

NBER WORKING PAPER SERIES

MACRO AND MICRO DYNAMICS OF PRODUCTIVITY:
FROM DEVILISH DETAILS TO INSIGHTS

Lucia S. Foster
Cheryl A. Grim
John Haltiwanger
Zoltan Wolf

Working Paper 23666
<http://www.nber.org/papers/w23666>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2017

We thank Jan DeLoecker, Ron Jarmin, Kirk White and conference participants at the 2013 Comparative Analysis of Enterprise Data in Atlanta, 2014 Research Data Center Annual Conference and the 2015 NBER CRIW workshop for valuable comments. We are grateful to Kirk White 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 or the National Bureau of Economic Research. All results have been reviewed to ensure that no confidential information is disclosed.

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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2017 by Lucia S. Foster, Cheryl A. Grim, John Haltiwanger, and Zoltan Wolf. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Macro and Micro Dynamics of Productivity: From Devilish Details to Insights
Lucia S. Foster, Cheryl A. Grim, John Haltiwanger, and Zoltan Wolf
NBER Working Paper No. 23666
August 2017
JEL No. E24,L22,O4

ABSTRACT

Researchers use a variety of methods to estimate total factor productivity (TFP) at the firm level and, while these may seem broadly equivalent, how the resulting measures relate to the TFP concept in theoretical models depends on the assumptions about the environment in which firms operate. Interpreting these measures and drawing insights based upon their characteristics thus must take into account these conceptual differences. Absent data on prices and quantities, most methods yield "revenue productivity" measures. We focus on two broad classes of revenue productivity measures in our examination of the relationship between measured and conceptual TFP (TFPQ). The first measure has been increasingly used as a measure of idiosyncratic distortions and to assess the degree of misallocation. The second measure is, under standard assumptions, a function of fundamentals (e.g., TFPQ). Using plant-level U.S. manufacturing data, we find these alternative measures are (i) highly correlated; (ii) exhibit similar dispersion; and (iii) have similar relationships with growth and survival. These findings raise questions about interpreting the first measure as a measure of idiosyncratic distortions. We also explore the sensitivity of estimates of the contribution of reallocation to aggregate productivity growth to these alternative approaches. We use recently developed structural decompositions of aggregate productivity growth that depend critically on estimates of output versus revenue elasticities. We find alternative approaches all yield a significant contribution of reallocation to productivity growth (although the quantitative contribution varies across approaches).

Lucia S. Foster
Center for Economic Studies
Census Bureau
Room 2K124
Washington, DC 20233-6300
lucia.s.foster@census.gov

John Haltiwanger
Department of Economics
University of Maryland
College Park, MD 20742
and NBER
haltiw@econ.umd.edu

Cheryl A. Grim
US Census Bureau
Center for Economic Studies
4600 Silver Hill Road
Washington, DC 20233
Cheryl.Ann.Grim@census.gov

Zoltan Wolf
Westat
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, implying that the reallocation of inputs observed across plants is 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 recent theoretical and empirical literature hypothesizes that gains from opening markets to trade are due to the improved allocation of resources induced by trade.¹

While there is considerable consensus that accounting for the dispersion of productivity and its connection to the allocation of activity are important for understanding differences in economic performance, there is not a consensus about the basics of estimating plant-level productivity. Given that typical micro datasets contain information about revenues and input expenditures but not quantities, the majority of results are based on what have become known as revenue productivity measures. How these measures are related to TFPQ, the standard concept in theoretical models of technical efficiency, depends on the assumptions about the environment in which establishments operate. It has become increasingly recognized that there is firm-level price heterogeneity within narrow sectors reflecting at least in part product differentiation and thus some degree of market power at the firm level. In the presence of such heterogeneity, the relationship between TFPQ and revenue productivity reflects the nature and extent to which firm-level prices are endogenous.

A variety of methods are available to researchers to estimate firm-level revenue productivity in the absence of direct measures of prices and quantities. One approach, actively used by the statistical agencies in official aggregate and industry-level productivity statistics, is based on the cost-shares of input expenditures. We call the productivity measure implied by this approach tfpr^{cs} (where cs denotes cost-shares). If plants are cost-minimizing and the assumption of constant returns to scale holds, then cost-shares are valid estimates of output elasticities and the implied revenue productivity measure will be a product of output prices and physical productivity. Other approaches estimate elasticities using regression techniques. We call the implied revenue productivity measure tfpr^{rr} (where rr denotes regression residual). Econometric issues aside², the most important property that distinguishes cost-share-based methods from regression-based techniques is the fact that absent data on quantities or further assumptions about demand, the latter yields revenue elasticities of the revenue function, while the former yields output elasticities of the production function.

While the point that regression methods identify revenue elasticities has been recognized

¹See the 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.

²We discuss these issues in detail below.

before, the implications for the conceptual differences in the interpretation of revenue productivity measures that emerge from these alternatives have not been widely recognized.³ In a highly influential paper that has led to a burgeoning literature on measuring and exploring the implications of misallocation, Hsieh and Klenow (2009) show, under specific assumptions about demand and production technology, that tfpr^{cs} should exhibit no dispersion if marginal revenue products are being equalized.⁴ This insight led to the identifying assumption that tfpr^{cs} can be interpreted as a measure of idiosyncratic firm-level distortions. That is, dispersion in such distortions can be used to measure the extent of misallocation – more dispersion in tfpr^{cs} translates into lower allocative efficiency and productivity. In contrast, we show below that, under the same assumptions that yields this interpretation of tfpr^{cs} , tfpr^{rr} is conceptually different. Specifically, tfpr^{rr} is a measure reflecting fundamentals – for example, TFPQ.

Since both of these approaches are empirically implementable with firm-level data on revenues and inputs, exploring the relationship between these alternative measures provides a means for evaluating the relationship between what has increasingly been used as a measure of distortions and a measure of fundamentals.⁵ We find that that tfpr^{cs} and tfpr^{rr} are highly correlated and exhibit similar dispersion. In addition, we explore the relationship between both measures and survival and growth. Theory implies that firms with higher realizations of fundamentals in a period should, holding initial employment in the period constant, be less likely to exit and more likely to exhibit employment growth. Hence the relationship between survival and growth and realizations of productivity should hold for tfpr^{rr} . However, under the interpretation that tfpr^{cs} reflects distortions there is no inherent reason that these predictions carry over to this measure. We find the relationship between productivity, growth and survival holds equally well for both measures.

These findings raise questions about the interpretation of tfpr^{cs} as an idiosyncratic measure of firm-level distortions. Under this interpretation, our findings imply that firms with high idiosyncratic distortions must also have high idiosyncratic fundamentals. In turn, firms with high idiosyncratic distortions are more likely to survive and grow. Indeed, tfpr^{cs} predicts growth and survival at least as well as tfpr^{rr} . There are some antecedents for our findings in the literature. Using price and quantity data for a small number of products in the U.S., Foster, Haltiwanger,

³De Loecker (2011) amongst others has noted the factor versus revenue elasticity distinction. A few recent papers provide some discussion and evidence on these issues, see Foster, Grim, Haltiwanger, and Wolf (2016b), Haltiwanger (2016) and Decker, Haltiwanger, Jarmin, and Miranda (2017). Foster, Grim, Haltiwanger, and Wolf (2016b) is a short companion paper that focuses on the relationship between tfpr^{rr} and the measure of tfpr that uses output elasticities without imposing constant returns to scale. Haltiwanger (2016) develops the theoretical relationship between tfpr^{cs} and tfpr^{rr} but has no independent empirical analysis. Decker, Haltiwanger, Jarmin, and Miranda (2017) investigate the declining business dynamism as possibly being driven by a rise in adjustment costs. They show that their findings are robust to using either tfpr^{cs} or tfpr^{rr} .

⁴The connection between idiosyncratic distortions and misallocation has been pursued by many researchers. The framework itself was developed by Restuccia and Rogerson (2008). However, Hsieh and Klenow (2009) developed the identifying assumption that tfpr^{cs} reflects distortions. Others have used alternative identification assumptions. For example, see Bartelsman, Haltiwanger, and Scarpetta (2013). By far the most popular methodology for identifying distortions is that developed by Hsieh and Klenow (2009). See, for example, recent work using this approach to understand the slowdown in productivity in Europe (Gopinath, Kalemli-Ozcan, Karabarbounis, and Villegas-Sanchez (2015)).

⁵A related but distinct approach to explore these issues is taken by Haltiwanger, Kulick, and Syverson (2017) who use the Foster, Haltiwanger, and Syverson (2008) data on prices and quantities for a selected set of products to measure tfpr^{cs} , TFPQ and demand shocks directly. The direct approach has many advantages but can only be explored for a limited set of products in the U.S. Haltiwanger, Kulick, and Syverson (2017) also take advantage of the price and quantity data to test the demand and technology assumptions of Hsieh and Klenow (2009). Their findings raise further questions about interpreting tfpr^{cs} as a measure of distortions.

and Syverson (2008), Foster, Haltiwanger, and Syverson (2016c) and Haltiwanger, Kulick, and Syverson (2017) find that tfpr^{cs} is highly correlated with direct measures of TFPQ and demand shocks. These studies also find that survival is increasing in tfpr^{cs} .⁶ A core contribution of this paper is our use of methods that can be widely applied to firm-level databases in many countries since our approach only requires revenue productivity measures. Specifically, our approach focuses on the distinction between tfpr^{cs} and tfpr^{rr} . Given this focus, we conduct extensive sensitivity and robustness analysis to estimating revenue productivity under these approaches.

While our broad conclusion is that tfpr^{cs} and tfpr^{rr} are closely connected with similar properties and implications, there are non-trivial differences in the estimated revenue elasticities across estimation methods. Moreover, examination of the differences between the factor vs. revenue elasticities yields additional insights. Under constant returns to scale (CRTS) and isoelastic demand (common assumptions in the literature), output elasticities estimated via cost shares should exceed revenue elasticities. However, this pattern does not hold generally. We investigate this further by using the methods described in Klette and Griliches (1996) to determine output elasticities by jointly estimating revenue elasticities and demand elasticities. Although this approach is pushing the data quite hard (given that we do not observe plant-level prices and quantities), we find output elasticities do exceed revenue elasticities. Reconciling these disparate findings requires that there be evidence of increasing returns. While the estimates of returns to scale vary substantially across industries, we find that for the average industry there is mild increasing returns. Note this finding does not invalidate the interpretation of tfpr^{rr} as a measure of fundamentals, but it does raise further questions about the interpretation of tfpr^{cs} as a measure of misallocation.

We develop further insights about the importance of these alternative estimation methods by exploring their implications for the contribution of reallocation to productivity growth. For this purpose, we use the structural decompositions developed and implemented by Petrin and Levinsohn (2012) and Petrin, White, and Reiter (2011) (PWR hereafter). These structural decompositions are an ideal setting for this purpose. In these decompositions, the contribution of reallocation to aggregate productivity growth (APG) depends on the relationship between input growth and gaps between marginal value products and marginal costs across firms. In turn, the empirical measurement of these gaps depends critically on the estimated output versus revenue elasticities. In particular, Petrin and Levinsohn (2012) and PWR estimate revenue functions and revenue elasticities rather than output elasticities while their decomposition requires estimates of output elasticities. Moreover, the measure of within plant productivity growth requires adjustment for the presence of the demand elasticity in the tfpr^{rr} measure that emerges from the revenue function estimation. We explore the sensitivity of these decompositions to these issues. Empirically, we find that all methods yield a substantial contribution of reallocation to APG. However, given quantitative differences in the estimated revenue elasticities across methods, we find the quantitative contribution of reallocation, within-plant productivity growth and fixed costs and frictions varies across specifications. In addition, we find that the implementation that makes the adjustment from revenue to output elasticities increases the contribution of the reallocation components. That is, the exact quantitative contribution of reallocation depends critically on these *devilish* details.

The paper is organized as follows. We discuss our methodology and data in Sections 2 and 3. Section 4 describes the effect of estimation methods on the distribution of elasticity estimates. Section 5 describes the implications of the differences in elasticity distributions on productiv-

⁶Eslava (2013) finds similar results using price and quantity data for Columbia.

ity dispersion, plant growth and survival, and aggregate productivity growth decompositions. Section 6 concludes.

2 Methodology

2.1 Productivity measures

A useful starting point is a conceptual measure of revenue per composite input under Cobb-Douglas technology, allowing the decomposition of demand side and supply side factors. This measure can be written in logs as

$$\text{tfpr}_i = p_i + q_i - \sum_j \alpha_j x_{ij} = p_i + \text{tfpq}_i, \quad (1)$$

where i indexes plants, p , q and x denote output prices, quantities and inputs, respectively (with lower case indicating logs of all of the variables), and α_j are the output elasticities of the Cobb-Douglas production function. Time subscripts are omitted in this and subsequent equations for expositional convenience. Equation (1) makes explicit that tfpr_i confounds the effect of output prices and physical productivity, denoted by tfpq_i . Given that typical micro datasets contain information about revenues and input expenditures but not quantities, this property has the important implication that the majority of results in the empirical productivity literature are based on what have become known as revenue productivity measures.

