

NBER WORKING PAPER SERIES

PRODUCTIVITY, PLACE, AND PLANTS

Benjamin Schoefer
Oren Ziv

Working Paper 28772
<http://www.nber.org/papers/w28772>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2021

We thank Luke Watson for excellent research assistance. We thank Dominick Bartelme, Gabriel Chodorow-Reich, Jonathan Dingel, Cecile Gaubert, Steven Haider, Andrei Levchenko, Martin Rotemberg, Juan Carlos Suárez Serrato, Yue Yu, and Ben Zou for useful discussions. We thank participants of the NBER Summer Institute Urban Economics Meeting, the Third International Workshop “Market studies and spatial economics,” the 10th European Meeting of the Urban Economics Association, and at Michigan State University for feedback. For financial support, we thank the UC Berkeley Clausen Center for International Business and Policy, and the Ewing Marion Kauffman Foundation. The paper draws on confidential data from the US Census Bureau, which requires the following disclaimer: “Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the US Census Bureau. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 1890. All results have been reviewed to ensure that no confidential information is disclosed.” The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by Benjamin Schoefer and Oren Ziv. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Productivity, Place, and Plants
Benjamin Schoefer and Oren Ziv
NBER Working Paper No. 28772
May 2021
JEL No. D22,D24,E23,R0,R12

ABSTRACT

Why do cities differ so much in productivity? Using a split-sample IV strategy, we document that up to three quarters of the large measured dispersion in productivity across US cities is spurious and reflects the "luck of the draw" of idiosyncratically heterogeneous plants. Due to this granularity bias, economies with randomly reallocated plants exhibit nearly as high a variance as the empirical economy. For new plants, four fifths of the raw dispersion reflects granularity bias, and their productivity is only imperfectly correlated with that of older plants, which dominate measured productivity levels. These US-based patterns broadly extend to European countries.

Benjamin Schoefer
Department of Economics
University of California, Berkeley
530 Evans Hall #3880
Berkeley, CA 94720-3880
and NBER
schoefer@berkeley.edu

Oren Ziv
Michigan State University
Department of Economics
orenziv@msu.edu

1 Introduction

The sources and consequences of spatial differences in productivity are a frequent focus in urban economics, starting with the state-level analysis of Ciccone and Hall (1996). Spatial differences in productivity contribute to regional dispersion in wages and employment (Caliendo, Parro, Rossi-Hansberg, and Sarte, 2018; Hornbeck and Moretti, 2018), drive location choices of plants (Ellison, Glaeser, and Kerr, 2010; Gaubert, 2018) and of people and hence the city size distribution (Desmet and Rossi-Hansberg, 2013). They are also testable implications of core urban economics models (e.g., Gaubert, 2018; Davis and Dingel, 2020). Furthermore, these differences motivate studies of spatial misallocation and the design of spatial policies (Moretti, 2012; Boeri, Ichino, Moretti, and Posch, 2019; Fajgelbaum, Morales, Suárez Serrato, and Zidar, 2019; Hsieh and Moretti, 2019; Fajgelbaum and Gaubert, 2020). Often, the underlying assumption of this literature is the presence of total factor productivity (TFP) differences in place-specific production functions, which shift productivity of all local plants, incumbent and new.

Yet, due to measurement challenges, existing research does not typically assert and dissect the raw empirical dispersion of city-specific productivity levels directly. Instead, productivity differences are indirectly measured on the basis of observable symptoms such as city size (Gabaix, 1999b,a), city growth (Glaeser, Kallal, Scheinkman, and Shleifer, 1992; Henderson, 1994), rent levels (Dekle and Eaton, 1999), or wage levels (Glaeser and Mare, 2001; Hsieh and Moretti, 2019; Ehrlich and Overman, 2020). Other research studies city-level productivity directly, but does so either as a dependent variable (Sveikauskas, 1975; Combes, Duranton, Gobillon, and Roux, 2010; Combes, Duranton, Gobillon, Puga, and Roux, 2012), or, if as an explanatory variable, uses instruments for productivity growth rather than raw productivity levels (Hornbeck and Moretti, 2018).

We provide a direct measure of the spatial dispersion in productivity in the United States and 15 European countries. Besides providing the raw dispersion measures, we estimate the dispersion of *true place effects*, which we define as the *expected* productivity of a randomly drawn plant from a potentially place-specific distribution. This statistical definition of place effects is agnostic to their sources, encompassing exogenous or endogenous causal effects, sorting—the economic sources that are the object of much research in urban economics—or even spatially correlated mismeasurement.

In our framework, true place effects stand in contrast to statistical differences in observed *averages* of productivity of finite populations of plants, which are finite draws from the latent productivity distribution of each location’s infinite superpopulation of plants. Of course, the resulting raw, prevailing productivity conditions of the actually existing plants

are real in that they matter for core outcomes. They drive, for instance, the dispersion in wage levels (see, e.g., Figure 1 plotting region averages of TFP and wages against density in Combes et al., 2010), and may hence be revealed in the observed wage or rent levels (Albouy, 2016). Moreover, even randomly placed plants of course have real effects on local economies (Greenstone, Hornbeck, and Moretti, 2010). Finally, raw productivity averages may also indicate the degree of spatial misallocation (Hsieh and Moretti, 2019).

However, rather than true place effects, these commonly measured productivity averages may simply reflect the “luck of the draw” of finite counts of plants, tremendously heterogeneous in idiosyncratic productivity and size, that happen to be located in a given place. That is, our setting is *granular* in the sense that a given plant need not wash out in the average.¹ Hence, the average productivity level in a given location indeed features considerable sampling variability—which enters the raw variance of place averages as a term we call *granularity bias*. In the extreme case, all of the observed spatial differences would reflect granularity bias, which would leave no room for the other economic sources that have been heretofore marshalled to explain this variation. In the data, we find that granularity bias accounts for two thirds to four fifths of the spatial raw differences in productivity—accordingly shifting the goalpost for what economic sources captured by true place effects should explain to one fifth to one third of the variance of standard average-based productivity measures.

Indeed, we document large dispersion in the raw variance of average productivity across US cities (metropolitan statistical areas, MSAs) demeaned within the national 4-digit industry (which nets out mechanical productivity differences working through industry composition). Our main measure of plant-level productivity is log of revenue total factor productivity (TFPr), although we also study log value added per worker. Our headline number of the raw variance of average-based location effects is 0.026—the dissection of which is the focus of the paper. That is, a city one standard deviation above the mean has about a 16 log points (17%) higher average productivity. Manufacturing plants in the top 10% most productive MSAs are on average 48 log points (61%) more productive than plants in the bottom 10%. While MSA-level results involve averages across many plants (the average plant count per MSA is around 300 in our sample), we also construct finer industry-specific location effects, which exhibit four times the variance, at 0.110.

We assess, and ultimately strip out, the scope of granularity bias in three comple-

¹See Syverson (2011) for a canonical overview of the degree and potential determinants of between-plant differences in productivity. We take this heterogeneity as given. Gaubert and Itskhoki (forthcoming) explore comparative advantage, of specific firms rather than all firms in a sector, in international trade. Dingel and Tintelnot (2020) study the impact of small and zero numbers of commuter flows on the estimation of general-equilibrium spatial models.

mentary strategies. First, we provide a nonparametric permutation test by calculating the probability that the observed level of spatial dispersion would emerge if plant-level productivity were independent of place. We simulate 1,000 US economies that randomly reassign the empirical plants across MSAs, preserving the empirical plant count in each MSA-industry cell. The resulting series of raw variances in the average-based place effects of these simulated economies makes up our nonparametric sampling distribution. The average simulated variance is 0.017. This dispersion is high, but one third less than the empirical number (0.026)—which is in the 99th percentile of the distribution of the randomized economies. For industry-specific location effects, the empirical variance of 0.110 is close to the permutation-based average variance of 0.096. But the empirical variance is in the 98th percentile of the distribution of permuted variances. Hence, while granularity can generate substantial *spurious* variation, the US economy does exhibit a statistically significant degree of dispersion in true place effects.²

Second, we constructively estimate the bias-corrected variance of true place effects. We employ a split-sample approach. In every MSA-industry cell, we randomly split plants into two equally sized groups. We then estimate average-based place effects for each of the two resulting economies. The covariance between the average-based place effects of these paired cities is an unbiased estimator of the variance of the true place effects, which, by definition, are the common components of both split samples. This bias-correction shrinks the variance substantially, from 0.110 to 0.023 for the industry-specific location effects, and even for the MSA-level location effects, from 0.026 to 0.008. In robustness checks, we find similar reductions in the raw variance when studying labor productivity (value added per worker) rather than TFP, as well as when using different industry definitions. In sum, when identifying the variance of true place effects, places are much more similar than the raw variance of average productivity levels suggests, and the majority of spatial differences in realized productivity reflect largely the luck of the draw from highly dispersed plant distributions.

Some immediate implications of the large role of granularity are that counterfactuals involving spatial reallocation of inputs from low to high TFP places (see, e.g., Desmet and Rossi-Hansberg, 2013; Hsieh and Moretti, 2019) would have lower effects on aggregate output, to the degree that new plant draws absorb the inputs (rather than through intensive-margin expansion of existing plant draws). Moreover, empirical correlations and variance decompositions accounting for spatial productivity differences between places are subject to granularity bias.

²Broadly, our test is in the dartboard spirit of Ellison and Glaeser (1997), who study whether the observed geographical concentration is statistically different from randomly located plants.

Third, we assess the role of three sources of granularity bias: idiosyncratic plant-level heterogeneity in productivity, finite plant counts, and plant size (as our baseline specification weights by plant employment). Regarding idiosyncratic variance, we winsorize plant-level TFP by 0% and 2.5% rather than our baseline of 1%, and find similar results. Regarding plant size, we also provide unweighted results, which imply smaller raw variances that nonetheless continue to overstate the variance of true place effects considerably. Regarding finite plant counts, we replicate our analysis but raise the minimum plant count for each location-industry cell incrementally from two (our baseline) to 150. While indeed the raw variance falls somewhat, it never reaches the bias-adjusted variance. In fact, requiring more populous cells throws out the baby with the bathwater: the number of places included in the sample falls more than proportionately with the raw variance. By design, this seeming solution to the granularity problem changes the sample, and, we find, does so towards a more homogeneous one with a lower bias-corrected variance—leaving the raw variance overstating the true place effects by a fairly stable amount. At a cell-size cutoff of 100, the remaining sample features only 10% of MSAs (and a much lower fraction of industry-MSA cells), fails the permutation test, and yields a bias-corrected variance of zero. Overall, therefore, granularity bias is a robust feature of the data, and alternative solutions that build on adjusting the sample and other dimensions of the data may defeat the purpose of quantifying the dispersion in the true and complete economy. Furthermore, our results persists even in our MSA-level analysis, where averages are taken over large numbers of plants (so that idiosyncratic dispersion and size skew contribute to granularity bias). Instead, our strategy reveals the true variance net of granularity that takes into account the, we find dominant, contribution from smaller—most granular—cells.

Hence, granularity bias accounts for two thirds of the raw variance of MSA effects, and fourth fifths of that of industry-specific MSA effects. The remaining dispersion in true place effects arises from at least three sources: causal effects of place, sorting among and selection of ex-ante differentiated plants, and remaining measurement error common to a place. Focusing on differences in expected values of productivity, we do not speak to spatial differences in other moments of the distribution of plant-level productivity. By focusing on differences in location averages and taking plant-level productivity measures as given, our approach also does not resolve the largely orthogonal but important issue of plant-level measurement of productivity (see, e.g., Haltiwanger, Kulick, and Syverson, 2018; Bils, Klenow, and Ruane, 2020; Rotemberg and White, 2020).

We close with three further applications of our method. First, we zoom in on the dispersion in place effects for new plants. We are motivated by a specific implication of our main findings, that productivity measures of existing plants need not be informative

for and carry over to the outcomes of new plants and projects in the location. Indeed, raw productivity averages measured from new plants (at most five years old) have markedly higher raw variances of 0.084 and 0.152 for location and industry-specific location effects. However, the bias-corrected variances drop by dramatically more, to 0.015 and 0.019, respectively. Additionally, we quantify systematic divergence of true place effects of new plants from those of incumbent, older plants, which dominate the raw measures. Using our split-sample method to instrument for the relationship between the true place effects of new and old plants, we find that places that are 10% more productive for old plants are only 6% more productive for new plants. This considerable divergence again implies that aggregate effects of counterfactual distributions of inputs across seemingly heterogeneous places are attenuated when absorbed at the extensive, plant margin. The imperfect relationship may also caution against using TFP levels implied by all plants as those affecting entrepreneurship, plant location choices, and investment decisions, or as a guide for place-based policies targeting those outcomes, in favor of measures estimated off new plants.

Second, we also study log value added per worker (labor productivity), which may be a less demanding productivity concept, which is the symptom of spatial misallocation of labor (Hsieh and Moretti, 2019), and which, mapping into the marginal product of labor, is the canonical labor demand factor determining wages. We find that granularity bias is if anything more pronounced for this productivity measure.

Third, we show that the role of granularity bias extends to the within-country regional dispersion of 15 European countries, drawing on Bureau van Dijk firm-level data.

Section 2 defines the conceptual and statistical framework of our analysis. Section 3 presents the data, and the sample and variable construction. Section 4 presents the main results, covering the United States, TFP, and all plants in a place. Section 5 presents the analysis of new plants. Section 6 presents the results for labor productivity. Section 7 presents the analysis for 15 European countries.

2 Statistical Basics: Place Effects Under Granularity

We start by defining true place effects, then clarify the pitfalls of estimating the dispersion of place effects, and present our strategies to overcome these measurement challenges.

2.1 Formal Framework

We formulate a statistical definition of true place effects of productivity—the expected value of productivity of a randomly drawn plant from a potentially place-specific distribution—and clarify how those relate to measured raw averages of finite counts of idiosyncratically heterogeneous plants.

Setting The economy is characterized by a set L of count N_L locations indexed by $l \in L$. Each location has a count of N^{P_l} plants, which are indexed by $p \in P_l$, where P_l denotes the population of location- l plants; we will also consider subsets $S_l \subseteq P_l$ of size N^{S_l} .

Plant p in location l has log total factor productivity $a_{pl} = \ln A_{pl}$ —but our derivations below would also apply to, e.g., average labor productivity and alternative productivity concepts. In fact, while we refer to a_{pl} as plant productivity, we also leave open the possibility that it also reflects mismeasurement.

Plants are heterogeneous in productivity in ways that potentially depend on their location l . We agnostically describe this property in the statistical form of a latent data-generating process $a \sim F_l^a(a)$, of plant-level productivity in a location l .

Importantly, this latent data generating process does not describe the given finite population of existing plants (as in the real economy, or our census data). In statistical terms, $F_l^a(a)$ describes the location’s infinite superpopulation of plants, from which the finite observed population is drawn. The draws from $F_l^a(a)$ give the plans that would, if drawn, be active (unlike the potential-entrants distribution that include non-viable plants as in plant selection models such as in Melitz, 2003), and hence the distribution may reflect a host of underlying economic forces—e.g., selection on entry and exit, and sorting.

In addition, a plant is characterized by size e_p (e.g., employment); the exposition below starts with unweighted (or equally sized) plants within locations, and no weighting across locations. We discuss the extension to weighted averages and heterogeneity in size in Section 2.6.

True Place Effects We define true place effects as *expected values* of plant-level productivity in location l :

$$\tau_l = \mathbb{E}[a_{pl} | l] = \int a dF_l(a). \quad (1)$$

This *statistical* definition of place effects is agnostic to and captures a variety of specific economic mechanisms that manifest themselves in the expected value of productivity of plants in a location. Causal effects of place on productivity, including from agglomeration

effects (including productivity spillovers), would affect (not necessarily exclusively) this expected value. Systematic sorting or collocation of plants into places by productivity would show up in this place effect. Location-specific mismeasurement of productivity (e.g., in the production functions, input and output prices, quantities or qualities—or the presence of multi-plant firms with their constituent plants having similar productivity) can be reflected in this place effect. By drawing on plants after plant location choices and entry and exit dynamics, it may also capture productivity-relevant spatial differences in these selection margins. Of course, our formulation of place effects as expected values would not sufficiently characterize any given specific model or mechanism (see Combes, Duranton, Gobillon, Puga, and Roux, 2012, for a discussion of potential effects on other moments); instead, other moments of the latent location-specific productivity distribution $F_l^a(a)$ may differ across places $l \in L$ due to the aforementioned factors.

