

Demand or productivity: What determines firm growth?*

Andrea Pozzi
EIEF

Fabiano Schivardi
University of Cagliari,
EIEF and CEPR

October 11, 2012

Abstract

We disentangle the contribution of unobserved heterogeneity in idiosyncratic demand and productivity to firm growth. We use a model of monopolistic competition with Cobb-Douglas production and a dataset of Italian manufacturing firms containing unique information on firm-level prices to reach three main conclusions. First, demand shocks are at least as important as productivity shocks for firm growth. Second, firms respond to shocks less than a frictionless model would predict, suggesting the existence of adjustment frictions. Finally, the degree of under-response is much larger for TFP shocks. This implies the existence of frictions with differential effects according to the nature of the shock, unlike the typical frictions studied by the literature on factor misallocation. We consider hurdles to firm reorganization as one such friction and show that they hamper firms' responses to TFP shocks but not to demand shocks.

JEL classification: D24, L11.

Key words: TFP, demand heterogeneity, firm growth, misallocation.

*We have benefitted from comments by Fernando Alvarez, Fabio Canova, Greg Crawford, Jan De Loecker, Alon Eizenberg, Bob Hall, John Haltiwanger, Hugo Hopenhayn, Rosa Matzkin, Ariel Pakes, Ed Vytlačil. We are also grateful to participants at the following seminars, workshops and conferences: Cagliari, EIEF, EUI, IFN, Oxford, Pompeu Fabra, Scuola Superiore Sant'Anna, Stockholm School of Economics, Mannheim, Central European University, Workshop on the productivity of the Italian system, EIEF-UNIBO Workshop on Industrial Organization, the 2nd EIEF workshop on structural approaches to productivity and industry dynamics, the 11th CAED Conference, the CEPR/JIE conference 2012, the SED Annual Meeting 2012, and the EARIE 2012.

1 Introduction

Modern theories of industry dynamics assume that firms are heterogeneous along a single unobserved dimension, productivity, which determines the firm's performance and growth (Jovanovic 1982, Hopenhayn 1992). The empirical literature on the topic has followed this view, tracing back firms' growth to the evolution of productivity (see Syverson (2011) for a comprehensive survey). However, several other dimensions of heterogeneity may make a difference to growth. In particular, the assumption that all firms look alike to consumers fails to capture an important ingredient of firm performance. Differences in the effectiveness in marketing, in developing relationships with customers, in maintaining brand image and in generating word-of-mouth are only some potential elements leading to heterogeneity across firms on the demand side. There is no reason to believe that demand factors are less important than productive efficiency in shaping a firm's success and its growth. For example, in many sectors marketing and advertising budgets are larger than research and development ones.

To study the relative importance of demand and productivity in determining firm growth, we model firms as characterized by two unobserved idiosyncratic variables, *market appeal* and *TFP*, that shift the demand and the production function respectively. As first pointed out by Klette and Griliches (1996), not accounting for heterogeneity in demand leads to productivity estimates that are a mix of true productivity and demand effects. However, distinguishing between demand-side and TFP shocks matters for reasons beyond those of simple measurement. We show that heterogeneity in market appeal is an interesting dimension to study in its own right and that it is quantitatively important. Furthermore, new insights can be derived from jointly considering two types of unobserved heterogeneity, which could not be captured in the standard scalar heterogeneity framework. Specifically, we find that the reallocation of factors of production following changes in productivity or demand appeal is imperfect, generating misallocation of resources, and that distortions in reallocation are more severe after productivity shocks. This departs from the approach of the previous literature, where the frictions that generate misallocation, such as firing costs (Hopenhayn and Rogerson 1993) or bribes and political favoritism (Hsieh and Klenow 2009), have effects that are independent from the nature of the shock.

The relevance of demand factors in shaping industry dynamics is hardly disputable. Foster, Haltiwanger and Syverson (2008) were the first to document its importance, showing that heterogeneity in demand affects firms' chances of survival. Empirical evidence on the relationship between idiosyncratic demand and firm performance is, however, still scant. In

fact, identifying this component requires firm-level price data typically not available in the datasets used to study firm performance. We use a survey of a representative panel of Italian manufacturing firms with at least 50 employees (INVIND), administered yearly by the Bank of Italy since 1984. Among other things, firms are asked about the average percentage change in prices of goods and services sold; this allows us to identify and distinguish market appeal and TFP.

To flesh out the assumptions needed for correctly identifying the two shocks, we set up a standard model of monopolistic competition on the demand side and Cobb-Douglas technology on the production side, each with its own stochastic shifter. We begin by backing out the unobserved demand component as the residual of the demand equation. We circumvent the usual simultaneity problem in demand estimation (Trajtenberg 1989, Berry 1994) by using a direct assessment of the elasticity of demand provided by the managers in the survey. Productivity shocks are then identified as residuals of the production function equation. To address the endogeneity of input choice, we adapt the Olley and Pakes (1996) procedure to accommodate for non scalar unobserved heterogeneity.

Armed with the estimates of demand and productivity shocks, we study their effects on firms' growth. We can do this simply by regressing measures of output and inputs growth on the estimated shocks, since we treat both demand appeal and TFP as exogenous processes. The exercise reveals that demand factors play an important role. An increase of one standard deviation in market appeal generates a 13% increase in nominal sales, against 8% for TFP. As expected, productivity enhancements also lead to a decrease in prices, while positive demand shocks trigger price increases. Finally, TFP shocks have negligible impact on inputs (number of hours worked, capital used in production and intermediate goods) while demand shocks trigger changes in inputs usage.

