as firms where the modal dollar of revenue is generated at a manufacturing establishment. Weighing establishments by employment does not change classifications significantly and does not change results. As a further robustness check, we divide firms into firm-sector groupings and repeat our empirical exercise on these sub-firm units. Our results are robust to this cut of the data as well.

As a control, we use age group (0-5, 5-15, and 15+ years) for establishments and firms (measured as the maximum age group of the firms' establishments).

**Size classes** We present all of our results in Section 3 by firm size by grouping to explore heterogeneity across groups of firms with different numbers of establishments. We group firms according to 22 size categories: one unique group for each firm with between 2 and 10 establishments, then by 5 for firms with between 11 and 40 establishments, then by 10 for firms with up to 100 establishments, and a final group for firms with more than 100 establishments. The size classes are chosen in part to satisfy Census disclosure requirements.

## 3 Facts

This section documents four patterns in the internal geography of firms. First, small and medium-size firms exhibit spatial clustering in their establishments, relative to a counterfactual with random establishment placement, while large firms do not. Second, this result also holds for the growth of firms through new establishments: small and medium firms locate new establishments close to old ones while large firms do not. Third, conditional on the number of plants, more successful firms have more dispersed plants. This relationship is again most pronounced for small and medium-sized firms. Fourth, establishments that are more distant from the firm center have lower sales and employment. We find significant heterogeneity in these patterns across size classes.

### 3.1 Geographic Clustering of Establishments

Our first result concerns the cross-sectional patterns in firm dispersion. Figure 1a plots the average log dispersion by size class for all firms in our sample with 2 or more plants, with standard errors clustered at the industry level, while Figure 1b plots the means for the manufacturing only sub-sample. Small firms are highly geographically clustered: for firms with only two plants, the average distance to the firm center is less than two miles. Geographic dispersion rises steadily with the number of plants: firms with 41-45 plants have an average distance of about 150 miles. A similar pattern holds for the manufacturing only sub-sample, although manufacturing firms are significantly more dispersed than the average firm along the entire size class distribution.

How different is this pattern from the geographic clustering we see in plant location across firms? That is, how tightly clustered are a firm's plants relative to the clustering behavior of a "similar" group of plants outside the firm? To answer this question, we construct a matching set of synthetic firms with plant locations drawn randomly from the set of "similar" plants, and use the dispersion observed in the synthetic firms as a baseline for comparison to the data.

For each 6-digit NAICS code $k$ and year $t$, and each firm $i$ with 2 or more plants, we group all the plants in the sample that belong to industry $k$ at time $t$ to create sets of plants $P_{kt}$ with the observed geographic locations and a set of firms $F_t$ with their observed number of plants, their industries and their locations. For each actual firm in $f_{it} \in F_t$, we create a synthetic firm $\tilde{f}_{it} \in \tilde{F}_t$ with the same number of plants in the same industries, but with the geographic locations of its plants randomly chosen, with equal probability and with replacement, from the relevant $P_{kt}$. We then compute the average log distance from the firm's center by size class for the synthetic firms, exactly as we did with the actual firms, and compare the results. This procedure controls for any industry agglomeration and co-agglomeration patterns in the data, allowing us to separate the within-firm component of geographic clustering from the across-firm components.[4]

Figures 1a and 1b plot the results from the synthetic baseline alongside the actual data. In contrast to the pattern of rising dispersion found in the data, the average dispersion in the synthetic baseline is relatively flat across size classes and similar for both sub-samples, averaging about 450 miles. For small and medium-sized firms, the baseline dispersion is far greater than the actual dispersion, indicating that these firms are highly geographically clustered relative to the typical group of plants in their industries. The gap between the baseline and the actual shrinks substantially as we move up the size class distribution, disappearing almost entirely for manufacturing firms with more than 10 plants and for the very largest non-manufacturing firms. Large firms "defy gra-

[4]Note that if firms are large in their industries, then the observed industry agglomeration patterns will partly be driven by the within-firm patterns. As a consequence, our procedure underestimates the true counterfactual dispersion, and therefore underestimates the degree to which firms are clustered relative to the counterfactual.

**Figure 1: Cross-Sectional Firm Dispersion**

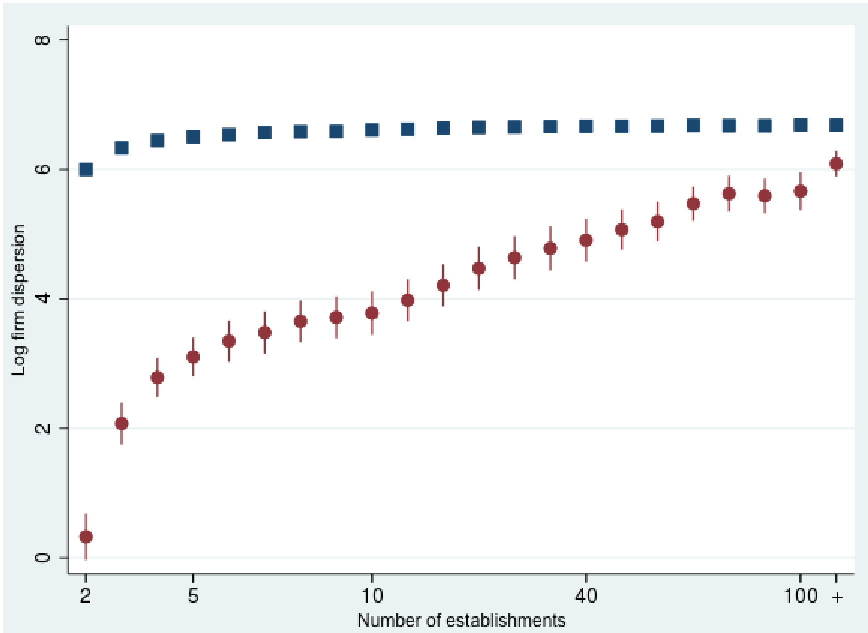


(a) All firms

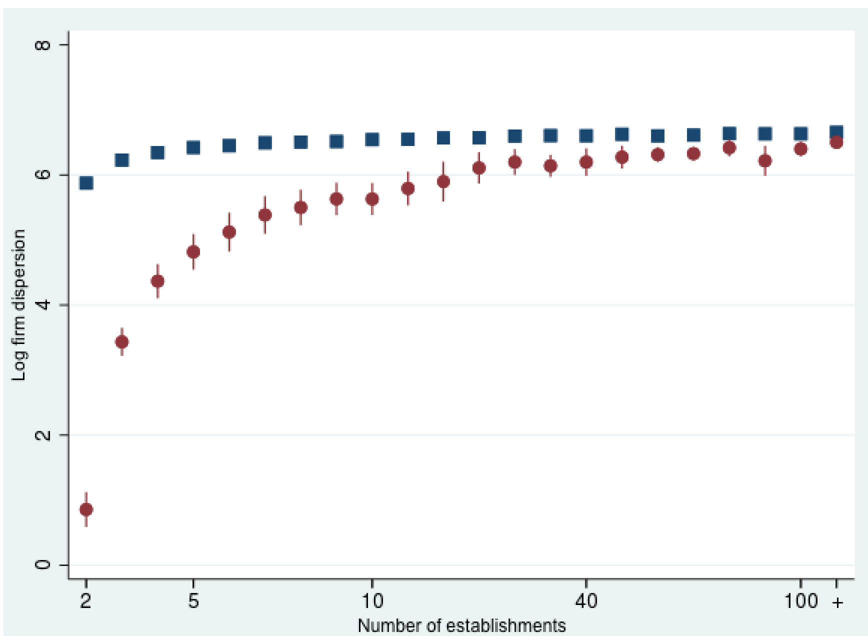


(b) Manufactures

*Note*: Red circles plot the average log mean establishment distance from firm centroid for each firm size category. Blue squares plot the corresponding measure for synthetically constructed firms. Estimates include all firm-year observations across five Economic Census waves between 1992-2012. All standard errors are clustered by firm modal 4-digit industry.

vity" in their location choices, particularly in manufacturing.

