

w o r k i n g
p a p e r

05 09

**The Importance of Reallocations
in Cyclical Productivity and
Returns to Scale: Evidence
from Plant-level Data**

by Yoonsoo Lee



FEDERAL RESERVE BANK OF CLEVELAND

Working papers of the Federal Reserve Bank of Cleveland are preliminary materials circulated to stimulate discussion and critical comment on research in progress. They may not have been subject to the formal editorial review accorded official Federal Reserve Bank of Cleveland publications. The views stated herein are those of the authors and are not necessarily those of the Federal Reserve Bank of Cleveland or of the Board of Governors of the Federal Reserve System.

Working papers are now available electronically through the Cleveland Fed's site on the World Wide Web: **www.clevelandfed.org/Research**.

The Importance of Reallocations in Cyclical Productivity and Returns to Scale: Evidence from Plant-level Data

By Yoonsoo Lee

Procyclical productivity plays an important role in many models of aggregate fluctuations. However, recent studies using aggregate data to directly measure technology shocks in the Solow residual find that technology shocks are not procyclical. This paper provides new evidence that, due to countercyclical composition changes between producers, the procyclicality of productivity observed in aggregate data may be understated. Using plant-level microdata, this paper finds that the reallocation of output shares across continuing plants, as well as the entry and exit of plants, creates a countercyclical component in aggregate productivity. This paper shows that such composition changes may cause a downward bias in industry-level estimates of returns to scale. The findings of this paper suggest that, without correcting for the countercyclical effects of reallocations, estimates based on aggregate data may not reflect the true cyclicity of technology shocks, which a representative agent faces over the business cycle.

JEL Classification: D24, E32, O47

Keywords: Entry, Exit, Productivity, Returns to Scale

Yoonsoo Lee is at the Federal Reserve Bank of Cleveland and may be reached at Yoonsoo.Lee@clev.frb.org or (216) 579-2970. The author thanks Mark Bills for his valuable advice and encouragement. He also thanks Eric Bartelsman, Paul Bauer, Yongsung Chang, Ben Craig, Paul Gomme, Lance Lochner, Peter Rupert, and Mark Schweitzer for very helpful conversations and suggestions, as well as Randy Becker, Renee Fields, Shan Klimek, and Christopher Kurz who were especially generous with their time in helping with the Census data. The research in this paper was conducted while the author was a Census Bureau research associate at the Carnegie Mellon University Research Data Center and at the Michigan Census Research Data Center.

1 Introduction

Explaining the cyclical behavior of productivity is essential for understanding business cycles. Based on the Solow residual as a measure of productivity changes, many researchers have given a central role to technology shocks in explaining the procyclical behavior of productivity and aggregate fluctuations (See, for example, papers in Cooley, 1995). However, in a series of papers for which they use a modification of Solow's methodology, Hall (1988, 1990) and others claim that real business cycle analysis overestimates the magnitude of productivity shocks and, thus, the contribution of these changes to aggregate fluctuations. Furthermore, Basu, Fernald, and Kimball (2005) argue that correctly measured technology shocks are not correlated with output or with the business cycle.¹

Although their findings vary, previous studies are similar in their use of aggregated data. In this paper, I argue that cyclical reallocations across producers with different levels of productivity cause a countercyclical bias in aggregate productivity and obscure the true cyclicity in productivity, which a representative firm or plant faces over the business cycle. I use detailed plant-level data on U.S. manufacturing to quantify and assess the empirical importance of the aggregation issues implicit in previous studies that use aggregated data to measure short-run productivity changes. Using the Longitudinal Research Database (LRD), I examine how entry and exit, as well as reallocations among plants of different efficiencies, change the composition of producers over the business cycle.

In general, I find that cyclical reallocations between plants create a countercyclical composition bias in aggregate total factor productivity (TFP), reducing the procyclicality of aggregate TFP. Output shares are reallocated from less-productive to more-productive plants during recessions. Furthermore, plants entering and exiting during a boom are less productive than those entering and exiting during a recession.

The countercyclical effect of composition bias on aggregate TFP is similar to the effect of composition bias on the cyclicity of real wages (Stockman, 1983; Bils, 1985; Solon et al, 1994; Chang, 2000). In contrast to the labor market, in which changes in the composition of workers are mostly explained by changes that occur on the extensive margin, i.e., entry and exit of

¹ Using a very different approach to identify technology shocks, Gali (1998) and Kiley (1997) also reach the similar conclusion that technology shocks are very negatively correlated with inputs.

workers, changes in composition of producers are mostly explained by changes that occur on the intensive margin of production, i.e., reallocations between continuing plants. The magnitude of composition bias caused by entry and exit of plants is relatively small, because of the small output share of the industry which entering and exiting plants account for.

Having established the importance of composition changes over the business cycle in understanding the cyclical behavior of aggregate productivity, I proceed to show that such composition changes may bias aggregate estimates of returns to scale. Increasing returns to scale is put forward as an explanation for the procyclical behavior of productivity and as an important propagation mechanism in models of the business cycle. However, recent studies based on industry-level data, such as Basu and Fernald (1997), find decreasing, rather than increasing returns to scale. What matters for macroeconomic models are returns to scale across firms, appropriately aggregated. Many researchers have interpreted returns-to-scale estimates based on aggregated data as the returns to scale of a representative firm. However, estimates from aggregated data may not serve as reliable estimates of the average firm-level parameters if the composition of producers with different levels of productivity changes over the business cycle. As inputs are reallocated toward less-productive firms during booms, for example, the marginal response of output to input changes may appear lower in aggregate data than the marginal increase in output of a typical firm, leading to smaller estimates of returns to scale in aggregate data. Directly assessing the effects of composition changes, I find that a composition bias helps explain the finding of decreasing returns to scale. In most two- and four-digit SIC industries with significantly decreasing returns to scale, returns-to-scale estimates *decrease* as the plant-level data are aggregated to the industry level.

This finding of the countercyclical effects of reallocations between plants sharply contrasts with previous findings on the effects of reallocations between industries, such as those of Basu and Fernald (2002) and Basu, Fernald, and Kimball (2005). Basu and Fernald (2002) claim that reallocation between two-digit industries with different marginal products explains much of the cyclicity of aggregated productivity. However, the plant-level evidence in this paper points to some potential problems in such studies based on aggregated data. First, the differences in true marginal products across industries, once corrected for composition bias, may not be as large as they appear in industry-level data. Compared to industry-level estimates that vary substantially across industries, plant-level estimates are rather closer to constant returns to scale.

Furthermore, the previous mentioned studies, which are based on macro-aggregates built from industry-level data, ignore important reallocations that occur within a detailed industry. The evidence from plant-level data clearly shows that, correcting for aggregation effects running from plants to the manufacturing industry, aggregated productivity would be more procyclical.