One typical measure is equivalent to tfpr_i if one is willing to assume that plants minimize costs under constant returns to scale technology. Under these assumptions, the shares of input expenditures in total costs, or cost shares, equal their respective output elasticities at least on average across plants and in the long run.⁷ Formally,

$$\text{tfpr}_i^{cs} = p_i + q_i - \sum_j \alpha_j^{cs} x_{ij} = \text{tfpr}_i. \quad (2)$$

Other estimation methods rely on regression techniques to obtain revenue productivity. We denote these measures by tfpr_i^{rr} in order to distinguish them from tfpr_i^{cs} . Absent data on prices and/or quantities, these methods yield revenue function estimates or elasticities, which are a function of output elasticities and demand parameters, in general. As a result, tfpr_i^{rr} depends on physical productivity (and if present demand shocks) and is not equal to either tfpr_i or its estimate, tfpr_i^{cs} :

$$\text{tfpr}_i^{rr} = p_i + q_i - \sum_j \beta_j x_{ij} \neq \begin{cases} \text{tfpr}_i \\ \text{tfpr}_i^{cs} \end{cases}. \quad (3)$$

To highlight the importance of the difference relative to tfpr_i^{cs} and tfpr_i , we denote revenue elasticities by β_j in equation (3). To gain insight into the differences, it is instructive to assume a constant elasticity of substitution (CES) demand structure which is a key assumption in the Hsieh and Klenow (2009) identification approach. Specifically, assume firms are not price takers in output markets but instead operate under isoelastic, downward sloping demand conditions. The inverse demand function can be written as $P_i = Q_i^{\rho-1}$ where $\rho - 1$ is the inverse of the

⁷We discuss how reasonable these assumptions are below.

price elasticity of demand with $\rho < 1$.⁸ Also, let the firm-level production function be given by $Q_{it} = A_i \Pi_j X_{ij}^{\alpha_j}$. tfpq_i (the log of TFPQ_i) is equal to a_i in this notation. Under these assumptions, revenue elasticities are equal to product of output elasticities and $\rho: \beta_j = \alpha_j \rho$, implying revenue elasticities are lower than output elasticities.

In this CES demand case, there are also implications for the interpretation of the relationship between tfpr_i^{cs} and tfpr_i^{rr} . In the absence of any frictions or distortions, marginal revenue products for each factor will be equalized, and there will be no dispersion in tfpr_i^{cs} within sectors. Since the latter is counterfactual, Hsieh and Klenow (2009) posit the presence of idiosyncratic distortions that can account for such dispersion. Such idiosyncratic distortions reflect firm-specific factors that prevent the equalization of marginal revenue products. To illustrate this point, we consider scale (output) idiosyncratic distortions such that firms are maximizing static profits given by $(1 - \tau_{Q_i})P_{it}Q_{it} - \sum_j w_j x_{ij}$ and neglect relative input distortions. This implies:⁹

$$\text{tfpr}_i^{cs} = \ln(1/(1 - \tau_{Q_i})). \quad (4)$$

In contrast, these same assumptions imply:¹⁰

$$\text{tfpr}_i^{rr} = \rho a_i. \quad (5)$$

The key implication for our purposes is that under these assumptions tfpr_i^{cs} is not inherently a function of idiosyncratic firm-level fundamentals but tfpr_i^{rr} is only a function of such fundamentals.¹¹ It is this conceptual difference that motivates much of the empirical analysis below. Specifically, we investigate the empirical relationship between tfpr_i^{cs} and tfpr_i^{rr} . Doing so requires constructing empirical estimates of these two measures to which we turn in the next section.

Before proceeding, we note that the assumptions under which tfpr_i^{cs} reflects only distortions are very strong. If returns to scale are not constant or demand is not CES, this finding does not hold (see, e.g., Haltiwanger, Kulick, and Syverson (2017), Haltiwanger (2016) and Foster, Grim, Haltiwanger, and Wolf (2016b)). In contrast, the finding that tfpr_i^{rr} is only a function of fundamentals is robust to allowing returns to scale to differ from one. In the empirical analysis that follows, we explore the issue of constant returns to scale in this context. We do not explicitly explore non-CES demand structures but discuss our findings below in light of the

⁸This is a simplified version of a more general specification given by $P_i = P(Q_i/Q)^{\rho-1} \xi_i$, where ξ_i is an idiosyncratic demand shifter, P and P_i denote aggregate and plant-level product prices, and Q and Q_i denote aggregate and plant-level product quantities, respectively. This specification is consistent with assuming that the final good is a CES aggregator of intermediate goods produced by individual firms. The final goods sector is perfectly competitive with the only inputs coming from intermediate goods. See Bartelsman, Haltiwanger, and Scarpetta (2013) and Foster, Grim, Haltiwanger, and Wolf (2016b) for more details. The specification in the main text abstracts from idiosyncratic demand shifters and also terms that vary at the industry and aggregate. The latter terms are not important for the analysis in this paper since we typically control for industry by year effects.

⁹This is consistent with footnote 10 of Hsieh and Klenow (2009) with only scale distortions. Our specification neglects industry-level factors of proportionality that don't matter for our analysis.

¹⁰See appendix A.1. We again are neglecting industry-level factors of proportionality that are present in a more general specification. These industry-level factors include input prices assumed to be the same at the industry or even the economy-wide level. If they are idiosyncratic then they are another possible source of idiosyncratic firm-level fundamentals. They can also lead to dispersion in the tfpr_i^{cs} measures.

¹¹In a more general specification, tfpr_i^{rr} is also a function of idiosyncratic demand shocks.

studies that consider this possibility.

2.2 Estimation methods

There are various estimation methods available to researchers; table 1 summarizes the procedures we use throughout the paper. We discuss the strengths and weaknesses of the most popular methods from a practical point of view.

Table 1: Estimation methods.

Method	Description	Proxy	Estimator	Productivity measure
GA	Foster, Haltiwanger, and Krizan (2001)		Cost-shares	tfpr^{cs}
OLS	Ordinary Least Squares		LS	tfpr^{rr}
OP	Olley and Pakes (1996)	Investment	NLS	tfpr^{rr}
LPGR	Levinsohn and Petrin (2003)	Materials	GMM	tfpr^{rr}
WLPE	Wooldridge (2009)	Energy	Efficient GMM	tfpr^{rr}
WLPM	Wooldridge (2009)	Materials	Efficient GMM	tfpr^{rr}

All productivity measures in the table are estimated using an output-based production function.

Cost-share-based or growth accounting methods (GA) exploit first order conditions from the firm’s cost-minimization problem. There are several advantages to this procedure. For example, this is the only estimator that - conditional on an assumption about returns to scale - provides direct estimates of output elasticities without data on prices and quantities. This is a useful property because it implies cost-share-based elasticities are robust to alternative demand structures. In addition, using the share of input expenditures in total costs rather than in value added has the advantage that the assumption of perfectly competitive markets is not required. In addition, it allows the exact shape of technology to vary down to the plant-level.¹² A potential caveat is that an estimate of, or an assumption about, returns to scale is necessary. In addition, this cost share approach requires an assumption that first order conditions hold at least on average. This property is relevant because it is unlikely that first order conditions hold for every plant at all points in time. Therefore, common output elasticities across plants in the same industry and/or over time are frequently imposed (see Syverson (2011)), which is the approach we take in this paper. We explore the implications of these assumptions further in appendix A.7.3 and note that all alternative estimation methods discussed in this paper also require such restrictions.

Ordinary least squares (OLS) estimates of elasticities are inconsistent because unobserved productivity shocks affect the decision problem of plants, see Marschak and Andrews (1944). Additional biases may result if prices are unobserved and correlated with factor inputs, see Klette and Griliches (1996). We include OLS in our analysis as a point of reference but focus on regression methods that are often referred to as proxy methods. They are denoted as such because they use firm-level indicators to control for the effect of unobserved productivity during estimation. The original logic, developed in Olley and Pakes (1996) (OP hereafter), is based on assuming a monotonic and increasing relationship between the proxy and productivity. On condition that productivity is the only unobserved state variable, plant-level variation in the proxy can be interpreted as a reflection of the variation in productivity shocks. This conclusion

¹²Further, evidence in Van Biesebroeck (2007) suggests that this method is accurate if the data are not subject to much measurement error.

together with timing assumptions about the plant’s decisions, can be used to recover all the elasticities.¹³ OP propose investment as a proxy. There is ample evidence that plant-level investment is lumpy (see for example Cooper and Haltiwanger (2006)), and lumpiness means bursts of investment activity are followed by inactive periods where observed net investment is zero rendering these observations uninformative for OP. To avoid the resulting drop in sample size, Levinsohn and Petrin (2003) (LP hereafter) advocate using intermediate input expenditures or energy costs.¹⁴ LP argue that if adjusting intermediate inputs is less costly, then they are likely to be more responsive to productivity shocks. This is especially relevant in the presence of non-convexities in capital adjustment.

The identifying assumptions regarding the timing of plants’ input decisions have been criticized by Akerberg, Caves, and Frazer (2015) (ACF hereafter). ACF argue that the optimal labor allocation is also a deterministic function of productivity and therefore the labor elasticity is not identified in the first step. They approach the identification problem by applying a two step procedure that does not try to identify any of the 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 that arises from two-step estimation is eliminated.

In closing this section, we note that estimating output elasticities under endogenous plant-level prices requires additional assumptions about demand, see Klette and Griliches (1996) or De Loecker (2011) as examples. Section 4 shows such considerations are empirically relevant for inferences about returns to scale.

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 database¹⁵, the Bureau of Labor Statistics and the 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 in years ending in ’2’ and ’7’ for roughly 180,000 - 240,000 plants. Establishments with less than five employees are not sent forms. Payroll and employment data for these very small plants are imputed using administrative records.¹⁶ The ASM surveys 50,000-70,000 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 (2016a) (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

¹³More details on these estimation methods can be found in appendix A.2.

¹⁴LP highlight that firms almost always report positive use of these variables in their data implying truncation due to zero proxy values is less severe.

¹⁵The NBER-CES Manufacturing Industry database is available at <http://www.nber.org/nberces>. An earlier version is documented in Bartelsman and Gray (1996).

¹⁶We drop administrative records cases.

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 the total wage bill to production worker wages. Our intermediate input variable is given by the the sum of three items: 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 needed 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.¹⁷

3.2 Analysis samples

The analyses in this paper exploit three different samples. For questions about the distributions of elasticities, productivity dispersion and 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 industry should be large enough so that elasticities can be estimated by all reviewed methods. Second, industries should be defined narrowly enough so that we can plausibly assume elasticities are constant across 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. Since we estimate elasticities on an industry-by-industry basis, changes in the classification system entail spurious breaks in plant-level time series and a drop in sample size. In the first part of the analysis we address these issues by selecting sets of 4-digit SIC industries which were not affected by classification changes or which were mapped one-to-one into another industry. There are 292 such industries of which we selected the first 10 and 50 based on the number of observations.

We create a third dataset to test whether the implications of the decomposition of aggregate productivity growth described in PWR are sensitive to the way productivity 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. More details about the

¹⁷Employment 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 productivity level of an establishment is calculated from imputed data. The use of inverse propensity score weights has advantages relative to ASM sample weights. It is important to note that using sample weighted totals does not correspond to published ASM totals. There are further adjustments for the non-mail universe and other factors. In addition, the treatment of the latter has varied over time so the relationship between sample weighted and published totals varies over time.

assignment procedure can be found in appendix A.3.

4 Elasticity distributions

We start by discussing differences in the distribution of capital elasticities. Next, we check whether the elasticity-rank of industries varies with estimator choice. We conclude the section by looking at the implications of estimator choice on returns to scale.

Figure 1 plots the density estimates of capital elasticities in the 50 most populous industries and table 2 shows basic descriptive statistics. There are non-trivial differences in the mean, dispersion and general shape of the distributions. Most notably, GA-based estimates tend to be generally smaller than the rest. At first glance, this is contrary to expectations since under CRTS technology and CES demand the cost share output elasticities should exceed revenue elasticities. We explore this issue further below.

We now turn our attention to the revenue elasticity distributions and focus on indications of possible biases. Unfortunately, the differences in estimated distributions alone do not tell much about the direction, much less the magnitude, of any bias. However, in light of what we know about the way these methods address endogeneity, the differences may give us clues as to whether or not they correct it in the right direction. The direction of the bias in OLS-estimates is determined by several factors. First, since input demand functions are increasing in productivity, OLS estimates are biased upward. If this is important in our data and proxy methods correct for it, then we should see proxy-based distributions to the left of OLS. Figure 1 suggests that only LPGR is likely to yield lower $\hat{\beta}_k$. However, the direction of the bias depends on additional factors. For example, Levinsohn and Petrin (2003) show that positive correlation between capital and labor may cause $\hat{\beta}_k$ to be biased downward, implying the elasticity distribution may emerge to the left of OLS not because LP corrects an upward 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 the right of OLS.

Table 2: Descriptive statistics of the between-industry distribution of $\hat{\beta}_k$.