The results in Figures 1a and 1b are robust to a number of sample selection criteria and specification changes. One concern is that, while our synthetic baseline controls for the industry composition of each firm's plants, there may be other omitted plant characteristics that drive agglomeration patterns both within and across firms. In a robustness check, we use the Census of Manufactures to match plants based on the full set of products produced and inputs used in production when constructing the synthetic firms. We find very similar results to those reported in Figures 1a and 1b. The details and results of this and other robustness checks and the inclusion of additional controls are reported in Appendix B.

### 3.2 Spatial Growth of Firms

Our second set of findings regards the spatial pattern in which firms add new establishments as they expand. When a firm decides to add a plant, how does the location of its existing plants influence the placement of the new plant? Are firms with $n$ establishments more agglomerated than what we would find if a firm with $n-1$ establishments randomly chose a location for its $n-th$ establishment? Figures 1a and 1b already show that, on average, larger firms are more dispersed, and one might think that the spatial pattern of firm growth is already revealed by differencing the coefficients across size classes. However, as we explain below, there are three issues with this naive approach.

First, the set of firms that expand by a size class are not randomly selected from the firms in the previous size class. Expanding firms may be systematically different in their initial location choices from the firms that do not expand. Second, since plants tend to be quite long-lived, the cross-sectional results in Figures 1a and 1b reflect location choices made in both the recent and more distant past. Given the sweeping technological and economic changes of the past 50+ years, it would hardly be surprising to find that the location choices made by firms today are systematically different than those of the past. Third, the synthetic baseline used in the cross-section is not useful for studying spatial growth patterns. We would like to compare the change in dispersion associated with the actual location choice made by an expanding firm to the change associated with a counterfactual choice, *taking its existing plant locations as given.* There is no way to do that with the information in Figures 1a and 1b.[5]

To address these issues, we proceed as follows. For each Census year[6] and initial size class, we find every firm that moved up by exactly one size class relative to the previous Census year and compute the change in its log dispersion.[7] We then average across firms and years, by size class, to get the average change in log dispersion for this subset of expanding firms, for each size class.[8]

The circles in Figures 2a and 2b plot the results for all firms and manufacturing, respectively. In both samples, firms that begin with roughly 7 or fewer plants tend to increase their average dispersion when they expand, at a rate decreasing in firm size. Firms that begin with roughly 8 or more plants do not significantly increase their geographic footprint when they expand. This result is different than what one would get by differencing the cross-sectional averages, which imply increasing dispersion across a wider range of firm size classes than found in Figures 2a and 2b. This indicates that either differential selection or history, or both, are at work.

How do these expansion patterns compare to those that would be observed if firms did not consider the location of their existing plants when expanding? We construct a new synthetic baseline by starting with the same underlying sample in Figures 2a and 2b, the set $F_{zt}$ of firms that have $z$ plants in Census year $t$ and had $z-1$ plants in Census year $t-1$, where $z$ refers to the size class, for each year and size class. For each firm $f_{it}$ in this set, we construct a synthetic firm $\tilde{f}_{it}$ by keeping the continuing plants of $f_{it}$, dropping the new plants and replacing each one with a plant drawn randomly (with replacement) from the set of plants $P_{kt}$ in the same 6-digit NAICS code that were active in that year.[9] We then compute the average increase in log dispersion across years, within

[5]The synthetic baselines plotted in Figures 1a and 1b show that, in a world where all plant location choices are random, firm expansion is not associated with any significant change in dispersion.

[6]We use Census waves because establishment birth years between Census waves are imputed for multi-unit firms in the LBD. The total number of establishments born between Census waves is not. We thank Martha Stinson for alerting us to this feature of the data.

[7]We observe some firms that simultaneously open and close multiple establishments. We use the net change the number of establishments to classify firms as having moved up one size class.

[8]We discuss the robustness of our results to moves of more than one size class below.

[9]From the set of expanding firms, some subset had also closed establishments between years. These

**Figure 2: Time Series Firm Dispersion**

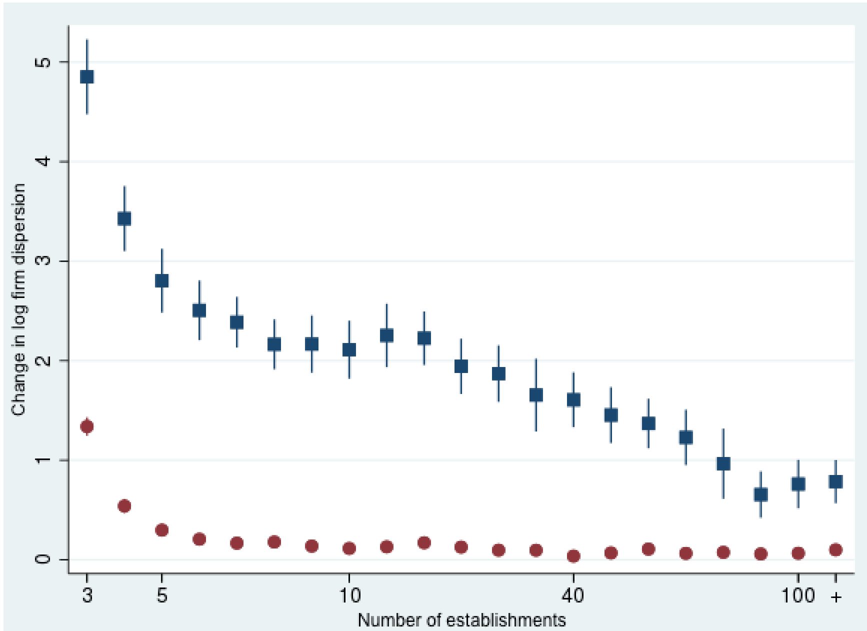


(a) All firms

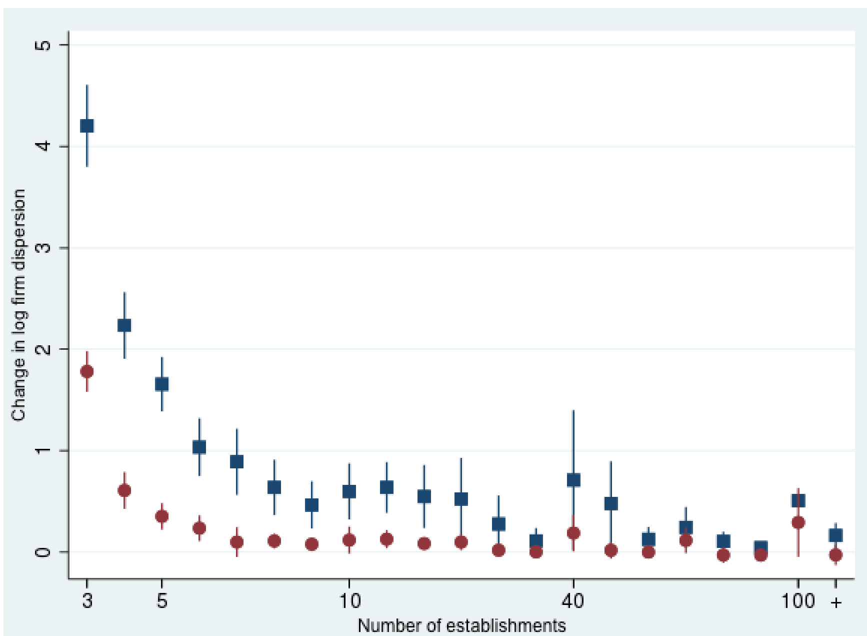


(b) Manufactures

*Note*: Red circles plot growth in log average establishment distance to firm centroid for firms observed moving up one firm size category between 5-year Census waves from 1992-2012. Blue squares plot the corresponding measure for firms with synthetically constructed expansions. All standard errors are clustered by firm modal 4-digit industry.

size class, for the synthetic baseline in the same way as we did for the data.

The results of the synthetic baseline are plotted using squares in Figures 2a and 2b. We see that for small and medium firms in both samples, the increase in dispersion predicted by the random location model is substantially greater than what is actually observed in the data. The gap between the random location model and the data decreases with firm size, much more rapidly for manufacturing firms.

