



## **Dispersion in Dispersion: Measuring Establishment-Level Differences in Productivity**

Sabrina Wulff Pabilonia (Bureau of Labor Statistics, United States), Lucia Foster (United States Census Bureau), Cheryl Grim (United States Census Bureau), Jay Stewart (Bureau of Labor Statistics, United States), Zoltan Wolf (United States Census Bureau), and Cindy Zoghi (Bureau of Labor Statistics, United States)

Paper prepared for the 34<sup>th</sup> IARIW General Conference

Dresden, Germany, August 21-27, 2016

Session 7C: Explaining Productivity Trends II

Time: Friday, August 26, 2016 [Morning]

## **Dispersion in Dispersion: Measuring Establishment-Level Differences in Productivity**

Lucia Foster, Cheryl Grim, Sabrina Wulff Pabilonia, Jay Stewart, Zoltan Wolf, and Cindy Zoghi\*

August 4, 2016

### **Abstract**

Productivity measures are critical for understanding economic performance; in this paper, we present work on supplementing official statistics on industry productivity growth with statistics on the second moments of productivity. Official productivity statistics are produced by Bureau of Labor Statistics (BLS) using industry level data. Those official statistics cannot provide insight on the *within-industry* variation in productivity, limiting our understanding of the rich productivity dynamics in the U.S. economy. Research has shown there are large and persistent productivity differences across businesses even within narrowly-defined industries. These differences vary across industries and time and are related to productivity-enhancing reallocation. Dispersion in productivity across businesses is informative about the nature of competition and frictions within sectors. Productivity differences across businesses are also related to wage differences across businesses; growth in productivity dispersion is related to the rising wage inequality across businesses. BLS and the Census Bureau are collaborating to create measures of within-industry productivity dispersion with the goal of developing public-use and restricted-use statistics that complement existing statistics. We construct establishment-level labor productivity using Census Bureau microdata, compare micro-aggregated industry-level measures to BLS industry-level measures, and examine the variation in our industry-level productivity dispersion measures across industries and time.

\* Foster and Grim: Center for Economic Studies, U.S. Census Bureau; Pabilonia, Stewart, and Zoghi: Bureau of Labor Statistics; Wolf: Westat. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau or the Bureau of Labor Statistics. All results have been reviewed to ensure that no confidential information is disclosed. We thank John Earle, John Haltiwanger, Bart Hobijn, Mark Roberts, Chad Syverson, the Federal Economic Statistics Advisory Committee, the BLS Technical Advisory Committee, and participants at the 2015 Federal Statistical Research Data Center Conference for their helpful comments.

## 1. Introduction

Productivity measures are critical for understanding economic performance in the U.S. economy. The Bureau of Labor Statistics (BLS) produces official labor and multifactor productivity statistics for major sectors and industries in the U.S. using aggregate industry level data. These statistics show how productivity changes *on average* within sectors and industries, but cannot provide insight on productivity variation across establishments within sectors or industries thus limiting our understanding of the rich productivity dynamics in the U.S. economy.<sup>1</sup>

BLS and the Census Bureau have formed the Collaborative Micro-productivity Project (CMP) to address this gap by creating measures of within-industry productivity dispersion with the goal of producing public-use and restricted-use statistics. The public-use measures would cover industries in the manufacturing sector and would be published jointly by BLS and the Census Bureau. Restricted-use establishment-level data measures would be made available in secure Federal Statistical Research Data Centers (FSRDCs).<sup>2</sup>

Why do we need measures of within-industry productivity dispersion? One of the most important findings in the literature on micro-level productivity is large productivity differences across establishments within narrowly-defined industries (Syverson (2011) provides a survey). For example, using data from the 1977 Census of Manufactures (CM), Syverson (2004b) found establishments at the 90<sup>th</sup> percentile of the within-4-digit-SIC productivity distribution are nearly twice as productive as those at the 10<sup>th</sup> percentile. This finding has generated much analysis of the causes and consequences of such dispersion.

Explanations of possible causes include curvature in the profit function that prevents

---

<sup>1</sup> Although usually referred to as industry productivity growth or aggregate productivity growth, these statistics can be thought of as the weighted average of within-industry growth rates.

<sup>2</sup> For more information on the FSRDCs, see the FSRDC website: <http://www.census.gov/fsrdc>.

the most productive business from taking over an industry, frictions in factor adjustments and barriers to the entry and exit of establishments, and distortions that drive wedges in the forces pushing towards the equalization of marginal products across businesses. Within-business drivers of productivity differences are differences in management skills, the quality of production factors, innovation, and R&D. Other factors, external to the business, include competition, productivity spillovers, and the degree of regulation on and flexibility of input markets.

Moreover, the dispersion in productivity across businesses varies across sector, time and geographic area. Syverson (2004a, 2004b) shows that variation in dispersion measures across industries and geographic areas are related to indicators of product substitutability, market structure, and competition. Hsieh and Klenow (2009) argue that cross-country variation as well as within-country variation over time in dispersion of productivity is related to distortions to productivity-enhancing reallocation. Asker, Collard-Wexler and De Loecker (2014) provide evidence that the patterns of dispersion reflect the dynamic adjustment frictions of factors within sectors. The findings in Foster, Grim, Haltiwanger and Wolf (2016) suggest that measured productivity differences across establishments may in fact be generated by differences in efficiency levels, demand shocks, frictions/distortions or all of the above.

These differences are correlated with important economic outcomes at the micro level, such as the growth and survival of establishments. Dispersion is important not only as a measure of heterogeneity but also because it is relevant for business dynamism and growth. There is a large literature on the connection between productivity, reallocation and growth (e.g., Griliches and Regev, 1992; Baily, Hulten, and Campbell, 1992; Baily, Bartelsman, and Haltiwanger, 2001; Petrin, White, and Reiter, 2011; Hsieh and Klenow, 2009, 2014). The main finding in these studies is that more productive businesses are more likely to grow and survive, with the accompanying implication that reallocation is an important contributor to

aggregate productivity growth.

In turn, covariances between micro-level growth and productivity are important for macro outcomes. The pace and extent to which low-productivity businesses contract and exit while high-productivity businesses expand is important for aggregate productivity (see, e.g., Foster, Haltiwanger and Krizan, 2001). Productivity differentials across businesses in the same sector are an indicator of the dynamic and static frictions that prevent resources from being quickly reallocated to the most productive businesses within a sector.

Dispersion is important for understanding other key economic issues as well. For example, rising wage inequality has been shown to be a between-firm phenomenon (e.g., Davis and Haltiwanger, 1991; Barth et al., 2014; Bloom et al., 2015). In addition, a number of studies have found that high-wage establishments are highly-productive establishments (e.g., Dunne et al., 2004). Dunne et al. (2004) also show that rising *between* establishment dispersion in wages is closely associated with rising *between* establishment dispersion in productivity. Economic theories of search and matching provide theoretical justification for why firm productivity dispersion should be closely tied to firm wage dispersion (e.g., Burdett and Mortensen, 1998). High-productivity businesses should pay high wages given that search and matching frictions create quasi-rents for matches between businesses and workers.

These considerations yield a broad motivation for this project along the following lines. Official industry-level productivity statistics can be interpreted as the change in measures of the first moment of micro-level productivity (appropriately weighted). Our objective is to develop second moment measures of micro-level productivity because theory and recent empirical evidence suggest tracking the second moments is informative on a number of important dimensions.