	Mean	Median	IQR	Mean	Median	IQR
	50 largest industries			10 largest industries		
OLS	.13	.10	.10	.14	.12	.15
OP	.19	.14	.10	.21	.15	.27
LPGR	.09	.06	.09	.11	.06	.09
WLPE	.16	.15	.11	.17	.14	.07
WLPM	.16	.12	.14	.15	.09	.14
GA	.08	.07	.05	.09	.10	.06

See notes to table 1 for method definitions. 10 and 50 largest industries: most populous 4-digit industries industries which were mapped 1-to-1 between classification systems.

We find that 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




Figure 1: Cross-industry distributions of $\hat{\beta}_k$.
The sample is 50 largest industries, estimators are described in table 1.

similar to those under OP. These results tell us that, in addition to proxy choice and addressing selection, the structure of the estimator - non-parametric two-step estimator versus efficient GMM - may also have important consequences for capital elasticities. Similar conclusions hold for the second moments of $\hat{\beta}_k$. The interquartile range measures, see the last column of table 2, suggest differences in the dispersion of these distributions are also non-trivial. For example, the means of GA and LPGR are very close (.08 and .09 among the 50 largest industries) but the dispersion of the latter is almost twofold (.05 and .09). Finally, the elasticities of variable inputs show stronger clustering, especially $\hat{\beta}_l$ (figure 2). This is partly explained by the fact that proxy methods estimate β_l in an OLS step (the exception is the Wooldridge GMM method). The main conclusions about $\hat{\beta}_e$ and $\hat{\beta}_m$ are the same, (figure A4): there are numerical differences in these distributions but they look generally similar across estimation methods. For both $\hat{\beta}_l$ and $\hat{\beta}_m$, the estimates of the revenue elasticities tend to be lower than the cost shares based estimates of output elasticities consistent with the hypothesis discussed above.¹⁸

In closing this section, we highlight the different implications estimation methods have for returns to scale. GA-based elasticities are consistent with constant returns to scale by construction. In contrast, regression based revenue function estimators do not impose constant returns to scale but do not provide direct returns to scale estimates either. However, if one is willing to accept the assumption that plants face isoelastic product demand, joint estimation of revenue elasticities (β_j) and demand characteristics (ρ) can be used to make inference about output elasticities and therefore returns to scale. We follow Klette and Griliches (1996), and jointly estimate β_j and ρ by regressing plant-level revenues on plant-level inputs and an indicator of industry-level output. This approach, while feasible, has relatively low power for estimating

¹⁸We did some further exploration of the sensitivity of the patterns of the estimators by examining how the estimator choice affects the ranking of industries by $\hat{\beta}_k/\hat{\beta}_l$. We find there is a positive probability that different estimators imply different industry rankings, see table A10 for details. We further investigate the empirical properties of elasticity distributions in Appendix A.4. Those findings highlight that some methods yield systematically more zero/negative elasticities than others. We find evidence that, in certain cases, this can be attributed to small sample size and/or proxy choice.

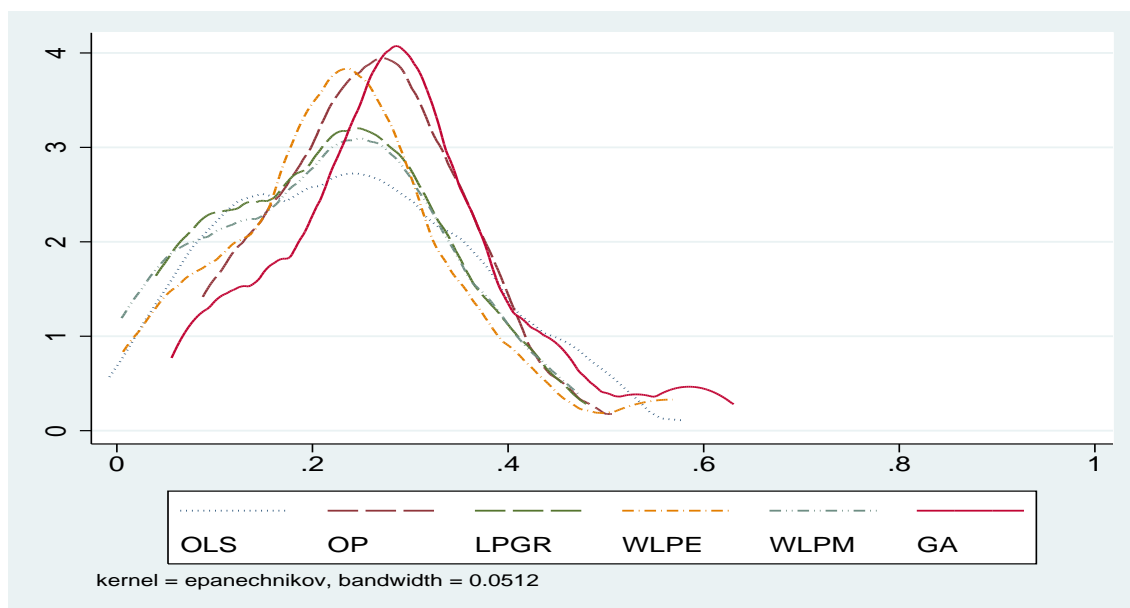


Figure 2: Cross-industry distributions of $\hat{\beta}_i$.
The sample is 50 largest industries, estimators are described in table 1.

demand parameters compared to methods that can be used when both price and quantity data are available, see Foster, Haltiwanger, and Syverson (2008). Our results, based on OP’s method,¹⁹ indicate that the elasticity of firm revenues with respect to aggregate demand is a small positive number, implying a negative average price elasticity and greater-than-1 markup value, see table A11 for details. To further illustrate the variation in these parameters, we compared the densities of the sums of revenue and output elasticities, see figure 3. The sum of the revenue elasticities has a mean of 0.97 while the sum of the output elasticities has a mean of 1.09. The latter suggests mild increasing returns. However, the sum of estimated output elasticities shows large dispersion reflecting the variation both in revenue elasticities and demand parameters suggesting that appropriate caution is needed when interpreting results, especially on an industry-by-industry basis.

¹⁹We chose OP because adding more state variables for this procedure is straightforward.

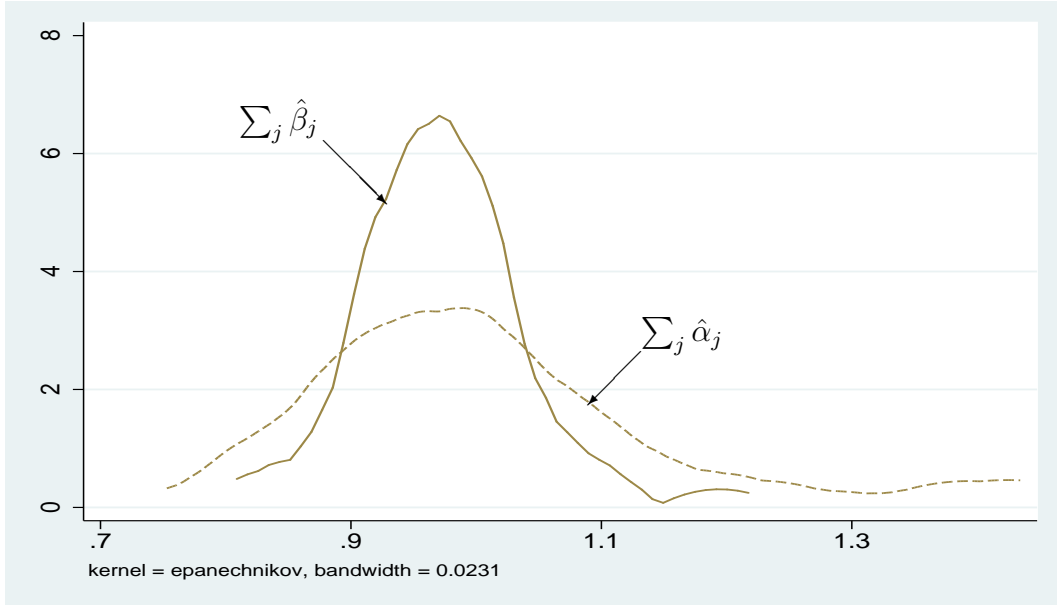


Figure 3: Cross-industry distributions of $\sum_j \hat{\beta}_j$ and $\sum_j \hat{\alpha}_j$ and the demand parameter are jointly estimated using the method described in Olley-Pakes (1996).

5 Implications of the differences in elasticity distributions

In this section, we turn to the main questions of interest. Specifically, we examine the relationship between the alternative measures of revenue productivity in terms of dispersion, correlations and the relationship between productivity, growth and survival. We use this investigation to assess the extent to which these conceptually different measures are indeed different empirically. We also explore the sensitivity of structural decompositions of aggregate productivity growth (APG) to these different approaches to estimating factor and revenue elasticities.

5.1 Productivity dispersion

Does the choice of estimator affect the second moment of the within-industry productivity distribution? The interquartile range (IQR) and standard deviation, averaged over industries and time, are presented in table 3.²⁰ The average IQR of the log-productivity distribution varies between 0.24 and 0.40 across methods. The dispersion in $\text{tfpr}_i^{\text{cs}}$ is lower than the dispersion in $\text{tfpr}_i^{\text{rr}}$ (from the alternative proxy methods). The WLPM method yields the highest dispersion. Recall from figure 1 that WLPM has the most extreme outliers for β_k . These outliers mean WLPM is a repeated outlier in the following analyses. The proxy methods including WLPM yield less variation in dispersion for the ten largest industries compared to the fifty largest industries.

Table 3 demonstrates that all methods yield substantial within industry dispersion in revenue productivity but there are numerical differences. Should we think of these as substantial? One way to answer this question is to test their equality by comparing confidence bands around average dispersion or comparing industry-specific results. Figure 4 shows simulated two-standard-error-wide confidence bands (narrow bars) around our dispersion estimates. If we interpret the narrow bars as approximately 95% confidence intervals, we can conclude that

²⁰More details are available in table A12.

Table 3: Descriptive statistics of productivity distributions.

	N (1000)	IQR	SD	N (1000)	IQR	SD
	50 largest industries			10 largest industries		
OLS	455	0.26	0.26	185	0.22	0.20
OP	380	0.32	0.38	152	0.31	0.39
LPGR	457	0.29	0.31	188	0.29	0.28
WLPE	457	0.34	0.36	187	0.31	0.28
WLPM	457	0.40	1.88	188	0.33	0.35
GA	433	0.24	0.22	177	0.23	0.21

See notes to tables 1 and 2 for the description of methods and samples. All statistics are based on deviations of plant-level log-productivity from industry- and time-specific means. All results shown were calculated using non-outlier observations only (pre-, post-estimation). A version of the table including pre-estimation outlier observations can be found in the appendix (table A12); the results show little change.

dispersion measures under proxy methods are not significantly different from each other.²¹ In other words, the sampling variation in revenue function estimates has similar effects across proxy-based estimates.²² The exception is GA, which generally yields lower dispersion.

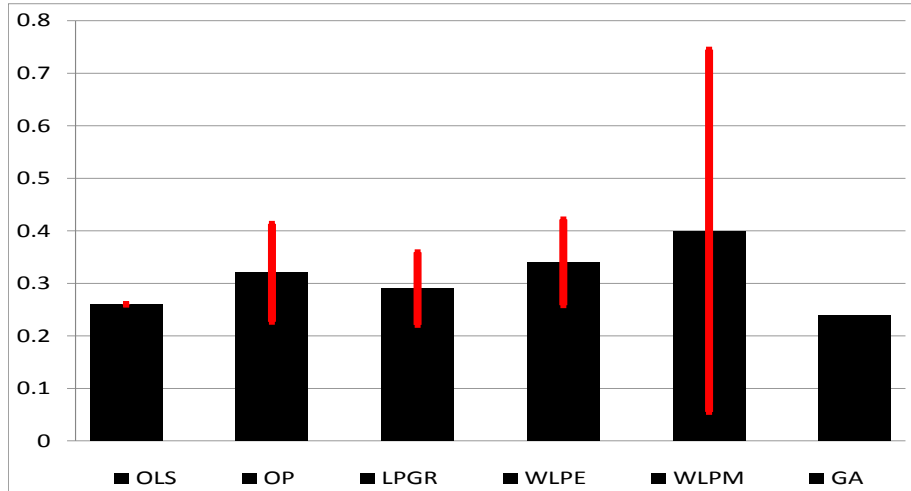


Figure 4: Bootstrapped standard errors of dispersion under selected methods.

The black bars are mean dispersion estimates from table 3, the narrow bars represent confidence bands based on bootstrapped standard errors of dispersion statistics, see also table A16. GA has no interval around the mean because coefficients are estimated using industry-level data on input costs.

We next investigate whether the choice of estimation method has consequences also for the within-industry productivity rank of establishments. Table 4 shows both the Pearson and Spearman rank correlations. Of particular interest is the correlation between tfpr_i^{cs} and tfpr_i^{rr} . Using the Spearman (rank) correlations, all pairwise correlations except for the WLPM are 0.60 or higher. All Pearson correlations exceed 0.46 except for WLPM.²³ In general, there is weaker

²¹Further analysis shows that differences may be greater at the industry-level, see figure A1 in Appendix A.6.

²²More details can be found in figure A16.