As was the case for Fact 1, these results are robust to alternative specifications and samples. Appendix B describes the results of alternative construction of the synthetic baseline including alternative treatments of mergers and additional samples and specifications. In one set of exercises we examined firms that moved more than one size class between Census waves. There are fewer of these firms and describing the full transition matrix is prohibitive because of small sample size. At the same time, the heterogeneity in dispersion between size classes suggests pooling results across different size class changes could be misleading. What we can clear for release is that, for the sample that includes all sectors, the confidence intervals of the data and synthetic baseline did not overlap for any transition between size classes with fewer than 25 establishments. The same statement applies to the manufacturing sub-sample for transitions across up to three size classes with fewer than 10 establishments.

Do the differences between the cross-sectional and time series results reflect differential selection of growing firms, the influence of historical plant location patterns, or both? Figures 3a and 3b plot the cross-section of firm dispersion by size class (analogous to Figures 1a and 1b) for 1976 (the orange triangles) and 2012 (the red circles). While some plants that were active in 1976 are still active in 2012, major changes in location patterns should be reflected in differences in cross-sectional log dispersions over this time period. The figures show some evidence for modest increases in the geographic dispersion of firms over time, especially for the smallest size categories and for manufacturing firms. However, the differences over time do not fully account for the differences between Figures 1 and 2.

choices also impact the footprint of the firm, so to simulate the new footprint of the firm with randomized location choices, we also randomize the establishments which closed from the firms' base-year locations.

**Figure 3: Firm Dispersion in 1976 and 2012**

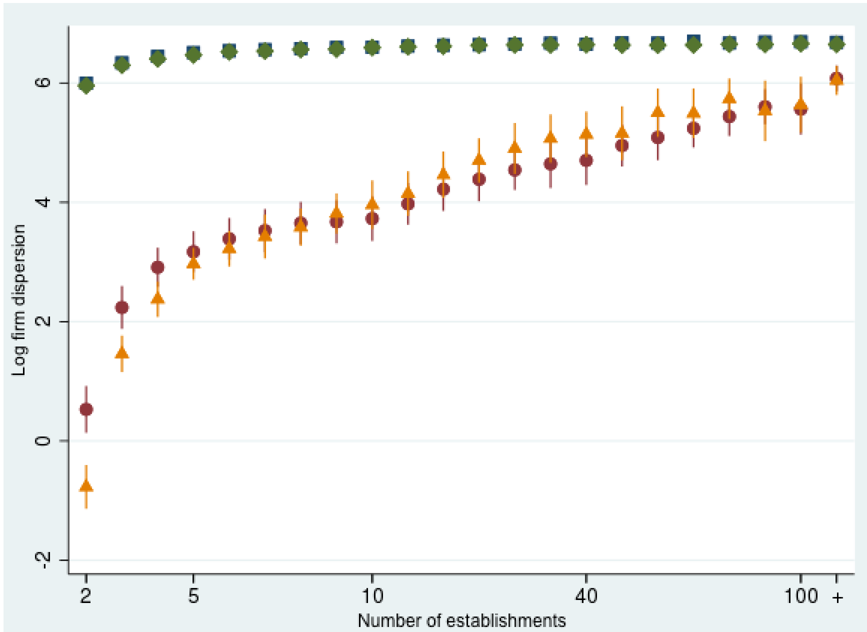


(a) All firms

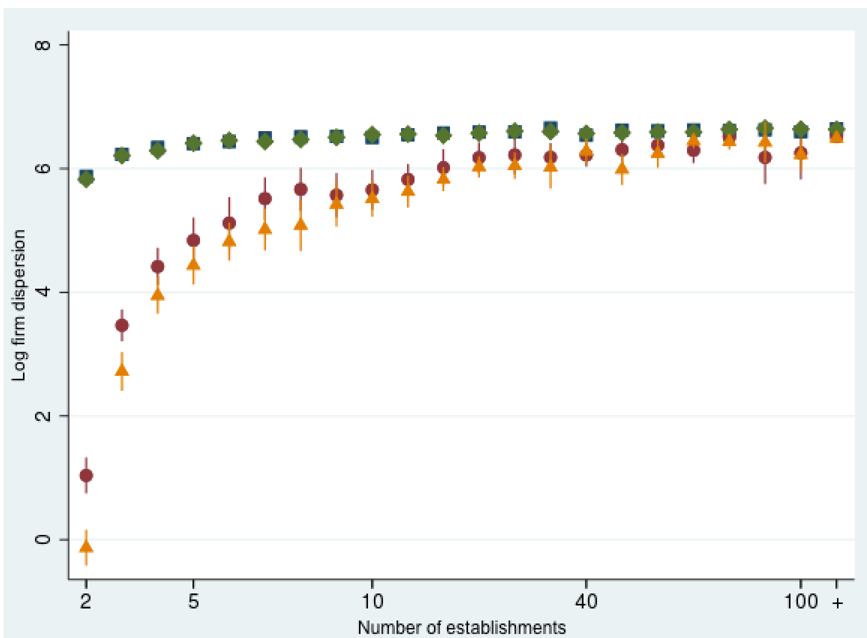


(b) Manufactures

*Note*: Red circles plot the average log mean establishment distance from firm centroid for each firm size category for firms in the 2012 Economic Census. Orange triangles plot the corresponding measure for firms in the 1976 Economic Census. Blue squares and green diamonds plot the same measures of synthetic firms in 2012 and 1976 respectively. All standard errors are clustered by firm modal 4-digit industry.

### 3.3 Firm Geography and Productivity

Our third set of facts relate the geographic footprint of the firm to measurable outcomes, chiefly employment and sales per worker. We interpret these variables as jointly yielding a robust qualitative measure of firm productivity based on the idea that, all else equal, firms exhibiting both higher employment and higher sales per worker must be more profitable. Relative profitability, in turn, is a measure of relative productivity in demand systems (such as CES) that link relative profitability to the relative Marshallian surplus generated by the firms.[10] These measures have their limits, especially as quantitative indicators, but involve far fewer conceptual and practical difficulties than computing measures of firm-level revenue TFP, especially for non-manufacturing firms.

We find that firms with higher employment and sales per worker tend to be more dispersed geographically. That this relationship holds unconditionally for total employment is not surprising, since it follows from combining our finding that firms with more plants are more dispersed with the trivial observation that firms with higher employment tend to have more plants. What is more interesting is the extent to which these relationships also hold after conditioning on the number of plants in the firm.

We document these patterns by using regressions of the form

$$y_{it} = \alpha_{kt} + \alpha_z + \beta_z x_{it} + \gamma w_{it} + \epsilon_{ikt} \tag{1}$$

where $y$ is log firm dispersion, $i$ indexes firm, $k$ indexes the firm's 4-digit modal industry, $z$ indexes the size class of the firm and $t$ indexes time. Here $x_{it}$ is the firm characteristic of interest and $w_{it}$ is a vector of other firm characteristics that serve as controls, which in our baseline includes age groups. As in our previous analyses, we allow the coefficient of interest $\beta_z$ to vary by the size class of the firm.

Figures 4a through 4d report $\beta_z$ for each size class using employment and sales per worker as outcomes, for all firms and manufacturing firms separately. On average, firms with higher employment and average labor revenue productivity are more dispersed than similar firms with the same number of plants, across all samples. However, these

[10]We provide a formal treatment in Appendix C. We show there that, in general, sales per worker by itself provides little information regarding relative productivity.

