Macro and Micro Dynamics of Productivity: Is the Devil in the Details?*

Lucia S. Foster[a], Cheryl A. Grim[a], John Haltiwanger[b,a], Zoltan Wolf[c†]

[a]U.S. Census Bureau, [b]University of Maryland, [c]Westat

March 30, 2015

Abstract

Does the method of estimating plant-level productivity matter? The paper attempts to answer this question in the context of key stylized facts and popular estimation methods. Using data from the Annual Survey of Manufactures and Census of Manufactures between 1972-2010, we test the robustness of results on four dimensions. First, we find non-trivial differences in estimated factor elasticities, especially for capital, across commonly used methods. Second, the variation in elasticities maps into differences in (total factor) productivity dispersion but does not invalidate the general conclusion that productivity differences across establishments within the same industry are large. However, the quantitative differences in within industry dispersion are substantial and the ranking of plants by productivity within industries is also sensitive to method. Third, more productive plants are shown to be more likely to grow and survive in our data, no matter how we estimate productivity. However, outliers in factor elasticities that arise more frequently from some methods non-trivially impact the quantitative margin effects of productivity on growth and survival. Fourth, all our total factor productivity variants confirm the main conclusions from the structural productivity decomposition literature: reallocation is productivity enhancing, and within-plant productivity seems more important in terms of cyclical fluctuations by all methods considered. However, here again there are non-trivial quantitative differences across methods in the contribution of reallocation to aggregate productivity growth. Some methods imply that all or even more than all of aggregate productivity growth is due to reallocation while other methods imply only 25 percent is due to reallocation. Finally, we find that the basic patterns are robust to the use of imputed plant-level data.

*We would like to thank conference participants at the CAED 2013 Atlanta, RDC 2014 for valuable comments. We are grateful to Kirk White for useful discussions and for making his code available to us. Any remaining errors are our own. Any conclusions expressed herein are those of the authors and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed.

†Corresponding author: zoltan.wolf@census.gov.

1 Introduction

A ubiquitous and influential finding in the empirical literature on firm dynamics is that there is large dispersion in measured productivity across establishments within narrowly defined industries. This finding has generated much analysis of the causes and consequences of such dispersion. In terms of causes, explanations range from curvature in the profit function so that the most productive firm does not take over an industry, to frictions in adjustment of factors and the entry and exit of plants, and to distortions that drive wedges in the forces pushing towards the equalization of marginal products across plants. In terms of consequences, there is a burgeoning literature on the connection between reallocation dynamics, growth and productivity. Many papers have found that more productive plants are more likely to grow and less likely to exit. This implies that the high pace of reallocation of inputs observed across plants has been found to be productivity enhancing. In like fashion, there is increased attention to reasons why these reallocation dynamics may vary over the cycle and across countries and in turn how this account for differences in economic performance across time and countries. In an important related area of inquiry, a new theoretical and empirical literature has developed that hypothesizes that gains from opening markets to trade are due to the improved allocation of resources across plants from the reallocation induced by trade.[1]

While there is considerable consensus that taking into account the dispersion of productivity and its connection to reallocation are important for variation in economic performance across countries, industries and time, there is not a consensus about the basics of estimating plant-level productivity. One core issue that is still being debated is the most appropriate method for estimating productivity – specifically, the method for estimating factor elasticities. One commonly used approach is to use growth accounting methods that are actively used by the statistical agencies in official aggregate and industry level productivity statistics. Such methods have advantages in terms of ease of computation but rely on strong assumptions. But another commonly used approach by researchers is to use econometric estimation methods. Since OLS is evidently problematic given endogeneity of factor outputs, alternative estimation methods have been developed to address this endogeneity. While instrumental variable methods are attractive in principle, they are not commonly used given the lack of plausible instruments on a wide scale basis to cover all industries over all time periods. Instead, methods have been developed that seek to deal with the endogeneity issues using proxies for the productivity residuals. Since this is an indirect method, numerous alternative proxy estimation approaches have been developed with a robust debate about the respective merits of the each approach. Different approaches make alternative assumptions about the timing of input decisions with respect to the realization of shocks and the relative importance of addressing the impact of selection on estimation.

[1] See the recent survey in Syverson (2011) for relevant cites to the findings in the literature and the theoretical and empirical literature that has developed in light of the large dispersion in productivity and and its relationship to reallocation, growth and aggregate productivity.

In this paper, we focus our attention on these estimation issues. Our objective is to assess how sensitive the basic facts about productivity dispersion and the connection between dispersion, reallocation, growth and productivity are to alternative methods for estimating productivity. Our main findings are summarized as follows. First, we find that there are large differences in estimated factor elasticities across methods. This variation is particularly evident in capital elasticities which in turn yields much variation in estimated returns to scale. Differences across methods are in both the first and second moments of the distribution of elasticities across industries. Second, in spite of these differences we find that basic facts about productivity dispersion are largely robust across methods. All methods yield considerable within industry dispersion of productivity. But there are quantitative differences and since the precise estimate of productivity dispersion matters for some purposes this variation is potentially quite important. A related finding is that the productivity rank of plants within industries varies across estimation methods. Third, the relationship between growth, survival and productivity at the establishment level is largely robust across methods. But even here we find that some of the methods yield large outliers in elasticities for specific industries. For example, some proxy methods yield negative factor elasticities for specific factors and industry combinations. Not surprisingly the relationship between growth, survival and plant level productivity is sensitive to the treatment of these outliers. Fourth, all methods of estimation yield a substantial contribution of reallocation to aggregate productivity growth using the structural decompositions developed by Petrin and Levinsohn (2012). But details again differ in terms of quantitative significance. Some methods yield that all or even more than all of the aggregate productivity growth is accounted for by reallocation effects while other methods yield that only about 25 percent is accounted for by reallocation. It is not surprising that there is quantitative variation here since measuring the quantitative contribution of reallocation relies on deviations of plant-level cost shares from estimated factor elasticities. Variation in the latter yields variation in the contribution of reallocation.

We think the robustness of the core findings in the literature to these estimation methods have not been settled. The literature offers guidelines, based on Monte Carlo evidence[2], as to which method is optimal in the presence of certain types of measurement and specification error. Important examples include heterogeneity in output and input prices, heterogeneity in technologies or measurement error in output, primary inputs, or some combination of these. These experiments tell us which estimation method is optimal conditional on isolated factors. While such evidence is valuable, in practice there may be a variety of sources of error. An equally important point is that existing Monte Carlo studies do not take into account the wide disparities in sample size that is available for empirical estimation. A limitation of our approach is that we cannot and do not take a stand on the optimal method. Instead, our objective is to

[2]One example is Biesebroeck (2007).

identify the nature of the sensitivity of key results to these estimation issues.

The paper is organized as follows. We discuss our methodology and data in Sections 2 and 3. Section 4 summarizes our results and section 5 concludes.

2 Methodology

2.1 Definition of TFP

As most empirical studies, we start with a single-output production function:

$$Q_{it} = K_{it}^{\beta_K} L_{it}^{\beta_L} E_{it}^{\beta_E} M_{it}^{\beta_M} \Omega_{it}, \tag{1}$$

where Q, K,L,E,M denote output, capital stock, labor, energy and material inputs, respectively. i and t index plants and time periods. The β-s denote the elasticity of Q with respect to respective inputs. It is then straightforward to define total factor productivity (TFP) as a ratio of output and an index of inputs: $\Omega_{it} = Q_{it}/(K_{it}^{\beta_K} L_{it}^{\beta_L} E_{it}^{\beta_E} M_{it}^{\beta_M})$. The input index is a weighted average of primary input factors where the weights are the elasticities of output with respect to the appropriate input factor.

2.2 Elasticities

Elasticities are unobserved and therefore have be estimated. There are various methods available to researchers, table 1 summarizes the methods we use throughout the paper. We discuss the strengths and weaknesses of the most popular ones from a practical point of view.

Table 1: Production function estimation methods.

Method	Description	Dependent variable	Proxy	Estimator	Numerical procedure*
OLS	Ordinary Least Squares	Output		LS	
GA	Foster et al. (2001)	Output		Cost-shares	
OP	Olley and Pakes (1996)	Output	investment	LS	NL
LPVA	Levinsohn and Petrin (2003)	Value Added	materials	GMM	GSS
LPNL	Levinsohn and Petrin (2003)	Output	materials	GMM	NL
LPGR	Levinsohn and Petrin (2003)	Output	materials	GMM	GS
LPGSS	Levinsohn and Petrin (2003)	Output	materials	GMM	GSS
WLPE	Wooldridge (2009)	Output	energy	efficient GMM**	
WLPM	Wooldridge (2009)	Output	materials	efficient GMM	

*The last column lists optimization procedures used in the paper. NL: gradient-based nonlinear technique, GSS: Golden Section Search, GR: Grid Search. **Efficient GMM implements the two-step efficient GMM estimator which minimizes the GMM criterion function $Q = Nm'Wm$, where N denotes sample size, m denotes the matrix of orthogonality conditions and W is an optimal weighting matrix. In 2-step GMM, W is chosen to be the inverse of an estimate of the covariance matrix of moment conditions.

Ordinary least squares (OLS) is a straightforward but naïve method. OLS-based estimates of the elasticities are inconsistent because TFP, unobserved by the econometrician, is a state variable in the decision problem of plants. As first pointed out by Marschak and Andrews (1944), an endogeneity problem emerges because unobserved TFP is incorporated in the error

term, which renders OLS estimates biased.

Other approaches use more explicit assumptions about the environment in which the plant operates. A frequently used non-statistical method in this class exploits the first order condition of a decision problem where the plant minimizes production costs given output and input prices. The first order condition of this problem is used to rewrite elasticities as respective cost shares of input factors in the plant's revenues. This method is often referred to as growth accounting method (GA) after the seminal work of Solow (1957). The main advantages of this procedure are: allows for plant-level heterogeneity in elasticities, it is easy to implement and flexible about the exact shape of production technology, accurate if data are not subject to much measurement error[3]. An important practical advantage is that it is free of statistical problems related to the sensitivity of estimates to sample size. A potential caveat is that elasticity estimates may be biased if the first order conditions are violated. Since there are many frictions that likely yield that the first order conditions do not hold at all instants in time at least at the plant level, it is common (see Syverson (2011)) to impose common factor elasticities across plants in the same industry and/or common factor elasticities over time within the same industry. We discuss the implications of such homogeneity assumptions further below.

The third group is commonly referred to as proxy methods. The original idea of using proxies in production function estimation was developed in Olley and Pakes (1996) (OP hereafter) in order to analyze the dynamics of the telecommunications equipment industry. OP take account of the previously mentioned endogeneity problem by including an investment proxy in the estimation process. Assuming that investment is a monotonic and increasing function of productivity and that productivity is the only unobserved state variable, including investment controls for unobserved TFP developments, and the variation in investment can be used invert out TFP shocks. OP focus on the period between the early 1970s and the mid-1980s during which period the telecommunications equipment industry saw large changes in the size of plants and significant entry and exit. Therefore they model plants' entry and exit decisions which depend on productivity.

The algorithm consists of multiple steps. The first step provides consistent OLS estimates of variable input elasticities because even though variable inputs respond to unobserved TFP developments the proxy controls for their effect during the estimation procedure. However, the coefficient of capital is not identified in this step because the control is included as a polynomial of the proxy and capital. Hence the need for a second step. Once variable input elasticities are obtained, the algorithm exploits a Markovian assumption about the plant-level productivity process to extract the unanticipated TFP shocks. Since capital is assumed to be predetermined[4], its value is orthogonal to the innovation in TFP. The algorithm exploits this condition to identify the elasticity of capital using lagged capital values as instrumental

[3]See **?**.

[4]Time-to-build lags may justify this assumption.

variables.[5]

Proxy methods use polynomial series at two points of the estimation algorithm. First, a polynomial of the state variables and the proxy is included in the first step to approximate unobserved productivity. Second, to determine the expected component of TFP, its estimated value is projected on a polynomial expansion of its past values[6]. While polynomial series provide flexible approximations, the higher order terms are also likely to exacerbate measurement error present in microdata.

OP propose using investment to proxy for unobserved productivity. There is ample evidence that plant-level investment is lumpy.[7] Lumpiness means bursts of investment activity are followed by inactive periods where observed net investment is zero. It is a consequence of the presence of non-convexities in capital adjustment. Unfortunately, zero investment observations are not informative for OP and are dropped, which may negatively affect efficiency if truncation significantly decreases sample size. In addition, since in OP's model entry and exit has a central role, their approach works only if we observe entrants and exiters. Therefore cannot be used in industries without data on entrants and exiters. As we will see in section 4.1, there may be a significant number of industries where these issues become relevant.

In order to eliminate the efficiency loss caused by dropping zero-investment observations, Levinsohn and Petrin (2003) (LP hereafter) advocate the use of intermediate input cost or electricity instead of investment. LP discuss the conditions which must hold if the intermediate input is to be used as a proxy. The basis of the argument is that if intermediate inputs are less costly to adjust than investment, especially in the presence of non-convex capital adjustment costs, it is likely to respond more to productivity. They also highlight that firms almost always report positive use of these variables in their data[8] implying truncation due to zero proxy values is less severe.

OP's effort to control for selection was motivated by the importance of not restricting ourselves to using only continuing establishments because plants' exit decisions depend on productivity. LP do not focus on selection issues because their panel is unbalanced.[9] As we will discuss in more detail in section 3, our data is subject to some degree of non-randomness because larger establishments are more likely to be sampled in the ASM. If size and productivity are

[5]We note that available Stata routines were programmed to use different numerical techniques to estimate β_k. OP's Stata commands are based on nonlinear least squares (NL). LP's implementation offers three options to minimize the GMM criterion function. The default setting assumes value added is the dependent variable and uses the so-called golden section search algorithm (GSS). If the dependent variable is revenue, nonlinear least squares (NL) or grid search (GR) can be requested. Both NL and GSS guarantee to find optimum points if the objective function is unimodal. If the criterion function has multiple modes, GR can be used to confirm global optimum. However, GR is more demanding computationally, especially if we are to estimate elasticities for many industries.

[6]Exploiting the assumption about its Markovian property.

[7]See for example, Cooper and Haltiwanger (2006).

[8]Their 8-year panel contains a census of Chilean manufacturing plants with at least 10 employees.

[9]They also note their sample is representative of the Chilean Manufacturing sector.

correlated, OP's arguments about selection become relevant and we may expect that controlling for it has an effect on our results.

The identifying assumptions regarding the timing of plants' input decisions have been criticized by Ackerberg et al. (2006) (ACF). ACF highlight that the optimal labor allocation is also a deterministic function of TFP and therefore the elasticity of labor is not identified in the first step. They propose a hybrid approach and offer structural assumptions on the timing of decisions concerning firms' input choices. They go about the identification problem by applying a two step procedure that does not try to identify any of elasticities in the first stage. Wooldridge (2009) proposed to circumvent the identification problem by estimating all the coefficients in a single GMM step and using earlier outcomes of both capital and variable inputs as instrumental variables. His approach is advantageous because it is robust to the ACF critique and because the efficiency loss due to two-step estimation is eliminated.[10]

3 Data

Our industry-level data, including deflators, capital rental prices and depreciation rates, source from the NBER-CES Manufacturing database[11], the Bureau of Labor Statistics and the Bureau of Economic Analysis. We extract information on establishments' output and inputs from the Annual Survey of Manufactures (ASM), Census of Manufactures (CM) and the Longitudinal Business Database (LBD).

Contrary to what the name suggests, CM does not collect data for the universe of manufacturing plants because establishments with less the five employees are not sent forms. Payroll and employment data for these very small plants are imputed using administrative records.[12] The ASM surveys 50,000-70,000 manufacturing establishments in non-Census years and is part of the CM in Census years. It is a rotating panel re-defined two years after the latest Census. The LBD contains the universe of non-agricultural business establishments with paid employees and is based on both survey information and administrative records. Appendix A in Foster et al. (2013) (FGH, hereafter) describes these data in more detail. Our initial dataset includes approximately 3.5 million plant-year observations between 1972-2010.

We use the ASM and CM to construct plant-level measures of inputs and output. Output is measured as a deflated value of total value of shipments, corrected for the change in finished goods and work-in-process inventories. Labor input (total hours worked) is constructed as the product of production worker hours and the ratio of salaries and wages to production worker wages. Our intermediate input variable is given by the deflated sum of cost of parts, contracted work and goods resold. The energy input consist of deflated electricity and fuel costs. We

[10] Two step estimators are inefficient because the correlations in the errors across equations and possible heteroskedasticity are ignored.