This finding of countercyclical composition bias in aggregate TFP is consistent with the findings of Baily, Bartelsman, and Haltiwanger (2001) on the way reallocations between plants affect the cyclical behavior of aggregate labor productivity. Whereas their study focuses on labor productivity, a key aspect of the present study is an exploration of the effects of cyclical changes in producer composition on aggregate TFP and estimates of returns to scale. As discussed in Chang and Hong (2005), labor productivity reflects changes in the input mix, in addition to technological changes; moreover, TFP is the right concept for studying technology shocks.

Section 2 of this paper provides a description of the data and empirical evidence of composition changes over business cycles. Section 3 examines how these changes in the composition of producers may affect returns-to-scale estimates for different levels of aggregation. Conclusions are presented in the last section.

2 Composition Changes and the Cyclicity of Productivity

2.1 Measurement of Productivity and Data Description

The plant-level data used in this study are taken from the LRD maintained by the Center for Economic Studies at the U.S. Bureau of the Census. The LRD is constructed by linking individual establishment records from the Census of Manufactures, which is taken every five years, and the Annual Survey of Manufactures (ASM), which is taken every non-census year. In this study, I use the ASM portion of the LRD for the years 1972 through 1997. Because the entire ASM comprises a representative sample of manufacturing plants (Davis, Haltiwanger, and Schuh, 1996), the survey allows me to assess the contribution of entry and exit to the cyclical behavior of productivity as well as the impact of output reallocation across plants.

Plant-level productivity is measured using a standard total factor productivity index similar to that used by Baily, Hulten, and Campbell (1992) and Foster, Haltiwanger, and Krizan (2001). The TFP index of plant j is computed as follows:

$$\ln tfp_{jt} = \ln Y_{jt} - \alpha_l \ln L_{jt} - \alpha_m \ln M_{jt} - \alpha_k \ln K_{jt},$$

where Y_{jt} is real gross output, L_{jt} is labor input, M_{jt} is real materials, and K_{jt} is real capital stock. The input cost shares for four-digit industries are used as the measure of the corresponding factor elasticities.² There are two problems in measuring cost shares in the ASM. First, the ASM only includes the wage and salary costs of labor. In calculating labor's share, I follow Bils and Chang (2000), magnifying each four-digit industry's wage and salary payments to reflect other labor payments, such as fringe payments and employer FICA payments.³ Another problem is that capital expenditures are not available for the ASM. Given that previous studies by Rotemberg and Woodford (1995) and Basu and Fernald (1997) find small profits in manufacturing, I assume zero profit at the industry level so that total revenue will be equal to total cost. Next, I calculate the share of costs for input J in the total revenue from the four-digit industry-level data, aggregated from the ASM panels. For these computations, I consider capital expenditure shares to be residuals. Basu, Fernald, and Kimball (2005) used the same strategy and found a result similar to direct attempts at measuring capital's share.

Outputs and inputs are measured in 1987 constant dollars. Real gross output is measured as the total value of shipments, deflated by the four-digit industry output deflator for the industry into which the plant is classified. All output, materials, and investment deflators are from the NBER manufacturing productivity data set (Bartelsman and Gray, 1996).⁴ Labor input is measured as total hours for production and non-production workers. Because hours for non-production worker are not collected in the ASM, the total hours are estimated following the method used in Baily, Hulten, and Campbell (1992), in which *total hours* represents total hours for production workers multiplied by the ratio of total payroll for all workers to the payroll for

² This procedure implicitly assumes that all plants in the industry operate with the same production technology, a common assumption in such studies.

³ Bils and Chang (2000) use information from the National Income and Product Accounts to calculate the ratio of these other labor payments to wages and salaries at the two-digit industry level. I thank Yongsung Chang for providing the data.

⁴ See Bartelsman and Doms (2000) for the drawback to using deflated production to measure productivity. Some caution is needed in interpreting the results. Ignoring any quality improvement in output that is not reflected in the deflator may result in a downward bias in productivity growth. If new plants enter a market with new products having higher prices, and the number of new plants increases during a boom, the use of a single industry-level deflator may lead to overestimate the procyclicality of aggregate productivity. As Klette and Griliches (1996) pointed out, returns to scale estimates obtained from production function regressions might be biased downward if firms sell outputs at different prices (imperfect competition) but firm-level outputs are deflated based on a common output deflator.

production workers. Material input is measured as the cost of materials deflated by the four-digit industry materials deflator. Capital stocks for equipment and structures are constructed using the perpetual inventory method.

2.2 Patterns of Entry and Exit over the Business Cycle

Previous studies on the entry and exit of producers document considerable fluctuations in entry and exit rates over the business cycle. Campbell (1998) finds that the quarterly entry rate exhibits procyclical behavior, whereas the quarterly exit rate is countercyclical and is positively correlated with future GDP growth. Cooper and Chatterjee (1993) and Devereux, Head, and Lapham (1996) find that net business formation shows a strong procyclical movement. Figure 1 shows the annual entry and exit rates measured as the share of entering or exiting plants in the manufacturing industry within a given year. In this study, entering plants are either new plants, which appeared in the LRD for the first time, or plants that restarted production after a certain period of inactivity. Similarly, exiting plants include those that stopped producing during the following period and stayed inactive for a certain period of time, as well as those that permanently shut down.⁵ As discussed in detail in Davis, Haltiwanger, and Schuh (1996), samples in the ASM panels are rotated every five years. Only large “certainty” plants are continuously observed across different ASM panels; moreover, it is very difficult to measure entry and exit between the two years in which the panels are rotated. In order to avoid measurement errors in entry and exit caused by the panel rotations, the results reported exclude entries and exits measured between two different ASM panels, namely 1973–74, 1978–79, 1983–84, 1988–89, and 1993–94. Figure 1 presents the interpolated values for these missing years (i.e., the first ASM years in each rotation, 1974, 1979, 1984, 1989, and 1994).

The entry rate rises during economic booms and falls during recessions. The correlation between the annual entry rate and the annual growth rate of real GDP (excluding the first ASM panel years) is .242. The procyclical behavior of the entry rate is consistent with the findings of previous studies. The annual exit rate in Figure 1 covaries positively with the entry rate. This counterintuitive, procyclical behavior of the exit rate is the result of the fact that, in this study, the category of exiting plants includes those that stop production temporarily. Most of these plants enter during the boom part of the cycle, operate for a short period of time, and stop

⁵ I classify plants that have zero employees or produce zero output as inactive plants.

operating until they reenter the market. Such temporary exits increase during booms, explaining more than 50% of plant exits in certain years. In contrast, most of the exits during recessions consist of plants that shutdown permanently. When measured on the basis of the number of *permanent* shutdowns (i.e., excluding temporary shutdowns), the annual exit rate does not show procyclical behavior.⁶

Table 1 summarizes the shares and relative TFP indexes of entering and exiting plants for the sample period. The shares are measured in terms of the numbers (the entry and exit rates in Figure 1), employment, and total output accounted for by entering and exiting plants. Entering plants account for about 7% of plants in a given year, while 10% of plants in a given year stop producing during the following year. Entering and exiting plants tend to be smaller than continuing plants, as reflected in their generally smaller shares of employment and output (2~3%). These small contributions contrast with previous studies, such as that of Foster, Haltiwanger, and Krizan (2001), which finds that entries and exits make significant contributions to aggregate productivity growth over a longer (five- or 10-year) time horizon. This difference results from the difference in the time horizon over which entry and exit are measured. As the period becomes longer, the number of plants that have entered or exited during the given time period increases. Consequently, the output and labor shares accounted for by entering and exiting plants can be much larger when measured over a longer time horizon.