**Figure 4: Distance vs Productivity by Firm Size**

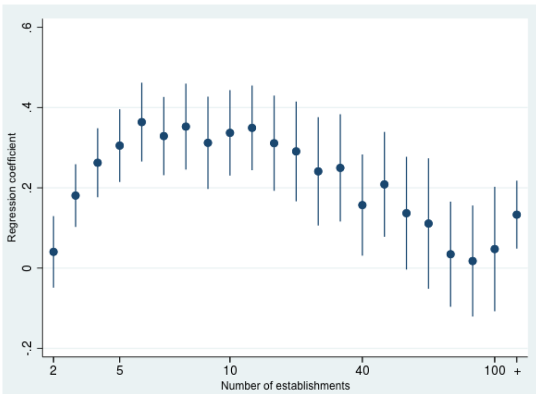


(a) Log employment, all firms

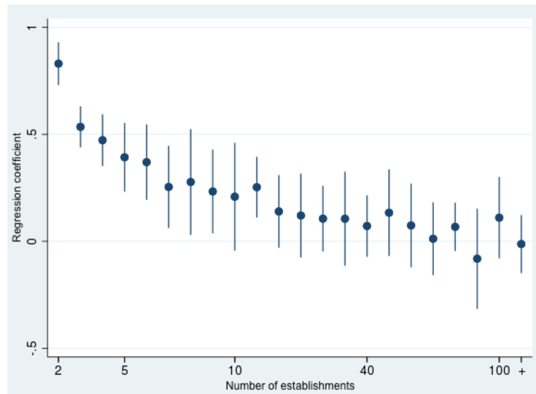


(c) Log employment, manufactures

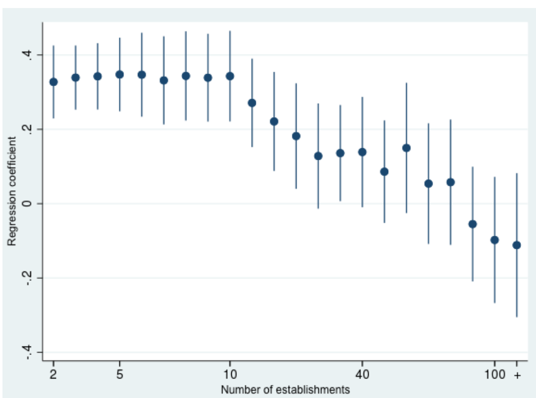


(b) Log sales per worker, all firms

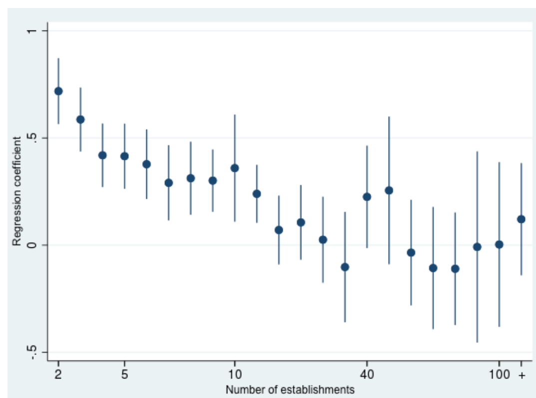


(d) Log sales per worker, manufactures

*Note*: Plotted points are firm-category specific slopes describing the relationship between firms' log employment (top) or sales per worker (bottom) and log average distance to firm center by firm size category as described by the specification in equation 1, which includes controls for firm age category fixed effects and 4-digit industry-year fixed-effects. Data include firm-year observations for firms across five Economic Census waves between 1992-2012. All standard errors are clustered by firm modal 4-digit industry.

average effects are driven primarily by small and medium-size firms, and tend to diminish or disappear as the firm size class increases. This pattern is particularly evident in the manufacturing sub-sample, where the effects are not statistically distinguishable from zero for firms with more than 10-15 plants, depending on the measure.

Appendix B discusses several robustness checks we perform. Our results are robust to including firm geographic controls, to measuring dispersion as distance to headquarters and to the firm's oldest establishment, to measuring dispersion and firm outcomes within firm-by-sector units, and to restricting the sample to establishments with five or more employees. Our results are also robust to alternative firm outcome measures including sales, value added per worker and average wage.

### 3.4 The distance-productivity relationship within the firm

Our final set of results concerns the within-firm relationship between plant characteristics, such as employment and sales per worker, and their distance from the geographic center of the firm. As in the across-firm comparisons, all else equal plants that exhibit both higher employment and sales per worker will be more profitable for the firm. However, within-firm comparisons face the additional challenge of the possibility of joint production and/or shared inputs across plants. Even more caution is required in interpreting differences in employment and sales per worker across plants as indicators of productivity differences in this case.

We run regressions of the form

$$y_{ijkt} = \alpha_{ikt} + \beta_z x_{ijkt} + \gamma w_{ijkt} + \epsilon_{ijkt} \tag{2}$$

where $y$ is the log distance of the plant from the firm's centroid, $x$ is an establishment characteristic, $i$ indexes firm, $j$ indexes establishments, $k$ indexes the establishment's 4-digit industry, $z$ indexes size class and $t$ indexes years. Fixed effects at the firm-industry-year levels are included. Plant-level controls $w_{ijkt}$ are age groups in our main specification.

Figures 5a through 5d reports the resulting $\beta_z$s using establishment employment and sales per worker as outcomes. In small and medium-sized firms, plants that are further

**Figure 5: Distance vs Productivity by Firm Size, Within Firms**

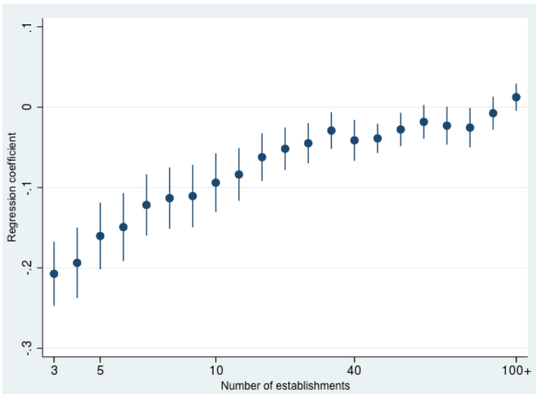


(a) Log employment, all firms

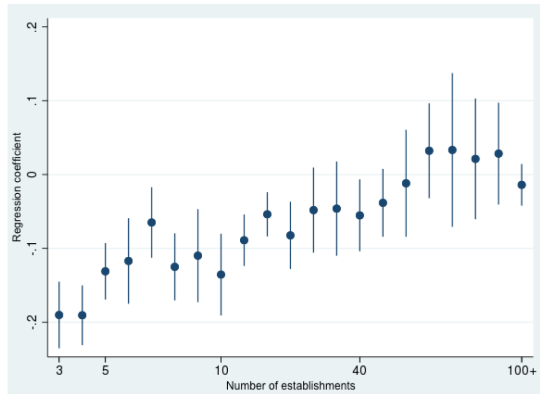


(c) Log employment, manufactures

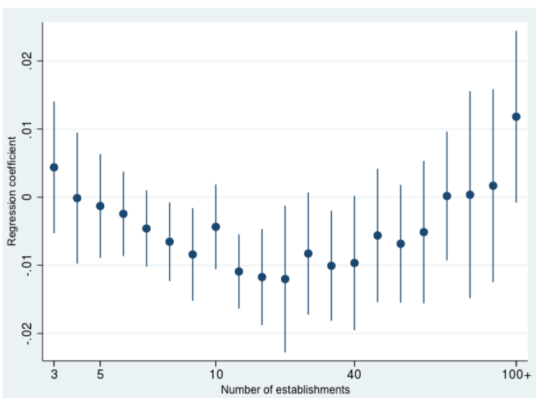


(b) Log sales per worker, all firms

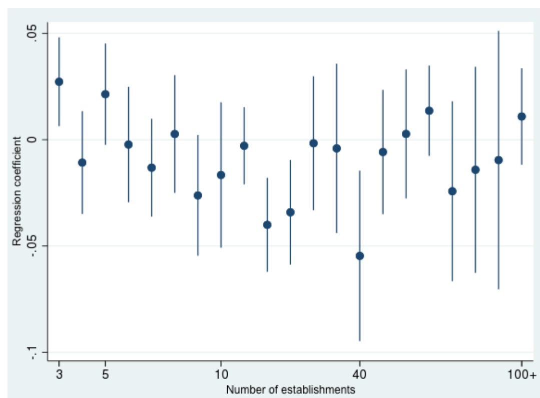


(d) Log sales per worker, manufactures

*Note*: Plotted points are firm-category specific slopes describing the relationship between establishments' log employment (top) or sales per worker (bottom) and log distance to firm centroid by firm size category as described by the specification in equation 2, which includes controls for a firm by (4-digit establishment) industry by year fixed effects as well as establishment age category fixed effects. Data include establishment-year observations for establishments across five Economic Census waves between 1992-2012. All standard errors are clustered by establishment 4-digit industry.