We construct establishment-level labor productivity using microdata from the Annual

Survey of Manufactures (ASM) and the Census of Manufactures (CM) and compare this to BLS industry-level measures.<sup>3</sup> Based upon earlier work comparing similar business data across the two agencies, we expect that there will be systematic differences (see Elvery et al., 2006). We examine a number of data and measurement issues, including those related to the imputation and weighting of the microdata. Since the objective of this project is to provide public-use statistics on productivity dispersion, it will be critical for potential users of these new statistics to understand the relationship between the second moments of the micro-level productivity distribution built up from Census Bureau microdata and the first moment statistics from BLS built from industry-level aggregates. Finally, we explore the variation in our industry-level labor productivity dispersion measures across industries and time.

The paper proceeds as follows. We provide a brief description of the data in Section 2. Section 3 discusses productivity measurement. Section 4 compares inputs, output, and productivity measures. Section 5 looks at both the magnitude of productivity dispersion and patterns in the dispersion in within-industry productivity dispersion across industries and over time. Section 6 summarizes our conclusions and describes plans for future work.

## **2. Data**

In this paper, we compare available BLS aggregate productivity statistics with our derived micro-level productivity statistics. In order to calculate establishment-level productivity, we combine information from three Census Bureau restricted-use microdata files with public-use industry-level data from the Bureau of Economic Analysis (BEA), BLS, and the NBER-CES Manufacturing Industry Database. We obtain deflators and information on the ratio of supervisory to non-supervisory worker hours from industry-level data.

---

<sup>3</sup> In the next draft of this paper, we will also construct establishment-level multifactor productivity (MFP) measures. In general terms, MFP is defined as output per unit of an input index where the latter is a weighted average of input factors. Assuming a Cobb-Douglas production function, the weights are the factor elasticities.

We use establishment-level microdata from the CM, the ASM, and the Longitudinal Business Database (LBD). The CM is collected every 5 years in years ending in ‘2’ or ‘7’. The CM collects data from all manufacturing establishments, except those that are very small. Data for very small establishments is imputed using information from administrative records. The ASM sample is a 5-year panel of manufacturing establishments collected annually. In CM years, ASM data are collected as part of the CM. ASM panels begin in years ending in ‘4’ and ‘9’ and the probability of selection into the ASM sample is a function of both industry and size. The LBD is a longitudinally-linked version of the Census Bureau’s Business Register. The LBD covers the non-agricultural employer universe of business establishments (see Jarmin and Miranda, 2002). The LBD provides us with both high-quality longitudinal links and information on the universe of manufacturing establishments.

Establishments in the CM that are not part of the ASM sample are more likely to be imputed than establishments in the CM that are part of the ASM sample. As we will discuss briefly in the next section, we are concerned that, given the methods used for imputation, imputed data may drive down measured productivity dispersion. While we calculate establishment-level productivity for all establishments where we have the data in the full CM, we create the within-industry dispersion measures shown later in this paper using only ASM establishments even in CM years.<sup>4</sup>

### 3. Productivity Measurement

In this section, we start by describing how BLS measures industry-level productivity, then continue with how we measure micro-level productivity, and end by addressing some more detailed measurement issues.

#### 3.1. *BLS Industry-level Productivity*

---

<sup>4</sup> The microdata made available in the FSRDCs will contain productivity measures for all CM establishments when productivity calculation is possible.

BLS publishes quarterly and annual measures of labor productivity for major sectors; annual measures of labor productivity for 199 3-digit and 4-digit NAICS industries; and annual measures of multifactor productivity for major sectors, 18 3-digit NAICS manufacturing industries, 86 4-digit NAICS manufacturing industries, the air transportation industry, and the line-haul railroad industry. The industry productivity measures published by BLS are constructed with the purpose of measuring changes in the first moment of the within-industry productivity distribution (weighted by establishment size). These indicators measure productivity growth as the percentage change in an index between time periods. Industry productivity *levels* are not published by BLS (but are available on request).

BLS industry output is based on a sectoral concept, which measures the value of goods produced for sale outside the industry. BLS uses published ASM/CM data and adjusts their output measure to remove intrasectoral transactions and resales, and account for changes in finished goods and work-in-process inventories. This nominal output measure is then distributed to detailed output categories of products and services, and revenues from each output category are deflated using the appropriate detailed producer price indexes (from the BLS price program). Deflated values of individual products and services are then Tornqvist aggregated into industry output indexes. For a small number of industries, output is measured using “physical” quantities.<sup>5</sup> Self-employment revenues for manufacturing firms from Internal Revenue Service data sources are also added back to these output measures.

The BLS hours measure is equal to the total annual hours worked by all persons in an industry. This includes paid employees as well as unincorporated self-employed workers and unpaid family workers. These data are obtained from the monthly BLS Current Employment Statistics (CES) and from the Current Population Survey (CPS). The former provides detailed

---

<sup>5</sup> These are mostly service-providing industries, such as commercial banking or airlines, where quantities (transactions or passenger miles) are available.



information on the employment and average weekly hours paid for production and non-supervisory employees. The National Compensation Survey (NCS) is used to adjust hours paid to an hours-at-work basis by removing paid vacation and sick leave. Information from the CPS on the relative average weekly hours of supervisory to non-supervisory workers is used to calculate hours for non-production and supervisory workers as well as to directly obtain hours worked by self-employed and unpaid family workers (Eldridge, Manser, and Otto, 2004). Specifically, non-production and supervisory worker hours are estimated as:

$$\text{Emp}_{\text{NP}}^{\text{CES}} * \text{AWH}_{\text{P}}^{\text{CES}} * \frac{\text{AWH}_{\text{NP}}^{\text{CPS}}}{\text{AWH}_{\text{P}}^{\text{CPS}}} \quad (1)$$

i.e., by applying the supervisory/nonsupervisory hours ratio  $\left(\frac{\text{AWH}_{\text{NP}}^{\text{CPS}}}{\text{AWH}_{\text{P}}^{\text{CPS}}}\right)$  to production worker average weekly hours ( $\text{AWH}_{\text{P}}^{\text{CES}}$ ), and multiplying by non-production (supervisory) employment ( $\text{Emp}_{\text{NP}}^{\text{CES}}$ ).

### 3.2. Establishment-level Productivity

We now describe how we measure labor productivity at the establishment level. Labor is measured as total hours, calculated as shown in (2) using the BLS non-production to production workers ratio.

If  $PH_{et} > 0$  and  $PW_{et} > 0$  and  $TE_{et} > 0$ , then

$$TH_{et} = PH_{et} + \left( \frac{PH_{et}}{PW_{et}} * (TE_{et} - PW_{et}) * \left( \frac{\text{AWH}_{\text{NP}}^{\text{CPS}}}{\text{AWH}_{\text{P}}^{\text{CPS}}} \right) \right) \quad (2)$$

where  $PH$  = production worker hours,  $PW$  = average number of production workers,  $TE$  = total employment,  $\frac{\text{AWH}_{\text{NP}}^{\text{CPS}}}{\text{AWH}_{\text{P}}^{\text{CPS}}}$  = supervisory (non-production) to nonsupervisory (production) hours ratio,  $e$  indexes establishments, and  $t$  indexes years.

Real plant-level total output is calculated as shown in (3). We subtract cost of resales

from output for consistency with the BLS output measure, but are still considering whether or not we will do this in later versions of this paper.

$$\begin{aligned} \text{If the resulting } Q \text{ is positive, then } Q_{et} &= (TVS_{et} + DF_{et} + DW_{et} - CR_{et})/PISHIP_{it} \\ \text{else } Q_{et} &= TVS_{et}/PISHIP_{it} \end{aligned} \quad (3)$$

where  $DF_{et} = FIE_{et} - FIB_{et}$  and  $DW_{et} = WIE_{et} - WIB_{et}$ .