The last column of the table reports the relative TFP indexes for entering and exiting plants. These indexes consist of the weighted averages of TFPs for entering or exiting plants, divided by the weighted averages of TFPs for continuing plants in the same four-digit industry during the same year. In general, entrants are more productive than exiting plants. This result is consistent with the vintage capital model, in which new plants with new technology replace older, less productive plants. I find that within the same four-digit industry, entrants are relatively more productive than continuing plants, while exiting plants are less productive than continuing plants. While the finding of exiting plants' lower productivity is consistent with the those of previous studies, such as Foster, Haltiwanger, and Krizan (2001), the finding of

⁶ These permanent shutdowns are often called plant deaths in the literature. In a similar way, new plants that appear in the LRD for the first time are called plant births. Measuring entry and exit rates as the share of plant births and deaths in the total number of plants in a given year, I find a negative correlation between the annual entry rate and the annual exit rate. The contemporaneous correlation between the entry rate and the log change of real GDP increases to .289. The correlation between the exit rate and the log change of real GDP is .05.

entrants' higher productivity differs from previous research. This difference is mainly explained by the following facts: First, the relative productivity of new plants in the ASM panel years is slightly higher than in Census of Manufactures years. Second, new plants that entered in the 1990s have relatively higher productivity than earlier cohorts of entrants. These recent cohorts were not examined in Foster, Haltiwanger, and Krizan (2001).⁷

The same statistics are separately reported for economic boom and recession periods in order to illustrate how the contributions of entering and exiting plants change over time. The second row summarizes the shares and relative TFP for entering and exiting plants for periods when the growth rate of real GDP was greater than 4% (i.e., a boom). The third row provides the same statistics for periods when the growth rate of real GDP was less than 1% (i.e., a recession). As described in Figure 1, the share of entrants, measured as the number of entering plants divided by the number of all plants (i.e., the entry rate), increases during a boom and decreases during a recession. However, although the number of entering plants increases during a boom, the output share accounted for by entering plants does not increase to a significant degree. This finding is partly explained by the relatively low productivity of entrants during a boom. Although overall productivity is higher for entrants than for continuing plants, plants that enter during a boom are less productive than continuing plants in the same industry. In contrast, plants that enter during a recession or in normal times are more productive than continuing plants. Although the magnitudes are relatively small, these differences in productivity over the business cycle are also found for exiting plants. Plants that exit during a recession are more productive than those that exit during a boom.

This difference in the relative productivity of plants that move in and out of production suggests that aggregate productivity is subject to composition effects. Because overall productivity is lower for plants that enter during a boom than for those that enter during a recession, plant entry may create a countercyclical composition bias in aggregate productivity. The differences in exiting plants' productivity across the cycle may make the behavior of aggregate productivity look more procyclical. Whether the composition bias caused by the entry

⁷ While the results are reported for the case that uses real capital stocks, obtained using the perpetual inventory method, the deflated book value for capital was also used for the purpose of comparison with previous studies. The results were similar to those of the previously mentioned authors when the deflated book value of capital was used.

and exit is countercyclical or procyclical will depend on the relative share and productivity of the entering and exiting plants over the business cycle.

2.3 Decomposition of Aggregate Productivity Changes

Using plant-level data, I examine the extent to which such changes in the composition of producers or shifts in the share of outputs across plants affect the cyclical patterns of aggregate productivity. Following Baily, Bartelsman, and Haltiwanger (2001), the time series changes in aggregate productivity are decomposed into components that reflect a within-plant component (holding output shares fixed) and other effects that reflect the reallocation of shares across plants including the effect of entry and exit:⁸

$$\begin{aligned} \Delta \ln TFP_t = & \sum_{j \in \text{Conti}} \bar{s}_j \Delta \ln tfp_{jt} + \sum_{j \in \text{Conti}} \Delta s_{jt} (\overline{\ln tfp_j} - \overline{\ln TFP}) \\ & + \sum_{j \in \text{Entry}} s_{jt} (\ln tfp_{jt} - \overline{\ln TFP}) - \sum_{j \in \text{Exit}} s_{j,t-1} (\ln tfp_{j,t-1} - \overline{\ln TFP}), \end{aligned}$$

where $\ln tfp_{jt}$ is TFP index for plant j at time t , $\ln TFP_t$ is the aggregate TFP index at time t , s_{jt} is the share of output at plant j at time t , and a bar over a variable indicates the average of the variable over the base and end years ($t - 1$ and t). Because a sample of plants from the ASM is used, the share is further inflated by the ASM sampling weight. The first term reflects changes in productivity from continuing plants holding output shares (often interpreted as a “within” effect). This term is measured as the weighted sum of productivity changes with the weights equal to the average output shares across time. The second term reflects changes in output share from continuing plants for fixed levels of productivity (often interpreted as a “between” effect). The last two terms represent the contribution of entering and exiting plants, respectively.

In this decomposition, the change in shares in the second (between-plant) term is weighted by the deviation of plant-level productivity from the average of aggregate productivity, so that an increase in the share of output for a plant contributes positively only if the plant has higher productivity than the average aggregate productivity. In a similar manner, a new plant contributes positively to the aggregate change only if its productivity is above the average, while an exiting plant contributes positively only if its productivity is below the average.

⁸ This modifies the decomposition method used by Griliches and Regev (1995) to allow for entry and exit. Foster, Haltiwanger, and Krizan (2001) provide excellent reviews of previous studies using different decomposition methodologies and measurement issues.

The results of these decompositions are reported in Figure 2. As in Section 2.2, the results reported exclude the first ASM years in each panel in order to avoid measurement errors due to sample rotations in the ASM panels. The values for these missing years in the figure are interpolated. The decomposition components for total factor productivity reveal cyclical patterns similar to those found for labor productivity in Baily, Bartelsman, and Haltiwanger (2001). The within-plant component shows clear procyclical behavior. It increased sharply during the booms of 1976 and 1983 and decreased markedly during the recessions of 1980 and 1991. Excluding the first ASM panel years, the contemporaneous correlation between the within-plant component and real GDP growth is 0.69. Whereas the within-plant term is very procyclical, the between-plant term moves in a countercyclical direction. The between-plant component increased during the recessions of 1975, 1982, and 1991 and decreased during the recovery years of 1976, 1983, and 1992. Although except during the 1990s the contribution of plant entry and exit to the annual change in aggregate productivity growth was relatively small, the net entry component also moved in a countercyclical direction. The correlations of the between-plant component and of the net entry component with the annual change in real GDP are $-.16$ and $-.11$, respectively.