Next we turn to our model for guidance in evaluating these findings. In fact, given the estimates of the parameters of the demand and production functions, our theoretical framework delivers quantitative predictions on the impact of the shocks on firms' growth. We contrast the figures implied by the model with those emerging from the empirical exercise. The comparison offers two main insights. First, the model predicts elasticities larger than those estimated in the reduced form regressions. This suggests the existence of adjustment frictions not accounted for by our theoretical framework. While highlighting that the introduction of frictions is important to explain the results, we also show that their presence does not invalidate our estimation procedure. Second, and more surprisingly, we find that the deviation between the model's predictions and the reduced form regressions is much larger

for TFP than for demand shocks. This indicates that adjustment costs have differential effects according to the nature of the shock.

Finally we investigate a potential cause of differential responses to demand and TFP shocks. We start from the premise that re-optimizing input choice can be more complicated following a TFP shock than after a demand shock. Whereas the latter involves simply scaling up operations, moving along a given production function, the former is triggered by a shift in the production technology itself. Responding to such a shift might entail some reorganization of business operations: a different skill mix, different types of capital inputs etc. These are more complex tasks than proportionally changing all inputs. In order to test this hypothesis we use heterogeneity in firms' organizational flexibility and managerial quality. We show that firms reporting in the survey that they are more constrained in reorganizing the production flow are also less responsive to TFP shocks but not to demand shocks. The same is true when comparing family firms to firms controlled by a financial institution or by a conglomerate. In fact, there is ample evidence that family firms tend to be characterized by less efficient managerial practices (Bloom and Van Reenen 2007) and could therefore be less effective in managing the reorganization and restructuring activities entailed by TFP shock.

This study relates to a vast literature interested in understanding the determinants of firm growth (Dunne, Roberts and Samuelson 1988, Dunne, Roberts and Samuelson 1989, Evans 1987*a*, Evans 1987*b*). We expand this literature by considering multiple sources of unobserved heterogeneity. The importance of disentangling demand and productivity heterogeneity has been stressed by recent literature. Foster et al. (2008) use data on homogeneous products, for which quantities can be meaningfully defined, to derive a price index from the value of sales and physical production. They show that failing to disentangle demand and TFP shocks leads to an underestimation of new entrants' contribution to productivity growth. De Loecker (2011) exploits theoretical restrictions to isolate physical productivity from confounding demand factors in estimating the effects of trade barriers on productivity for Belgian textile firms. We advance this literature by considering jointly demand and production heterogeneity in the context of firm growth. Our study is one of a number of recent contributions that take advantage of opportunities to observe firms' prices directly (De Loecker, Goldberg, Khandelwal and Pavcnik 2012, Fan, Roberts, Xu and Zhang 2012). We exploit it to consider industries characterized by product differentiation, without relying entirely on functional form for identification, as we are forced to do when firms' prices are not available. Finally, we contribute to the literature on the inefficient allocation of resources

across firms (Hsieh and Klenow 2009, Midrigan and Xu 2010, Collard-Wexler, De Loecker and Asker 2011, Yang 2011, Bartelsman, Haltiwanger and Scarpetta Forthcoming). Our results imply that factors are not allocated efficiently across firms, in line with the findings of this growing literature. The larger deviation for productivity than for demand shocks, however, suggests the existence of barriers to the efficient allocation of resources that cannot be only of regulatory nature, as studies on this subject typically assume (Hopenhayn and Rogerson 1993, Restuccia and Rogerson 2008, Hsieh and Klenow 2009). Those types of hurdles should, in fact, affect in the same way the two shocks. Our results call for the introduction of frictions that have differential effects on the response to demand and TFP shocks.

The rest of the paper is organized as follows. Section 2 presents a standard model of a monopolistic competitive firm characterized by demand and productivity shifters. Section 3 introduces the data, while Section 4 presents the estimation approach. Section 5 discusses the effects of the shocks on firm growth and points out their divergence from the theoretical predictions of the model. Section 6 analyzes the implications of our findings for misallocation and proposes hurdles to reorganization as an example of a friction that generates it. Section 7 concludes.

2 The model

Our theoretical framework relies on a model of monopolistic competition where firms choose inputs to produce output, subject to a CES demand and a Cobb-Douglas production function as in Melitz (2000). This standard setup serves the empirical analysis along three dimensions. Firstly, it formalizes the assumptions needed for consistently estimating the parameters of the production and the demand functions. Secondly, it illustrates the consequences for estimated productivity of ignoring firm prices. Finally, it supplies a benchmark against which to evaluate the results of the growth regressions we will perform in the second part of the paper.

Firm i faces a constant elasticity of demand function:

$$Q_{it} = P_{it}^{-\sigma} \Xi_{it} \tag{1}$$

where $\sigma > 1$ is the elasticity of demand and Ξ_{it} is a demand shifter, observed by the firm (but not by the econometrician) when choosing output. Other time-specific factors, constant across firms, can be ignored without loss of generality as they will be captured by time dummies in the empirical specification.

The *market appeal* component (Ξ_{it}) picks up heterogeneity in firms' demand driven by differences in the perceived quality of the product, controlling for its physical attributes. It relates to similar concepts introduced by Foster, Haltiwanger and Syverson (2012) and Gourio and Rudanko (2012) who link it to the stock of consumers who have tried the product in the past (the "customer base"). Other instances of demand shocks consistent with our setting are spreading of good word-of-mouth, improvements in the brand image and the perception or the visibility of the products, for example as a result of advertising.