In (3),  $TVS$  = total value of shipments,  $CR$  = cost of resales,  $FIB$  = beginning of year finished goods inventories,  $FIE$  = end of year finished goods inventories,  $WIB$  = beginning of year work-in-process inventories,  $WIE$  = end of year work-in-process inventories,  $PISHIP$  = BLS deflator,  $e$  indexes establishments,  $i$  indexes industries, and  $t$  indexes years. Note that when components of  $DF$  or  $DW$  are missing, they are set to zero.

Establishment-level log labor productivity is calculated as

$$lp_{et} = \log(Q_{et}) - \log(TH_{et}) \quad (4)$$

where  $Q$  is output,  $TH$  is total hours,  $e$  indexes establishments, and  $t$  indexes years.

This micro data is subject to item non-response and these missing values are imputed by the Census Bureau. There is evidence that certain imputation methods may affect microdata analyses that use Census Bureau completed data because imputation is non-random. Census Bureau imputation methods are designed to yield accurate published aggregates and do not necessarily preserve the distribution of the underlying microdata. There are various techniques available to account for the effect of imputation on dispersion. For example, White, Petrin, and Reiter (2012) analyze dispersion statistics using classification and regression-tree methods. Foster et al. (2015) follow a different approach and address imputation by dropping imputed observations and simultaneously correcting for possible selection issues by calculating weighted statistics. The results from these studies suggest that

imputation has a negative effect on within-industry dispersion. We intend to address these issues in future work.

#### 4. Comparison of Productivity Measures

We start by comparing BLS and micro-aggregated input and output measures and then compare labor productivity measures. Aggregating micro-based levels into industry-level indices is not straightforward. In principle, it is possible to generate micro-level index numbers in a manner similar to the approach used by BLS. This amounts to calculating a productivity index for each establishment relative to a reference value from a given year. However, our panel is unbalanced implying that the index may not be defined for all establishments. In addition, using establishment-specific reference values in this manner could potentially change establishments' rankings in the productivity distribution.<sup>6</sup> For these reasons, we calculate the productivity differential of establishments relative to an industry- and year-specific average. Although this approach does not replicate the approach underlying published aggregates, it preserves the number of establishment-year observations we can use for analysis, leaving the productivity rankings intact. The interpretation of establishment-specific productivity levels is intuitive: it tells us how far above or below the mean the establishment sits in the productivity distribution.

##### 4.1. *Comparing Input and Output Measures*

As noted in the introduction, it is important to know how the first moments of our data compare to other data sources. Although differences in the levels of these variables

---

<sup>6</sup> Generating BLS-type indices at the establishment level requires the existence of a value of a variable in a chosen base year. Since the panel is unbalanced, time series are longer for some establishments than others and there are establishments with missing data in the chosen base year. One example is all the establishments that were born after the chosen base year. The implication is that the productivity distribution would need to be truncated for purely technical reasons if we were to follow this approach.

across datasets do not affect dispersion, it is important to have a sense of how our micro-aggregated data compare to aggregate measures underlying published statistics in order to have confidence that measures of dispersion calculated based on Census Bureau microdata may be viewed as dispersion around the published aggregates. There are several sources of data that can be used to measure productivity. Below, we compare our estimates to published estimates from the ASM, data from BLS productivity program, and data from the NBER-CES productivity project. The NBER-CES database is used for comparisons in this section and should be thought of as equivalent to the published ASM and CM data on which it is based.<sup>7</sup>

In Figure 1, we begin by showing the total number of employees in the manufacturing sector from the different series. Employment levels based on aggregated microdata from the ASM using the ASM sample weight are significantly lower than the other series because they do not include the non-mail stratum. The published ASM series includes adjustments for the non-mail stratum and it is slightly below the BLS series. Note, the published ASM series is approximately equal to the NBER-CES series. For this project, we use inverse propensity score weights (see Appendix A for a full discussion of the weighting procedures).<sup>8</sup> All series trend similarly.

We next compare the level and growth in total manufacturing hours, which will be used in calculating labor productivity, between the BLS series based upon employees' hours only, the CMP (micro-aggregated) series, and the NBER-CES series (Figure 2).<sup>9</sup> Recall that the BLS labor input includes the hours worked by all employees, unincorporated self-

---

<sup>7</sup> More information on the NBER-CES Manufacturing Industry Database can be found at <http://www.nber.org/nberces/>. The NBER-CES series we use ends in 2010 and was downloaded June 2016.

<sup>8</sup> We explored the possibility of benchmarking CMP employment (based on the manufacturing universe in the Census Business Register) to BLS employment (based on the manufacturing universe in the BLS Quarterly Census of Employment and Wages). While this benchmarking, by definition, improves the correlation between labor (employment and hours) measures between BLS and CMP, it actually decreases the correlation between BLS and CMP output and measures of other inputs, which are both based on the manufacturing universe in the Census Business Register.

<sup>9</sup> The BLS employee hours series is aggregated from the BLS industry productivity program's data and will not match publicly available data because hours for the self-employed and unpaid family workers have been removed.

employed workers, and unpaid family workers. The ASM collects information on the total number of employees and, for production workers, it collects both hours worked and the number of employees. Thus, prior to comparing BLS annual hours to those derived primarily from the ASM, we remove hours worked by the unincorporated self-employed and unpaid family workers. For the CMP labor hours, we calculate employee annual hours by taking the sum of ASM production worker hours and an estimate of nonproduction worker hours, where the latter is calculated as the product of the number of non-production workers from the ASM, the average weekly hours worked by production workers from the ASM, and the BLS supervisory/nonsupervisory ratio (similar to expression 1). The NBER hours series has also been adjusted using the CPS supervisory/non-supervisory ratio. Again, the levels trend similarly; however, the CMP hours series lies slightly above the BLS and NBER series. Employee hours growth rates trend similarly except for 2005-2007 when the adjusted NBER-CES series diverges.

We compare the three real output series both in levels and growth rates in Figure 3. The ASM does not track the industries to which establishments sell their products, therefore, it not possible to eliminate the effect of intrasectoral transactions at the establishment-level. However, the ASM contains information on finished goods and work-in-process inventories and resales so we can adjust the Total Value of Shipments (TVS) using these variables. The adjusted TVS is then deflated using the BLS deflator as is the NBER series. Removing resales from CMP data should bring it closer to the BLS output measure, or vice versa, but the room for improvement is limited because resales are a small fraction of output. Preliminary evidence suggests correcting for resales marginally improves the match between micro-aggregated output and BLS sectoral output. We add intrasectoral transactions back into the BLS output series for the comparisons.<sup>10</sup> The three level series have similar trends though

---

<sup>10</sup> Another difference that we do not control for is the inclusion of imputed unincorporated self-employed output

the CMP series is much higher than the BLS and NBER-CES series.

The top panel of Table 1 shows correlations between the three series for the total manufacturing sector over the 14 years covered by our sample. Hours data and output data, both in levels and growth rates, are highly correlated across data series, ranging from 0.92 to 0.99, despite the different sources of data used to measure labor hours. The bottom half of the table previews the within-industry results that we will ultimately develop, showing the average correlations across 4-digit NAICS industries. The bottom panel of Table 1 shows the average of the correlations for 4-digit NAICS industries, where the correlations for each industry are calculated over the 14 years of the sample and those correlations are weighted using the proportion of BLS employees in each industry. Again, these correlations are above 0.7 for labor and output, both in levels and growth rates. These correlations indicate that despite differences in data sources and methodologies, micro-aggregated data largely resemble published aggregate data.

#### *4.2. Comparing Labor Productivity*

For the labor productivity comparisons, we calculate labor productivity growth as the change in the log-level labor productivity. For the BLS series, in this instance only, we use output and hours that includes data on the unincorporated self-employed and unpaid family workers.<sup>11</sup> Figure 4 shows that labor productivity growth rates are similar from 1998-2002 and 2010, but differ from 2003-2009. This is unsurprising given the difference in output trends discussed in the previous section. Table 2 shows that the correlations between the BLS and CMP series are not as strong for labor productivity as for the underlying data.