[11] http://www.nber.org/nberces and earlier version is documented in Bartelsman and Gray (1996).

[12] We drop these cases because TFP cannot be calculated for them.

created establishment-level capital stock measures using a version of the Perpetual Inventory Method, which calculates current capital as a sum of depreciated stock and current investment. We set plants' initial capital stock to a deflated book value taken from the ASM and CM. More details on the construction of input and output measures can be found in appendix B of FGH.

The LBD serves two purposes in our analysis. First, high-quality longitudinal identifiers help us determine the accurate time of establishments' exit which is a necessary indicator to estimate the relationship between productivity, growth and exit. Second, we use employment and establishment age data to construct inverse propensity scores weights that control for non-randomness in our sample.[13]

The sample for the exercises about dispersion and growth should fulfill two potentially contradicting requirements. First, the number of plant-year observations within each 4-digit industry should be large enough so that elasticities can be estimated by all reviewed methods. Second, industries should be defined narrowly enough so that we can plausibly assume elasticities are constant among establishments.

An additional complicating issue is related to the changes in industry identifiers. Since we estimate elasticities on an industry-by-industry basis, changes[14] in the classification system entail spurious breaks in plants' timeseries and a drop in sample size. In the first part of the analysis we address these issues by selecting multiple sets of 4-digit industries which were not affected by changes or which were mapped one-to-one into another industry.[15] There are 292 such industries of which we selected the 10 and 50 largest.

One of the objectives in this paper is to test whether the implications of the decomposition of aggregate productivity growth described in Petrin et al. (2011) (PWR hereafter) are sensitive to the way TFP is estimated. PWR's data spans the period between 1976-1996 so the 1987-change in SIC classification is relevant. To correct for these breaks, we follow the first step of PWR's procedure and assign the SIC code to any establishment observed between 1987-1996. However, we deviate from their approach for cases observed prior to 1987. If a plant is not assigned an industry code in the previous step, we apply a random assignment procedure. The basic idea is to choose from among plants such that the share of re-assigned plants matches the appropriate share in the concordance. Randomness is necessary in order to avoid that the procedure is dominated by a few large establishments. As an illustration, suppose the mapping says 10% of industry i's (SIC 1972) TVS should be mapped into industry j (SIC 1987). First, compute the time-average of each plant's share in the total value of shipments of industry i and then randomly sort them by these averages. Next, calculate the cumulative sum of shares

[13]Employment data is useful to determine the probability of size-based selection into the ASM and CM. Establishment age is an important determinant of the probability that the TFP of an establishment is calculated from imputed data.

[14]Changes in SIC in 1987, and the change to NAICS in 1997.

[15]The NBER-CES crosswalk lists how much of industry i's total value of shipments should be mapped into industry j.

and find the first n plants for which the sum does not exceed 10%. These establishments are classified in industry j.[16]

4 Results

The main question in this paper is whether the method of estimating TFP matters. This section breaks down our answers by questions about: elasticities; TFP dispersion; the relationship between productivity, growth and survival; and structural decompositions of aggregate productivity growth (APG). The final question is whether our findings are robust to the presence of imputed data. We organize the discussion around the various estimators shown in table 1.

4.1 Distributions of elasticities

The exercises in this section attempt to gauge the effect of the estimation method on the basic characteristics of elasticity distributions. One way to evaluate the methods would be to compare the sign and magnitude of the bias they entail. Unfortunately, the differences among empirical distributions alone do not tell much about the magnitude of the bias because the population distributions are unknown. However, in light of what we know about these methods address the bias in OLS estimates, the differences may give us clues whether they correct it in the right direction. We start by discussing standard descriptive statistics. Next, we check whether the ranking of industries is invariant to the estimator choice. We conclude the section by commenting on the relationship between sample size and plausibility of elasticities and, finally, their implications for returns-to-scale.

The direction of the bias in OLS estimates is determined by several factors. First, since input demand functions are increasing in productivity we may expect OLS estimates to be biased upward. If this is important in our data and proxy methods correct for it then we should see proxy-based $\hat{\beta}_k$ distributions to the left of OLS. The second factor is the covariance among factor inputs. It can be shown that the bias may also be negative if capital and labor are correlated.[17] If this factor is prevalent in our data then we should see proxy-based elasticity distributions to the right of OLS. In addition, OP argue that since plants' profit and value functions are increasing in capital, larger establishments anticipate larger future returns and

[16]Table A1 shows frequency counts from the assignment. Panel 1 summarizes the initial sample. About 66% of establishments show up in the years between 1987-1996, the remaining observations need industry assignment (about 34%). Panel 2 shows statistics about instances where we observed a switch in the industry identifier. Our assignment procedure implies that approximately 29% of the 130 thousand original switching instances disappear. As a cross-check, we compared the average shares our procedure implies to those in the crosswalk. The results of this latter exercise, not disclosed here, suggest that random assignment approximately replicates the NBER-CES mappings.

[17]LP illustrate this in the two-input case. If labor responds to productivity and capital and labor are uncorrelated the bias in $\hat{\beta}_l$ is positive and $\hat{\beta}_k$ remains unbiased. In the more likely case when capital and labor are correlated both $\hat{\beta}_l$ and $\hat{\beta}_k$ may be biased downward. Further, if labor is more correlated to productivity than capital, the bias in $\hat{\beta}_l$ is positive while it is negative for $\hat{\beta}_k$. They point out that the direction of the bias in OLS estimates are more difficult to determine with more variable inputs.

therefore can operate at lower productivity levels, which entails a negative bias in OLS. This last point implies that if OP corrects for selection-induced negative bias then the OP-based $\hat{\beta}_k$ distribution should be to right of OLS.

Figure 1(a) plots the densities of the distribution of capital elasticities in the 50 most populous industries ($\hat{\beta}_k$) and table 2 shows descriptive statistics. There are non-trivial differences in the mean, dispersion and also the general shape. First, LPVA yields significantly larger elasticities, by construction. The elasticity of value added with respect to capital and labor is inherently larger than the elasticity of gross output. We include the LPVA results in our analysis since it is a commonly used method but appropriate caution should be used in comparing the elasticities from gross output and value added production functions.[18] Second, OP tends to result in higher $\hat{\beta}_k$ than LP-variants and OLS. The distributions in figure 1(a) also suggest LPNL and LPGR are more likely to yield lower, and OP tends to yield higher $\hat{\beta}_k$ than OLS. WLPE and WLPM have fatter tails but the majority of their probability mass is located close to OP.

The elasticities of variable inputs show stronger clustering, especially $\hat{\beta}_l$ (figure 1(b)). Part of it is explained by the fact that proxy methods estimate β_l in an OLS step. The main conclusions about $\hat{\beta}_e$ and $\hat{\beta}_m$ are the same, (figure A1): there are numerical differences in these distributions but they generally look more similar.[19]

These results tell us it matters as to what proxy we use to control for the simultaneity of factor inputs and productivity. LP variants tend to yield smaller $\hat{\beta}_k$-s than OLS, which may be interpreted as a sign that the intermediate input proxy controls for possible upward biases in OLS estimates. On the other hand, assuming both investment and intermediate input costs are valid proxies, the fact that OP-based capital elasticities are more likely to be higher than OLS-based ones suggests that controlling for selection-induced bias may also be important. Finally, the difference in WLPE/WLPM-based $\hat{\beta}_k$ distributions relative to other methods suggest the choice of the estimation algorithm - non-parametric two-step estimator versus efficient GMM - also affects the level of $\hat{\beta}_k$. Similar conclusions hold for the second moments of the $\hat{\beta}_k$ distributions. The interquartile range measures in the last column of table 2 suggest that differences in the dispersion of these distributions are also non-trivial. For instance, although the means of GA and LPGR are very close (.08 and .09 in the 50 largest industries) the difference in dispersion is almost twofold (.05 and .09).

In the next exercise we explore whether the choice of estimator affects the ranking of industries by $\hat{\beta}_k/\hat{\beta}_l$.[20] If the estimation method is not important then the ranking should not vary with the estimation method. Our results suggest the opposite. Comparing the rankings across

[18] Using value added as the dependent variable should yield systematically larger elasticities. See section 3 in Petrin and Levinsohn (2012) for more details.

[19] An interesting point is that the location and scale of these distributions hardly change across samples, except the mean of LPNL.

[20] Normalizing is useful because our prior on capital intensity may be stronger than on β_L or β_K.

Table 2: Descriptives of the between-industry distribution of $\hat{\beta}_K$ under TFP estimator variants.

	N	mean	med	iqr
50 largest				
OLS	48	.13	.10	.10
OP	49	.19	.14	.10
LPVA	50	.34	.33	.28
LPNL	50	.11	.04	.18
LPGR	50	.09	.06	.09
LPGSS	50	.18	.13	.15
WLPE	47	.16	.15	.11
WLPM	33	.16	.12	.14
GA	50	.08	.07	.05
10 largest				
OLS	10	.14	.12	.15
OP	10	.21	.15	.27
LPVA	10	.37	.33	.21
LPNL	10	.22	.09	.27
LPGR	10	.11	.06	.09
LPGSS	10	.17	.10	.14
WLPE	10	.17	.14	.07
WLPM	9	.15	.09	.14
GA	10	.09	.10	.06

See notes to table 1 for legends. 50 largest: 4-digit industries which were mapped 1-to-1 between classification systems. Industries were ordered by the within-industry number of plant-year observations.

estimation methods we find there is a positive probability that different estimators will imply different industry rankings (see table A2 for details).[21]

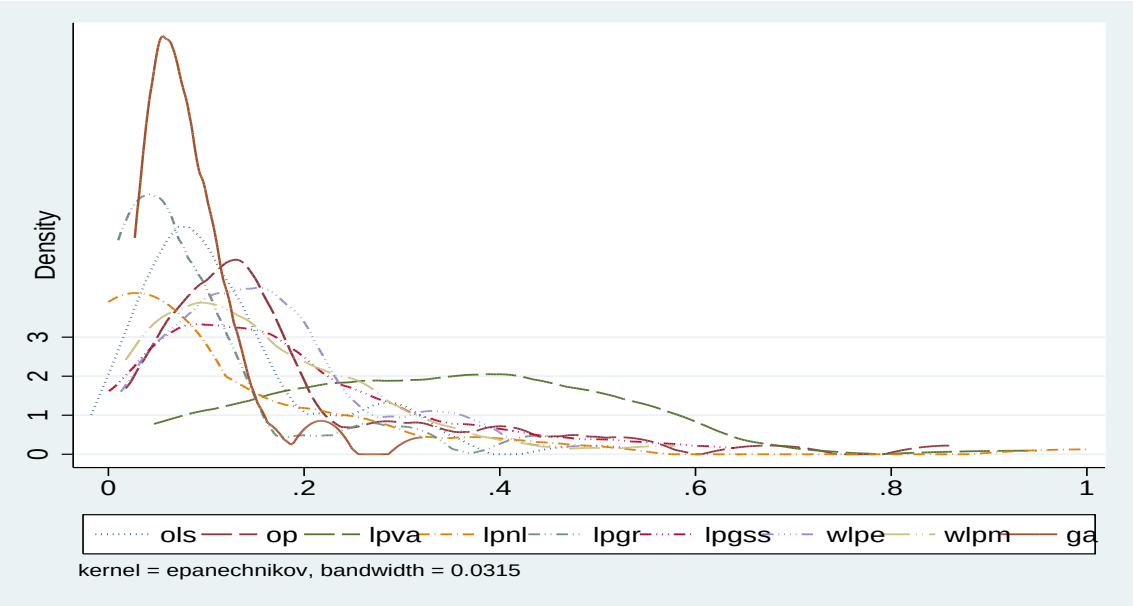
The distributions discussed above are based on elasticities from the 50 most populous 4-digit industries. While 50 observations seem sufficient to estimate elasticity distributions, including all 50 of them also means including estimates from industries with varying sample size. This is important because the elasticities of smaller industries are more likely to be estimated less efficiently.[22]

What is the consequence of such variation? Are some methods more likely to yield nonpositive $\hat{\beta}$-s than others? In order to answer these questions, we re-estimated the above elasticity distributions using all 459 industries - the sample for this exercise is described in more detail at the end of section 3 -, and counted the cases with positive, zero, and negative elasticities for every method.

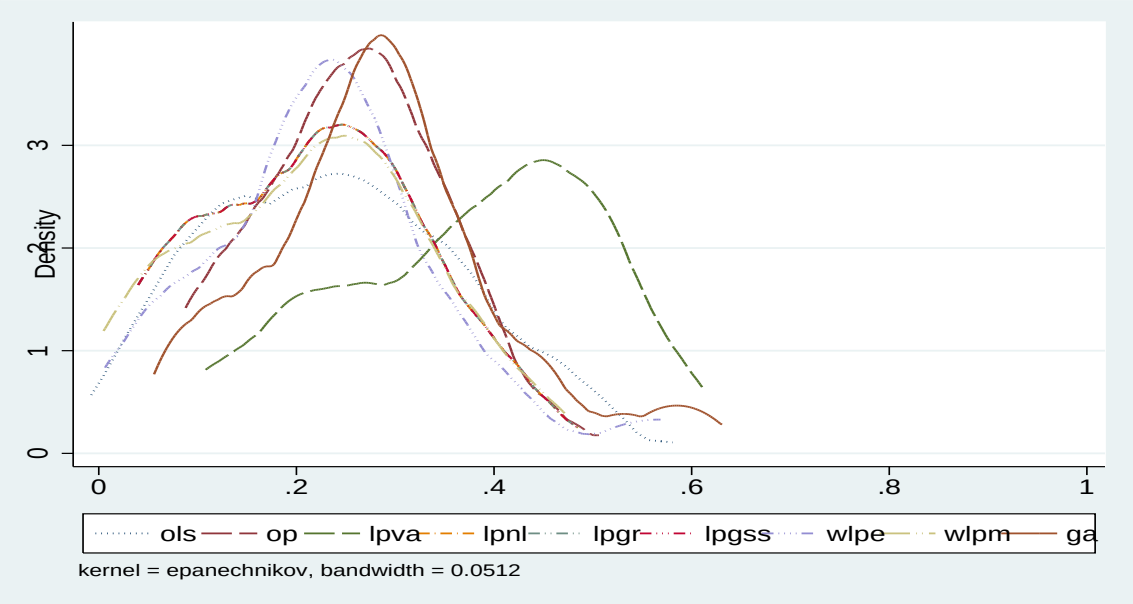
All methods result in positive $\hat{\beta}_l$ and $\hat{\beta}_m$ in most industries (table A4, columns 6 and 9),

[21] We compare rankings by quintiles. If an industry is belongs to the same quintile under different estimators, then we can say results are generally not sensitive to the choice of estimation method. Table A2 describes the changes in distributions as we move from LPGR to OP, GA and WLPE. The first entry in the table says that half of the industries in the lowest quintile under LPGR are also classified in the lowest quintile under OP. If two methods imply similar industry rankings, the (off-)diagonal elements of the corresponding matrix should be close to one (zero). The table shows this is not the case.

[22] Table A3 sorts industries by sample size as measured by number of plant-year observations. The first entry in column 3 says the largest industry has about 39 thousand plant-year observations. The sample size drops by more than 50% in the 10th, and by 80% in the 50th industry.



(a) $\hat{\beta}_K$



(b) $\hat{\beta}_L$

Figure 1: Between-industry distributions of elasticites. TFP estimators in table 1, 50 largest 4-D industries.

while negative β_k and β_e estimates are generally more likely to occur (smaller percentages in columns 3 and 12). There are differences across estimation methods, as well. For example, LPGR always delivers positive $\hat{\beta}_k$, while using other methods yield negative $\hat{\beta}_k$ with a positive probability.[23] OP's algorithm stops in 18% of the cases (columns 2, 5, 8 and 11) due to the lack of information on exiters and LPNL yields zero $\hat{\beta}_k$ in 16% of the industries (column 2).

[23]This is because LPGR searches over a pre-defined grid between .01 and .99. In projection-based methods like OLS, WLPE or WLPM, there is no such 'guarantee' positive $\hat{\beta}_k$.

We obtain negative $\hat{\beta}_E$-s with especially high probability when using WLPE (column 10).