The firm enters the period with a given level of capital stock \bar{K}_{it} , accumulated through investment up to period $t - 1$:

$$\bar{K}_{it} = (1 - \delta)\bar{K}_{it-1} + I_{it-1} \quad (2)$$

where δ is the depreciation rate. Although the firm cannot modify the capital stock in place for the current period, it decides the degree of capital utilization U_{it} . The effective capital used for production is then:

$$K_{it} = U_{it}\bar{K}_{it}, \quad 0 \leq U_{it} \leq 1. \quad (3)$$

We assume that using capital is costly¹ so that it may be optimal to use less than the whole installed capacity. For simplicity, we assume that capital depreciation is independent from usage.² The firm produces output combining utilized capital, intermediate inputs and labor with a Cobb-Douglas production function

$$Q_{it} = \Omega_{it}K_{it}^\alpha L_{it}^\beta M_{it}^\gamma \quad (4)$$

where Ω_{it} is firm TFP, observed before choosing inputs. Labor (L) and intermediates (M) can be chosen freely and have no dynamic implications, whereas capital input can be varied through the degree of utilization, up to full utilization. Given \bar{K}_{it} and after observing Ω_{it}, Ξ_{it} , the firm chooses inputs to maximize profits:

$$\text{Max}_{\{K_{it}, L_{it}, M_{it}\}} P_{it}Q_{it} - p_K K_{it} - p_L L_{it} - p_M M_{it} \quad (5)$$

subject to the demand equation (1), the capital constraint (3) and the production function (4), where p_X is the cost of utilizing input X .

¹For example, if capital must be used in a fixed proportion $1/a$ with energy, and the price of energy is p^e ; then the cost of using capital is defined as $r = a * p^e$, where we use r as the standard notation for the cost of capital usage.

²For some types of capital, such as buildings, this seems the most natural assumption. In general, a component of depreciation is clearly linked to time, independently from usage. Moreover, when capital is used it might be easier to maintain it in an efficient state.

Assume first that the capital constraint is not binding. In this case, equilibrium quantities do not depend on the capital stock in place. Using lowercase letters to denote logs, the optimal quantity, price and inputs demand functions are:

$$q_{it}^* = c_q + \frac{\sigma}{\theta}\omega_{it} + \frac{(\alpha + \beta + \gamma)}{\theta}\xi_{it} \quad (6)$$

$$p_{it}^* = c_P - \frac{1}{\theta}\omega_{it} + \frac{(1 - \alpha - \beta - \gamma)}{\theta}\xi_{it} \quad (7)$$

$$x_{it}^* = c_x + \frac{(\sigma - 1)}{\theta}\omega_{it} + \frac{1}{\theta}\xi_{it} \quad (8)$$

where $\theta \equiv \alpha + \beta + \gamma + \sigma(1 - \alpha - \beta - \gamma)$, $x = k, l, m$ and c_q, c_p, c_x are constants. Under decreasing returns to scale ($\alpha + \beta + \gamma < 1$), output increases with both productivity and demand shocks; whereas price decreases with productivity and increases with market appeal, and inputs demand increases with both. With constant returns to scale, however, p^* depends only on costs parameters and not on demand ones. Note that, although the markup is constant at $\frac{\sigma}{\sigma - 1}$, prices differ across firms. In fact, firms' marginal costs differ for two reasons. Firstly, they are characterized by different efficiency levels ω_{it} , which directly affect marginal costs given output. Secondly, if the production function displays non constant returns to scale, different levels of ω and ξ entail different level of output, and therefore, of marginal costs.

In terms of the capital constraint, from equation (8) it follows that the firm uses its full capacity, that is $U_{it} = 1$, if and only if

$$\bar{k}_{it} \leq c_k + \frac{(\sigma - 1)}{\theta}\omega_{it} + \frac{1}{\theta}\xi_{it} \quad (9)$$

Condition (9) states that the capital stock in place is not a binding constraint as long as the productivity and demand shocks are not too large. In fact, as shown above, output is increasing in both shocks. We analyze the case where the constraint binds in Appendix A.1.³

Measuring TFP using physical output (TFPQ, in the language of Foster et al. (2008)) or output deflated with the sectoral price deflator (TFPR, where \mathbf{R} stands for revenues) leads to identify different objects. In our setting, $\text{TFPR}_{it} = \omega_{it} + p_{it} - \bar{p}_t$, where \bar{p}_t is the

³In our data, only 2% of the observations pertain to firms that report full capacity utilization. Note that hitting the capital constraint does not affect the demand or the production function estimation. In fact, in the demand equation price depends only on output, independently from how it is produced. In the production function, output depends on the input combination, independently of whether the firm is at the corner in terms of capital utilization. Therefore, we can use the entire sample to estimate the demand and production functions. However, in the appendix we show that the relationship between output and input demand and the shocks does depend on whether the capital constraint is binding. As a consequence, we exclude observations where the capital constraint has been met in the second part of the paper, where we look at the elasticity of output and input to shocks.

(log of the) sectoral price deflator. Using equation (7) and suppressing the constant, we can show that not accounting for firm level prices introduces a bias in the estimate of TFP:

$$\text{TFPR}_{it} = (1 - \frac{1}{\theta})\omega_{it} + \frac{(1 - \alpha - \beta - \gamma)}{\theta}\xi_{it} - \tilde{p}_t$$

The bias has two sources. First, true TFP (ω) enters with a coefficient smaller than one, as higher productivity in part translates into lower prices (see equation 7) and therefore lower revenues. The effect is stronger the higher the degree of returns to scale. In fact, with constant returns to scale, θ is equal to 1 and any change in ω is reflected one-to-one on prices while TFPR is completely unaffected. Second, TFPR also depends on demand shocks. Take the case of a firm that receives a positive demand shock and rises price as a consequence. Since the sectoral deflator would not capture this idiosyncratic change, under the conventional approach we would mistakenly conclude that the firm has increased its produced quantity and, therefore, its productivity. This bias is stronger the lower the demand elasticity and the lower the degree of returns to scale. With constant returns to scale the firm level price is unaffected by demand shocks and the effect disappears.