This comparison of inputs, output, and labor productivity serves as an important

---

in the BLS series.

<sup>11</sup> The official BLS productivity series are calculated using percentage changes in the index and thus the BLS series we refer to here differs. Additionally, the total manufacturing productivity series is published by the BLS Division of Major Sector Productivity whereas the data here are aggregated from industry data provided by the BLS Division of Industry Productivity Statistics.

backdrop to the statistics on within-industry dispersion we wish to produce. Although there remains some work to be done to fully understand the differences between the micro-aggregated data and published aggregates and their implications for productivity dispersion, we now proceed with an examination of productivity dispersion.

## 5. Productivity Dispersion

We measure dispersion using the interquartile range (IQR) and the 90-10 differential, which are easy to interpret. The IQR measure shows how much more productive an establishment at the 75<sup>th</sup> percentile of the productivity distribution is than an establishment at the 25<sup>th</sup> percentile. Although the standard deviation could be an obvious alternative measure of dispersion, it is known to be more sensitive to outliers than quantile-based dispersion measures.

To examine how productivity varies across establishments within industries, we measure the dispersion within 4-digit NAICS industry in labor productivity, in terms of the interquartile range and the 90-10 differential. Summary statistics of these measures are shown in Table 3. On average, we find that establishments at the 75<sup>th</sup> percentile are about 2.3 times as productive as those at the 25<sup>th</sup> percentile (the mean IQR is 0.85). And establishments in the 90<sup>th</sup> percentile are almost 6 times as productive as those in the 10<sup>th</sup> percentile (the mean 90-10 differential is 1.78). For value-added per hour, Syverson (2004b) finds that establishments at the 75<sup>th</sup> percentile are about 1.9 times as productive as those at the 25<sup>th</sup> percentile. He finds that the average 90-10 ratio is about 4. These are somewhat smaller than we are finding for gross output per worker.

Furthermore, we find that the amount of within-industry dispersion in labor productivity varies across industries. Figure 5 shows the average within-industry dispersion by year, as well as its IQR spread for the 1997-2010 period. The large differences between the 25<sup>th</sup>, 50<sup>th</sup>, and 75<sup>th</sup> percentiles indicate that there is a lot of variation in the IQR across

industries. This suggests that there are factors that differ by industry, yielding the observed “dispersion in dispersion.” These factors could include differences in shocks, adjustment costs, distortions, technology, and distributions of capital intensities. Further, the variation over time is small relative to the variation across industries. Still, we do find time series variation, and it is likely important. For example, the median IQR in 1998-1999 is about the same as the 75<sup>th</sup> percentile IQR in 2003-2006. And the 25<sup>th</sup> percentile IQR in 1998-1999 is about the same as the median IQR in 2003-2006. Kehrig (2015) finds that the cross-sectional dispersion in productivity is countercyclical and argues this is important for understanding the business cycle. Decker et al. (2015) find a rising trend in productivity dispersion and note that the change in the *responsiveness* to idiosyncratic differences is important for helping to understand the changing pace of dynamism in the U.S. economy.

One purpose of producing these dispersion statistics is to examine how dispersion changes over time, as well as by other industry characteristics. We have shown there do not appear to be any discernible changes in dispersion over the period; we now look at one industry characteristic, establishment-size dispersion. Figure 6 compares within-industry productivity dispersion with within-industry establishment-size dispersion. The figure does appear to show a positive relationship, suggesting that industries with larger establishment size dispersion appear more likely to have greater productivity dispersion. The correlation between the within-industry interquartile range of labor productivity and the establishment size dispersion is 0.3. While we may have expected average establishment size in an industry to have more of an impact, the size distribution of the industry only accounts for a small fraction of the variation across industries. Many other factors are at work, such as those suggested by Syverson (2004a, 2004b). As we note below, we hope to explore these further. It is our objective to release dispersion series into the public domain so that the user community can explore the causes and consequences of the variation in dispersion.



## 6. Conclusions and Future Work

A growing literature uses micro-level data to examine establishment-level productivity dynamics and finds substantial within-industry productivity dispersion. This paper is part of a larger project that seeks to systematically measure and disseminate information about this dispersion within the context of broader statistical measures currently available.

Thus, we start our paper by comparing inputs and output aggregated from micro-level data to BLS aggregates. We start with labor inputs and output (future versions of this paper will add capital, materials, and energy comparisons). Not surprisingly, we find some differences between BLS industry-level data and micro-aggregated ASM data; however, correlations between BLS and micro-aggregated hours and output range from 0.8064 to 0.9877.

Using these measures of labor inputs and output, we develop a measure of labor productivity growth and examine some of its properties. Correlations between BLS and micro-aggregated labor productivity growth are lower than the corresponding hours and output correlations, around 0.5. We find large within-industry dispersion in labor productivity - establishments at the 75th percentile are about 2.3 times as productive as those at the 25<sup>th</sup> percentile on average. In addition, we find significant dispersion in that within-industry dispersion across industries. Dispersion in dispersion over time is small by comparison, but it is likely still important. In a future draft, we will perform similar analyses with a MFP measure.

In future work, we also plan to explore using information from the LBD on entrants and exiters to produce public-use measures of the within-industry productivity differences between entering and exiting plants (i.e., the net entry component of industry-level productivity). This would give researchers new data to leverage when trying to understand

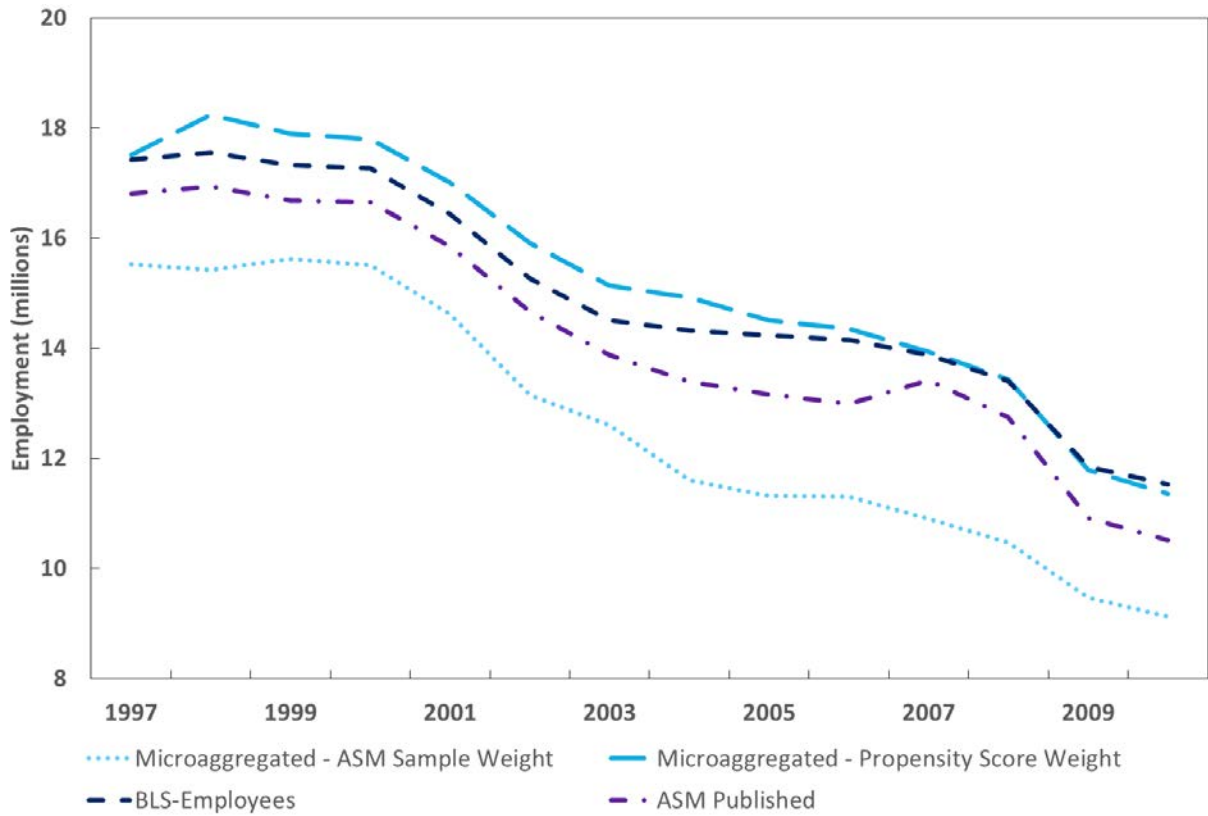
differences over time and across industries in the relative productivity of entrants and exiters.

