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

July 2, 2015

Abstract

Does the method of estimating plant-level productivity matter? We attempt to answer this question in the context of key stylized facts and popular estimation methods. Using plant-level manufacturing data for the U.S., we test the robustness of results on five dimensions. First, we find non-trivial differences in estimated factor elasticities, especially for capital, across commonly used methods. These differences yield considerable variation in estimated returns to scale across 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. In addition, 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, no matter how we estimate productivity. However, outliers in factor elasticities that arise more frequently from some methods non-trivially impact the quantitative marginal effects of productivity on growth and survival. Fourth, all our productivity variants confirm the main conclusions from the structural productivity decomposition literature: reallocation is productivity enhancing, and variation in within-plant productivity seems more important in terms of cyclical fluctuations of aggregate productivity 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 look at the robustness of productivity dispersion and growth and survival results to imputation and the assumption that elasticities are homogenous within industries. Dispersion is negatively influenced by imputation and the homogeneity assumption. Growth and survival results are also affected but the effect of these factors is more in line with the variation we found in previous exercises.

---

*We would like to thank conference participants at the 2013 Comparative Analysis of Enterprise Data in Atlanta and the 2014 Research Data Center Annual Conference 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. Explanations of possible causes include curvature in the profit function that prevents the most productive firm from taking over an industry, frictions in adjustment of factors and the entry and exit of plants, and 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 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 business cycle and across countries and in turn how these 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 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 accounting for 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 common 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. Another commonly used approach by researchers is to use econometric estimation methods. Since OLS is problematic given endogeneity of factor inputs, 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 by 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 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 reflected 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. However, 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). However, there is variation across methods 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 has not been settled. In his recent survey, Syverson (2011) discusses this debate and observes that many papers in the literature explore the robustness of their findings to alternative estimation methods. Our reading, like his, is that many papers report that results are reasonably robust to alternative estimation methods. But most papers focus on a specific question often for a narrow set of industries. It is less clear to us how robust core findings in the literature are to these issues. The literature offers some 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 input prices, technologies or measurement

[2]Van Biesebroeck (2007) explores the behavior of a number of estimators in the presence of measurement error and parameter heterogeneity. Martin (2008) investigates a case when variable inputs are measured with error.

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

A source of caution in interpreting our findings stems from common data limitations in this literature. Our sample has no information about the quantities used in the production process. We follow standard practice using industry-specific deflators to calculate the constant-dollar values of revenues and input expenditures. This approach implies that we are focusing on what have become known as revenue productivity measures in the literature. The interpretation of revenue productivity measures is sensitive to the assumptions about the environment in which establishments operate. The simplest but perhaps least likely case is for output to be homogeneous within narrow sectors and plants to be price takers in output and input markets so there is no output price heterogeneity. In this case, revenue productivity differences across plants within industries can be interpreted as differences in technical efficiency. If plants are price takers but there is output and input price heterogeneity reflecting exogenous quality differences in outputs and inputs, then revenue productivity can be interpreted as differences in technical efficiency adjusted for quality differences. But increasingly the evidence suggests that output price heterogeneity within narrow sectors reflects both exogenous and endogenous idiosyncratic demand side factors. Exploring the role of the latter has become an active area of research in recent years (see, e.g., Foster, Haltiwanger, and Syverson (2008) and De Loecker (2011)). This implies our revenue productivity measures reflect both technical efficiency and demand side factors. In addition, the estimated factor elasticities should, as appropriate, be considered to reflect both the technology and the demand structure. Exactly how these demand side factors impact the interpretation of exercises resulting from a specific estimation method will depend on the assumed demand structure including functional forms. Much of the literature assumes some form of isoelastic demand structure which is readily tractable but involves very strong assumptions. In what follows, we provide some limited discussion of the likely impact of these issues on the intepretration of the results.

The paper is organized as follows. We discuss our methodology and data in Sections 2 and 3. Section 4 describes the effect of estimation method choice on the distribution of elasticity estimates. Section 5 describes the implications of the differences in elasticity distributions on TFP dispersion, plant growth and survival, and aggregate productivity growth decompositions. Section 6 discusses the robustness of our dispersion and growth results to concerns about the effect of imputated data and the restrictiveness of the homogeneity assumption. Finally, section 7 concludes.

2 Methodology

2.1 Definition of TFP

As in 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 the 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 Estimation of elasticities

Elasticities are unobserved and therefore have to 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 methods 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, Haltiwanger, and Krizan (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	GR
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. In the case of endogenous plant-level prices, additional biases result from the error term including those prices and being correlated with the factor inputs (Klette and Griliches (1996)).

In addition to OLS, we look at the growth accounting method (GA) based on the seminal work of Solow (1957). GA is a frequently used non-statistical method, which uses more explicit assumptions about the environment in which the plant operates. We use the version of GA that 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 shares of input factors in the plant's total cost. The main advantages of this procedure are: it allows for plant-level heterogeneity in elasticities, it is easy to implement and is flexible about the exact shape of production technology, and it is accurate if the data are not subject to much measurement error[3]. An important practical advantage is that GA is free of statistical problems related to the sensitivity of estimates to sample size. Using the cost share of total costs rather than of total value has the advantage that we do not require the assumption of perfect competition. This implies that another advantage of the GA based factor elasticities using cost shares of total costs is that they are robust to alternative demand structures. A potential caveat is that elasticity estimates may be biased if the first order conditions are violated. Since there are many frictions, it is likely the first order conditions do not hold at all points in time at the plant level. Therefore, common factor elasticities across plants in the same industry and/or common factor elasticities over time within the same industry are frequently imposed (see Syverson (2011)). We discuss the implications of such homogeneity assumptions further below. But we also note that most of the alternative estimation methods assume common factor elasticties over time within the same industry.

The remainder of the methods we examine belong to a class of methods often 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. Then the variation in investment can be used back out plant-level 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. Under the assumption that investment is a monotonic and increasing function of productivity and that productivity is the only unobserved state variable, the first step provides consistent OLS estimates of variable input elasticities because the proxy controls for plant-level TFP shocks during the estimation procedure. However, the coefficient of capital is not identified in this step because TFP shocks are controlled for by in-

[3]See Van Biesebroeck (2007).

cluding 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 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, OP works only if we observe entrants and exiters.[8] Therefore, OP cannot be used in industries without data on entrants and exiters. As we will see in section 4, 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, they are likely to respond more to productivity. This is especially relevant in the presence of non-convexities in capital adjustment. LP also highlight that firms almost always report positive use of these variables in their data implying truncation

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

[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]This exploits the assumption about plant-level TFP's Markovian property.

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

[8]Recall, entry and exit was important in the telecommunications industry and therefore have a central role in the OP model designed to analyze that industry.

due to zero proxy values is less severe.[9]

As mentioned above, OP's effort to control for selection was motivated by the fact that plants' entry and exit decisions depend on productivity in their model. By contrast, LP do not focus on selection issues because their panel is unbalanced and is representative of the Chilean Manufacturing sector. 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 survey data we use. If size and productivity are correlated, OP's arguments about selection[10] 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, Caves, and Frazer (2006) (ACF, not included in our analysis). 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 approach 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.[11]

For all the proxy methods, if estimating the revenue function in the presence of endogenous plant-level prices, the estimated coefficients are not factor elastiticities of the production function but rather of the revenue function. Extracting the production function elasticities requires additional structure as in De Loecker (2011). Using the assumptions of the latter paper implies that the revenue function elasticities are under-estimates of the the production function elasticities. Likewise, estimated returns to scale of the revenue function under-estimate the returns to scale of the production function. As we compare estimates across methods, these issues should at least be kept in the background.

