Production Function Estimation Controlling for Endogenous Productivity Disruptions

Plutarchos Sakellaris* and Dimitris Zaverdas†

Athens University of Economics and Business

May 19, 2022

Abstract

We modify the standard production function estimation framework to incorporate endogenous disruptions in the productivity process due to lumpy firm investment. We investigate the differences between the existing (baseline) approach and our own (disruption) on a large proprietary panel of Greek Manufacturing firms. We find significantly different production function estimates and different results for subsequent inference. The implied average levels of productivity and magnitude of endogenous disruptions are different. The baseline model cannot capture the full dynamics of disruption costs, which extend across time. The decomposition of Aggregate Productivity Growth is substantially different under the disruption model.

Keywords: Matching, Investment spike, Production function estimation, Productivity

Acknowledgements

We thank the AUEB Laboratory of Econometrics and the Laboratory of Economic Policy Studies (EMOP) for their hospitality, the Hellenic Statistical Authority (ELSTAT) for providing us with 2-digit NACE rev. 2 Producer Price Indices for Greece, Vadym Volosovych for providing us with the NACE rev. 1 to (unique) NACE rev. 2 correspondence table used in Kalemli-Ozcan, Sorensen, Villegas-Sanchez, Volosovych, and Yesiltas (2015), and Dimitris Smyrnakis for illuminating us on the topic of threshold model regressions. We thank for their comments Michał Gradzewicz, Chad Syverson, and the participants of the RCEA Conference on Recent Developments in Economics, Econometrics and Finance 2022, and the PhD - Economics Virtual Seminar, which helped improve this manuscript.

*E-mail: plutarch@aueb.gr
Address: Athens University of Economics and Business, Derigny Wing, 4th Floor
Phone: +30 210 8203 353

†Zaverdas gratefully acknowledges financial support received for this research, in the form of a scholarship by the Academy of Athens.

1 Introduction

Studies on a variety of economic topics, such as productivity and markups, as well as applications that use structural models of firms, all rely on the accurate estimation of production functions. The current standard for production function estimation is to assume that productivity is completely exogenous. However, the literature studying lumpy adjustment has shown that investment spikes are associated with disruptions in productivity. In this paper, we propose a modification to the standard production function estimation framework that incorporates endogenous disruptions in the productivity process due to lumpy firm investment. We investigate the differences between our novel (disruption) framework and the existing (baseline) approach on a large proprietary panel of Greek Manufacturing firms.

We find that estimates of production function coefficients differ significantly. The disruption model estimates the capital coefficient to be lower and the labor coefficient and the returns to scale higher. This, in turn, affects the results of subsequent inference. In particular, the disruption models implies that the sample average level of productivity is 1.10% higher than that found by the baseline model. For the long-run average, the two models produce much more different results. The disruption model finds a long-run average that is 41.60% higher than that of the baseline.

The measured magnitude of endogenous disruptions is also different. We explore the different implications of the baseline and the disruption model contrasting a sub-sample of investment spike episodes with a matched sample that did not display lumpy investment. The sample matching is performed following one of the methods in Rosenbaum and Rubin (1985). We identify a spike observation as the case where the investment rate is above a large multiple of the expected investment rate conditional on current capital and a minimum threshold. Both models produce evidence of disruptions, in the year after the spike (t=1).

However, the baseline model gives an average drop in TFP that is approximately 7.07%, whereas the disruption model gives 5.75%. Two years on (at t=2), both models show that the average TFP of spike firms starts to recover towards that of the non-spike firms.

In addition, the baseline model cannot capture the full pattern of disruption costs. In contrast, the disruption model allows for a persistent effect of the disruption on productivity that extends across time and is found to be non-trivial. Ceteris paribus, the firm incurs an implicit adjustment cost of approximately 24.98% of its future output across time, in the form of forgone output. In practice, output does not fall by this entire amount, the negative effects of an investment spike output are counteracted by the increased levels of capital. This significant implicit extra adjustment cost that the firm has to bear when choosing to adjust in a lumpy fashion, rather than smoothly, is not limited to one period, and is completely missed by the baseline model.

We are, also, the first to obtain measures of the individual components of TFP, an observable and an unobservable (by the firm) component, and repeat the above analysis per component. We find that disruptions affect the firm through both parts of TFP. The observable part is solely responsible for the differences between the models (since the unobservable component is identical to both models by construction). In the matching experiment, the baseline model shows a drop in observable productivity by approximately 3.28%, while the disruption model shows an approximate 1.95% drop.

Finally, the discrepancies between models have an important effect when analyzing the decomposition of Aggregate Productivity Growth (APG). We construct APG for the inference sample and decompose it to its underlying components, following the methodology of Petrin and Levinsohn (2012). Our sample period extends from 2001 to 2017. This is a turbulent period for Greek manufacturing. It encompasses three distinct subperiods in terms of economic activity. The first subperiod is between 2001 and 2007, which is the pre-financial crisis period when the sector was booming. The second subperiod is from

2008 to 2015, the crisis period. Finally, from 2016 to 2017 the sector was under recovery.

Both models find technical efficiency as the more significant factor of APG, with reallocation efficiency being generally of the same sign but smaller. There is, however, the exception of the financial crisis period, where reallocation efficiency operates in a positive direction (demonstrating the cleansing effect of the crisis), but is overcome by the larger negative effect of technical efficiency. Although the qualitative conclusions are similar for both models, there are significant differences concerning the magnitudes of the components, with the disruption model finding the average values of all components smaller that those of the baseline.

In the pre-crisis period, the disruption model estimates average technical efficiency to be, in absolute terms, 20.43% lower than that of the baseline model and reallocation efficiency 11.39% lower. During the crisis, the disruption model gives a 59.29% lower magnitude for average technical efficiency and 79.28% lower reallocation efficiency. During recovery, the disruption model's average technical efficiency is lower by 15.36% and reallocation efficiency lower by 79.49%.

The two gross-output models we consider are based on the production function model of Ackerberg, Caves, and Frazer (2015) (henceforth ACF). The use of gross output in this paradigm may seem erroneous at first, as Bond and Söderbom (2005) have shown that such a combination does not produce uniquely identifiable production functions, in the presence of fully flexible inputs (i.e. the intermediate inputs present in gross-output specifications). This has led many authors to estimate value-added functions, instead, whose implicit assumptions, however, have received their own skepticism[1]. In response to all this, we follow a new method developed by Gandhi, Navarro, and Rivers (2020), which enables the estimation of gross-output production functions in the ACF setting. In particular, we use a simplified version found in Collard-Wexler and De Loecker (2016).

[1] See Bruno (1978), Diewert (1978), and Basu and Fernald (1995).

Our key modification of this framework is to allow the observable shock to productivity to follow a first order Markov process that depends also on the first lag of an investment spike dummy, in addition to its own lag.

There is a literature that documents investment spike induced disruptions in firm productivity. Hugget and Ospina (2001) find that TFP growth drops following an investment spike, in Colombian Manufacturing. Sakellaris (2004) finds that investment spikes on a sample of US Manufacturing plants are followed by TFP drops. Cooper and Haltiwanger (2006) estimate various models of capital adjustment on US Manufacturing plants and find that a model with mixed convex and non-convex adjustment costs, the latter being triggered by investment spikes, to be the best fit for their data. Shima (2010) finds a negative relationship between lumpy adjustment and TFP in Japanese Manufacturing. Gradzewicz (2020) finds disruption patterns, similar to those discussed in this paper, in a large panel of the Polish economy.

Some studies have analyzed the possibility of endogenous productivity biasing the production function estimates and subsequent inference on productivity. Doraszelski and Jaumandreu (2013) use R&D expenditures. De Loecker (2013) uses variables related to exports. Khan and Khederlarian (2021) use inventory costs. To the best of our knowledge, this paper is the first to incorporate the findings of the lumpy adjustment literature into the production function estimation framework. Our own modification of the productivity process, founded on the observed spike-induced productivity disruptions in our dataset, should allow for more accurate production function estimates and greater reliability on results that depend on them. In addition, we further investigate the behavior of individual components of productivity, which is also something that has been missing from the literature.

The rest of this paper proceeds as follows. Section 2 presents the dataset we employ, discusses some variable definitions, and presents evidence of the disruption effect in our sample. Section 3 presents a disruption model, where the productivity process is affected by investment spikes. Section 4 discusses the differences in estimates and conclusions between this and the baseline model, and presents the empirical application on Aggregate Productivity Growth. Section 5 concludes. Our online appendix provides further useful details.

2 Data and Variable Definitions

We employ a large proprietary unbalanced dataset from the ICAP database on Greek firm annual financial statements and employment data from 1998 to 2017. After the basic dropping of duplicates and observations without financial data, the ICAP dataset contains about 750,000 unique firm-year observations for about 100,000 Greek firms. However, we restrict to a sample of roughly 125,000 observations for the Manufacturing sector, covering about 15,000 firms. We focus on Manufacturing because it reports the most reliable data and because there is better availability for the necessary deflators and other sub-sector specific variables from other sources.

Our key variables are output, the production inputs, and investment in physical capital. For output, we use the book values of net sales, other operating income, and the change in inventories of finished goods, which we deflate by a producer price index provided to us by ELSTAT[2]. For intermediate inputs, we use the book values of the cost of goods sold and other operating expenses minus their depreciation and the cost of labour. Then, we deflate by the intermediate inputs price index reported by EUKLEMS. For investment and capital, we follow the perpetual inventory method (PIM) per asset class, and then add together the corresponding series for Buildings and Machinery & Equipment to get the aggregate

[2] In a few cases where that is not available, we use the Harmonized Consumer Price index for the whole economy, reported by Eurostat.

Table 1: Summary Statistics

	Output	Capital	Labour	Int. Inputs	Inv. Rate
Min	0.2838	-0.2007	0.0000	2.5493	-0.9993
Median	14.6503	13.6829	3.2226	14.2217	0.0671
Mean	14.7694	13.6066	3.2599	14.2988	0.1832
Max	22.9690	21.6615	10.1701	23.0460	11.9640
S.D.	1.5343	1.8786	1.2658	1.7040	0.6093
Skewness	0.3257	-0.2614	0.2111	0.1463	7.7747
Kurtosis	4.3863	4.1542	3.4011	3.9975	95.3947

$N = 41768$

Output, capital, labor, and intermediate inputs are in logs. The investment rate is defined as total real investment in physical capital over total real physical capital.

measures.