Our goal is to characterize the dispersion of true place effects τ_l across locations $l \in L$. We now contrast the true place effect defined in Equation (1) with the average productivity level of a finite set of plants in a place.

Idiosyncratic Plant-Level Productivity Using our definition of place effects in Equation (1), we can rewrite a plant’s log productivity a_{pl} as the sum of the place effect τ_l and an idiosyncratic component u_{pl} :

$$a_{pl} = \tau_l + u_{pl}. \quad (2)$$

That is, plants’ idiosyncratic residuals u_{pl} within a location l are simply deviations around the true place effect, expected value τ_l , and the residuals too are hence drawn from a potentially location-specific distribution $u \sim F_l^u(u)$. Hence, their expected value is zero: $E[u_{pl}|l] = 0$. Moreover, $F_l^a(a) = F_l^u(a - \tau_l)$. Again, just as with place effects, this statistical definition of the idiosyncratic deviations u_{pl} takes no stance on their economic origins. Like plant productivity a_{pl} , idiosyncratic deviation u_{pl} may capture actual productivity differences between heterogeneous plants in a location, mean-reverting shocks (although we will study one cross section), or plant-specific measurement error (e.g., Bils, Klenow, and Ruane, 2020; Rotemberg and White, 2020).

Average-Based Place Effects Since $E[u_{pl}|p \in S_l] = 0$ for any random sample S_l , the average productivity $\widehat{\tau}_l^{S_l}$ of any given single and finite set of plants $p \in S_l$ is an unbiased and

consistent estimator of the true place effect of location l , the expected value τ_l :

$$\widehat{\tau}^{S_l} = \frac{1}{N^{S_l}} \sum_{p \in S_l} a_{pl}. \quad (3)$$

For the rest of the paper, unless we consider specific samples, we will consider the—finite—*population* of plants $S_l = P_l$ (rather than a sample $S_l \subsetneq P_l$ in each location), as with our census data. We then we denote the population average by $\widehat{\tau}_l$:

$$\widehat{\tau}_l = \frac{1}{N^{P_l}} \sum_{p \in P_l} a_{pl}. \quad (4)$$

We will also consider and label specific samples (split samples, old or new plants etc), for which we index place effects by a superscript X , i.e., τ_l^X . We will also denote cell-level plant counts by N_l (population) or N_l^X (for subsamples indexed by X).

However, even the population average defined in Equation (4) generally differs from our object of interest, namely the expected value $E[a_{pl}|l]$ defined in Equation (1). That expectation is taken over the latent data-generating process $F_l^a(a)$ of plant productivities, from which the real economy and the census data draw a finite set of plants. We next discuss the pitfalls of estimating the true place effects on the basis of average productivities taken over finite sets of plants, and then our identification strategy stripping out the associated measurement error, which we characterize as granularity bias below.

2.2 Pitfalls of Estimating Dispersion in Place Effects

While each place effect $\widehat{\tau}_l$ is estimated without bias, dispersion measures based on averages $\widehat{\tau}_l$ are upward-biased estimates of the dispersion of true place effects τ_l , for reasons we jointly label as "granularity." We formalize this bias with the example of our leading dispersion statistic, namely the variance.

Variations of Place Averages The following equation, written in the notation for the population of finite plants, clarifies the pitfalls of estimating the variance of place effects

τ_l on the basis of location averages $\widehat{\tau}_l^{S_l}$:

$$\overbrace{\text{Var}(\widehat{\tau}_l)}^{\text{Raw Var of Place Averages}} = \text{Var}\left(\frac{1}{N_l} \sum_{p \in P_l} [\tau_l + u_{pl}]\right) = \text{Var}\left(\tau_l + \frac{1}{N_l} \sum_{p \in P_l} u_{pl}\right) \quad (5)$$

$$= \underbrace{\text{Var}(\tau_l)}_{\text{Var of True Place Effects}} + \underbrace{\frac{1}{L} \sum_{l \in L} \frac{\sigma_l(u)^2}{N_l}}_{\substack{>0 \text{ if } N_l < \infty \wedge \sigma(u) > 0 \\ \text{Bias from Granularity:} \\ \text{Var of Sample Means}}} + \underbrace{2 \text{Cov}\left(\tau_l, \frac{1}{N_l} \sum_{p \in P_l} u_{pl}\right)}_{=0 \text{ Orthogonal by Construction}} \quad (6)$$

where $\sigma_l(u)^2$ is the variance of plant-specific deviations u_{pl} from place effect τ_l in location l . The l -index permits heteroskedasticity in the distribution of plant-level idiosyncratic productivities $F_l(u)$ across locations. The third term in the second line is zero because the idiosyncratic deviations from the expected value are orthogonal to the expected value.

In the second line, the first term is the variance of true place effects. The second term is a term that biases upward the raw variance of productivity averages as an estimate of the variance of true place effects due to the granular nature of plants in places. We call this term granularity bias, and characterize it below.

2.3 Granularity and its Sources

While the variance of the place averages $\widehat{\tau}_l$ is a consistent estimator of the variance of true place effects τ_l , with finite populations of plants within locations it is *biased upward* by the second term in the second line of Equation (6): the weighted average of within-location variances divided by location count of plants. This term reflects "granularity" in the sense that a given plant need not wash out in the average. It arises under finite cell counts of plants N_l combined with large idiosyncratic variance $\sigma_l(u)^2$ within a cell l , so it is present even when $S_l = P_l$ as the populations of existing plants are finite. Intuitively, these factors generate realized deviations of sample averages from expected values τ_l , raising the observed variance of place averages above the variance of true place effects. Below we dissect each source of granularity and discuss the potential empirical relevance of each.

Plant Counts per Cell First, to gauge the empirical range of cell sizes (plant counts) N_l , we plot the CDF of cell sizes from public-use data on manufacturing plants in the US County Business Patterns in Appendix Figure A.1. Panel (a) plots the CDF of plant counts for cells defined as MSAs (pooling all industries); Panel (b) does so for MSA-industry cells, at the level of the 86 4-digit NAICS manufacturing industries. While small cells are

an obvious issue in measuring variance across location-industries, for which 60% have fewer than 5 plants, there are more plant observations at the MSA level when pooling all industries. Our empirical implementation will exclude varying degrees of small cells from our analysis.

Idiosyncratic Dispersion Second, Equation (6) clarifies that even for larger N_l , granularity bias can be large if plants exhibit large idiosyncratic variance in at least some locations. In the national data, within-industry dispersion of productivity across manufacturing plants is indeed tremendous. Cunningham, Foster, Grim, Haltiwanger, Pablonia, Stewart, and Wolf (2020) report standard deviations of 0.460 and 0.684 for log TFP and log labor productivity (revenue per hour worked), respectively, covering 1997-2016 and the US manufacturing sector, using US Census data similar to ours, and studying 4-digit NAICS industries (with equal weights across industries and years). The 90th percentile of plants in a 4-digit NAICS industry are, nationally, 193% (1.078 in logs) or 490% (1.773 in logs) more productive than the 10th percentile, for TFP and labor productivity respectively.³ The important work by Syverson (2004) reports similar statistics for an earlier period.

Large, Dominant Plants A third source of granularity arises in the common specification in which plants are weighted by size such as employment, as we will do in our empirical implementation (although our theoretical exposition above is written in terms of unweighted (or equally sized) plants). Then, large plants can dominate plant averages. (We additionally discuss weighting in Section 2.6.) In Appendix Figure A.2, we present Lorenz curves and Gini coefficients of employment-plant concentration at the MSA and MSA-industry (4-digit NAICS) cells. We do so for the p10, p25, p50, p75 and p90 of cells in terms of their Gini coefficient. We drop MSAs and MSA-industry cells with fewer than ten plants, which zooms into the cells least likely to be subject to granularity bias (and for the Lorenz curve to have a clear interpretation). This restriction does not bind when we define cells as MSAs. Again we draw on public-use CBP data, with further details on the construction of the indices contained in the figure note.

In manufacturing, the typical (median) MSA and MSA-industry cell are tremendously concentrated. Pooling all industries, the Gini coefficient for the median MSA is 0.76 and ranges from 0.70 to 0.82 for the 10th and 90th percentile MSAs. In the median MSA, the top 5% of plants employ 46% of manufacturing workers. The top 20% of plants account for 80% of manufacturing employees. For the MSA-industries, where the restriction of at least ten plants drops 75% of cells and hence selects cells least subject to granularity, the

³The corresponding 75th to 25th percentile differences are 68% (0.520 in logs) and 145% (0.898 in logs).

median Gini coefficient is 0.63, and the 10th and 90th percentile Gini coefficients are 0.49 and 0.77. In the MSA-industry cell with the median Gini coefficient among even these remaining cells, the top 5% of plants account for 24% of employment, and the top 20% account for 64%. But even in the bottom 10th percentile of the MSAs and MSA-industries in terms of concentration (Gini coefficient), 72% and 44% of employment are accounted for by the largest 20% of plants.

2.4 Permutation Test: Pure Granularity and No Place Effects

To assess the scope for granularity bias in the data, our first strategy is to consider an extreme benchmark for the distribution of place effects: that all locations $l \in L$ have the same data-generating process for plant productivity $F_l^a(a) = F^a(a) \forall l \in L$. We test a more specific version, namely that all places have the same expected value $\tau_l = \tau \forall l \in L$, so that the variance of true place effects is zero. Then, dispersion in measured place averages arises solely as an artifact of grouping heterogeneous plants, i.e., from granularity bias.

We implement a nonparametric (or randomization), exact test of this hypothesis in the spirit of permutation tests.⁴ We construct the sampling distribution of our test statistic of interest under the following procedure: plants are randomly distributed across space into places. Specifically, we preserve the count of plants in each place (in practice in each location-industry cell). Under this procedure, the rank of a given empirical dispersion static in the CDF of those of the random economies gives the nonparametric p -value corresponding to that null hypothesis.

Broadly, by referencing a random-location benchmark, our test of productivity place effects is in the spirit of Ellison and Glaeser (1997), who study whether the observed geographical concentration is statistically different from randomly located plants.

2.5 Bias Correction of Variance: Split Samples

To provide a constructive measure of the dispersion in true place effects, we implement a split-sample procedure to remove the granularity bias. We construct an average using one half of the plants within a given location and use it as an instrument for the average in the other half. Formally, we split the plants into two random and equally sized subsamples indexed by superscript A and B in each location l . We then estimate subsample-specific place effects $(\hat{\tau}_l^A, \hat{\tau}_l^B)$ for all cells (l, A) and (l, B) , now described by place times half-sample. We then calculate the *covariance* of the two separate sets of fixed effects across

⁴A potential alternative statistical test, however requiring parametric assumptions, is an F -test of all place averages (for example, estimated as fixed effects in a regression) being statistically different from zero.

locations l between half-samples A and B :

$$\text{Cov}(\widehat{\tau}_l^A, \widehat{\tau}_l^B) = \text{Cov}(\tau_l + \bar{u}_l^A, \tau_l + \bar{u}_l^B) \quad (7)$$

$$= \text{Var}(\tau_l) + \underbrace{\text{Cov}(\tau_l, \bar{u}_l^A)}_{=0} + \underbrace{\text{Cov}(\tau_l, \bar{u}_l^B)}_{=0} + \underbrace{\text{Cov}(\bar{u}_l^A, \bar{u}_l^B)}_{=0}, \quad (8)$$

where we have introduced the notation of $\bar{u}_l^S = \frac{1}{N_l^S} \sum_{p \in S_l} u_{pl}$ as the sample average of deviations u_{pl} in a sample S_l of location l . Because subsamples are drawn randomly in our procedure, the second and third terms are zero. To ensure the fourth term is zero, we introduce a, we believe plausible, assumption, that errors $u_{pl}, u_{p'}$ are independent within locations l .

Hence, the covariance of averages of randomly chosen subsamples is an unbiased estimator of the variance of the true place effects, eliminating granularity bias.⁵

2.6 Weighting and Industry Variation

There are two ways in which weighting might enter and affect the above exposition. First, the exposition above does not weight locations differently when constructing the cross-regional variance. That is, the bias term is the unweighted average over all location-specific bias terms. Consistent with this specification, in our empirical implementation, we weight MSAs equally (rather than giving larger MSAs a larger weight).

Second, the expressions above have presented the case where place effects τ_l equally weight (or consider equally sized) plants within a location. In practice, we weight plants by a plant employment e_{pl} when constructing cell-level averages. The true place effect then takes the weighted expectation of productivity over the joint distribution of plant productivity and size, $F^l(a, e)$. The main implication is that the bias from granularity now also encompasses the potential dominance of large plants. Importantly, the covariance remains the unbiased estimator of the variance of true place effects also in the weighted case. (For the permutation tests, we do not reassign plants based on size.) We additionally present specifications without weighting by plant size.

⁵This adjustment has been used in the context of group-level wage differences (namely, AKM firm wage fixed effects, by Gerard, Lagos, Severnini, and Card, 2018; Drenik, Jäger, Plotkin, and Schoefer, 2020; Kline, Saggio, and Sølvssten, 2020), or to estimate peer effects in personnel economics (Silver, 2020). We are not aware of applications to the measurement of group-level averages of productivity. Bils, Klenow, and Ruane (2020) use a within-firm IV strategy to adjust for measurement error in plant-level productivity, i.e., a separate method that reduces idiosyncratic measurement error but could not be used to estimate place effects on plant groups. An alternative method to remove measurement error may be shrinkage (see Chandra, Finkelstein, Sacarny, and Syverson, 2016, for an application to hospital fixed effects). We are not aware of shrinkage estimators in the context of productivity place effects.

Finally, the above exposition sidesteps industry differences, being written as if only one industry existed. We account for differences between industries in two ways. First, we implement our approach within industries and report our results as averages across industries. Second, we define plant productivity as deviation from industry average and pool across industries. These measures are presented in Section 3.

3 Data and Construction of Place Effects

We now describe our empirical implementation of the framework developed in Section 2.

Plant-Level Data: US Census of Manufactures Our primary data set is the US Census of Manufactures (CMF), which provides plant-level data on production and plant characteristics for the universe of US manufacturing plants. We use the most recently available wave, 2012.⁶ As our location measure, we use plants' MSA (dropping plants outside of MSAs). While the data contain 6-digit NAICS industry codes (we use industry definitions from Fort and Klimek, 2018), we will coarsen the measure to 4-digit for most purposes below.

The CMF contains information on revenue, employment and payroll (separately for production and non-production workers), production worker hours, material and energy expenditures, and capital expenditures. Value added is revenue minus non-labor/capital inputs. Since we include an extension studying place effects for new and old plants in Section 5, we also construct plant age as the difference between 2012 and the first time the plant enters the Longitudinal Business Database (LBD), which we merge onto the CMF for this purpose.

We require plants in our sample to have an industry code, to be located in an MSA, and to have positive value added as well as TFP, described below, and remove administrative records.

Plant-Level Productivity: TFP and Labor Productivity Our primary measure of plant productivity is revenue-based TFP ("TFPr"). For this productivity measure, we draw on the TFPr measures constructed by Foster, Grim, and Haltiwanger (2016) and updated in Decker, Haltiwanger, Jarmin, and Miranda (forthcoming) through 2013 and covering our 2012 CMF wave. The construction assumes a standard Cobb-Douglas production function

⁶While 2012 is in the aftermath of the Great Recession and Kehrig (2015) documents that TFP dispersion is countercyclical in the US manufacturing sector, we have found that our findings are overall robust to cyclical properties, e.g., by pooling multiple Census waves.