Can the variation in sample size explain these patterns? Comparing the average number of plant-year observations in the problematic group (negative, zero or non-estimable) to that of the positive group suggest the answer is partial 'yes'. Problematic industries are generally smaller, their average size is between 26-70% of the positive group (see table A5 for more details). For example, in the industries where OP stops the average number of observations is less than half of that in the positive group. There are two distinct cases where erratic estimates are unlikely to be related to sample size. First, the zero-$\hat{\beta}_k$ group with LPNL is of very similar size as the positive group (second entry in column 2). It is more likely that LPNL's gradient-based numerical procedure stops at a local optimum point at zero. The second exception is the negative $\hat{\beta}_e$ group with WLPE where average sample size is similar to that in the positive group (column 7). These results are interesting because they suggest outliers may arise more frequently from some methods than others.

In practice, what should be done with negative elasticity estimates that emerge from some methods? One approach would be to exclude such industries since such estimates are implausible but that would obviously raise questions of selection bias. An alternative approach we explore below is to make stronger assumptions of homogeneity so that for example industries within a 3-digit NAICS group share the same factor elasticities. We show below that this approach leads to more plausible elasticity estimates and implied firm dynamics in many cases.

In the last exercise of this section we validate elasticity estimates by examining the plausibility of their implications for returns-to-scale (RTS). RTS is a useful metric because it concisely captures the various effects estimation methods may have on elasticities. We use the revenue-share of profits to define the plausible RTS range and choose the interval $[.7, 1.3]$.[24] This range is consistent with profit shares between $[-25\%, +25\%]$, which does not seem too restrictive considering our data encompasses the time period between 1972 and 2010. In order to allow for uncertainty in the definition of the plausible range, we consider plausible RTS intervals in both higher and lower ranges. Then, we count the number of industries with plausible RTS values for each interval. Our results can be summarized as follows (see table A6 for more details). First, the number of plausible RTS measures decreases as the plausible range narrows for all methods except OLS.[25] This is expected considering the variation in elasticity estimates. Given the widest range, LPGR is very likely to deliver plausible RTS. Second, all estimators are more likely to yield less-than-1 RTS measures, especially LPVA (the entries in columns 5-7 of Panel B are greater than those in Panel A, and more so for LPVA). This is interesting because LPVA-based RTS measures are shown to be relatively low despite the fact that the elasticity distributions under LPVA are located to the right of the others. As a final point, we note that in only less than 30% of the industries do all methods deliver plausible RTS values

[24] It can be shown that RTS can be written as a function of the revenue share of profits.

[25] For GA, $RTS = 1$ by construction, therefore it is not shown in these calculations.

(last rows in the panels). This finding highlights the quantitative differences in the elasticities mentioned earlier.

4.2 TFP dispersion

Does the choice of the estimator affect the second moment of the within-industry TFP distribution? The answer is yes and no. The interquartile range (IQR) and standard deviation, averaged over industries and time, suggest that differences in dispersion are nontrivial (table 3, more details in table A7). The average IQR of the log-TFP distribution varies between .24 and .4 across methods (GA and WLPM, not counting LPVA). This says that the plant at the 75th percentile generates approximately 27-49% more revenue with the same amount of inputs than the plant at the 25th percentile. The differences suggest one must exercise caution when using results for economic analysis where small changes in magnitudes may make a big difference. The differences notwithstanding, these results also confirm that productivity dispersion is invariably large no matter how we measure it.

We also find however that the productivity rank of plants within industries is impacted by estimation methods as shown in table 4. Both the Pearson and Spearman rank correlations are substantially smaller than 1 implying that where plants sit in the productivity distribution is sensitive to method. The WLPM method yields especially low correlations with other methods.

TFP dispersion varies with the estimation method. It varies with industries, too.[26] Taking GA-based TFP as an example, productivity dispersion in the industry one standard deviation above (below) the sample mean is .35 (.13). These results suggests the effect of industry differences are at least as important as estimation methods.

So what determines within-industry dispersion from a measurement perspective? Does it depend positively on the estimated level of elasticities? Are elasticities more important than input-output characteristics?

Equation (1) implies it depends on elasticities, output and factor variances and the covariances among them.[27] It can be shown that the theoretical relationship between TFP dispersion and β_j depends on the covariance between output and the j-th input, the variance of the j-th input, the covariances of the jth input with the other inputs, and the elasticities.[28] However, analytical expressions become difficult to evaluate if one intends to assess the relative

[26]Results not shown in this paper indicate the cross-method correlation between within-industry dispersion and its dispersion across industries is positive. This means that if an estimation method implies greater dispersion within industries then between-industry differences in dispersion are also greater.

[27]Applying the definition of the variance to log TFP we can write $\sigma^2_\omega = var[q - \sum_j \beta_j x_j] = \sigma^2_q + \sum_j \beta^2_j \sigma^2_{x_j}$ - $2\sum_j \beta_j \sigma_{qx_j} + 2\sum_j \sum_{i \neq j} \beta_j \beta_i \sigma_{x_j x_i}$.

[28]The partial derivatives of σ^2_ω wrt. σ_q and β_j are given by $\frac{\partial \sigma^2_\omega}{\partial \sigma_q} = 2\sigma_q$ and $\frac{\partial \sigma^2_\omega}{\partial \beta_j} = 2\beta_j \sigma^2_{x_j}$ - $2\sigma_{qx_j}$ + $2\sum_{i \neq j} \beta_i \sigma_{x_j x_i}$. $\frac{\partial \sigma^2_\omega}{\partial \sigma_q} = 2\sigma_q$. The former expression implies we may expect TFP dispersion to be positively associated with output dispersion. The condition for the sign of $\frac{\partial \sigma^2_\omega}{\partial \beta_j}$ is given by: $\frac{\partial \sigma^2_\omega}{\partial \beta_j} \lessgtr 0$ iff. $\beta_j \lessgtr \frac{\sigma_{qx_j} - \sum_{i \neq j} \beta_i \sigma_{x_j x_i}}{\sigma^2_{x_j}}$, which shows the relationship depends on σ_{qx_j}, $\sigma^2_{x_j}$, the $\beta_i \sigma_{x_j x_i}$ terms and β_j itself.

Table 3: Descriptive statistics of tfp-distributions. Largest industries which were mapped 1-to-1 between different industry classification systems. Industries were ordered by the within-industry number of plant-year observations. Trim: IQR.

	N (1000)	iqr	sd
50 largest industries			
OLS	455	0.26	0.26
OP	380	0.32	0.38
LPVA	458	0.68	0.57
LPNL	457	0.30	0.38
LPGR	457	0.29	0.31
LPGSS	455	0.27	0.27
WLPE	457	0.34	0.36
WLPM	457	0.40	1.88
GA	433	0.24	0.22
10 largest industries			
OLS	185	0.22	0.20
OP	152	0.31	0.39
LPVA	191	0.65	0.52
LPNL	188	0.30	0.37
LPGR	188	0.29	0.28
LPGSS	188	0.27	0.25
WLPE	187	0.31	0.28
WLPM	188	0.33	0.35
GA	177	0.23	0.21

See notes to table 1 for legends. Statistics are calculated using the deviation of plant-level log-TFP from industry- and time-specific means. All results shown were calculated using non-outlier observations only (pre-, post-estimation). A version including pre-estimation outlier observations can be found in the appendix, results barely change.

importance of elasticities and input-output characteristics for two reasons. First, β_j and the previously mentioned covariances are endogenous because those covariances are used to estimate the β_j. Second, the expressions are not informative about the effect of empirical issues such as measurement error.

We carried out counterfactual exercises in order to assess the relative importance of these factors (see table A11 for more details). Our results suggest that measured productivity dispersion is increasing in both $\hat{\beta}_k$ and output dispersion (σ_q). Specifically, TFP dispersion approximately doubles after a one-standard-deviation increase in $\hat{\beta}_k$. A similar change can also be achieved by increasing σ_q. However, the required increase in σ_q is relatively large[29], which indicates that typical changes in $\hat{\beta}_k$ may have larger effects on TFP dispersion than typical changes in σ_q.

4.3 Growth and survival

In this section we explore whether one of the most important predictions from standard models of firm dynamics is robust to the way TFP is estimated. The most influential theories of firm dynamics are described in the classic models by Jovanovic (1982), Hopenhayn (1992), Ericson

[29]The necessary increase in σ_q is unusually large in the sense that it amounts to approximately 3.5-standard-deviations.

Table 4: Correlations among within-industry tfp-distributions, 50 largest 4-D industries.

	OLS	OP	LPVA	LPNL	LPGR	LPGSS	WLPE	WLPM	GA
Pearson									
OLS	1								
OP	0.51	1							
LPVA	0.69	0.43	1						
LPNL	0.64	0.33	0.68	1					
LPGR	0.82	0.43	0.8	0.84	1				
LPGSS	0.81	0.54	0.76	0.73	0.88	1			
WLPE	0.51	0.46	0.49	0.49	0.52	0.63	1		
WLPM	0.02	0.02	-0.16	-0.12	-0.15	-0.05	0.15	1	
GA	0.79	0.46	0.57	0.56	0.68	0.71	0.51	0.09	1
Spearman									
OLS	1								
OP	0.68	1							
LPVA	0.79	0.62	1						
LPNL	0.78	0.56	0.73	1					
LPGR	0.87	0.61	0.81	0.88	1				
LPGSS	0.82	0.69	0.76	0.8	0.87	1			
WLPE	0.59	0.61	0.55	0.52	0.56	0.69	1		
WLPM	0.35	0.36	0.19	0.23	0.26	0.38	0.49	1	
GA	0.81	0.63	0.61	0.66	0.7	0.73	0.6	0.43	1

Correlations reflect distributional differences discussed above: proxy methods show greater similarity, while WLPM seems more different. Rank correlations confirm. Including pre-estimation outliers barely changes correlations (see table A8).

and Pakes (1995). These models all share an important common ingredient: firms decide on exit or growth upon learning their ex ante uncertain productivity level. That is, firm dynamics are determined endogenously and firms' decisions are based on a firm-specific productivity shock. A common prediction of these models is that more productive firms are more likely to grow and survive than their less productive competitors.[30] Empirical work on the connection between growth and productivity also relies on the results of these models. A study by Foster et al. (2013) (FGH hereafter) is a recent example of the research in this context.[31] A comprehensive

[30]These models also have implications for the behavior of aggregate variables. At first sight, they seem different in the way they look at the relationship between microeconomic and aggregate behavior. Hopenhayn (1992) studies competitive equilibria with an infinite number of small firms, where equilibrium conditions are a result of asymptotic approximations and assuming a constant industry state. In contrast, Ericson and Pakes (1995) analyze a small (finite) number of firms because calculating equilibrium becomes computationally challenging when this number is infinite. As shown in Weintraub et al. (2011), these seemingly different frameworks are asymptotically equivalent if the firm size distribution satisfies a certain light-tail condition. If this condition is satisfied, i.e. firms in the right tail of the distribution are not large enough to exert influence on the aggregate industry state, then firms only have to take the aggregate industry state into account and we are free to assume a constant industry state. This is an interesting result because although other literatures offer important conclusions about why distributional characteristics matter for aggregate dynamics (see, for example, Caballero et al. (1997) on nonlinear labor adjustment, or Gabaix (2011) on the granular origins of aggregate fluctuations), these studies offer consistent analytical frameworks to understand the role of distributional characteristics using canonical models of firm dynamics.

[31]They investigate the relationship between productivity and reallocation in their paper by associating job creation and destruction with reallocation. We repeatedly use the term growth and survival instead of reallocation but essentially mean the same thing.

survey of the literature from the past decade can be found in Syverson (2011).

We now turn to our own empirical analysis of the relationship between productivity, growth and survival. Our discussion is centered on, similarly to sections 4.1 and 4.2, whether the most commonly found patterns hold up across TFP estimators. We build upon the existing literature concerning the properties of productivity dynamics and look at simple regression models linking a set of outcomes to productivity and plant-level controls. As mentioned above, canonical models of firm dynamics describe growth and survival as functions of idiosyncratic productivity shocks. The main prediction from these models is that plants with positive shocks expand while plants with negative shocks shrink and/or exit. We test the robustness of these predictions using various productivity measures.

We follow FGH when considering the relationship between productivity and growth of all establishments, exiters and incumbents separately. This approach is justified by theoretical and empirical considerations. First, the basic models of firm dynamics themselves analyze these margins separately. Second, earlier empirical research found that there are differences in the productivity levels of continuers, entrants and exiters.[32]

Equation (2) describes our empirical specification, analogous to that in FGH:

$$Y_{i,t+1} = \gamma_1 \theta_{state} + \gamma_2 \theta_{year} + \gamma_3 \theta_{size} + \gamma_4 \omega_{it} + \gamma_5 u_{s,t+1} + \epsilon_{i,t+1}, \tag{2}$$

where $Y_{i,t+1}$ is the outcome of interest such as growth between t and $t+1$, ω and u denote a plant-level measure of TFP and state-level measure of change in unemployment from t to $t+1$. The θ-s denote state-, year- and establishment size-class effects. i, t and s index plants, time periods and states, respectively. This specification differs from the one in FGH in that age effects, the Great Recession dummy and its interactions are omitted for simplicity. It is important to note that this specification relates growth and survival outcomes from t to $t+1$ based on productivity in period t.

Table 5 shows $\hat{\gamma}_4$ from equation (2) based on a sample of the 50 largest industries. Each row lists the effect of productivity on the outcome shown in the first column. The three outcomes are: employment growth among all establishments, exit, and employment growth among continuers. The columns correspond to the TFP estimator variants. For example, the first entry in column 1 says a plant is estimated to grow .16% faster if it is 1% more productive when we measure TFP as an OLS residual. All other entries are to read analogously. Point estimates suggest there are non-trivial differences in the measured effect of TFP. For example, the first entry in column 2 says using OP-based TFP among all establishments implies the growth-effect of productivity is less than half of that of OLS-based TFP. The difference between the results of these two estimators encompasses the variation in coefficients, not counting WLPE and WLPM, which we will come back to shortly. Despite the non-negligible differences, our

[32]See, for example, Baily et al. (1992), Foster et al. (2001), Foster et al. (2006).

estimates support the earlier finding that more productive plants grow significantly faster than their less productive competitors. The estimates in row 3 show that exit is significantly more likely for low-productivity establishments compared to high-productivity establishments

Table 5: The effect of TFP on outcomes, 50 largest industries. Outcomes are: employment growth among all establishments (row 1), exit (row 3), employment growth among continuers (row 5).

	ols	op	lpva	lpnl	lpgr	lpgss	wlpe	wlpm	ga
4-digit elasticities									
$dlnE_{it}$	0.163***	0.072***	0.09***	0.089***	0.139***	0.129***	0.06***	-0.006***	0.190***
	(0.015)	(0.012)	(0.006)	(0.008)	(0.01)	(0.011)	(0.007)	(0.002)	(0.014)
exit	-0.05***	-0.018***	-0.028***	-0.035***	-0.046***	-0.047***	-0.02***	0.002**	-0.064***
	(0.006)	(0.005)	(0.003)	(0.004)	(0.004)	(0.005)	(0.004)	(0.001)	(0.006)
$dlnE_{it}^{cont}$	0.068***	0.039***	0.037***	0.022***	0.053***	0.04***	0.022***	-0.003***	0.067***
	(0.007)	(0.004)	(0.005)	(0.005)	(0.006)	(0.006)	(0.006)	(0.001)	(0.006)
3-digit elasticities									
overall growth	0.19***	0.155***	0.101***	0.126***	0.152***	0.154***	0.1***	-0.017***	0.183***
	(0.015)	(0.015)	(0.007)	(0.015)	(0.012)	(0.016)	(0.01)	(0.006)	(0.014)
exit	-0.066***	-0.048***	-0.03***	-0.036***	-0.044***	-0.054***	-0.03***	0.005**	-0.062***
	(0.007)	(0.006)	(0.003)	(0.005)	(0.006)	(0.006)	(0.004)	(0.002)	(0.006)
conditional growth	0.061***	0.065***	0.045***	0.06***	0.071***	0.05***	0.043***	-0.008***	0.064***
	(0.006)	(0.007)	(0.005)	(0.007)	(0.008)	(0.007)	(0.007)	(0.003)	(0.006)
3-digit elasticities, industries with negative or non-estimable elasticities dropped									
overall growth	0.197***	0.153***	0.101***	0.123***	0.164***	0.152***	0.104***	0.076***	0.183***
	(0.017)	(0.015)	(0.007)	(0.015)	(0.013)	(0.016)	(0.013)	(0.018)	(0.014)
exit	-0.069***	-0.044***	-0.03***	-0.033***	-0.048***	-0.05***	-0.036***	-0.028***	-0.062***
	(0.007)	(0.005)	(0.003)	(0.006)	(0.005)	(0.006)	(0.005)	(0.006)	(0.006)
conditional growth	0.064***	0.072***	0.045***	0.063***	0.074***	0.056***	0.035***	0.023**	0.064***
	(0.008)	(0.009)	(0.005)	(0.011)	(0.009)	(0.009)	(0.009)	(0.011)	(0.006)

Estimates are taken from regressions (equation (2)) of three outcomes (employment growth among all establishments, exit, and employment growth among continuers) on a plant-level measure of TFP (columns), a state-level measure of unemployment growth, year-, sizeclass- and state-fixed effects. Standard errors (in parentheses) are clustered at the state level. All regressions are based on trimmed TFP distributions. Sample size information can be found in table A12. Results for two additional industry sets can be found in tables A13-A14.