More details on variable construction and data cleaning are given in online appendix B. After these tasks, we are left with a final sample of 41768 observations for 14636 Manufacturing firms. The panel is unbalanced, as some firms are present in the sample for most of the time range covered, while others may only appear for a brief period. We will be referring to this as the inference sample, as it comprises of all of the observations that are involved in the various inference steps to follow.

Table 1 reports some summary statistics for output, capital, labor, and intermediate inputs in logs and for the investment rate, in the inference sample. All variables appear to be slightly asymmetrically distributed with somewhat fat tails, with the notable exception of the investment rate. Table 2 reports the cross correlations and autocorrelations of the same variables. The production function variables seem to display a significant positive

Table 2: Correlations and Autocorrelations

	Output	Capital	Labor	Int. Inputs	Inv. Rate
Output	1.0000				
Capital	0.7399	1.0000			
Labour	0.7741	0.6531	1.0000		
Int. Inputs	0.9571	0.7026	0.6857	1.0000	
Inv. Rate	0.0322	-0.1082	0.0032	0.0324	1.0000
Autocorrelation	0.9697	0.9836	0.8886	0.9438	0.0564

$N = 41768$

Output, capital, labor, and intermediate inputs are in logs. The investment rate is defined as total real investment in physical capital over total real physical capital.

relationship and strong autocorrelations. In particular, the large positive autocorrelation of log labor allows us to use its first lag as its instrument when appropriate, which is needed for the purposes of the estimation of the model presented in Section 3. Again, the investment rate is the odd one out, being generally independent of the rest of the variables and not particularly autocorrelated.

In table A.1 of online appendix A, we report the per year average growth rates of output and the production inputs, and the Aggregate Productivity Growth (APG), for the inference sample, the latter of which we calculate in Section 4.3 based on the methodology of Petrin and Levinsohn (2012). We observe that for the period between 2001 and 2007 the average growth rate of output is 2.56%, which is driven by the positive growth in both aggregate productivity and its inputs. From 2008 to 2015, during the 2008 Greek financial crisis, the average output growth is -6.80%, due to the negative growth in aggregate productivity and its inputs except for capital, which merely slows down. During the short recovery period in our sample, all variables show positive growth rates.

Table 3: Investment Rate Statistics

Number of observations with:	
$I_{it-1}/K_{it-1} \leq -0.01$	6062
$\lvert I_{it-1}/K_{it-1} \rvert < 0.01$	4128
$I_{it-1}/K_{it-1} \geq 0.01$	31578
$I_{it-1}/K_{it-1} > 0.20$	11416
Number of spikes	3661

$N = 41768$

2.1 Investment Spike Definition

The statistics in table 1 imply that the empirical distribution of the investment rate in our sample is not symmetric. Table 3 gives further insights on this by reporting the number of observations in the sample for which the first lag of the investment rate falls within certain intervals of interest[3].

Out of the 41768 observations, only 6062 report disinvestment and 4128 are within the interval of -0.01 and 0.01, which is associated with investment inactivity in the literature. The rest report positive investment. This is evidence of the presence of non-convex adjustment costs in capital and justifies the need to investigate the effect of investment spikes on the firm. We see that 11416 observations lie above the 0.20 threshold for the investment rate, which is a typical lower bound used to characterize investment spikes in the literature. However, because our definition of an investment spike sets additional criteria, only 3661 of those are identified as such.

We define investment spikes similarly to Gradzewicz (2020). We use a combination of the definitions in Power (1998) and Sakellaris (2004). Let K_{it} be the level of physical capital of firm i in period t, I_{it} its level of investment, and d_{it}^{I} be the spike dummy, then

[3] We report for the first lag because that is what is used at the estimation stage.

we define

$$d_{it}^{I} = \begin{cases} 1, & \dfrac{I_{it}}{K_{it}} > \max\left\{\alpha E\left(\left.\dfrac{I_{it}}{K_{it}}\right| K_{it}\right), \gamma\right\} \\ 0, & \text{otherwise} \end{cases} \tag{1}$$

with $\alpha = 2.75$ and $\gamma = 0.20$. $E\left(I_{it}/K_{it}\right| K_{it})$ is the expected investment rate for firm i in period t conditional on its level of capital. We construct it as the fitted value of the ordinary least squares regression of I_{it}/K_{it} on a third degree polynomial of the natural logarithm of K_{it} and a complete set of time dummies. To account for sector heterogeneity and the possibility of a structural break caused by the 2008 Greek financial crisis, we perform a different regression per 2-digit NACE rev. 2 sector and before and after 2008. Finally, if the criterion is satisfied for consecutive years for the same firm, we only identify the first instance as a spike.

This definition results in 3661 spike events being identified in the inference sample, which is about a third of the original 11416 candidates. Furthermore, table A.2 of online appendix A reports the average investment rate and percentage of identified investment spikes in the inference sample per year. We observe that average investment rates decline during the 2008 financial crisis relative to the pre-crisis period[4]. Because we control for the crisis, it does not affect the proportion of observations we identify as investment spikes each year, always being roughly between 5% - 10% of all contemporaneous observations in the sample.

2.2 Sample Matching Method

To study the effect investment spikes have on productivity, we define the sub-sample of spike observations, as the treatment group, and an equally sized matched sub-sample of non-spike firms with similar characteristics, as the control group. We follow a variant of the

[4] Fakos, Sakellaris, and Tavares (2022) document and analyze the investment slump in Greek Manufacturing during the financial crisis.

Mahalanobis distance within calipers approach described in Rosenbaum and Rubin (1985). Because we need our observations to have available data on productivity for five consecutive years, we defer the construction of the spike and matched sub-samples until after the models have been estimated (see Section 4) and the desired measures of productivity have been calculated. We, then, restrict to the observations with the data availability we require.

The matching process works with a set of approximate matching variables and a set of exact matching variables. We select firm age, log real sales, real sales growth (the first difference in log sales), and lagged leverage (total liabilities over total assets) as our approximate matching variables. We use the observation's year, the firm's legal form, its 2-digit NACE rev. 2 sector, and its trading status (public or private) for exact matching. We then regress a logit model with the spike dummy as the dependent variable and the approximate matching variables, a set of time dummies, and a set of 2-digit sector dummies as the explanatory variables. The fitted log odds ratio is called the propensity score and is used to determine the closeness of the observations between samples.

Both the sample of spikes and candidate matches are sorted randomly. For each observation in the spike sample, we find its match by restricting to the candidates that exactly match the exact matching variables and are within a 0.5 range above or below of the propensity score of the spike. Then, we do not just match by the closest propensity score, but by the Mahalanobis distance between the spike and all of the candidate matches. This distance is calculated according to the approximate matching variables and the propensity score. The observation with the minimum Mahalanobis distance is designated a match and it is removed from the sample of candidates (matching without replacement). Once every spike has its match, we drop any pairs with a Mahalanobis distance greater than 3. The two final sub-samples cover 591 spike episodes.

Figure 1 about here.

Figure 1 illustrates the average performance of TFP of a firm that undergoes an investment spike, for a time window of two periods before to two periods after a spike event, compared to its statistical match that does not. The measure of TFP used is derived based on the estimates of a baseline gross-output production function estimated according to GNR. Here, we have a first look at the figure and go into greater detail in Section 4.2. One can see that on average productivity drops by approximately[5] 7.07%, following an investment spike. Compared to the overall variability of the series, this drop is too large to be dismissed as a random occurrence. A Wilcoxon type test also rejects the null of no shift in the spike sample from $t = 0$ to $t = 1$, with a p-value of zero. This means that the distribution of productivity of spike firms is shifted downwards after a spike. In addition, no similar behavior is observed in the matched sample. This is evidence of endogenous productivity disruptions caused by lumpy adjustment. However, this phenomenon is not reflected in the standard assumptions of the ACF/GNR model. Thus, we find it necessary to align it with the evidence and develop a new version which allows for investment spikes to affect productivity, which is what we do in Section 3.

3 The Disruption Model

A lot of research that uses production functions focuses on value added. However, there has been skepticism regarding its validity and the validity of the necessary conditions for it to give equivalent results to gross-output[6]. There is also the structural real value-added specification, used for example in Ackerberg et al. (2015), where the production function is assumed to be separable in the primary and the intermediate inputs by a Leontief function. The practical difference between (deflated) value-added and structural value-added comes down to if and how to deduct the value of the intermediate inputs from output before

[5] Since the log of TFP drops by 0.0707, the percentage drop is approximately 7.07%.

[6] See Bruno (1978) and Diewert (1978).

regressing. In both cases, one estimates a model of some measure of output on the primary inputs alone.

Given that intermediate inputs are generally freely adjustable and the non-identifiability critique of Cobb-Douglas functions with flexible inputs by Bond and Söderbom (2005), it would seem that value added is the only option. Even with the added assumptions it requires. However, the new approach by GNR allows for the estimation of gross output production functions with fully flexible inputs. So, we present a disruption model, which builds upon the baseline model of ACF and GNR, but includes investment spike decisions as a determinant of productivity. Its estimation is practically identical to that of the baseline. The model and its assumptions are as follows.

Consider a set of firms, with a gross-output Cobb-Douglas production function in logs

$$y_{it} = \beta_K k_{it} + \beta_L l_{it} + \beta_M m_{it} + \omega_{it} + \varepsilon_{it} \tag{2}$$

where the i subscript indicates the firm in question and t the time period. y_{it} is the natural logarithm of the firm's gross output and k_{it}, l_{it}, and m_{it} are the logs of capital, labor, and other intermediate inputs (raw materials, energy, etc), respectively. β_K, β_L, and β_M are the respective Cobb-Douglas parameters.

The production process is subject to two types of shock, ω_{it} and ε_{it}, which combined give its total factor productivity. Both variables are unobservable to the econometrician. However, ω_{it} is observed by the firm before production begins, so we may also refer to it as the observable shock. ε_{it} is not observed by the firm until after production, thus we may also refer to it as the unobservable shock. It may incorporate highly unpredictable factors that affect output or even the measurement errors of some variables, all of which may or may not display some serial correlation.

Let $\mathcal{I}_{it}$ denote firm i's information set in period t.