3 Data

3.1 Source data

Our industry-level data, including deflators, capital rental prices and depreciation rates, are taken from the NBER-CES Manufacturing database[12], the Bureau of Labor Statistics and the

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

[10] See section 4 for more details.

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

[12] The NBER-CES Manufacturing Industry database is available at http://www.nber.org/nberces. An earlier version is documented in Bartelsman and Gray (1996).

Bureau of Economic Analysis. We use establishment-level information from the Annual Survey of Manufactures (ASM), Census of Manufactures (CM) and the Longitudinal Business Database (LBD).

The CM collects data every five years in years ending in '2' and '7' for roughly 180,000 - 240,000 manufacturing plants. Establishments with less the five employees are not sent forms. Payroll and employment data for these very small plants are imputed using administrative records.[13] 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, Grim, and Haltiwanger (2014) (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 consists of deflated electricity and fuel costs. We create establishment-level capital stock measures using a version of the Perpetual Inventory Method, which calculates current capital as a sum of the 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, the LBD acts as a universe file; we use employment and establishment age data from the LBD to construct inverse propensity score weights that control for non-randomness in our sample.[14]

3.2 Analysis samples

The analyses in this paper exploit three different samples. For our analyses of distributions of elasticities, TFP dispersion and plant growth and survival, we look simultaneously at two samples, which we refer to as the 10 and 50 largest industry samples. These samples need to 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

[13] We drop administrative records cases.

[14] 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.

all reviewed methods. Second, industries should be defined narrowly enough so that we can plausibly assume elasticities are constant among establishments.

Changes in industry classification systems over time make defining these samples more complicated than simply choosing the 10 and 50 industries with the largest number of plant-year observations.[15] Since we estimate elasticities on an industry-by-industry basis, changes in the classification system entail spurious breaks in plants' time series and a drop in sample size. In the first part of the analysis we address these issues by selecting multiple sets of 4-digit SIC industries which were not affected by changes or which were mapped one-to-one into another industry.[16] There are 292 such industries of which we selected the first 10 and 50 industries with the largest number of observations.

We create a third dataset to test whether the implications of the decomposition of aggregate productivity growth described in Petrin, White, and Reiter (2011) (PWR hereafter) are sensitive to the way TFP is estimated. Since we are attempting in part to replicate the results in PWR, we create a roughly comparable dataset. 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 only observed prior to 1987. If a plant is not assigned an industry code in the previous step, we apply a random assignment procedure based on the share of shipments mapped from the 1972 to 1987 SIC industry code.[17]

4 Distributions of elasticities

In this section, we attempt to gauge the effect of the estimation method on the basic characteristics of elasticity distributions.[18] We start by discussing differences in the distribution of capital elasticities. Next, we check whether the ranking of industries by the ratio of the estimated capital elasticity to the estimated labor elasticity varies with estimator choice. We then examine the relationship between sample size and the plausibility of elasticities. We conclude the section by looking at the implications of estimator choice on returns-to-scale.

Figure 1(a) plots the densities of the distribution of capital elasticities in the 50 most populous industries and table 2 shows the corresponding descriptive statistics. There are non-trivial differences in the mean, dispersion and general shape of the distributions. Most notably, LPVA yields significantly larger elasticities. This is because, by construction, the elasticity of value added with respect to capital and labor is inherently larger than the elasticity of gross

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

[16] The NBER-CES 1972 SIC to 1987 SIC and 1987 SIC to 1997 NAICS concordances list list how much of industry i's total value of shipments should be mapped into industry j.

[17] See appendix 8.2 for more information on our random assignment procedure.

[18] Since population distributions of elasticities are unknown, the differences among empirical distributions alone do not tell much about the magnitude of any bias. However, in light of what we know about the way these methods address the bias in OLS estimates, the differences may give us clues to whether or not they correct it in the right direction.

output.[19] 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.

Note that growth accounting-based elasticities are calculated using the cost shares of input factors. If the economic assumptions underlying the method are satisfied, the resulting coefficients are valid measures of the elasticity of output with respect to inputs. There is no need to estimate the coefficients in the econometric sense. Therefore, the elasticities under growth accounting are free of the biases that may be present in elasticites estimated using statistical methods.[20] This fact implies that, ceteris paribus, we should expect to see differences between the elasticity distributions under growth accounting and other methods.

The direction of the bias in OLS-based β_k-estimates is determined by several factors. First, since input demand functions are increasing in productivity we may expect OLS estimates to be biased upward.[21] 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 distributions in figure 1(a) suggest that only LPNL and LPGR are more likely to yield lower $\hat{\beta}_k$ than OLS. Does the relative position of these distributions indicate the expected bias-correction? The answer is no because the direction of the bias depends on additional factors. For one, there may be an offsetting price effect bias. In addition, LP show that positive correlation between capital and labor may cause these estimates to be biased downward.[22] This implies that the $\hat{\beta}_k$ distribution may emerge to the left of OLS not because LP corrects a positive bias but because it includes a downward bias. A further complicating factor is selection. OP argue that since plants' profit and value functions are increasing in capital, larger establishments anticipate larger future returns and therefore can operate at lower current productivity levels, which also entails a negative bias in OLS. If OP corrects for such selection-induced negative bias, and this effect is important in our data, then the OP-based $\hat{\beta}_k$ distribution should be to right of OLS.

[19] See section 3 in Petrin and Levinsohn (2012) for more details.

[20] Other types of specification error may be present, see section 2 for details.

[21] But as discussed above there is a potential offsetting negative bias from neglecting endogeneous plant-level prices.

[22] LP illustrate the effect of the covariance between inputs 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 is more difficult to determine with more variable inputs.