Without knowledge of firm prices the coefficients estimated using revenue data are also inconsistent (Klette and Griliches 1996). In fact, using (1) and (4) and taking logs, it is straightforward to show that a revenue production function can be expressed as:

$$q_{it} + p_{it} = \frac{\sigma - 1}{\sigma}\alpha k_{it} + \frac{\sigma - 1}{\sigma}\beta l_{it} + \frac{\sigma - 1}{\sigma}\gamma m_{it} + \frac{\sigma - 1}{\sigma}\omega_{it} + \frac{1}{\sigma}\xi_{it} \quad (10)$$

Even if we accounted for the endogeneity of inputs, the coefficients of a revenue function underestimate the true degree of returns to scale. As in Melitz (2000), the size of the bias is $\frac{\sigma-1}{\sigma}$. Intuitively, when a firm expands its output, it must decrease the price to move down the demand curve, so that the increase in revenues is less than proportional to the increase of physical output.

The only dynamic choice the firm faces in our setting is investment, through which the firm can increase the stock of capital in place next period. Following Olley and Pakes (1996), we use investment to control for unobserved productivity in the estimation of the production function. In Appendix A.2 we set up the dynamic problem and show that our investment function depends not only on the capital in place \bar{k} and productivity ω , as in the standard case, but also on the firm's market appeal ξ , violating the assumption of scalar unobservability. Akerberg, Benkard, Berry, and Pakes (2007) argue that the Olley and Pakes (1996) procedure can be generalized to this case by including the demand shifter in the control function:

$$\omega_{it} = h(i_{it}, \xi_{it}, \bar{k}_{it}) \quad (11)$$

We assume that market appeal and TFP follow first order Markov processes such that $F_{\Omega}(\cdot|\Omega'_{it-1})$ first order stochastically dominates $F_{\Omega}(\cdot|\Omega_{it-1})$ for $\Omega'_{it-1} > \Omega_{it-1}$ and $F_{\Xi}(\cdot|\Xi'_{it-1})$ first order stochastically dominates $F_{\Xi}(\cdot|\Xi_{it-1})$ for $\Xi'_{it-1} > \Xi_{it-1}$. High TFP (market appeal) today implies high expected TFP (market appeal) tomorrow. This assumption is important for the invertibility of the investment function. Intuitively the policy function is invertible if, given two firms with the same installed capital and demand shock, investment is strictly higher in the firm with the higher productivity shock.

We have so far assumed that firms produce a single product. Multi-product firms pose some additional challenges for the estimation as our data report average price changes, total output and input usage at firm level with no disaggregation for single product lines. In Appendix A.3 we extend the theoretical framework to the case of multi-product firms and show that demand and productivity shocks can be recovered also under this scenario. In particular, if demand and productivity shocks are identical across products, as typically assumed in empirical work (Foster et al. 2008, De Loecker 2011),⁴ the distinction between working with product or firm level data becomes blurred. Our methodology works even if demand shocks are specific to individual products, as long as there is a unique production function for all products at the firm level. The use of aggregate firm level data is, however, problematic when there are product-specific productivity shocks. As far as we know, such a case has not yet been addressed in the empirical literature.

3 Data description

The data used in this study come from the “Indagine sugli investimenti delle imprese manifatturiere” (Inquiry into investments of manufacturing firms; henceforth, INVIND), a survey collected yearly since 1984 by the Bank of Italy. The survey is a panel representative of Italian manufacturing firms (no plant level information is available) with at least 50 employees⁵ and contains detailed information on revenues, ownership, capital and debt structure, as well as on usage of production factors. Additional firm information is drawn from “Centrale dei Bilanci” (Company Accounts Data Service; henceforth, CB), which contains balance sheets data of around 30,000 Italian firms. Firms in INVIND can be matched to their balance

⁴An important exception is De Loecker et al. (2012), who use a unique dataset of Indian firms with information on prices and sales at the product level to estimate marginal costs at the product level. They assume that each product has its own production function, but that there is a unique productivity shock common to all products within the firm.

⁵From 2002 the survey was extended to service firms and the employment threshold lowered to 20. However, these firms are given a shorter questionnaire, which excludes some of the key variables for our analysis. We therefore focus throughout on manufacturing firms with at least 50 employees.

sheet data in CB using their tax identification number.

To ensure homogeneity of the final good produced we group firms into sectors. We use an aggregation of the ATECO 2002 classification of economic activities, giving us seven sectors, listed in Table 1 (we exclude sectors where sample size was too small). We drop observations pre-1988, since prior waves of the survey do not contain information on firm-level prices. We also ignore firms not matched with CB (25% of the INVIND respondents) as well as those not surveyed for at least two consecutive years (22% of the residual sample). After applying these refinements, we are left with a pooled sample of 12,102 firm-years over the period 1988-2007.

The information on firm prices contained in the INVIND survey is key to our goal of disentangling demand from TFP shocks. Foster et al. (2008) attained the same goal using directly firm-level quantities. Their strategy can only be applied to industries producing homogenous output; direct observation of firm level prices, however, allows us to extend the analysis to industries where product differentiation is important.⁶

Firms are asked to state the “*average percentage change in the prices of goods sold*”. This implies that we will only be able to exploit this information estimating the model in first differences.⁷ Using the average price change is problematic in cases of the introduction of new products and the dropping of old ones, for which the price change is not defined. We implicitly assume that the share of products introduced or retired by any firm in a given year is small enough not to affect significantly the average growth rate of price. At the same time, using growth rates also delivers some advantages. For example, for multi-product firms the average growth in prices is a more meaningful object than the average price level. Leaving aside the introduction of new products, using first differences nets out any fixed unobserved heterogeneity that might distort the estimates. This is important because in the model we posited that market appeal does not pick up quality differences embodied in physical attributes of the product. Exploiting only variation within firms ensures that: i)

⁶Since the information in the INVIND survey, and in particular the price data, is self-reported by the interviewee, we perform several checks to validate the variable. First, for a number of variables (e.g. revenues, investments, etc.) appearing both in the INVIND survey and balance sheet data, we find that the figures match well. Therefore, there is no indication that entrepreneurs are more inclined to lie or to provide inaccurate answers in the survey than they are when compiling official documents. Furthermore, the Bank of Italy itself relies on the INVIND pricing information for its official reports. Finally, in Appendix B.1 we compare a price index built upon INVIND prices with that constructed by the national statistical office (ISTAT). The two series are highly correlated.