Contingent on internal and external review, the CMP plans to produce public-use measures of within-industry productivity dispersion for industries in the manufacturing sector. A key benefit of making these data available will be to allow researchers without access to the confidential microdata to explore the various possible causes – and effects - of the differences in within-industry dispersion across industries and over time.

## References

- Asker, John, Allan Collard-Wexler, and Jan De Loecker. 2014. "Demand Fluctuations in the Ready-Mix Concrete Industry." *Journal of Political Economy*, 122(5), 1013-63.
- Barth, Erling, Alex Bryson, James C. Davis, and Richard Freeman. 2014. "It's Where You Work: Increases in Earnings Dispersion across Establishments and Individuals in the U.S," NBER Working Papers 20447, National Bureau of Economic Research, Inc.
- Bloom, Nicholas, Fatih Guvenen, David J. Price, Jae Song, and Till von Wachter. 2015. "Firming Up Inequality," NBER Working Papers 21199, National Bureau of Economic Research, Inc.
- Burdett, Kenneth, and Dale T. Mortensen. 1998. "Wage Differentials, Employer Size, and Unemployment," *International Economic Review*, Department of Economics, University of Pennsylvania and Osaka University Institute of Social and Economic Research Association, 39(2), 257-73.
- Davis, Steven J., and John Haltiwanger. 1991. "Wage Dispersion Between and Within U.S. Manufacturing Plants, 1963-1986," NBER Working Papers 3722, National Bureau of Economic Research, Inc.
- Davis, Steven J., John C. Haltiwanger, and Scott Schuh. 1996. *Job Creation and Destruction*. Cambridge MA: MIT Press.
- Decker, Ryan, John Haltiwanger, Ron S. Jarmin, and Javier Miranda. 2015. "Changing Business Dynamism: Volatility of vs. Responsiveness to Shocks?"
- Dunne, Timothy, Lucia Foster, John Haltiwanger and Kenneth R. Troske. 2004. "Wage and Productivity Dispersion in United States Manufacturing: The Role of Computer Investment," *Journal of Labor Economics*, 22(2), 397-430.
- Eldridge, Lucy P., Marilyn E. Manser, and Phyllis F. Otto. 2004. "Alternative Measures of Supervisory Employee Hours and Productivity Growth." *Monthly Labor Review*, 27(4), 9-28.
- Elvery, Joel, Lucia Foster, C.J. Krizan, and David Talan. 2006. "Preliminary Micro Data Results from the Business List Comparison Project," 2006 Proceedings of the American Statistical Associations, Business and Economics Statistics Section [CD-ROM], Alexandria, VA: American Statistical Association.
- Foster, Lucia, Cheryl Grim, and John Haltiwanger. 2014. "Reallocation in the Great Recession: Cleansing or Not?" NBER Working Paper, No. 20427.
- Foster, Lucia, Cheryl Grim, John Haltiwanger, and Zoltan Wolf. 2016. "Firm-Level Dispersion in Productivity: Is the Devil in the Details?" *American Economic Review* 106(5), 95-98.
- Foster, Lucia, Cheryl Grim, John Haltiwanger, and Zoltan Wolf. 2015. "Micro and Macro Dynamics of Productivity: From Devilish Details to Insights" Working Paper, NBER CRIW Workshop.
- Foster, Lucia, John Haltiwanger, and C.J. Krizan. 2001. "Aggregate Productivity Growth: Lessons from Microeconomic Evidence." In *New Developments in Productivity Analysis*, ed. Charles R. Hulten, Edward R. Dean, and Michael J. Harper. Chicago and London: University of Chicago Press, 303-72.

- Jarmin, Ron S., and Javier Miranda. 2002. "The Longitudinal Business Database." Center for Economic Studies Discussion Paper, No. 02-17.
- Kehrig, Matthias. 2015. "The Cyclical Nature of the Productivity Distribution." Working Paper. Available at [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=1854401](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1854401).
- Syverson, Chad. 2004a. "Market Structure and Productivity: A Concrete Example." *Journal of Political Economy*, 112(6), 1181-222.
- Syverson, Chad. 2004b. "Product Substitutability and Productivity Dispersion." *The Review of Economics and Statistics*, 86(2), 534-550.
- Syverson, Chad. 2011. "What Determines Productivity?" *Journal of Economic Literature*, 49(2), 326-365.
- White, T. Kirk, Jerome P. Reiter, and Amil Petrin. 2012. "Plant-level Productivity and Imputation of Missing Data in U.S. Census Manufacturing Data." NBER Working Paper, No. 17816.



Source: “BLS-Employees” is the annual average of the not seasonally adjusted employment in manufacturing [CEU300000001, Current Employment Statistics program]. “ASM Published” is the published aggregate employment series from the ASM. Microaggregated series are the authors’ calculations on the ASM.