The results in the middle and lower panels highlight the sensitivity of some methods to changes in sample definition. The middle panel of the table shows results when elasticities are estimated pooling data from all 4-digit industries within the same 3-digit industry. In theoretical terms, pooling amounts to maintaining stricter assumptions about the homogeneity of elasticities. But, as we will demonstrate, it may be necessary in order to use more data for estimation. Comparing estimates between the top and middle panels suggests that some methods are prone to such changes. For instance, the absolute values of OP-based coefficients of all three regressions increase. This suggests pooling data is beneficial because if there is no statistical association between two variables, their partial correlation coefficient would tend towards zero. The other estimates change somewhat but most of them remain in a comparable range.

Unfortunately, increasing sample size by pooling data does not seem to help with WLPM. The relationship between WLPM-based productivity and growth remains reversed. However, if we drop industries with negative elasticities, the signs appear to be more in line with other results. In addition, the absolute value of the point estimates also increase and/or standard

errors decrease, despite the decrease in sample size.[33]

These findings are intuitive. The OP example highlights a case where there is better chance of estimating plausible elasticities by pooling data from neighboring industries. In such cases, sacrificing homogeneity and pooling data amounts to replacing uninformative TFP observations with informative ones. WLPM is an example where only pooling and also dropping 3-digit implausibles help exclude observations which are less informative. We carried out similar exercises using other samples from our default industry set. Without discussing them in detail, we only note here that a combination of using 3-digit elasticities and dropping industries with implausible elasticities yields expected growth and exit coefficients under all estimation methods in all the samples we considered.[34]

To sum up, our estimates show positive (negative) and significant association between TFP and growth (exit), irrespective of how TFP is estimated. Some methods yield implausible (i.e., negative) factor elasticities that weaken the estimated relationship between productivity, growth and survival. If we use larger sample sizes from broader industry definitions to overcome implausible elasticity estimates we find results for the relationship between productivity, growth and survival that line up better with other methods.

4.4 Structural decompositions of aggregate productivity growth

Productivity decompositions are identities that parse aggregate productivity growth (APG) into components that are assumed to capture different sources of growth. Numerous papers explored how APG is tied to the behavior of microeconomic agents, some examples are Olley and Pakes (1996), Baily et al. (2001), Petrin et al. (2011) (PWR hereafter), Petrin and Levinsohn (2012). These decompositions are different in many respects but there is a common thread: they are all based on the general idea that APG sources either from the productivity gains within, or the more efficient allocation of factor inputs between, establishments.

We implement a recent decomposition proposed by PWR because the role of factor elasticities is made explicit in their model. They show that the contribution to APG by the reallocation of each factor input is a weighted average of input growth rates where the weights are determined by the difference between the marginal revenue product of inputs and their marginal cost. The elasticities affect this difference, and therefore reallocation, via the marginal products of inputs. Another nice feature of PWR's approach is that the measure of APG itself does not depend on the way TFP is estimated. This is because APG, the left-hand-side variable of the equation, is defined as the growth in final demand in excess of capital and labor growth.[35]

[33]More information about the sample size can be found in table A12.

[34]See tables A13-A14 for more details. We also note there is some sensitivity to outliers (pre- and post-estimation), but this variation does not invalidate the main point. The latter results are not included in the paper for brevity.

[35]Since final demand is not observed, they measure its growth using value added. This approximation is exact at economy-level only. It is therefore important to realize that implementing the decomposition for a subset

Our main results are shown in table 6. The first four lines contain the elements of PWR's APG definition, the rest of the table summarizes the details of decomposing this growth into contributions by reallocation, within-plant productivity growth and fixed costs. All the entries in the table are calculated as averages of the yearly contributions between 1977-1996.

Since the definition of APG does not depend on elasticities, lines 1-4 are identical across our estimation methods (columns 1-10). Our annual average growth rate for labor (-.3%) is similar to that in PWR (-.2%, last column[36]) but our value added and capital growth rates are shown to be somewhat smaller implying similar APG measures. These differences are due partly to measurement differences and partly to an issue with PWR's source data for implicit deflators.[37] Value added growth looks more similar if we ignore the first years of ASM panels (shown in the lower panel of table 6), when calculating aggregate growth is problematic. In these years, frequency-weights cannot correct for the selection problem in the ASM because we do not observe growth rates for smaller establishments, which were not included in the previous survey. As a consequence, average growth rates will reflect the growth rate of the largest plants only.[38]

The rest of the table describes the decomposition of APG into reallocation and within-plant growth. The contribution of fixed costs is obtained as a residual. Since these components depend on the estimation method, we expect to see differences across the columns of the table. Indeed, the implications of the quantitative differences in the results are substantial. They suggest that the choice of estimation method matters. Considering the entire sample (upper panel of table 6), our estimate of the annual contribution of reallocation is between .7-1.5 percentage points, which amounts to 30-70% of the average annual aggregate productivity growth. The results for within-plant productivity growth are 0-1.4 percentage points and 0-64%. Note these ranges do not account for WLPM (column 7), which will be discussed later in more detail. We further illustrate method-related variation using results which ignore the first years of ASM panels. These are shown in the lower panel of table 6. For example, OLS, OP, LP variants and GA result in .5-.9 percentage point average annual contributions for reallocation

of plants we compute the contribution of that subset to APG, not the APG of the subset. To calculate the subset's exact APG, we would have to observe final demand for that subset.

[36]The last column repeats the results from tables 1,2 and 3a in PWR for convenience.

[37]The ASM includes information on hours worked for production workers only. PWR generate nonproduction worker hours data using variation in the average number of nonproduction workers and assuming a 40-hour working week and 50 weeks: $H_{np} = (E_{total} - E_{prod})\frac{50*40}{1000} = 2(E_{total} - E_{prod}) = 2E_{np}$. Our approach estimates total hours as a function of the number of production worker hours and the ratio of the total wage bill to the wages of production workers: $H_{total} = H_{prod}\frac{SW}{WW}$. The difference in capital growth is explained by a labeling mistake in PWR's source data for implicit deflators. As a last point, we note there is a small difference in the way we calculate value added. While we use energy deflators to calculate constant-dollar energy costs, PWR deflate energy costs together with other intermediate inputs using material deflators, which may affect aggregate growth rates.

[38]It is impossible to calculate growth rates for smaller establishments which were just rotated into the panel. The ASM sample weights are therefore not useful because there is no information we can use to correct for size-related non-randomness. The calculation of level-based aggregates is free of such problems.

Table 6: Time-averages (1977-1996) of APG contributions (in percentage points). No pre- or post-estimation outlier filters applied. ASM (continuers), weighted.

	ols	op	lpnl	lpgr	lpgss	wlpe	wlpm	wlpe*	wlpm*	ga	PWR
	1	2	3	4	5	6	7	8	9	10	11
VA growth	1.9	1.9	1.9	1.9	1.9	1.9	1.9	1.9	1.9	1.9	2.3
Capital growth	0	0	0	0	0	0	0	0	0	0	0.3
Labor growth	-0.3	-0.3	-0.3	-0.3	-0.3	-0.3	-0.3	-0.3	-0.3	-0.3	-0.2
APG= dVA-dK-dL	2.2	2.2	2.2	2.2	2.2	2.2	2.2	2.2	2.2	2.2	2.2
Total reallocation	1	0.8	0.7	0.9	1.1	1	4.9	1.4	1.5	0.7	2.1
Capital reallocation	0.2	0.2	0.1	0.3	0.5	0.8	-3.6	0.8	0.6	0.1	0.4
Labor reallocation	0	0	0	0	0	0	0	0	0	0.1	0.4
Materials reallocation	0.6	0.4	0.5	0.5	0.5	0.6	8.3	0.6	0.8	0.5	0.7
Energy reallocation	0.2	0.2	0.1	0.1	0.1	-0.3	0.1	0	0.1	0	0.2
Technical efficiency ter	0.9	1	1.3	1.4	0.8	0.4	-3	0.2	0	1.2	0.2
Fixed cost residual term	-0.2	-0.3	-0.2	0.1	-0.3	-0.8	-0.3	-0.5	-0.6	-0.3	0.1
σ_{RE}	0.9	0.6	0.8	0.8	0.9	3.8	13	1	1.4	1.4	1.7
σ_{Within}	2.2	2.7	2.7	2.9	2.6	6.4	8.6	2.4	2.3	2.4	2.7
Time-averages ignoring first years of ASM panels: 1979, 1984, 1989, 1994***											
VA growth	1.4	1.4	1.4	1.4	1.4	1.4	1.4	1.4	1.4	1.4	1.6
Capital growth	0	0	0	0	0	0	0	0	0	0	0.4
Labor growth	-0.4	-0.5	-0.5	-0.5	-0.4	-0.5	-0.4	-0.5	-0.5	-0.5	-0.4
APG= dVA-dK-dL	1.9	1.9	1.9	1.9	1.9	1.9	1.9	1.9	1.9	1.9	1.6
Total reallocation	0.8	0.7	0.6	0.7	0.9	1.1	3.9	1.3	1.3	0.5	2
Capital reallocation	0.2	0.2	0.1	0.3	0.5	0.7	-3.4	0.8	0.6	0.1	0.9
Labor reallocation	0	0	-0.1	-0.1	0	-0.1	0.1	0	0	0.1	0.4
Materials reallocation	0.4	0.4	0.4	0.4	0.3	0.4	7	0.5	0.5	0.3	0.4
Energy reallocation	0.2	0.1	0.1	0.1	0.1	0	0.1	0	0.1	0	0.3
Technical efficiency ter	0.8	0.7	1.3	1.4	0.7	-0.7	-0.6	-0.1	-0.1	1.1	-0.1
Fixed cost residual term	-0.3	-0.5	0	0.2	-0.4	-1.5	1.4	-0.8	-0.7	-0.3	0.2
σ_{RE}	0.8	0.5	0.7	0.7	0.7	3.8	13.3	0.9	0.8	1.2	1.7
σ_{Within}	2.3	2.8	2.7	3.1	2.8	6.6	6.1	2.5	2.5	2.4	2.7
N** (1000)	4.1	3.5	2.3	2.3	2.3	1.9	1.9	1.5	1.5	4.1	n.a.

*Instrument set as in PWR: second and third lags. **Average number of plant-year observations per industry (in thousands).*** First years of ASM panels are problematic: weighted growth rates are based on only large plants since growth rates do not exist for establishments just rotated in.

(columns 1-5 and 10) implying that reallocation accounted for less than half of the aggregate productivity growth between 1977-1996. These contributions are 1.3 percentage point in the case of both WLPE and WLPM (columns 8-9) indicating that reallocation was more important, it explained almost 70% of average annual productivity growth.

WLPM-based results highlight a more technical point related to the sensitivity of GMM-based estimators to the choice of the instrument matrix. The growth rates in column 7 are based on elasticities estimated using instrument sets containing only the most recent values of hours (t-1), energy (t-1) and materials (t-2).[39] Using these instruments implies the reallocation and within-plant contributions appear one order of magnitude greater compared to other methods.

[39]This instrument set is smaller than the one PWR used in their paper. A smaller instrument set may be justified by the fact that using earlier lags leaves less variation for estimation due to the loss of observations at the beginning of plants' timeseries.

Increasing the size of the instrument matrix by including earlier lags (1-3), as in in PWR makes the extreme contributions disappear (column 9). This finding indicates that, at least in the case of WLPM, the benefits of using more lags outweigh the costs due to the loss of observations at the beginning of plants' timeseries. Repeating this exercise using energy costs to control for TFP (columns 6 and 8) suggests the energy proxy is not as sensitive to the choice of the instrument matrix.

With these qualifications in mind, we conclude that all estimators imply positive average contributions for reallocation indicating that reallocation was productivity-enhancing in U.S. Manufacturing industries between 1977-1996. This is in line with both existing theories of firm dynamics and our intuition about well functioning market economies because they tell us that the movement of resources across plants should in general contribute positively to aggregate productivity growth.[40] As a final point, we note the standard deviation of reallocation is small relative to within plant productivity reinforcing previous findings (last rows of the panels in table 6). This indicates within-plant growth is more important for aggregate fluctuations than reallocation.

To sum up, we found sensitivity in terms of the relative importance of reallocation and within-plant productivity growth. However, our overall results suggest that the previously observed main patterns in the decomposition literature seem to be robust to the way TFP is estimated. Reallocation is shown to be productivity enhancing and within-plant growth proves more important for aggregate fluctuations. We also found that choosing materials as a proxy entails additional sensitivity to the choice of the instrument matrix while choosing energy costs does not have such implications.

4.5 Imputation, dispersion and growth

U.S. Manufacturing data collected by the Census Bureau is subject to item nonresponse. This means respondents answer some questions but not others. Such missing values are imputed for by the Bureau. Recent research found that certain imputation methods can impact analyses that use such completed data because imputation is non-random (see White et al. (2012), WRP hereafter). Some of the methods use industry level data or fitted values from regression models implying that the variation in completed data can be assumed to be smaller relative to what observed responses would imply. One consequence of this is that within industry dispersion statistics based on completed data may be biased downward (WRP give an example). The direction of bias is less clear for the relationship between productivity, growth and survival.

One may address imputation in a variety of ways. One way would be to just use non-imputed observations. The problem with that approach is that imputed observations may not be random so selection biases would result. To correct for the latter, one can use selection

[40] Table A15 lists the timeseries underlying columns 8-9 of table 6. The contribution of reallocation is positive in the majority of years which reinforce what our aggregate results suggest.

bias correction methods such as weighting non-imputed observations by inverse probability weights. We pursue that approach as we describe below. Another way would be to use multiple imputation methods to improve on the methods that have been used to impute the plant level data. For example, the classification and regression tree method (CART) as used by WRP. Calculating our results in this manner is beyond the scope of this paper but we do compare the patterns of our findings to those in WRP.

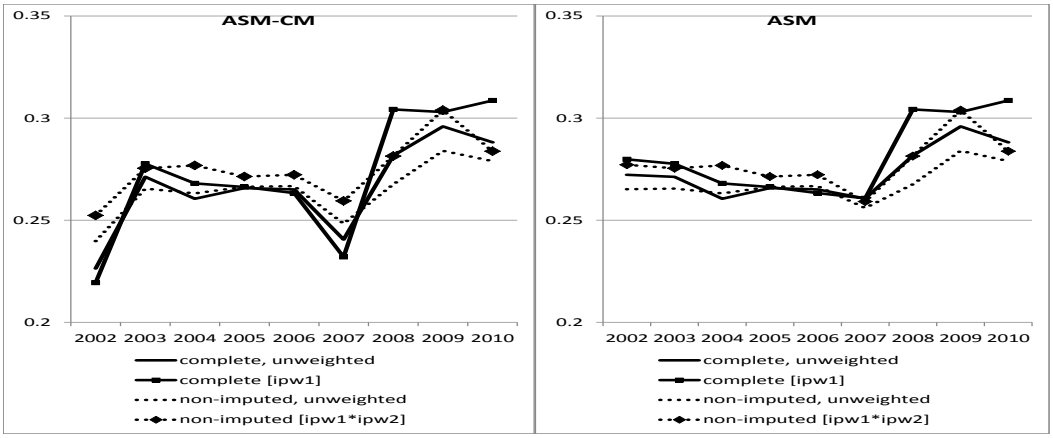
Our approach in this section is different from earlier sections in that what we compare here is not various TFP measures in the same sample but the same TFP measure across complete and non-imputed samples. In order to account for nonrandom imputation we try to measure the relationship between the probability of imputation and plant characteristics. If we can successfully quantify this relationship the probability weights from these regressions can be used to control for non-randomness when calculating dispersion statistics and regression coefficients. The weighting scheme is simple: weights are inversely proportional to the probability that the plant's TFP is calculated using non-imputed data. More details about the probability models can be found in appendix 6.2.