These countercyclical reallocation terms suggest that output shares shift from less-productive toward more-productive plants during recessions. While the contribution of net entry is relatively small, the magnitude of the impact of reallocations between continuing plants is significant, with the countercyclical effect of the between-plant component occasionally dwarfing the procyclical effect of the within-plant component. As a result of these countercyclical tendencies, aggregate productivity may look less procyclical than the true procyclicality of productivity that is typically observed among individual plants.

2.4 Implications of Returns-to-Scale Estimates Based on Aggregated Data

The decomposition results of aggregate productivity confirm the importance of heterogeneity and aggregation issues for aggregate fluctuations. In the remainder of the paper, I argue that composition changes across plants with different productivity levels may bias estimates of returns to scale based on aggregated data.

The decomposition results suggest two channels through which aggregation might affect the aggregate returns-to-scale estimates in equation (4). First, the aggregate returns-to-scale estimates may be affected by cyclical changes in the output shares between continuing plants.

Although aggregate inputs increase during booms and decrease during recessions, the extent to which inputs change over the business cycle varies across plants. The relatively large between-plant component, which exhibits countercyclical behavior, suggests that the shift of output shares from less-productive to more-productive plants during a recession may cause a downward bias in the returns-to-scale estimates. Because the shares of more-productive plants increase during recessions, the extent to which aggregate output decreases would be smaller than the decreases that would have been observed in a representative plant. And because the shares of less-productive plants increase more during booms, the increase in output would look smaller in aggregate data than the marginal increase in output of a representative plant.

Second, the entry and exit of plants with different levels of productivity may also affect the aggregate returns-to-scale estimates. Plants that enter during booms are less productive, so entry may cause marginal increases in the outputs of continuing plants to be understated during booms, leading to smaller aggregate estimates of returns to scale. Conversely, because plants that exit during booms are less productive, exit may cause overstated marginal decreases in outputs during booms, leading to larger aggregate estimates of returns to scale.

3 Estimating Returns to Scale: Industry- vs. Plant-level Data

3.1 Estimating Returns to Scale and the Effect of Composition Bias

In this section, I assess the size of the potential bias caused by composition changes. The baseline model for estimating returns to scale follows Basu and Fernald (1997). I assume that a firm's production function for gross output, Y , can be written as a function of labor, L , capital, K , intermediate inputs of materials and energy, M , and the state of technology, T ,

$$Y = F(L, K, M, T). \quad (2)$$

The logarithmic differences of the production function lead to the following equation:

$$d \ln Y = \frac{L}{Y} \frac{\partial F}{\partial L} d \ln L + \frac{K}{Y} \frac{\partial F}{\partial K} d \ln K + \frac{M}{Y} \frac{\partial F}{\partial M} d \ln M + \frac{T}{Y} \frac{\partial F}{\partial T} d \ln T.$$

I further assume that factor markets are competitive and that firms minimize costs. Cost minimization implies that returns to scale (γ) equals the ratio of average to marginal cost; therefore, the equation above can be rewritten as:

$$\begin{aligned}
dy &= \gamma[c_L dl + (1 - c_L - c_M)dk + c_M dm] + dt \\
&= \gamma dx + dt,
\end{aligned} \tag{3}$$

where dl , dk , and dm are the growth rates of L , K , and M , respectively, and c_j is the share of costs for input J in total cost. That is, the growth rate of output, dy , equals the returns to scale multiplied by the cost-share-weighted growth in inputs, dx , plus the productivity growth, dt . Although inputs are plant-specific, I use industry-level input cost shares, averaged over the beginning and ending years of the period of change.

Many researchers, relying on a representative firm framework, have used aggregated data and run regressions similar to equation (3) to estimate returns to scale. To illustrate that these estimates from aggregated data may be subject to a composition bias, consider the case in which equation (3) is estimated using aggregated data. This procedure implicitly assumes the existence of an aggregate production function similar to that of (2), while dt now reflects aggregate productivity changes. As discussed in detail in the previous section, aggregate productivity changes can be decomposed into components that reflect productivity changes within the plant and other components that reflect reallocations of shares across plants, including entry and exit:⁹

$$\begin{aligned}
dy &= \gamma dx + dt \cong \gamma dx + d \ln TFP \\
&= \gamma dx + \sum_{j \in \text{Conti}} \bar{s}_j \Delta \ln tfp_{jt} + \sum_{j \in \text{Conti}} \Delta s_{jt} (\overline{\ln tfp}_j - \overline{\ln TFP}) \\
&\quad + \sum_{j \in \text{Entry}} s_{jt} (\ln tfp_{jt} - \overline{\ln TFP}) - \sum_{j \in \text{Exit}} s_{j,t-1} (\ln tfp_{j,t-1} - \overline{\ln TFP}) \\
&= \gamma dx + \sum_{j \in \text{Conti}} \bar{s}_j \Delta \ln tfp_{jt} + \text{"composition changes"} \\
&= \gamma dx + \sum_{j \in \text{Conti}} \bar{s}_j \Delta \ln tfp_{jt} + \delta dx + \varepsilon.
\end{aligned} \tag{4}$$

The aggregate output growth in equation (4) depends not only on aggregate input growth and aggregated changes in plant-level productivity, but also on the composition changes of producers with different efficiencies. Given that reallocations are negatively correlated with aggregate input changes, a regression run on aggregated data may be subject to a bias caused by composition changes.

⁹ Although the unit of the discussion in this section is “a firm”, the unit of empirical analysis is “a plant”, which refers to a physical location where production takes place. The distinction between firm and plant may not be important if a firm has only one plant. In the case of multi-unit firms, the analysis implicitly assumes that decisions of production are made at the plant level.

In order to assess the size of the bias that may be present in studies using aggregated data, the “composition changes” term calculated from the LRD is run on aggregated input changes, dx , constructed from the publicly available aggregated data. For the total manufacturing samples from the NBER manufacturing database, the regression coefficient, δ , is $-.077$ (see Table 2). As expected, because of the relatively small contribution of plant entry and exit to aggregate productivity growth, the bias caused by net entry (Column 3) is much smaller than the bias caused by between-plant reallocations (Column 2).

Overall, the results in the first row suggest that the effects of composition bias on the returns-to-scale estimates for manufacturing as a whole might not be large. However, they do not necessarily imply that the effect of the bias would be smaller at the disaggregated level. The results for durables and nondurables suggest that the effect of composition bias may be larger, as well as significant at more disaggregated levels of data. Given that previous findings of decreasing returns to scale are based on two-digit industry-level data, this finding suggests that composition bias may be more important at more disaggregated levels of data.