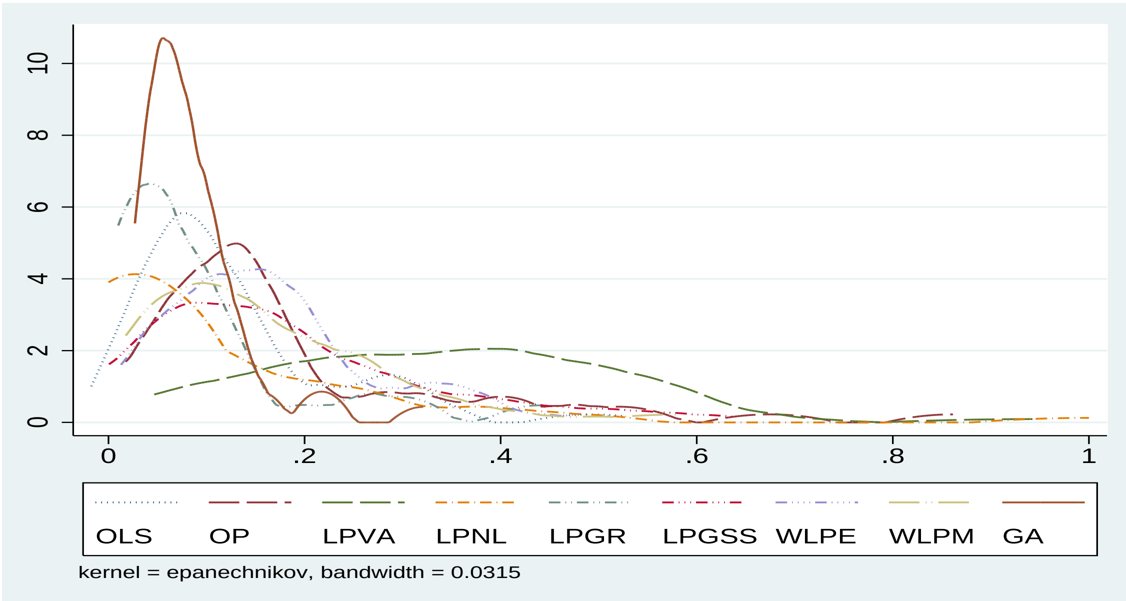

(a) $\hat{\beta}_K$

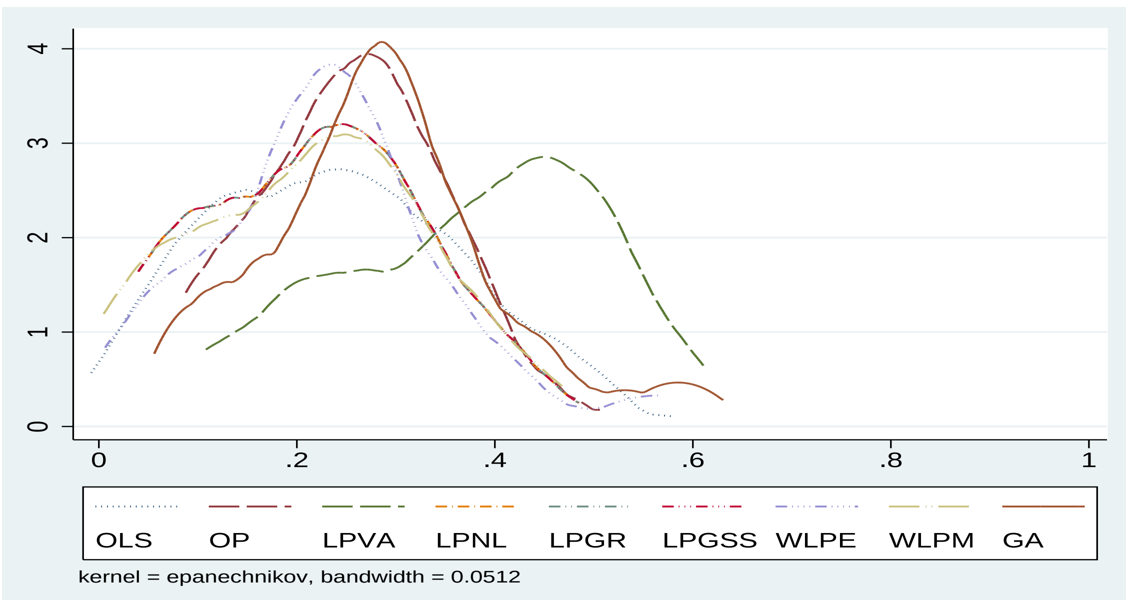

(b) $\hat{\beta}_L$

Figure 1: Between-industry distributions of capital and labor elasticites. TFP estimators described in table 1, sample is 50 largest industries.

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

	N	Mean	Median	IQR
50 largest industries				
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 industries				
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 method definitions. 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.

In practice, OP tends to result in higher $\hat{\beta}_k$ than OLS suggesting that controlling for selection-induced bias may be important. As for other proxy methods, WLPE and WLPM are more likely to yield extreme $\hat{\beta}_k$ even though the typical elasticities under these methods are similar to those under OP. This finding indicates the choice of the estimation algorithm - non-parametric two-step estimator versus efficient GMM - also affects the level of $\hat{\beta}_k$. These results tell us that not only the choice of the proxy and addressing selection have important consequences but the choice of the estimation algorithm also affect the level of capital elasticities. 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). Finally, the elasticities of variable inputs show stronger clustering, especially $\hat{\beta}_l$ (figure 1(b)). This is partly 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 across estimation methods than capital elasticity distributions.

Next, we explore whether the choice of estimator affects the ranking of industries by $\hat{\beta}_k/\hat{\beta}_l$.[23]

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

If the estimation method is not important then industries' rank should not vary with the estimation method. Our results suggest the opposite. Comparing the rankings across estimation methods we find there is a positive probability that different estimators imply different industry rankings (see table A2 for details).[24]

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.[25]

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-estimate the above elasticity distributions using all 459 industries and count the cases with positive, zero, and negative elasticities for every method.[26] All methods result in positive $\hat{\beta}_l$ and $\hat{\beta}_m$ in most industries (table A4, columns 6 and 9), 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.[27] OP's algorithm stops in 18% of the industries (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). 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 at least partially '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. However, we find two distinct cases where erratic estimates are unlikely to be related to sample size. First, the zero-$\hat{\beta}_k$ group for LPNL is of very similar size as the positive group (second entry in column 2 of table A5). It is 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 for WLPE where average sample size is similar

[24]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.

[25]Table A3 sorts industries by sample size as measured by the number of plant-year observations. The first entry in column 3 says the largest industry has about 39,000 plant-year observations. The sample size drops by more than 50% in the 10th, and by 80% in the 50th industry.

[26]The sample for this exercise is described in more detail at the end of section 3.2.

[27]This is because LPGR searches over a pre-defined grid between .01 and .99.

to that in the positive group (column 7 of table A5). These results 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 industries with negative elasticity estimates since such estimates are implausible but that would raise issues of selection bias. An alternative approach we explore below is to make stronger assumptions of homogeneity - for example assuming plants within 3-digit SIC industries share the same factor elasticities. We show later that this approach leads to more plausible elasticity estimates and implied firm dynamics in many cases.

In the last exercise of this section we look at the implications of estimation method choice for returns-to-scale (RTS). For reasons outlined in section 1, we measure RTS as the sum of estimated revenue elasticities. It is a useful metric because it concisely captures the various effects estimation methods may have on elasticities.[28] For purposes of exposition, we choose the interval $[.7, 1.3]$ to illustrate which methods yield more observations in this range. This range is arbitrary but is consistent with profit shares between -25% and $+25\%$, which does not seem too restrictive.[29] Our results can be summarized as follows (see table A6 for more details). First, all methods except for LPVA yield more than 75 percent of the industries with RTS in the specified range. There is, however, considerable variation across methods with only the LPGR method yielding more than 90 percent of the industries within this range. Second, the proxy methods are more likely to differ from 1 on the low side rather than on the high side. This follows from the entries in columns 5-7 of Panel B being greater than those in Panel A, and more so for LPVA). Estimated RTS more likely being below 1 using the proxy methods is consistent with the predicted relationship between revenue and production function elasticities.

5 Implications of the differences in elasticity distributions

In the last section, we showed the choice of estimation method matters for answering questions that rely heavily on elasticity estimates. Here we investigate whether or not such differences in elasticities matter for core findings in the productivity literature. Specifically, we examine: TFP dispersion; the relationship between productivity, growth and survival; and structural decompositions of aggregate productivity growth (APG).

5.1 TFP dispersion

Does the choice of 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

[28]Note that growth accounting-based elasticities are calculated assumping constant returns to scale, and therefore $RTS = 1$ by construction. For this reason, GA-based results are not shown in these calculations.