⁷Firms report % price change $\equiv \frac{P_{it}}{P_{it-1}} - 1$. We use this figure to obtain the first difference in the logarithm of price Δp . We obtain the growth rate of the logged prices using the transformation $\Delta p = \ln(1 + \% \text{ price change})$. All the variables reported in the survey as percentage changes are transformed in the same way.

cross-firm differences in product quality do not contribute to identification; ii) the bulk of the variation in Δp for a given firm relates to a given set of products with fixed physical attributes.

Nominal output is obtained from balance sheets data in CB. We deflate its growth rate using firm level price changes to obtain real output growth. Labor input is measured as hours worked. Intermediate inputs come from CB and are deflated with sectoral prices. To measure capital inputs, we exploit questions in INVIND on both production capacity⁸ and the degree of capacity utilization. Direct assessment of the change in installed productive capacity $\Delta \bar{k}$ ⁹ helps us to circumvent issues of measurement error linked to standard measures of capital based on book values or permanent inventory method. Utilized capital is a better measure of capital services in the production function than installed capital, which implicitly assumes that the degree of capital utilization is 100%, something clearly not borne out in the data. The average degree of capacity utilization of 81%, with a standard deviation of 13%; the 5th and the 95th percentile are 60% and 98% respectively. Moreover, utilized capital displays additional variation that is useful for identification. Estimating the production function in first differences, we rely exclusively on the within-firm variation in the capital input. This poses a challenge for the estimation of its coefficient, as capital in place tends to have limited within-firm variability. Utilized capital displays greater within-firm variation than capital in place.

Table 1 consists of descriptive statistics for our key variables both in levels (Panel A) and in growth rates (Panel B). Textile and leather and Mechanical machinery are the industries most represented, reflecting the reality of Italian sectoral specialization. There is substantial cross-industry variation in sales, which stretch from an average of around 60m euros in Textile and leather up to almost 500m euros in Vehicles. Variation in the average number of employees is more limited, ranging between 300 and 600 workers, with Vehicles being the outlier at almost 2,000 workers.

A first look at growth rates shows that real sales and output grew on average 2% per year over the sample period, with a standard deviation of 6%. Labor input contracted slightly, whereas capital input grew at 4% yearly. The average firm in the sample raises

⁸This variable corresponds to \bar{K}_{it} in the notation of Section 2 and is defined as “*the maximum output that can be obtained using the plants at full capacity, without changing the organization of the work shifts*”.

⁹Note that the question asks about the change in the maximum output obtained using the plants at full capacity, “*without changing the organization of the work shifts*”. This excludes the possibility that the measure of capital so obtained already incorporates changes in productivity. Any TFP gain should in fact entail a certain degree of work reorganization. We have also experimented with traditional measures of the capital stock, constructed with the permanent inventory method using sectoral deflators and depreciations rates. Results are robust.

prices by 2% per year. Average price growth shows little cross-sectoral dispersion ranging between 1.6% (Paper) and 2.7% (Metal). Figure 1 shows the distribution of price changes for each of the seven sectors in one year of our data. The picture confirms that there is substantial dispersion around the sectoral average and reaffirms the importance of having information on firm level price adjustments.

4 Demand and TFP Estimation

4.1 Demand estimation

Firms face a CES demand function of the form expressed in equation (1). We estimate demand separately for each of the ATECO sectors in our sample:

$$\Delta q_{it} = \sigma \Delta p_{it} + \Delta \xi_{it} \quad (12)$$

The shock to *market appeal*, $\Delta \xi_{it}$, is known to the firm but unobserved to the econometrician. If we obtained consistent estimates of the price elasticity (σ), we could estimate $\Delta \xi_{it}$ as follows:

$$\widehat{\Delta \xi_{it}} = \Delta q_{it} - \hat{\sigma} \Delta p_{it} \quad (13)$$

Estimation of equation (12) is complicated by the familiar simultaneity problem. Positive shocks to market appeal lead producers to raise prices, as shown in equation (7), making Δp and $\Delta \xi$ positively correlated. Therefore, estimating the equation by OLS would understate demand elasticity. In our context, finding valid instruments for price constitutes a challenge. To solve this problem, we exploit a unique piece of information included in our data. In 1996, and again in 2007, the interviewed managers were directly asked to report the elasticity of the demand faced by their firm through the following question:

“Consider the following thought experiment: if your firm increased prices by 10% today, what would be the percentage variation in its nominal sales, provided that competitors did not adjust their pricing and all other things being equal?”.

Since managers are explicitly asked to perform a thought exercise isolating the effect of price changes on demand, the estimates we derive from their answers should not be affected by simultaneity. Therefore, we choose to rely upon answers to this question to estimate a sector-specific demand elasticity as the average of the elasticities reported by firms belonging to a given sector. Figure 2 reports kernel densities by sector for the distribution of self-reported elasticities in the two waves. They look similar, and a Kolmogorov-Smirnov test

does not reject the equality of the distribution in the two waves for five of our seven sectors. We use elasticities reported by the cross-section of representative firms interviewed in 1996 to estimate σ since this wave falls in the middle of our sample period and the response rate is higher than in 2007 (over 80%).