3.2 Estimates of Returns to Scale at Various Levels of Aggregation

One straightforward method of avoiding composition bias is to measure returns to scale at the plant level, giving fixed weights to the exact same plants over time. In this section, the baseline model in Equation (3) is estimated using an ordinary least squares (OLS) regression at various levels of aggregation, from the plant level to the two-digit SIC industries, and to manufacturing as a whole.¹⁰ Assuming that the specification is correct for both the plant- and industry-level regressions, a direct comparison of plant- and industry-level estimates allows an assessment of the size of aggregation bias in estimates obtained from industry-level data. Appropriate caution should be used in interpreting these results, because this assumption may not be warranted. For instance, measurement errors in plant-level variables may cause returns to scale to be understated; if industry-level variation in inputs is correlated with technology changes, the estimates of returns will be biased. The size of a bias, caused by measurement errors,

¹⁰ As pointed out by a number of researchers since the classic study by Marschak and Andrews (1944), production function estimates obtained by the OLS are subject to a simultaneity bias generated by the relationship between productivity and input demands. Since Hall (1988, 1990) and Basu and Fernald (1997), a growing number of studies have generalized the methodology to avoid either measurement errors in factor utilization or the simultaneity bias problem. However, because the study focuses on the effects of composition bias and its primary concern is the differences between estimates at different levels of aggregation, the OLS estimation serves the primary purpose of this study.

misspecification, the effects of technology shocks, or endogeneity in regressions, may vary across different levels of aggregation.

Plant-level Estimates—Correcting Measurement Errors

A potential problem of plant-level analysis is the attenuation bias caused by measurement errors. Previous studies suggest that plant-level returns to scale might be understated by measurement errors present in plant-level hours or capital stocks.¹¹ Because the specification requires measuring changes in inputs and outputs, first-differencing variables may magnify the attenuation bias, leading to a much smaller returns-to-scale estimate.

As a standard response to errors in variables, two instruments are introduced. The first is cost-share-weighted growth in inputs, dx , measured over $t + 1$ and $t - 2$. Given that a firm's input decisions are highly correlated, plant-level input changes between t and $t - 1$ and those between $t + 1$ and $t - 2$ should be highly correlated as well. If measurement errors are not serially correlated, an IV estimation using the instrument will yield consistent estimates of returns to scale. The second instrument is obtained by aggregating the plant-level data. Assuming that measurement errors are canceled out at higher levels of aggregation, plant-level inputs may be aggregated at the four-digit-industry level to calculate cost-share-weighted growth in inputs at that level. Although IV estimation may help reduce the attenuation bias caused by measurement errors, it does not take account of the endogeneity of inputs.

The plant-level estimates are reported in the first three columns of Table 3. Column (1) presents the plant-level OLS result for a pooled sample, which includes plants that have operated for two consecutive years in which they produced nonzero output. The first ASM panel years are excluded in order to avoid sampling issues from panel rotation. Throughout the paper, all plant-level regression results are obtained from weighted regressions using the ASM sampling weight so that the sample is representative of US manufacturing as a whole.

The second and third columns report the results from the IV estimation. Column (2) reports these results using the first instrument, (i.e., the cost-share-weighted input growth

¹¹ See Westbrook and Tybout (1993) and Becker et al. (2005) for evidence of measurement errors in capital. Hansen and Lindstrom (2004) argue that measurement errors in factor inputs can explain the puzzles of rising returns-to-scale estimates at higher levels of aggregation and decreasing returns to scale at the firm level. In the appendix, I present evidence of the attenuation bias caused by measurement errors, following the method of Griliches and Hausman (1986) and Goolsbee (2000). In Table A5, returns-to-scale estimates rise (even within the same set of plants), as changes in inputs and outputs are measured over a longer period of time.

between $t + 1$ and $t - 2$ at the plant level), while Column (3) reports equivalent results using the second instrument, (i.e., the cost-share-weighted input growth between t and $t - 1$ at the four-digit-industry level). Overall, the returns-to-scale estimates are higher than the OLS estimates for all three groups of samples. In most cases, the Hausman specification test rejects the null hypothesis, which would hold that the OLS estimate is consistent. For durables, I find statistically significant increasing returns to scale when the industry-level instrument is used.

Industry-level Estimates and the Effect of Composition Bias

In order to examine the effects of aggregation, the same equation (3) is estimated using the industry-level data, created by aggregating all plants in the industry for a given year. The ASM sampling weight is used to make the aggregated data to mimic the data used in aggregate studies representing the entire industry. Because the industry-level estimation is less likely to be subject to measurement errors, OLS estimates are used as the industry-level estimates.¹²

The productivity decomposition in the previous section provides two channels through which changes in composition can affect aggregate statistics: (1) the reallocation of output shares among continuing plants; and (2) the entry and exit of plants. In order to examine the effects of these two channels separately, plant-level data are aggregated in two different ways: (1) aggregating only continuing plants *excluding* entering and exiting plants; and (2) aggregating all plants *including* entering and exiting plants. The returns-to-scale estimates from the two different sets of aggregated data are presented in the last two columns of Table 3.

First, by comparing plant-level estimates to industry-level estimates based on aggregated data excluding entering and exiting plants (Column 4), I assess the effects of composition bias caused by share changes among continuing plants (i.e., the first channel). Whereas the regression coefficients obtained from the aggregate data reflect the effect of share changes among continuing plants, a regression run on plant-level data would give the same weight to all continuing plants. The countercyclical behavior of the between-plant component in the productivity decomposition predicts a downward bias in industry-level estimates. The industry-level, returns-to-scale estimate, 1.049, is slightly lower than the IV estimates in Column 3, which use the second (industry-level) instrument (1.077). Although the difference between the

¹² As it turns out, the Hausman test suggests that the OLS estimates are not statistically different from the industry-level estimates obtained from IV estimations, using an instrument corresponding to the first instrument in the plant-level analysis (i.e., the plant-level cost-share-weighted input growth between $t + 1$ and $t - 2$).

estimates is not big enough to change the implication of returns to scale for total manufacturing, this downward composition bias seems to have a more significant impact on durables, a category in which changes in productivity due to between-plant reallocations exhibit more countercyclical behavior than do such changes in manufacturing as a whole.¹³ This finding is consistent with the results in Table 2.

Next, I compare the returns-to-scale estimates obtained from aggregate data *excluding* entering and exiting plants (Column 4) to those obtained from aggregate data *including* entering and exiting plants (Column 5), in order to determine the size of the composition bias caused by the second channel of entry and exit of plants. The relatively small contribution of entering and exiting plants to aggregate productivity growth, along with their relatively small output shares, suggests that the biases caused by entry and exit may not be very large. However, there exists a relatively substantial difference between estimates from aggregate data, excluding entry and exit (1.049 for manufacturing, Column 4) on one hand, and estimates from aggregate data covering all manufacturing plants (1.384, Column 5) on the other. Setting aside the small shares of entering and exiting plants, the direction of the bias, which contradicts the prediction of the productivity decomposition exercise, suggests that composition bias alone is unlikely to explain the difference in the returns-to-scale estimates for these two different aggregate data sets.¹⁴

The Impact of Composition Bias within Two-digit Industries

The previous section discussed the effect of composition bias at a higher level of aggregation than that of the two-digit SIC industries. Thus, the effect of reallocations between plants includes the effects of reallocations both *between* and *within* two-digit industries. Although the plant-level evidence suggests that reallocations within industries are countercyclical, previous studies such as Basu and Fernald (1997) find that reallocations between industries are procyclical overall; as inputs are reallocated toward industries with higher returns to scale during a boom, the estimate of returns to scale is higher at the higher level of data aggregation. Given the opposite effect of between-industry reallocations, offsetting the effect of the composition bias caused by within-industry reallocations across plants, it is more relevant to examine the effect of

¹³ The contemporaneous correlations of the between-plant component with changes in real GDP are $-.155$ for the entire manufacturing sector, $-.583$ for durable goods, and $-.082$ for nondurable goods.