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

in table A7). Ignoring LPVA because it is intrinsically different from the other methods, the average IQR of the log-TFP distribution varies between .24 (GA) and .4 (WLPM) across methods. These values mean 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. TFP dispersion varies also with industries.[30] Taking GA-based TFP as an example, productivity dispersion in the industry one standard deviation above (below) the sample mean is .35 (.13) indicating that the effect of industry differences are at least as important as estimation methods. Additionally, table 4 shows the productivity rank of plants within industries is impacted by estimation methods. Both the Pearson and Spearman rank correlations are substantially smaller than 1 implying that where plants sit in the productivity distribution is sensitive to choice of estimation method. The WLPM method yields especially low correlations with other methods.

Table 3: Descriptive statistics of TFP distributions. TFP estimators described in table 1, sample is 50 largest industries.

	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 method descriptions. 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 (table A7), results barely change.

[30] 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.

Table 4: Correlations among within-industry TFP distributions, sample is 50 largest 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).

So what determines within-industry TFP 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 within-industry TFP dispersion depends on elasticities (β_j), output and factor variances (σ_q, σ_{x_j}) and the covariances between them ($\sigma_{qx_j}, \sigma_{x_jx_i}$).[31] In particular, the relationship between β_j and TFP dispersion depends on σ_{qx_j}, $\sigma^2_{x_j}$, $\sigma_{x_jx_i}$, and the β_j.[32] However, analytical expressions become difficult to evaluate if one intends to assess the relative importance of elasticities and input-output characteristics for two reasons. First, β_j and the previously mentioned covariances are endogenous because those covariances are exploited 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). The results suggest that measured productivity dispersion is increasing in both $\hat{\beta}_k$ and σ_q. Specifically, TFP dispersion approximately doubles after a one-standard-deviation increase in

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

[32] The partial derivatives of σ^2_ω with respect to σ_q and β_j are given by $\frac{\partial \sigma^2_\omega}{\partial \sigma_q} = 2\sigma_q$ and $\frac{\partial \sigma^2_\omega}{\partial \beta_j} = 2\beta_j\sigma^2_{x_j}$ - $2\sigma_{qx_j} + 2\sum_{i\neq j} \beta_i \sigma_{x_jx_i}$. $\frac{\partial \sigma^2_\omega}{\partial \sigma_q} = 2\sigma_q$. 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: $\frac{\partial \sigma^2_\omega}{\partial \beta_j} \lesseqqgtr 0$ iff. $\beta_j \lesseqqgtr \frac{\sigma_{qx_j} - \sum_{i\neq j} \beta_i \sigma_{x_jx_i}}{\sigma^2_{x_j}}$, which shows the relationship depends on σ_{qx_j}, $\sigma^2_{x_j}$, the $\beta_i \sigma_{x_jx_i}$ terms and β_j itself.

$\hat{\beta}_k$. A similar change can also be achieved by increasing σ_q. However, the required increase in σ_q is relatively large[33], which indicates that typical changes in $\hat{\beta}_k$ may have larger effects on TFP dispersion than typical changes in σ_q.

In closing this section, we note that the available evidence (see, e.g., Foster, Haltiwanger, and Syverson (2008)) is that revenue productivity dispersion is lower than physical productivity dispersion. This reflects the inverse correlation between productivity and prices. But Foster, Haltiwanger, and Syverson (2008) also highlight that demand shocks exhibit high dispersion relative to physical productivity dispersion. As such, dispersion in revenue productivity likely reflects both dispersion in physical productivity and in demand shocks but this is tempered by the inverse correlation between prices and productivity. Many other factors are potentially important sources of revenue productivity dispersion over and above physical productivity dispersion. Hsieh and Klenow (2009) highlight the role of distortions in generating dispersion in revenue productivity. Others highlight the role of frictions such as overhead factor costs (see, e.g., Bartelsman, Haltiwanger, and Scarpetta (2013)) and adjustment frictions (see, e.g., Asker, Collard-Wexler, and De Loecker (2014)).

5.2 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 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.[34] Empirical work on the connection between growth and productivity also relies on the results of these models. A study by Foster, Grim,

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

[34]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, Benkard, and Roy (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. Although other literatures offer important conclusions about why distributional characteristics matter for aggregate dynamics (see, for example, Caballero, Engel, and Haltiwanger (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.

and Haltiwanger (2014) (FGH hereafter) is a recent example of the research in this context.[35] A comprehensive 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 whether the most commonly found patterns hold across TFP estimators. 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 build upon the existing literature concerning the properties of productivity dynamics and test the robustness of these predictions using simple regression models linking a set of outcomes to productivity and plant-level controls. 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.[36]

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 and i, t and s index plants, time periods and states, respectively.[37] It is important to note that this specification relates growth and survival outcomes from t to $t+1$ based on productivity in period t.[38]

Table 5 shows $\hat{\gamma}_4$ from equation (2) using our 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 shows 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 analogous. Point estimates suggest there are non-trivial differences in the measured effect of TFP. For example, the first entry in column 2 shows

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

[36]See, for example, Baily, Hulten, and Campbell (1992), Foster, Haltiwanger, and Krizan (2001), Foster, Haltiwanger, and Krizan (2006).

[37]This specification differs from the one in FGH in that age effects, the Great Recession dummy and its interactions are omitted for simplicity.

[38]We follow FGH by using the integrated LBD with the ASM data for this analysis. The ASM data provides the distribution of plant-level productivity in any given year and the LBD provides the growth and survival outcomes for the full set of plants in the ASM in that year between t and t+1.

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, ignoring for now WLPE and WLPM. Despite the non-negligible differences, our 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, sample is 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									
overall growth	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)
conditional growth	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 sensitive to such changes. For instance, the absolute values of OP-based coefficients of all three regressions increase. This suggests pooling data may be beneficial because if there is no statistical association between two variables, their partial correlation coefficient would tend towards zero. Estimates for other methods change somewhat but most of them remain in a comparable range.

WLPM yields the counter-intuitive results that more productive plants are less likely to grow and more likely to exit. Increasing sample size by pooling data does not reverse this result. However, if we drop industries with negative elasticities (see the bottom panel), the signs appear to be more in line with other results. In addition, the absolute values of point estimates also increase and/or standard errors decrease, despite the decrease in sample size.[39]

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 heterogeneity and pooling data amounts to replacing less informative TFP observations with more 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 note 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.[40]

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 (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 across methods.

Finally, we return to the consequences of using revenue productivity as opposed to physical productivity in these exercises. Foster, Haltiwanger, and Syverson (2008) find that the marginal response of exit to revenue productivity is actually larger than that for physical productivity. They show that this is because revenue productivity includes both the effects of physical productivity and demand effects. We anticipate that similar remarks likely hold in this context.

5.3 Structural decompositions of aggregate productivity growth

In this section, we examine whether choice of method affects 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 explore how APG is tied to the behavior of microeconomic agents; some examples are Olley and Pakes (1996), Baily, Bartelsman, and Haltiwanger (2001), Petrin, White, and Reiter (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

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

[40]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.

idea that APG sources either from the productivity gains *within* plants or the more efficient allocation of factor inputs *between* plants.