Table 2 presents estimated demand elasticities for each of our seven sectors. In the first column, we list average sectoral self-reported demand elasticities from the 1996 wave of INVIND. Textile and leather and Chemical products are the least elastic with a σ of 4.5 and 4.7, respectively. Firms in the Vehicles sector face the most elastic demand ($\sigma=6$). These values are in the range of those found of the literature. For instance, the average elasticity for Vehicles is close to the price elasticity found by Berry, Levinsohn and Pakes (1995) for compact cars, which make up most of the market of Italian car producers (mostly Fiat and its suppliers). Our estimate for textile is in the range found by De Loecker (2011), who looks at several segments within the textile sector in Belgium. Hsieh and Klenow (2009) in their calibration exercise use what they refer to as a conservative value of 3 and check the robustness of their results with an alternative value of 5. In appendix B.2, we comment on additional results reported in Table 2, where we check whether the presence of multiproduct, multiplant and exporting firms affects our estimates of price elasticity. We also compare the results based on self-reported information with a more traditional approach to estimating demand using OLS and IV techniques. The findings are in line with those in Column 1.

4.2 TFP estimation

We directly estimate a *quantity* production function (as opposed to revenues) in first differences as in the equation below:

$$\Delta q_{it} = \alpha \Delta k_{it} + \beta \Delta l_{it} + \gamma \Delta m_{it} + \Delta \omega_{it} + \epsilon_{it} \quad (14)$$

where ϵ_{it} is an iid random shock unobserved to the firm when choosing inputs, or measurement error. We compute the growth rate of real output by subtracting the price change from the nominal output.¹⁰

¹⁰Using output instead of value added delivers several advantages. First, the use of value added implicitly imposes strong assumptions on the degree of substitutability between intermediates and other inputs. Second, Gandhi, Navarro and Rivers (2011) have shown that estimating TFP using value added can lead to overstate the productivity dispersion. Third, and most important, we want to ensure comparability between shocks to market appeal and to TFP. Shocks to market appeal are computed using sales. Estimating TFP from output therefore ensures that shocks are computed from comparable quantities, as sales and output only differ due to inventories, while value added also subtract intermediates.

To account for the endogeneity of inputs, we follow the control function approach introduced by Olley and Pakes (1996). Our setting differs from the standard one along three dimensions: first, we estimate the production function in first differences rather than in levels. Second, firms are characterized by two unobservables, ω and ξ , rather than just one, so that we need to relax the scalar unobservability assumption. Finally, we measure the capital input as utilized capital, so that the capital actually used in production can be adjusted after observing the shocks, rather than being pre-determined. These modifications to the basic setting raise technical issues that we discuss in detail in Appendix B.3. We use as control function a third degree polynomial in $\Delta\xi_{it}, \Delta i_{it}, \Delta\bar{k}_{it}$ and estimate the production function using the following regression equation (year dummies are included in each specification):

$$\Delta q_{it} = \alpha\Delta k_{it} + \beta\Delta l_{it} + \gamma\Delta m_{it} + h(\Delta\xi_{it}, \Delta i_{it}, \Delta\bar{k}_{it}) + \epsilon_{it} \quad (15)$$

Once we have estimated the coefficients, we recover the changes in TFP (up to the random component ϵ_{it}) as

$$\widehat{\Delta TFP}_{it} = \Delta q_{it} - \hat{\alpha}\Delta k_{it} - \hat{\beta}\Delta l_{it} - \hat{\gamma}\Delta m_{it}$$

Table 3 reports sector-by-sector estimates of the coefficients of the production function. To reduce the effects of extreme values on the estimates, we exclude the observations the top and bottom centiles of the distribution of $\Delta q_{it}, \Delta k_{it}, \Delta l_{it}$ and Δm_{it} . Panel A shows the baseline results. We find evidence of decreasing returns to scale for all sectors: the degree of returns to scale $\alpha + \beta + \gamma$, reported in the last row of the panel, ranges between .74 in Minerals and .92 in Chemicals.¹¹

In Panel B we run the estimation procedure using output deflated with sectoral prices rather than with firm level prices. For all sectors, the real output based estimates are larger than the revenue based estimates, as predicted by Klette and Griliches (1996). As shown in equation (10), the relation between the true parameters of the production function and the estimates derived using revenue based measures is as follows: $\alpha + \beta + \gamma = \frac{\sigma}{\sigma-1}(\tilde{\alpha} + \tilde{\beta} + \tilde{\gamma})$, where $\tilde{\alpha}, \tilde{\beta}, \tilde{\gamma}$ are the estimates that do not correct for the own price deflator. In the last

¹¹These figures are lower than those typically estimated with level production functions. For example, Levinsohn and Petrin (2003) report returns to scale close to 1. Compared to their estimates, we find a lower elasticity of the capital coefficient: their sectoral estimates vary between .19 and .29, while ours are between .1 and .2. A low elasticity of output to capital is typically found in fixed effects estimations, which are known to give low and imprecise estimates of the capital coefficient. Olley and Pakes (1996) attribute this to the fact that the capital stock has little within firm variability. Such a critique is less likely to apply in our setting. In fact, with capital utilization we can compute utilized capital, which is more variable than the capital stock. Doraszelski and Jaumandreu (2012) find evidence of decreasing returns to scale for similar sectors at the same level of aggregation as ours.

row of the table we compute the implied returns to scale, using the sectoral self-reported elasticities reported in Table 2. Applying the correction to the estimates based on sectoral deflators brings them close to those obtained using output deflated with firm level prices, although for four sectors the implied coefficients become larger than those in Panel A.