from the firm center tend to have lower employment, with the effect attenuating or disappearing for the largest firms. In contrast, we see a weak and inconsistent relationship between distance and plant-level sales per worker along the size class distribution, with a shallow U-shape for the sample with all firms and no clear pattern for manufacturing firms. These results imply that, for small and medium sized firms, plants that are closer to the firm center have higher employment and similar sales per worker to plants that are further away. If we put aside the additional measurement issues raised by attributing inputs and outputs across plants within the firm, this implies that the closer plants are on average more profitable and hence, under some additional assumptions, more productive.[11]

Appendix B discusses several robustness checks we perform on these specifications. Our results are robust to including finer 6-digit industry controls interacted with firm fixed effects, geographic controls for establishments, to measuring distance as the establishment's distance to headquarters and the firm's oldest establishment rather than firm centroid, to measuring distance to the firm-sector center, to restricting the sample to establishments with five or more employees, and to alternative establishment outcome measures including sales, value added per worker, and average wage.

## 4 Concluding Discussion

Taken together, these facts are consistent with the hypothesis that establishment location decisions are significantly influenced by costs of geographic dispersion that induce firms to agglomerate in space. These costs are generally thought to be supply side frictions in moving goods, people or ideas across space (via trade, travel and communication costs) that impede coordination between far-flung production units (Giroud, 2013; Kalnins & Lafontaine, 2013; Eichholtz et al., 2015; Alcacer & Delgado, 2016; Charnoz et al., 2018; Atalay et al., 2017). While there are other explanations that may contribute to the empirical patterns, such as spatial correlation in idiosyncratic supply and demand shocks at the firm level, the most natural explanation is that these supply frictions are substantial. Our findings suggest that these forces have pervasive effects on plant loca-

[11] See Appendix C for a formal discussion.

tion choices across most or all sectors of the economy.

However, these costs appear to be significant influences on small and medium-sized firms only. Larger and more productive firms appear to be less affected, either because the costs themselves are lower for larger firms or because the benefits of dispersion are correspondingly greater for those firms. It could be that larger, more productive firms invest in lowering the costs of dispersion, or they might be "born" with lower costs which in turn contribute to their measured productivity. Differentiating between these explanations using purely empirical methods would require detailed data on the inner workings of the firm over a substantial period of time (to capture long-run relocation effects) and quasi-experimental variation in firm productivity and coordination frictions, neither of which are available in our data.

We have shown that costs of dispersion play a quantitatively significant role on the location choices of small and medium sized firms, which raises questions about their nature and size. Both the costs and the benefits of dispersion might be small, or they might both be large. Costs could be "constants of nature" that are largely invariant to policies such as regulations and infrastructure, or they might be quite responsive to the right policy changes. Figure 3 shows that the cross-sectional patterns of firm dispersion have changed modestly since 1976, especially for the smallest firms. Along with the evidence from changes in travel costs in Giroud (2013) and Charnoz et al. (2018) this suggests that the costs are at least somewhat mutable. If such costs are both large and responsive to policy, there is scope for policies aimed at reducing these costs to generate significant welfare gains. It therefore seems important for future research to explore the size of the costs and what policy levers they might respond to.

While beyond the scope of this paper, a natural direction for future research is to estimate a structural model of plant location decisions in general equilibrium. This exercise could be useful for learning both about the mechanisms underlying our findings as well as the aggregate implications for the location of economic activity and welfare. However, despite recent developments (see e.g. (Jia, 2008; Tintelnot, 2017; Arkolakis & Eckert, 2017; Antras et al., 2017)), modeling multiple discrete choices with interdependence and substitutability across many choices remains very challenging. Finding the structure that provides the right balance between tractability and flexibility will be the key to making progress on this problem.

## References


J. Alcacer & M. Delgado (2016). 'Spatial organization of firms and location choices through the value chain'. *Management Science* **62**(11):3213–3234.

L. Alfaro & M. X. Chen (2018). '12. transportation cost and the geography of foreign'. *Handbook of International Trade and Transportation* p. 369.

P. Antras, T. C. Fort, & F. Tintelnot (2017). 'The margins of global sourcing: Theory and evidence from us firms'. *American Economic Review* **107**(9):2514–64.

P. Antràs & S. R. Yeaple (2014). 'Multinational firms and the structure of international trade'. In *Handbook of international economics*, vol. 4, pp. 55–130. Elsevier.

C. Arkolakis & F. Eckert (2017). 'Combinatorial discrete choice'. Tech. rep., Yale University.

E. Atalay, et al. (2017). 'How wide is the firm border?'. Tech. rep., National Bureau of Economic Research.

D. Bahar (2014). 'Heavier than air? knowledge transmission within the multinational firm'. Tech. rep., mimeo, Harvard University.

K. Behrens & V. Sharunova (2015). 'Inter-and intra-firm linkages: Evidence from microgeographic location patterns'. *CEPR Discussion Papers* (10921).

P. Charnoz, C. Lelarge, & C. Trevien (2018). 'Communication costs and the internal organisation of multi-plant businesses: Evidence from the impact of the french high-speed rail'. *The Economic Journal* **128**(610):949–994.

G. Duranton & H. G. Overman (2005). 'Testing for localization using micro-geographic data'. *The Review of Economic Studies* **72**(4):1077–1106.

P. Eichholtz, R. Holtermans, & E. Yönder (2015). 'The economic effects of owner distance and local property management in us office markets'. *Journal of Economic Geography* **16**(4):781–803.

T. C. Fort & S. D. Klimek (2016). 'The effects of industry classification changes on us employment composition'. Tech. rep., Tuck School at Dartmouth.

X. Giroud (2013). 'Proximity and investment: Evidence from plant-level data'. *The Quarterly Journal of Economics* **128**(2):861–915.

A. Gumpert (2018). 'The organization of knowledge in multinational firms'. *Journal of the European Economic Association* **16**(6):1929–1976.

J. V. Henderson & Y. Ono (2008). 'Where do manufacturing firms locate their headquarters?'. *Journal of Urban Economics* **63**(2):431–450.

T. J. Holmes (2011). 'The diffusion of wal-mart and economies of density'. *Econometrica* **79**(1):253–302.

A. Irarrazabal, A. Moxnes, & L. D. Opromolla (2013). 'The margins of multinational production and the role of intrafirm trade'. *Journal of Political Economy* **121**(1):74–126.

P. Jia (2008). 'What happens when wal-mart comes to town: An empirical analysis of the discount retailing industry'. *Econometrica* **76**(6):1263–1316.

A. Kalnins & F. Lafontaine (2013). 'Too far away? the effect of distance to headquarters on business establishment performance'. *American Economic Journal: Microeconomics* **5**(3):157–79.

W. Keller & S. R. Yeaple (2013). 'The gravity of knowledge'. *American Economic Review* **103**(4):1414–44.

F. Tintelnot (2017). 'Global production with export platforms'. *The Quarterly Journal of Economics* **132**(1):157–209.

E. Tomiura (2007). 'Foreign outsourcing, exporting, and fdi: A productivity comparison at the firm level'. *Journal of International Economics* **72**(1):113–127.

S. R. Yeaple (2009). 'Firm heterogeneity and the structure of us multinational activity'. *Journal of International Economics* **78**(2):206–215.

# A Additional Results: Not For Publication

In this appendix, we report some additional statistics and results.

## A.1 Description of Samples

Our first fact uses the average distance between establishments in each firm and the firm's centroid as a measure of the firm's geographic dispersion, and compares this measure in the data to that of a synthetic control constructed as described in Section 3.