We use a recent decomposition proposed by PWR primarily 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 is defined as the growth in final demand in excess of capital and labor growth.[41]

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[42]) but our value added and capital growth rates are shown to be somewhat smaller. These differences are due partly to measurement differences and partly to an issue with PWR's source data for implicit deflators.[43] We note our value added growth growth measure is closer to the PWR measure if we ignore the first years of ASM panels (shown in the lower panel of table 6). Calculating aggregate growth in these years is problematic because we do not observe growth rates for smaller establishments which were not included in the previous ASM panel. As a consequence, average growth rates in the first year of an ASM panel will reflect the growth rate of the largest plants only.

Table 6, lines 5-11 describe the decomposition of APG into reallocation and within-plant growth. The contribution of fixed costs is obtained as a residual. These components depend on the estimation method and we see nontrivial differences. The implications of the such differences are substantial. Considering the entire sample (upper panel of table 6), our estimate of the

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

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

[43]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.

Table 6: Time-averages (1977-1996) of weighted contributions to aggregate productivity growth, in percentage points. The sample contains only ASM continuers.

	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
$\bar{N}$** (1000)	4.1	3.5	2.3	2.3	2.3	1.9	1.9	1.5	1.5	4.1	n.a.
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

*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.

annual contribution of reallocation is between .7-1.5 percentage points, which amounts to 30-70% of the average annual aggregate productivity growth. The contribution by within-plant productivity growth is at most 1.4 percentage points, about 64%. Note these ranges do not include WLPM (column 7), which we will discuss later in more detail. Method-related variation is further illustrated by the results which ignore the first years of ASM panels (lower panel). For example, OLS, OP, LP variants and GA result in .5-.9 percentage point average annual contributions by reallocation (columns 1-5 and 10) implying that reallocation accounted for 33-64% of aggregate productivity growth between 1977-1996.

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 $(t-2)$.[44] Using these instruments implies reallocation

[44]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' time series.

and within-plant contributions 3-4 times larger (4.9% and -3%) compared to other methods. Increasing the size of the instrument matrix by including earlier lags (1-3), as in in PWR makes the extreme contributions disappear (column 9). With both WLPE and WLPM, the measured contributions by reallocation are 1.3 percentage point (columns 8-9) indicating that reallocation explains almost 70% of average annual productivity growth. 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' time series. 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.[45] As a final point, we note the time series standard deviation of the total reallocation contribution is small relative to the within-plant productivity contribution reinforcing previous findings (last rows of the panels in table 6). This indicates time series variation in within-plant growth is more important for aggregate (cyclical) fluctuations than reallocation.

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

We refrain from speculating on the impact of using revenue functions to estimate factor elastities in this context because we think the issues are not yet well understood. In applying this decomposition empirically, we follow PWR and Petrin and Levinsohn (2012) by estimating revenue factor elasticities when using proxy methods. However, as discussed above, revenue factor elastities will not in general be equal to production factor elastiticies. PWR and Petrin and Levinsohn (2012) argue that some of the factors that might prevent marginal products from being equal to marginal costs at all moments in time are markups and adjustment frictions. It is such wedges between marginal products and marginal costs that underlie the reallocation terms in their structural decomposition. But if it is markups driving the wedges it is not clear that the reallocation terms in the decomposition have been correctly estimated. Further investigation of these issues seems warranted.

[45]Table A15 lists the time series 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.

6 Robustness

We next examine the robustness of our results about TFP dispersion and growth and survival to concerns about imputation methods used in the underlying microdata and our assumption that elasticities are homogeneous within industries. Our approach in this section is different from earlier sections. Here we compare the same TFP measure across completed and non-imputed samples rather than different TFP measures in the same sample.

6.1 Imputation, dispersion and growth

U.S. Manufacturing data collected by the Census Bureau is subject to item nonresponse where respondents answer some questions but not others. Such missing values are imputed by the Census Bureau. Recent research found that certain imputation methods can impact analyses that use such completed data because imputation is non-random (see White, Reiter, and Petrin (2012), WRP hereafter). Some of the imputation methods employed 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 (see WRP for 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 approach is to drop imputed observations, but this results in selection bias if the probability of imputation instances is correlated with establishment characteristicts. We take this approach below, correcting for selection bias by using inverse propensity score weights. Another option is to use multiple imputation methods (e.g., the classification and regression tree method (CART) used by WRP) to improve on the methods that have been used to impute the plant-level data. 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.

To account for nonrandom imputation, we 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 using only non-imputed observations. 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 8.3.

We carry out two exercises. First, we compare weighted and unweighted statistics of within-industry dispersion which are based either on completed data or only non-imputed observations. Second, we 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. Figure 2 summarizes our results.[46] Solid and dotted lines denote statistics which are calculated using either completed or non-imputed data. Weighted statistics are denoted by marks on the corresponding graphs. In the completed sample, weighted statistics are based on weights ($ipw1$) which account for the fact that selection into the ASM and CM is non-random.[47] In the non-imputed sample, the weights are a composite of $ipw1$ and a weight that is inversely proportional to the probability that the plant's TFP is calculated using non-imputed data ($ipw2$). The graphs on the right-hand-side panels of the figures are based on ASM observations only.

We find the standard deviation is not very very sensitive to imputation, while the quantile-based dispersion measures are affected. The interquantile-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 lack of dips in the solid lines in 2002 and 2007 in the right-hand-side charts.[48] If we weight non-imputed observations by $ipw1 \times ipw2$, measured dispersion is higher (marked and unmarked dotted lines in figures 2(b)-2(c)). This is an important finding because results using the CART multiple imputation method in WRP suggest dispersion measures based on imputed data tend to be smaller than those that take imputation into account.[49] If imputation causes a downward bias in dispersion, our weighting scheme corrects it in the right direction.

We find 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 5.2.[50] 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 $ipw1 \times ipw2$. In the last set of regressions, we repeat calculations excluding non-ASM cases in Census years. Table 7 summarizes our results. The effect of productivity on growth among all establishments (column 1) and on the probability of exit (column 2) seems less in the non-imputed sample, regardless of weighting. Point estimates of the effect on growth (exit) in the non-imputed sample are about half (third)

[46]The dispersion estimates underlying figure 2 are shown in table A16 in the appendix.

[47]We address this issue by weighting observations with a propensity score weight where the propensity scores are inversely proportional to the probability that a plant is selected into the ASM-CM. More details about this selection issue and the technique we use to address it can be found in FGH.

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

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

[50]Note that as the estimates in section 5.2 were calculated using $ipw1$ weights, all regressions in this exercise - based on either completed or non-imputed observations - are also weighted by $ipw1$.