Assumption 1. *The current and all past levels of observable productivity,*

$\omega_{it-s} \forall s \in \{0, 1, 2, \dots\}$, *are included in* $\mathcal{I}_{it}$, *but any future values,* $\omega_{it+s} \forall s \in \{1, 2, \dots\}$, *are not. For the unobservable shock it holds that* $E(\varepsilon_{it}|\mathcal{I}_{it}) = 0$.

Assumption 1 is identical to that of the ACF/GNR framework. Our point of differentiation is in assumption 2.

Assumption 2. *The observable shock,* ω_{it}, *follows a first order Markovian process*

$$\omega_{it} = g(\omega_{it-1}, d^I_{it-1}) + \xi_{it} \tag{3}$$

where d^I_{it-1} *is the first lag of the disruption dummy defined in Section 2.1 and* ξ_{it} *is such that* $E(\xi_{it}|\mathcal{I}_{it-1}) = 0$.

The choice of d^I_{it-1} instead of d^I_{it}, follows from our findings in Section 4 that disruption takes effect one period after the investment spike. In the baseline approach, which uses a restricted version of (3), g is only a function of ω_{it-1}. We will be estimating both models in order to juxtapose them in Section 4.

Because we are interested in the long-run average and the persistence of productivity, we parameterize g to be linear and have ω_{it} following an AR(1) process in the baseline model and a controlled AR(1) in the disruption model. Thus, (3) becomes

$$\omega_{it} = c + \rho\omega_{it-1} + \xi_{it} \tag{4}$$

for the baseline model and

$$\omega_{it} = c + \rho\omega_{it-1} + \delta d^I_{it-1} + \xi_{it} \tag{5}$$

for the disruption model, with $0 < \rho < 1$. δ measures the effect of lagged investment spikes on current productivity. Notice how (4) is a restricted version of (5) with δ fixed at zero.

We also follow Olley and Pakes (1996) and, during estimation, we augment (4) and (5) with a survival probability term, to account for firm attrition, which can bias the results. We obtain the survival probabilities for the next period by logit regression of a dummy variable indicating survival in the next period on third degree polynomials (without cross terms) of the investment rate, the logs of output, capital, labor, and the intermediate inputs, and a complete set of time dummies. We then lag the fitted probabilities to get the current period value and use it in the regressions.

The rest of the assumptions follow ACF/GNR.

Assumption 3. *Capital is a fixed input and its level, k_{it}, is determined by the previous period's investment, at time $t-1$, and some law of motion. The level of labor, l_{it}, is chosen at time $t-b$ with $0 < b < 1$[7], when ω_{it} is partially observed. Thus, labor is considered a quasi-fixed input.*

By the above assumptions, l_{it} is correlated with ω_{it} and thus ξ_{it}, whereas k_{it} is not. However, the relative rigidity of labor means that it should display some positive correlation with its first lag, which is independent of ξ_{it} and can be used as its instrument. Both k_{it} and l_{it} are uncorrelated with ε_{it}. From all this, we can get the following unconditional moment conditions

$$E(k_{it}\xi_{it}) = 0 \tag{6}$$

$$E(l_{it-1}\xi_{it}) = 0 \tag{7}$$

$$E(k_{it}\varepsilon_{it}) = 0 \tag{8}$$

$$E(l_{it}\varepsilon_{it}) = 0 \tag{9}$$

Assumption 4. *The other intermediate inputs are fully flexible and their level, m_{it}, is determined just before production starts in period t. There exists an intermediate input*

[7] Depending on the particularities of the data, the model, and the general context, the researcher may want to set b equal to exactly zero or exactly one. This can be accommodated by the GNR method with the necessary modifications to the estimation process, by treating labor same as capital or the intermediate inputs, depending on the situation.

demand function

$$m_{it} = f(k_{it}, l_{it}, \omega_{it}) \tag{10}$$

This means that m_{it} is chosen freely and optimally, up to the unobservable shock, ε_{it}.

Additionally, assumptions 1 and 4 imply that

$$E(m_{it}\varepsilon_{it}) = 0 \tag{11}$$

Assumption 5. *f is strictly increasing and invertible in ω_{it}.*

The invertibility of f with respect to ω_{it} implies that there exists a function, f^{-1}, such that

$$\omega_{it} = f^{-1}(k_{it}, l_{it}, m_{it}) \tag{12}$$

The main argument of GNR, which we also follow, is that the flexibility of m_{it} permits the use of revenue share regressions to identify β_M. Using an estimate of β_M, we can define a measure of output net of the effect of intermediate inputs

$$\widetilde{y}_{it} = y_{it} - \widehat{\beta_M} m_{it} = \beta_K k_{it} + \beta_L l_{it} + \omega_{it} + \varepsilon_{it} \tag{13}$$

Given that the rest of the inputs are not fully flexible, one can use a revenue share regression for the flexible inputs in conjunction with the ACF framework for $\widetilde{y}_{it}$ and equation (13). Thus, one can retrieve estimates for every parameter in (2) and circumvent the unidentifiability issues underlined by Bond and Söderbom (2005).

The ACF two-stage procedure substitutes (12) in (13) to get

$$\widetilde{y}_{it} = \beta_K k_{it} + \beta_L l_{it} + f^{-1}(k_{it}, l_{it}, m_{it}) + \varepsilon_{it} = \Phi(k_{it}, l_{it}, m_{it}) + \varepsilon_{it} \tag{14}$$

where Φ is an unknown function of observable variables. At the first stage, Φ is approximated by least squares on a polynomial of the observables. At the second stage, the fitted values, $\widehat{\Phi}_{it}$, are used with candidate estimates of (β_K, β_L) to get estimates for ω

with which to estimate (or approximate) g. The residuals, $\widehat{\xi}_{it}$, are combined with a set of instruments to give moment conditions for which $(\widehat{\beta_K}, \widehat{\beta_L})$ is chosen optimally.

We use the revenue share of intermediate inputs and the moment conditions derived above, following Collard-Wexler and De Loecker (2016), to estimate the production function in (2). However, do not account for measurement errors, as they do, to avoid the systematic loss of a considerable number of observations. We employ a linear approximation for Φ and the linear functional form we have assumed for g. As a side effect, we can also get estimates for the parameters in (5) and (4), as well as approximate the unobservable (to the econometrician) variables ω_{it} and ε_{it}. The estimation essentially follows from the procedure described in the appendix of Collard-Wexler and De Loecker (2016). Nevertheless, we provide detailed estimation steps in online appendix C. LP, ACF, and GNR also provide their own detailed appendices.

4 Estimation Results

4.1 Model Estimates

We proceed by estimating the production function for the baseline and the disruption model described in Section 3. Table 4 reports the estimated parameters for both models. The standard errors for the production function parameters are calculated by 1000 bootstrap samples, following Levinsohn and Petrin (2003)[8]. The standard errors for the parameters in the productivity process, g, are heteroskedasticity and autocorrelation robust errors.

We first examine the estimates for the production function. Given how relatively small

[8] According to Hahn and Liao (2021), standard error estimates produced via bootstrap can be more conservative (i.e. greater) than the true standard errors. This is not a problem for us, however, as it only means that our parameter estimates are more accurate than we report. For a more detailed analysis on the bootstrap, also see Horowitz (2001).

Table 4: Production Function and Productivity Process Estimates

	Production Function			Productivity Process		
	β_K	β_L	β_M	c	ρ	δ
Baseline	0.1029	0.1449	0.6777	0.4812	0.8427	
	(0.0098)	(0.0225)	(0.0016)	(0.0147)	(0.0048)	
Disruption	0.0799	0.2299	0.6777	0.3109	0.9284	-0.0179
	(0.0046)	(0.0061)	(0.0016)	(0.0080)	(0.0025)	(0.0004)

$N = 41768$

Values in parentheses report standard errors. For the production function, they are calculated using 1000 bootstrap samples. For the productivity process, they are heteroskedasticity and autocorrelation robust errors.

the estimated standard errors are, we can confidently say that all coefficients are statistically significant for both models. We also observe that the estimates for some coefficients change significantly. The coefficient for capital decreases from 0.1029 to 0.0799, which using the standard deviation of the baseline model is 2.3302 standard deviations lower. The coefficient for labor increases from 0.1449 to 0.2299, which is 3.7753 baseline standard deviations higher. That these two coefficients adjust between models in the opposite direction is not surprising. According to Collard-Wexler and De Loecker (2016), because capital and labor are positively correlated, an downwards correction to the capital coefficient leads to a upwards adaptation of the coefficient for labor and the overall effect on returns to scale is unknown. In this case, returns to scale increase, form 0.9255 to 0.9875. Also, the disruption model seems to produce estimates with smaller standard errors. Moreover, the recent findings by Hahn and Liao (2021) suggest that second moment estimates using the bootstrap are biased upwards. So, the standard deviations for the production function coefficients may be smaller than estimated and the models are more precise and more statistically distinct than they appear to be. Of course, the estimate for the intermediate inputs coefficient remains unchanged, since it is calculated before the inclusion of the spike dummy takes place. The important conclusion is that including the spike dummy definitely has an effect on the production function estimates.

We now have a look at the estimates for the productivity process. Again, the relatively small standard errors imply that all coefficients are statistically significant. In both models, productivity is strongly persistent, with an autocorrelation coefficient of 0.8427 and 0.9284 for the baseline and the disruption model, respectively. In both cases, the estimated processes are well bounded away from unit roots. Both models have a positive and statistically significant constant, implying a positive long-run average for productivity. In the disruption model, we can see that the investment spike dummy has a negative and statistically significant coefficient of -0.0179. This suggests a short-term negative effect of investment spikes on output of approximately 1.79%. Furthermore, the persistent nature of the productivity process means that investment spikes also have their mark in later periods, and productivity needs more time to recover. We return to this idea in the next subsection.

Finally, the constant and persistence parameters between both models imply a very different long-run average of productivity. For the baseline model it can be estimated as

$$\bar{\omega} = \frac{\hat{c}}{1 - \hat{\rho}} \tag{15}$$

which gives an estimate of 3.0594. For the disruption model, the calculation becomes

$$\bar{\omega} = \frac{\hat{c} + \hat{\delta} \cdot P(\widehat{d^I_{it-1} = 1})}{1 - \hat{\rho}} \tag{16}$$

In the inference sample, 8.77% of observations report an investment spike (in first lags). Using this percentage as the probability of an investment spike, the estimated $\bar{\omega}$ for the disruption model is 4.3187. So, the disruption ω_{it} is calculated to be greater by 41.16%

compared to the baseline estimate. Since the average value of the unobservable shock is zero for both models, the difference in the long-run averages of log TFP is the same.