Panel A of Appendix Table A1 summarizes these measures for the 650,000 firm-by-year observations in this sample as well as the 49,000 firm-year observations in the sub-sample of manufacturing firms. Panel A also reports, for both all and the sub-sample of manufacturing firms, firm log employment and firm log output per worker, used for our third fact.

Our second fact uses a subset of the firms in Fact 1 which can be observed moving up exactly one firm size category between Economic Censuses. In our main sample this is just 44,000 firm-year observations from the full sample and 4,900 firm-year observations from the manufacturing firm sub-sample, or around 10% of each initial sample. Panel A reports changes in real and synthetic firm data for these sub-samples.

Our fourth and final fact examines relationships between establishments' log miles to the firm center and establishment log employment or sales per worker. Summary statistics for the 6,413,000 establishment-year observations in the full sample and 270,000 establishment-year observations for the sub-sample of manufacturing establishments are reported in Panel B.

Table A1: Summary statistics

Panel A: Firm-level variables

| Variable | Obs | Mean | St. Dev. |
|---|---|---|---|
| | All firms | | |
| Number of establishments | 650,000 | 9.38 | 86.21 |
| Log average miles to centroid | 650,000 | 2.43 | 3.60 |
| Log average miles to centroid, synthetic | 650,000 | 6.51 | 6.77 |
| Log employment | 650,000 | 2.83 | 1.09 |
| Log sales per worker | 650,000 | 2.43 | 3.60 |
| Change in log average miles to centroid | 44,000 | 0.60 | 2.35 |
| Change in log average miles to centroid, synthetic | 44,000 | 3.27 | 3.70 |
| | Manufactures | | |
| Number of establishments | 49,000 | 8.51 | 42.13 |
| Log average miles to centroid | 49,000 | 3.65 | 4.08 |
| Log average miles to centroid, synthetic | 49,000 | 6.52 | 3.65 |
| Log employment | 49,000 | 3.98 | 1.06 |
| Log sales per worker | 49,000 | 3.65 | 4.08 |
| Change in log average miles to centroid | 4,900 | 0.90 | 3.14 |
| Change in log average miles to centroid, synthetic | 4,900 | 2.47 | 4.32 |

Panel B: Establishment-level variables

| Variable | Obs | Mean | St. Dev. |
|---|---|---|---|
| | All establishments | | |
| Log miles to firm centroid | 6,413,000 | 4.50 | 2.58 |
| Log employment | 6,413,000 | 2.62 | 1.25 |
| Log sales per worker | 6,413,000 | 4.80 | 1.13 |
| | Manufactures | | |
| Log miles to firm centroid | 270,000 | 4.47 | 3.10 |
| Log employment | 270,000 | 4.05 | 1.43 |
| Log sales per worker | 270,000 | 5.32 | 0.96 |

### A.2 Duranton Overman vs Distance to Centroid

Here we discuss how our measure of within-firm dispersion relates to a more common measure of agglomeration (and dispersion), the mean of the CDF of bilateral plant distances as proposed by discussed in Duranton & Overman (2005).

In the context of the firm, the Duranton-Overman measure is the log of the mean of the distance between any two plant pairs in the firm. For a firm $f$ with $N$ constituent plants $p_1, p_2, p_N$, for all pairs $p_i, p_j, i \neq j$, and a distance measured in miles of $d_{ij}$ between the plants, the measure is:

$$\frac{\sum_{i=1}^{N} \sum_{j=1}^{N} \mathbf{1}_{i \neq j} 2 \cdot d_{ij}}{N \cdot (N-1)}.$$

Our measure is:

$$\frac{\sum_{i=1}^{N} d_{ic}}{N}$$

where $d_{ic}$ is the distance between plant $i$ and firm centroid $c = (c_x, c_y)$, which we approximate as the average latitude and longitude of all the establishments in the firm

$$c_x = \frac{\sum_{j \in J} \mathbf{x}_j}{J},\ c_y = \frac{\sum_{j \in J} \mathbf{y}_j}{J}$$

where $\mathbf{x}_j$ and $\mathbf{y}_j$ are coordinates of the $j$th point.

Both measures are multilateral: a single plant's 'closeness' is a product of its relationship to all other plants. In a single dimension, these measures are identical, however because the average distance to other establishments' locations is not identical to the distance to the average location of the plants in two dimensions, these measures will be slightly different. However, as we show below in Appendix Figure 1a, the two measures are extremely highly correlated.

In this figure, we generate firms with a variable number of establishments drawn from zipcodes using the 2012 Zipcode Business Patterns data. We then calculate our firm dispersion measure and the Duranton Overman measure for these firms and plot both measures. Below is a binned scatterplot of the two measures. While the measure are not identical, the relationship between them is tight and close to, but not precisely, linear. The corresponding univariate regression has an Adjusted $R^2$ value of .946. We conclude that both measures capture a similar notion of overall dispersion.

**Appendix Figure 1a: Duranton Overman vs Distance to Centroid**

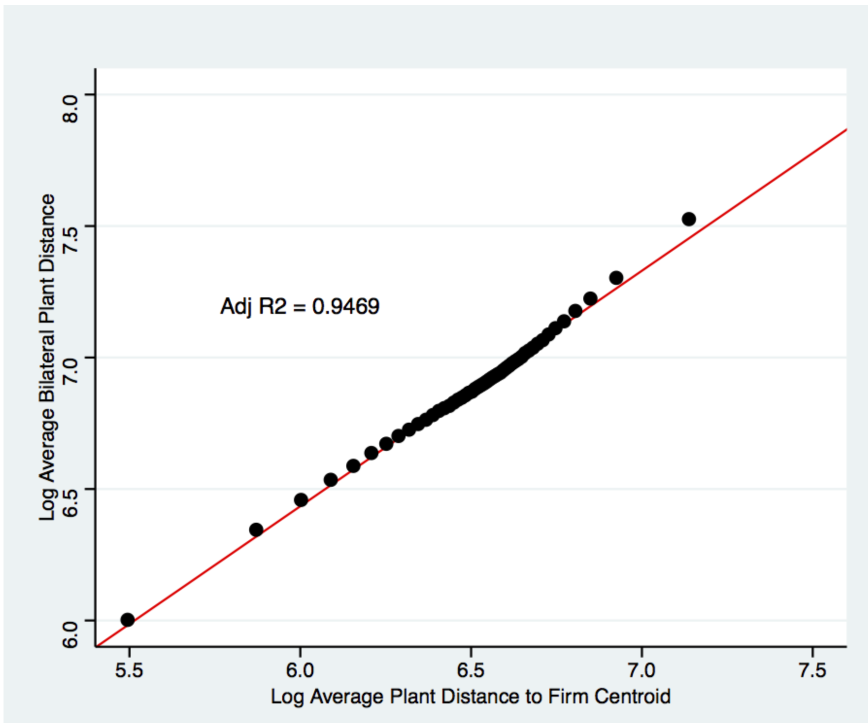


(a)

Note. Y-axis is log mean bilateral distance between establishments and X-axis is log mean distance to firm centroid for firms constructed from 2012 Zipcode Business Pattern data. Adjusted $R^2$ reported for univariate regression of the two measures.

## B Robustness Checks: Not For Publication

Here we report robustness checks.

### B.1 Fact 1

**Matching firms on inputs and products** A key assumption in Figure 1 is that for each firm, the set of establishments selected for the synthetic baseline are identical to the firm's actual establishments in all but location. This assumption may be violated if there are systematic differences between establishments within the same 6-digit industry. For instance, establishments in the same industry code may sell the same product but differ in their production processes. Or they may have the same modal product but one may be a multi-product firm while the other produces only a single product. If these differences affect the establishment's location, the null will reflect any corresponding geographic differences in production in addition to the underlying cross-firm dispersion within industries.