¹⁴ While the role of entry and exit merits further investigation, a possible explanation would be that these entering and exiting plants are inefficient, producing in a region in which average and marginal cost are declining (i.e., average cost > marginal cost).

composition changes within an industry. Considering that previous findings of decreasing returns to scale are based on data at the two-digit-industry level, I focus on analysis at this level.¹⁵

Table 4 provides plant-level and industry-level estimates of returns to scale for each two-digit sector, in a manner similar to that used in Table 3. The IV estimates in Column 2 appear in bold if the Hausman specification test rejects the null hypothesis, i.e., consistency of the OLS at the 5% level of significance.¹⁶ The industry-level estimates show wide variation, ranging from .413 for Tobacco (SIC 21) to 1.621 for Electrical Machinery (SIC 36), whereas the plant-level estimates are rather closer to constant returns to scale. Compared to the industry-level estimates (Column 4), the plant-level estimates are smaller in industries with industry-level estimates larger than 1 and larger in industries with industry-level estimates smaller than 1.

The bias implicit in aggregating plant-level data might help resolve the puzzling finding of the decreasing returns to scale in previous studies that used industry-level data. For example, the statistically significant, decreasing returns-to-scale estimates in the Petroleum (SIC 29) and Leather (SIC 31) industries suggest the existence of relatively large positive profits, which seems to contradict previous studies' empirical evidence of a low profit level (Rotemberg and Woodford, 1995; Basu and Fernald, 1997). However, this does not necessarily imply that a typical plant in these industries has decreasing returns to scale, making positive pure profits. Even if an average plant in these industries produces with constant returns to scale (as plant-level estimates suggest), aggregation may create a bias in the aggregate estimates and lead to a different implication than would the true returns to scale of an average plant. This finding suggests that differences in industry-level estimates of returns to scale across industries may reflect differences in the size of the bias caused by within-industry reallocations, rather than the between-industry differences in returns to scale of an average plant. Whether the industry-level estimate is larger or smaller than true returns to scale will depend on the cyclical behavior of reallocations within the industry.

¹⁵ Basu and Fernald also find that correcting for reallocations in aggregate data does not fully recover industry-level averages of returns-to-scale estimates. They interpret this to mean that there exist additional aggregation effects, running from the level of firms to the level of individual two-digit industries.

¹⁶ Because the Hausman test suggests that the IV estimates using the industry-level instruments are not statistically different from the OLS in most two-digit industries, these estimates are not reported.

4 Conclusion

In examining longitudinal plant-level data in U.S. manufacturing, I find that actual productivity may be more procyclical than observed aggregate productivity. As reallocations among producers over the business cycle create a countercyclical component of aggregate productivity, aggregate productivity exhibits less procyclicality than the true procyclicality of productivity observed in the case of a typical producer. Without correcting for such a countercyclical composition bias, technology shock measures based on aggregated data may understate the cyclicity of the technology shocks that a representative agent experiences over the business cycle.

Composition bias, caused by countercyclical reallocations within an industry, helps explain the finding of decreasing returns to scale at the industry level of data. However, the lack of evidence for important increasing returns to scale suggests that other factors such as technology shocks or cyclical utilization may be more important factors in the cyclical behavior of productivity.

Caution should be used in interpreting differences in plant-level productivity, because it is uncertain how much of the difference in TFP across plants is explained by differences in the quality of inputs. Further investigation on this issue will illuminate how factors are reallocated over the business cycle.

REFERENCES

- Baily, Martin N., Bartelsman, Eric J., and Haltiwanger, John C. "Labor Productivity: Structural Change and Cyclical Dynamics," *Review of Economics and Statistics*, 83(3), August 2001: 420–33.
- Baily, Martin N., Hulten, Charles, and Campbell, David. "Productivity Dynamics in Manufacturing Plants," *Brookings Papers on Economic Activity: Microeconomics* 1992: 187–249.
- Bartelsman, Eric J., and Doms, Mark. "Understanding Productivity: Lessons from Longitudinal Micro Datasets," *Journal of Economic Literature* 38(3), September 2000: 569–94.
- Bartelsman, Eric J. and Gray, Wayne. "The NBER Manufacturing Productivity Database," NBER Technical Working Paper 205, 1996.
- Basu, Susanto, and Fernald, John G. "Returns to Scale in U.S. Production: Estimates and Implications," *Journal of Political Economy* 105(2), April 1997: 249–83.
- Basu, Susanto, and Fernald, John G. "Aggregate Productivity and Aggregate Technology," *European Economic Review* 46(6), June 2002: 963–91.
- Basu, Susanto, Fernald, John G., and Kimball, Miles "Are Technology Improvements Contractionary?" *American Economic Review*, forthcoming.
- Bils, Mark. "Real Wages Over the Business Cycles: Evidence From Panel Data." *Journal of Political Economy* 93(4), August 1985: 666–89.
- Bils, Mark, and Chan, Yongsung. "Understanding How Price Responds to Costs and Production." *Carnegie-Rochester Conference Series on Public Policy* 52, June 2000: 33–78.
- Campbell, Jeffrey R. "Entry, Exit, Technology and Business Cycles," *Review of Economic Dynamics* 1(2), April 1998: 371–408.
- Chang, Yongsung. "Wages, Business Cycles, and Comparative Advantage," *Journal of Monetary Economics*, 46(1), August 2000: 143–71.
- Chang, Yongsung and Hong, Jay H. "Do Technological Improvements in the Manufacturing Sector Raise or Lower Employment?" *American Economic Review*, forthcoming.
- Cooley, T., ed. *Frontiers of Business Cycle Research*. Princeton: Princeton University Press, 1995.
- Cooper, R., and Chatterjee, S. "Entry and Exit, Product Variety and the Business Cycle," NBER Working Paper 4562, 1993.