²³We put little weight on the patterns for the WLPM since this procedure, as noted earlier, exhibits large outliers in revenue elasticities.

pairwise correlation across the proxy methods than between each of the proxy methods and the GA approach. For example, the LPGR has a Pearson correlation with WLPE of 0.52 and a 0.68 correlation with the GA approach (this pattern carries over to the Spearman correlations).

We interpret these initial exercises as implying that tfpr_i^{cs} and tfpr_i^{rr} exhibit broadly similar within industry dispersion and are strongly correlated. The dispersion in tfpr_i^{rr} tends to be somewhat larger than tfpr_i^{cs} . The pairwise correlations are quite strong particularly for the rank (Spearman) correlations. Under the assumption of isoelastic (CES) demand, tfpr_i^{rr} is a measure of fundamentals. Our findings imply that whether or not tfpr_i^{cs} is an appropriate measure of distortions, it is positively correlated to, and similarly dispersed as, fundamentals.

Our findings are consistent with the recent literature using direct measures of prices and quantities. Foster, Haltiwanger, and Syverson (2008), Foster, Haltiwanger, and Syverson (2016c), Haltiwanger, Kulick, and Syverson (2017), and Eslava (2013) provide evidence that tfpr_i^{cs} is highly correlated with direct measures of tfpq_i and positively correlated with demand shock measures that can be estimated with price and quantity data. Moreover, these studies find the dispersion of tfpr_i^{cs} is slightly lower than tfpq_i . For the US, these findings are restricted to a relatively small (11) number of products.

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

	OLS	OP	LPGR	WLPE	WLPM	GA
Pearson						
OLS	1					
OP	0.51	1				
LPGR	0.82	0.43	1			
WLPE	0.51	0.46	0.52	1		
WLPM	0.02	0.02	-0.15	0.15	1	
GA	0.79	0.46	0.68	0.51	0.09	1
Spearman						
OLS	1					
OP	0.68	1				
LPGR	0.87	0.61	1			
WLPE	0.59	0.61	0.56	1		
WLPM	0.35	0.36	0.26	0.49	1	
GA	0.81	0.63	0.70	0.60	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 has a minor effect on correlations (see table A13).

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 productivity is estimated. For this purpose, we are particularly interested in models where firms face adjustment frictions on both the entry/exit margins as well as on the intensive margin.²⁴ We focus our attention on incumbents (existing firms) and thus on the exit and intensive margin of adjustment where firms face fixed costs of operating each period and adjustment costs for changing the scale of operations. For purposes

²⁴For example, on the entry/exit margins see, e.g., Hopenhayn (1992) and Hopenhayn and Rogerson (1993). For adjustment cost models at the firm-level on employment, see, e.g., Cooper, Haltiwanger, and Willis (2007) and Elsby and Michaels (2013).

of exposition, we focus on describing the implications of these models with one factor of production subject to adjustment costs: employment. In this class of models, incumbent firms have two key state variables each period: the prior period level of employment and the realization of productivity/profitability in the period. Relating this to the discussion above, the latter are the realizations of TFPQ (and potentially demand shocks) in the period. These models predict that firms with sufficiently low draws of productivity/profitability shocks should exit. Moreover, holding the prior period level of employment constant, firms with higher realizations of productivity/profitability should grow.

Syverson (2011) highlights that these broad predictions of growth and survival being closely related to productivity are ubiquitous findings in the literature. However, in the context of the above briefly described models, these predictions should hold with respect to measures of revenue productivity that reflect fundamentals. That is, the predictions should hold for tfpr_i^{rr} but there is no inherent reason that the predictions should hold for tfpr_i^{cs} . To investigate these issues, we consider the relationship between productivity and growth of all establishments, exiters and incumbents separately. We estimate specifications given by:²⁵

$$Y_{it+1} = \gamma_1 \omega_{it} + \gamma_2 \theta_{size_{it}} + X'_{it} \beta + \epsilon_{it+1}, \quad (6)$$

where Y_{it+1} is the outcome of interest such as growth between periods t and $t + 1$ at plant i , ω is a plant-level measure of productivity, θ_{size} is the control for the initial size (employment) in the period, and X_{it} is a vector of additional controls including year effects, state effects, and cyclical controls (the change in the unemployment rate at the state level).²⁶

Table 5 shows $\hat{\gamma}_1$ from equation (6) using our sample of the 50 most populous industries. Each row lists the effect of productivity on a specific outcome, shown in the row header. The three outcomes are: employment growth among all establishments, the probability of exit, and employment growth among continuers. The columns show results under productivity estimator variants. For example, the first entry in column 1 says that a plant is estimated to grow approximately 0.16% faster if it is 1% more productive under an OLS-based estimator. All other entries are analogous. Point estimates suggest there are non-trivial differences in the productivity effect. For example, the first entry in column 2 shows that the OP-based productivity effect is less than half of the OLS-based one. The difference between the results of these two estimators encompasses the variation in coefficients, ignoring for now WLPE and WLPM. Despite these non-negligible differences, the estimates support the earlier finding that more productive plants grow significantly faster than their less productive competitors. The estimates in row 3 show low-productivity establishments are significantly more likely to exit than high-productivity establishments.

The results in Panels B and C highlight the sensitivity of some methods to changes in sample definition. Panel B shows results when elasticities are estimated pooling data from all 4-digit industries within the same 3-digit industry. Pooling amounts to allowing stricter assumptions about the homogeneity of elasticities, which may be justified in order to increase precision. 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 indicating that pooling may be beneficial because if there is no

²⁵This is a simplified version of the specification considered by FGH.

²⁶We follow FGH by using the integrated LBD with the ASM data for this analysis. The ASM 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 $t + 1$.

Table 5: The effect of productivity on outcomes, sample is 50 largest industries. Outcomes are: employment growth among all establishments (row 1), exit (row 2), employment growth among continuers (row 3).

	OLS	OP	LPGR	WLPE	WLPM	GA
Panel A. 4-digit elasticities						
overall growth	0.163***	0.072***	0.139***	0.06***	-0.006***	0.190***
exit	-0.050***	-0.018***	-0.046***	-0.020***	0.002**	-0.064***
conditional growth	0.068***	0.039***	0.053***	0.022***	-0.003***	0.067***
Panel B. 3-digit elasticities						
overall growth	0.190***	0.155***	0.152***	0.100***	-0.017***	0.183***
exit	-0.066***	-0.048***	-0.030***	-0.030***	0.005**	-0.062***
conditional growth	0.061***	0.065***	0.071***	0.043***	-0.008***	0.064***
Panel C. 3-digit elasticities, industries with negative or non-estimable elasticities dropped						
overall growth	0.197***	0.153***	0.164***	0.104***	0.076***	0.183***
exit	-0.069***	-0.044***	-0.030***	-0.036***	-0.028***	-0.062***
conditional growth	0.064***	0.072***	0.074***	0.035***	0.023**	0.064***

Estimates are taken from regressions of three outcomes (employment growth among all establishments, exit, and employment growth among continuers) on a plant-level measure of productivity (columns), a state-level measure of unemployment growth, year-, sizeclass- and state-fixed effects, see equation (6). Standard errors are clustered at the state level.

** and *** denote 5% and 1% significance levels, respectively. All regressions are based on trimmed productivity distributions. Sample size information can be found in table A17. Results for two additional industry sets can be found in tables A18-A19.

statistical association between two variables, their partial correlation coefficient would tend towards zero. Other estimators seem less prone to such changes and remain in a comparable range. WLPM yields the counter-intuitive result 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, our earlier results already highlight that WLPM produces in some cases extreme values of elasticities (see figure 1). As further sensitivity analysis, Panel C drops industries with negative revenue elasticities. This has a modest impact on all methods except for the results using the WLPM. In Panel C, the estimated marginal effects of revenue productivity on growth and survival are more in line with the other methods.²⁷

To sum up, our estimates show positive (negative) and significant association between productivity and growth (exit), irrespective of how productivity is estimated. Some methods yield outlier elasticities that weaken the estimated relationship between productivity, growth and survival. In these cases, increasing sample size by using broader industry definitions yields results that line up better across methods.²⁸ The similarity of conclusions discussed in this

²⁷We carried out similar exercises using other samples from our default industry set. Without discussing them in detail, we note 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. See tables A18-A19 for more details.

²⁸A related question concerns the effects of allowing for not less but more heterogeneity in elasticities. Appendix A.7.3 shows that using productivity numbers that are based on plant-specific elasticities result in significantly higher productivity dispersion. The increase also implies weaker relationships between productivity, growth and survival. We interpret the latter finding as indirect evidence that the increased dispersion is due at

section is a good example of the reasons why researchers have not focused on the differences between estimation methods.

5.3 TFPR: Fundamentals versus Distortions?

Taking stock of the results so far, tfpr_i^{cs} and tfpr_i^{rr} (i) exhibit similar dispersion, (ii) are strongly positively correlated and (iii) exhibit very similar relationships with growth and survival. In many cases, the quantitative differences in results are greater across methods for estimating tfpr_i^{rr} than between tfpr_i^{rr} and tfpr_i^{cs} . This is partly because the proxy methods in some cases yield extreme estimates of elasticities. Our findings suggest that pooling data across plants within the broader industry categories (e.g., estimating elasticities at the 3-digit as opposed to the 4-digit level) mitigates this problem.

What do these results imply about interpreting tfpr_i^{cs} as a measure of idiosyncratic distortions? To maintain that assumption, idiosyncratic distortions would need to exhibit similar dispersion and be highly correlated with idiosyncratic fundamentals. Moreover, idiosyncratic distortions would need to be as tightly connected to survival and growth as idiosyncratic fundamentals. Indeed, our findings suggest that growth and survival are as closely linked to tfpr_i^{cs} as to tfpr_i^{rr} . It might be feasible to reconcile these findings with a particular pattern of idiosyncratic distortions. However, there are alternative interpretations of these measures that are arguably more plausible. Specifically, adjustment frictions of capital and labor can potentially reconcile these findings as well as departures from the Cobb-Douglas CRTS technology and CES demand.

We do not formally develop these alternatives here since our focus is on estimating and investigating the empirical properties of the alternative revenue productivity measures. However, there are a number of studies exploring these alternatives and we summarize the intuition from one of the leading alternatives that has been suggested. Specifically, consider the potential role of adjustment costs for employment that we used to motivate the empirical specifications above. A firm with a positive realization of log-TFPQ_i (a_i) will exhibit an increase in desired employment. In a frictionless environment the firm increases employment to the point where marginal revenue products are equal to the input factor cost (assumed to be the same across firms). The increase in employment (and other factors) will yield an increase in firm-level output and a decline in firm-level prices. The firm-level full adjustment in a frictionless environment implies that prices will decline just enough so that $\text{tfpr}_i^{cs} = p_i + a_i$ will remain unchanged. But in the presence of adjustment frictions, the increase in inputs and output will be smaller so the decline in prices will be smaller. Accordingly, the positive realization of a_i will yield an increase in tfpr_i^{cs} . It can be shown that this positive correlation also implies that realizations of tfpr_i^{cs} will be related to input growth in a very similar manner as tfpr_i^{rr} .²⁹ In short, adjustment frictions have implications that match the core findings of our analysis.

One possible perspective is that interpreting dispersion of TFPR as reflecting distortions is a reduced-form means of capturing sources of dispersion in marginal revenue products across firms including the role of adjustment costs. This perspective has a number of limitations. First, distinguishing between frictions that the social planner cannot avoid (such as some aspects of

least partly to measurement error implying that using plant-level shares is unlikely to be optimal. The effect of imputation on these coefficients is similar. However, despite attenuated coefficients, the main conclusions hold; see appendix A.7.2 for more details.

²⁹See Haltiwanger, Kulick, and Syverson (2017) and Decker, Haltiwanger, Jarmin, and Miranda (2017) for formal development of this intuition. Asker, Collard-Wexler, and De Loecker (2014) also explore the role of adjustment frictions in accounting for dispersion in firm-level productivity measures.

adjustment frictions of capital and labor) and distortions seems important. Second, adjustment frictions have inherently dynamic implications while a distribution of idiosyncratic distortions does not. That is, with adjustment frictions, firms will be adjusting to their optimal level of activity over time while with idiosyncratic distortions this is not inherently apparent. An associated implication is the point made above – adjustment frictions inherently yield a strong positive correlation between fundamentals and marginal revenue products as well as TFPR. This strong positive correlation arguably underlies many of the findings in the literature using TFPR as a proxy for fundamentals (see Syverson (2011)).

5.4 Structural decompositions of aggregate productivity growth

We now examine whether the estimation method affects the results from structural decompositions of aggregate productivity growth (APG). We focus on these structural decompositions for a number of reasons. First, they provide a decomposition of the contribution of within firm versus reallocation for productivity growth using an internally consistent theoretical framework. Second, these decompositions rely critically on estimates of output and revenue elasticities. Petrin and Levinsohn (2012) and PWR show that the reallocation contribution terms depend on the relationship between input growth and gaps between marginal value products of inputs and the marginal cost of inputs. Measurement of these gaps depends directly on the estimation of output elasticities. Using the proxy methods estimation of revenue functions (as Petrin and Levinsohn (2012) and PWR do) yields revenue instead of output elasticities. Moreover, the residual from the revenue function is a function of fundamentals like TFPQ but its variation (and therefore its growth) is scaled by a demand elasticity parameter. Thus, in principle, both the within and reallocation components in these decompositions require taking these issues into account.