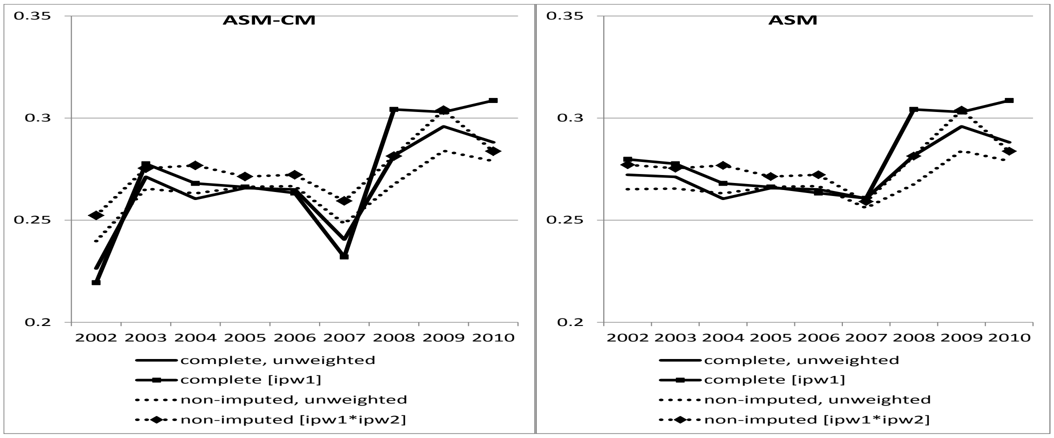

(a) Dispersion measure: standard deviation

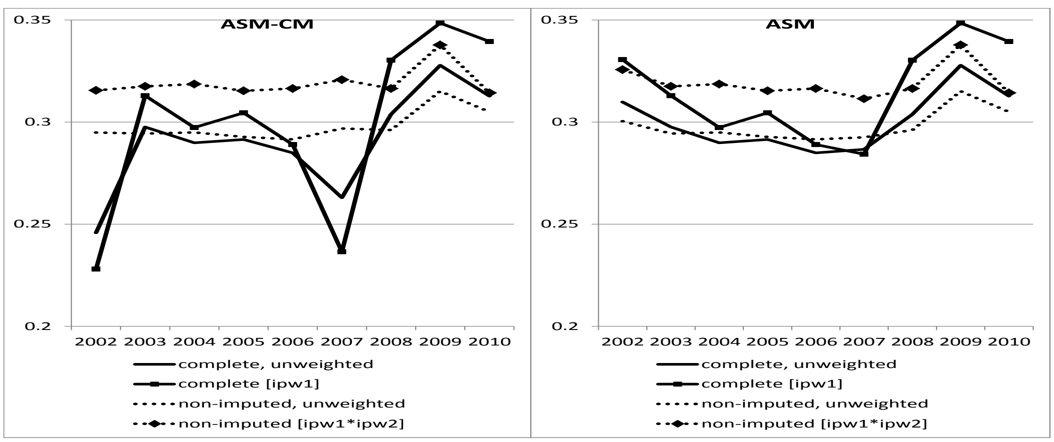

(b) Dispersion measure: inter-quartile range

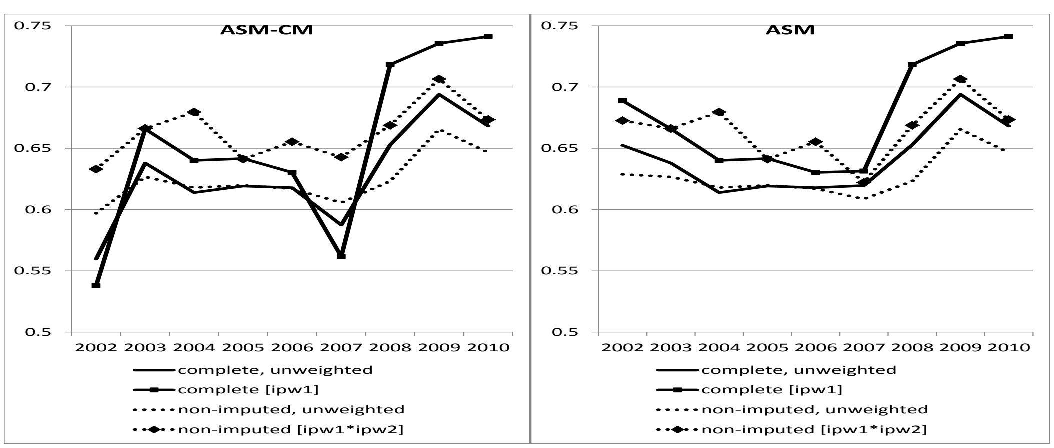

(c) Dispersion measure: 90th to 10th percentile range

Figure 2: Time series of average within-industry dispersion measures (table A16) in various samples, ASM/CM 2002-2010, all industries.

the size of those in the completed data. We do not detect such attenuation among continuers (column 3). Such variation in the coefficients is by no means negligible, particularly for economic analyses where differences in magnitudes may matter a lot. But it is not unfamiliar either. Similarly to the results in section 5.2 our last exercise offers evidence that, at least in terms of sign and order of magnitude, this relationship is robust to imputation issues.

Table 7: The effect of TFP on outcomes between 2002-2010 in the ASM/CM, various samples within the ASM/CM, growth accounting based TFP.

	sample (1)	overall growth (2)	exit (3)	conditional growth (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 parentheses. †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.

6.2 Homogeneity, dispersion and growth

Throughout this paper, we assume elasticities are homogeneous within industries and constant over time. In this section, we look at the effects of the homogeneity assumption on our results on TFP dispersion and plant growth and survival.

Whether or not such homogeneity assumptions are restrictive depends on at least two properties: the underlying within-industry differences in technology and sample size. 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.[51] 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.[52] 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.[53] 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 is typically present in microdata.

We assess the consequences of allowing for plant-level heterogeneity in the elasticities by comparing TFP 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[54] 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 two (see the difference between the thin and thick solid lines in the figure). Our results also indicate that allowing for time series 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 dashed and dotted lines). The overall conclusion is that the effect of time series smoothing is dwarfed by that of cross-section smoothing.

Both theory and earlier research[55] indicate 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 8 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

[51] 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.

[52] We saw in section 5.2 that for some estimation methods even cross-industry pooling may be necessary.

[53] See section 4 for examples.

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

[55] For example, Syverson (2004).

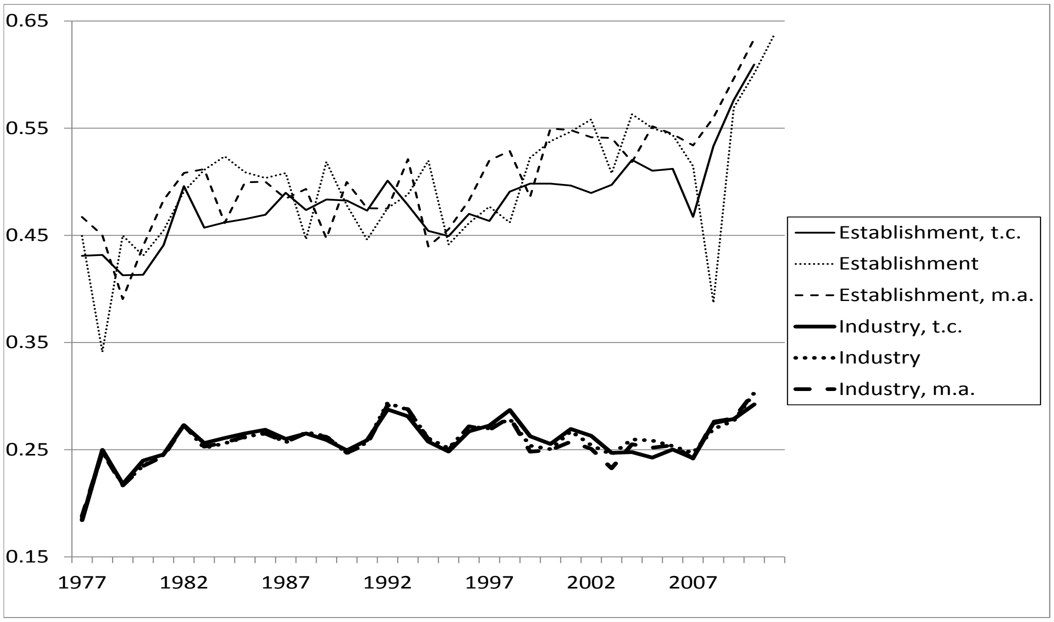

Figure 3: Within-industry interquartile range measures of growth accounting based log-TFP, 10 largest industries. Thin and thick lines denote dispersion measures based on shares calculated at the establishment- and industry-levels, respectively. *t.c.* and *m.a.* label cases where statistics are based 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.

shares leaves the effect of productivity virtually unchanged. In contrast, using plant-level shares (columns 3-5) reduces the magnitude of estimates - although they remain statistically significant. The consequence of using plant-level shares is that the growth effect of productivity drops by almost 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 time series 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 time series variation of establishments.