$Y_p = A_p \prod_{\iota} Q_p^{\iota c_i^{\iota}(p)}$ with constant returns in input quantities Q_p^{ι} of type ι , each with industry-specific factor shares c_i^{ι} . A plant p 's TFP is the residual of its inputs (capital, labor, materials, and energy) subtracted from revenue output, with industry-level factor shares $i(p)$:

$$a_p = \ln Y_p - [c_i^s \ln S_p + c_i^k \ln K_p + c_i^h \ln H_p + c_i^m \ln M_p + c_i^e \ln E_p] \quad (9)$$

Details are provided in Foster, Grim, and Haltiwanger (2016) and Decker, Haltiwanger, Jarmin, and Miranda (forthcoming), which we summarize here. Revenue-based output Y is the real value of shipments plus changes in inventory (or value of shipments if the difference is negative), deflated using a 6-digit NAICS industry output price deflator from the NBER-CES Manufacturing Industry Database. For lack of comprehensive plant-level price data, between-plant demand factors in the form of plant-level product price differences show up in this revenue-based TFP measure (unlike “TFPq”). The labor input H is total hours of production workers (marked up by the ratio of total to production worker payroll if both are nonmissing, otherwise production worker hours). The capital stock construction also draws on the Annual Survey of Manufacturers (ASM), consists of structures S and equipment K , and is, in most cases, obtained from the perpetual inventory method separately for equipment and structures, with initial values given by the book value (adjusted by the ratio of real to book value from BEA data, at the 3-digit NAICS level), and then evolves using capital expenditure data from the ASM where applicable. Materials M are the cost of materials plus the cost of resales plus the work done by other plants on the materials, deflated by the 6-digit NAICS input price deflator from the NBER-CES Manufacturing Industry Database, similarly for energy costs E as the costs of electricity and fuels. For each factor (with buildings and equipment separately), industry cost shares c_i^{ι} are at the 6-digit NAICS level constructed in the NBER-CES Manufacturing Industry Database.

We complement our primary productivity measure, TFP, with value added per worker (labor productivity), where our labor input concept is total employment (Foster, Haltiwanger, and Krizan, 2001, find that using total hours and employment yield similar results). This measure is less demanding, requiring neither specifying a production function nor comprehensive input measures. As a benchmark, with Cobb-Douglas production, the marginal product of labor corresponds to the labor share in production times output per worker. As we take logs and include industry effects, we would net out industry labor shares, such that one can think of this alternative measure as marginal product of labor.

Winsorization In our primary specifications, we winsorize the final plant-level productivity measures at 1% and 99%. We also probe robustness to no as well as 2.5%

winsorization.⁷

Defining Places Our primary definition of places are location-industry cells, by MSA and 4-digit NAICS industry. We also probe robustness to 6-digit industry classification.

We keep location-industry cells with at least two plants, the minimum number required for our split-sample strategy described in Section 2.5. We will probe robustness to varying this cutoff between two to 150 plants.

Our main sample consists of around 120,000 plants in around 11,500 cells. (Plant, MSA, and cell counts must be rounded to meet Census disclosure requirements.) There are 384 US MSAs and 86 4-digit NAICS manufacturing industries.

Industry-Specific Location Effects (Averages) We define industry-specific location effects as in Equation (2) (plus an index for industry i), where a plant p is characterized by TFP a_{pli} , 4-digit NAICS industry $i(p)$, and location $l(p)$. Industry-specific location effects are estimated as industry-demeaned averages:

$$\widehat{\tau}_{l,i} = \text{Avg}[a_{pli}|l, i] - \text{Avg}[a_{pli}|i]. \quad (10)$$

By subtracting industry mean $\text{Avg}[a_{pli}|i]$, we construct the industry-specific location effect $\widehat{\tau}_{i,l}$ as the location-industry's plant-level average log TFP premium over the national industry average of plant-level log TFP.⁸ (Thanks to the dramatically larger sample size, we sidestep granularity bias in the *national* industry average $\text{Avg}[a_{pli}|i]$.) We follow the literature in weighting both these averages by plant employment (see, e.g. Hornbeck and Moretti, 2018). We will also probe robustness to unweighted specifications.

Since the place effects are demeaned by industry, we can pool all industries' place effects and report the resulting variances across all industries, corresponding to the average of within-industry variances, and our headline number. Whenever we refer to dispersion statistics or plot distributions of location-industry effects, we do so weighted by the share of the industry in total employment within its respective location, for comparability with the location effects that pool all industries in a location, described below.

⁷Foster, Grim, and Haltiwanger (2016) also drop a small number of plants with imputed data as well as plants with TFP more than 200 log points above or below the industry-year mean.

⁸Our choice to demean within industry is also taken by, for instance, Cunningham, Foster, Grim, Haltiwanger, Pablonia, Stewart, and Wolf (2020), who express plant-level TFP from an industry mean when studying between-plant dispersion. They do not study geographical dispersion, nor do they touch on granularity bias.

Location Effects (Averages) As our most comprehensive place effect measure, we collapse (average) the industry-specific location effects $\widehat{\tau}_{l,i}$ into one location effect for each location l , defined as

$$\widehat{\xi}_l = \text{Avg}[\widehat{\tau}_{l,i}|l]. \quad (11)$$

Here, we weight each industry-specific location effect $\widehat{\tau}_{l,i}$ by the industry’s local employment share (among the cells surviving sample restrictions i.e., with non-missing $\widehat{\tau}_{l,i}$) within location l . This weighting mimics the plant-employment measure of the industry-specific location effects and hence the scenario of a single national industry. Again, we will also probe robustness to specifications equally weighting plants. Of course, the industry-MSA panel will not be balanced if some locations do not have (at least two) plants in a given industry, such that this value is constructed among an MSA’s filled industries. The aforementioned demeaning of the industry-specific location effect $\widehat{\tau}_{l,i}$ ensures that we have removed industry composition and industry-level TFP levels as confounders in the location effects $\widehat{\xi}_l$. In a final step, we center the national (unweighted across MSAs) mean of these location effects too around zero.

Hence, location effects $\widehat{\xi}_l$ can be thought of as the common location component across all industries in a given location, whereas the industry-specific location effects $\widehat{\tau}_{l,i}$ essentially treat location-industries as independent entities with no connection to neighboring industries. When place effects are not perfectly correlated across industries, the dispersion of location-industry effects will be larger than that of location effects.⁹

European Countries: Bureau van Dijk Firm-Level Data We complement our primary analysis, of US plant-level data, with firm-level data from Bureau van Dijk (BvD) covering 15 European countries. On a country-by-country basis, we replicate our analysis using NUTS-2 within-country regional divisions, which most closely resemble US MSAs, with each containing between 800,000 and 3 million inhabitants. We construct TFP measures for the manufacturing sector, at the 2-digit NACE industry level (due to the lower number of observations in BvD rather than census data), which gives us 23 industries. We obtain the industry-country-specific labor share by dividing the sum of payroll at the industry level of all sample firms by the corresponding sum of value added for firms with nonmissing observations on both variables. We then construct firm-specific TFP by assuming a Cobb-Douglas production function, and subtracting from log value added labor-share-weighted employment and one-minus-labor-share-weighted log capital. We use fixed tangible assets as the capital stock measure. We winsorize the resulting TFP measures at 1% and 99%.

⁹Our location effect ξ_l hence resembles the “TFP2” measure in Combes et al. (2010), which, at the location level, averages residuals of TFP after controlling for sector.

We again keep all location-industry cells with at least 2 firms. Appendix Table A.2 lists number of regions, cells, and firm counts in each country.

These data have several drawbacks for our purposes compared to the US plant census. For instance, coverage is imperfect, and data quality varies across countries specifically regarding value added (see, e.g., Gopinath, Kalemli-Özcan, Karabarbounis, and Villegas-Sanchez, 2017; Jäger, Schoefer, and Heining, 2019). To maximize coverage and mimic a census, we keep each firm’s most recent observation, implying that most observations come from the late 2010s. The capital stock measure is based on book value of assets. We do not apply industry-specific input price indices (industry fixed effects absorb national output price indices). Since BvD is at the firm rather than plant level, all production units of multi-plant firms are assigned to a single industry and headquarter location.

4 Results: Productivity Dispersion Across US Cities

We first measure the raw geographic dispersion in productivity across US cities. We then implement our permutation test of the null hypothesis that this empirical variance is entirely spurious and would arise even if plants were randomly allocated across places. We implement our split-sample strategy to cleanse the naive dispersion of granularity bias and provide the unbiased estimate of the variance of true place effects. Finally, we run a series of robustness checks which dissect the sources of the granularity bias we uncover. Table 1 reports the key numbers cited here. In all of our figures and dispersion statistics, we weight MSAs equally, and MSA-industries by local employment share, as discussed in Section 2.6.

4.1 Raw Dispersion

Figure 1 plots the distribution of average-based productivity place effects for locations, $\widehat{\xi}_l$, and for location-industries, $\widehat{\pi}_{l,i}$. In accordance with Census disclosure requirements, the graphs present kernel densities rather than histograms and censors the density plots at the 5th and 95th percentiles. Hence, tails and potential skewness are not depicted (so the mode need not be centered at zero although we have centered place effects at zero over the full support).

Location Effects (Averages) $\widehat{\xi}_l$ Figure 1 Panel (a) plots the distribution of location-specific effects based on averages, $\widehat{\xi}_l$. The focus of this section is the black solid line, representing the raw distribution of average-based place effects. The location effects $\widehat{\xi}_l$ trace out a bell-

Table 1: Spatial Dispersion of Productivity (TFP) in the United States (x100 For All Dispersion Statistics)

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) |
|---|---------|------------|--------------|---------|------------|---------------|-------------|-----------|--------|--------|
| | Main | ≥10 Plants | 6-Digit Ind. | Unwins. | 2.5% Wins. | Plant Weights | Labor Prod. | New & Old | New | Old |
| Panel A: Location Effects | | | | | | | | | | |
| Empirical: | | | | | | | | | | |
| Raw $\text{Var}(\hat{\xi}_i)$ | 2.55 | 2.07 | 1.38 | 2.86 | 2.21 | 0.50 | 6.32 | 3.19 | 8.40 | 3.26 |
| 90th – 10th Percentile | 47.48 | 46.26 | 36.95 | 51.08 | 45.33 | 23.46 | 85.20 | 56.87 | 74.43 | 60.66 |
| Permutations: $\text{Var}(\hat{\xi}_i)$ | | | | | | | | | | |
| Mean | 1.71 | 1.44 | 1.37 | 1.93 | 1.45 | 0.30 | 4.62 | 2.47 | 6.17 | 3.21 |
| Standard Deviation | 0.28 | 0.24 | 0.28 | 0.36 | 0.21 | 0.04 | 0.61 | 0.41 | 0.96 | 0.61 |
| p-value | 0.010 | 0.017 | 0.425 | 0.017 | 0.005 | 0.004 | 0.009 | 0.056 | 0.022 | 0.420 |
| Bias-Corrected: $\text{Cov}(\hat{\xi}_i^A, \hat{\xi}_i^B) = \text{Var}(\xi_i)$ | | | | | | | | | | |
| Mean | 0.81 | 0.81 | 0.33 | 0.91 | 0.66 | 0.21 | 1.58 | 0.87 | 1.49 | 0.84 |
| 97.5th Percentile | 0.98 | 1.05 | 0.46 | 1.12 | 0.82 | 0.27 | 1.97 | 1.30 | 2.38 | 1.22 |
| 2.5th Percentile | 0.61 | 0.55 | 0.19 | 0.66 | 0.49 | 0.14 | 1.17 | 0.46 | 0.44 | 0.40 |
| Panel B: Industry-Specific Location Effects | | | | | | | | | | |
| Empirical: | | | | | | | | | | |
| Raw $\text{Var}(\hat{\tau}_{i,i})$ | 10.96 | 4.45 | 7.23 | 12.37 | 9.41 | 4.88 | 28.96 | 6.39 | 15.24 | 6.59 |
| 90th – 10th Percentile | 98.01 | 66.14 | 75.89 | 102.3 | 93.13 | 68.81 | 188.9 | 81.89 | 111.0 | 84.96 |
| Permutations: $\text{Var}(\hat{\tau}_{i,i})$ | | | | | | | | | | |
| Mean | 9.64 | 3.47 | 7.08 | 10.95 | 8.25 | 4.24 | 26.75 | 5.04 | 12.71 | 6.13 |
| Standard Deviation | 0.53 | 0.31 | 0.48 | 0.71 | 0.41 | 0.13 | 1.12 | 0.48 | 1.09 | 0.69 |
| p-value | 0.016 | 0.005 | 0.352 | 0.037 | 0.006 | 0.002 | 0.029 | 0.009 | 0.019 | 0.234 |
| Bias-Corrected: $\text{Cov}(\hat{\tau}_{i,i}^A, \hat{\tau}_{i,i}^B) = \text{Var}(\tau_{i,i})$ | | | | | | | | | | |
| Mean | 2.32 | 1.39 | 0.88 | 2.63 | 1.87 | 0.84 | 4.05 | 1.71 | 1.90 | 1.59 |
| 97.5th Percentile | 2.70 | 1.74 | 1.08 | 3.12 | 2.19 | 0.96 | 4.84 | 2.21 | 2.87 | 2.06 |
| 2.5th Percentile | 1.86 | 1.00 | 0.65 | 2.07 | 1.50 | 0.73 | 3.24 | 1.20 | 0.70 | 1.04 |
| N, MSAs | 400 | 250 | 400 | 400 | 400 | 400 | 400 | 300 | 300 | 300 |
| N, MSA-Industries | 11,500 | 2,800 | 18,000 | 11,500 | 11,500 | 11,500 | 11,500 | 2,800 | 2,800 | 2,800 |
| N, Plants | 120,000 | 86,000 | 105,000 | 120,000 | 120,000 | 120,000 | 120,000 | 78,000 | 14,000 | 64,000 |

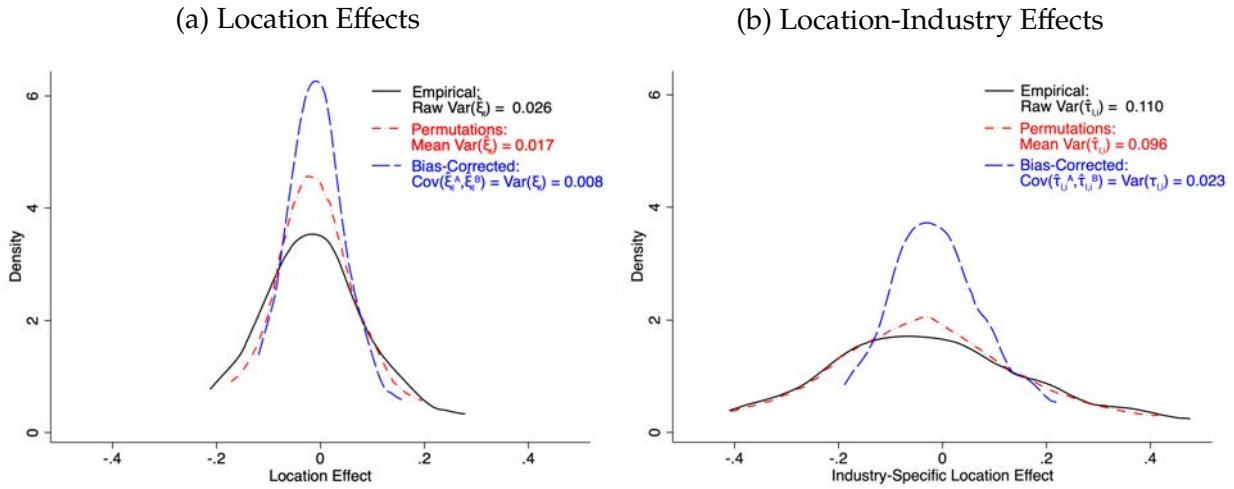
Note: The table reports statistics on the dispersion of productivity of MSA effects (Panel A) and MSA-industry (4-digit NAICS) effects (Panel B), estimated in the 2012 US Census of Manufacturers. In order to satisfy Census disclosure requirements, 90th–10th percentile differences are calculated by reporting the difference in the average of place effect from the 85th to 95th percentiles, and the average of place effects from the 5th to 15th percentiles. All specifications use TFP_r as the productivity measure, except for Column (7), which studies labor productivity (log value added per worker). Column (1) reports on our baseline sample and specification. All other columns report specifications that each change one aspect of Column (1), as follows: Column (2) requires at least ten (rather than two) plants per industry-MSA cell. Column (3) defines industries at the 6-digit rather than 4-digit NAICS level. Column (4) uses plant-level data that has not been winsorized, as opposed to our main specification with 1% winsorization. Column (5) uses 2.5% winsorization. Column (6) uses identical plant weights (rather than weighting by plant employment). Column (7) repeats our main specification using log value added per worker. Column (8) requires at least two plants each five years or younger (new plants) and two plants older than five years (old plants). It is stricter than the baseline specification of at least two plants, and is the pooled comparison for the following two columns. Column (9) includes only the new plants from the sample of Column (8). Column (10) includes only the old plants from the sample of Column (8).

shaped distribution. (The censoring by the Census disclosure process masks potential skewness.) As printed into Figure 1, the variance is 0.026. This statistic is our headline number reported in the introduction in Section 1. That is, plants in MSAs with location averages one standard deviation above the mean have on average around $\sqrt{0.026} = 0.16$ higher productivity (log TFP) than plants in their peer industries nationally.