4.2 Productivity Component Performance

We now study the effects of investment spikes on log TFP and its components, ω_{it} and ε_{it}. We retrieve the three variables for each model as follows: for $log(TFP_{it})$ we use the estimated production function parameters with the production function in logs to get $log(TFP_{it}) = y_{it} - \widehat{\beta_K} k_{it} - \widehat{\beta_L} l_{it} - \widehat{\beta_M} m_{it}$, for ω_{it} we use the values produced in the second stage of the estimation, and we get ε_{it} as the difference between $log(TFP_{it})$ and ω_{it}[9].

In the sample, both the average logarithm of TFP and the average of ω_{it} for the baseline model are 3.2056 (compared to the long-run estimate of 3.0594, calculated in Section 4.1). For the disruption model, they are both at 3.2408 (compared to a long-run average of 4.3187 in Section 4.1). Thus, the disruption model places average productivity 1.10% higher than the baseline model. These sample averages and their difference are very different than the corresponding long-run estimates of Section 4.1. Technically, the reason for this is the inclusion of firm exit correction term in the estimation step. Excluding this, would bring both statistics for each model in agreement. However, that would also introduce selection bias in the results. For ε_{it} in both models the average is zero, which is to be expected since it is equivalent to the OLS residual from the first stage of the estimation.

We select a sub-sample of 591 cases, in our sample, that experience an investment spike and have available data for a time window of two periods before and after the spike. Following the methodology described in Section 2.2, we select an equally sized sub-sample of cases of firms that are statistically similar, but do not exhibit a spike, and have the same

[9] Equivalently, we could use the residuals from the first stage of the estimation. We have found that the two alternatives only differ at the very low digits, which could be attributed to representational rounding errors.

time window of available data. We draw a series of plots for these sub-samples, regarding the performance of average TFP and its components in the sample around investment spike episodes. All investment spike episodes are grouped together, regardless of the time period they occurred, thus mitigating any time effects.

We can also use the reported averages to conduct simple difference-in-difference analyses. Let d be a dummy variable indicating observations in the spike sub-sample. The difference-in-difference statistic measuring the average effect of being a spike firm on a firm specific variable, X, between periods $s-1$ and s is given by

$$\begin{aligned} &(E(X|t=s,\, d=1) - E(X|t=s-1,\, d=1)) \\ -&(E(X|t=s,\, d=0) - E(X|t=s-1,\, d=0)) \end{aligned} \tag{17}$$

Figure 2 about here.

Figure 2, plots the average values of $log(TFP_{it})$, ω_{it}, and ε_{it} separately for the spike and the matched sub-samples, for a time window of two periods before to two periods after an investment spike. Observe that, in these sub-samples also, the averages for both $log(TFP_{it})$ and ω_{it} are at visibly higher levels for the disruption model compared to the baseline model, same as it is found for the sample averages. That is not the case for ε_{it}, however, which is identical for both models. This makes sense, since the residuals of the first stage, with which ε_{it} is equivalent, are calculated before any assumption on the process of ω_{it} is involved.

For both models, two periods before the spike (at $t=-2$), the averages of log TFP (subfigures (a) and (b)) evolve mostly in parallel towards the period before the spike ($t=-1$). The difference-in-difference statistic for log TFP from $t=-2$ to $t=-1$ is -0.0066 for the baseline model and -0.0076 for the disruption model. This reassures that the matching process did a good job at finding appropriate statistical twins of the spike firms, even though productivity and its components are not targeted variables of the matching algorithm.

During the spike period (at $t = 0$), in both models, the average spike firm enjoys a relatively higher TFP than its non-spike counterpart. For the baseline model the baseline difference-in-difference statistic is 0.0343 and the disruption statistic is 0.0332, introducing a gap between spike and non-spike firms. This productivity advantage of spike firms may be part of what motivates a firm to adjust via spike. The increased productivity may be seen by the firm as an opportunity to mitigate some of the expected disruptions of the next period. Alternatively, the firm may itself be increasing the utilization of its inputs to prepare for future disruption, making it appear as though its productivity is increased. Testing the latter hypothesis would require a more that allows for variable input utilization, however, our dataset does not permit for the estimation of such a model.

The first period after the spike (at $t = 1$), we can see a stark difference in productivity growth rates, between spike and non-spike firms. In the baseline model, the difference-in-difference statistic is -0.0702 and for the disruption model is it slightly more conservative at -0.0552. So, it seems that both models show signs of a disruption effect, but the baseline model may be overestimating its magnitude. Such productivity drops, which are correlated with lumpy adjustment behavior, constitute an extra adjustment cost that the firm has to face for choosing spiky adjustment, compared to smoother adjustment, on top of any usual convex costs. These costs may appear because significant adjustment may require production units to halt operations while the new capital is installed, or may be the result of the production workers slowly adapting to the new capital they work with (learning-by-doing). As we will see below, such productivity disruptions have a lingering effect on the firm's future output, due to the persistent nature of ω_{it}.

Finally, at $t = 2$, the log TFP of spike firms starts to recover, with a difference-in-difference statistic of 0.0185 and 0.0084, for the baseline and the disruption model, respectively. All of the above observations are in line with the general findings of the lumpy adjustment literature.

We now move on to dissecting TFP into its individual components, the observable ω_{it} component and the unobservable ε_{it} component (subfigures (c), (d), (e), and (f)). Firstly, we observe in both figure that ω_{it} is much less variable for both models, compared to ε_{it}, which appears to drive a lot of the variability in TFP. Before the spike (at $t = -2$ and $t = -1$), ω_{it} and ε_{it} move relatively in parallel between sub-samples, same as with TFP, with the baseline statistic being 0.0015 and -0.0081, for ω and ε respectively, and 0.0005 and -0.0081 in the disruption model.

The relatively increased productivity of spike firms at the spike period ($t = 0$) is almost exclusively due to the unobservable shock, ε_{it}, with a difference-in-difference statistic of 0.0312 for both models. The corresponding statistics for the observable shock, ω, are 0.0031 and 0.0020. That ε is the main driver of the boosted productivity of the spike period discussed above, is in line with both interpretation given. That of a random occurrence and that of variable utilization.

In the first period after the spike ($t = 1$), the situation is different. Now, both components respond to the spike with a significant drop. The difference-in-difference statistics for the observable ω_{it} are -0.0335 for the baseline model and -0.0186 for the disruption model. Even though the sub-samples used in these experiments are rather limited in size, which could lead to biased and inaccurate conclusions, the statistic for the disruption model is relatively close to the estimated disruption coefficient of -0.0179, based on 3661 spike events in the inference sample. Also, the larger statistic found for the baseline model reveals that failing to account for endogenous productivity disruptions due to investment spikes can lead to a false measurement of the disruption effect's size itself. The statistic for the unobservable ε_{it} is -0.0367, possibly implying that there are additional sources of disruption, which act though channels unobservable by the firm. Finally, at $t = 2$, both productivity types begin recovery.

We are interested in examining whether our estimate of the disruption effect on the observable shock, $\widehat{\delta}$, is backed by the data. We do this via a paired Wilcoxon signed rank test[10]. We take 2283 cases for which ω_{it} is available during and after an investment spike and consider two samples, one containing the observations during the spike and one for the next period. We use the ω_{it} that is produced by the disruption model.

We conduct a two-sided Wilcoxon test, where the null hypothesis is that the location shift between the two samples is equal to $-\widehat{\delta}$. Our estimate for δ is -0.0179, implying a location shift of 0.0179 from the spike period to the next. The test's p-value is 0.8283, which accepts the null at the 5% confidence level. Furthermore, the 95% confidence interval is [0.0220, 0.0249], which excludes zero. This reinforces the evidence for the presence of a disruption effect in the examined sample. Furthermore, the test reports a pseudomedian for the location shift of 0.0178, which is the test's estimate of the shift. This is practically identical to our estimate and adds to the evidence in favor of the existence of a disruption effect, which should not be ignored when conducting any kind of inference.

Finally, using the model parameter estimates, we can get an idea of the overall effect a single investment spike episode has on output across time. Due to the persistence of ω_{it}, the overall effect of an investment spike on it from the time disruptions take effect and across all future periods is given by $\dfrac{\hat{\delta}}{1-\hat{\rho}} = -0.2498$. This means that, ceteris paribus, lumpy adjustment incurs the firm with the loss of approximately 24.98% of its future output. Notice here how the baseline model provides no structural way of measuring this long-term effect of lumpy investment. Thus, this significant implicit extra adjustment cost the firm has to bear, when choosing to adjust in spikes rather than smoothly, is not limited to one period, and, in addition, is completely missed by the baseline model.

[10] Because the exact Wilcoxon test can be computationally demanding even for relatively small samples, we use a continuity corrected normal approximation.

4.3 Aggregate Productivity Growth

We further illustrate the importance of the difference in estimates between the baseline and the disruption model with an empirical application. We construct Aggregate Productivity Growth (APG) for the inference sample and decompose it to its underlying components, following the methodology of Petrin and Levinsohn (2012). We deviate slightly from their definition, to reflect our gross-output assumption, and calculate the discrete time version of APG in each time period as

$$APG_t = \sum_i \overline{D}_{it}\Delta y_{it} - \sum_i \overline{D}_{it}\overline{s}_{ikt}\Delta k_{it} - \sum_i \overline{D}_{it}\overline{s}_{ilt}\Delta l_{it} - \sum_i \overline{D}_{it}\overline{s}_{imt}\Delta m_{it} \tag{18}$$

where summation is across all firms present in the sample in period t, $\overline{D}_{it}$ is the firm's Domar weight, which we calculate as total revenue over the sum of all revenues across all firms in period t, and $\overline{s}_{ikt}$ is the revenue share of capital and similarly with $\overline{s}_{ilt}$ and $\overline{s}_{imt}$ for labor and intermediate inputs, respectively. The bars over $\overline{D}_{it}$, $\overline{s}_{ikt}$, $\overline{s}_{ilt}$, and $\overline{s}_{imt}$ indicate average values between periods $t-1$ and t. Δy_{it}, Δk_{it}, Δl_{it}, and Δm_{it} are the first differences of log output and the respective log inputs. Before constructing APG and any of its components, we drop observations that lie at the top and bottom 0.5% of the empirical distribution of the sum of revenue shares.