Our overall conclusion is that calculating cost shares at the plant level is unlikely to be optimal. 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 8: 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)
conditional 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.

7 Concluding remarks

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

The choice of estimation method affects 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 a plant at the 75th percentile generates between 25%-50% percent more revenue with the same amount of inputs than a plant at the 25th percentile. This is an average calculated over industries and time, which suggests 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 that quantitative results are sensitive to the treatment of these outliers.

Does the choice of estimation method have implications for the relative importance of the sources of aggregate productivity growth? We find that although some methods seem to be more sensitive than others and 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, 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 suggests some of the increase in dispersion is due to measurement error, which implies using plant-level shares is unlikely to be optimal.

In sum, it is important to understand when the devil is in the details. Finally, one devil that may remain in the details that we have not directly investigated is the impact of heterogeneous and endogenous plant-level prices. We have commented on the likely impact of endogenous demand side factors throughout but it would be of interest to consider this issue in more depth. We think that exploring the role of endogenous demand side factors in the current context will require comparing and contrasting approaches that include direct measurement of prices and quantities (for the limited number of products with such information) vs. methods that impose strong functional form assumptions (i.e., isoelastic demand structures) to deal with these issues.

References

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

Asker, John, Allan Collard-Wexler, and Jan De Loecker. Dynamic Inputs and Resource (Mis)Allocation. *Journal of Political Economy*, 122(5):1013 – 1063, 2014.

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

Baily, Martin N., 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.

Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta. Cross-Country Differences in Productivity: The Role of Allocation and Selection. *American Economic Review*, 103(1): 305–34, February 2013.

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

Caballero, Ricardo J., 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.

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

De Loecker, Jan. Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity. *Econometrica*, 79(5):1407–1451, 09 2011.

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

Foster, Lucia S., John C. Haltiwanger, and Cornell 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.

Foster, Lucia S., John C. Haltiwanger, and Cornell 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.

Foster, Lucia S., John C. Haltiwanger, and Chad Syverson. Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *American Economic Review*, 98(1): 394–425, 2008.

Foster, Lucia S., Cheryl A. Grim, and John C. Haltiwanger. Reallocation in the great recession: Cleansing or not? Working Paper 20427, National Bureau of Economic Research, August 2014.

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

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

Hsieh, Chang-Tai and Peter J. Klenow. Misallocation and manufacturing TFP in china and india. *Quarterly Journal of Economics*, 124(4):1403–1448, November 2009.

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

Klette, Tor Jakob and Zvi Griliches. The Inconsistency of Common Scale Estimators When Output Prices Are Unobserved and Endogenous. *Journal of Applied Econometrics*, 11(4): 343–61, July-Aug. 1996.

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

Marschak, Jacob and William H. Andrews. Random simultaneous equation and the theory of production. *Econometrica*, 12(3/4):143–205, 09 1944.

Martin, Ralf. Productivity dispersion, competition and productivity measurement. LSE Research Online Documents on Economics 19573, London School of Economics and Political Science, LSE Library, June 2008.

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

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

Petrin, Amil, T. Kirk White, and Jerome Reiter. 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.

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

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

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

Weintraub, Gabriel Y., 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.

White, T. Kirk, 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.

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

8 Appendix

8.1 Remarks on proxy methods

The two-step procedure described in Levinsohn and Petrin (2003) has become widely used among practitioners in the past decade. Its advantages and caveats have been discussed in the literature, see for example Ackerberg, Caves, and Frazer (2006) and 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 'golden section search' algorithm (GSS).[56] 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 guarantee finding the optimum points of unimodal functions they may get stuck at local optima. 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 depending on the scale of estimation.[57]

8.2 Random assignment of industry codes

As described in Section 3.2, we correct for the 1972 to 1987 SIC change in our third analysis sample by following PWR by assigning the observed 1987 SIC code to the 1976-1986 observation for any establishment observed between 1987-1996. However, we deviate from their approach for cases only observed prior to 1987. If a plant is not assigned an industry code in the first step, we apply a random assignment procedure. The basic idea of the random assignment procedure is to choose from among plants such that the share of reassigned plants matches the appropriate share in the concordance. Randomness is necessary to ensure the procedure is not dominated by a few large establishments. As an illustration, suppose the mapping says 10% of industry i's (SIC 1972) total value of shipments should be mapped into industry j (SIC 1987). First, we compute the time-average of each plant's share in the shipments of industry i and then we randomly sort them by these averages. Next, we calculate the cumulative sum of shares and find the first n plants for which the sum does not exceed 10%. These establishments are classified in industry j. 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 shown here, suggest that random assignment approximately replicates the NBER-CES mappings.

[56] The idea of GSS is successively narrowing the range of values inside which the extremum is known to exist. The name of the algorithm sources from the fact that the procedure maintaings the function values for triples of points whose distances form a golden ratio.

[57] Estimating the 459 LPGR β-sets for the APG comparisons of section 5.3 took about 28 days given our computing resources.

8.3 Imputation and plant characteristics

Evidence in White, Reiter, and Petrin (2012) (WRP hereafter) 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. See table A18 for more details.[58] 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 WRP). Yet another approach is to use non-imputed observations only, which 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 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:

$$log \frac{p(X_{it})}{1 - p(X_{it})} = X_{it}\theta_t + \epsilon_{it},$$

which amounts modeling the probability $Pr(I_{it} = 1|X_t) = E[I_{it}|X_{it}] = F(X_{it}\theta_t) + \epsilon_{it}$, where $F(x) = \frac{1}{1+e^{-x}}$ and I_{it} denotes an indicator variable equal to 1 if any of the components of plant-level 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 θ_t denote a vector of controls and coefficients. 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,

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

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. 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.

Imputation rates differ across variables. The upper panel of table A17 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.[59] 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.[60] 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 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[61] 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.[62] In section 6.1, 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.[63] The difference in the AIC is about a factor of 3 and 2 in 2002 and 2007, respectively.

[59] 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).

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

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

[62] 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.

[63] Available upon request.

9 Appendix Tables

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 time series) 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.

1	2	3	4	5
rank	SIC 1987	N	MeanL	Nplants
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: SIC 1987 code; column 3: number of plant-year observations in thousands; column 4: average number of observations per plant; column 5: number of plants in thousands.

Table A4: Industries with positive, negative and non-estimable elasticities, in percent 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 returns to scale measures (RTS) in alternative ranges, by estimation method, as in table A4. 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 RTS in the specified range; the small numbers highlight the differences in elasticity distributions, and quantify the qualitative differences in the elasticity distributions mentioned in section 4.

Table A7: Descriptive statistics of TFP distributions. TFP estimators described in table 1, sample is 50 largest industries, pre-estimation outliers included.