We carry out two exercises. In the first one, we compare weighted and unweighted statistics of within-industry dispersion which are based either on completed data or only non-imputed observations. In the second exercise, the objective is to assess whether accounting for imputation affects our results on the relationship between productivity, growth and survival.

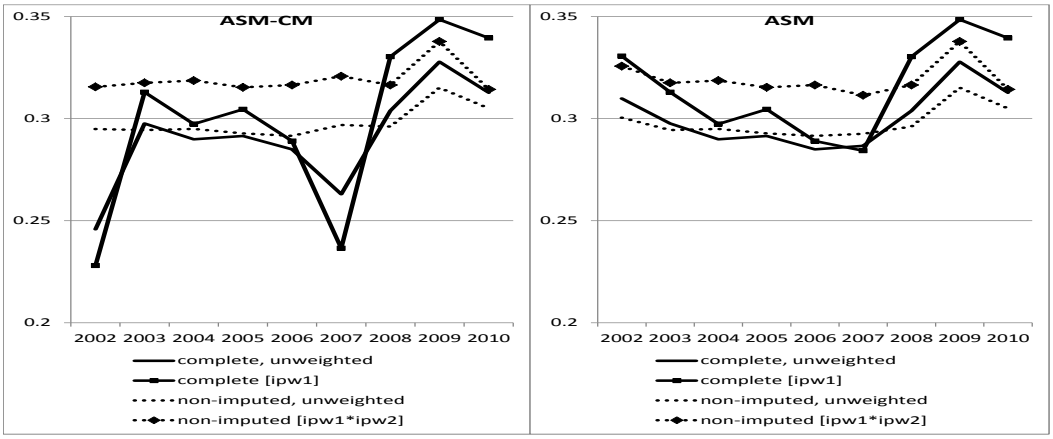
We consider three dispersion measures (standard deviation (sd), interquartile range (iqr), 90th-10th percentile range (9010)) in four different samples. Table 7 summarizes our results. The statistics in columns 1-3 in the upper section of Panel 1 are unweighted and based on the complete sample including imputed and non-imputed observations. Columns 4-6 show weighted statistics based on the complete sample where the weights $ipw1$ are inversely proportional to the probability that a plant is selected into the ASM-CM.[41] Note that $ipw1$ is not designed to control for non-randomness in imputation but to replicate the size distribution of LBD. The lower section of Panel 1 presents the statistics which are calculated using only non-imputed observations. Columns 1-3 in the lower section show unweighted statistics and columns 4-6 show statistics in the same non-imputed sample but now observations are weighted by a composite of $pw1$ and a weight that is inversely proportional to the probability that the plant's TFP is calculated using non-imputed data ($pw2$). The statistics in Panel 2 are based on ASM cases only. Figure 2 visualizes these timeseries. Solid and dotted lines graph statistics which are based on complete and non-imputed data, respectively. Weighted statistics are denoted by marks on the corresponding graphs.

Our conclusions are the following. The standard deviation does not seem to be very sensitive to imputation, while the quantile-based dispersion measures are affected. The interquantile-

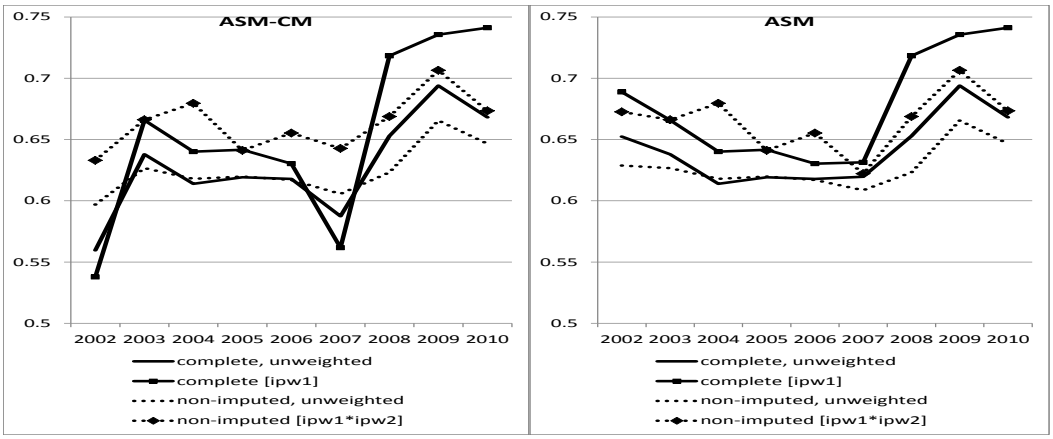
[41]Selection into the ASM and CM is also non-random. More details about this selection issue and the technique we use to address it can be found in FGH.



(a) Dispersion measure: standard deviation



(b) Dispersion measure: inter-quartile range



(c) Dispersion measure: 90th to 10th percentile range

Figure 2: Timeseries of average within-industry dispersion measures (table 7) in various samples, ASM/CM 2002-2010, all industries.

Table 7: The effect of imputation on within-industry dispersion in the ASM/CM between 2002-2010, by dispersion measures and samples. See figure 2 for visualizations.

year	sd [1]	iqr [2]	9010 [3]	sd [4]	iqr [5]	9010 [6]
Panel 1. ASM-CM						
	Complete data, unweighted*			Complete data, weighted [ipw1]**		
2002	0.2264	0.2458	0.5596	0.2194	0.228	0.5378
2003	0.2712	0.2976	0.638	0.2776	0.3129	0.6658
2004	0.2605	0.2899	0.6139	0.268	0.2974	0.6401
2005	0.2657	0.2915	0.6192	0.2662	0.3045	0.6416
2006	0.2651	0.2849	0.6179	0.2633	0.289	0.6303
2007	0.2405	0.263	0.5874	0.232	0.2365	0.5617
2008	0.2816	0.3039	0.6529	0.3043	0.3303	0.7183
2009	0.2959	0.3278	0.694	0.303	0.3485	0.7356
2010	0.2881	0.3127	0.6682	0.3086	0.3395	0.7413
	Non-imputed data, unweighted			Non-imputed data, weighted [ipw1×ipw2]***		
2002	0.2398	0.2949	0.597	0.2523	0.3155	0.633
2003	0.2655	0.2944	0.6267	0.2754	0.3175	0.6662
2004	0.2631	0.295	0.6179	0.2768	0.3186	0.6796
2005	0.2663	0.2928	0.6198	0.2714	0.3153	0.6412
2006	0.2667	0.2916	0.617	0.2723	0.3164	0.6553
2007	0.2483	0.2969	0.6056	0.2594	0.3207	0.6428
2008	0.2676	0.2961	0.6232	0.2814	0.3164	0.6687
2009	0.284	0.3151	0.6655	0.3039	0.3378	0.7065
2010	0.279	0.305	0.6465	0.2838	0.3143	0.6734
Panel 2. ASM cases						
	Complete data, unweighted			Complete data, weighted [ipw1]		
2002	0.2722	0.3098	0.6524	0.2798	0.3306	0.6888
2007	0.2607	0.2867	0.6196	0.2607	0.2843	0.6314
	Non-imputed data, unweighted			Non-imputed data, weighted [ipw1×ipw2]		
	sd	iqr	9010	sd	iqr	9010
2002	0.2652	0.3004	0.6288	0.2771	0.3257	0.6726
2007	0.256	0.2926	0.6086	0.2591	0.3114	0.6221

*Unweighted data: the entire sample including both imputed and non-imputed observations, first and second moments are unweighted.
**Weighted data [ipw1]: propensity score weighted observations, where the weight $ipw1$ is inversely proportional to the probability that the plant is selected into the ASM/CM.
***Non-imputed, weighted data [pw1,pw2]: observations in the non-imputed sample, weighted by a composite propensity score $ipw1*ipw2$ where $ipw2$ denotes a weight inversely proportional to the probability that the plant's TFP is calculated using non-imputed data. We calculate the probability using a logistic regression of TFP being non-imputed on industry fixed effects, employment size and age classes and payroll deciles.

range and the 90-10 range in non-imputed data appear generally smaller relative to those calculated using completed data in non-Census years. In figures 2(b)-2(c) the unmarked dotted line tends to be below the marked solid line. The reverse holds in Census years: using the entire CM in Census years yields smaller dispersion in 2002 and 2007 than using only ASM plants. This is shown in figure 2 by the missing dips in the solid lines in 2002 and 2007 in the right-hand-side charts.[42]

If we weight non-imputed observations by $ipw1 \times ipw2$, measured dispersion is higher

[42]This reflects the effect of smallest plants which are not selected into the ASM.

(marked and unmarked dotted lines in figures 2(b)-2(c)). This is an important finding because results based on classification and regression trees (CART) in WRP suggest that dispersion measures based on imputed data tend to be smaller than those that take imputation into account.[43] If imputation causes a downward bias in dispersion, our weighting scheme corrects it in the right direction.

We found that imputation has a negative effect on within industry dispersion. Does it affect the relationship between productivity, growth and survival? Our second exercise sheds light on this question by re-estimating the growth and exit regressions using data between 2002 and 2010. The empirical model is the same as what we described in section 4.3.[44] We focus on three cases. In the first one, we use all observations from the CM. In the second case, only non-imputed observations are included but our non-impute weighting scheme is not applied. In the last regressions we weight non-imputed observations by $ipw1 \times ipw2$. Similarly to the previous exercise, we repeat calculations excluding non-ASM cases in Census years.

Table 8 summarizes our results. Imputation seems to decrease the effect of productivity on growth among all establishments (column 1) and on the probability of exit (column 2), regardless on weighting. Point estimates of the effect on growth (exit) in the non-imputed sample are about half (third) the size of those in the completed data. We do not detect such attenuation among incumbents (column 3).

Such variation in the coefficients is by no means negligible, especially if one attempts to answer, say, economic analyses where differences in magnitudes may matter a lot. However, it is not unfamiliar either: we have seen similar quantitative differences when we assessed the effect of TFP estimation methods (see tables 5, A13 and A14). As we mentioned earlier, one of the most common findings in the productivity literature is that more productive plants are more likely to grow and survive. Our last exercise offers evidence that, at least in terms of sign and order of magnitude, this relationship is robust to imputation issues.

[43] See columns 1 and 2 of table 4 in their paper.

[44] Note that as those estimates were calculated using $ipw1$ weights, all regressions in this exercise - based on either completed or non-imputed observations - are also weighted by $ipw1$.

Table 8: The effect of TFP on outcomes between 2002-2010 in the ASM/CM, various samples within the ASM/CM,g rowth accounting based TFP.

	smpl	g	exit	gc
	(1)	(2)	(3)	(4)
ASM-CM				
	Complete data [ipw1]†	0.165***	-0.065***	0.04***
		(0.0152)	(0.0077)	(0.0054)
	Non-imputed data [ipw1]††	0.083***	-0.019***	0.047***
		(0.0088)	(0.0035)	(0.0049)
	Non-imputed data [ipw1×ipw2]†††	0.079***	-0.021***	0.041***
		(0.0108)	(0.0036)	(0.0079)
ASM cases				
	Complete data [ipw1]	0.149***	-0.06***	0.035***
		(0.0141)	(0.0071)	(0.0059)
	Non-imputed data [ipw1]	0.086***	-0.022***	0.044***
		(0.0094)	(0.0033)	(0.006)
	Non-imputed data [ipw1×ipw2]	0.086***	-0.024***	0.041***
		(0.011)	(0.0035)	(0.0084)
Sample size (in thousands)				
ASMCM				
	Complete data [ipw1]	594	594	570
	Non-imputed data [ipw1]	263	263	258
	Non-imputed data [ipw1×ipw2]	263	263	258
ASM cases				
	Complete data [ipw1]	400	400	386
	Non-imputed data [ipw1]	218	218	214
	Non-imputed data [ipw1×ipw2]	218	218	213

Standard errors are in parantheses. †Weighted Census data [pw1]: All observations, propensity score weighted, where the weight is inversely proportional to the probability that the plant is selected into the ASM/CM (see FGH). ††Non-imputed[pw1]: non-imputed subset of the same sample. †††Non-imputed [pw1,pw2]: non-imputed subset of the sample, where observations are weighted by a composite propensity score, where $pw2$ is inversely proportional to the probability that a plant's TFP is calculate using non-imputed data. Probabilities were estimated separately for each year and are based on industry-, size-, age-class and payroll-decile fixed effects. See appendix for more details.

4.6 Homogeneity, dispersion and growth

In the preceding analysis, we calculated TFP using a single value for each elasticity in each industry over the entire time period implying we smoothed out all differences in the $\hat{\beta}$-s that may be present across establishments and over time. This approach is the empirical equivalent of assuming that production function parameters are time invariant and constant across establishments within an industry. Whether or not such homogeneity assumptions are restrictive depends on at least two properties: the underlying within-industry differences in technology and sample size.

For example, if there are within-industry differences in plants' technologies then pooling the data from an entire industry of plants may be too restrictive. In this case, allowing for plant-level heterogeneity in the β-s, empirical feasibility aside, better accounts for within-industry differences in factor intensities. Unfortunately, testing for the true degree of heterogeneity is not straightforward because the results on which we base inference about heterogeneity are endogenous to both the estimation method and the homogeneity assumption. In principle, the optimal way would be to carry out our own taxonomy, which is beyond the scope of this paper. On the other hand, pooling data may in general be necessary to increase sample size in order to reduce finite-sample bias and increase precision.[45] Some of the estimators reviewed in this paper require more data than others meaning that for these methods pooling is not a question of bias and precision but feasibility.[46] A more general but equally important point in this regard is that pooling also implies results are less likely to be sensitive to measurement error, which typically is present in microdata.

We assess the consequences of allowing for plant-level heterogeneity in the elasticities by comparing dispersion statistics which are based on industry- and establishment-specific $\hat{\beta}$-s. We use the growth accounting framework because calculating measures of plant-level elasticities is straightforward in this approach. Figure 3 shows our main result. We find that dispersion[47] is substantially higher if we use establishment-level cost shares to calculate TFP compared to when it is based on industry-level shares. In particular, plant-specific shares increase the interquartile range by a factor of almost 2 (see the difference between the thin and thick solid lines in the figure). This finding indicates that the assumption of establishment-level heterogeneity in elasticities is not innocuous. Our results also indicate that allowing for timeseries variation in the $\hat{\beta}$-s is more likely to affect the volatility, than the level of dispersion. Moreover, such an effect is quantitatively significant only if TFP is based on establishment-level shares (thin dased and dotted lines). The overall conclusion is that the effect of timeseries smoothing is dwarfed by that of cross-section smoothing.

[45] We saw in section 4.3 that sometimes even cross-industry pooling may be necessary.

[46] See section 4.1 for examples.

[47] All statistics are based on the deviation of log TFP from its industry- and year-specific mean.

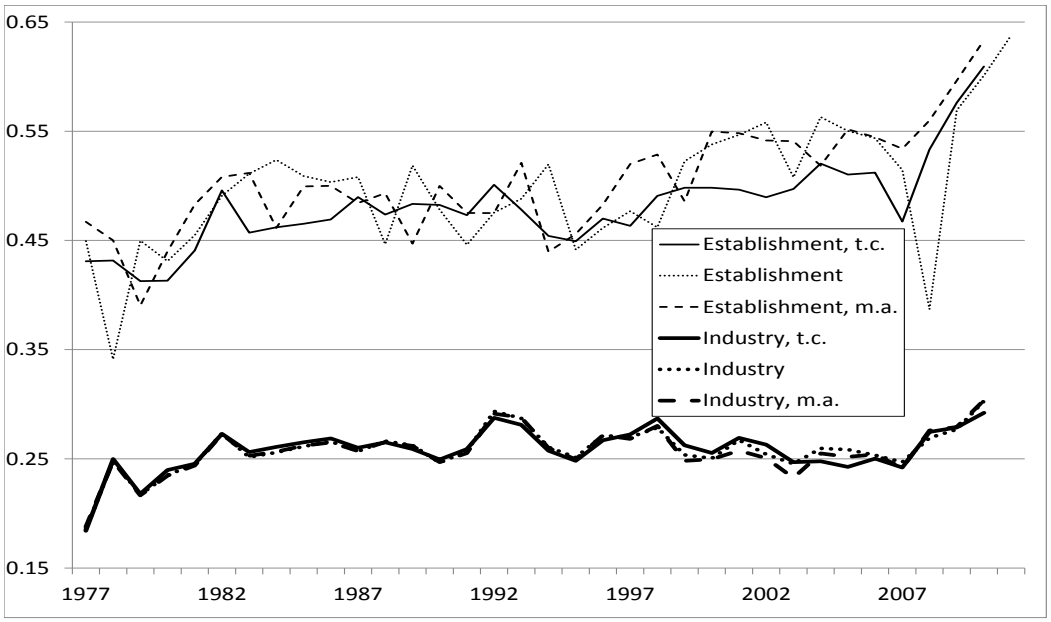


Figure 3: Within-industry interquartile range measures of growth accounting based log-TFP, 10 largest industries. Thin and thick lines denote dispersion measures based on shares calculated at the establishment- and industry-levels, respectively. *t.c.* and *m.a.* label cases where statistics are based on time-averages of share. *t.c.* denotes a case where the average is taken over all available time-periods, i.e. it is constant. *m.a.* is based on the moving average of periods t and t-1.