- Davis, Steven J., Haltiwanger, John C., and Schuh, Scott. *Job Creation and Destruction*, Cambridge, MA. The MIT Press, 1996.
- Devereux, M., Head, A., and Lapham, B. "Aggregate Fluctuations with Increasing Returns to Specialization and Scale," *Journal of Economics Dynamics and Control*, 20(4), April 1996: 627–57.
- Foster, Lucia, Haltiwanger, John C., and Krizan, C.J. "Aggregate Productivity Growth: Lessons from Microeconomic Evidence," pp. 303–63, in Hulten, Charles R.; Dean, Edwin R.; and Harper, Michael J., eds., *New Developments in Productivity Analysis*, University of Chicago Press, 2001.
- Goolsbee, Austan. "The Importance of Measurement Error in the Cost of Capital," *National Tax Journal*, 53(2), June 2000: 215–28.
- Griliches, Zvi, and Hausman, Jerry. "Errors in Variables in Panel Data." *Journal of Econometrics*, 31(1), February 1986: 93–118.
- Griliches, Zvi, and Regev, Haim. "Firm Productivity in Israeli Industry: 1979–1988," *Journal of Econometrics*, 65, January 1995: 175–203.
- Hall, Robert E. "The Relation between Price and Marginal Cost in U.S. Industry," *Journal of Political Economy*, 96(5), October 1988: 921–47.
- Hall, Robert E. "Invariance Properties of Solow's Productivity Residual," in P. Diamond, ed., *Growth/Productivity/Unemployment*, M.I.T. Press: Cambridge, 1990: 71–112.
- Haltiwanger, John. "Measuring and Analyzing Aggregate Fluctuations: The Importance of Building from Microeconomic Evidence," Federal Reserve Bank of St. Louis, *Economic Review* 79(3), May–June 1997: 55–77.
- Klette, Tor J. and Griliches, Zvi. "The Inconsistency of Common Scale Estimators when Output Prices are Unobserved and Endogenous," *Journal of Applied. Econometrics* 11(4), July–August 1996: 343–61.
- Marschak, Jacob, and Andrews, William H. "Random Simultaneous Equations and the Theory of Production." *Econometrica* 12(3/4), July–October 1944: 143–205.
- Rotemberg, J., J., and Woodford, M.: "Dynamic General Equilibrium Models with Imperfectly Competitive Product Markets," In *Frontiers of Business Cycle Research*, ed. T.F. Cooley. Princeton University Press :Princeton, N.J, 1995.
- Solon, G., Barsky, R., and Parker, J. "Measuring the Cyclicalities of Real Wages: How Important is Compositional Bias?" *Quarterly Journal of Economics* 109(1), February 1994:587–616.

Stockman, A., "Aggregation Bias and the Cyclical Behavior of Real Wages." Manuscript, University of Rochester, 1983.

Westbrook, M.D., and J.R. Tybout, "Estimating Returns to Scale with Large Imperfect Panels: An Application to Chilean Manufacturing Industries," *World Bank Economic Review*, January 1993.

Figure 1: Entry and Exit Rate

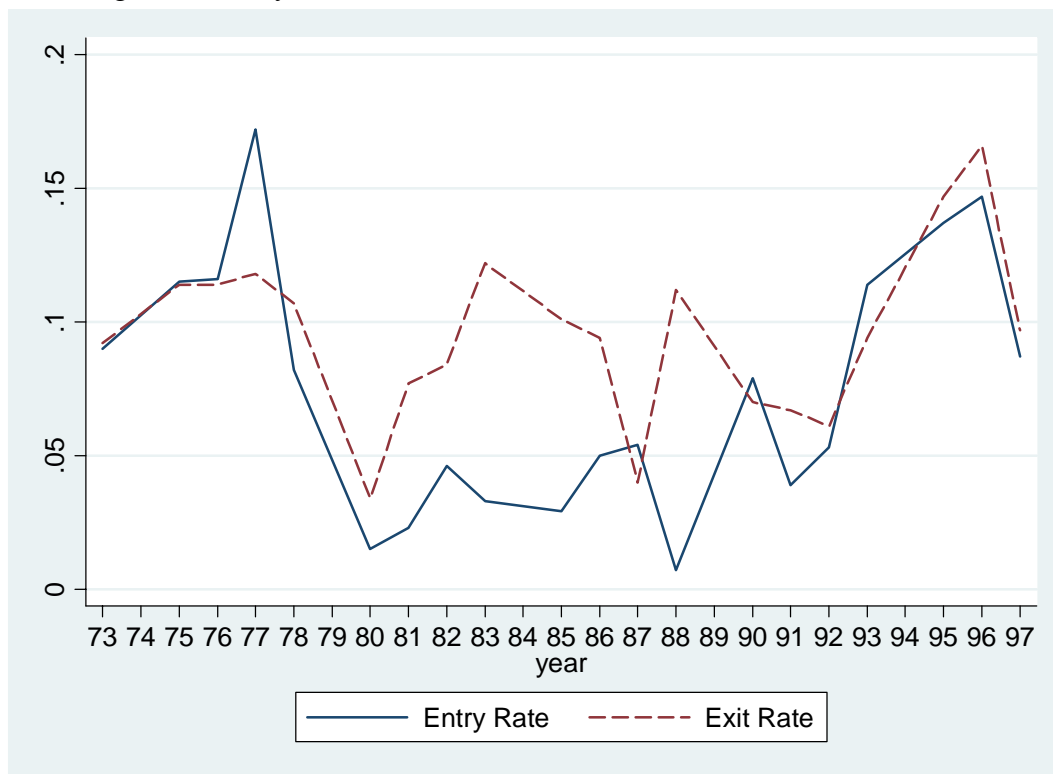



Table 1: Shares and Relative TFP of Entrants and Exiting Plants

	Shares						Relative TFP	
	Number (entry/exit rate)		Employment		Output			
	Entrants	Exiting Plants	Entrants	Exiting Plants	Entrants	Exiting Plants	Entrants	Exiting Plants
All sample years	.074	.096	.029	.023	.024	.031	1.094	.977
Boom	.077	.108	.025	.019	.021	.034	.981	.935
Recession	.054	.075	.026	.026	.020	.026	1.072	.974

Note: Boom: log change of real GDP > 4%

Recession: log change of real GDP < 1%

Figure 2: Productivity Decompositions



Note 74, 79, 84, 89, & 94 interpolated in the graph (panel rotation)

Table 2: Composition Bias in Returns-to-Scale Estimates

	[1]	[2]	[3]
	<i>Composition Changes</i> (Between & Net Entry)	$\sum_{j \in \text{Conti}} \Delta s_{jt} (\ln \overline{tfp}_j - \ln \overline{TFP})$ (Between-plant)	$\sum_{j \in \text{Entry}} s_{jt} (\ln tfp_{jt} - \ln \overline{TFP})$ $- \sum_{j \in \text{Exit}} s_{j,t-1} (\ln tfp_{j,t-1} - \ln \overline{TFP})$ (Net Entry)
Manufacturing			
δ (Std. Err)	-.077 (.097)	-.098 (.082)	.020 (.032)
Num. of obs.	19	19	19
Nondurables			
δ (Std. Err)	-.128 (.218)	-.144 (.204)	.017 (.046)
Num. of obs.	19	19	19
Durables			
δ (Std. Err)	-.182 (.070)	-.192 (.060)	.010 (.025)
Num. of obs.	19	19	19

Note: The independent variable is a subset of “*composition changes*” term in Equation (4), stated in the column heading. This measure is calculated from the LRD. The dependent variable is the cost-share weighted change in inputs (dx), measured in the NBER manufacturing database. The sample period is 1972–96, excluding the years 1974, 1979, 1984, 1989, and 1994. The coefficients of the constant terms are not reported.