An attractive property of this approach is that the measure of APG itself does not depend on the way productivity is estimated nor on the presence of imperfect competition because APG is defined as the growth in final demand in excess of capital and labor growth.³⁰ This implies that we can readily compare the contribution of within and reallocation components across estimation methods and across assumptions about price taking versus price setting behavior.

The implementation of these decompositions by Petrin and Levinsohn (2012) and PWR requires no adjustment if one assumes homogenous products within industries and price taking behavior. Under these assumptions, revenue elasticities are equal to output elasticities. Even though much of this paper has focused on exploring departures from this assumption, we begin by considering this case. This is a useful starting point because then our results can be directly compared to PWR and also we can assess how the different methods yield differences in contributions even under this assumption.

Table 6 summarizes our results under the price taking assumption. Panel A shows the elements of the APG definition in PWR. The annual average growth rate for labor (-0.3%) in our sample is similar to that in PWR (-0.2%) but our value added and capital growth rates are smaller. These differences are due partly to measurement differences and partly to an issue with PWR’s source data for implicit deflators (see the notes to table 6 for more details). However, value added growth is closer to PWR’s measure if we ignore the first years of ASM panels, where calculating aggregate growth is problematic given panel rotation issues

³⁰Since final demand is not observed, PWR 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 by implementing the decomposition for a subset of plants we compute the contribution of the 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.

Table 6: Aggregate productivity growth and its decomposition. Annual averages calculated between 1977 and 1996 using ASM continuers.

Panel A. APG definition (%)							
	All years		Ignoring first years*				
	Our sample	PWR	Our sample	PWR			
VA	1.9	2.3	1.4	1.6			
Capital	0.0	0.3	0.0	0.4			
Labor	-0.3	-0.2	-0.5	-0.4			
APG	2.2	2.2	1.9	1.6			

Panel B. Annual average contributions by components, ignoring first years.							
	OLS	OP	LPGR	WLPE**	WLPM**	GA	PWR
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Total RE:	0.8	0.7	0.7	1.3	1.3	0.5	2
Capital	0.2	0.2	0.3	0.8	0.6	0.1	0.9
Labor	0.0	0.0	-0.1	0.0	0.0	0.1	0.4
Materials	0.4	0.4	0.4	0.5	0.5	0.3	0.4
Energy	0.2	0.1	0.1	0.0	0.1	0.0	0.3
Within	0.8	0.7	1.4	-0.1	-0.1	1.1	-0.1
Fixed costs	0.3	0.5	-0.2	0.8	0.7	0.3	-0.2
σ_{RE}	0.8	0.5	0.7	0.9	0.8	1.2	1.7
σ_{Within}	2.3	2.8	3.1	2.5	2.5	2.4	2.7

*Measuring growth in first years of ASM panels is problematic because weighted growth rates are based on only large plants since growth rates do not exist for establishments just rotated in. **Instruments as in PWR: second and third lags. The last column in each panel is based on tables 1, 2 and 3a in PWR.

PWR generate data on non-production worker hours using variation in the average number of non-production workers and assuming a 40-hour working week and 50 weeks, while we estimate total hours as a function of production worker hours and the ratio of the total wage bill to the wages of production workers. The difference in capital growth is explained by a labeling issue 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.

in the ASM (see the rightmost columns in Panel A of table 6). Panel B lists the results of decomposing APG ignoring the first years of ASM panels. Each row corresponds to the contribution by a distinct component, calculated as an annual average between 1977 and 1996. The contribution of reallocation is shown in the row labeled as Total RE. Our estimate of the annual contribution by reallocation falls between 0.7-1.3 percentage points. These numbers imply that approximately 25-70% of annual APG is attributed to reallocation, depending on the chosen estimation method. The contribution of within-plant growth, shown in the row labeled as Within, is at most 1.4 percentage points, about 74% of total APG.³¹ While there is quantitative variation in these results all of the approaches yield a substantial contribution of reallocation to productivity growth.³²

We now turn to consider price setting behavior with markups. As Petrin and Levinsohn (2012) and PWR highlight, one possible source of gaps between the value of marginal products and marginal costs is markups. As they note, moving resources from low markup plants to

³¹Including the first years of ASM panels yield similar results for reallocation and somewhat lower contributions by within-plant growth.

³²The time series underlying table 6 indicate the contribution of reallocation is positive in the majority of years which reinforces what our annual average results suggest. In addition, the time series standard deviation of the total reallocation contribution is small relative to the within-plant productivity contribution.

high markup plants will increase APG through the reallocation components. Also, as is evident from equation (5), under CES demand, tfpr^{rr} is a function of technical efficiency but that term is scaled by ρ . This implies that the within plant productivity growth measured by changes in tfpr^{rr} needs to be adjusted for ρ .³³ The contribution of fixed costs (residual) also needs to be adjusted accordingly.

To explore these issues, we consider two approaches. First, to illustrate the impact of the scaling issue, we consider two common values of ρ : 0.9 and 0.8 corresponding to markups of about 10 percent and 25 percent.³⁴ Second, we consider the more interesting case where markups that vary across plants. We explore this issue using the tools introduced in section 4. Particularly, we can estimate ρ at the detailed industry-level and apply the appropriate adjustments. Table 7 shows the results using the estimation with such adjustments. The $\rho = 1$ column shows the results of using the OP method presuming price taking behavior and the other columns show the results of the decomposition using the alternative estimates of ρ and appropriate adjustments. As the table shows, such distinctions can be relevant empirically.

For common markups, the within and reallocation terms tend to be scaled up and the fixed costs term becomes less important the larger the markup. For the estimated markups case, reallocation is also scaled up, while the within term is significantly lower. The latter finding reflects the combined effect of the variation and interaction in estimated ρ 's and tfpr_i^{rr} growth rates. In general, incorporating the markup adjustment makes the within plant component more volatile over time. Moreover, with estimates of ρ that vary across industries, the increase in this volatility does not need to be symmetric across upturns and downturns. In unreported results, we find that the early 1980s downturn has an especially large decline in measured within plant productivity growth with the variable markups. This finding accounts for the large decline in the within plant productivity growth contribution under this approach. These findings on the within plant contribution with variable markups are interesting but caution should be used in this assessment since the estimates of ρ without price and quantity data are noisy. Put differently, this highlights the need for precise estimates of demand elasticities. In addition, measuring and considering plant-specific markups within industries would be interesting in this setting. In spite of this sensitivity, one robust implication of this analysis is that adjusting revenue elasticities to output elasticities increases the reallocation contribution.

³³See appendix A.5 for more details.

³⁴Basu, Pascali, Schiantarelli, and Serven (2010) suggest that with common markups it is not clear that the reallocation terms should be adjusted upwards but rather separate terms included in this type of decomposition. With a common markup, these extra terms would capture the fact that inputs and outputs are too low under imperfect competition so input growth adds to APG.

Table 7: The effect of controlling for demand characteristics in PWR’s decomposition of APG.

	$\rho = 1$	$\rho = 0.9$	$\rho = 0.8$	ρ =KG estimates
Total RE	0.7	0.9	1.1	1.4
Capital RE	0.2	0.3	0.3	0.3
Labor RE	0.0	-0.1	-0.1	0.0
Materials RE	0.4	0.5	0.7	0.9
Energy RE	0.1	0.2	0.2	0.2
Within	0.7	0.8	0.9	-2.7
Fixed costs	0.5	0.1	-0.1	3.1
σ_{RE}	0.5	1.0	1.8	1.3
σ_{Within}	2.8	3.1	3.5	5.7

Results are based on OP and using data between 1977-1996, ignoring first years of ASM panels. Entries in the column labeled ρ are identical to those in column 2 of Panel B in table 6. The remaining columns show results permitting markups. The KG estimates are the Klette-Griliches estimates under joint estimation of revenue elasticities and demand parameters, see section 4.

6 Concluding remarks

Researchers have been using a variety of methods for estimating productivity at the firm-level. Absent data on prices and quantities, these methods yield what have become known as revenue productivity measures. How these measures are related to physical productivity depends on the assumptions about the environment in which establishments operate. It is perhaps less recognized that the differences across estimation methods have important consequences for interpretation since the alternative measures are different conceptually. Cost-share-based coefficients are, in principle, equivalent to output elasticities, while regression-based estimates are revenue elasticities reflecting both output elasticities and demand parameters. This implies that revenue residuals are conceptually different under these two broad approaches. One measure, $tfpr^{cs}$, has increasingly become used as a measure of distortions. Alternatively, under the same assumptions that yield this interpretation of $tfpr^{cs}$, $tfpr^{rr}$ will reflect fundamentals such as TFPQ and demand shocks.

In spite of these conceptual differences, we find that these alternative measures (i) exhibit similar dispersion, (ii) are strongly positively correlated and (iii) exhibit very similar relationships with growth and survival. What do these results imply about interpreting $tfpr^{cs}$ as a measure of idiosyncratic distortions? To maintain that assumption, idiosyncratic distortions would need to exhibit similar dispersion and be highly correlated with idiosyncratic fundamentals. Given the tight relationship between $tfpr^{cs}$ and growth and survival, this poses the challenge of explaining why the most distorted firms are the most likely to grow and survive.

In developing these findings, we conduct an in-depth sensitivity analysis of the methods used to estimate $tfpr^{rr}$. We find the alternative estimation methods yield broadly similar results but there are some non-trivial quantitative differences. Interestingly, we find that the variation in patterns for $tfpr^{rr}$ are at least as great as the variation in patterns between $tfpr^{cs}$ and $tfpr^{rr}$. Partly this reflects the fact that the proxy methods use higher order polynomials. Our findings suggest these methods are more robust with large sample sizes. This pattern is consistent with the observation that many researchers using these methods pool data across plants within a 2 or 3-digit level of industry aggregation in the estimation.

We also use this in-depth analysis of sensitivity to estimation methodologies to quantify the sensitivity of the contribution of reallocation to productivity growth to alternative approaches. We use the structural decompositions that have been developed recently for this purpose. This

is an ideal setting for exploring sensitivity to alternative methods for revenue and output elasticities because the reallocation terms in the structural decompositions depend critically on these estimates. With these methods, the gap between the value of the marginal product and marginal costs depend on the estimates of such elasticities. Implementation of these decompositions in the literature has typically not taken into account the distinction between revenue and output elasticities. We find that adjusting the revenue elasticities so that the appropriate output elasticities are used tends to increase the contribution of the reallocation component. Moreover, while there is non-trivial quantitative variation in the implied contribution of reallocation to aggregate productivity growth across methods, we find all methods yield a sizable contribution.

In sum, it is important to understand when the devil is in the details. One potential remaining devil in the details is the impact of heterogeneous and endogenous plant-level product and input prices. The results from our (admittedly restrictive) demand analysis can be used to make inferences about output elasticities and returns to scale; however, without plant-level data on prices and/or quantities, the effects of prices on other key stylized facts are difficult to quantify. 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 measures of prices and quantities (for the limited number of products with such information) with methods that impose strong functional form assumptions (e.g., isoelastic demand structures) to deal with these issues. We also neglect the impact of input price heterogeneity which may be another source of dispersion in micro-level productivity measures.

References

- Akerberg, Daniel A., Kevin Caves, and Garth Frazer. Identification Properties of Recent Production Function Estimators. *Econometrica*, 83:2411–2451, November 2015.
- Asker, John, Allan Collard-Wexler, and Jan De Loecker. Dynamic Inputs and Resource (Mis)Allocation. *Journal of Political Economy*, 122(5):1013 – 1063, 2014.
- 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.
- Basu, Susanto, Luigi Pascali, Fabio Schiantarelli, and Luis Servén. Productivity, Welfare and Reallocation: Theory and Firm-Level Evidence. Technical Report 728, 2010.
- Cooper, Russell W. and John C. Haltiwanger. On the Nature of Capital Adjustment Costs. *Review of Economic Studies*, 73(3):611–633, 2006.
- Cooper, Russell W., John C. Haltiwanger, and Jonathan Wiliis. Search frictions: Matching aggregate and establishment observations. *Journal of Monetary Economics*, 54:56–78, 2007.
- De Loecker, Jan. Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity. *Econometrica*, 79(5):1407–1451, 09 2011.

- Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda. Changing business dynamics and productivity: Shocks vs. responsiveness. 2017.
- Elsby, Michael and Ryan Michaels. Marginal Jobs, Heterogeneous Firms, and Unemployment Flows. *American Economic Journal: Macroeconomics*, 5(1):1–48, 2013.
- Eslava, John Haltiwanger Adriana Kugler Maurice Kugler, Marcela. Trade and market selection: Evidence from manufacturing plants in colombia. *Review of Economic Dynamics*, 16(1):135–158, 2013.
- Foster, Lucia, Cheryl Grim, and John Haltiwanger. Reallocation in the Great Recession: Cleansing or Not? *Journal of Labor Economics*, 34, 2016a.
- 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 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, John C. Haltiwanger, and Zoltan Wolf. Firm-Level Dispersion in Productivity: Is the Devil in the Details? *American Economic Review*, 106(5):95–98, May 2016b.
- Foster, Lucia S., John C. Haltiwanger, and Chad Syverson. The slow growth of new plants: Learning about demand? *Economica*, 83(1):91–129, 2016c.
- Gopinath, Gita, Sebnem Kalemli-Ozcan, Loukas Karabarbounis, and Carolina Villegas-Sanchez. Capital allocation and productivity in south europe. *NBER Working Paper Series*, (21453), 2015.
- Haltiwanger, John. Firm Dynamics and Productivity: TFPQ, TFPR, and Demand Side Factors. *Economia*, 0:3–26, Fall 2016.
- Haltiwanger, John, Robert Kulick, and Chad Syverson. Misallocation measures: Glowing like the metal on the edge of a knife. 2017.
- Hopenhayn, Hugo A. Entry, Exit, and Firm Dynamics in Long Run Equilibrium. *Econometrica*, 60(5):1127–50, September 1992.
- Hopenhayn, Hugo A. and Richard Rogerson. Job turnover and policy evaluation: A general equilibrium analysis. *Journal of Political Economy*, 101(5):915–938, 1993.
- 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.
- 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.
- 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.
- Restuccia, Diego and Richard Rogerson. Policy distortions and aggregate productivity with heterogeneous establishments. *Review of Economic Dynamics*, 11(4):707–720, October 2008.
- 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.
- 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.

A Appendix

A.1 tfpr^{rr} and fundamentals

We derive the properties of tfpr^{rr} under the more general specification of demand given by $P_i Q_i = P_s Q_s^{1-\rho} Q_i^{\rho-1} \xi_i Q_i = P_s Q_s^{1-\rho} Q_i^{\rho} \xi_i$. In this case, log plant level revenues can be written as

$$p_i + q_i = \rho q_i + (1 - \rho) q_s + p_s + \ln \xi_i = \rho \left(\sum_j \alpha_j x_{ij} + a_i \right) + (1 - \rho) q_s + p_s + \ln \xi_i.$$

This permits characterizing tfpr_i^{rr} as:

$$p_i + q_i - \rho \sum_j \alpha_j x_{ij} = p_i + q_i - \sum_j \beta_j x_{ij} = \rho a_i + \ln \xi_i + (1 - \rho) q_s + p_s,$$

which says that tfpr_i^{rr} is a function of tfpq_i (a_i), demand shocks (ξ_i) and sectoral level factors (q_s, p_s and ρ). In the main text, we abstract from idiosyncratic demand shocks and sectoral level factors for transparency. We estimate the β_j using regression (proxy) methods. In some robustness analysis, we also use the Klette and Griliches (1996) approach to jointly estimate β_j and ρ by including a measure of industry-level output as a regressor. This permits us to back out the α_j from the combined estimates and provides an alternative method to cost shares for estimating α_j . The advantage of this approach is that it does not impose CRTS. The disadvantage of this approach is that in the absence of data on plant level prices and quantities, this is pushing the data quite hard. Foster, Grim, Haltiwanger, and Wolf (2016b) discuss the latter limitations in more depth.

A.2 Remarks on proxy methods

The details of the procedures described in Olley and Pakes (1996) and Levinsohn and Petrin (2003) have been discussed in the literature, see for example Akerberg, Caves, and Frazer (2015). Caveats notwithstanding, these procedures have become widely used among practitioners in the past decade.

Most proxy methods are multi-step procedures (the exception being the Wooldridge GMM procedure as discussed in the main text), where the first-step-estimates of variable input elasticities are determined by OLS and the coefficients of quasi-fixed inputs such as capital are separately identified using moment conditions formed by estimated productivity innovations and lagged capital values. These methods rely on polynomial approximations at two points of the estimation algorithm. First, a polynomial of the state variables and the proxy is included to approximate unobserved productivity. Second, the algorithm exploits a Markovian assumption about the plant-level productivity process in order to extract productivity innovations. It is important to note that while polynomial series approximations are flexible, higher order terms may exacerbate measurement error likely present in microdata.

The available empirical implementations differ according to the production function type and the numerical procedure used to minimize the objective function. For example, the Stata implementation of OP's approach is based on nonlinear least squares; while LP's implementation offers three options to minimize the GMM criterion function. If the dependent variable is value added then the golden section search algorithm (GSS) is applied.³⁵ For output-based

³⁵The idea of GSS is successively narrowing the range of values inside which the extremum is known to exist.

specifications, either a gradient-based or a grid search based routine may be chosen. Both the gradient-based and the golden section search guarantee to find optimum points if the objective function is unimodal. If the criterion function has multiple modes, grid-based search procedures can be used to confirm global optimum but often at the cost of computational burden, especially if one intends to estimate elasticities for many industries.³⁶ Our results in section 4 suggest these choices entail non-trivial differences in the distributions of elasticities.

Table A1: Variants of LP.

Method	Description	Dependent variable	Proxy	Estimator	Numerical procedure*
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

*NL: gradient-based technique, GSS: Golden Section Search, GR: Grid Search.

A.3 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 A2 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 130,000 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 mappings in available concordances.

A.4 Empirical properties of estimators

The distributions discussed in section 4 are based on elasticities from the 50 most populous 4-digit industries. While 50 observations may seem sufficient to estimate cross-industry elasticity

The name of the algorithm sources from the fact that the procedure maintains the function values for triples of points whose distances form a golden ratio.

³⁶Estimating the elasticities by grid-search for all 459 industries in section 5.4 took about 28 days given our computing resources.

Table A2: Descriptive statistics on industry assignment.

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.			
Original industry	Assigned 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.

distributions, it also means estimates are based on varying sample size. This is important because the estimates from smaller industries are more likely to be less precise. Table A3 shows that the sample size drops by more than 50% in the 10th, and by 80% in the 50th most populous industry. What is the consequence of such variation? Are some methods more likely to yield non-positive $\hat{\beta}$ -s than others?

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(1000)	\bar{T}_i	N_j
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
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.

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.³⁷ All methods result in positive $\hat{\beta}_l$ and $\hat{\beta}_m$ in most industries (table A4, columns 6 and 9). Negative β_k and β_e estimates are generally more likely to occur, as shown by 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 other methods yield negative $\hat{\beta}_k$ with a positive probability. 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 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 is to make stronger assumptions of homogeneity – for example, assuming plants within 3-digit SIC industries share the same output elasticities. We explore this approach in section 5.2 and show that this approach leads to more plausible elasticity estimates and implied firm dynamics in many cases.

³⁷The sample for this exercise is described in more detail at the end of section 3.2.

Table A4: The number of 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, see section 3 for more details.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	$\hat{\beta}_k$			$\hat{\beta}_l$			$\hat{\beta}_m$			$\hat{\beta}_e$		
	-	0* or n.e.**	+	-	0 or n.e.	+	-	0 or n.e.	+	-	0 or n.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*: $\hat{\beta}$ is zero. n.e.**: $\hat{\beta}$ could not be estimated.

For example, the first entry in column (2) says OP delivers error in 18% of industries. This happens because the algorithm stops in industries with insufficient information on exit (or investment). The second entry in column (2) says gradient-based optimization in LPNL yields $\hat{\beta}_k = 0$ 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)
	$\hat{\beta}_k$		$\hat{\beta}_L$		$\hat{\beta}_M$		$\hat{\beta}_E$	
	-	0 or n.e.	-	0 or n.e.	-	0 or n.e.	-	0 or n.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 $\hat{\beta}_j$. See also the notes below table A4.

A.5 APG decomposition under downward sloping isoelastic demand

In this section we discuss the adjustments to the implementation of Petrin and Levinsohn (2012) and PWR’s APG decomposition required when using revenue function estimates of the revenue elasticities and the residual from this function under CES demand. APG is measured as the change in the value of final demand minus the value of the change in primary inputs. Specifically:

$$APG = \sum_i dVA_i - \sum_i \sum_k W_{ik} d\ln X_{ik}$$

where VA_i is value added, X_{ik} are the primary inputs and W_{ik} are the factor prices of the primary inputs. The APG decomposition is given by (using the notation from section 2):

$$\sum_i D_i d\ln a_i + \sum_i D_i \sum_k (\alpha_k - s_{ik}^{\text{rev}}) d\ln X_{ik} + \sum_i D_i \sum_j (\alpha_j - s_{ij}^{\text{rev}}) d\ln M_{ij} - \sum_i D_i d\ln F_i$$


where M_{ij} are the intermediate inputs, s_{ik} and s_{ij} are the revenue shares of the primary and intermediate inputs respectively, D_i are domar weights and F_i denote fixed costs. To implement this decomposition empirically requires measures of APG, the output elasticities (α_k), the domar weights, the revenue shares and the within plant growth in productivity ($d\ln a_i$). With these measures, the contribution of fixed costs to APG can be measured as a residual. Estimating the revenue function yields estimates of β_k which with CES demand imply that $\alpha_k = \beta_k/\rho$. Thus, the revenue elasticities need to be adjusted by the scaling factor ρ to recover the output elasticities. Moreover, the within plant growth in the revenue residual is given by $\rho d\ln a_i$ so that to recover the within plant growth in productivity the growth of the residual must be adjusted accordingly. These adjustments require estimates of ρ . This is the approach taken in the main text. We also note that in a more general specification that permitted demand shocks that further adjustments would be required depending on the properties of these demand shocks.

A.6 Industry comparisons of dispersion


Note that even if different methods yield similar average dispersion, industry-specific implications may differ. To highlight the implications of estimators in economic analyses where differences in magnitudes may matter a lot, we compare these measures across a selected set of estimators in figure A1. In panels (a)-(b), almost all points are located above the 45-degree line, which means that proxy methods (OP, LPGR) yield systematically larger dispersion than cost-share-based methods (GA). One explanation behind such divergence may be related to the conceptual differences mentioned earlier. On the other hand, comparing dispersion under LPGR and OP suggests that differences between two methods of the same class may also exist, see panel (c). These differences, smaller and less clear-cut than in the previous two panels, are still non-negligible. They can be attributed to technical aspects like the chosen econometric procedure, and/or more substantive factors like the choice of proxy.

A.7 Robustness


We next examine the robustness of our results about productivity dispersion and growth and survival to concerns about imputation methods used in the underlying microdata and our assumption that elasticities are homogeneous within industries. In contrast to previous exer-



(a) OP and GA



(b) LPGR and GA



(c) LPGR and OP

Figure A1: Industry-specific dispersion under selected methods.

Each point shows a pair of industry-specific dispersion measures and the line represents the 45-degree line in each chart.

cises, our approach here is to compare the same productivity measure across completed and non-imputed samples rather than different productivity measures in the same sample.

A.7.1 Imputation and plant characteristics

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 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 presence or direction of possible biases is less clear for the relationship between productivity, growth and survival 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 productivity, see table A6 for more details.³⁸ 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.

Table A6: 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.

	tv	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 typically based on univariate regressions, while ASM imputes are much more likely to be based multivariate regressions.

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

³⁸We do not discuss individual imputation models further but more details are available upon request.

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 productivity 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 productivity 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 define 10 size classes based on employment (1-9, 10-109, 20-29, 30-49, 50-99, 100-149, 150-249, 250-499, 500-999, 1000+) in addition to the payroll deciles. Finally, we classify establishments into 9 age classes (births, 1, 2, 3, 4, 5, 6-10, 11-15, 16+ years). As mentioned in the main text, data for these variables source from the LBD. We use the Census Bureau's impute flags to determine whether an item is imputed. More details on imputation procedures at the Census Bureau are available from the authors. See also table A1 in WPR.

Imputation rates differ across variables. The upper panel of table A7 shows that from among the main components of productivity, PH and CP are imputed the most, and SW the least frequently. Imputation rates vary not only with variables but also sample definition.³⁹ 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 productivity (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.⁴⁰ Productivity 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

³⁹We 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).

⁴⁰These results are not shown here.

Table A7: Imputation rates of the main components of productivity in the ASM/CM, as a per cent of plant-year observations between 2002-2010.

Average imputation rates (2002-2010)							
	productivity	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 productivity, 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

Productivity 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.

provides a somewhat better fit⁴¹ than any of the remaining five. Therefore, we present results based on propensity scores from a model with establishment size, payroll and age.

Figure A5 plots the point estimates from this model. We conclude that productivity data for smaller and younger plants with less payroll are significantly more likely to be imputed.⁴² In section A.7.2, 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.⁴³ The difference in the AIC is about a factor of 3 and 2 in 2002 and 2007, respectively.

A.7.2 Imputation, dispersion and growth

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 characteristics. We take this approach below and attempt to correct for selection bias by using inverse propensity score weights discussed in section A.7.1. 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.

We carry out two exercises to assess the effect of imputation. First, we compare dispersion in the completed sample to statistics which are based on non-imputed observations only. Second,

⁴¹As measured by the AIC, not shown here.

⁴²Point estimates are precise enough such that we can confirm the positive relationship between plant-size, -age and the probability of productivity being non-imputed. More details on point estimates and standard errors are available upon request.