Both theory and earlier research[48] indicates that at least part of this increase in dispersion is spurious. Our own analysis also offers indirect evidence that the increased variation is at least partly noise. Table 9 shows results from growth and exit regressions based on growth accounting TFP variants. Columns 1-2 show that using time-varying instead of constant industry-level shares leaves the effect of productivity virtually unchanged. In contrast, using plant-level shares (columns 3-5) make point estimates shrink - although they remain statistically significant. The consequence of using plant-level shares is that the growth effect of productivity drops by approximately 50% (columns 1 and 4) among all establishments (first row) or continuers (last row). Among continuers, an additional 25% of this effect disappears if we also allow for timeseries variation in the shares (columns 3 and 5). The effect on the exit probability is smaller, about 33%. We interpret this attenuation as a sign that the data contains more noise with respect to the relationship between productivity and growth. Note that constant plant-level shares imply a higher point estimate among continuers (0.032^{***} in column 4 relative to 0.018^{***} in columns 3 and 5) suggesting there may be noise not only in the cross-section but also in the timeseries variation of establishments.

Our overall conclusion is that calculating cost shares at the plant level is unlikely to be opti-

[48]For example, Syverson (2004).

mal. Although plant-level shares may better capture within-industry differences in technology, they are also prone to measurement error. This is reflected in the smaller partial correlations between productivity, growth and the survival probability of establishments.

Table 9: The effect of growth accounting TFP variants on outcomes, 50 largest industries. Outcomes are: employment growth among all establishments (rows 1-2), exit (rows 3-4), employment growth among continuers (rows 5-6).

	Industry-level shares (1-2)		Plant-level shares (3-5)		
	(1)	(2)	(3)	(4)	(5)
	constant $\left[\bar{s}^j\right]$	time-varying $\left[\frac{s_t^j+s_{t-1}^j}{2}\right]$	time-varying $\left[s_{it}^j\right]$	constant $\left[\bar{s}_i^j\right]$	time-varying $\left[\frac{s_{it}^j+s_{it-1}^j}{2}\right]$
overall growth	0.190***	0.193***	0.1***	0.114***	0.096***
	(0.014)	(0.014)	(0.008)	(0.01)	(0.007)
exit	-0.064***	-0.066***	-0.041***	-0.042***	-0.039***
	(0.006)	(0.005)	(0.003)	(0.004)	(0.004)
cont. growth	0.067***	0.067***	0.018***	0.032***	0.018***
	(0.006)	(0.006)	(0.003)	(0.006)	(0.004)

These estimates are based on the same 50 largest industries used to generate the results in table 5. The entries in column 1 above are identical to the last column of the first panel in table 5.

5 Concluding remarks

Does the method of estimating productivity matter for the main conclusions of the literature? The answer is yes and no, depending on how we ask the question.

The choice of estimation method seems to affect elasticity distributions, especially that of capital. This in turn yields substantial variation in returns to scale across methods. One needs to take this variation into account when answering questions related to the magnitude of elasticities.

The differences in elasticities map into differences in within-industry productivity dispersion. We estimate that the plant at the 75th percentile generates between 25%-50% percent more revenue with the same amount of inputs than the plant at the 25th percentile. This is an average calculated over industries and time, which suggest one must exercise caution when using industry-specific dispersion measures for economic analyses where small changes may make a big difference. However, these results also confirm that there is enormous heterogeneity in establishments' productivity levels, which is in line with what recent microeconomic research found.

Canonical models of firm dynamics describe growth and survival decisions of plants' as functions of idiosyncratic productivity shocks. The main prediction from these models is that plants with positive shocks expand while plants with negative shocks downsize and/or exit. We test this prediction and find that more productive establishments grow significantly faster and are more likely to survive than their less productive competitors. We find that some estimation methods yield outliers including negative factor elasticities and quantitative results are sensitive to the treatment of these outliers.

Does the choice of estimation method have implications for the sources of aggregate productivity growth? We find that despite the fact that some methods seem to be more sensitive than others, and that the relative importance of reallocation and within-plant growth depends on the estimator, the main conclusions of the decomposition literature hold under different estimator variants: reallocation is productivity enhancing and within-plant growth seems to be more important for cyclical fluctuations. However, again here the quantitative details differ substantially. Some methods imply virtually all of the aggregate productivity growth is due to reallocation while others yield that only 25 percent is due to reallocation.

Recent research found that dispersion analyses based on imputed data may underestimate true productivity differences. Our results confirm that accounting for imputation implies greater dispersion. As for the growth effect of productivity, imputation seems to have some attenuating effect, but the general conclusions hold.

Since pooling data is necessary for the econometric estimation methods, we have to be willing to accept homogeneity assumptions on factor elasticities. This implies elasticities are interpreted as within-industry averages of plant-level elasticities. One can argue that plant-

specific elasticities better capture technological differences between establishments and also affect productivity differences across plants. We find that TFP dispersion is about two times higher if we allow for such heterogeneity. Using this more dispersed TFP in growth and exit regressions implies weaker relationships between productivity, growth and survival. This indicates at least some of the increase in dispersion likely is due to measurement error, which implies using plant-level shares is unlikely to be optimal.

References

Daniel Ackerberg, Kevin Caves, and Garth Frazer. Structural identification of production functions. MPRA Paper 38349, University Library of Munich, Germany, December 2006.

Martin N. Baily, Charles R. Hulten, and David Campbell. Productivity dynamics in manufacturing plants. In Martin N. Baily and C. Winston, editors, *Brookings Papers on Economic Activity: Microeconomics*, volume 4. DC: Brookings Institute, 1992.

Martin N. Baily, Eric J. Bartelsman, and John C. Haltiwanger. Labor productivity: Structural change and cyclical dynamics. *The Review of Economics and Statistics*, 83(3):420–433, 2001.

Eric J. Bartelsman and Wayne Gray. The nber manufacturing productivity database. Working Paper 205, National Bureau of Economic Research, October 1996.

Johannes Van Biesebroeck. Robustness of Productivity Estimates. *Journal of Industrial Economics*, 55(3):529–569, 09 2007.

Ricardo J. Caballero, Eduardo M. R. A. Engel, and John C. Haltiwanger. Aggregate Employment Dynamics: Building from Microeconomic Evidence. *American Economic Review*, 87 (1):115–37, March 1997.

Russell W. Cooper and John C. Haltiwanger. On the Nature of Capital Adjustment Costs. *Review of Economic Studies*, 73(3):611–633, 2006.

Richard Ericson and Ariel Pakes. Markov-Perfect industry dynamics: A framework for empirical work. *Review of Economic Studies*, 62(210):53–82, 1995.

Lucia S. Foster, John C. Haltiwanger, and C. J. Krizan. Aggregate Productivity Growth. Lessons from Microeconomic Evidence. In *New Developments in Productivity Analysis*, NBER Chapters, pages 303–372. National Bureau of Economic Research, Inc, 2001.

Lucia S. Foster, John C. Haltiwanger, and C. J. Krizan. Market Selection, Reallocation, and Restructuring in the U.S. Retail Trade Sector in the 1990s. *The Review of Economics and Statistics*, 88(4):748–758, November 2006.

Lucia S. Foster, Cheryl A. Grim, and John C. Haltiwanger. Reallocation In The Great Recession: Cleansing Or Not? Working Papers 13-42, Center for Economic Studies, U.S. Census Bureau, August 2013.

Xavier Gabaix. The Granular Origins of Aggregate Fluctuations. *Econometrica*, 79(3):733–772, 05 2011.

Hugo A. Hopenhayn. Entry, Exit, and Firm Dynamics in Long Run Equilibrium. *Econometrica*, 60(5):1127–50, September 1992.

Boyan Jovanovic. Selection and the Evolution of Industry. *Econometrica*, 50(3):649–70, May 1982.

James A. Levinsohn and Amil Petrin. Estimating production functions using inputs to control for unobservables. *The Review of Economic Studies*, 70(2):317–341, April 2003.

Steven G. Olley and Ariel Pakes. The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263–1297, 1996.

Amil Petrin and James A. Levinsohn. Measuring aggregate productivity growth using plant-level data. *The RAND Journal of Economics*, 43(4):705–725, 2012.

Amil Petrin, Jerome Reiter, and T. Kirk White. The Impact of Plant-level Resource Reallocations and Technical Progress on U.S. Macroeconomic Growth. *Review of Economic Dynamics*, 14(1):3–26, January 2011.

Robert M. Solow. Technical Change and the Aggregate Production Function. *Review of Economics and Statistics*, 39:312–320, August 1957.

Chad Syverson. What determines productivity. *Journal of Economic Literature*, 49(2):326–365, 2011.

Gabriel Y. Weintraub, C. Lanier Benkard, and Benjamin Van Roy. Industry dynamics: Foundations for models with an infinite number of firms. *Journal of Economic Theory*, 146(5): 1965–1994, September 2011.

T. Kirk White, Jerome P. Reiter, and Amil Petrin. Plant-level Productivity and Imputation of Missing Data in U.S. Census Manufacturing Data. NBER Working Papers 17816, National Bureau of Economic Research, Inc, February 2012.

Jeffrey M. Wooldridge. On estimating firm-level production functions using proxy variables to control for unobservables. *Economics Letters*, 104(3):112–114, September 2009.

6 Appendix

6.1 Remarks on proxy methods

The two-step procedure described in Levinsohn and Petrin (2003) became widely used among practitioners in the past decade. Its advantages and caveats have been discussed in the literature, see for example Ackerberg-Caves-Frazer (2006). Part of the reasons why it gained popularity likely is the fact that easy-to-use routines have been developed to implement it in Stata. The available versions of the algorithm differ according to the production function type and the numerical procedure used to minimize the GMM criterion function. If the production function is defined in terms of value added then the criterion function is minimized using the so-called golden section search algorithm (GSS). For output-based functions, the user may choose between a non-linear gradient-based (NL) routine and a grid search (GR).

Our results suggest these choices entail non-trivial differences in the distributions of elasticities. This is because while GSS and NL are guaranteed to find the optimum points of unimodal functions they may get stuck at local optima if there are multiple modes in the parameter distribution. GR searches over the state space on a pre-defined grid, which is safe if the grid is not too coarse. However, GR may also require long computing time.[49]

6.2 Imputation and plant characteristics

Evidence in WPR suggests dispersion measures based on imputed data tend to underestimate true productivity dispersion. The effect of imputation on growth coefficients is less obvious for reasons we will outline shortly.

The effect of imputation on dispersion is easiest to illustrate by considering the results of regression-based imputation, which also happens to be one of the most frequently used methods at the Census Bureau to impute the components of TFP.[50] Regression-based imputation amounts to substituting fitted values of a regression for values of the underlying distribution. That is, using the regression line, $E(y_i|x_i)$, instead of draws from the conditional distribution of y_i. This essentially means a collapse in variation in the data.

How can we approximate the unobserved conditional distribution of y_i? One way is to use a known distribution to simulate data from it. Another is to draw from the set of non-imputed observations with similar characteristics to the ones in the imputed sample, which is what classification and regression tree (CART) methods do (see WPR). Yet another approach is to use non-imputed observations only. This works if imputation is random. However, if the probability of imputation is correlated with plant characteristics, excluding imputed observations generates selection issues.

Both earlier evidence and our analysis - as we will see shortly - indicate that imputation

[49] Estimating the 459 LPGR β-sets for the APG comparisons of section 4.4 took about 28 days.

[50] We do not discuss individual imputation models further but more details are available upon request.

is not random in our data. Consequently, empirical models based on non-imputed data must take this issue into account otherwise selection renders results biased.

We address imputation in two steps. First, we use a logistic regression to describe the relationship between plant characteristics and the probability that TFP is calculated using non-imputed data. Next, we use inverse propensity scores to weight observations in the non-imputed sample to calculate dispersion statistics and growth coefficients. To be specific, we estimate the following equation separately for each year between 2002 and 2010:

$$Pr(I_{it} = 1|X_t) = E[I_{it}|X_{it}] = F(X_{it}\beta_t),$$
$$F(x) = \frac{1}{1 + e^{-x}}.$$

where I_{it} denotes an indicator variable equal to 1 if any of the components of TFP is non-imputed. The main components are: plants' total value of shipments (TVS), production hours (PH), salaries and wages (SW), production workers wages (WW), cost of parts (CP).

X_{it} and β_t denote a vector of controls and coefficients. We use the Census Bureau's impute flags to determine whether an item is imputed. More details on imputation procedures at the Bureau are available from the authors. See also table A1 in WPR. Control variables X_{it} are included to capture plant characteristics: industry effects, employment size class, payroll deciles and age class fixed effects. We have 86 4-digit NAICS industries. To control for size, we defined 10 size classes based on employment (1-9, 10-109, 20-29, 30-49, 50-99, 100-149 150-249, 250-499, 500-999, 1000-) in addition to the payroll deciles. Finally, we classify establishments into 9 age classes (births, 1, 2, 3, 4, 5, 6-10, 11-15, 16- years). As mentioned in the main text, data for these variables source from the LBD.

Imputation rates differ across variables. The upper panel of table A16 shows that from among the main components of TFP, PH and CP are imputed the most, and SW the least frequently. Imputation rates vary not only with variables but also sample definition.[51] The last row in the upper panel indicates imputation tends to be less frequent among ASM establishments. Overall, less frequent imputation implies almost 10-percentage-point smaller imputation rate for TFP (last column in the lower panel). This suggests it may be worth exploring the effects of restricting ourselves to ASM cases when estimating propensity scores. Therefore, we present results also for a scenario where only ASM establishments are included in the analysis.

Comparing size-, age-, and payroll distributions across non-imputed and completed samples suggests imputation instances are correlated with plant characteristics.[52] TFP components are more likely to be imputed for smaller and younger establishments with less payroll. These three characteristics give a multitude of possible regressor sets for the probability model. We

[51] We can measure imputation rates in the entire ASM/CM (row 1), restricting ourselves to using observations for which there exists size and age information in the LBD (row 2), in the ASM only (row 3).

[52] These are undisclosed results and not shown here.

experimented with six of those ([1] employment size; [2] payroll; [3] employment size, payroll; [4] employment size, age; [5] payroll, age; [6] employment size, payroll, age) and found that the basic implications do not change. However, we also found that including all three variables provides a somewhat better fit[53] than any of the remaining five. Therefore, we present results based on propensity scores from a model with establishment size, payroll and age.

Figure A2 plots the point estimates from this model. We conclude that TFP data for smaller and younger plants with less payroll are significantly more likely to be imputed.[54] In section 4.5, when calculating weighted dispersion measures and regression coefficients, we use the inverse of probabilities implied by these models as weights.

As a final point, we mention that diagnostics indicate logistic regressions fit ASM establishments better than ASM/CM establishments.[55] The difference in the AIC is about a factor of 3 and 2 in 2002 and 2007, respectively.

[53] As measured by the AIC, not shown here.

[54] Point estimates are precise enough such that we can confirm the positive relationship between plant-size, -age and the probability of TFP being non-imputed. More details on point estimates and standard errors are available upon request.

[55] Available upon request.

7 Tables in Appendix

Table A1: Descriptive statistics on industry assigment.