**Figure 1.** Manufacturing Employment Levels, 1997-2010



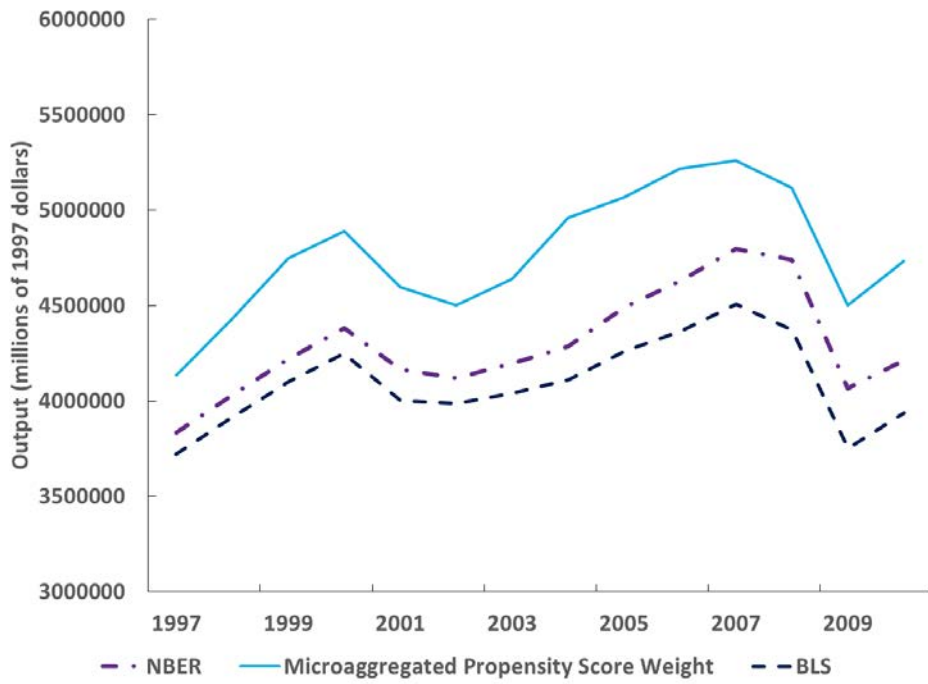
(a) Levels



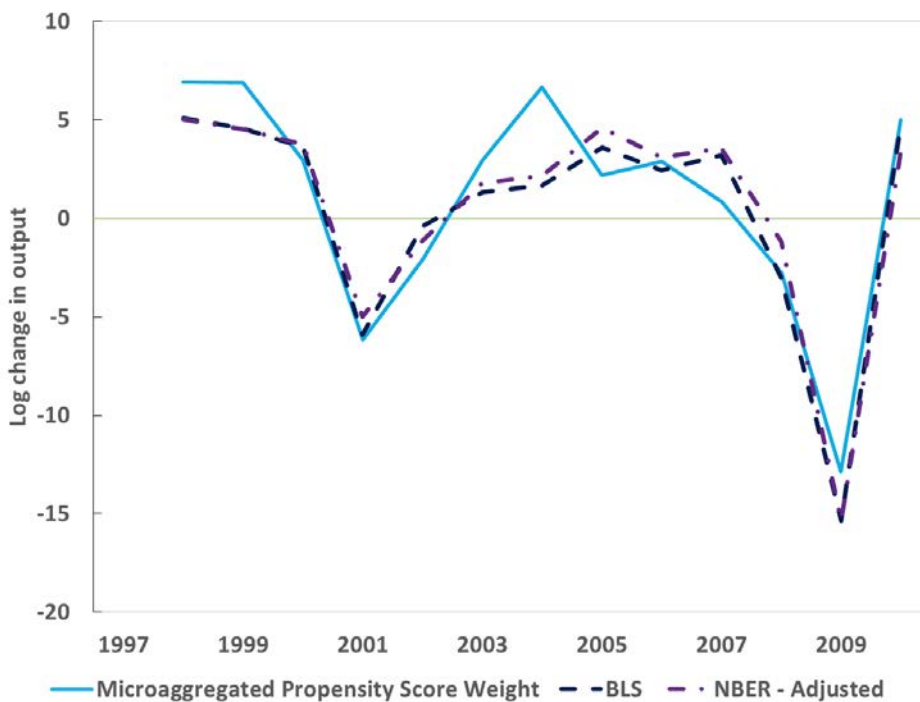
(b) Growth

Source: “BLS-Employees only” is the authors’ calculations from Industry Productivity Program data. “Microaggregated – Propensity Score Weight” is the authors’ calculations on the ASM. “NBER – Adjusted” is the authors’ calculations on the NBER Productivity Database.

**Figure 2.** Growth in Manufacturing Hours Worked, 1997-2010



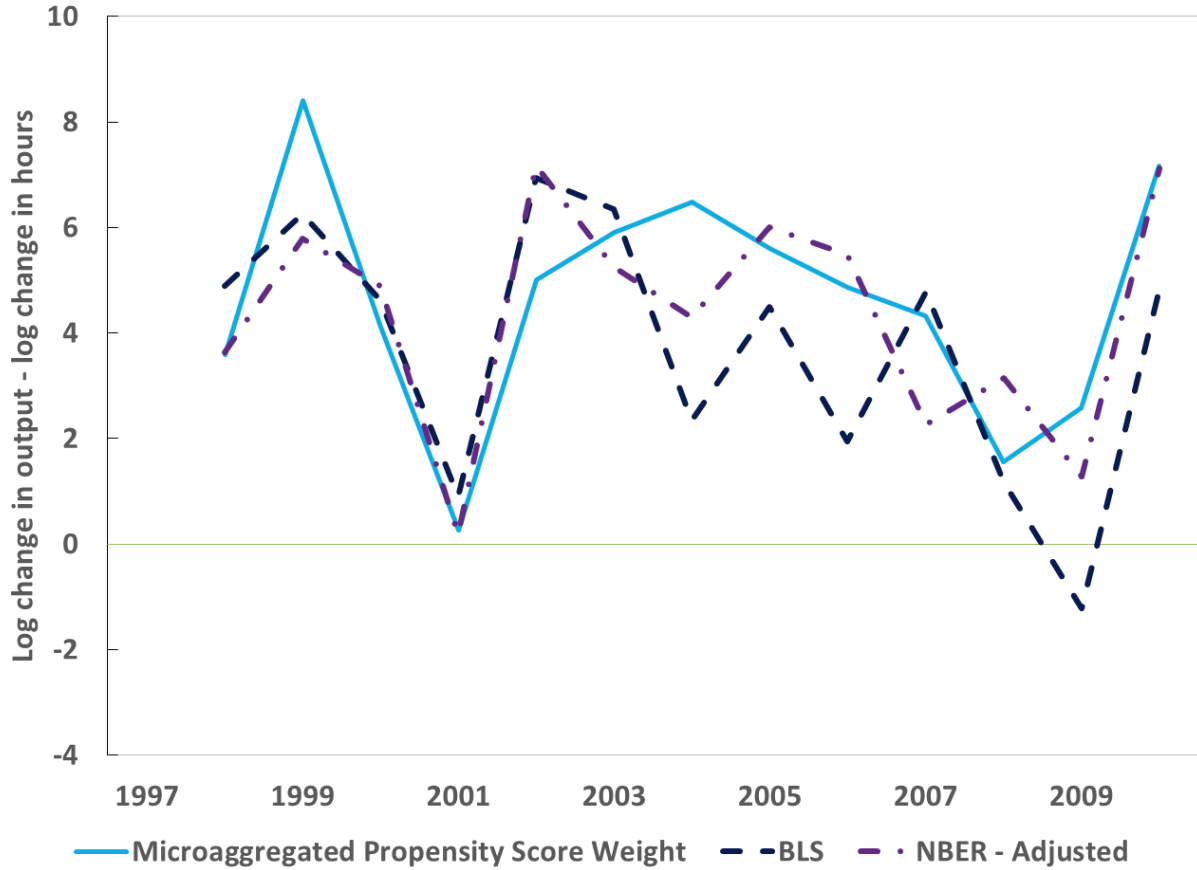
(a) Levels in 1997 dollars



(b) Growth

Source: “BLS” is from the Industry Productivity Program. “Microaggregated – Propensity Score Weight” is the authors’ calculations on the ASM. “NBER” is the NBER Productivity Database.