While our main dispersion statistic is the variance, we also provide the productivity difference between the 90th and 10th percentiles in the second row of Table 1.¹⁰ This difference in log TFP is 0.47 (61%)—on average, plants in the 90th percentile MSA are 61%

¹⁰In order to avoid disclosure of identifiable data, we approximate the 90-10 dispersion ratio using the difference in the averages of the 85th-95th percentiles and 5th-15th percentiles.

Figure 1: Spatial Dispersion in TFP Across US MSAs: Raw Empirical Place Averages, Benchmark from 1,000 Random Permutations of Plants Across Places, and Corrected for Granularity Bias



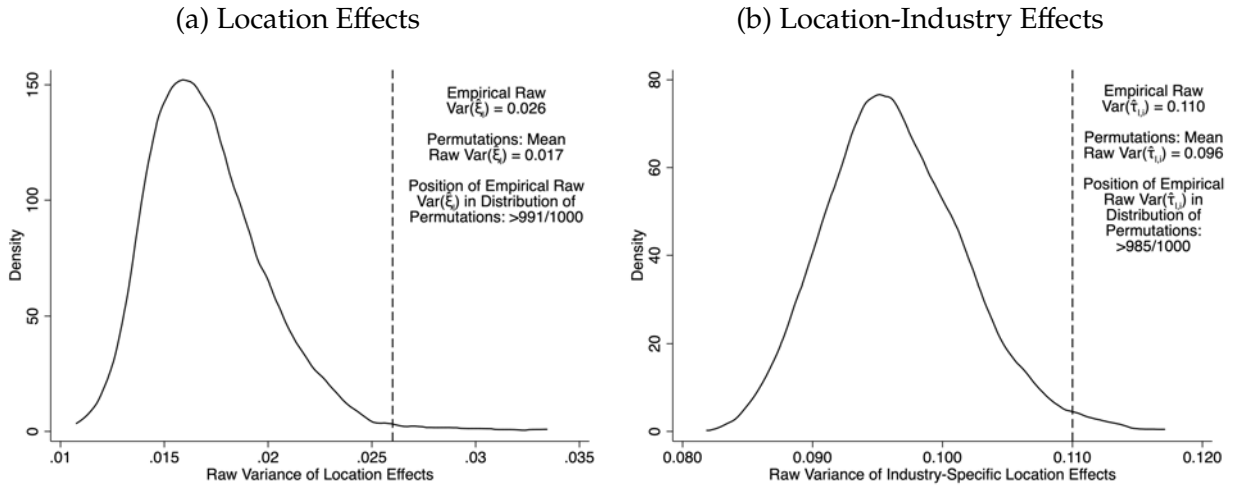
Note: The figure reports kernel density plots representing the distribution of location effects (Panel (a)) and industry-specific location effects (Panel (b)) estimated across US MSAs and MSA-industry (4-digit NAICS) cells, in the 2012 US Census of Manufacturers. The solid black line plots the distribution of the raw, average-based place effects in the actual US data. The dash-dotted red line shows the distribution of analogous average-based place effects from a permutation tests, randomly allocating the empirical plants across US MSA-industry cells (plotting the representative randomized economy with the raw variance that is closest to the mean raw variances of all 1,000 permutations). The blue dashed line illustrates the distribution corrected for granularity bias. It does so by applying a mean-preserving variance-adjustment using a linear transformation of the original distribution as described in Footnote 12. Panel (a) is not weighted; Panel (b) weights location-industry effects by the industry’s employment shares within the location. In accordance with Census disclosure requirements, the figure presents kernel densities rather than histograms and censors the density plots at the 5th and 95th percentiles.

more productive than plants in the MSA at the 10th percentile.

Location-Industry Effects (Averages) $\hat{\tau}_{l,i}$ Panel (b) replicates this analysis for the location-industry effects $\hat{\tau}_{l,i}$, which permit place effects to vary by industry. Recall that, when plotting the variance and constructing dispersion statistics, we weight each location-industry effect by the share of the industry in total employment within its respective location, for comparability with the location effects $\hat{\xi}_l$ plotted in Panel (a), which at the MSA level weighted industries by employment shares.

As a benchmark, if all industries in a location deviated by the same percent from their national industry benchmark, then the location and location-industry effects would exhibit the same dispersion. If productivity premia are imperfectly correlated across industry within a location, then industry-specific location effects may exhibit more dispersion than location effects. We find a considerably more dispersed distribution of the industry-specific location effects, with a variance of 0.110. The difference in log TFP between the 90th and 10th percentile is 0.98 (i.e., 166%).

Figure 2: Permutation Test: Sampling Distribution of Raw Variances of Place Effects From 1,000 Random Allocations of Plants Over Places



Note: The figure reports kernel density plots corresponding to the sampling distribution of 1,000 economies' raw variances of average-based place effects, for location effects (Panel (a)) and industry-specific location effects (Panel (b)). The randomization procedure reassigns MSA IDs over the empirical plant observations within an industry, preserving the plant count distribution in each MSA-industry cell. The vertical dashed line denotes the empirical raw variance.

4.2 Permutation Test: Pure Granularity and No Place Effects

Section 4.1 has quantified the degree to which places differ in their observed average productivity. However, some of this dispersion could be spurious, reflecting dispersion in idiosyncratic productivity levels of heterogeneous plants rather than true place effects. We now test the extreme hypothesis that there is no variation in place effects, and that the raw variance of the measured average-based place effects reflects idiosyncratic plant heterogeneity only.

Implementation Implementing the general methodology in Section 2.4, we randomly relabel each plant's location $l(p)$ while preserving the empirical plant-count distribution in each location-industry $l \in L$, ensuring MSAs' industry structures are unaltered in terms of plant counts. Since we draw without replacement, all plants are used exactly once per randomization. We generate 1,000 randomized economies. We treat each randomized economy as we did the empirical one, calculating place averages of productivity and their raw variance. We also construct 95% "confidence intervals" by extracting the 25th and 975th ranked observations in this sampling distribution; the position of the empirical raw variance in this sampling distribution gives the p -value, which we also report.

Results Across the 1,000 randomized US economies, the mean variances are 0.017 for the location effects, and 0.096 for the location-industry effects, compared to 0.026 and

0.110 for the empirical raw variances calculated and discussed above in Section 4.1. That is, granularity on its own already generates tremendous—purely spurious—variation in average-based place effects.

For illustration, Figure 1 also includes, as dot-dashed red lines, the distribution of place effects in the specific randomized economy most closely matching the mean permuted variance. Intuitively, the closer that red dot-dashed line is to the black line representing the actual US economy, the more similar the random economies are to the empirical distribution, and the more granularity alone could account for the observed dispersion.

In Figure 2, we present the distribution of raw variances of the 1,000 randomized US economies, the mean of which we reported above. (Here, Census disclosure guidelines do permit us to release the uncensored distribution.) This distribution is the nonparametric sampling distribution of the variance of place averages under the assumption of no place effects. The vertical dashed line denotes the level of the empirical variance. For location effects, Panel (a) clarifies that the empirical variance is above 991 of the 1000 permutation values given the sampling distribution; the location-industry effects in Panel (b) puts the data above the 985th observation of the 1000 permutation values. These values imply one-sided p -values of 0.010 and 0.016. That is, the empirical variance is statistically more dispersed than would be expected from a null hypothesis of no place effects, even though such economies would generate substantial purely spurious variance.

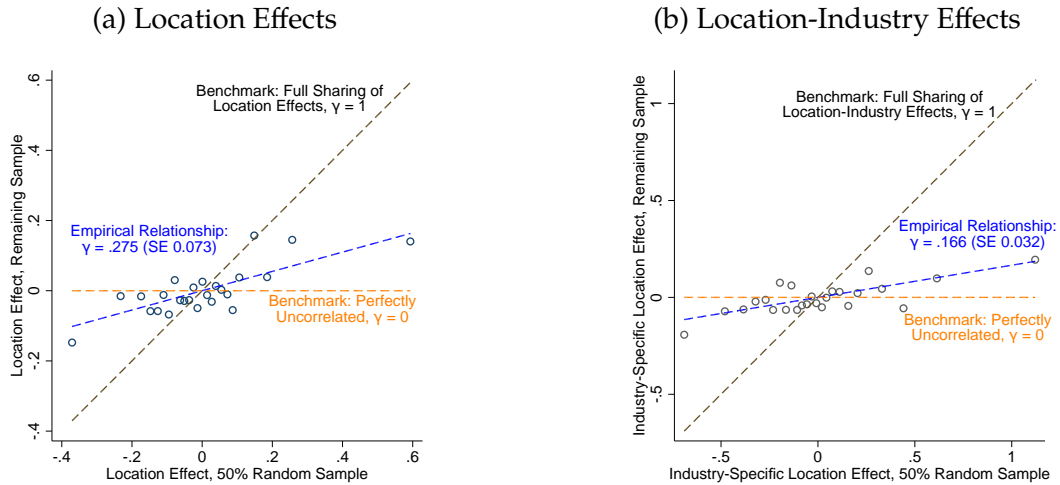
This test provides a nonparametric statistical rejection of the absence of true place effects, which hence contribute to the observed variance of place averages of productivity. But to obtain an estimate of the variance of true place effects, we cannot simply subtract the raw variance from the counterfactual mean variance of the random economies, for instance, as, in the permutation test, plants' productivity levels a_{pli} take with them their potential true place effects τ_l . Instead, we next implement our split-sample correction to constructively quantify the variance of true place effects.

4.3 True Place Effects: Bias Correction of Variance With a Split Sample Instrumental Variable Strategy

We now implement our split-sample procedure in order to constructively quantify the variance of true place effects, by removing the granularity bias.

Implementation The concrete implementation of the methodology laid out in Section 2.5 takes the following steps. While the covariance is an unbiased and consistent estimator of the variance, granularity—the very reason we draw on this method—may still imply

Figure 3: Illustration of Method: Split-Sample Correction of Raw Variance Removing Granularity Bias



Note: The figure presents scatter plots juxtaposing, along the common location or location-industry ID, the estimated place effects from one split sample on place effects of the other split sample. The observations are binned into 25 equally sized bins, although the regressions are run on the underlying observations. We also plot two benchmarks. $\gamma = 0$ represents the scenario of no place effects whatsoever, i.e., no relationship between place effects of the split samples; $\gamma = 1$ represents the scenario in which place effects feature no attenuation bias from measurement error such as granularity. The blue line traces out the linear slope from the regression of the y-axis effects (one split sample) against those of their x-axis neighbor (other split sample). Since the underlying univariate regression coefficient represents the covariance of the variables on the two axes divided by the variance of the x-axis variable, and since the split-sample covariance is the bias-corrected estimator of the variance, this coefficient also represents the share of the variance (of the half sample depicted on the x-axis) surviving the bias correction. Panel (a) is not weighted; Panel (b) weights location-industry effects by the industry’s employment shares within the location.

substantial error in one given random sample split. We therefore implement 1,000 random sample splits, and extract the resulting distribution of covariances. We report the mean as our preferred statistic, but additionally provide information on the distribution. Specifically, we extract the, informal, nonparametric 95% “confidence intervals” implied by the sampling distribution of the resulting 1,000 covariances.¹¹ When constructing location-industry effects and computing covariances, we weight by the location-industry employment share of the industry in the MSA of the full rather than split sample employment shares. (As before, we do not weight across MSAs.)

Results The bias correction dramatically reduces the variance of productivity across US regions, by more than two thirds for location effects, from 0.026 to 0.008, with 95% of our draws falling in the range of 0.006 to 0.010. For industry-specific location effects, the variance is reduced by four fifths, from 0.110 to 0.023, with 95% of the draws falling in the range of 0.019 to 0.027. Figure 1 illustrates the effect of the bias-correction on the distribution of place effects as a blue dashed line, drawing on a simple mean-preserving

¹¹See Kline, Saggio, and Sølvsten (2020) for a formal econometric framework for leave-out estimation of variance components with an application to firm fixed effects in worker-level wages.

linear transformation of the raw distribution to match the bias-corrected variance.¹² That is, this distribution corresponds to one that matches the variance of true place effects, which is dramatically more compressed than the raw distributions. Hence, on the one hand, two thirds to four fifths of the cross-regional variation in productivity reflects the bias arising from granularity unrelated to true place effects. On the other hand, and consistent with our permutation test, the remaining variation constitutes the still economically and statistically significant variance in true place effects.

Visualization of Method We visualize the bias correction in Figure 3. The binned scatter plot depicts one specific split-sample economy, selected among the 1,000 randomizations to have its covariance-to-variance ratio be closest to the corresponding mean value of our 1000 split-sample economies' coefficients, as described below. Panel (a) reports on the location effects, Panel (b) on the industry-specific location effects. The panels are scatter plots, juxtaposing, for each place, its pair of place effects computed separately on the basis of the samples split in half within each place. The graph bins the underlying observations into 25 equally sized bins, but the regressions are run on the underlying individual place effects. The x-axis captures the means of each bin from one split sample, and hence traces out the raw dispersion of average-based place effects (in the half sample). As throughout, we do not weight across MSAs, although, for industry-MSA effects, we weight by the industry's employment share in the MSA (here again, as above, weights are constructed off the pooled sample, rather than on the basis of a specific sample).

The figure also includes two extreme benchmarks. Intuitively, if there were no relationship in productivity whatsoever between the two split samples A and B within a place, each split sample average would reflect idiosyncratic effects only, and a line fitted to the scatter plot would have a slope of zero through the origin. As another benchmark, in the absence of idiosyncratic effects, place effects, common to both split samples, would dominate the split samples' averages, such that the scatter would align along the 45 degree line.

The empirical effects, depicted as the binned scatter points, fall somewhere in between these two extremes, providing a striking visual clarification that the empirical economy is characterized by a large degree of granularity in productivity differences across places. Besides plotting the binned scatter points, the figure plots the linear regression line implied

¹² That is, we construct $x' = a + bx$ and $f_{x'} = f((x-a)/b)/b$. We set $b = \sqrt{\frac{\text{Var}(x')}{\text{Var}(x)}}$ to match the desired variance of the transformed distribution, and $a = (1 - b) E[x]$ to preserve the mean. We resort to this procedure as an illustration of the split sample method, which provides an estimate of the variance, but of no other moment of the distribution.

by the data (where estimation is on the actual underlying data rather than the binned data). In fact, out of the 1,000 split samples, we have chosen the representative split sample depicted in the figure as the one with the slope that is closest to the mean slope of the 1,000 split-sample economies. We estimate a slope of 0.275 for the location effects, and a slope of 0.166 for industry-specific location effects.¹³ The linear regression coefficient corresponds to $\gamma = \frac{\text{Cov}(\widehat{\xi}_l^A, \widehat{\xi}_l^B)}{\text{Var}(\widehat{\xi}_l^A)}$ for the location effects, and analogously for the location-industry effects. Since $\text{Cov}(\widehat{\xi}_l^A, \widehat{\xi}_l^B)$ is our bias-corrected estimate of the (full sample $A \cup B$) variance $\text{Var}(\xi_l^{A \cup B})$, the regression coefficient represents the share of the raw variance (of the half sample making up the variation traced out on the x-axis) that survives bias correction. These positive slopes therefore confirm the attenuated, yet existent, presence of true place effects we have discussed above by comparing the mean covariance of the split samples with the population raw variance. Of course, the slope actually implies a slightly smaller fraction because the split-sample variance (i.e., the denominator corresponding to the estimated coefficient, $\text{Var}(\widehat{\xi}_l^A)$) is slightly larger exactly due to heightened granularity from halving the sample size.¹⁴

4.4 Assessing the Sources of Granularity

We now provide a series of alternative specifications with the goal of tracing out the sources of granularity bias and explore, but ultimately dismiss, potential alternative strategies to reduce it. Mirroring the discussion in Section 2.3 and guided by our core Equation 6, we dissect the three potential sources: finite plant counts per place, large idiosyncratic dispersion in plant-level productivity unrelated to place, and heterogeneity in plant size. Our main exhibits for these checks are Table 1 and Figure 4.