For the revenue shares of the inputs, we need data on their costs. For labor and intermediate inputs, we simply multiply the real variables in levels with the average cost of labor and the intermediate inputs price index, respectively (see online appendix B). Of course, data on capital use costs are not reported in any of the firm's financial statements. Basu and Fernald (1995) and Balk (2021) provide some measures one can construct, but the necessary data for Greece are not easily available. Petrin and Levinsohn (2012), who face the same problem, opt to not account for capital in their data and provide an APG measure which is accurate "up to an adjustment for capital expenditures". We prefer to use a simple, but attainable, approximation used in Bitros and Panas (2001) for Greek

Manufacturing. We define the cost of capital as

$$r_{it}^K = P_{it}^I(r_{it} + \delta_{it}) \tag{19}$$

where P_{it}^I is the investment deflator, r_{it} is the intertemporal interest rate, and δ_{it} the capital depreciation rate. We have data for the investment deflator and the capital depreciation rates, which we use for our variable construction. We calibrate the intertemporal interest rate at a fixed 0.025 for all periods, taken from table 9 of the online appendix of Fakos et al. (2022), who study the effects of credit supply shocks on Greek Manufacturing firms' investment before and during the Greek 2008 financial crisis. We calculate the unit cost both for Buildings and Machinery & Equipment and then appropriately calculate the revenue share of aggregate capital.

Notice how the calculation of APG relies only on observable data and does not depend on any estimated parameters. Production function estimates are needed, however, to get a breakdown of the forces acting on APG. Petrin and Levinsohn (2012) split APG into changes in Technical Efficiency (TE), Reallocation Efficiency (RE) for each input, and fixed and sunk costs (F). Thus,

$$APG_t = TE_t + RE_t + F_t \tag{20}$$

with

$$TE_t = \sum_i \overline{D}_{it} \Delta ln(TFP_{it}) \tag{21}$$

$$RE_t = \sum_i \overline{D}_{it}(\beta_K - \overline{s}_{ikt})\Delta k_{it} + \sum_i \overline{D}_{it}(\beta_L - \overline{s}_{ilt})\Delta l_{it} + \sum_i \overline{D}_{it}(\beta_M - \overline{s}_{imt})\Delta m_{it} \tag{22}$$

$$F_t = -\sum_i \overline{D}_{it} \Delta ln(F_{it}) \tag{23}$$

and F_{it} being the fixed and sunk costs of firm i at time t. F_t is calculated residually as $F_t = APG_t - TE_t - RE_t$. TFP is calculated as in Section 4.2.

Table 5 shows the evolution of APG and its components, during three time periods in our sample, according to both the baseline and the disruption model. The values reported

Table 5: Aggregate Productivity Growth breakdown according to the each model. Averages per subperiod.

Period	Model	APG	TE	RE	RE_K	RE_L	RE_M	N
2001-2007	Baseline	0.0016	-0.0020	0.0036	0.0073	-0.0016	-0.0021	20954
	Disruption	0.0016	-0.0016	0.0032	0.0050	0.0003	-0.0021	
2008-2015	Baseline	-0.0034	-0.0136	0.0102	0.0041	0.0073	-0.0013	11331
	Disruption	-0.0034	-0.0055	0.0021	0.0031	0.0003	-0.0013	
2016-2017	Baseline	0.0088	0.0109	-0.0021	0.0014	0.0005	-0.0040	5701
	Disruption	0.0088	0.0092	-0.0004	0.0011	0.0025	-0.0040	

Averages of the annual values of APG and its breakdown for the specified subperiods (annual values are found in tables A.3 and A.4 in online appendix A).

The estimates for fixed and sunk costs (F) are not reported, because they are zero everywhere.

N is the number of observations.

are the averages of the corresponding annual values, given by the Petrin and Levinsohn (2012) methodology, for each subperiod we define. The per year values are reported in online appendix A (tables A.3 and A.4). Our first subperiod is between 2001 and 2007, which is the pre-financial crisis period. The second subperiod is from 2008 to 2015, which we consider to be the crisis period. Finally, from 2016 to 2017 is our recovery period. The values of APG are common for both models, since its calculation is model independent. We observe that the average of APG per subperiod implies an average increase of aggregate productivity by 0.16% per year in the pre-crisis period, an average decrease by 0.34% per year during the crisis, and an average increase by 0.88% during recovery.

Moving on to the breakdowns by component, superficially it appears that both models mostly agree on the signs of each part of APG. Both models find that technical efficiency had a negative effect on APG before the crisis, which only got more negative in the crisis period.

Table 6: Aggregate Productivity Growth breakdown as percentage of annual level and ratios between models. Averages per subperiod.

	Baseline		Disruption		Ratio		N
Period	TE	RE	TE	RE	TE	RE	
2001-2007	-1.2515	2.2515	-1.0016	2.0016	0.8003	0.8890	20954
2008-2015	3.9585	-2.9585	1.6145	-0.6145	0.4079	0.2077	11331
2016-2017	1.2393	-0.2393	1.0492	-0.0492	0.8466	0.2055	5701

The columns under Baseline and Disruption report the ratios of TE/APG and RE/APG in each time period for the baseline and the disruption model, respectively. The columns under Ratio report the ratio of the disruption estimate over the corresponding baseline estimate.

N is the number of observations.

During the crisis, reallocation efficiency operates in a positive direction, demonstrating a cleansing effect of the crisis, but is overcome by the larger negative effect of technical efficiency. Then, technical efficiency drives the recovery of APG.

Although the qualitative conclusions are similar for both models, there is significant disagreement about the magnitudes of the components. Table 6 reports the ratio of the corresponding values of TE and RE over APG for the baseline and the disruption model, as well as the ratio between the models (as disruption over baseline). Notice that for each grouping the values of TE and RE sum up to one. We can observe that the models assign significantly different contributions of TE and RE to APG in all subperiods. In the pre-crisis period the disruption model estimates average technical efficiency to be in absolute terms 19.97% lower and reallocation efficiency 11.10% lower[11], compared to the baseline model. During the crisis, the disruption model gives a 59.21% lower magnitude

[11] These values are found by subtracting one from the corresponding columns under Ratio, which report the value estimated by the disruption model over the corresponding value of the baseline model.

for average technical efficiency and 79.23% lower reallocation efficiency. During recovery, technical efficiency is 15.34% lower and reallocation efficiency 79.45% lower. Overall, the evidence we have reviewed suggests that the choice of model can paint a substantially different picture of the underpinnings of APG.

5 Summary and Conclusions

In this paper, we study the effects of investment spikes on firm-level productivity on a panel of Greek Manufacturing firms from the ICAP database. Motivated by the evidence from the lumpy investment literature, we argue that production function estimation needs to take into account endogenous productivity disruptions caused by investment spikes. We contribute to the production function estimation literature, by providing a modified structural model that incorporates spike induced disruptions to production and can be estimated by the same method.

We find that our modification is meaningful and produces significantly different production function estimates in our sample, compared to the standard exogeneity assumption for productivity. Furthermore, we break down TFP in an observable (by the firm) part and an unobservable part. We find that the disruptions examined have an effect on output both in the next period and a gradually waning effect in subsequent periods, the latter of which can only be structurally estimated using our model. These constitute an additional adjustment cost the the firm incurs for lumpy adjustment, in contrast to smoother adjustment where only the typical convex costs may apply. Additionally, we calculate Aggregate Productivity Growth and its implied components according to each model. Again, we find significant differences.

Firm-level datasets may display patterns of endogenous productivity disruptions. If ignored, they can skew production function or other estimates and produce erroneous inference results, as we have seen here. Researchers should try to control for such disruptions.

Conflict of Interest Statement

The authors report there are no competing interests to declare.

References

Ackerberg, D., Benkard, L. C., Berry, S., & Pakes, A. (2007). Econometric Tools for Analyzing Market Outcomes. *Handbook of econometrics* (pp. 4171–4276). https://doi.org/10.1016/S1573-4412(07)06063-1

Ackerberg, D. A., Caves, K., & Frazer, G. (2015). Identification Properties of Recent Production Function Estimators. *Econometrica*, *83*(6), 2411–2451.

Balk, B. M. (2021). *Productivity* (1st ed.). Springer, Cham. https://doi.org/10.1007/978-3-030-75448-8

Basu, S., & Fernald, J. G. (1995). Are apparent productive spillovers a figment of specification error? *Journal of Monetary Economics*, *36*(1), 165–188. https://doi.org/10.1016/0304-3932(95)01208-6

Bitros, G. C., & Panas, E. E. (2001). Is there an inflation-productivity trade-off? Some evidence from the manufacturing sector in Greece. *Applied Economics*, *33*(15), 1961–1969. https://doi.org/10.1080/00036840110043730

Bond, S., & Söderbom, M. (2005). Adjustment Costs and the Identification of Cobb Douglas Production Functions. *The Institute for Fiscal Studies*, *Working Paper*(WP05/04).

Bruno, M. (1978). Duality, Intermediate Inputs and Value-Added. In M. Fuss and D. McFadden (Eds.), *Production economics: A dual approach to theory and* *applications* (pp. 3–16). McMaster University Archive for the History of Economic Thought. https://eml.berkeley.edu/prodecon/apps/ch8.pdf

Collard-Wexler, A., & De Loecker, J. (2016). *Production Function Estimation and Capital Measurement Error* (Working Paper No. 22437). National Bureau of Economic Research. https://doi.org/10.3386/w22437

Cooper, R. W., & Haltiwanger, J. C. (2006). On the Nature of Capital Adjustment Costs. *Review of Economic Studies*, *73*(3), 611–633.

De Loecker, J. (2013). Detecting Learning by Exporting. *American Economic Journal: Microeconomics*, *5*(3), 1–21. https://doi.org/10.1257/mic.5.3.1

De Loecker, J., Eeckhout, J., & Unger, G. (2020). The Rise of Market Power and the Macroeconomic Implications. *Quarterly Journal of Economics*, *135*(2), 561–644.