We have also considered the possibility that first differencing introduces downward bias due to measurement error in the independent variable. Although we cannot run the regressions in levels, we can increase the length of the lag on which the production function is estimated. As the lag increases the extent of the measurement error should decline, as we consider lower frequency movements in inputs and output. We therefore estimate the production function using differences over a three year period and compare results with those obtained in the baseline setup where we used first differences, excluding the Paper and Vehicles sector, for which the number of observation becomes a concern when using longer lags. The results are reported in Appendix Table A-1; we find that the degree of returns to scale slightly increases but the resulting TFP estimates are similar to the basic ones, with a correlation of .97.

4.3 Descriptive statistics on Δ TFP and $\Delta\xi$

Panel A of Table 4 shows descriptive statistics for Δ TFP. The average growth rate is below 1%, consistently with the well documented low productivity growth that has characterized the Italian economy since the early nineties (Brandolini and Cipollone 2001). There is also substantial dispersion in TFP growth (the standard deviation is .14). The second row reports the distribution of TFP computed using the estimates of the production function based on three year differences. The two distributions are virtually identical.

Panel B reports analogous information for $\Delta\xi$. The row labeled “ $\Delta\xi$ sector” reports estimates for $\Delta\xi$ based on the self-reported elasticities contained in the 1996 wave of the INVIND survey, averaged at the sectoral level. We also report estimates of $\Delta\xi$ based on elasticities averaged at the ATECO class level, a much finer definition of the area of activity.¹² The estimates labeled “ $\Delta\xi$ individual” are obtained using the individually reported estimates, rather than sectoral averages. In that case, we can only use the firms that directly answered the question in 1996. There are some differences in the mean of the distribution of the $\Delta\xi$ estimated using different level of aggregation. However, these discrepancies are entirely due to outliers. If we compare the quintiles of the distributions, the estimates are

¹²As an example, production of iron and non iron metals belong to different classes of activities within the sector Metals. Similarly, the classes within the Chemicals sector distinguish between firms producing paint and those producing soap and detergents.

nearly identical.¹³

The correlation between $\Delta\xi$ and ΔTFP is nearly zero, validating the assumption of independence made by Foster et al. (2008) when using TFP as an instrument for ξ in the pricing equation. Estimating the degree of serial correlation under the assumption that both ξ and TFP are AR(1) processes is more complicated, as first differencing invalidates OLS regressions of each variable on its lag. In Appendix B.5 we discuss our preferred IV specification, based on further lags of the shocks as instruments. We find a degree of serial correlation of .76 for TFP and of .25 for ξ . The first number is consistent with that of Foster et al. (2008), while the latter is substantially lower. Both should be taken with caution, as they are sensitive to the choice of instruments and to the treatment of outliers.

5 Shocks and firm growth

In this section we quantify the importance of changes in productivity and market appeal in driving firm growth. Furthermore, we show that the results imply the existence of major adjustment costs and that such costs are different for TFP and market appeal shocks.

5.1 Measurement

Under the assumption that the process generating shocks to TFP and market appeal is exogenous, we can assess the elasticity of growth measures to these shocks by estimating regressions of the following form

$$\Delta y_{it} = a_0 + a_1 \Delta TFP_{it} + a_2 \Delta \xi_{it} + a_3 X_{it} + e_{it} \quad (16)$$

where Δy_{it} is the growth rate of some variable of interest (sales and output, prices, inputs), ΔTFP and $\Delta \xi$ are the estimated idiosyncratic shocks and X_{it} are additional controls.

Table 5 reports the results of a set of such regressions for output and price. For reasons of brevity, we only report pooled cross sectoral estimates. We account flexibly for cross sectoral heterogeneity through a full set of time-sector dummies and also include location dummies for five macro-regions of Italy. Sectoral estimates, reported in the appendix, are fully in line with the pooled ones.¹⁴ We account for the fact that ΔTFP and $\Delta \xi$ are generated regressors by bootstrapping the standard errors.

¹³We have also looked at the distribution of $\Delta \xi$ implied by using only single product firms or only firms that do not export, and found no substantial differences.

¹⁴We have also performed firm fixed effects regressions to control for unobserved heterogeneity even within sector, finding no significant variation in the results. We take this to be an indication that, since our analysis involves first differences, we are already purging unobserved firm heterogeneity that might affect both shocks and sales.

Column (1) shows that shocks to demand and to TFP have a positive impact on *nominal* sales growth.¹⁵ The elasticity of sales to TFP shocks is larger than that to market appeal (0.6 vs 0.41). Once we factor in dispersion of the shocks, however, we find that one standard deviation change in ΔTFP would increase sales by 8%, whereas a similar change in $\Delta\xi$ would have an impact of 13%. Demand shocks, therefore, are more important than productivity shocks in determining the evolution of market shares. Once we move from nominal to real sales (Column 2), however, the elasticity to TFP grows and that to demand shrinks: a standard deviation increase in TFP increases real sales by 10% against 8% for market appeal. This is not surprising since we are removing the price effect that tends to inflate the response of revenue after demand shocks and reduce that following productivity gains. In fact, a firm experiencing an increase in $\Delta\xi$ should not only increase the quantity sold but also the price, while the opposite should occur for TFP. Accordingly, in Column (3) we show that positive shocks to TFP lead to price cuts and improvements in demand appeal trigger price raises. The positive effect of demand shocks on prices is also consistent with our findings of decreasing returns to scale. In fact, in a constant returns to scale scenario, variations in market appeal should not affect the price.