Plant Counts per Cell As a direct way to gauge the role of cell counts, we vary the minimum (location-industry) cell size cutoff. Table 1 Column (1) reports the baseline specification with minimum size requirement of at least two plants per cells. Column (2) reports the statistics for a sample restriction with at least ten plants. In Figure 4, we vary this restriction incrementally from two to 150 (with spacing determined by Census disclosure rules).

¹³The robust standard errors of those slope estimates are 0.073 and 0.032, respectively; we do not emphasize the standard errors and sidestep that the fixed effects are generated regressors.

¹⁴We find that the split sample raw variance is 133% as large as the population variance for the location effects ($\text{Var}(\widehat{\xi}_l^A) = 0.034$ vs. $\text{Var}(\widehat{\xi}_l^{A \cup B}) = 0.026$) and 126% as large for location-industry effects ($\text{Var}(\widehat{\tau}_{l,i}^A) = 0.138$ vs. $\text{Var}(\widehat{\tau}_{l,i}^{A \cup B}) = 0.110$). On average across our 1,000 split economies these numbers are 0.030 and 0.140 respectively. Another granularity factor (besides idiosyncratic heterogeneity) is skewed plant size, which we discuss in Section 4.4.

Consistent with a decline in granularity bias, the raw variance falls when we raise the minimum cell count. Column (2) of Table 1 clarifies that moving to ten rather than two plants per cell, the raw variance falls from 0.026 to 0.021 for location effects, and, more dramatically, from 0.110 to 0.045 for location-industry effects. As depicted with the solid black line in the figure, the raw variance falls as the minimum plant count per cell increases, and more so for the industry-specific location effects (Panel B in the table, and Panel (b) in the figure) than for the location effects (Panel A and panel (a)).

Yet, inevitably, restricting the sample to larger and larger cells has compositional effects beyond granularity bias.¹⁵ It is possible that the remaining places have place effects that are more and more similar. As shown in Table 1 Column (2) Panel A, starting from the threshold two and moving to ten, the covariance for location effects is essentially unchanged. Figure 4 Panel (a), which plots the (mean) covariance in the blue dashed-dotted line, reveals that the covariance fluctuates around the original value until around 40 plants as the minimum sample restriction, and then the covariance starts dropping gradually, reaching zero at around 100 plants. For industry-specific location effects, Table 1 Column (2) Panel B reports a drop in the covariance to 0.014 with at least ten plants compared to 0.023 with at least two plants. Figure 4 Panel (b) reports this moderately steeper gradient of the covariance to the minimum cell count for industry-specific place effects.

The p -values of the permutation tests for the specification with at least 10 plants, reported in Table 1 Column (2), continue to indicate that even with more cells, the empirical economies exhibit a raw variance that remains squarely statistically different from what a random allocation of plants across places would have predicted. This fact is indicated by the red dashed line and the associated 95% confident intervals, making clear that the empirical raw variance only crosses these confidence bands at around 100 plants, for both location and location-industry effects.

The figure also makes clear the catch-22 that granularity bias provides: the easiest “solution” to reduce granularity bias is to restrict one’s analysis sample to cells that are less subject to it—throwing the baby out with the bathwater. Comparing Column (1) (our baseline specification with at least two plants per cell) with Column (2) (where we require at least ten plants) in Table 1, illustrates this trade-off: only about 250 out of our initial set of MSAs remain, 2,800 out of the 11,500 MSA-industry cells, and around 86,000 out of the

¹⁵An alternative route is to hold the MSA-industry sample constant (at a new baseline sample with more than two plants per cell) but randomly drop plants in cells. In this alternative route, the covariance would stay stable, and the raw variance and permutation results would reflect purely plant counts. Our split-sample method implements exactly this approach in the context of cutting the sample randomly in half and permits an analysis of the raw variance (but not the covariance), on which we report in Footnote 14.

120,000 plants (where Census disclosure requirements force us to round all counts). The black dots in Figure 4 report on the steep sample count (cells) trade-off for the interval between two and 150 plants, which is most dramatic when going from two to around 20-30, and then continues more gradually. At a cell-size cutoff of 100, the remaining sample features only 10% of MSAs (and a much lower fraction of industry-MSA cells), fails the permutation test, and features a bias-corrected variance of zero.

Another way to adjust plant cell counts is to redefine cells. Our results for location effects ξ_i as aggregates of industry-specific location effects $\tau_{l,i}$ already speak to the impact of aggregating industry cells. Alternatively, we now redefine industry cells from 4-digit to 6-digit NAICS, reporting results in Column (3) of Table 1. There are 359 6-digit NAICS manufacturing industries nested in the 86 4-digit industries, and indeed, Column (3) reveals that the number of plants per location-industry cell falls while the number of location-industry cells increase. A pure granularity perspective, holding the composition of place effects constant, would predict an increase in raw variances. Yet, there is a composition shift as well, as some 6-digit cells do not have at least two plants at the 6-digit industry-location level, as indicated by the lower plant count. Indeed, the remaining cells appear more homogeneous: though the raw variance falls by less than half, the covariance falls by around two thirds.¹⁶ Moreover, this sample and specification fails the permutation test, with average permuted variances being close to the empirical raw variance. Overall, we here conclude that the composition differences that go along with the cell definition appear to again prevent us from isolating granularity.¹⁷

Besides the limitation that the above exercises inevitably change the underlying fundamentals of the remaining cells in the analysis sample, they do not address the other two dimensions of the bias, namely the degree of idiosyncratic plant heterogeneity and plant size differences, which we separately study next.


Idiosyncratic Dispersion As described in Equation (6), granularity can also reflect large plant-level, idiosyncratic within-location variance, even if cells have relatively many plants. Our first lever to study this source of granularity bias is to vary the winsorization of plant-level TFP, winsorizing by 0% and 2.5% rather than, as in our baseline specification, 1% (at the national level). The results are reported in Column (4) of Table 1 for the 0% winsorization, and in Column (5) for the 2.5% winsorization. While indeed winsorization and hence extreme values of plant-level productivity play a role in the dispersion measures, the effect

¹⁶Serving as another source of mechanical attenuation of the dispersion, the industry demeaning is now conducted at a lower level, perhaps removing a source of compositional dispersion.


¹⁷An alternative route is to redefine industries at lower rather than higher levels of aggregation, or to aggregate to larger geographical units such as the state level.

Figure 4: Varying Cell Thresholds of Plant Counts

(a) Location Effects



(b) Industry-Specific Location Effects



Note: The figure reports the raw variances (solid black line), permutation variances (dashed red line), and mean covariances (dot-dashed blue line) of the US place effects (TFP) for various sample restrictions on the minimum levels of plant counts per location-industry cell. Panel (a) does so for location effects; Panel (b) does so for industry-specific location effects. The red and blue shaded regions report 95% confidence intervals for variances from 1,000 permuted economies and 1,000 split-samples, respectively. Referring to the secondary y-axis, the black dotted lines plot in Panel (a) and (b) the number of MSAs and the number of MSA-industries, respectively.

of our exercise is somewhat limited. Studying variances, abandoning winsorization leads to a small increase from 0.026 to 0.029 for location effects, and 0.011 to 0.012 for industry-specific location effects; raising the threshold to a symmetric 2.5% winsorization leads to

a modest attenuation to 0.022 for location effects and 0.094 for location-specific industry effects. The 90th-10th percentile difference is also quite robust to winsorization.

Turning to the permutation tests, the p -values remain clearly below the 5% threshold, although, if anything, winsorization boosts the statistical significance of the empirical raw variance compared to the randomization benchmark, consistent with granularity bias.

A priori, winsorization should probably lower the bias-corrected variance, if extreme plant values are clustered in specific places and reflect true place specific effects. Indeed, the covariances are 0.009, 0.008, and 0.006 for the location effects for the 0%, 1% and 2.5% winsorizations, and, respectively, 0.026, 0.023 and 0.019 for the location-industry effects. In sum, granularity bias withstands our attenuation strategy of adjusting tails in the national data, even as this strategy itself alters the sample and erodes the true variance.

Large, Dominant Plants The third source of granularity bias is heterogeneity in plant size, by which place effects are weighted. (For exposition, Equation (6) is written with equal plant weights.) Large plants will dominate average-based raw place effects, and may generate much of the spurious dispersion in the permutation tests. To gauge the role of large plants, we also present results that weight each plant equally, with results reported in Column (6) of Table 1. Consistent with the role of large plants in granularity bias, the raw variances fall by around four fifths for the location effects, and slightly more than half for the location-industry effects. The permutation test reveals a tantamount decline (in percent terms) for the mean raw variance of the randomized economies. But the p -values of the empirical raw variance become even lower, falling below 1%, indicating if anything that the role of granularity has not declined. Congruently, a similar scaling down occurs for the bias-corrected variance of true place effects, which falls by three quarters for location effects and by less than two thirds for industry-location effects, compared to the weighted specifications in Column (1). Hence, we conclude that the unweighted specifications yield a similar picture for the share of the raw variance reflecting granularity bias, while scaling down the overall level of dispersion. Again, we caution that this specification check inevitably entails a substantive redefinition of productivity place effects and the underlying (weighted) sample, so that, even if granularity bias had been less pronounced in the unweighted specification, equal weights naturally do not provide a solution if the preferred specification is weighted (for it to be consistent with aggregation, for instance).

Assessment Overall, we conclude that, granularity bias is a robust feature of the data, which dominates the raw variances of average-based place effects, and that potential alternatives to our covariance-split-sample strategy that build on adjusting the core dispersion-

relevant fundamentals of the data run the risk of throwing out the baby with the bathwater.

5 Additional Application I: The Productivity of New Plants

Our main findings have revealed that due to plant idiosyncracies, systematic place effects are considerably less pronounced than raw averages suggest. An interesting specific question is the degree to which *new plants* inherit the place effects of the old, incumbent plants. For instance, a common assumption is that spatial TFP differences across places are a fixed property of the place-specific production function that would also determine productivity “at the margin” for counterfactual input reallocations (e.g., Desmet and Rossi-Hansberg, 2013; Hsieh and Moretti, 2019). To the degree that the extensive margin (new plants) absorb reallocated inputs (rather than incumbent plants scaling up or down), our strategy can quantify the degree to which TFP effects indeed carry over to new plants. As one extreme, if new plants’ productivity is considerably more compressed across places, or unrelated to that of incumbent plants, such counterfactual reallocation would not generate the gains from reallocation implied by the productivity average of existing plants.

A priori, older and larger incumbent plants dominate the pooled averages constructed in Section 4, leaving room for the place effects for new plants to differ. In 2012, plants older than five years made up 78% of manufacturing plants, and 91% of total employment in the manufacturing sector (US Census Business Dynamics Statistics, 2014 Release). On the theoretical side, in models of embodied technological change (as in, e.g., Sakellaris and Wilson, 2004) new projects reflect the frontier technology while incumbent, old projects reflect legacy technologies, so that measured place averages may reflect age composition, or place effects would show up among new plants. Some models (e.g., Duranton and Puga, 2001) specifically predict that some cities are better environments for entrepreneurship than others.¹⁸ The entry and location choices of new plants can also be considered revealed-preference proxies for productivity differences (Henderson, 1994).

Strategy We partition the population of plants into new plants aged five years and younger (superscript Y for “young,” not N , which we use for counts), and old plants (O) aged six years and older. To implement our split-sample method, we now require at least two plants per age group per location-industry cell, rather than two plants of any age, as before. This restriction has limited effects on dispersion statistics for the pooled (i.e., not age-specific) sample, which we report in Table 1 Column (8). Compared to the baseline

¹⁸See Foster, Haltiwanger, and Syverson (2016) and Pugsley, Sedlacek, and Sterk (forthcoming) for dedicated studies of the growth dynamics and selection of new plants.

sample requiring two plants of any age (Column (1)), the new restrictions leave the raw variance of the location effects (Panel A) fairly stable and essentially leaves the covariance unaffected, despite dropping of around a quarter of MSAs. For the industry-specific location effects in Panel B, the restriction drops around three quarters of cells. The raw variance drops by around 40%, but with a much smaller drop (a quarter) in the covariance, suggesting that the restriction drops particularly small and noisy location-industry cells.

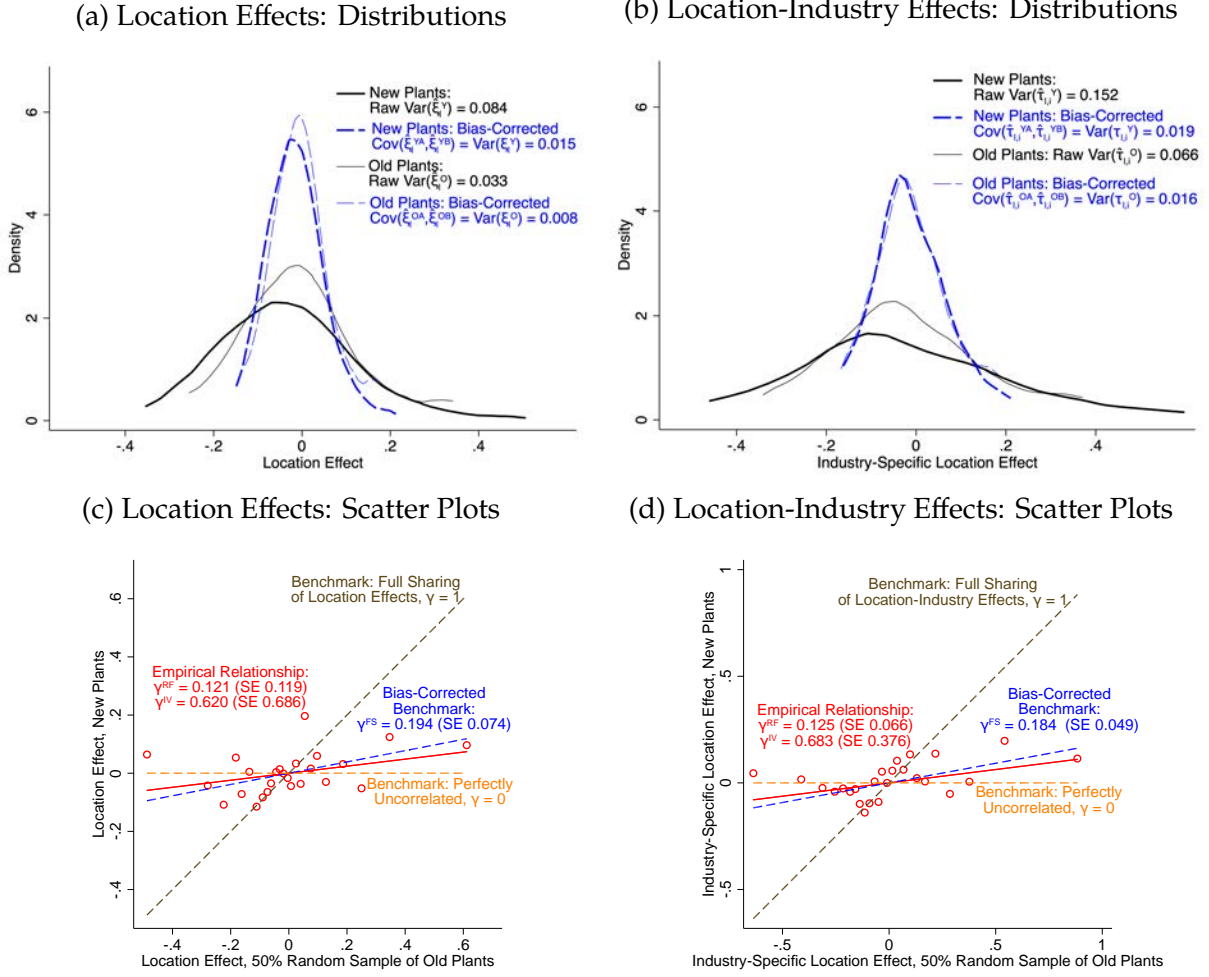
Dispersion Figure 5 Panels (a) and (b) plot, in thick black lines, the distribution of raw location effects for new plants, $\widehat{\xi}_l^Y$, and, respectively, location-industry effects, $\widehat{\tau}_{l,i}^Y$. Table 1 reports the full set of dispersion statistics for the new plants only (Column (9)). In thinner black lines, the figure plots the distributions of old plants' place effects, for which Table 1 Column (10) reports the dispersion statistics for the sample of old plants. In these two panels, as well as the dispersion statistics reported, we now weight location-industry effects by the industry's employment share in total employment in the location with both numbers separately computed for the old and new plants. We continue to not weight across locations.