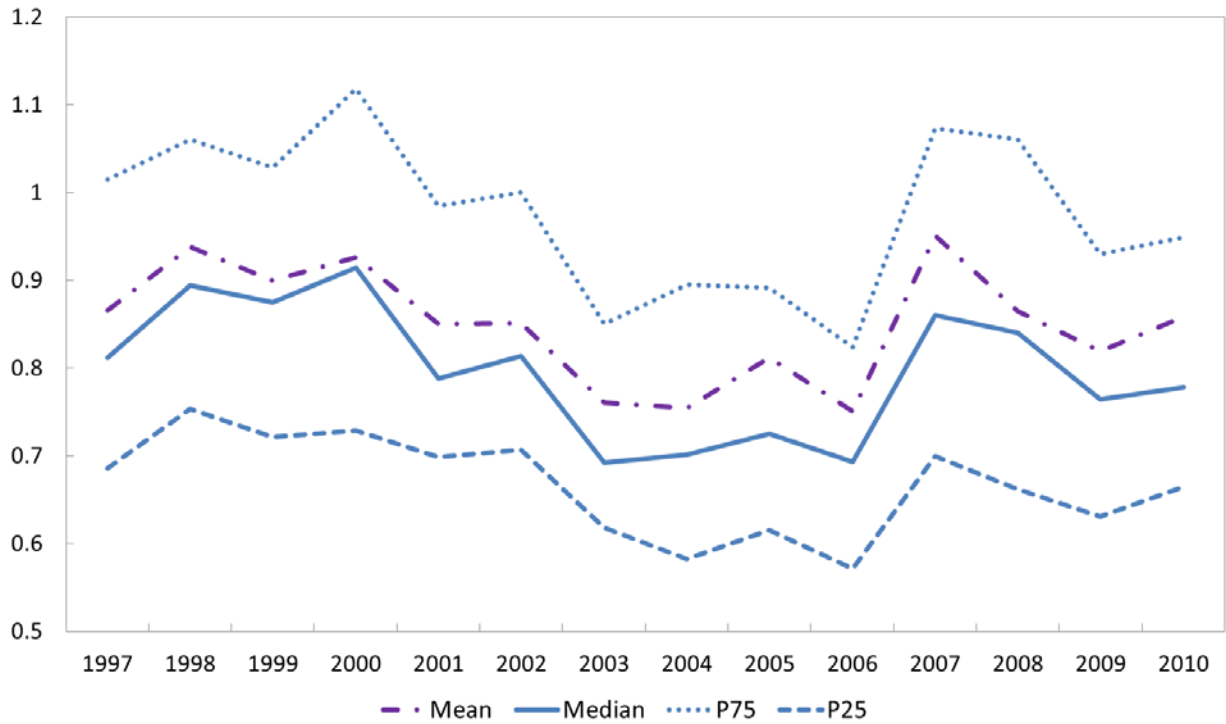
**Figure 3.** Manufacturing Output, 1997-2010



Source: “BLS – Employees only” is the authors’ calculations from Industry Productivity Program data. “Microaggregated – Propensity Score Weight” is the authors’ calculations on the ASM. “NBER - Adjusted” is the authors’ calculations from NBER Productivity Database data.

**Figure 4.** Growth in Labor Productivity, 1997-2010

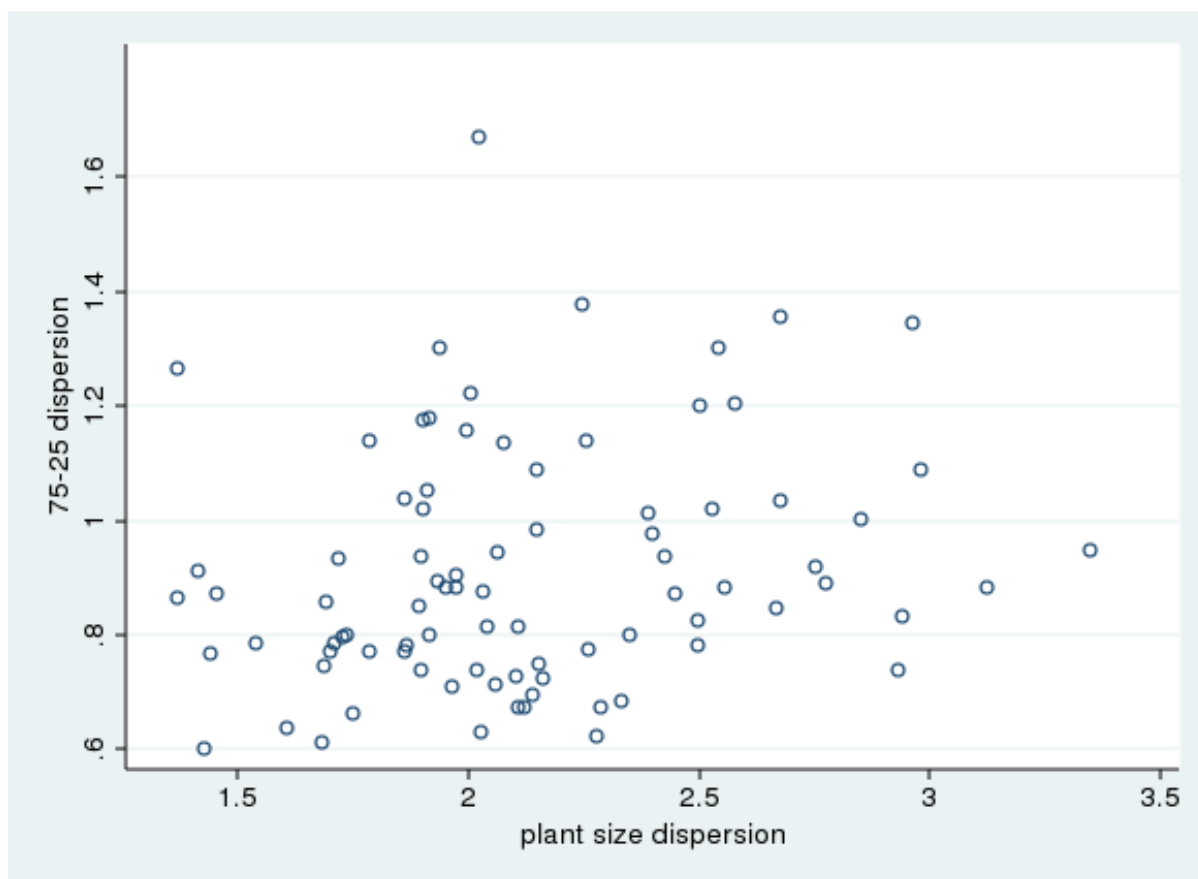




Source: Authors' calculations on the ASM.

Notes: The 6-digit NAICS industry mean is subtracted off establishment log labor productivity. Within-industry productivity moments are created at the 4-digit NAICS level, weighted by our propensity score weight. Annual summary statistics of industry statistics are unweighted.

**Figure 5.** Dispersion in Within-industry IQR of Labor Productivity, 1997-2010



Source: Authors' calculations on the ASM.

Notes: Each point represents the average over 1997 to 2010 of 4-digit NAICS industry-year within-industry IQR of labor productivity. One outlier has been removed for scaling reasons.

**Figure 6.** Within-industry IQR of Labor Productivity by Average Establishment Size, 1997-2010

**Table 1.** Correlations between BLS, CMP, and NBER Hours Worked and Output, Total Manufacturing and Average of 4-Digit NAICS, 1997-2010

	<b>BLS/CMP Correlation</b>	<b>BLS/NBER Correlation</b>	<b>CMP/NBER Correlation</b>
<b>Total Manufacturing</b>			
Hours worked, levels	0.9877	0.9952	0.9922
Hours worked, growth	0.9198	0.9356	0.9196
Output, levels	0.9388	0.9657	0.9618
Output, growth	0.9378	0.9938	0.9267
<b>Average of 4-Digit NAICS</b>			
Hours worked, levels	0.8392	0.9078	0.9029
Hours worked, growth	0.9374	0.7108	0.7708
Output, levels	0.8860	0.9714	0.8821
Output, growth	0.8064	0.9337	0.7579

Source: Authors' calculations on the ASM.

Notes: Resales are removed from output. The BLS output deflator is used.

**Table 2.** Correlations between BLS, CMP, and NBER Labor Productivity Growth, 1997-2010

	<b>BLS/CMP Correlation</b>	<b>BLS/NBER Correlation</b>	<b>CMP/NBER Correlation</b>
Labor productivity growth, total manufacturing	0.6454	0.7136	0.7826
Labor productivity growth, average of 4-digit NAICS	0.5385	0.6000	0.6480

Source: Authors' calculations on the ASM.