⁴³Available upon request.

we assess whether accounting for imputation affects our results on the relationship between productivity, growth and survival. We address imputation by first constructing a set of propensity score weights, which are inversely proportional to the probability of imputation and then using these weights to calculate weighted dispersion measures. Note that we use another set of weights to control for the fact that selection into the ASM is also non-random. We distinguish between the two sets by labeling non-impute weights as $ipw2$ and the ASM-weights as $ipw1$. Appendix A.7.1 describes how $ipw2$ was constructed. More details about $ipw1$ can be found in Foster, Grim, and Haltiwanger (2016a).

Figure A2 summarizes the results of the first exercise. We find that the interquartile range seems to be smaller in non-imputed data. This can be seen in panel (a) of A2 by comparing the dotted and solid lines. However, if we weight non-imputed observations by the composite weight $[ipw1*ipw2]$, measured dispersion is higher.⁴⁴ 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.⁴⁵ If imputation causes a downward bias in dispersion, our weighting scheme corrects for it in the right direction. We note that imputation has similar effects on the standard deviation but this measure seems to be less sensitive to these issues (see panel b of figure A2).

In the second exercise, we revisit the relationship between productivity, growth and survival and re-estimate the growth and exit regressions using non-imputed data between 2002 and 2010. The empirical model is the same as that described in section 5.2. We consider three samples. In the first one, we use all observations from the CM.⁴⁶ 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 $[ipw1*ipw2]$. Table A8 summarizes our results. The effect of productivity on growth among all establishments (column 1) and on the probability of exit (column 2) seems smaller 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) 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 also familiar from earlier tables. 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.

A.7.3 Homogeneity, dispersion and growth


Throughout this paper, we assume that elasticities are homogeneous within industries and constant over time. Whether or not such an assumption is restrictive depends on at least two properties: the underlying within-industry differences in technology and sample size.⁴⁷ If there are within-industry differences in technology then pooling the data from an entire industry of establishments may be too restrictive. In this case, allowing for plant-level heterogeneity in the elasticities, empirical feasibility aside, better accounts for within-industry differences in factor intensities. On the other hand, pooling data may be necessary to increase sample size in order to

⁴⁴Similar patterns hold for the 90-10 ratio, not shown here.


⁴⁵See columns 1 and 2 of table 4 in their paper.

⁴⁶Note 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 $ipw1$.

⁴⁷It is not straightforward to test for the true degree of heterogeneity because the results on which we base inference about heterogeneity are endogenous to both the estimation method and the homogeneity assumption.



(a) Inter-quartile range



(b) Standard deviation

Figure A2: Time series of average within-industry dispersion measures in various samples, ASM 2002-2010, all industries.

Solid and dotted lines denote statistics which are calculated using completed and non-imputed data, respectively. In the completed sample, weighted statistics are based on weights which account for the fact that selection into the ASM and CM is non-random (*ipw1*). 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 productivity is calculated using non-imputed data (*ipw2*).

Table A8: The effect of productivity on outcomes between 2002-2010 in the ASM/CM, various samples within the ASM/CM, growth accounting based productivity.

sample	overall growth (1)	exit (2)	conditional growth (3)
ASM-CM			
Completed [ipw1] [†]	0.165***	-0.065***	0.04***
Non-imputed [ipw1] ^{††}	0.083***	-0.019***	0.047***
Non-imputed [ipw1*ipw2] ^{†††}	0.079***	-0.021***	0.041***
ASM cases only			
Completed [ipw1]	0.149***	-0.06***	0.035***
Non-imputed [ipw1]	0.086***	-0.022***	0.044***
Non-imputed [ipw1*ipw2]	0.086***	-0.024***	0.041***
Sample size (in thousands)			
ASM-CM			
Completed data [ipw1]	594	594	570
Non-imputed data [ipw1]	263	263	258
Non-imputed data [ipw1*ipw2]	263	263	258
ASM cases only			
Completed data [ipw1]	400	400	386
Non-imputed data [ipw1]	218	218	214
Non-imputed data [ipw1*ipw2]	218	218	213

Triple-asterisks denote 1% significance levels. [†]Weighted Census data [ipw1]: 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 [ipw1]: non-imputed subset of the same sample. ^{†††}Non-imputed [ipw1*ipw2]: non-imputed subset of the sample, where observations are weighted by a composite propensity score, where *ipw2* is inversely proportional to the probability that a plant's productivity is calculated using non-imputed data. Probabilities were estimated separately for each year and are based on industry-, size-, age-class and payroll-decile fixed effects.

reduce finite-sample bias and increase precision.⁴⁸ 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 micro data.

In this section, we relax the homogeneity assumption and assess the consequences for productivity dispersion and growth and exit regressions. Figure A3 shows that dispersion is substantially higher if we allow cost-shares to vary across establishments compared to when it is constant within an industry. In particular, plant-specific shares increase the interquartile range by a factor of almost two (see the difference between the thin and thick lines in the figure). Results also indicate that allowing for time varying elasticities affects only the volatility of dispersion but not its level. In figure A3 this is shown by observing the difference between the thin dashed and solid lines. There is no significant difference in variation with industry-level shares implying the overall conclusion that the effect of time series smoothing is dwarfed by that of cross-section smoothing.

Both theory and earlier research⁴⁹ indicate that part of this increase in dispersion may be spurious. Our own analysis also offers indirect evidence that the increased variation is at least partly noise. Table A9 shows results from growth and exit regressions based on growth

⁴⁸Moreover, for some of the estimators reviewed in this paper whether one pools data from different industries is not a matter of bias and precision but feasibility. See sections 4 and 5.2 for examples.

⁴⁹See, for example, Syverson (2004).

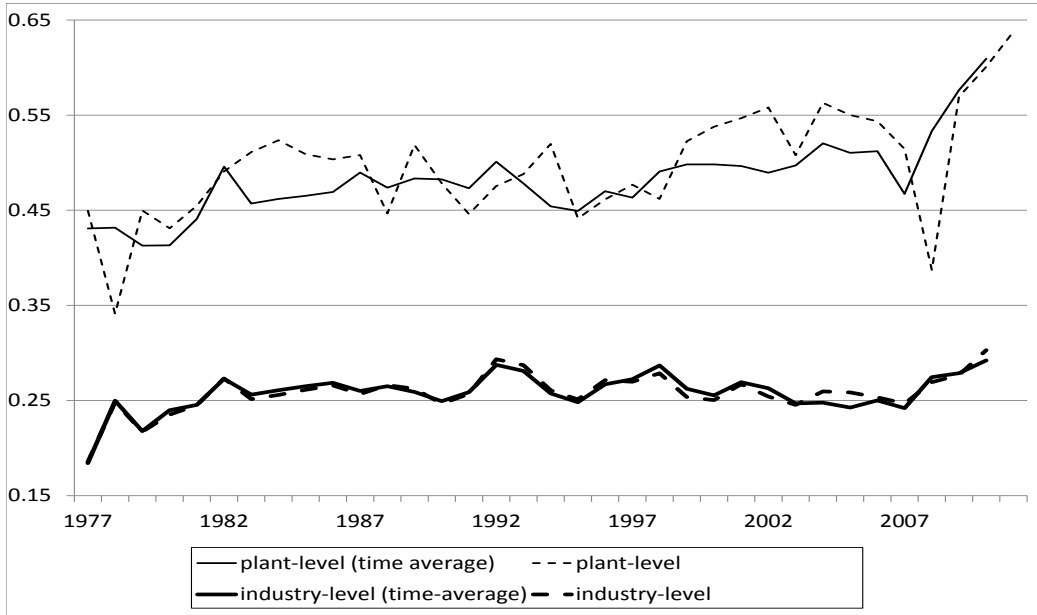


Figure A3: Within-industry interquartile range measures of growth accounting based log-productivity. Sample: 10 largest industries. Thin and thick lines denote dispersion measures based on shares calculated at the establishment- and industry-levels, using the growth accounting framework. Solid lines denote dispersion measures based on cost shares that are constant over time. Dashed lines denote statistics calculated using time-varying cost shares.

accounting productivity estimates. Comparing columns 1-2 shows that using time-varying industry-level 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 3) 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 3-4). 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 (column 3 relative to 4 and 5) suggesting there may be noise not only in the cross-section but also in the time series variation of establishment-specific cost shares. 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 A9: The effect of growth accounting productivity variants on outcomes, 50 largest industries.

	Industry-level shares		Plant-level shares		
	(1)	(2)	(3)	(4)	(5)
	constant	time-varying	constant	time-varying	time-varying
	$[\bar{S}^j]$	$[\frac{s_t^j + s_{t-1}^j}{2}]$	$[\bar{S}_i^j]$	$[S_{it}^j]$	$[\frac{s_{it}^j + s_{it-1}^j}{2}]$
overall growth	0.190***	0.193***	0.114***	0.1***	0.096***
exit	-0.064***	-0.066***	-0.042***	-0.041***	-0.039***
conditional growth	0.067***	0.067***	0.032***	0.018***	0.018***

Outcomes are: employment growth among all establishments (rows 1-2), exit (rows 3-4), employment growth among continuers (rows 5-6). 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.

A.8 Additional Tables

Table A10: 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

See notes to table 1 for legends. The first entry in the table says that half of the industries in the lowest quintile under LPGR (row) are also classified in the lowest quintile under OP (column). All other entries are analogous.

Table A11: Descriptive statistics of parameter estimates under isoelastic demand.

Panel A. Demand parameters.								
	$\hat{\beta}_q$		$\hat{\rho} = 1 - \hat{\beta}_q$		$\varepsilon_{P_i}^{Q_i} = \frac{1}{\rho-1} = -\frac{1}{\hat{\beta}_q}$		$\frac{1}{\rho} = \frac{1}{1-\hat{\beta}_q}$	
	OLS	OP	OLS	OP	OLS	OP	OLS	OP
mean	0.09	0.05	0.91	0.95	-11.83	-19.53	1.2	1.11
stdev	0.19	0.16	0.19	0.16	46.26	126.2	0.51	0.45
Panel B. Revenue elasticities: $\hat{\beta}_j$								
	$\hat{\beta}_K$		$\hat{\beta}_L$		$\hat{\beta}_M$		$\hat{\beta}_E$	
	OLS	OP	OLS	OP	OLS	OP	OLS	OP
mean	0.1	0.11	0.29	0.27	0.5	0.49	0.11	0.1
stdev	0.06	0.09	0.1	0.09	0.13	0.13	0.05	0.05
Panel C. output elasticities: $\hat{\alpha}_j = \hat{\beta}_j/\hat{\rho} = \hat{\beta}_j/(1 - \hat{\beta}_q)$								
	$\hat{\alpha}_K$		$\hat{\alpha}_L$		$\hat{\alpha}_M$		$\hat{\alpha}_E$	
	OLS	OP	OLS	OP	OLS	OP	OLS	OP
mean	0.13	0.13	0.35	0.31	0.58	0.54	0.14	0.11
stdev	0.12	0.14	0.2	0.2	0.23	0.23	0.12	0.07

We assume $P_i = P(Q/Q_i)^{1-\rho}$ in the estimation. The parameters related to demand exhibit large standard deviation which is a sign that these parameters are poorly estimated in a non-trivial number of industries (Panel A). The elasticity of plant-level revenues with respect to industry-revenues ($\hat{\beta}_q$) is on average a small positive number, but varies a lot across industries. Consequently, $\hat{\rho}$ and the demand elasticity ($\varepsilon_{P_i}^{Q_i}$) and the markup ($1/\hat{\rho}$), also exhibit large variation. $\hat{\alpha}_j$ -s are greater than $\hat{\beta}_j$ -s in the average industry (see Panels B and C), which is to be expected under isoelastic demand. However, $\text{sd}(\hat{\alpha}_j) > \text{sd}(\hat{\beta}_j)$. Overall, these patterns suggests that although the variation in output elasticities reflects $\text{var}(\hat{\beta}_j)$ and $\text{var}(\hat{\rho})$, the contribution of the latter likely dominates.

Table A12: Descriptive statistics of productivity distributions. Estimators are described in table 1, pre-estimation outliers included.

	N (1000)	IQR	SD	N (1000)	IQR	SD
	50 largest industries			10 largest industries		
OLS	568	0.29	0.28	230	0.23	0.21
OP	567	0.35	0.41	233	0.34	0.41
LPGR	572	0.33	0.35	235	0.32	0.31
WLPE	575	0.37	0.41	236	0.34	0.30
WLPM	575	0.46	2.21	236	0.37	0.39
GA	561	0.25	0.24	231	0.25	0.22

Table A13: Correlations among within-industry tfp-distributions, 50 largest 4-digit industries. Pre-estimation outliers included.

	OLS	OP	LPGR	WLPE	WLPM	GA
Pearson						
OLS	1					
OP	0.52	1				
LPGR	0.8	0.42	1			
WLPE	0.49	0.47	0.5	1		
WLPM	0.01	0.02	-0.15	0.15	1	
GA	0.77	0.47	0.63	0.49	0.09	1
Spearman						
OLS	1					
OP	0.67	1				
LPGR	0.86	0.6	1			
WLPE	0.58	0.61	0.55	1		
WLPM	0.34	0.36	0.25	0.48	1	
GA	0.8	0.62	0.67	0.58	0.42	1

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

Table A14: Correlations among within-industry tfp-distributions, 10 largest 4-digit industries.