Panel 1. Distribution of plants, by year.

observed year	frequency (1000)	percent
1987	1156	51.3
1988	5.1	0.2
1989	7.6	0.3
1990	9.8	0.4
1991	3.5	0.2
1992	230	10.2
1993	20.8	0.9
1994	30.3	1.3
1995	19.4	0.9
1996	15.1	0.7
1987-1996	1499	66.5
1972-1986 (random assignment)	754	33.5
Total	2252	100

Panel 2. The effect of industry assignment on switchers.

	Assigned industry		
Original industry	no switch	switch	Total
frequency (1000)	38	92	130
percent	29	71	100

Panel 1. shows frequency counts of time periods which were used to assign plants into industries. Panel 2. breaks down switching instances (plants with 2 or more SIC codes in their timeseries) under the original classification. The first entry in the last row says that 29% of the switching instances in the original classification system disappear under random assignment.

Table A2: Conditional probabilities of an industry moving from quintile i (row) to quintile j (column) of the $\hat{\beta}_K/\hat{\beta}_L$-distribution, 50 largest industries.

	OP	1	2	3	4	5
LPGR	1	0.5	0.2	0.1	0.1	0.1
	2	0.3	0.3	0.2	0.2	0
	3	0.1	0.3	0.3	0.1	0.2
	4	0.1	0.1	0.3	0.2	0.3
	5	0.1	0.1	0.4	0.3	0
	GA	1	2	3	4	5
LPGR	1	0.1	0.2	0.3	0.2	0.2
	2	0.3	0.2	0.3	0.2	0
	3	0.3	0.2	0.2	0.2	0.1
	4	0.2	0.3	0.1	0.3	0.1
	5	0.1	0.1	0.1	0.3	0.4
	WLPM	1	2	3	4	5
LPGR	1	0.4	0.3	0.1	0.1	0
	2	0.2	0.1	0.3	0.2	0.1
	3	0.3	0.2	0.3	0.1	0.1
	4	0.1	0.3	0.2	0.2	0.2
	5	0.1	0.3	0.6	0	0

Each entry shows the fraction of industries which moved from a quintile under LPGR (row) to a quintile under another estimator (column). See notes to table 1 for legends. The entries in each row sum to 1, except rounding.

Table A3: Descriptive statistics of 50 largest industries ordered by the within-industry number of plant-year observations between 1976-1996. The industry classification contains SIC 1987 industry codes concorded 1-to-1 SIC 1972 and NAICS 1997.

rank	SIC 1987	N	MeanL	Nplants
1	2	3	4	5
1	2711	39	9.9	3.9
2	3273	37	10.1	3.7
3	2411	37	4.8	7.6
4	3441	32	11.8	2.6
5	2653	30	22.1	1.3
6	3471	29	10.1	2.8
7	2051	26	16.1	1.6
8	3451	23	10.7	2.1
9	2951	22	12.8	1.7
10	3442	16	12.4	1.3
20	3231	12	11.9	1
50	3613	8	18	0.4
1-50 total		710	11	67

Column 1: rank; column 2: number of plant-year observations in thousands; column 3: average number of observations per plant; column 4: number of plants in thousands.

Table A4: Industries with positive, negative and non-estimable elasticities, as a per cent of the total number of industries. Results are based on 459 4-digit industries between 1972-2010, described in more detail at the end of section 3.

	1	2	3	4	5	6	7	8	9	10	11	12
	-	0* or n.e.**	+	-	0* or n.e.**	+	-	0* or n.e.**	+	-	0* or n.e.**	+
	$\widehat{\beta}_k$			$\widehat{\beta}_L$			$\widehat{\beta}_M$			$\widehat{\beta}_E$		
ols	2		98	2		98			100	4		96
op	3	18*	79	2	18*	80		18*	82	3	18*	79
lpva	4		96	3		97						
lpnl		16**	84	4		96			100	6		94
lpgss	8		92	4		96			100	6		94
lpgr			100	4		96			100	6		94
wlpe	10		90	6		94			100	32		68
wlpm	16		84	7		93	2		98	11		89

0* or n.e.**: $\widehat{\beta}$ is zero or could not be estimated. *: The first entry in column 2 says OP delivers error in 18% of industries. This happens because the algorithm stalls in industries with insufficient information on exit (or investment). **: The second entry in column 2 says gradient-based optimization in LPNL yields $\widehat{\beta}_K = 0$ in 16% of the industries.

Table A5: Average number of within-industry observations by groups, as in table A4.

	1	2	3	4	5	6	7	8
	-	0 or n.e.*	-	0 or n.e.*	-	0 or n.e.*	-	0 or n.e.*
	$\widehat{\beta}_k$		$\widehat{\beta}_L$		$\widehat{\beta}_M$		$\widehat{\beta}_E$	
ols	0.59		0.46				0.32	
op	0.56	0.39	0.58	0.39		0.4	0.42	0.39
lpva	0.26		0.29					
lpnl		0.96	0.55				0.42	
lpgss	0.35		0.55				0.42	
lpgr			0.55				0.42	
wlpe	0.4		0.56			0.14	0.93	
wlpm	0.7		0.52		0.61		0.44	

All entries are calculated respectively as $\bar{N}_-/\bar{N}_+$ and $\bar{N}_{0 \text{ or n.e.}}/\bar{N}_+$, where $\bar{N}_-$, $\bar{N}_{0 \text{ or n.e.}}$ and $\bar{N}_+$ denote the average number of observations in industries with negative, zero or non-estimable and positive $\widehat{\beta}_j$.

Table A6: Fraction of industries with plausible returns to scale measures (RTS), by estimation method, as in table A4. Plausible RTS is defined as $RTS \in [L, H]$, $RTS = \sum_j \hat{\beta}_j$, where $[L, H]$ are as shown in the table.

	1	2	3	4	5	6	7
Panel A. $H = 1.3$							
	L						
	0.7	0.75	0.8	0.85	0.9	0.95	average
ols	1	1	1	1	1	0.94	0.99
op	0.76	0.75	0.72	0.69	0.58	0.39	0.65
lpva	0.53	0.46	0.37	0.31	0.23	0.15	0.34
lpnl	0.85	0.78	0.7	0.62	0.53	0.4	0.65
lpgr	0.92	0.89	0.85	0.78	0.66	0.54	0.77
lpgss	0.89	0.82	0.72	0.6	0.48	0.32	0.64
wlpe	0.83	0.81	0.79	0.75	0.69	0.62	0.75
wlpm	0.83	0.81	0.78	0.73	0.68	0.61	0.74
all*	0.27	0.22	0.15	0.09	0.04	0.02	0.13
Panel B. $L = 0.7$							
	H						
	1.3	1.25	1.2	1.15	1.1	1.05	average
ols	1	1	1	1	0.99	0.95	0.99
op	0.76	0.74	0.72	0.71	0.69	0.64	0.71
lpva	0.53	0.52	0.5	0.48	0.47	0.45	0.49
lpnl	0.85	0.85	0.83	0.8	0.76	0.68	0.8
lpgr	0.92	0.9	0.85	0.81	0.74	0.65	0.81
lpgss	0.89	0.88	0.87	0.85	0.82	0.78	0.85
wlpe	0.83	0.78	0.7	0.62	0.52	0.41	0.64
wlpm	0.83	0.78	0.72	0.64	0.55	0.44	0.66
all*	0.27	0.24	0.19	0.14	0.09	0.04	0.16

Each entry in the table shows the fraction of industries with RTS falling into the interval as defined by the upper (H) and lower (L) bounds in the header of the panels. This widest interval ($[.7, 1.3]$ in column 1) is consistent with the average revenue-share of profits being between $[-25\%, +25\%]$, which does not seem to be too restrictive. Columns 2-6 in Panel A show RTS intervals as we increase the lower bound while holding $H = 1.3$. Panel B shows intervals as we decrease the upper bound while holding $L = .7$. *all: this line counts industries where all methods simultaneously imply plausible RTS; the small numbers highlight the differences in elasticity distributions, and quantify the qualitative differences in the elasticity distributions mentioned in section 4.1.

Table A7: Descriptive statistics of tfp-distributions, pre-estimation outliers included.

	N (1000)	iqr	std
10 largest industries			
OLS	230	0.23	0.21
OP	233	0.34	0.41
LPVA	236	0.73	0.60
LPNL	235	0.34	0.41
LPGR	235	0.32	0.31
LPGSS	234	0.30	0.28
WLPE	236	0.34	0.30
WLPM	236	0.37	0.39
GA	231	0.25	0.22
50 largest industries			
OLS	568	0.29	0.28
OP	567	0.35	0.41
LPVA	569	0.77	0.66
LPNL	572	0.34	0.44
LPGR	572	0.33	0.35
LPGSS	569	0.3	0.31
WLPE	575	0.37	0.41
WLPM	575	0.46	2.21
GA	561	0.25	0.24

Table A8: Correlations among within-industry tfp-distributions, 50 largest 4-D industries. Pre-estimation outliers included.

	OLS	OP	LPVA	LPNL	LPGR	LPGSS	WLPE	WLPM	GA
Pearson									
OLS	1								
OP	0.52	1							
LPVA	0.68	0.44	1						
LPNL	0.61	0.31	0.66	1					
LPGR	0.8	0.42	0.78	0.83	1				
LPGSS	0.8	0.56	0.74	0.71	0.87	1			
WLPE	0.49	0.47	0.48	0.47	0.5	0.62	1		
WLPM	0.01	0.02	-0.15	-0.12	-0.15	-0.05	0.15	1	
GA	0.77	0.47	0.55	0.52	0.63	0.68	0.49	0.09	
Spearman									
OLS	1								
OP	0.67	1							
LPVA	0.77	0.62	1						
LPNL	0.77	0.55	0.71	1					
LPGR	0.86	0.6	0.79	0.87	1				
LPGSS	0.81	0.69	0.75	0.79	0.86	1			
WLPE	0.58	0.61	0.55	0.51	0.55	0.68	1		
WLPM	0.34	0.36	0.2	0.23	0.25	0.37	0.48	1	
GA	0.8	0.62	0.59	0.64	0.67	0.71	0.58	0.42	1

This table differs from table 4 in that we computed correlations including pre-estimation outlier observations.

Table A9: Correlations among within-industry tfp-distributions, 10 largest 4-D industries.

	OLS	OP	LPVA	LPNL	LPGR	LPGSS	WLPE	WLPM	GA
Pearson									
OLS	1								
OP	0.48	1							
LPVA	0.61	0.44	1						
LPNL	0.53	0.15	0.41	1					
LPGR	0.71	0.43	0.88	0.48	1				
LPGSS	0.8	0.5	0.85	0.5	0.96	1			
WLPE	0.71	0.54	0.78	0.25	0.8	0.86	1		
WLPM	0.52	0.27	0.44	0.13	0.41	0.54	0.66	1	
GA	0.88	0.42	0.59	0.51	0.68	0.72	0.66	0.5	1
Spearman									
OP	0.71	1							
LPVA	0.61	0.64	1						
LPNL	0.66	0.39	0.62	1					
LPGR	0.71	0.65	0.89	0.72	1				
LPGSS	0.79	0.7	0.85	0.71	0.96	1			
WLPE	0.71	0.73	0.77	0.48	0.79	0.85	1		
WLPM	0.6	0.49	0.52	0.28	0.51	0.58	0.69	1	
GA	0.88	0.63	0.6	0.63	0.68	0.72	0.67	0.56	1

Including pre-estimation outliers barely changes correlations.

Table A10: Correlations among within-industry tfp-distributions, 10 largest 4-D industries. Pre-estimation outliers included.

	OLS	OP	LPVA	LPNL	LPGR	LPGSS	WLPE	WLPM	GA
Pearson									
OLS	1								
OP	0.51	1							
LPVA	0.59	0.45	1						
LPNL	0.53	0.16	0.44	1					
LPGR	0.69	0.43	0.85	0.53	1				
LPGSS	0.79	0.52	0.83	0.54	0.95	1			
WLPE	0.69	0.57	0.76	0.28	0.78	0.85	1		
WLPM	0.5	0.27	0.42	0.12	0.37	0.5	0.64	1	
GA	0.85	0.43	0.57	0.5	0.64	0.68	0.63	0.48	1
Spearman									
OLS	1								
OP	0.7	1							
LPVA	0.6	0.64	1						
LPNL	0.66	0.39	0.63	1					
LPGR	0.71	0.63	0.87	0.74	1				
LPGSS	0.78	0.68	0.84	0.73	0.96	1			
WLPE	0.7	0.73	0.77	0.49	0.78	0.84	1		
WLPM	0.58	0.48	0.5	0.28	0.47	0.54	0.67	1	
GA	0.86	0.61	0.58	0.62	0.66	0.7	0.65	0.54	1

This table differs from table A9 in that we computed correlations including pre-estimation outlier observations.

Table A11: The effect of industry characteristics and β_k on TFP dispersion.

		σ_q	iqr(ω)
Panel 1. Effect of $\hat{\beta}_k$			
baseline	$\hat{\beta}_K = 0.26$	0.81 (min)	0.19
cf 1	$\hat{\beta}_K = 0.17$	0.81 (min)	0.15
cf 2	$\hat{\beta}_K + sd(\hat{\beta}_K) = 0.53$	0.81 (min)	0.42
cf 3	$max(\hat{\beta}_K) = 0.75$	0.81 (min)	0.52
Panel 2. Effect of inputs and output			
cf 4	$\beta_j = \hat{\beta}_j, j = \{K, L, E, M\}$	1.68 (max)	0.34

σ_q denotes the within-industry dispersion of output, β_j denotes the value of elasticity used in the exercise. TFP is based on LPGR in the group 10 largest industries. Row "baseline" represents the minimum output-dispersion industry where the estimated value of β_K was used to calculate TFP and its dispersion. Row "cf 1" represents counterfactual 1, where input/output data are taken from the minimum output-dispersion industry but β_K is set to its sample mean. Rows "cf 2" and "cf 3" represent industries where β_K is set to 1 standard deviation above its estimated value and the sample maximum, respectively. Row "cf 4" represents an industry where input and output data are taken from the maximum output-dispersion industry and all elasticities are taken from the minimum dispersion industry. Since $\sigma_{\sigma_q} = 0.25$, this change amounts to increasing σ_q by about 3.5 standard deviations ($(1.68 - 0.81)/0.25 = 3.48$).

Table A12: Sample size in the specifications shown in table 5. Sample size is measured as the total number of plant-year observations used in a regression (in thousands).

	ols	op	lpva	lpnl	lpgr	lpgss	wlpe	wlpm	ga
4-digit elasticities									
N1	410	415	411	413	413	410	414	414	405
N2	410	415	411	413	413	410	414	414	405
N3	393	398	395	396	396	394	397	397	388
Total N	440	440	440	440	440	440	440	440	440
3-digit elasticities									
N1	405	413	413	413	414	411	413	413	405
N2	405	413	413	413	414	411	413	413	405
N3	389	396	396	396	398	395	396	396	388
Total N	440	440	440	440	440	440	440	440	440
3-digit elasticities, industries with negative or non-estimable elasticities dropped									
N1	391	406	413	352	408	405	289	310	405
N2	391	406	413	352	408	405	289	310	405
N3	375	390	396	337	391	388	277	298	388
Total N	440	440	440	440	440	440	440	440	440

N1-N3 denote the number of plant-year observations used in the three regressions. Total N denotes the original number of plant-year observations before estimation and post-estimation outlier trimming.