Diewert, W. (1978). Hick's Aggregation Theorem and the Existence of a Real Value Added Function. In M. Fuss and D. McFadden (Eds.), *Production economics: A dual approach to theory and applications*. McMaster University Archive for the History of Economic Thought. https://eml.berkeley.edu/prodecon/apps/ch9.pdf

Doraszelski, U., & Jaumandreu, J. (2013). R&D and Productivity: Estimating Endogenous Productivity. *The Review of Economic Studies*, *80*(4 (285)), 1338–1383. http://www.jstor.org/stable/43551561

Fakos, A., Sakellaris, P., & Tavares, T. (2022). Investment Slumps During Financial Crises: The Real Effects of Credit Supply. *Journal of Financial Economics*, *145*(1), 29–44. https://doi.org/10.1016/j.jfineco.2022.04.004

Gandhi, A., Navarro, S., & Rivers, D. A. (2020). On the Identification of Gross Output Production Functions. *Journal of Political Economy*, *128*(8), 2973–3016. https://doi.org/10.1086/707736

Gradzewicz, M. (2020). What Happens After an Investment Spike—Investment Events and Firm Performance. *Journal of Business and Economic Statistics*, *00*(0), 1–16. https://doi.org/10.1080/07350015.2019.1708369

Hahn, J., & Liao, Z. (2021). Bootstrap Standard Error Estimates and Inference. *Econometrica*, *89*(4), 1963–1977. https://doi.org/10.3982/ECTA17912

Horowitz, J. L. (2001). The Bootstrap. In J. Heckman and E. Leamer (Eds.), *Handbook of Econometrics* (1st ed., pp. 3159–3228). Elsevier.

Hu, Y., Huang, G., & Sasaki, Y. (2020). Estimating production functions with robustness against errors in the proxy variables. *Journal of Econometrics*, *215*(2), 375–398. https://doi.org/10.1016/j.jeconom.2019.05.024

Hugget, M., & Ospina, S. (2001). Does Productivity Growth Fall After the Adoption of New Technology? *Journal of Monetary Economics*, *48*(1), 173–195. https://doi.org/10.1016/S0304-3932(01)00065-4

İmrohoroğlu, A., & Tüzel, Ş. (2014). Firm-Level Productivity, Risk, and Return. *Management Science*, *60*(8), 1861–2109.

Kalemli-Ozcan, S., Sorensen, B., Villegas-Sanchez, C., Volosovych, V., & Yesiltas, S. (2015). *How to Construct Nationally Representative Firm Level Data from the Orbis Global Database: New Facts and Aggregate Implications* (Working Paper No. 21558). National Bureau of Economic Research. https://doi.org/10.3386/w21558

Keller, W., & Yeaple, S. R. (2009). Multinational Enterprises, International Trade, and Productivity Growth: Firm-Level Evidence from the United States. *The Review of Economics and Statistics*, *91*(4), 821–831.

Khan, S. Y., & Khederlarian, A. (2021). *Inventories, Input Costs, and Productivity Gains from Trade Liberalizations* (Working Paper No. 9564). World Bank. https://openknowledge.worldbank.org/handle/10986/35213

Kim, K. i., Petrin, A., & Song, S. (2016). Estimating production functions with control functions when capital is measured with error. *Journal of Econometrics*, *190*(2), 267–279. https://doi.org/10.1016/j.jeconom.2015.06.016

Levinsohn, J., & Petrin, A. (2003). Estimating Production Functions Using Inputs to Control for Unobservables. *Review of Economic Studies*, *70*(2), 317–342.

Manjón, M., & Mañez, J. (2016). Production function estimation in Stata using the Ackerberg–Caves–Frazer method. *The Stata Journal*, *16*(4), 900–916. https://doi.org/10.1177/1536867X1601600406

Olley, S., & Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, *64*(6), 1263–1297.

Pakes, A., & Olley, S. (1995). A limit theorem for a smooth class of semiparametric estimators. *Journal of Econometrics*, *65*(1), 295–332.

Petrin, A., & Levinsohn, J. (2012). Measuring aggregate productivity growth using plant-level data. *RAND Journal of Economics*, *43*(4), 705–725. https://doi.org/10.1111/1756-2171.12005

Power, L. (1998). The Missing Link: Technology, Investment, and Productivity. *The Review of Economics and Statistics*, *80*(2), 300–313. https://doi.org/10.1162/003465398557393

Rosenbaum, P. R., & Rubin, D. B. (1985). Constructing a Control Group Using Multivariate Matched Sampling Methods That Incorporate the Propensity Score. *The American Statistician*, *39*(1), 33–38. https://doi.org/10.2307/2683903

Rovigatti, G. (2017). Production Function Estimation in R: The prodest Package. *The Journal of Open Source Software*, *2*(18). https://doi.org/10.21105/joss.00371

Rovigatti, G., & Mollisi, V. (2018). Theory and Practice of TFP Estimation: the Control Function Approach Using Stata. *The Stata Journal*, *18*(3), 618–662. https://doi.org/10.1177/1536867X1801800307

Sakellaris, P. (2004). Patterns of Plant Adjustment. *Journal of Monetary Economics*, *51*(2), 425–450.

Shima, K. (2010). Lumpy Capital Adjustment and Technical Efficiency. *Economics Bulletin*, *30*(4), 2817–2824.

Stehrer, R., Bykova, A., Jäger, K., Reiter, O., & Schwarzhappel, M. (2019). *Industry Level Growth and Productivity Data with Special Focus on Intangible Assets: Report on methodologies and data construction for the EU KLEMS Release 2019* (No. 8). The Vienna Institute for International Economic Studies (wiiw).

Wooldridge, J. M. (2009). On estimating firm-level production functions using proxy variables to control for unobservables. *Economics Letters*, *104*(3), 112–114. https://doi.org/10.1016/j.econlet.2009.04.026

Figures

Figure 1: Baseline $log(TFP)$ averages around investment spikes

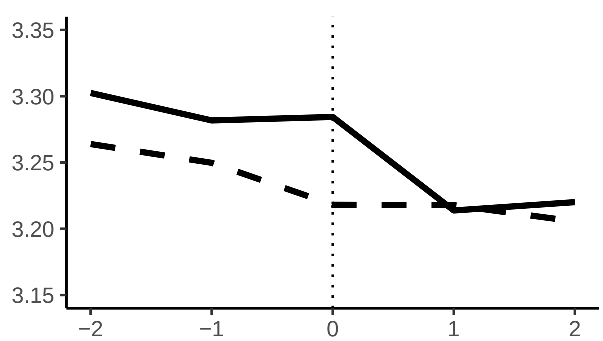

$N = 591$

The solid lines report for the spike sub-sample and the dashed lines report for the matched sub-sample.

Figure 2: Productivity component averages around investment spikes

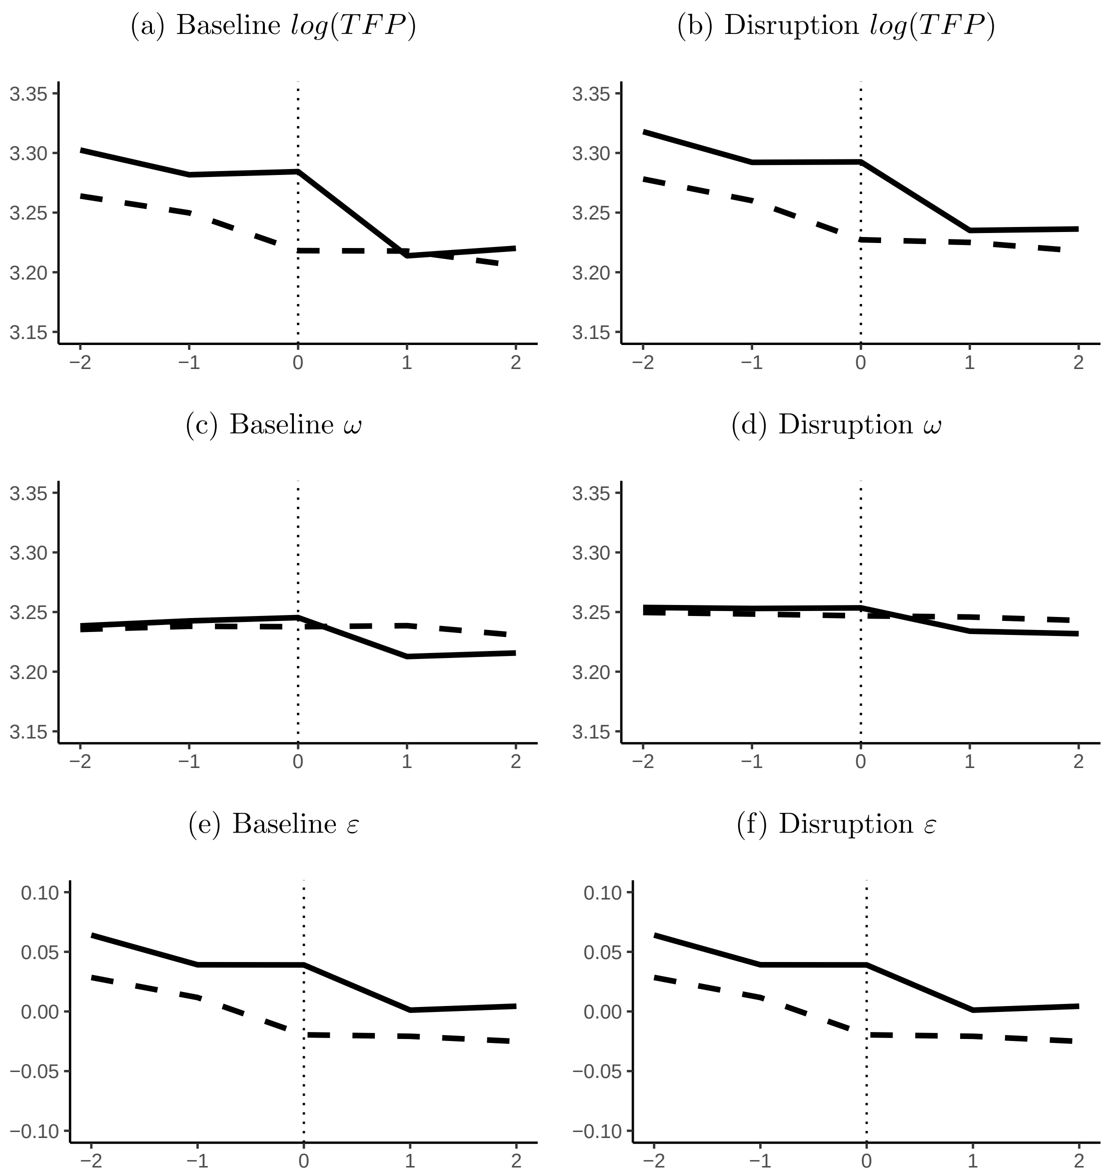

$N = 591$

The solid lines report for the spike sub-sample and the dashed lines report for the matched sub-sample.