At $\text{Var}(\widehat{\xi}_l^Y) = 0.084$, location effects of new plants are two-and-a-half times as dispersed as the old place location effects ($\text{Var}(\widehat{\xi}_l^O) = 0.033$), which in turn appear to dominate the pooled sample's raw variance (0.032, see Table 1 Column (8) Panel A, which reports dispersion statistics for the pooled sample). The new plants' location-industry effects are also two-and-a-half times more dispersed at $\text{Var}(\widehat{\tau}_{l,i}^Y) = 0.152$ compared to the (pooled) raw variance of 0.064, in turn again close to those for the old plants' place effects ($\text{Var}(\widehat{\tau}_{l,i}^O) = 0.066$). For the new plants, Table 1 Column (9) also reveal that the top 90th to 10th spread of location effects is 0.744, and 1.11 for location-industry effects.

Hence, taking the raw variances at face value, new plants appear dramatically more dispersed in their productivity than old plants or as would be suggested when pooling all plants. This increased dispersion would imply, for instance, that place matters much more for the productivity of marginal projects, or entrepreneurship, than would be suggested by a standard pooled measure, and that, potentially, forces leading to this dispersion, such as sorting or agglomeration forces, might be even more pronounced for such new projects.

However, much of this higher variance of the new plants' place effects may simply reflect heightened granularity bias, due to smaller populations and potentially even higher idiosyncratic dispersion in true or measured TFP. Our split-sample strategy permits us to again remove this bias, and to isolate the variance of true place effects for new plants. Indeed, the bias-corrected variances of the new plants drop dramatically to 0.015 for location effects and to 0.019 for location-industry effects. These corrections for new plants'

Figure 5: Place Effects of New vs. Old Plants



Note: The figure presents place effects for new and old plants. The top panels are kernel density plots representing the distribution of estimated location effects (Panel (a)) and industry-specific location effects (Panel (b)). In each panel, the solid black line plots the empirical raw distribution of the place averages. The blue dashed line illustrates the biased-corrected distribution, and reflects a mean-preserving variance-adjustment of the raw distribution using a linear transformation as described in Footnote 12. The thick pair of lines refers to place effects of new plants (five years and younger); the thin pair of lines refers to the effects of remaining plants, i.e., those older than five years. In accordance with Census disclosure requirements, Panels (a) and (b) present kernel densities rather than histograms and censors the density plots at the 5th and 95th percentiles. The bottom panels are scatter plots of the average-based location effects (Panel (c)) and industry-specific location effects (Panel (d)), with new plants on the y-axis and the old plants' place effects in the same place on the x-axis. The graph bins the underlying observations into 25 equally sized bins, but the regressions are run on the underlying observations. The red line represents the linear regression slope γ^{RF} , which is the reduced-form effect in the IV interpretation of the split-sample method of new against old plants' place effects. We also report the IV effect $\gamma^{IV} = \gamma^{RF} / \gamma^{FS}$ drawing on first-stage effect γ^{FS} , which we describe next. The blue line plots the regression slope γ^{FS} of the first stage, which is obtained by regressing the half sample of old plants' place effects on the y-axis with the place effects estimated in the complementary set of old plants on the x-axis. We also plot two benchmarks. $\gamma = 0$ represents no relationship between place effects of old and new plants; $\gamma = 1$ represents the naive effect assuming no attenuation bias from measurement error such as granularity. The appropriate comparison for full sharing of place effects is the first-stage slope. Intuitively, the IV effect measures the distance to that corrected benchmark of full sharing. Panels (a) and (c) are not weighted; Panels (b) and (d) weight location-industry effects by the industry's employment shares within the location.

place effects entail much larger reductions from the corresponding raw variances than for the pooled samples' place effects, while still leaving the dispersion of true place effects higher than that of the old plants (which exhibit bias-corrected variances of 0.008 and 0.016

for location and industry-specific effects, respectively).¹⁹

Are the Places that Appear Productive for Old Plants also More Productive For New Plants? True place effects for new and old plants may be distinct. For example, a nursery cities view (Duranton and Puga, 2001) would permit some cities to be particularly suitable for entrepreneurship, in ways that need not carry over to incumbent, large and old production units. Alternatively, place-based productivity differences could be entirely cohort-specific.²⁰

Augmenting our split sample method, we investigate this question. In Figure 5 Panels (c) (location effects) and (d) (location-industry effects), we juxtapose the new-plants place effect of a given place with the corresponding place effect of old plants only.²¹ As one benchmark, we plot a slope of one: place effects would show up for both new and old projects identically. As another benchmark, we plot a slope of zero, which would indicate no relationship between new and old plants' respective place effects.

The red hollow scatter points trace out the empirical relationship between the new and old place effects. They suggest that the place effects for new plants are much closer to a no-correlation benchmark than the 45 degree line. We also include, as a red solid line, the estimated linear regression slope γ^{RF} , which is the reduced-form effect in the IV interpretation of the split-sample method we describe below. The slopes reveal a small (and imprecisely estimated) elasticity of 0.121 (SE 0.119) and 0.125 (SE 0.066) for location and location-industry effects of new plants' to old plants' place effects.

However, the unity benchmark is inappropriate due to granularity bias, which shows up as attenuation bias in the regression estimate. To construct a bias-corrected benchmark, we implement a formal instrumental variables (IV) approach. We estimate a first stage using the split samples, regressing the place averages of a random half sample of old plants indexed by (O, B) (y-axis) on those on the other sample of old plants (O, A) (x-axis). (In fact, the aforementioned slope between new and old plants used that half sample on the x-axis rather than all old plants.) The blue line plots the resulting first-stage regression slope

¹⁹Because TFP may be especially noisily measured for new plants, Appendix Table A.1 Columns (7)-(9) additionally present the analogous results for log value added per worker.

²⁰A motivating finding, reported in Table 1 Columns (9) and (10), is that while location effects for the new remain significantly dispersed over a random location benchmark ($p = 0.022$) and location-industry benchmarks ($p = 0.019$), we can no longer reject that old plants differ from random allocations—an intriguing result that suggests that places are considerably more similar for old plants. However, since the permutation test carries over place effects (and the covariance adjustment yielded economically large biased-corrected variances for the old plants, with the confidence interval excluding zero), the results overall point towards significant place effects for the old plants too.

²¹For this exercise depicted in Panels (c) and (d), we now weight location-industries by their pooled employment share. For consistency with Panels (a) and (b) as well as Table 1 Columns (9) and (10), we demean location effects by industry-age group.

γ^{FS} , which provides benchmarks of 0.194 (SE 0.074) for location effects and 0.184 (SE 0.049) for industry-specific location effects, hence far from one.²² This first-stage relationship is analogous to the visualization of the overall bias correction in the full sample in Figure 3. With this bias-corrected benchmark, the new plants appear to inherit a larger—but far from perfect—share of the place effects of the old. The formal IV effect $\gamma^{\text{IV}} = \gamma^{\text{RF}}/\gamma^{\text{FS}}$ is 0.620 (SE 0.686) for the location effects and 0.683 (SE 0.376) for location-industry effects. Intuitively, the IV effect measures the distance of the reduced form effect from the first stage, i.e., the corrected benchmark of full sharing.

To generate the figure, we generate 1,000 split-sample economies, and then select, for visualization in the scatter plot, two economies with first-stage and reduced-form coefficients being closest to the mean coefficients across the 1,000 economies (putting twice as high a penalty on the error in the first-stage coefficients). Across the 1,000 split samples, the mean estimated first-stage, reduced-form, and IV coefficients are, respectively, 0.194, 0.113 and 0.681 for the location effects, and 0.183, 0.127 and 0.716 for the location-industry effects.

We conclude that while place effects for new plants appear to comove by about 62–68% (using the point estimates), with those of old plants, there is a substantial degree of independent variation in the new plants’ place effect, with estimates having wide confidence intervals.

6 Additional Application II: Labor Productivity

We complement our primary productivity measure, log TFP, with log value added per worker (labor productivity). This measure is less demanding, requiring neither specifying a production function nor comprehensive input measures. We restrict our sample to those plants that have TFP defined, so the samples of the TFP and labor productivity analyses are identical. As a benchmark, with Cobb-Douglas production, the dispersion of log labor productivity corresponds to that of the log marginal product of labor (which equals the log labor share in value added—netted out by the industry-level labor shares—plus the plant’s log value added per worker).

Across places, dispersion in marginal product of labor can indicate spatial misallocation (see Hsieh and Moretti, 2019, who study average wages to proxy for raw average labor productivity, but do not directly measure TFP or MPL across places). While a benchmark

²²The F -statistics for the first-stage regressions in Panels (c) and (d) are 7.0 and 13.8, respectively, which, besides the attenuated level of the IV coefficient below one, provides another caveat to the interpretation of a strong relationship between old and new plants. The F -tests are quite dispersed across the 1,000 simulated economies.

model of perfectly competitive factor markets would predict marginal products to be equalized within a market, which would suggest little room for idiosyncratic dispersion and hence for granularity bias therein, such idiosyncratic within-market between-firm dispersion can emerge with frictions and indicate another dimension of misallocation (Hsieh and Klenow, 2009).²³

Column (7) of Table 1 replicates our main specification using labor productivity in place of TFP. Appendix B contains replications of other main exhibits. In contrast to the prediction that marginal products are more compressed than TFP, we find that raw variances increase compared to those of TFP, by about 150%, from 0.026 to 0.063 and from 0.110 to 0.290 for location and location-industry effects. Given the tantamount increase in permutation variances, p -values remain similar. Finally, the bias-corrected variance estimates increase by about 100%, from 0.008 to 0.016 for location effects and from 0.023 to 0.041 for location-industry effects.

We conclude that the dispersion in true place effects is if anything more pronounced for marginal products than for TFP. Moreover, the relative share of the raw variance that reflects granularity bias is in fact somewhat higher: three quarters (rather than two thirds) for location effects, and sixth sevenths (rather than four fifths) for location-industry effects.²⁴

7 Additional Application III: The Countries of Europe

We close our empirical study by applying our analysis to the within-country, cross-regional dispersion of 15 European countries. We draw on internationally comparable firm-level data from Bureau van Dijk (BvD), construct firm-level TFP measures, as described in Section 3, and study regional dispersion among the NUTS-2 regions, which most closely correspond to US MSAs.

Figure 6 and Appendix Tables A.2 (TFP) report dispersion statistics country by country. The figure recapitulates the US Census based findings as the leftmost entry. For each country, it reports three dispersion statistics: first, the raw variance (solid black circles), second, the mean raw variance implied by 1,000 random allocations of plants across places (hollow red diamonds) along with the 95% confidence intervals (dashed red lines) taken from the sampling distribution given by the 1,000 randomizations, and, third, the

²³For misallocation measures, weighting MSAs by plant or employment may be an alternative specification of interest; for consistency with the TFP approach, we do not weight MSAs by size here either.

²⁴The increased dispersion of labor productivity place averages (and effects) compared to TFP-based dispersion measures is consistent with the studies of within-industry, between-plant dispersion in Syverson (2004); Cunningham, Foster, Grim, Haltiwanger, Pablonia, Stewart, and Wolf (2020).

mean covariance (blue triangles)—the bias-corrected estimate of the variance—of the 1,000 randomly split samples along with 2.5% and 97.5% confidence intervals implied by the sampling distribution (solid blue lines).

Consistent with the US findings, the European applications exhibit large—and quite heterogeneous—raw variances of average-based place effects. The raw variance of the location effects range from 0.005 for Austria to high values of around 0.024, 0.029, 0.037, and 0.053 for Norway, the United Kingdom, Italy, and Germany. It is tempting to ascribe these high raw variances to intuitive regional divergences in those countries (the high-unemployment regions in the North in Norway, the productive urban centers in the United Kingdom, the South-North gap in Italy, the East-West division in Germany). Yet, in the United Kingdom, random allocations would have yielded similar dispersion for location effects, which is the case for 7 countries (Austria, Denmark, France, Hungary, Spain, Sweden, and the United Kingdom). Encouragingly however, that null hypothesis of no-more-than-random productivity dispersion can be rejected at the 5% level in many but not all of those aforementioned countries with anecdotal divergences and high raw variances (the full list is Bulgaria, the Czech Republic, Germany, Italy, Norway, and Romania). Poland and Portugal exhibit marginally significant results.

The biased-corrected variances are lower than the raw variances in all cases, but their ratios vary substantially. The covariance is only a quarter lower in Bulgaria, Norway, Romania, Portugal, and Italy, compared to a 70% difference for Austria, France, the United Kingdom.


For the industry-specific location effects, granularity bias is amplified, as with the US data. As with the US data, the raw variances are dramatically larger, by an order of magnitude, but the bias-corrected estimates using the split-sample covariances settle in at quite similar (but generally slightly higher) levels to the location effects. However, reflecting the heightened granularity bias, the p -values of the permutation tests reveal that the industry-specific location effects are only insignificantly different from the random allocation benchmark in all countries but Germany and Romania.

We end our assessment of the cross-country context by noting the small amount of locations (NUTS-2 regions), which range between 5 and 42, industry-location cells, ranging from 132 to 717, and firms, ranging from 1,678 to 132,244 (reported in Appendix Table A.2).²⁵ Appendix Table A.3 and Appendix Figure A.4 report the results for value added per worker. The value added results, much like in the US context, in many instances double the dispersion statistics, although again with considerable heterogeneity.


²⁵We have also probed robustness to the, considerably smaller, NUTS-3 regional divisions, for which location counts increase but within-location firm counts drop, such that granularity bias would grow further.

Figure 6: Spatial Dispersion of TFP: United States and European Countries

(a) Location Effects



(b) Industry-Specific Location Effects



Note: The figure reports, for each country, the raw variance (solid black dot), the mean and 95% confidence intervals of the raw variances of 1,000 permutations, i.e., random allocations of firms across location-industry cells (hollow red diamond and dotted line, respectively), and the mean and 95% confidence intervals of the bias-corrected variances from 1,000 split samples (solid blue triangle and line, respectively). Panel (a) reports the location (NUTS-2) effects and Panel (b) reports the location (NUTS-2) \times industry (2-digit NACE) effects for European countries; the leftmost entry reiterates the US numbers using the MSA and 4-digit NAICS cell definitions. The place effects for the European countries are measured using Bureau van Dijk (BvD) firm-level data. Appendix Table A.2 details the statistics printed here.

We tentatively conclude that granularity leads to large—if anything more dramatic—upward bias in productivity differences across place in the 15 European countries we can study in internationally comparable firm-level micro data. Obvious challenges are the lower cell-level firm observation counts in the internationally comparable BvD data, and the series of data quality issues associated with the BvD data, which we have discussed in Section 3.

8 Conclusion

We have revisited the measurement of the dispersion in productivity across places. We have confirmed the high raw variance of productivity averages across cities, which has been a frequent focus of research in urban economics. We have investigated the possibility that some of this dispersion might reflect the “luck of the draw” of tremendously heterogeneous plants that happen to be located in a given location, especially when zooming into fine-grained industries. Indeed, we show that randomly allocating plants across places would generate nearly as much dispersion as the empirical economy—the dispersion of which however does differ statistically significantly from what random economies could have generated. We have also provided a bias-corrected estimate of the variation in true place effects, i.e., those differences in productivity that are systematically related to the place in which plants are located. We have found that the majority—around two thirds to four fifths—of the raw variation is an artifact of granular data: cells of finite plants with high idiosyncratic variation in productivity and size. This pattern holds when measuring productivity as TFP as well as log of average value added per worker, in US plant-level data and in firm-level data of 15 European countries. Furthermore, we uncover substantially more, and independent, variation in location effects measured from new plants, implying that place effects may not perfectly carry over to new plants.