	N (1000)	IQR	SD
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
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

Table A8: Correlations among within-industry tfp-distributions, 50 largest 4-digit 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-digit 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-digit 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	$\bar{\hat{\beta}}_K = 0.17$	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: Time series 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: 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						
	Completed data, unweighted*			Completed 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						
	Completed data, unweighted			Completed 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 *ipw*1 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 *ipw*1**ipw*2 where *ipw*2 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.

Table A17: 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' time series 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.

Table A18: Estimated rates of imputation models within the most frequent impute types, as a percentage of the total number of observations, ASM-CM 2007. The most frequent impute types are based on linear regressions (flagged "B") and/or historical (flagged "H") information.

	tvs	ph	cm
ASM CM			
Overall imputation rate (fact)	28.5	58.8	41.1
Types: B or H flags (fact)	20.6	55.1	33.3
Models (estimate)*			
multivariate regression	4.1	12.4	6.3
univariate regression	15	44.8	26.6
historical	0.9	5.8	0.5
total	20	62.9	33.4
ASM cases only			
Overall imputation rate (fact)	24.1	39	34.3
Types: B or H flags (fact)	18.2	36.1	25.7
Models (estimate)*			
multivariate regression	10.5	14.4	14.4
univariate regression	8.1	13.5	11.6
historical	1	9.7	1.4
total	19.6	37.6	27.4

*Census data indicates the impute type but not the model and there are more than one impute models within a type. For example, within type-B imputes, both univariate and multivariate regression models are used. We estimate impute model instances by evaluating the restriction each impute model implies. These results show regression-based imputes are most common. In the ASM-CM, imputes are typicall based on univariate regressions, while ASM imputes are much more likely to be based multivariate regressions.

10 Appendix Figures

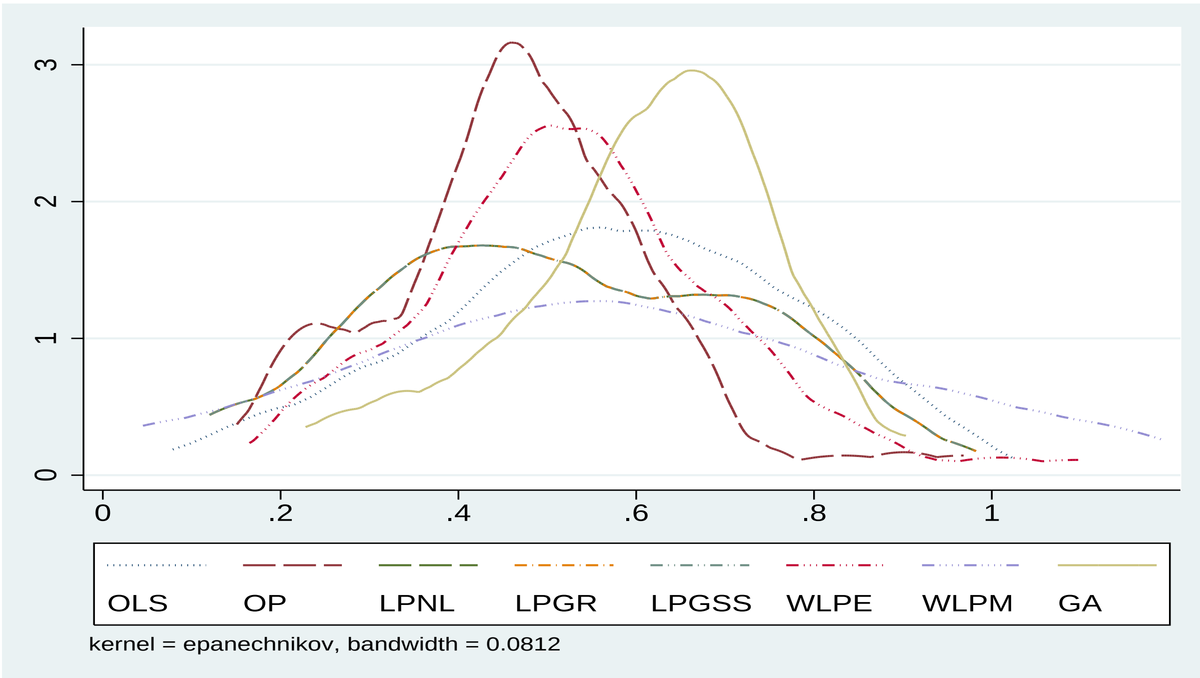

(a) $\hat{\beta}_M$, similar general shape but differences exist

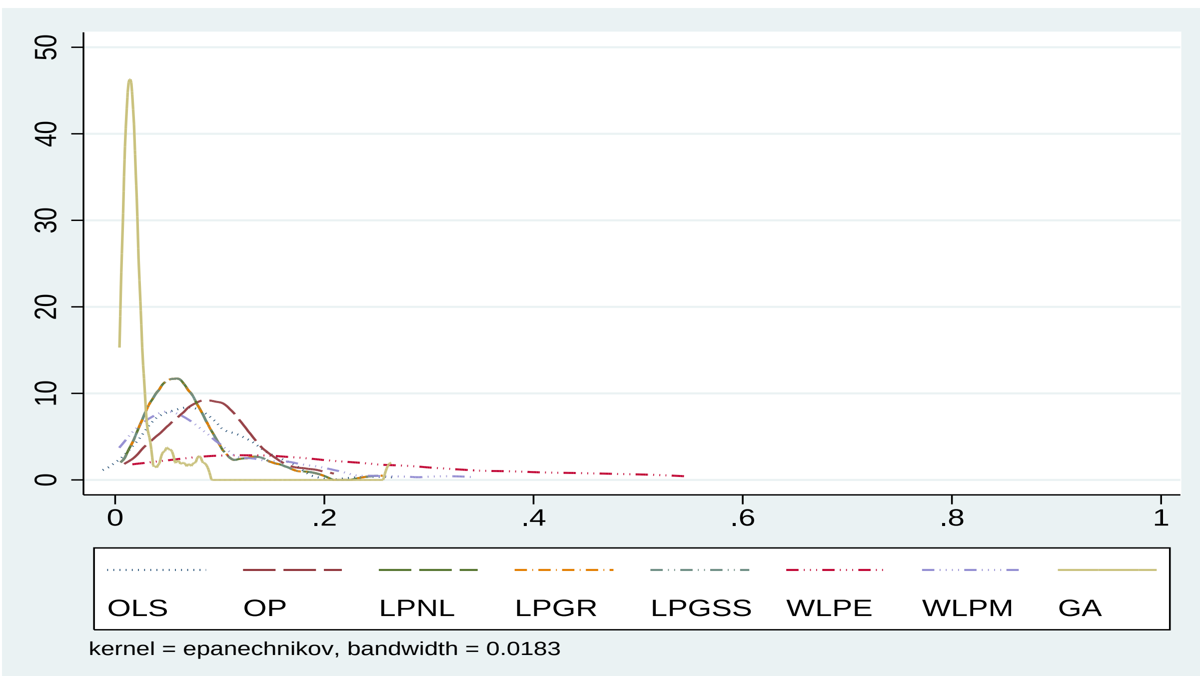

(b) $\hat{\beta}_E$, more similar except GA and WLPE

Figure A1: Between-industry distributions of elasticities under TFP estimator variants, 50 largest 4-digit 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. 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 classified establishments into 9 age classes (births, 1, 2, 3, 4, 5, 6-10, 11-15, 16+ years).