A Tables

Table A.1 reports for the inference sample, per year for the pre-crisis, crisis, and recovery periods: Aggregate Productivity Growth (APG), calculated according to Petrin and Levinsohn (2012), and the average growth rates of the output and input variables (capital, labor, and other intermediate inputs), calculated as the average of their first differences in logs.

Table A.2 reports, per year for the inference sample, the average investment rate (i.e. average I_{it}/K_{it} across i for each t) and the percentage of observations identified as investment spikes (i.e. when $d_{it} = 1$). Observe the noticeable drop in average investment behavior after the onset of the Greek financial crisis in 2008. Our definition of lumpy investment appears unaffected, however, because it takes the financial regime into account (pre 2008 or not). In fact, one could notice some increase in lumpy behavior when approaching the recovery period (circa post 2015).

Table A.3 reports Aggregate Productivity Growth (APG) and its components, technical efficiency (TE) and reallocation efficiency (RE), for the baseline model. RE is further broken down into its summands, RE due to capital ($\mathrm{RE_K}$), RE due to labor ($\mathrm{RE_L}$), and RE due to other intermediate inputs ($\mathrm{RE_M}$). Table A.4 reports the same for the disruption model.

Table A.5 reports, per year for the inference sample, technical efficiency and reallocation efficiency as a percentage of APG (TE/APG and RE/APG, respectively) for the baseline and the disruption model, and their ratio between models (as disruption over baseline). Notice how for each grouping TE and RE sum up to one.

Table A.1: Aggregate Productivity Growth and Average Growth Rates

Year	APG	Output	Capital	Labour	Int. Inputs
2001	-0.0299	0.0017	0.1381	0.1301	0.0536
2002	0.0148	0.0181	0.0919	0.0881	0.0335
2003	0.0007	0.0366	0.0971	0.0553	0.0572
2004	-0.0136	0.0381	0.1451	0.0751	0.0624
2005	0.0241	-0.0123	0.0980	-0.0863	-0.0650
2006	-0.0083	0.0300	0.0900	0.0162	0.0162
2007	0.0234	0.0668	0.1071	0.0254	0.0225
2008	-0.0418	-0.0175	0.1055	0.0201	0.0152
2009	0.0773	-0.1709	0.0380	-0.0785	-0.2302
2010	-0.0957	-0.0990	0.0234	-0.0837	-0.1373
2011	0.0020	-0.1581	-0.0127	-0.0984	-0.1379
2012	-0.0049	-0.1506	-0.0212	-0.2003	-0.1609
2013	-0.0148	-0.0257	0.0550	-0.0048	-0.0105
2014	0.0241	0.0499	0.0271	0.2750	0.0448
2015	0.0264	0.0283	0.1329	0.0595	0.0421
2016	0.0259	0.0103	0.0219	-0.0009	0.0497
2017	-0.0083	0.0118	0.0564	0.0541	0.0246
Average	0.0001	-0.0201	0.0702	0.0145	-0.0188
SD	0.0368	0.0762	0.0513	0.1071	0.0915
Average 2001-2007	0.0016	0.0256	0.1096	0.0434	0.0258
Average 2008-2015	-0.0034	-0.0680	0.0435	-0.0139	-0.0718
Average 2016-2017	0.0088	0.0111	0.0392	0.0266	0.0372

Aggregate Productivity Growth is calculated based on Petrin and Levinsohn (2012). The average growth rates for the rest of the variables are calculated as the average of their first differences in logs.

Table A.2: Average Investment Rate and Proportion of Investment Spikes per year

Year	Investment Rate	Investment Spike Percentage
2001	0.2188	0.0913
2002	0.2332	0.0841
2003	0.3013	0.0795
2004	0.2337	0.0868
2005	0.2166	0.1004
2006	0.2492	0.0800
2007	0.2419	0.0825
2008	0.1532	0.0789
2009	0.1369	0.0758
2010	0.0850	0.0697
2011	0.0802	0.0692
2012	0.1649	0.0640
2013	0.1374	0.0688
2014	0.3073	0.0953
2015	0.1475	0.1017
2016	0.1702	0.0887

Table A.3: Aggregate Productivity Growth breakdown according to the Baseline Model

Year	APG	TE	RE	RE_K	RE_L	RE_M
2001	-0.0299	-0.0372	0.0073	0.0095	-0.0020	-0.0003
2002	0.0148	0.0060	0.0088	0.0055	-0.0017	0.0050
2003	0.0007	-0.0006	0.0013	0.0069	0.0009	-0.0065
2004	-0.0136	-0.0169	0.0033	0.0143	-0.0024	-0.0085
2005	0.0241	0.0363	-0.0122	0.0044	-0.0120	-0.0045
2006	-0.0083	-0.0279	0.0195	0.0051	0.0144	0.0000
2007	0.0234	0.0262	-0.0028	0.0053	-0.0083	0.0002
2008	-0.0418	-0.1044	0.0626	0.0081	0.0687	-0.0142
2009	0.0773	0.0340	0.0433	0.0090	0.0010	0.0333
2010	-0.0957	-0.0630	-0.0327	0.0112	-0.0103	-0.0336
2011	0.0020	0.0238	-0.0217	0.0002	-0.0198	-0.0022
2012	-0.0049	-0.0351	0.0302	0.0004	0.0126	0.0172
2013	-0.0148	-0.0358	0.0210	0.0009	0.0106	0.0095
2014	0.0241	0.0558	-0.0317	0.0008	-0.0114	-0.0210
2015	0.0264	0.0160	0.0104	0.0022	0.0074	0.0008
2016	0.0259	0.0265	-0.0006	0.0008	-0.0002	-0.0012
2017	-0.0083	-0.0047	-0.0036	0.0020	0.0012	-0.0068
Average	0.0001	-0.0059	0.0060	0.0051	0.0029	-0.0019
SD	0.0368	0.0411	0.0248	0.0042	0.0193	0.0146
Average 2001-2007	0.0016	-0.0020	0.0036	0.0073	-0.0016	-0.0021
Average 2008-2015	-0.0034	-0.0136	0.0102	0.0041	0.0073	-0.0013
Average 2016-2017	0.0088	0.0109	-0.0021	0.0014	0.0005	-0.0040

The estimates for fixed and sunk costs (F) are not reported, because they are zero everywhere.

Table A.4: Aggregate Productivity Growth breakdown according to the Disruption Model

Year	APG	TE	RE	RE_K	RE_L	RE_M
2001	-0.0299	-0.0402	0.0103	0.0067	0.0039	-0.0003
2002	0.0148	0.0042	0.0106	0.0039	0.0017	0.0050
2003	0.0007	-0.0037	0.0044	0.0049	0.0061	-0.0065
2004	-0.0136	-0.0131	-0.0005	0.0092	-0.0012	-0.0085
2005	0.0241	0.0613	-0.0371	0.0031	-0.0357	-0.0045
2006	-0.0083	-0.0360	0.0277	0.0036	0.0241	0.0000
2007	0.0234	0.0163	0.0071	0.0039	0.0030	0.0002
2008	-0.0418	-0.1032	0.0614	0.0060	0.0695	-0.0142
2009	0.0773	0.0298	0.0474	0.0063	0.0078	0.0333
2010	-0.0957	-0.0400	-0.0557	0.0088	-0.0309	-0.0336
2011	0.0020	0.0402	-0.0382	0.0003	-0.0363	-0.0022
2012	-0.0049	-0.0286	0.0236	0.0004	0.0060	0.0172
2013	-0.0148	-0.0493	0.0345	0.0007	0.0243	0.0095
2014	0.0241	0.0989	-0.0748	0.0007	-0.0545	-0.0210
2015	0.0264	0.0078	0.0186	0.0013	0.0164	0.0008
2016	0.0259	0.0263	-0.0004	0.0007	0.0001	-0.0012
2017	-0.0083	-0.0078	-0.0005	0.0014	0.0049	-0.0068
Average	0.0001	-0.0022	0.0023	0.0036	0.0005	-0.0019
SD	0.0368	0.0472	0.0359	0.0030	0.0285	0.0146
Average 2001-2007	0.0016	-0.0016	0.0032	0.0050	0.0003	-0.0021
Average 2008-2015	-0.0034	-0.0055	0.0021	0.0031	0.0003	-0.0013
Average 2016-2017	0.0088	0.0092	-0.0004	0.0011	0.0025	-0.0040

The estimates for fixed and sunk costs (F) are not reported, because they are zero everywhere.

Table A.5: Aggregate Productivity Growth breakdown as percentage of annual level and ratios between models

	Baseline		Disruption		Ratio	
Year	TE	RE	TE	RE	TE	RE
2001	1.2437	-0.2437	1.3449	-0.3449	1.0813	1.4151
2002	0.4050	0.5950	0.2843	0.7157	0.7020	1.2029
2003	-0.8038	1.8038	-5.1678	6.1678	6.4293	3.4194
2004	1.2443	-0.2443	0.9628	0.0372	0.7737	-0.1523
2005	1.5045	-0.5045	2.5396	-1.5396	1.6879	3.0514
2006	3.3423	-2.3423	4.3176	-3.3176	1.2918	1.4164
2007	1.1183	-0.1183	0.6948	0.3052	0.6213	-2.5805
2008	2.4987	-1.4987	2.4699	-1.4699	0.9885	0.9808
2009	0.4398	0.5602	0.3862	0.6138	0.8783	1.0955
2010	0.6581	0.3419	0.4182	0.5818	0.6355	1.7014
2011	11.6695	-10.6695	19.7365	-18.7365	1.6913	1.7561
2012	7.0980	-6.0980	5.7782	-4.7782	0.8141	0.7836
2013	2.4176	-1.4176	3.3319	-2.3319	1.3782	1.6449
2014	2.3157	-1.3157	4.1055	-3.1055	1.7729	2.3604
2015	0.6050	0.3950	0.2968	0.7032	0.4905	1.7805
2016	1.0232	-0.0232	1.0144	-0.0144	0.9914	0.6212
2017	0.5627	0.4373	0.9404	0.0596	1.6711	0.1363
Average 2001-2007	-1.2515	2.2515	-1.0016	2.0016	0.8003	0.8890
Average 2008-2015	3.9585	-2.9585	1.6145	-0.6145	0.4079	0.2077
Average 2016-2017	1.2393	-0.2393	1.0492	-0.0492	0.8466	0.2055

The columns under Baseline and Disruption report the ratios of TE/APG and RE/APG in each time period for the baseline and the disruption model, respectively. The columns under Ratio report the ratio of the disruption estimate over the corresponding baseline estimate.