Importantly, by conducting within-industry comparisons across locations and demeaning plant productivity by national industry, our location measures remove industry composition as a factor in place effects (even when we aggregate location-industry effects into location effects). An interesting question beyond the scope of our paper is whether industry composition and location choice itself may be a reflection of a broader notion of place effects on productivity.

We close by reiterating that our bias-corrected numbers are model-agnostic and reduced-form. They may reflect a combination of causal effects of place on productivity (such as agglomeration forces), as well as sorting. Moreover, our method does not correct for even more fundamental sources of measurement error, for example arising from location-specific production function misspecification or mismeasurement of inputs or prices. Hence, the remaining true dispersion in place effects may be even lower than our bias-corrected estimates. Overall, our results suggest a substantially reduced upper bound for the potential productivity impact of the actual underlying spatial forces, to the degree that they show up through the kind of true place effects we study, namely place-specific expected values of plant productivity.²⁶

²⁶In a separate project, we are applying our strategy to measuring the between-firm dispersion in plant

References

- Albouy, David. 2016. "What Are Cities Worth? Land Rents, Local Productivity, and the Total Value of Amenities." *Review of Economics and Statistics* 98 (3):477–487.
- Bils, Mark, Peter Klenow, and Cian Ruane. 2020. "Misallocation or Mismeasurement?" *NBER Working Paper* .
- Boeri, Tito, Andrea Ichino, Enrico Moretti, and Johanna Posch. 2019. "Wage Equalization and Regional Misallocation: Evidence from Italian and German Provinces." *CEPR Discussion Paper* .
- Caliendo, Lorenzo, Fernando Parro, Esteban Rossi-Hansberg, and Pierre-Daniel Sarte. 2018. "The Impact of Regional and Sectoral Productivity Changes on the US Economy." *The Review of Economic Studies* 85 (4):2042–2096.
- Chandra, Amitabh, Amy Finkelstein, Adam Sacarny, and Chad Syverson. 2016. "Health Care Exceptionalism? Performance and Allocation in the US Health Care Sector." *American Economic Review* 106 (8):2110–44.
- Ciccone, Antonio and Robert Hall. 1996. "Productivity and the Density of Economic Activity." *The American Economic Review* 86 (1):54.
- Combes, Pierre-Philippe, Gilles Duranton, Laurent Gobillon, Diego Puga, and Sébastien Roux. 2012. "The Productivity Advantages of Large Cities: Distinguishing Agglomeration from Firm Selection." *Econometrica* 80 (6):2543–2594.
- Combes, Pierre-Philippe, Gilles Duranton, Laurent Gobillon, and Sébastien Roux. 2010. "Estimating Agglomeration Economies with History, Geology, and Worker Effects." In *Agglomeration Economics*. University of Chicago Press, 15–66.
- Cunningham, Cindy, Lucia Foster, Cheryl Grim, John Haltiwanger, Sabrina Wulff Pabilonia, Jay Stewart, and Zoltan Wolf. 2020. "Dispersion in Dispersion: Measuring Establishment-level Differences in Productivity." *BLS Working Paper* 530 .
- Davis, Donald and Jonathan Dingel. 2020. "The Comparative Advantage of Cities." *Journal of International Economics* 123:103291.
- Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda. forthcoming. "Changing Business Dynamism and Productivity: Shocks vs. Responsiveness." *The American Economic Review* .
- Dekle, Robert and Jonathan Eaton. 1999. "Agglomeration and Land Rents: Evidence from the Prefectures." *Journal of Urban Economics* 46 (2):200–214.

productivity studying multi-plant firms rather than locations (for a study of between-plant dispersion, see Kehrig and Vincent, 2020, who however do not account for measurement error using a split-sample method), a topic of active research in the misallocation literature (Bils, Klenow, and Ruane, 2020; Rotemberg and White, 2020; Haltiwanger, Kulick, and Syverson, 2018).

- Desmet, Klaus and Esteban Rossi-Hansberg. 2013. "Urban Accounting and Welfare." *American Economic Review* 103 (6):2296–2327.
- Dingel, Jonathan and Felix Tintelnot. 2020. "Spatial Economics for Granular Settings." *NBER Working Paper* .
- Drenik, Andres, Simon Jäger, Pascuel Plotkin, and Benjamin Schoefer. 2020. "Paying Outsourced Labor: Direct Evidence from Linked Temp Agency-Worker-Client Data." *NBER Working Paper* .
- Duranton, Gilles and Diego Puga. 2001. "Nursery Cities: Urban Diversity, Process Innovation, and the Life Cycle of Products." *American Economic Review* 91 (5):1454–1477.
- Ehrlich, Maximilian and Henry Overman. 2020. "Place-based Policies and Spatial Disparities Across European Cities." *Journal of Economic Perspectives* 34 (3):128–49.
- Ellison, Glenn and Edward Glaeser. 1997. "Geographic Concentration in US Manufacturing Industries: a Dartboard Approach." *Journal of Political Economy* 105 (5):889–927.
- Ellison, Glenn, Edward Glaeser, and William Kerr. 2010. "What Causes Industry Agglomeration? Evidence from Coagglomeration Patterns." *American Economic Review* 100 (3):1195–1213.
- Fajgelbaum, Pablo and Cecile Gaubert. 2020. "Optimal Spatial Policies, Geography, and Sorting." *The Quarterly Journal of Economics* 135 (2):959–1036.
- Fajgelbaum, Pablo, Eduardo Morales, Juan Carlos Suárez Serrato, and Owen Zidar. 2019. "State Taxes and Spatial Misallocation." *The Review of Economic Studies* 86 (1):333–376.
- Fort, Teresa and Shawn Klimek. 2018. "The Effects of Industry Classification Changes on US Employment Composition." *US Census Bureau Center for Economic Studies Paper No. CES-WP-18-28* .
- Foster, Lucia, Cheryl Grim, and John Haltiwanger. 2016. "Reallocation in the Great Recession: Cleansing or Not?" *Journal of Labor Economics* 34 (S1):S293–S331.
- Foster, Lucia, John Haltiwanger, and Cornell John Krizan. 2001. "Aggregate Productivity Growth: Lessons from Microeconomic Evidence." In *New Developments in Productivity Analysis*. University of Chicago Press, 303–372.
- Foster, Lucia, John Haltiwanger, and Chad Syverson. 2016. "The Slow Growth of New Plants: Learning about Demand?" *Economica* 83 (329):91–129.
- Gabaix, Xavier. 1999a. "Zipf's Law and the Growth of Cities." *American Economic Review* 89 (2):129–132.
- . 1999b. "Zipf's Law for Cities: an Explanation." *The Quarterly Journal of Economics* 114 (3):739–767.


- Gaubert, Cecile. 2018. "Firm Sorting and Agglomeration." *American Economic Review* 108 (11):3117–53.
- Gaubert, Cecile and Oleg Itskhoki. forthcoming. "Granular Comparative Advantage." *Journal of Political Economy* .
- Gerard, François, Lorenzo Lagos, Edson Severnini, and David Card. 2018. "Assortative Matching or Exclusionary Hiring? The Impact of Firm Policies on Racial Wage Differences in Brazil." *NBER Working Paper* .
- Glaeser, Edward, Hedi Kallal, Jose Scheinkman, and Andrei Shleifer. 1992. "Growth in Cities." *Journal of Political Economy* 100 (6):1126–1152.
- Glaeser, Edward and David Mare. 2001. "Cities and Skills." *Journal of Labor Economics* 19 (2):316–342.
- Gopinath, Gita, Şebnem Kalemli-Özcan, Loukas Karabarbounis, and Carolina Villegas-Sanchez. 2017. "Capital Allocation and Productivity in South Europe." *The Quarterly Journal of Economics* 132 (4):1915–1967.
- Greenstone, Michael, Richard Hornbeck, and Enrico Moretti. 2010. "Identifying Agglomeration Spillovers: Evidence from Winners and Losers of Large Plant Openings." *Journal of Political Economy* 118 (3):536–598.
- Haltiwanger, John, Robert Kulick, and Chad Syverson. 2018. "Misallocation Measures: The Distortion that Ate the Residual." *NBER Working Paper* .
- Henderson, Vernon. 1994. "Where Does an Industry Locate?" *Journal of Urban Economics* 35 (1):83–104.
- Hornbeck, Richard and Enrico Moretti. 2018. "Who Benefits from Productivity Growth? Direct and Indirect Effects of Local TFP Growth on Wages, Rents, and Inequality." *NBER Working Paper* .
- Hsieh, Chang-Tai and Peter J Klenow. 2009. "Misallocation and Manufacturing TFP in China and India." *The Quarterly Journal of Economics* 124 (4):1403–1448.
- Hsieh, Chang-Tai and Enrico Moretti. 2019. "Housing Constraints and Spatial Misallocation." *American Economic Journal: Macroeconomics* 11 (2):1–39.
- Jäger, Simon, Benjamin Schoefer, and Jörg Heining. 2019. "Labor in the Boardroom." *NBER Working Paper w26519* .
- Kehrig, Matthias. 2015. "The Cyclical Nature of the Productivity Distribution." *US Census Bureau Center for Economic Studies Paper No. CES-WP-11-15* .
- Kehrig, Matthias and Nicolas Vincent. 2020. "Good Dispersion, Bad Dispersion." *NBER Working Paper* .

- Kline, Patrick, Raffaele Saggio, and Mikkel Sølvsten. 2020. "Leave-out Estimation of Variance Components." *Econometrica* 88 (5):1859–1898.
- Melitz, Marc. 2003. "The Impact of Trade on Intra-industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71 (6):1695–1725.
- Moretti, Enrico. 2012. *The New Geography of Jobs*. Houghton Mifflin Harcourt.
- Pugsley, Benjamin, Petr Sedlacek, and Vincent Sterk. forthcoming. "The Nature of Firm Growth." *American Economic Review* .
- Rotemberg, Martin and Kirk White. 2020. "Measuring cross-country Differences in Misallocation." *NYU Working Paper* .
- Sakellaris, Plutarchos and Daniel Wilson. 2004. "Quantifying Embodied Technological Change." *Review of Economic Dynamics* 7 (1):1–26.
- Silver, David. 2020. "Haste or Waste? Peer Pressure and Productivity in the Emergency Department." *Working Paper* .
- Sveikauskas, Leo. 1975. "The Productivity of Cities." *The Quarterly Journal of Economics* 89 (3):393–413.
- Syverson, Chad. 2004. "Product Substitutability and Productivity Dispersion." *Review of Economics and Statistics* 86 (2):534–550.
- . 2011. "What Determines Productivity?" *Journal of Economic Literature* 49 (2):326–65.

**Online Appendix of
Productivity, Place, and Plants: Revisiting the Measurement
Benjamin Schoefer and Oren Ziv**

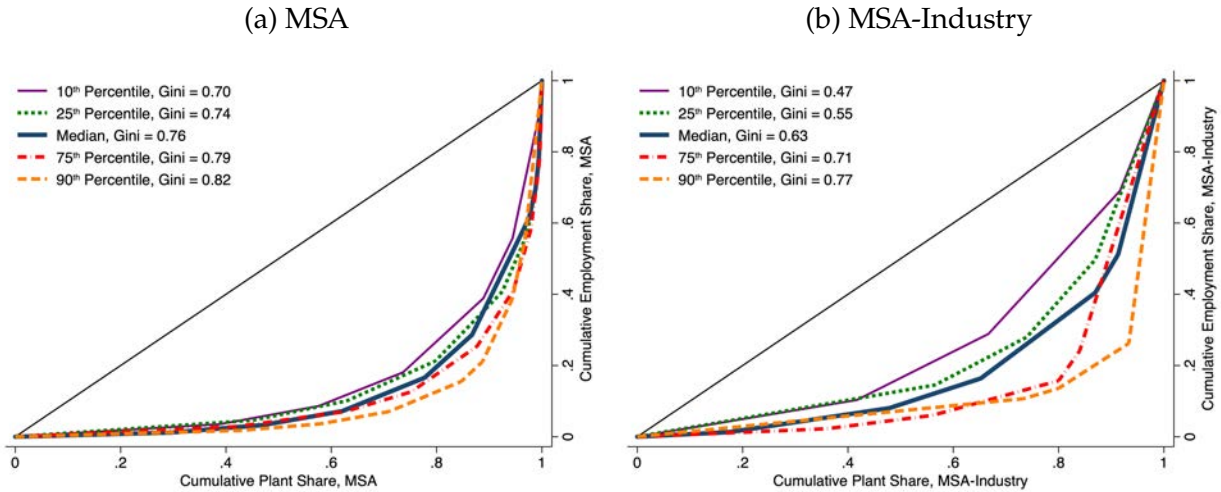
A Supplementary Facts: Distributions of Plants and Employment

Figure A.1: Plant Counts in U.S. MSAs and MSA-Industry Cells in Manufacturing (Public-Use Data)



Note: The figure presents data from the 2012 US County Business Patterns (CBP) files on the plant counts in manufacturing at the MSA and MSA-industry level. Panels (a) and (b) present the cumulative distribution of plant counts for US MSAs and MSA \times 4-digit NAICS cells, respectively. 37% of location-industry cells have zero plants.

Figure A.2: Concentration of Employment in U.S. MSAs and MSA-Industry Cells in Manufacturing (Public-Use Data)



Note: The figure presents data from the 2012 US County Business Patterns (CBP) files on employment concentration in manufacturing at the MSA and MSA-industry level. Both panels restrict samples to cells with at least 10 plants, which drops zero MSAs and 75% of MSA-industry cells, and thereby zooms into MSA-industry cells less likely to be concentrated and less likely to be subject to granularity bias. Panels (a) and (b) report Lorenz curves and Gini coefficients for, respectively MSAs and MSA-industries with Gini coefficients at the 90th, 75th, 50th, 25th, and 10th percentiles of their respective distributions. For these panels, plant-level employment is imputed as follows. The CBP reports plant counts by employment size bins (sizes range from plants with 1-5 employees to the top bin, plants with 5000 or more employees) for all MSA-industry cells. The CBP also reports total MSA-industry employment, which is censored for some MSA-industries to avoid disclosing data for individual companies or for data quality reasons. In addition, some numbers include noise. First, focusing only on MSA-industries where employment is not censored and where there are no plants in the top bin, we impute employment for each plant. For a plant in a bin ranging from employment size E_B to E_T , we impute plants' employment as $E_B + R \cdot (E_T - E_B)$, where R is an economy-wide constant between 0 and 1 chosen such that, in a regression of \log actual MSA-industry total employment on \log imputed MSA-industry total employment, the coefficient is closest to 1 and the intercept is closest to zero. Next, we use R (which we estimate to be 0.35) to impute plants' employment in MSA-industries with censored employment according to the same formula for all plants except plants in the top bin. For the remaining plants, those in the top bin, where MSA-industry employment is censored, we set their employment to 5,000. Where MSA-industry total employment is not censored, we estimate employment as the maximum of 5,000 and the difference between actual total employment and imputed employment outside of the top bin, divided by the number of top-bin plants. Because we impute all plants in the same bin to have equal employment, we likely understate the concentration of employment in an MSA or MSA-industry.