	OLS	OP	LPGR	WLPE	WLPM	GA
Pearson						
OLS	1					
OP	0.48	1				
LPGR	0.71	0.43	1			
WLPE	0.71	0.54	0.8	1		
WLPM	0.52	0.27	0.41	0.66	1	
GA	0.88	0.42	0.68	0.66	0.5	1
Spearman						
OP	0.71	1				
LPGR	0.71	0.65	1			
WLPE	0.71	0.73	0.79	1		
WLPM	0.6	0.49	0.51	0.69	1	
GA	0.88	0.63	0.68	0.67	0.56	1

Including pre-estimation outliers barely changes correlations.

Table A15: Correlations among within-industry tfp-distributions, 10 largest 4-digit industries. Pre-estimation outliers included.

	OLS	OP	LPGR	WLPE	WLPM	GA
Pearson						
OLS	1					
OP	0.51	1				
LPGR	0.69	0.43	1			
WLPE	0.69	0.57	0.78	1		
WLPM	0.5	0.27	0.37	0.64	1	
GA	0.85	0.43	0.64	0.63	0.48	1
Spearman						
OLS	1					
OP	0.7	1				
LPGR	0.71	0.63	1			
WLPE	0.7	0.73	0.78	1		
WLPM	0.58	0.48	0.47	0.67	1	
GA	0.86	0.61	0.66	0.65	0.54	1

This table differs from table A14 in that we computed correlations including pre-estimation outlier observations.

Table A16: Monte Carlo results on productivity dispersion. Dispersion is measured using the interquartile range. Standard errors of iqr measures are bootstrapped and are in parentheses.

	Panel 1. Pooled data			Panel 2. Unweighted cross-industry average			
	orig. smpl	R=50	R=100	orig. smpl	R=50	R=100	
OLS	0.22	0.22 (0.0004)	0.22 (0.0004)	OLS	0.26	0.26 (0.0009)	0.26 (0.0008)
OP	0.32	0.78 (0.1014)	0.68 (0.0933)	OP	0.36	0.39 (0.0145)	0.39 (0.0203)
LPGR	0.29	0.60 (0.0706)	0.65 (0.0691)	LPGR	0.36	0.38 (0.0099)	0.38 (0.0111)
WLPE	0.34	0.75 (0.0794)	0.75 (0.0817)	WLPE	0.58	1.02 (0.4614)	1.05 (0.5588)
WLPM	0.42	1.82 (0.3788)	1.80 (0.3449)	WLPM	1.37	1.94 (1.1976)	1.95 (1.0269)
GA	0.23	0.23 (0.001)	0.23 (0.001)	GA	0.26	0.26 (0.0011)	0.26 (0.0011)

The moments of dispersion statistics in Panel 1 are calculated pooling data from all industries. Mean dispersion calculated in this manner can be interpreted as a weighted average across industries where the weights depend on the number of plant-year observations in an industry. This procedure is consistent with the way we calculate dispersion statistics in table 3. In Panel 2, dispersion is calculated as an unweighted cross-industry average of within-industry dispersion measures. Interestingly if we bootstrap dispersion statistics industry-by-industry and then calculate average dispersion (Panel 2), simulated results under proxy methods are closer to those in the original sample. Our MC approach implies that in this setup, larger and more heterogenous industries have smaller influence. The only exceptions are WLPM and WLPE, which are more likely to yield extreme elasticity estimates (see section 4) and therefore dispersion. For GA, coefficients are held constant while input and output observations vary across bootstrap replications. This exercise helps us assess the effect of the variation in input and output observations. It is natural to use GA coefficients for this exercise because the cost-shares are calculated using industry level information.

Table A17: 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	LPGR	WLPE	WLPM	GA
4-digit elasticities						
$dlnE$	410	415	413	414	414	405
exit	410	415	413	414	414	405
$dlnE_{cont}$	393	398	396	397	397	388
3-digit elasticities						
$dlnE$	405	413	414	413	413	405
exit	405	413	414	413	413	405
$dlnE_{cont}$	389	396	398	396	396	388
3-digit elasticities, industries with negative or non-estimable elasticities are dropped						
$dlnE$	391	406	408	289	310	405
exit	391	406	408	289	310	405
$dlnE_{cont}$	375	390	391	277	298	388
Total N	440	440	440	440	440	440

Total N denotes the original number of plant-year observations before estimation and post-estimation outlier trimming.

Table A18: The effect of productivity on outcomes, all (292) largest industries. Outcomes are: employment growth among all establishments, exit and employment growth among continuers

	OLS	OP	LPGR	WLPE	WLPM	GA
4-digit elasticities						
$dlnE$	0.183***	0.001	0.128***	0.004**	-0.004***	0.203***
exit	-0.066***	0.001	-0.041***	-0.002*	0.001**	-0.068***
$dlnE_{cont}$	0.055***	0.002***	0.05***	0.001	-0.001*	0.071***
3-digit elasticities						
$dlnE$	0.202***	0.115***	0.15***	0.001	-0.012***	0.193***
exit	-0.072***	-0.037***	-0.045***	0.000	0.004**	-0.066***
$dlnE_{cont}$	0.062***	0.046***	0.067***	0.000	-0.005***	0.066***
3-digit elasticities, industries with negative or non-estimable elasticities are dropped						
$dlnE$	0.207***	0.159***	0.169***	0.078***	0.06***	0.193***
exit	-0.073***	-0.05***	-0.05***	-0.031***	-0.022***	-0.066***
$dlnE_{cont}$	0.065***	0.065***	0.075***	0.017***	0.017***	0.066***

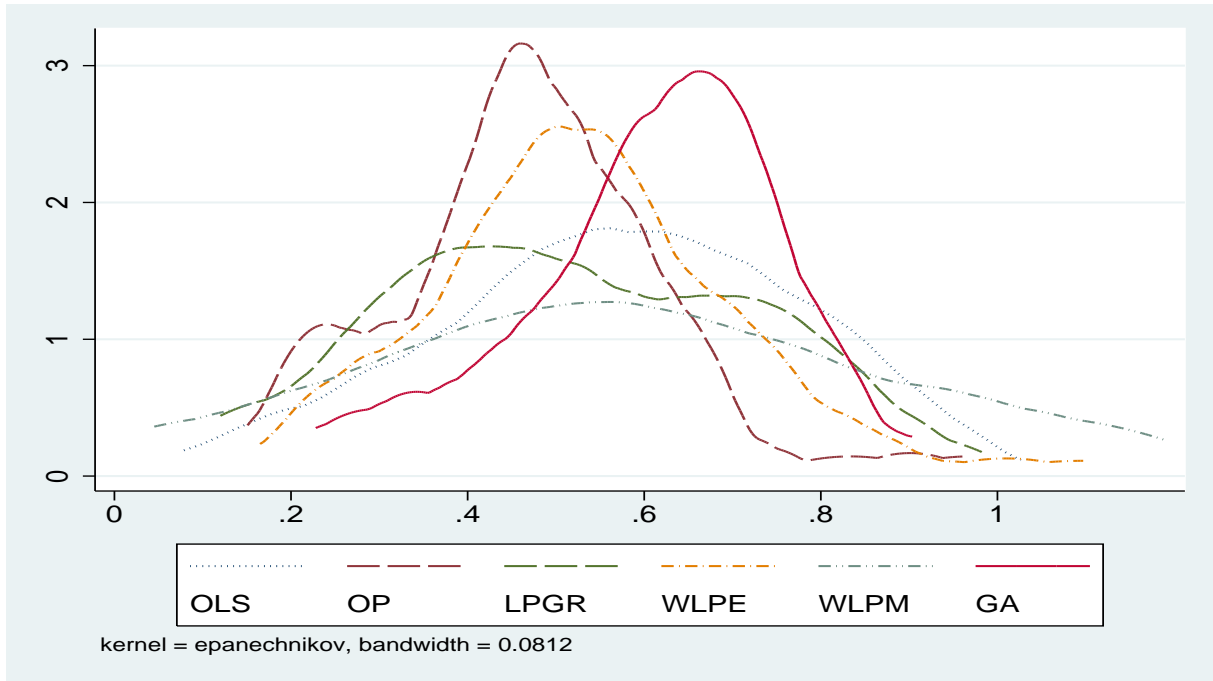
Estimates are taken from equations of three outcomes on plant-level measure of productivity, state-level measure of unemployment growth, and year-, sizeclass- and state-fixed effects. Standard errors are clustered at the state level, asterisks denote standard significance levels.

Table A19: The effect of productivity on outcomes, 10 largest industries. Outcomes are: employment growth among all establishments, exit and employment growth among continuers


	OLS	OP	LPGR	WLPE	WLPM
4-digit elasticities					
$dlnE$	0.206***	0.051***	0.154***	0.099***	-0.027**
exit	-0.067***	-0.002	-0.046***	-0.022***	0.014***
$dlnE_{cont}$	0.076***	0.054***	0.066***	0.06***	0.003
3-digit elasticities					
$dlnE$	0.208***	0.158***	0.164***	0.109***	0.016
exit	-0.068***	-0.032***	-0.049***	-0.026***	0.002
$dlnE_{cont}$	0.075***	0.102***	0.07***	0.061***	0.023*
3-digit elasticities, industries with negative or non-estimable elasticities are dropped					
$dlnE$	0.208***	0.158***	0.164***	0.13***	0.061***
exit	-0.068***	-0.032***	-0.049***	-0.043***	-0.018***
$dlnE_{cont}$	0.075***	0.102***	0.07***	0.046***	0.027*

Estimates are taken from equations of three outcomes on plant-level measure of productivity, state-level measure of unemployment growth, and year-, sizeclass- and state-fixed effects. Standard errors are clustered at the state level. *, ** and *** denote 10%, 5% and 1% significance levels.

A.9 Additional Figures



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



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

Figure A4: Cross-industry distributions of elasticities of $\hat{\beta}_m$ and $\hat{\beta}_e$. The sample is 50 largest industries, estimators are described in table 1. $\hat{\beta}_m$ -s show similar general shape but differences exist. For $\hat{\beta}_e$, results are also more similar except GA and WLPE.

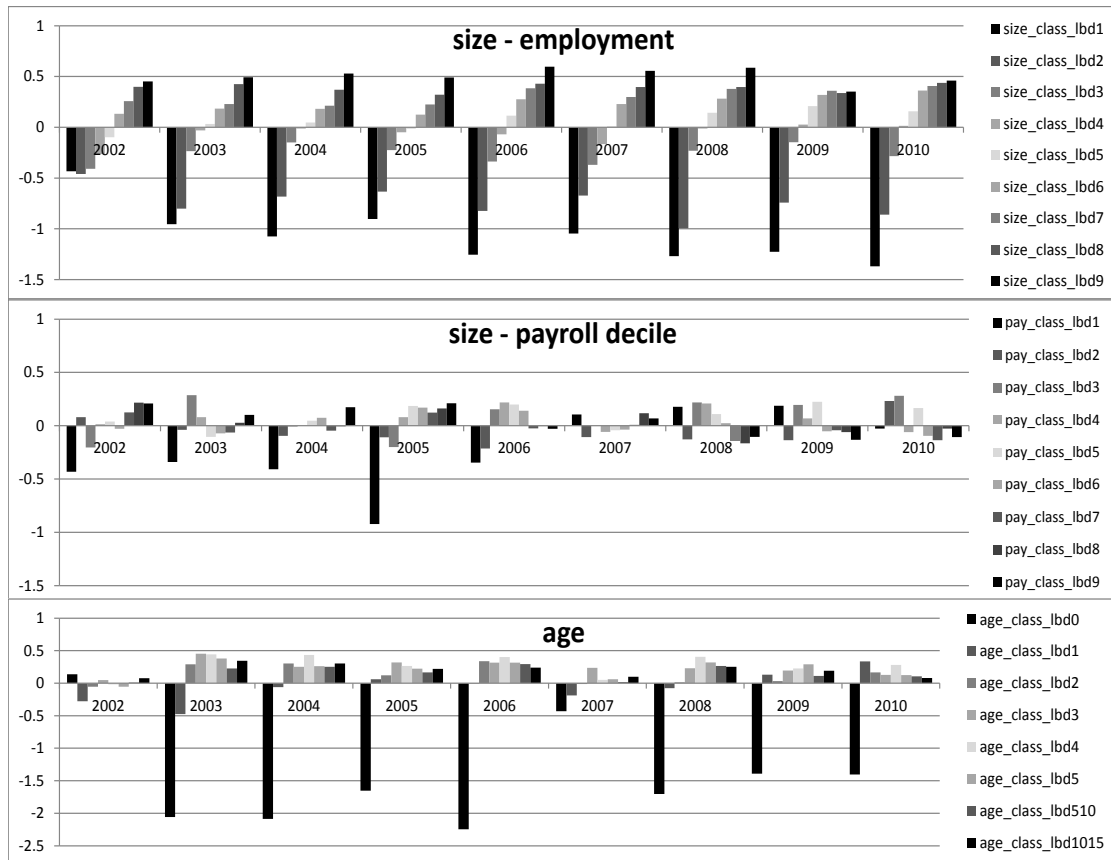


Figure A5: Point estimates from a logistic regression of the probability of productivity being non-imputed.

Regressors: 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).