B Variable Construction

We try to recover any missing information in our financial statement data based on accounting identities[12]. We then construct the variables we need as explained below. Finally, we take some additional steps to reach the final dataset we use for inference.

The inference variables are constructed as follows:

Firm Age, as the difference between the current year of each observation and the Year of Establishment reported for each firm.

Physical Capital and Investment Deflators (base year 2010), as the Implicit Price Index for Gross Fixed Capital Formation, provided by the World Bank.

Geometric Depreciation Rates by Asset Class per Year, δ_{jit}, as the weighted average of the geometric depreciation rates for each subclass of capital, provided by the EU KLEMS 2019 survey, weighted by the gross capital stocks reported by Eurostat (nama10 file) every year and for every 2-digit NACE 2 division. If data for a division are not reported, the data for the immediate parent grouping available are used. If some δ_{jit} is undefined due to a zero over zero division, this means that the amount particular asset class for that particular time and sector is zero, and thus δ_{jit} set to zero. Then firms are assigned the proper δ_{jit} according to their NACE 2 classification.

[12] We set to missing any value that is reported as negative where negative values do not apply. The only financial variables that we allow to be negative are: Cash and Deposits, Shareholders' Equity, Retained Earnings/Accumulated Losses, Gross Margin, Operating Margin, Profits Before Taxes, Income Taxes, and EBITDA. To account for rounding errors, we set every instance of the value of negative one to zero.

Deflated Book Value of Capital per Asset Class[13], as

$$K_{Book\ jit} = \frac{Book_{jit} - AccumulatedDepreciation_{jit}}{IPIGFCF_t}, \forall i, j, t \tag{B.1}$$

where $j \in \{Buildings,\ Machinery\ and\ Equipment\}$ and $IPIGFCF$ is the Implicit Price Index for Gross Fixed Capital Formation for Greece provided by the World Bank per asset class per year.

Real Investment per Asset Class[14], as

$$I_{jit} = \frac{\begin{array}{l}\Delta\left(Book_{jit+1} - AccumulatedDepreciation_{jit+1}\right)\\ + CurrentDepreciation_{jit}\end{array}}{IPIGFCF_t}, \forall i, j, t \tag{B.2}$$

where $j \in \{Buildings,\ Machinery\ and\ Equipment\}$.

Real Capital per Asset Class, according to the Perpetual Inventory Method (PIM) where for consecutive years

$$K_{jit} = I_{jit} + (1 - \delta_{jit})K_{jit-1}, \forall i, j, t \tag{B.3}$$

and K_{jit} is initialized using the corresponding Deflated Book Value, $K_{Book\ jit}$, when the firm first appears in the sample or if there is a skip in the series.

Aggregate Physical Capital and Aggregate Investment in Physical Capital, as the respective sums of Buildings, and Machinery and Equipment, $\sum_{j\in\{Buildings, Machinery\ and\ Equipment\}} K_{jit}$ and $\sum_{j\in\{Buildings, Machinery\ and\ Equipment\}} I_{jit}$.

[13] Here, we could have lagged the deflator by the average age of each asset class, as in İmrohoroğlu and Tüzel (2014). However, in our case the estimates for capital age were impossibly high for a considerable number of cases, so we opted not to use it.

[14] Our data do not report Current Depreciation for each asset class, but only in total. We decompose this to its per asset class components as best as possible, using a series of reasonable assumptions (e.g. setting current depreciation to zero when the associated book value is zero, etc).

Real Gross Output, as Net Sales plus Other Operating Income and the change in Inventories of Finished Goods deflated by the Producer Price Index (PPI) of the corresponding 2-digit NACE rev. 2 division (base year 2010) provided to us by ELSTAT. For a few cases where PPI is not available, we use Eurostat's Harmonized Consumer Price Index.

Labor, as the reported number of employees times the year and sector specific ratio of persons employed to number of employees reported by Eurostat. If possible, we fill any missing ratios by linear interpolation. We fill the rest with the ratio of persons employed to employees reported by the EU KLEMS survey. Because self employment may be reported as zero employees (as opposed to missing), we first set observations reporting zero employees to one for non S.A. firms.

Intermediate Inputs, according to Keller and Yeaple (2009) and De Loecker, Eeckhout, and Unger (2020), as Cost of Goods Sold and Other Operating Expenses less Depreciation in Cost of Finished Goods and wage expenditures. Wage expenditures are calculated as the Average Cost of Labor for the corresponding division and year times the original number of employees reported by the firm at that year. We get the average wage by dividing the total cost of employees by the number of employees, reported by Eurostat. If possible, we fill any missing average cost by linear interpolation. We fill the rest with the ratio of the EU KLEMS item COMP over EMPE.

Leverage, as $1 - \frac{TotalEquity}{TotalAssets}$.

We drop observations with negative values for age, output, real Buildings, real Machinery and Equipment, and intermediate inputs. We set to missing any investment rate that is below -1 and, following Gradzewicz (2020), we also set to missing any investment rate above 12, as unlikely.

C Estimation

The estimation procedure of the models in Section 3 is a simplified version of Gandhi, Navarro, and Rivers (2020), as described in the appendix of Collard-Wexler and De Loecker (2016) (less the instrumentation of log capital by log investment to control for measurement errors). The estimation process is identical for both models, with the difference that in the disruption model, g takes a not only lagged productivity as an input, but also a lagged investment spike dummy. We base our computer code for the estimation on the `prodest` R package by Rovigatti (2017). Here we describe the precise steps that yield the results in table 4.

Consider the Cobb-Douglas production function in logs in equation (2)

$$y_{it} = \beta_K k_{it} + \beta_L l_{it} + \beta_M m_{it} + \omega_{it} + \varepsilon_{it} \tag{2}$$

Produce an estimate of β_M as the median of the empirical distribution for the revenue share of intermediate inputs

$$\widehat{\beta_M} = median\left(\frac{p_{it}^M M_{it}}{P_{it} Y_{it}}\right) \tag{C.1}$$

where the capitalized variables, Y_{it} and M_{it}, are output and intermediate inputs in levels instead of logs. P_{it} is the price firm i receives at time t for its output and p_{it}^M the corresponding price it pays for the inputs. In practice, they are the corresponding deflators used when constructing the variables.

Subtract the effect of intermediate inputs from output to get

$$\widetilde{y}_{it} = y_{it} - \widehat{\beta_M} m_{it} \tag{C.2}$$

and use $\widetilde{y}_{it}$ in the Ackerberg, Caves, and Frazer (2015) framework to get estimates for β_K and β_L in the following specification

$$\widetilde{y}_{it} = \beta_K k_{it} + \beta_L l_{it} + \omega_{it} + \varepsilon_{it} \tag{13}$$

This allows for the unique identification of all Cobb-Douglas coefficients.

The Ackerberg et al. (2015) framework is as follows.

For the first stage, proxy for the unobservable ω_{it} by inverting f and substituting in equation (2)

$$\begin{aligned}\widetilde{y}_{it} &= \beta_K k_{it} + \beta_L l_{it} + \omega_{it} + \varepsilon_{it} \\ &= \beta_K k_{it} + \beta_L l_{it} + f^{-1}(k_{it}, l_{it}, m_{it}) + \varepsilon_{it} \\ &= \Phi(k_{it}, l_{it}, m_{it}) + \varepsilon_{it} \qquad \text{(C.3)}\end{aligned}$$

so that the production function can be written as an unknown function of observable variables, k_{it}, l_{it}, and m_{it}, and the unobservable shock, ε_{it}. To deal with the unknown function, Φ, approximate it with a linear function of its inputs. Because $E(k_{it}\varepsilon_{it}|\mathcal{I}_{it}) = E(l_{it}\varepsilon_{it}|\mathcal{I}_{it}) = E(m_{it}\varepsilon_{it}|\mathcal{I}_{it}) = 0$, this approximation can be estimated by OLS.

Then, take the fitted values of the regression, $\widehat{\Phi}_{it}$, to use in the second stage. This, essentially, purges ε_{it} from output.

In the second stage, choose candidates for $\widehat{\beta_K}$ and $\widehat{\beta_L}$ and get an estimate measure of ω_{it}

$$\widehat{\omega}_{it} = \widehat{\Phi}_{it} - \widehat{\beta_K} k_{it} - \widehat{\beta_L} l_{it} \qquad \text{(C.4)}$$

The next step is to estimate g. For our AR(1) process for the baseline model we get

$$\omega_{it} = c + \rho\omega_{it-1} + \xi_{it} \qquad (4)$$

and for the disruption model we get

$$\omega_{it} = c + \rho\omega_{it-1} + \delta d^I_{it-1} + \xi_{it} \qquad (5)$$

At this point, a selection correction term can be included as a regressor in the estimation of (4) and (5).

The model assumptions in Section 3 imply that $E(k_{it}\xi_{it}) = E(l_{it-1}\xi_{it}) = 0$.

Thus, use $\widehat{\omega}_{it}$ (and any other data necessary) to estimate either (4) or (5) by OLS, depending on the chosen model, and use the residuals, $\widehat{\xi}_{it}$ to form the above moment conditions. Essentially, k_{it} and l_{it-1} act as instruments of a non-linear GMM specification given by (C.4) and (4) or (C.4) and (5), respectively. Form the GMM objective as

$$J = \widehat{\xi}' Z' W Z \widehat{\xi} \tag{C.5}$$

where $Z = [k\, l_{-1}]$ and W a weighting matrix. We set $W = (Z'Z)^{-1}$.

We optimize the objective using the Newton-Raphson algorithm, but any appropriate optimization routine can equally be used.