To address this, we use the product and material files in the Census of Manufactures to match establishments within the same 6-digit industry based on their entire set of products sold and, separately, the entire set of inputs used in production. To perform such a match, we first have to measure the similarity between establishments based on the entire set of their product sales and inputs used. However, there is no obvious way to compare the similarity of two distinct inputs used or products sold. Our approach is to consider each input, and separately, each product, as a distinct dimension. Of $I$ possible inputs and $P$ possible products, each plant $i$ can then be described by an $I$ and $P$ dimensional vector of inputs $\vec{T_i} = [t_1, t_2, ..., t_I]$ and outputs $\vec{S_i} = [t_1, t_2, ..., t_O]$, respectively.

To measure the similarity of the two vectors of inputs, we use the angle of similarity defined as

$$sim_{T_i,T_j} = \frac{\vec{T_i} \cdot \vec{T_j}}{\|\vec{T_i}\|\|\vec{T_j}\|}$$

We then use $sim_{T_i,T_j}$ as a sampling weight. For each plant $i$, we draw a matched plant from the same 6-digit industry with $N$ plants such that the probability of drawing any particular plant $j$ as plant $i$'s match is

$$Pr(i,j) = \frac{sim_{T_i,T_j}}{\sum_{k=1}^{N} sim_{T_i,T_k}}$$

Using the matched plants, we recompute the firm distance to centroid as our new synthetic baseline. We repeat this for the vector of outputs $\vec{S_i} = [t_1, t_2, ..., t_O]$ and report the results below in Appendix Figure 2a.

Synthetic baseline firms constructed using this approach do show significantly less dispersion across the board. However, the qualitative results of Fact 1 are unchanged: firms on average are clustered and smaller firms are significantly more clustered relative to their baselines.

**Robustness to industry controls.** Figure 1 compares, within each size bin, the observed dispersion to the synthetically constructed null. By construction, the null accounts for the characteristics of the firm, including industry composition. However, the cross-group pattern of decreasing dispersion in unconditional means may in part reflect compositional differences in the groups' characteristics. We cannot simultaneously observe the unconditional difference in the means between the observed and synthetic data and control for these differences. So, without a synthetic baseline, we can estimate

$$y_{jkt} = \alpha_{kt} + \beta_Z \cdot \mathbf{1}_{estgroup=Z} + \alpha_a \cdot \mathbf{1}_{agegroup=a} + \epsilon_{jkt}$$

where $y_{jkt}$ is the dispersion measure of firm $j$ in industry $k$ at time $t$, $\alpha_{kt}$ is an industry-year fixed effect, $\beta_Z$ is the conditional mean of firms for each establishment group in the set of groups $Z$, and $\alpha_a$ is a fixed effect for firm age group. Our results are robust to this specification.

**Additional robustness checks** In order to ensure our results are not driven by our choice of dispersion measure, we use two additional measures of firm dispersion: the average distance to the oldest surviving establishment in the firm as well as the average distance to the firm's headquarters. These measures may be inherently interesting to readers who believe distance costs are specifically attributable to distance-related management costs, for example.

**Appendix Figure 2a: Synthetic controls using angle of product and input similarity**

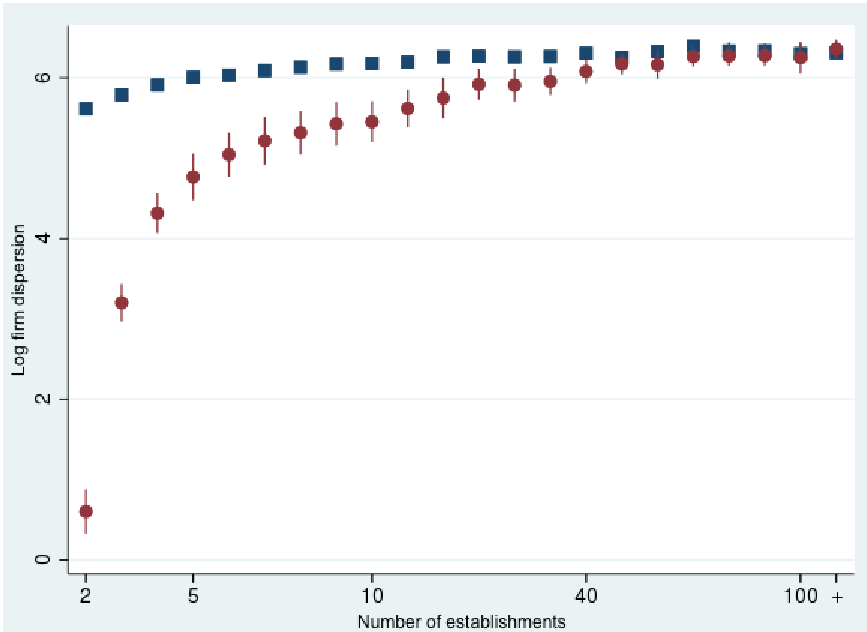


(a) Product mix

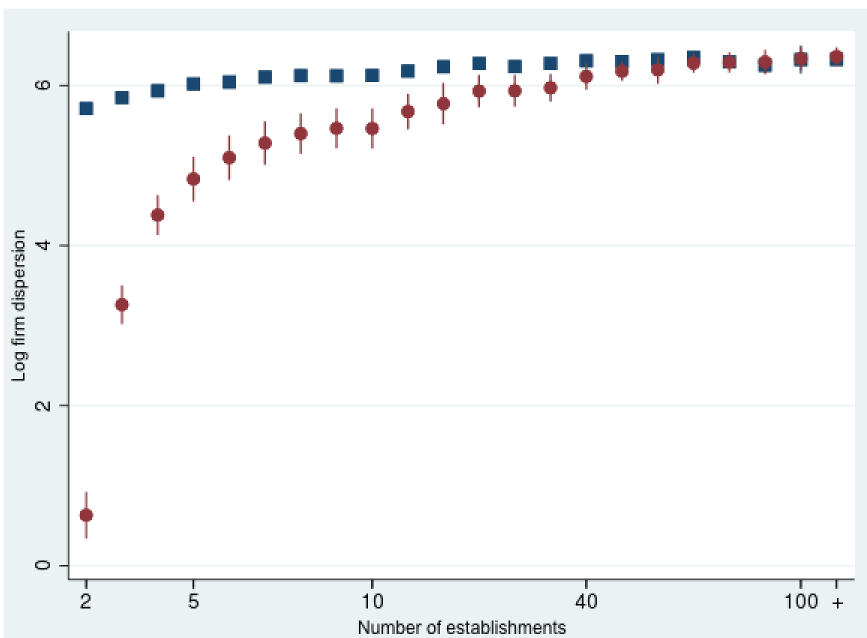


(b) Input mix

*Note*: Red circles plot the average log mean establishment distance from firm centroid for each firm size category. Estimates include all firm-year observations across five Economic Census waves between 1992-2012 for which product data (top panel) and input data (bottom panel) is available. Blue squares plot the corresponding measure for synthetically constructed firms where existing firms' plants are matched to and replaced with new plants based on plant output and input mix, respectively. Standard errors are clustered by firm modal 4-digit industry.

We recalculate observed firm dispersion using both distance to oldest establishment and distance to headquarters. For the latter, only about 42,000 firm-year observations across five Economic Census waves report unique firm headquarters.

To construct nulls, we randomly assign one of the synthetic firm's randomly selected establishments to that of a headquarters or oldest establishment and computer average distance to that establishment. For that reason, these two nulls will be identical. These do not meaningfully change our results.

Lastly, in order to rule out the possibility that our results are driven by extremely small establishments, we drop firms with five of fewer establishments. None of our results are affected by this change to the sample selection criteria.

### B.2 Fact 2

**Alternative synthetic baselines** To form the synthetic control, we match new establishments with randomly selected establishments in the same 6-digit industry. Changes in the spatial distribution of firms over time, which are documented in Figure 3, may mean we are selecting here from a systematically different set of establishments: those born within our time frame and those born before it. Ideally, we would address this by selecting establishments for the null that are both of the same age and industry. However, these constraints quickly reduce the full set of establishments from which to choose. An intermediate approach matches new establishments with those in the same industry that are born within a specific window. For instance, we can match new establishments with establishments that are five of fewer years of age. Our results are robust to such restrictions.