B Replication of Main Results with Value Added Per Worker


Table A.1: Spatial Dispersion of Labor Productivity in the United States (Value Added Per Worker) (x100 For All Dispersion Statistics)

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
|---|---------|------------|--------------|---------|------------|---------------|-----------|--------|--------|
| | Main | ≥10 Plants | 6-Digit Ind. | Unwins. | 2.5% Wins. | Plant Weights | New & Old | New | Old |
| Panel A: Location Effects | | | | | | | | | |
| Empirical: | | | | | | | | | |
| Raw $\text{Var}(\hat{\xi}_i)$ | 6.32 | 5.16 | 7.30 | 6.92 | 5.26 | 1.76 | 8.45 | 15.11 | 10.05 |
| 90th – 10th Percentile | 85.20 | 72.55 | 87.70 | 87.69 | 79.50 | 42.74 | 100.3 | 138.3 | 108.3 |
| Permutations: $\text{Var}(\hat{\xi}_i)$ | | | | | | | | | |
| Mean | 4.62 | 3.79 | 5.09 | 5.34 | 3.85 | 0.76 | 6.77 | 16.49 | 8.55 |
| Standard Deviation | 0.61 | 0.57 | 0.68 | 0.79 | 0.48 | 0.11 | 0.93 | 1.94 | 1.24 |
| p -value | 0.009 | 0.021 | 0.004 | 0.042 | 0.010 | 0.001 | 0.048 | 0.755 | 0.113 |
| Bias-Corrected: $\text{Cov}(\hat{\xi}_i^A, \hat{\xi}_i^B) = \text{Var}(\xi_i)$ | | | | | | | | | |
| Mean | 1.58 | 1.80 | 2.61 | 1.73 | 1.39 | 1.02 | 1.98 | 1.64 | 1.58 |
| 97.5th Percentile | 1.97 | 2.41 | 3.00 | 2.19 | 1.74 | 1.19 | 2.77 | 2.94 | 2.67 |
| 2.5th Percentile | 1.17 | 1.11 | 2.16 | 1.22 | 1.02 | 0.84 | 1.08 | 0.23 | 0.43 |
| Panel B: Industry-Specific Location Effects | | | | | | | | | |
| Empirical: | | | | | | | | | |
| Raw $\text{Var}(\hat{\tau}_{i,i})$ | 28.96 | 11.29 | 28.78 | 32.82 | 24.40 | 13.17 | 16.33 | 32.21 | 18.21 |
| 90th – 10th Percentile | 188.9 | 105.8 | 187.2 | 198.9 | 175.9 | 127.3 | 135.8 | 205.0 | 142.9 |
| Permutations: $\text{Var}(\hat{\tau}_{i,i})$ | | | | | | | | | |
| Mean | 26.75 | 9.22 | 25.90 | 31.03 | 22.64 | 10.97 | 13.84 | 33.80 | 16.39 |
| Standard Deviation | 1.12 | 0.69 | 1.14 | 1.59 | 0.84 | 0.30 | 1.13 | 2.21 | 1.42 |
| p -value | 0.029 | 0.004 | 0.008 | 0.132 | 0.021 | 0.001 | 0.029 | 0.755 | 0.092 |
| Bias-Corrected: $\text{Cov}(\hat{\tau}_{i,i}^A, \hat{\tau}_{i,i}^B) = \text{Var}(\tau_{i,i})$ | | | | | | | | | |
| Mean | 4.05 | 3.00 | 5.39 | 4.64 | 3.43 | 2.60 | 3.60 | 3.64 | 2.42 |
| 97.5th Percentile | 4.84 | 3.85 | 5.97 | 5.56 | 4.06 | 2.88 | 4.58 | 5.26 | 3.65 |
| 2.5th Percentile | 3.24 | 2.11 | 4.77 | 3.72 | 2.77 | 2.30 | 2.48 | 2.00 | 1.03 |
| Number of MSAs | 400 | 250 | 400 | 400 | 400 | 400 | 300 | 300 | 300 |
| Number of MSA-Industries | 11,500 | 2,800 | 18,000 | 11,500 | 11,500 | 11,500 | 2,800 | 2,800 | 2,800 |
| Number of Plants | 120,000 | 86,000 | 105,000 | 120,000 | 120,000 | 120,000 | 78,000 | 14,000 | 64,000 |


Note: The table repeats the analysis in Table 1 but uses log value added per worker rather than TFP (and skips Column (7) there, which is identical to Column (1) of the table at hand).

Figure A.3: Varying Cell Thresholds of Plant Counts (Value Added per Worker)

(a) Location Effects



(b) Industry-Specific Location Effects



Note: The figure replicates Figure 4 but uses log value added per worker rather than TFP.

C Spatial Dispersion of Productivity Within 15 European Countries (Bureau van Dijk Data)

Table A.2: Spatial Dispersion of TFP in 15 European Countries (x100 For All Dispersion Statistics)

| | Austria | Bulgaria | Czech Rep. | Denmark | France | Germany | Hungary | Italy | Norway | Poland | Portugal | Romania | Spain | Sweden | UK |
|---|---------|----------|------------|---------|--------|---------|---------|---------|--------|--------|----------|---------|--------|--------|--------|
| Panel A: Location Effects | | | | | | | | | | | | | | | |
| Empirical: | | | | | | | | | | | | | | | |
| Raw $\text{Var}(\widehat{\xi}_i)$ | 0.53 | 0.99 | 0.66 | 0.64 | 0.86 | 5.27 | 0.50 | 3.74 | 2.36 | 1.22 | 1.23 | 0.75 | 1.69 | 0.24 | 2.93 |
| 90th – 10th Percentile | 21.73 | 29.05 | 22.33 | 22.23 | 27.35 | 51.09 | 21.50 | 41.51 | 46.26 | 32.10 | 33.55 | 25.85 | 36.02 | 15.09 | 39.83 |
| Permutations: $\text{Var}(\widehat{\xi}_i)$ | | | | | | | | | | | | | | | |
| Mean | 0.42 | 0.22 | 0.23 | 1.06 | 1.90 | 2.23 | 0.47 | 0.95 | 0.61 | 0.64 | 0.44 | 0.19 | 1.63 | 0.38 | 4.32 |
| Standard Deviation | 0.23 | 0.14 | 0.12 | 0.73 | 1.13 | 0.88 | 0.24 | 0.98 | 0.41 | 0.32 | 0.38 | 0.10 | 2.07 | 0.26 | 1.40 |
| p-value | 0.263 | 0.001 | 0.002 | 0.663 | 0.930 | 0.002 | 0.376 | 0.012 | 0.005 | 0.058 | 0.053 | 0.001 | 0.321 | 0.656 | 0.863 |
| Bias-Corrected: $\text{Cov}(\widehat{\xi}_i^A, \widehat{\xi}_i^B) = \text{Var}(\xi_i)$ | | | | | | | | | | | | | | | |
| Mean | 0.12 | 0.92 | 0.39 | 0.26 | 0.25 | 2.14 | 0.24 | 2.96 | 1.80 | 0.56 | 0.96 | 0.58 | 1.23 | 0.16 | 0.82 |
| 97.5th Percentile | 0.59 | 1.14 | 0.61 | 0.99 | 0.45 | 2.69 | 0.69 | 3.73 | 2.38 | 0.84 | 1.21 | 0.85 | 1.54 | 0.34 | 1.47 |
| 2.5th Percentile | -0.56 | 0.65 | 0.11 | -0.73 | 0.02 | 1.55 | -0.27 | 1.86 | 1.15 | 0.24 | 0.56 | 0.30 | 0.81 | -0.09 | 0.14 |
| Panel B: Industry-Specific Location Effects | | | | | | | | | | | | | | | |
| Empirical: | | | | | | | | | | | | | | | |
| Raw $\text{Var}(\widehat{\tau}_{i,i})$ | 3.98 | 3.72 | 2.86 | 4.31 | 4.97 | 9.55 | 4.48 | 7.92 | 5.70 | 8.57 | 3.07 | 3.81 | 5.11 | 5.16 | 14.66 |
| 90th – 10th Percentile | 37.99 | 26.13 | 45.01 | 54.82 | 34.53 | 56.01 | 35.85 | 45.26 | 33.94 | 53.86 | 36.84 | 31.58 | 42.84 | 25.75 | 49.00 |
| Permutations: $\text{Var}(\widehat{\tau}_{i,i})$ | | | | | | | | | | | | | | | |
| Mean | 4.56 | 3.33 | 2.48 | 5.57 | 7.44 | 7.61 | 4.81 | 6.61 | 7.16 | 7.64 | 3.98 | 2.77 | 6.38 | 4.68 | 16.77 |
| Standard Deviation | 0.78 | 0.53 | 0.36 | 1.49 | 1.72 | 1.09 | 0.76 | 2.06 | 2.86 | 1.63 | 1.31 | 0.55 | 2.94 | 1.24 | 2.67 |
| p-value | 0.779 | 0.239 | 0.140 | 0.825 | 0.989 | 0.042 | 0.644 | 0.182 | 0.625 | 0.238 | 0.746 | 0.050 | 0.608 | 0.222 | 0.777 |
| Bias-Corrected: $\text{Cov}(\widehat{\tau}_{i,i}^A, \widehat{\tau}_{i,i}^B) = \text{Var}(\tau_{i,i})$ | | | | | | | | | | | | | | | |
| Mean | 0.15 | 0.42 | 0.35 | -0.42 | 1.49 | 1.99 | 0.42 | 3.31 | 1.07 | 0.83 | 1.04 | 1.03 | 1.82 | 0.07 | 2.26 |
| 97.5th Percentile | 1.14 | 1.14 | 0.92 | 0.93 | 2.08 | 2.93 | 1.72 | 4.73 | 2.48 | 1.52 | 1.81 | 1.68 | 2.53 | 1.03 | 4.54 |
| 2.5th Percentile | -1.17 | -0.36 | -0.26 | -2.19 | 0.91 | 1.05 | -1.10 | 1.01 | -0.60 | 0.06 | 0.22 | 0.29 | 1.09 | -0.91 | -0.19 |
| Number of NUTS-2 Regions | 9 | 6 | 8 | 5 | 27 | 38 | 8 | 21 | 7 | 17 | 7 | 8 | 19 | 8 | 42 |
| Number of NUTS-2 Region-Industries | 148 | 132 | 170 | 107 | 491 | 717 | 162 | 455 | 135 | 330 | 126 | 176 | 362 | 174 | 731 |
| Number of Firms | 1,678 | 14,625 | 13,091 | 5,829 | 32,050 | 13,244 | 2,768 | 116,918 | 5,189 | 6,650 | 26,019 | 21,279 | 60,375 | 13,916 | 10,125 |

Note: The table reports, for 15 European countries, dispersion measures for location effects (NUTS-2, Panel A) and industry-specific location effects (NUTS-2 x NACE 2-digit, Panel B), using Bureau van Dijk firm-level data. Each column otherwise replicates our main specification in Column (1) of Table 1 for one European country. It reports results for TFP.


Table A.3: Spatial Dispersion of Labor Productivity (Value Added per Worker) in 15 European Countries (x100 For All Dispersion Statistics)

| | Austria | Bulgaria | Czech Rep. | Denmark | France | Germany | Hungary | Italy | Norway | Poland | Portugal | Romania | Spain | Sweden | UK |
|---|---------|----------|------------|---------|--------|---------|---------|---------|--------|--------|----------|---------|--------|--------|--------|
| Panel A: Location Effects | | | | | | | | | | | | | | | |
| Empirical: | | | | | | | | | | | | | | | |
| Raw $\text{Var}(\hat{\xi}_i)$ | 1.26 | 0.70 | 1.05 | 1.08 | 1.54 | 9.19 | 1.65 | 5.58 | 3.90 | 2.61 | 3.13 | 1.38 | 3.27 | 0.12 | 5.10 |
| 90th – 10th Percentile | 39.46 | 23.79 | 32.87 | 26.93 | 31.55 | 73.61 | 41.51 | 56.90 | 64.44 | 37.22 | 49.51 | 41.21 | 51.61 | 10.10 | 55.45 |
| Permutations: $\text{Var}(\hat{\xi}_i)$ | | | | | | | | | | | | | | | |
| Mean | 0.56 | 0.27 | 0.45 | 2.33 | 2.48 | 5.44 | 1.55 | 1.03 | 0.93 | 0.90 | 0.66 | 0.42 | 1.93 | 0.32 | 5.22 |
| Standard Deviation | 0.29 | 0.17 | 0.22 | 1.65 | 1.12 | 1.46 | 0.74 | 1.34 | 0.48 | 0.37 | 0.52 | 0.21 | 1.36 | 0.20 | 1.42 |
| <i>p</i> -value | 0.031 | 0.029 | 0.015 | 0.744 | 0.863 | 0.012 | 0.390 | 0.008 | 0.001 | 0.003 | 0.005 | 0.002 | 0.126 | 0.915 | 0.497 |
| Bias-Corrected: $\text{Cov}(\hat{\xi}_i^A, \hat{\xi}_i^B) = \text{Var}(\xi_i)$ | | | | | | | | | | | | | | | |
| Mean | 0.60 | 0.64 | 0.51 | -0.12 | 0.68 | 3.69 | 0.79 | 5.13 | 2.68 | 1.86 | 2.78 | 1.09 | 2.53 | -0.06 | 2.52 |
| 97.5th Percentile | 1.03 | 0.90 | 0.89 | 1.12 | 1.00 | 5.43 | 1.87 | 5.70 | 3.80 | 2.25 | 3.13 | 1.54 | 3.04 | 0.10 | 3.20 |
| 2.5th Percentile | 0.12 | 0.29 | -0.05 | -1.87 | 0.29 | 1.65 | -0.47 | 4.41 | 1.45 | 1.40 | 2.24 | 0.52 | 1.91 | -0.44 | 1.76 |
| Panel B: Industry-Specific Location Effects | | | | | | | | | | | | | | | |
| Empirical: | | | | | | | | | | | | | | | |
| Raw $\text{Var}(\hat{\tau}_{i,i})$ | 4.91 | 5.64 | 5.20 | 5.79 | 7.17 | 16.32 | 14.74 | 10.88 | 8.67 | 11.89 | 5.37 | 6.51 | 8.11 | 5.57 | 20.34 |
| 90th – 10th Percentile | 47.22 | 41.87 | 87.94 | 66.51 | 52.44 | 68.46 | 58.19 | 86.85 | 66.63 | 72.20 | 50.99 | 35.24 | 55.82 | 55.16 | 69.80 |
| Permutations: $\text{Var}(\hat{\tau}_{i,i})$ | | | | | | | | | | | | | | | |
| Mean | 5.22 | 4.82 | 5.13 | 9.72 | 9.11 | 14.12 | 13.32 | 6.83 | 7.68 | 10.07 | 4.74 | 4.71 | 7.53 | 4.28 | 19.46 |
| Standard Deviation | 0.75 | 0.79 | 0.72 | 3.08 | 1.60 | 1.62 | 1.61 | 2.47 | 1.38 | 1.36 | 1.30 | 0.79 | 2.24 | 0.82 | 2.02 |
| <i>p</i> -value | 0.626 | 0.156 | 0.453 | 0.942 | 0.943 | 0.098 | 0.184 | 0.047 | 0.210 | 0.096 | 0.263 | 0.015 | 0.307 | 0.065 | 0.304 |
| Bias-Corrected: $\text{Cov}(\hat{\tau}_{i,i}^A, \hat{\tau}_{i,i}^B) = \text{Var}(\tau_{i,i})$ | | | | | | | | | | | | | | | |
| Mean | -0.14 | 1.32 | 0.69 | -1.40 | 2.52 | 3.85 | 1.87 | 5.99 | 3.55 | 2.88 | 3.50 | 1.76 | 3.49 | 1.57 | 6.47 |
| 97.5th Percentile | 0.88 | 2.38 | 1.88 | 1.20 | 3.28 | 6.21 | 4.57 | 6.90 | 5.38 | 3.95 | 4.18 | 3.07 | 4.46 | 2.58 | 8.12 |
| 2.5th Percentile | -1.23 | 0.16 | -0.73 | -5.21 | 1.69 | 1.16 | -1.46 | 4.17 | 1.51 | 1.69 | 2.64 | 0.41 | 2.47 | 0.38 | 4.70 |
| Number of NUTS-2 Regions | 9 | 6 | 8 | 5 | 27 | 38 | 8 | 21 | 7 | 17 | 7 | 8 | 19 | 8 | 42 |
| Number of NUTS-2 Region-Industries | 148 | 132 | 170 | 107 | 491 | 717 | 162 | 455 | 135 | 330 | 126 | 176 | 362 | 174 | 731 |
| Number of Firms | 1,678 | 14,625 | 13,091 | 5,829 | 32,050 | 13,244 | 2,768 | 116,918 | 5,189 | 6,650 | 26,019 | 21,279 | 60,375 | 13,916 | 10,125 |


Note: The table replicates Table A.2 but uses log value added per worker rather than TFP.

Figure A.4: Spatial Dispersion of Labor Productivity (Value Added per Worker): United States and European Countries

(a) Location Effects



(b) Industry-Specific Location Effects



Note: The figure replicates Figure 6 but uses log value added per worker rather than TFP.