**Table 3.** Summary of Within-Industry Productivity Distributions, 1997-2010

<b>Within-Industry Productivity Moment</b>	<b>Mean</b>	<b>Standard Deviation</b>	<b>IQR</b>
<i>Labor productivity</i>			
IQR	0.850	0.266	0.310
90-10 Differential	1.780	0.463	0.602
Standard Deviation	0.761	0.179	0.222

Source: Authors' calculations on the ASM.

Notes:

1. Log labor productivity is calculated as  $\log(\text{output}/\text{hours})$  where hours are BLS-adjusted total hours.
2. The 6-digit NAICS industry mean is subtracted off establishment-level log LP. Within-industry productivity moments are created at the 4-digit NAICS level using propensity score weights. Annual summary statistics of these industry statistics are then created (unweighted). Finally, the numbers shown are unweighted means of the annual summary statistic values from 1997-2010.
3. Resales have been removed from output.

## **Appendix: Weighting**

### *A.1. Properties of ASM samples*

The ASM is a rotating 5-year panel of roughly 50,000-70,000 manufacturing establishments and is a stratified random sample of establishments drawn from the manufacturing portion of the Census Bureau's Business Register.<sup>12</sup> The largest establishments are sampled with certainty and are included in every panel.<sup>13</sup> Smaller establishments are sampled with a probability less than 1, where the probability increases with establishment size (measured by shipments). The smallest establishments, which are part of the "non-mail" stratum, are not sampled at all.

The ASM sample is refreshed every 5 years. New ASM panels are drawn from the Economic Census and begin 2 years after the Census from which it was drawn (years ending in 4 and 9). The sample is also updated to include new establishments that were not present in the Economic Census. In the first year of a new panel, the sample is supplemented with records for new manufacturing establishments that were identified on the Census Bureau's Business Register for the previous year (the processing year following the Economic Census). In subsequent years, the sample is updated using administrative records from the IRS and data from the Company Organization Survey.

Data for the ASM are collected in all years except for years ending in 2 and 7 when the ASM data are collected as part of the Economic Census. Data on payroll, employment, industry, and geography for establishments in the "non-mail" stratum are obtained from administrative records.<sup>14</sup>

---

<sup>12</sup> More information about the ASM can be found at the Census Bureau's website at <http://www.census.gov/manufacturing/asm/>.

<sup>13</sup> Prior to 1999, certainty units were establishments with 250 or more employees. In 1999, the cutoff was increased to 500 employees, and in 2004, it was increased again to 1,000 employees.

<sup>14</sup> Federal regulations require the Census Bureau to limit small establishments' survey response burden.

As a result of this sample design, large establishments are over-represented in the ASM sample, and the smallest establishments are completely excluded from the sample. The ASM sample weights, which are inversely proportional to a shipments-based establishment size measure, could in principle be used to correct for the effects of the ASM sample design. However, because these weights do not account for establishments in the non-mail stratum, the weighted sum of shipments will not match published totals.<sup>15</sup>

Another important aspect of the sample design is the rotating nature of the ASM, which means that the composition of establishments changes over time and between panels. Any weighting procedure aiming at creating unbiased estimates should account for the fact that the sampling probabilities and therefore the composition of the ASM change every 5 years. In addition, sampling and non-mail stratum thresholds vary across years. This aspect of the ASM sample design could result in breaks in series that are not easily addressed with sample weights.

#### *A.2. Establishment characteristics and the probability of selection into the ASM*

The ASM's sample design has important implications for our analysis. Typically, sample weights are used to make stratified samples representative. For example, the sum of the sample weights of the establishments in the sample might equal total employment or total sales. However, it is not clear that the ASM sample weights are appropriate for our analysis. This section is devoted to describing our weighting procedure.

To correct for the effects of the ASM's sample design and to address the shortcomings of its sample weights, we construct propensity score weights using the Longitudinal Business Database (LBD). See Jarmin and Miranda (2002) for more details about the LBD. The propensity score weights are inversely proportional to the probability

---

<sup>15</sup> See Davis, Haltiwanger and Schuh (1996) for more details.

that an establishment is selected into the ASM. The probability of being selected into the ASM is determined using logistic regressions.

We start by matching establishments in the ASM to LBD establishments by year and “LBD Number.”<sup>16</sup> Our dependent variable is a dummy variable that equals one if the establishment is in both the ASM and the LBD for that year and zero if the establishment is only in the LBD.

Establishments that are in the non-mail stratum have largely imputed data, and we set the ASM dummy to 0 for such cases. The set of regressors consists of dummy variables that classify each establishment based on its employment and payroll size class, whether the establishment is part of a multi-unit entity, the establishment’s industry code, and the interaction between industry and employment size effects. Including industry-size interactions allows us to estimate industry-specific size distributions. These variables are obvious candidates for our logistic regressions because the probability of selection into the ASM sample and the cutoff for the non-mail stratum in the ASM vary by industry and size.

When determining weights, we define industries at the 3-digit NAICS level because the interaction of size indicators and more narrowly-defined industry codes leads to empty cells in smaller industries. Empty-size bins imply the size distribution cannot be estimated in these industries.<sup>17</sup> When the size distribution cannot be estimated for an industry, propensity scores cannot be calculated because maximum-likelihood estimates of the size effects do not exist. Empty cells can in principle be avoided by collapsing size bins, combining similar narrowly-defined industries, or allowing bin definitions to vary across industries. We experimented with the number and definition of the size bins, and also the level of industry

---

<sup>16</sup> The LBD Number is an establishment identifier that is consistently defined across both datasets. While linking the datasets by LBD Number is straightforward, there are a small percentage of establishment-year observations that do not match due to timing issues between the ASM and the LBD.

<sup>17</sup> The size distribution cannot be estimated if all establishments are in the same size bin.

aggregation and found that using 3-digit industry codes together with 4 size bins allows us to estimate the size distribution in every industry and year. Allowing for more heterogeneity by using either industry-specific size bins or more narrowly defined industries leads to feasibility problems with the logistic regression.

We defined the size bins so that the resulting distribution allows the lowest size bins to vary over time. That is, in every year and every industry, the 50 percentile of the establishments with less than 50 employees is used to define bins 1 and 2. For larger establishments, the following bins are defined: 50-99, 100-199, 200+.<sup>18</sup> There are 21 3-digit NAICS industries in the 2002 classification system, which results in 105 industry specific size distributions. We include a continuous size measure in order to allow the weights to vary within these cells. This is necessary to account for possible within-cell compositional changes. Adding 5 payroll classes and 2 groups related to multi-unit status increases the number of cells to 113.<sup>19</sup>

The 2002 change in the industry classification system resulted in missing NAICS-2002 codes for a nontrivial number of establishments in the LBD between 1997 and 2001. For example, the NAICS code is missing if an establishment exited prior to 2002. For these observations, we used imputed NAICS codes.<sup>20</sup> From 2002 on, NAICS codes are available for all establishments in the LBD.

Our inverse propensity score weights generate employment counts that do a good job of matching the trends and cyclical variation on BLS manufacturing employment, but they do not match BLS levels.

---

<sup>18</sup> The payroll size classes are 0-200, 201-500, 501-1000, 1001-5000, 5001+.

<sup>19</sup> If we were to use 4-digit industry, the number of cells would increase significantly. There are 86 4-digit NAICS industries implying 86 different size distributions and 430 industry-size cells. Such an increase in the number of cells yields empty size bins in several industries.

<sup>20</sup> NAICS codes are imputed using a method described in: Teresa Fort. 2013. "NAICS 2002 Code Assignments to the Longitudinal Business Database." Mimeo, Center for Economic Studies