Table 3: Returns-to-Scale Estimates at Different Levels of Aggregation

	[1]	[2]	[3]	[4]	[5]
	Plant level (pooled)	Plant level (pooled)	Plant level (pooled)	Aggregate (continuing plants only)	Aggregate (all plants)
	OLS	IV (plant level) ^a	IV (industry level) ^b	OLS	OLS
Manufacturing					
γ (Std. Err)	.828 (.006)	.910 (.012)	1.077 (.128)	1.049 (.078)	1.384 (.228)
Num. of obs.	1,078,471	655,350	1,078,471	20	20
Hasuman test statistics (chi2)		263.73	3.81		
Nondurables					
γ (Std. Err)	.790 (.008)	.862 (.019)	.781 (.278)	.990 (.060)	.770 (.196)
Num. of obs.	490,635	306,870	490,635	20	20
Hasuman test statistics (chi2)		79.47	0.00		
Durables					
γ (Std. Err)	.853 (.007)	.941 (.015)	1.242 (.124)	1.080 (.094)	1.503 (.216)
Num. of obs.	587,836	348,480	587,836	20	20
Hasuman test statistics (chi2)		186.21	9.88		

Note: ASM sample weight is used. The sample period is 1972–97, excluding the years 1974, 1979, 1984, 1989, and 1994.

a) Plant-level instrument: Plant-level changes in input between $t + 1$ & $t - 2$

b) Industry-level instrument: Four-digit, industry-level changes in input between t and $t - 1$

Table 4: Returns-to-Scale Estimates at Different Levels of Aggregation, Two-digit SIC

A. Nondurables

			[1]	[2]	[3]	[4]
			Plant level (pooled)	Plant level (pooled)	Aggregate (continuing plants only)	Aggregate (all manufacturing plants)
SIC code	Industry		OLS	IV	OLS	OLS
20	Food	γ (Std. Err)	.688 (.021)	.741 (.041)	1.002 (.044)	.632 (.154)
		Num. of obs.	118,839	79,337	20	20
21	Tobacco	γ (Std. Err)	.682 (.078)	.659 (.082)	.642 (.099)	.413 (.214)
		Num. of obs.	1,389	1,030	20	20
22	Textiles	γ (Std. Err)	.876 (.030)	.868 (.036)	1.083 (.076)	1.022 (.132)
		Num. of obs.	37,572	25,598	20	20
23	Apparel	γ (Std. Err)	.837 (.017)	.869 (.029)	1.040 (.035)	.833 (.110)
		Num. of obs.	62,086	33,192	20	20
26	Paper	γ (Std. Err)	.904 (.028)	.971 (.040)	.876 (.072)	1.053 (.212)
		Num. of obs.	45,584	32,434	20	20
27	Printing	γ (Std. Err)	.729 (.020)	.806 (.054)	.900 (.033)	.584 (.238)
		Num. of obs.	75,493	38,712	20	20
28	Chemicals	γ (Std. Err)	.864 (.034)	.972 (.096)	.737 (.200)	.465 (.303)
		Num. of obs.	68,289	45,441	20	20
29	Petroleum	γ (Std. Err)	.874 (.056)	.983 (.062)	.562 (.148)	.425 (.169)
		Num. of obs.	15,455	10,200	20	20
30	Rubber	γ (Std. Err)	.855 (.020)	.996 (.031)	1.065 (.038)	1.271 (.098)
		Num. of obs.	55,865	34,389	20	20
31	Leather	γ (Std. Err)	.897 (.059)	1.060 (.082)	1.007 (.068)	.836 (.136)
		Num. of obs.	10,066	6,537	20	20

* Plant-level IV estimates appear in bold if the Hausman test rejects the consistency of the corresponding OLS estimates at the 5% level of significance.

B. Durables

			[1]	[2]	[3]	[4]
			Plant level (pooled)	Plant level (pooled)	Aggregate (continuing plants only)	Aggregate (all manufacturing plants)
SIC code	Industry		OLS	IV	OLS	OLS
24	Lumber	γ (Std. Err)	0.745 (.017)	.803 (.048)	1.006 (.025)	.885 (.127)
		Num. of obs.	67,711	35,644	20	20
25	Furniture	γ (Std. Err)	0.919 (.022)	1.000 (.049)	1.046 (.029)	1.296 (.105)
		Num. of obs.	28,549	16,290	20	20
32	Stone, Clay, & Glass	γ (Std. Err)	0.990 (.021)	1.074 (.067)	.968 (.045)	1.109 (.219)
		Num. of obs.	28,549	29,721	20	20
33	Primary Metals	γ (Std. Err)	0.792 (.037)	.914 (.032)	1.125 (.100)	1.212 (.099)
		Num. of obs.	40,854	27,782	20	20
34	Fabricated Metals	γ (Std. Err)	0.840 (.016)	.952 (.032)	.979 (.042)	1.293 (.202)
		Num. of obs.	112,483	66,116	20	20
35	Nonelectrical Machinery	γ (Std. Err)	0.872 (.016)	.943 (.034)	1.137 (.104)	1.608 (.181)
		Num. of obs.	118,588	68,887	20	20
36	Electrical Machinery	γ (Std. Err)	0.927 (.018)	1.029 (.035)	1.601 (.189)	1.621 (.264)
		Num. of obs.	69,047	45,032	20	20
37	Transportation Equipment	γ (Std. Err)	0.889 (.023)	.997 (.066)	1.142 (.091)	1.207 (.077)
		Num. of obs.	39,173	25,331	20	20
38	Instruments	γ (Std. Err)	0.855 (.030)	.845 (.050)	.626 (.094)	.524 (.122)
		Num. of obs.	31,455	19,733	20	20
39	Miscellaneous Durables	γ (Std. Err)	0.905 (.043)	.978 (.056)	.986 (.032)	.887 (.175)
		Num. of obs.	26,491	13,944	20	20

* Plant-level IV estimates appear in bold if the Hausman test rejects the consistency of the corresponding OLS estimates at the 5% level of significance.

Appendix

Table A5: Estimates of Returns to Scale over Different Time Horizons

A. Plant-level pooled regressions

	[1]	[2]	[3]	[4]
	$t \& t-1$	$t \& t-2$	$t \& t-3$	$t \& t-4$
γ (Std. Err)	0.835 (0.005)	0.866 (0.005)	0.893 (0.006)	0.924 (0.005)
Num. of obs.	1,234,619	989,255	798,207	634,096

B. Plant-level pooled regressions with the same sample of continuing plants

	[1]	[2]	[3]	[4]
	$t \& t-1$	$t \& t-2$	$t \& t-3$	$t \& t-4$
γ (Std. Err)	0.889 (0.009)	0.915 (0.007)	0.924 (0.006)	0.930 (0.005)
Num. of obs.	632,268	632,268	632,268	632,268

Note: The dependent variable is the log change in real output measured over the stated time period in the column heading. The independent variable is the cost-share weighted change in inputs over the same, stated time period. ASM sample weight is used.

In panel B, the same regression is run for the same sample of plants that have operated for at least four consecutive years, to exclude the effects of sample changes due to changes in the time period.