Table A13: The effect of TFP on outcomes, all (292) largest industries. Outcomes are: employment growth among all establishments, exit and employment growth among continuers

	ols	op	lpva	lpnl	lpgr	lpgss	wlpe	wlpm	ga
4-digit elasticities									
$dlnE$	0.183***	0.001	0.086***	0.093***	0.128***	0.119***	0.004**	-0.004***	0.203***
	(0.013)	(0.003)	(0.004)	(0.007)	(0.007)	(0.008)	(0.002)	(0.001)	(0.011)
exit	-0.066***	0.001	-0.027***	-0.029***	-0.041***	-0.044***	-0.002*	0.001**	-0.068***
	(0.005)	(0.001)	(0.002)	(0.004)	(0.003)	(0.004)	(0.001)	(0.001)	(0.004)
$dlnE_{cont}$	0.055***	0.002***	0.037***	0.038***	0.05***	0.035***	0.001	-0.001*	0.071***
	(0.007)	(0.001)	(0.003)	(0.004)	(0.005)	(0.005)	(0.001)	(0.001)	(0.006)
3-digit elasticities									
$dlnE$	0.202***	0.115***	0.098***	0.13***	0.15***	0.162***	0.001	-0.012***	0.193***
	(0.011)	(0.012)	(0.006)	(0.011)	(0.008)	(0.013)	(0.001)	(0.002)	(0.01)
exit	-0.072***	-0.037***	-0.03***	-0.041***	-0.045***	-0.058***	0	0.004***	-0.066***
	(0.005)	(0.005)	(0.003)	(0.005)	(0.004)	(0.005)	(0)	(0.001)	(0.004)
$dlnE_{cont}$	0.062***	0.046***	0.043***	0.053***	0.067***	0.051***	0	-0.005***	0.066***
	(0.006)	(0.006)	(0.003)	(0.005)	(0.005)	(0.005)	(0)	(0.002)	(0.006)
3-digit elasticities, industries with implausible or non-estimable elasticities are dropped									
$dlnE$	0.207***	0.159***	0.1***	0.124***	0.169***	0.154***	0.078***	0.06***	0.193***
	(0.013)	(0.011)	(0.006)	(0.01)	(0.009)	(0.011)	(0.008)	(0.012)	(0.01)
exit	-0.073***	-0.05***	-0.03***	-0.037***	-0.05***	-0.052***	-0.031***	-0.022***	-0.066***
	(0.005)	(0.004)	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)
$dlnE_{cont}$	0.065***	0.065***	0.045***	0.054***	0.075***	0.054***	0.017***	0.017***	0.066***
	(0.006)	(0.006)	(0.003)	(0.006)	(0.006)	(0.006)	(0.005)	(0.005)	(0.006)

Estimates are taken from equations of three outcomes on plant-level measure of TFP, state-level measure of unemployment growth, and year-, sizeclass- and state-fixed effects. Standard errors (in parentheses) are clustered at the state level.

Table A14: The effect of TFP on outcomes, 10 largest industries. Outcomes are: employment growth among all establishments, exit and employment growth among continuers

	ols	op	lpva	lpnl	lpgr	lpgss	wlpe	wlpm
4-digit elasticities								
$dlnE$	0.206***	0.051***	0.088***	0.088***	0.154***	0.136***	0.099***	-0.027**
	(0.022)	(0.018)	(0.01)	(0.011)	(0.015)	(0.021)	(0.015)	(0.014)
exit	-0.067***	-0.002	-0.025***	-0.043***	-0.046***	-0.044***	-0.022***	0.014***
	(0.01)	(0.007)	(0.003)	(0.005)	(0.007)	(0.007)	(0.005)	(0.005)
$dlnE_{cont}$	0.076***	0.054***	0.042***	-0.001	0.066***	0.051***	0.06***	0.003
	(0.013)	(0.006)	(0.007)	(0.006)	(0.012)	(0.014)	(0.011)	(0.01)
3-digit elasticities								
$dlnE$	0.208***	0.158***	0.089***	0.112***	0.164***	0.142***	0.109***	0.016
	(0.023)	(0.021)	(0.010)	(0.019)	(0.016)	(0.022)	(0.015)	(0.018)
exit	-0.068***	-0.032***	-0.025***	-0.029***	-0.049***	-0.047***	-0.026***	0.002
	(0.010)	(0.008)	(0.004)	(0.006)	(0.007)	(0.008)	(0.005)	(0.006)
$dlnE_{cont}$	0.075***	0.102***	0.042***	0.058***	0.07***	0.051***	0.061***	0.023*
	(0.013)	(0.012)	(0.007)	(0.015)	(0.013)	(0.014)	(0.012)	(0.012)
3-digit elasticities, industries with negative or non-estimable elasticities dropped								
$dlnE$	0.208***	0.158***	0.089***	0.113***	0.164***	0.142***	0.13***	0.061***
	(0.023)	(0.021)	(0.010)	(0.021)	(0.016)	(0.022)	(0.019)	(0.021)
exit	-0.068***	-0.032***	-0.025***	-0.034***	-0.049***	-0.047***	-0.043***	-0.018***
	(0.010)	(0.008)	(0.004)	(0.007)	(0.007)	(0.008)	(0.008)	(0.007)
$dlnE_{cont}$	0.075***	0.102***	0.042***	0.047***	0.07***	0.051***	0.046***	0.027*
	(0.013)	(0.012)	(0.007)	(0.014)	(0.013)	(0.014)	(0.013)	(0.014)

Estimates are taken from equations of three outcomes on plant-level measure of TFP, state-level measure of unemployment growth, and year-, sizeclass- and state-fixed effects. Standard errors (in parentheses) are clustered at the state level.

Table A15: Timeseries of the contribution of reallocation (RE) and within-plant productivity growth, in per cent. These numbers were used to calculated averages in columns 8 and 9 in table 6.

	WLPE	WLPM	PWR	WLPE	WLPM	PWR	WLPE	WLPM	PWR
	1	2	3	4	5	6	7	8	9
	Total RE (1-3)			Capital RE (4-6)			Labor RE (7-9)		
1977	2.2	2.4	4.6	0.4	0.2	0.9	0.2	0.2	0.9
1978	1.8	1.9	2.4	0.8	0.5	0.2	0.2	0.2	0.6
1979	0.9	1.6	1	0.8	0.5	0.6	0.2	0.1	0.3
1980	0.3	-0.4	-0.3	1.3	0.8	1.1	-0.1	-0.2	-0.4
1981	1.5	1.8	1.4	1.2	0.8	0.7	0.3	0.2	0
1982	0.1	-1.1	-1.4	1	0.6	-0.4	-0.3	-0.3	-0.4
1983	1.4	2.1	1.6	0.9	0.7	-0.3	0	-0.1	0.5
1984	4	5.3	4.9	1.2	0.8	0.1	0.2	0.1	0.6
1985	2.1	1.8	3.5	1	0.7	1.3	0.3	0.2	0.2
1986	1.4	0.5	3.9	0.6	0.5	1.7	0.1	0.2	0.4
1987	1.3	2.1	2.9	0.7	0.6	1.2	-0.1	0	0.4
1988	1.7	2	2.4	0.6	0.5	0	0.3	0.3	0.6
1989	0.9	1.2	1.7	0.7	0.5	0.5	0.1	0.1	0.7
1990	-0.2	-0.5	-1.1	0.5	0.5	1.5	-0.7	-0.7	0.9
1991	0.2	-0.7	1.9	0.5	0.5	2	0.2	0.7	0.7
1992	1.3	1.9	1.2	0.6	0.5	0.5	0	0.1	0.5
1993	1.9	1.3	2.6	0.9	0.5	1.3	-0.1	-0.1	0.6
1994	2	2.6	3	0.7	0.6	0.8	0	0	0.1
1995	1	1.9	2.4	1.2	1.1	1.3	-0.3	-0.4	0.2
1996	2.6	3	3.2	1	0.8	1.2	-0.1	-0.2	0
mean(77-96)	1.4	1.5	2.1	0.8	0.6	0.4	0	0	0.4
std(77-96)	1	1.4	1.7	0.2	0.2	0.5	0.2	0.3	0.5
	Materials RE (1-3)			Energy RE (4-6)			Within-plnat growth (7-9)		
1977	1.5	1.9	2.2	0	0	0.6	-0.8	0.6	-0.5
1978	1	1.3	1.5	-0.1	-0.2	0	1.4	1.1	1
1979	0.7	0.9	0.5	-0.7	0.1	-0.4	3	2.1	3.1
1980	-0.2	-1	-0.5	-0.8	0	-0.4	-4.3	-3.4	-3.9
1981	0.3	0.4	-0.1	-0.3	0.4	0.8	0.4	-0.2	-0.1
1982	-1.2	-2	-2	0.6	0.5	1.4	-3.7	-4.3	-2.9
1983	0.7	1.5	0.2	-0.2	0	1.2	4.9	3.2	4.2
1984	2.6	4	3.4	0	0.3	0.7	0.6	-0.7	1.9
1985	0.6	0.7	0.9	0.3	0.1	1.1	-1.8	-3.5	-3.5
1986	0.1	-0.4	0.7	0.5	0.2	1.1	-2.2	-1.9	-4.3
1987	0.8	1.4	1	-0.2	0.1	0.3	3.6	3.1	3.1
1988	0.9	1.2	1.2	-0.1	0	0.6	1.5	2.5	2.1
1989	0.5	0.7	0.4	-0.3	-0.1	0.1	0.4	-1.7	-2.3
1990	-0.3	-0.4	-1.8	0.3	0.1	-1.7	-1.4	-1.5	-0.4
1991	-0.7	-1.8	-0.7	0.2	0	-0.1	-1.7	-0.7	-2.7
1992	0.7	1.1	0.9	0	0.2	-0.6	-0.4	1.8	3
1993	0.2	0.8	0.2	0.8	0.1	0.4	0.9	2.5	0.1
1994	1.1	1.8	1.9	0.1	0.2	0.2	3	1.8	3.9
1995	0.9	1.2	1.3	-0.9	0	-0.3	2.1	1.1	2.2
1996	1.9	2.3	2.1	-0.1	0	-0.1	-0.8	-1.7	0.6
mean(77-96)	0.6	0.8	0.7	0	0.1	0.2	0.2	0	0.2
std(77-96)	0.9	1.4	1.3	0.4	0.2	0.7	2.4	2.3	2.7

Table A16: Imputation rates of the main components of TFP in the ASM/CM, as a per cent of plant-year observations between 2002-2010.

Average imputation rates (2002-2010)

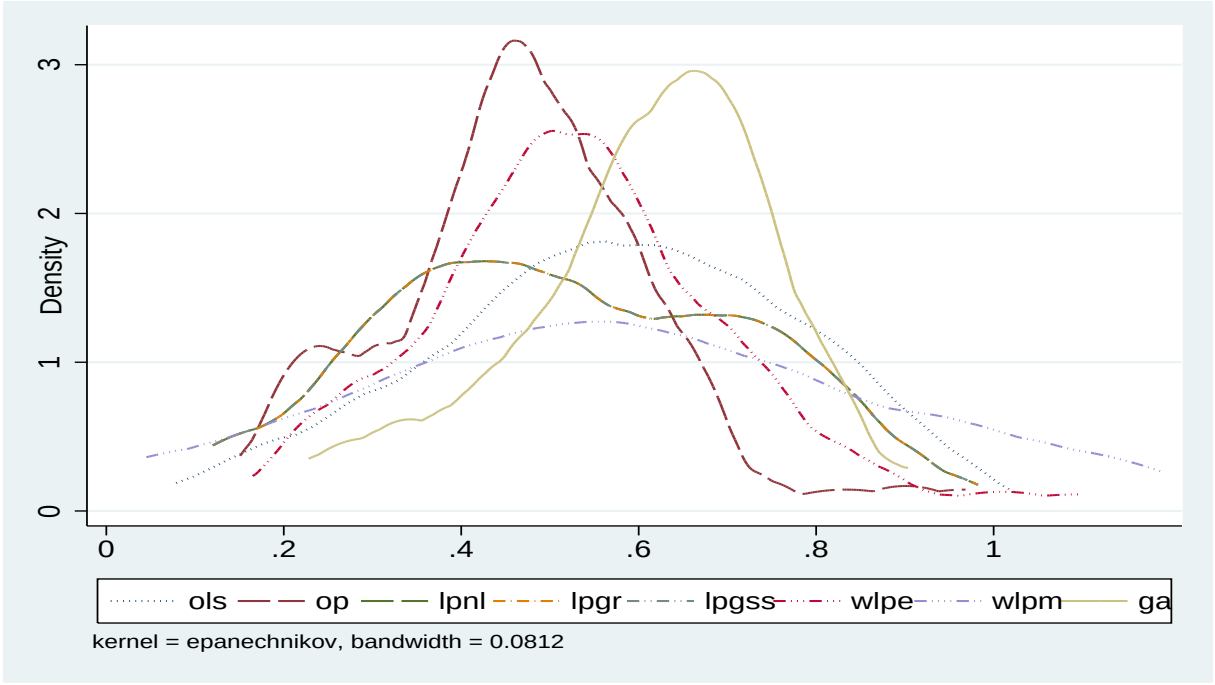
	TFP	TVS	PH	SW	WW	CP	CM
ASM/CM*	59.1	28.1	42.6	14.2	29.2	42.6	37.5
ASM/CM/LBD**	58.3	27.6	42	14.1	28.5	41.7	36.3
ASM/LBD***	48.9	25.5	35.3	13.2	25.1	34	31.9

Yearly imputation rates of TFP, per cent of plant-year observations

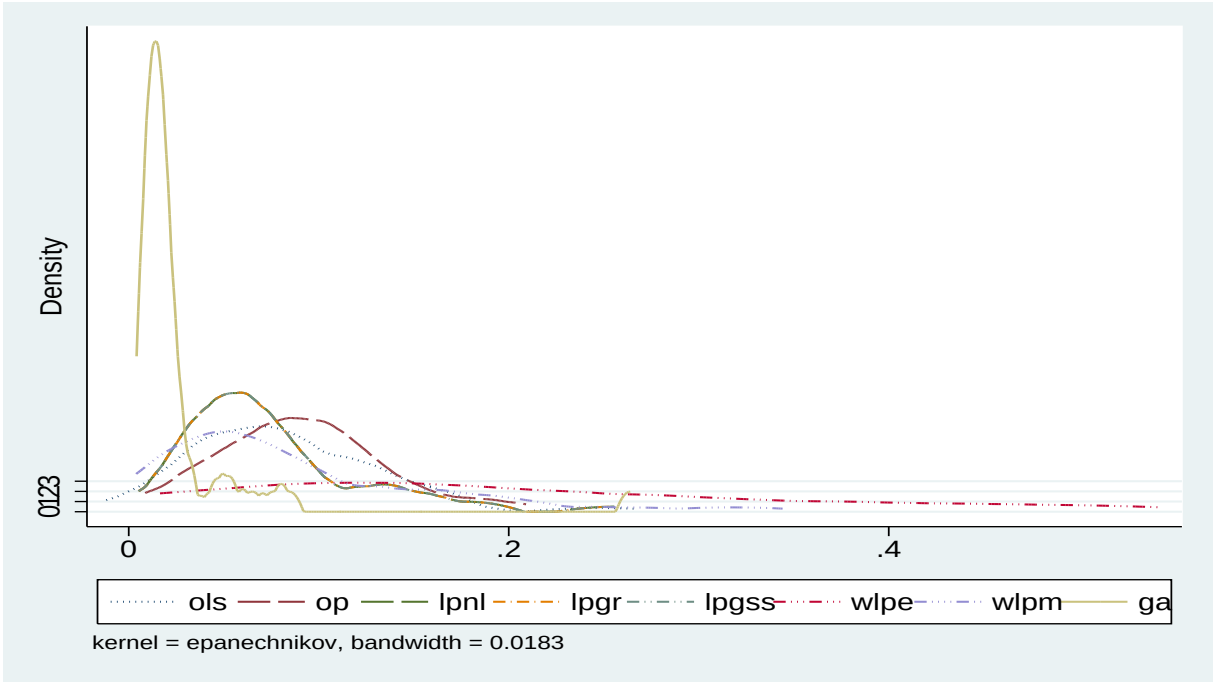
	2002	2003	2004	2005	2006	2007	2008	2009	2010	Total
ASM/CM	68.8	45.8	42.6	43.8	49.9	69.2	53.2	52.4	53.9	59.1
ASM/CM/LBD	68	45.1	41.5	42.8	49.1	68.1	52.7	52	53.7	58.3
ASM/LBD	48.1					52.8				48.9

TFP is considered imputed if at least one of its components is imputed by the Census Bureau. Components: total value of shipments (TVS - changes in inventories are not considered here); total hours, calculated as a product of production worker hours (PH) and the ratio of salaries and wages (SW) and production worker wages (WW); cost of materials (CM), calculated as a sum of the cost of parts (CP), resales (CR) and contract work (CW), but only CM and CP are included in the table. We excluded capital from the analysis because plants' timeseries on capital are created using the perpetual inventory method.
*ASM/CM: ASM and CM combined. **ASM/CM/LBD: observations in ASM and CM for which we observe employment in the LBD. ***ASM/LBD: observations in ASM for which observe employment in the LBD.

8 Figures in Appendix



(a) $\hat{\beta}_M$, similar general shape but differences exist



(b) $\hat{\beta}_E$, more similar except GA and WLPE

Figure A1: Between-industry distributions of elasticities under TFP estimator variants, 50 largest 4-D industries.



Figure A2: Point estimates from a logistic regression of TFP being non-imputed on industry fixed effects, employment-, payroll-size and age classes. ASM cases only. Size classes are based on data from the LBD.