Our results are also robust to different treatment of firm growth that is the result of merger activity. In our main results, we do not differentiate between newly born establishments and establishments that enter the firm as the result of a merger. Our synthetic baseline randomizes the location of both of these groups of new establishments. Alternatively, we can run our specification randomizing only new births. When computed in this method, the synthetic control is closer to the real data but similarly preserves economically and statistically meaningful growth above what is observed in real data for the same size categories as in our main specification.

Results using manufacturing establishments matched on input and output usage by firm size category do not pass disclosure avoidance requirements and cannot be reported.

**Transitions down and up by more than one category** Our main specification looks at firms that grow, moving up one size category. Alternatively, we can look at firms that move down one size category. Results are qualitatively similar for these firms.

Of course, firms may move more than one size category within 5-year windows. We can rerun our specification for the full transition matrix. Individual cells with few observations lose significance, but observed growth in dispersion is generally lower than synthetic baseline growth, and the difference is statistically significant for all growing firms for any transitions to under 25 establishments, as well as for transitions by up to three groups and to 10 establishments for manufacturing firms.

**Additional robustness checks** As with Fact 1, Fact 2 is robust to dropping plants with fewer than five employees, adding industry and age group controls, measuring distance to the oldest establishment and headquarters, and measuring within firm-sectors.

### B.3 Fact 3

One concern may be that there is geographic variation in the size of firms (both in the average number of establishments and their dispersion) that is correlated with outcomes through the firm's location in space, which is omitted. To address this, we control for the firm's general location. Because firms' footprints overlap with traditional geographic boundaries, there is no perfect way to assign firms to locations. Two ways to do this are to control for cubic polynomials of the the firm centroid latitude and longitude and to assign firms to Census regions based on their centroids and control for region fixed effects.

Our results are robust to both of these controls, as well as to dropping plants with fewer than five employees and to measuring geography and outcomes at the firm-sector level and measuring distance to the firm's oldest establishment or the firm's headquarters, as well as to using alternative firm outcome measures such as sales, value added per worker, or average wage.

### B.4 Fact 4

As with Fact 3, establishment locations may be correlated both with outcomes and distance to firm centroid. Here, controlling for establishment location is more obviously done by including MSA or county fixed-effects. Our results are robust to these additions.

Our goal is to compare the outcomes of establishments undertaking the same activity within the firm at different distances from the firm center. One issue is that establishments in the same 4-digit industry may be undertaking different activities within the same industry. To attempt to address this, we include finer 6-digit industry-by-firm controls. Our results are unaffected by these controls.

As above, different definitions of internal firm distance and definitions of the relevant firm boundary may impact results. Our results are robust to using distance to the oldest establishment in the firm and distance to the firm's headquarter, to measuring distance to firm-sector centroid or oldest establishment in the firm-sector, to dropping establishments with fewer than five employees, as well as to using alternative establishment outcome measures such as sales, value added per worker, or average wage.

## C Productivity comparisons across firms and plants: Not For Publication

In this section we provide a simple example with CES demand and a general cost function to establish the firm-level relationship between total sales, sales per worker and a notion of productivity which encompasses the effects of both supply and demand shocks.

Consider a set of firms $i, j \in J$ that hire factors $\mathbf{z}$ on competitive markets and choose output quantity $q$ in order to maximize profits. Production and demand functions across firms differ only by multiplicative constants, i.e.

$$p_i(q_i) = a_i q_i^{-\frac{1}{\sigma}}, \;\; q_i = b_i f(\mathbf{z}_i), \;\; \forall i \in J, \tag{3}$$

with $f$ differentiable and $\sigma > 1$.[12] We assume that firms face common factor prices $\mathbf{w}$. Under these assumptions, we have

$$c_i(\mathbf{w}, q_i/b_i) = c_j(\mathbf{w}, q_j/b_j) = c(\mathbf{w}, \tilde{q}), \;\; \forall \mathbf{w}, \; i, j \in J \tag{4}$$

for each firm's cost function. That is, each firm faces the exact same cost of "producing" $\tilde{q}$, although each firm has a different mapping from $\tilde{q}$ to actual quantity produced.

Letting $\ell$ be the quantity of labor hired, we further assume that

$$\frac{\partial c(\mathbf{w}, \tilde{q})}{\partial w_\ell} > 0, \;\; \forall \mathbf{w}, \tilde{q}. \tag{5}$$

By Shepard's lemma, this condition implies that each firm's conditional labor demand is strictly increasing, i.e. whenever the firm wants to increase quantity produced, its labor input strictly increases. This allows us to associate each "quantity" $\tilde{q}$ with a unique choice of labor input $\ell^*$, so we can write $\tilde{q}(\ell^*)$, with $\tilde{q}'(\ell^*) > 0$.

Using the notation above, firm $i$'s problem can be written as

$$\max_{\ell_i^*} \frac{\mu_i \cdot \tilde{q}(\ell_i^*)^{1-\frac{1}{\sigma}}}{\ell_i^*} \cdot \ell_i^* - c(\tilde{q}(\ell_i^*)). \tag{6}$$

where we have suppressed the dependence of both $c$ and $\ell^*$ on the common factor prices. Notice that only revenue per worker differs across firms, due to both supply and demand shocks, whose effects can be summarized by a single multiplicative parameter $\mu_i = a_i b_i^{1-\frac{1}{\sigma}}$ due to the assumption of CES demand. Firm productivity is then defined as the total Marshallian surplus generated by the firm, or

$$Prod_i = \int_o^{\ell_i^*} \left[ \frac{\mu_i \cdot \tilde{q}(\ell)^{1-\frac{1}{\sigma}}}{\ell} - c'(\tilde{q}(\ell))\tilde{q}'(\ell) \right] \, d\ell. \tag{7}$$

We are now ready to examine the relationship between total sales, sales per worker, and productivity. Using the Envelope Theorem, it is easy to show that firms with higher $\mu$ can have higher or lower revenue per worker, depending on the shape of $c(\tilde{q}(\ell^*))$.

[12] This analysis applies to multi-plant firms so long as we can aggregate the plant outputs into a quantity index with an associated price index.

For example, with constant marginal cost CES demand implies that all firms have the same revenue per worker regardless of the value of productivity $\mu_i$. Thus, in this model, sales per worker is not very informative regarding firm productivity. Either total sales or total employment is a better indicator of firm productivity, since $\partial \ell_i^* / \partial \mu_i > 0$ and $\partial Prod_i / \partial \mu_i > 0$.

### C.1 Departures from CES demand

We maintain the previous assumptions, but relax CES demand. In this setting, a firm with both (weakly) higher employment and (weakly) higher sales per worker (and hence weakly higher sales) must be (weakly) more profitable. Furthermore, neither sales per worker nor employment is sufficient to establish relative profitability on its own, so the two statistics together yield an increase in information. However, moving from firm profitability to Marshallian surplus cannot generally be done. While relative profitability can be inferred by comparing the height of the revenue per worker curve for a single employment level (the smaller firm's optimal employment), the relative consumer surplus depends on the global behavior of these curves as well as the marginal cost curve.[13] With both supply and demand shocks, this behavior can be complex. To the extent that relative profitability is informative about relative productivity, both sales per worker and employment together are better indicators of productivity than either alone.

### C.2 Comparing plants within firms

The results above apply equally well to comparisons between plants in the same firm. Marginal revenue products of labor must be equalized across plants, but average revenue products need not be. Larger plants may have higher or lower average revenue products of labor, depending on the shape of the cost function. Under CES demand, total sales or employment is positively related to productivity and profitability.

Within-firm comparisons face the additional challenge of the possibility of joint production and/or shared inputs across plants. To the extent that this measurement error is random, it will tend to bias toward a null finding. If geographically more central plants

[13]CES demand allows us to infer global behavior from relative height at a single point, which makes it convenient for this problem.

tend to provide more shared inputs (e.g. headquarter services), that will tend to mechanically produce a negative relationship between centrality and sales per worker even in the absence of true productivity differences. Thus, even more caution is required in interpreting differences in sales per worker across plants than across firms.