A lingering concern may be that we have used sales as a proxy for output. Whereas this is the measure we want to consider when thinking about demand and, therefore, the market appeal component, it could affect measurement when we turn to productivity. In fact, quantity sold and quantity produced do not have to coincide, due to inventories. Since we have information on quantity produced, in the last two columns of Table 5 we repeat the exercise using it as the dependent variable and check whether results are robust. We find that the elasticity of TFP shocks increases by about 0.2 relative to the sales regressions, both in the nominal and in the real output regressions, while the coefficients of demand shocks decrease slightly.

The figures presented above refer to the overall effect of productivity and market appeal on output. For TFP this is the sum of a direct and of an indirect effect. Positive changes in TFP directly increase the quantity produced or sold but also should have an indirect impact as they affect demand for factors of production: l, k, m . For ξ , however, the effect comes entirely through the indirect channel, as demand shocks have no direct contribution to the quantity produced. Total differentiation of equation (14) delivers a decomposition of

¹⁵Note that with sector-year dummies there is no difference between using nominal or real values obtained through sectoral price deflators.

the overall effect of the two shocks on output:

$$\frac{d\Delta q_{it}}{d\Delta\omega_{it}} = \underbrace{1}_{\text{direct effect}} + \underbrace{\alpha \frac{\partial\Delta k_{it}}{\partial\Delta\omega_{it}} + \beta \frac{\partial\Delta l_{it}}{\partial\Delta\omega_{it}} + \gamma \frac{\partial\Delta m_{it}}{\partial\Delta\omega_{it}}}_{\text{indirect effect}} \quad (17)$$

$$\frac{d\Delta q_{it}}{d\Delta\xi_{it}} = \underbrace{\alpha \frac{\partial\Delta k_{it}}{\partial\Delta\xi_{it}} + \beta \frac{\partial\Delta l_{it}}{\partial\Delta\xi_{it}} + \gamma \frac{\partial\Delta m_{it}}{\partial\Delta\xi_{it}}}_{\text{indirect effect}} \quad (18)$$

where, according to equation (8), $\frac{\partial\Delta x_{it}}{\partial\Delta\omega_{it}} = \frac{\sigma-1}{\theta}$ and $\frac{\partial\Delta x_{it}}{\partial\Delta\xi_{it}} = \frac{1}{\theta}$, for $x = \{k, l, m\}$. Given that we found a unit elasticity of output to TFP shocks in Table 5, we expect that the indirect effect of inputs demand is zero. Indeed, this is what we find. Panel A of Table 6 reports the growth of inputs on demand and TFP shocks. With the exception of intermediate goods, production inputs are not responsive to TFP shocks: productivity changes do not set in motion the changes in inputs that would compound their effect on output. All inputs, however, react to a demand shock, with intermediates again showing the higher elasticity.

We assumed that variable inputs (hours worked, utilized capital, intermediates) can be adjusted in the short run, thus representing an intensive margin. However, there is a limit to the number of hours that can be squeezed out of a fixed number of workers and a firm cannot use more than 100 percent of its installed capital. If firms experiencing improvements in productivity or market appeal want to increase their scale they must act on what we label *quasi-fixed* input of production: the number of workers and installed capital. The reallocation of inputs from less to more productive units (and, in our setting, from low to high market appeal units) is a major source of productivity growth in market economies. For example, Olley and Pakes (1996) attribute to capital reallocation most of the productivity growth that occurred after the deregulation of the US telecom sector. Moreover, a growing literature focuses on the obstacles to the efficient allocation of resources across production units as a major impediment to growth (Restuccia and Rogerson 2008, Hsieh and Klenow 2009). To study the effects of shocks on reallocation, we now consider the extensive margin of inputs adjustment. Panel B of Table 6 displays the correlation of growth in the number of workers and the investment rate with changes in TFP and market appeal. One standard deviation increase in Δ TFP and in $\Delta\xi$ produces a similar impact on investment rate of roughly 1.2 percentage points, that is 16 percent of the median investment rate in our sample (7.3%).¹⁶ Figures for the elasticity of growth in the (end of the year) labor force are similar. Breaking down the employment growth rate into its determinants, hiring and

¹⁶This is also an indirect test of the fact that investment increases with shocks, as assumed in the control function approach.

separation rates, provides additional insights (Columns (2) and (3), Table 6 - Panel B). Most of the action takes place on the hiring margin, whereas the elasticity of separations to demand appeal is low and that to TFP is not significant. This can be interpreted as evidence of adjustment costs on the firing margin, consistent with the fact that in the Italian labor market employment protection legislation imposes firing costs on firms (Schivardi and Torrini 2008).

5.2 Evidence of adjustment costs

We have measured the relative importance of productivity and market appeal changes in driving firm growth. However, in the absence of a benchmark it is hard to assess the size of these effects. Given estimates of the demand elasticity and of the production function coefficients, the theoretical framework we setup in Section 2 delivers quantitative predictions on the impact of demand and supply shocks on inputs and output growth. The relationship between output, price, and inputs growth and shocks to TFP and demand appeal can be obtained by first differencing equations (6), (7) and (8), respectively. To compute the model predictions, we assume returns to scale of .8 and an elasticity of demand of 5 (roughly the cross-sectoral averages from our estimates). The elasticities implied by the model are reported in Tables 7. For convenience, we also report there the corresponding estimates from Table 5 and 6. A one percent increase in TFP should bring about a 2.2 percent increase in nominal output and a 2.8 percent increase in real output, whereas price should go down by .56 percent. The predicted effects are smaller for demand shocks: the elasticity of nominal output is .56, that of real output .16. It is immediately evident that predicted elasticities are larger than those we measure empirically. For instance, our recovered elasticity of nominal sales to demand shocks is .41, whereas the model would predict an elasticity of .56. For TFP, the estimate of nominal output elasticity is .81 against a predicted figure of 2.2. These large differences survive in all the specifications and for all the outcome variables considered. We state this finding explicitly.

Finding 1. *The responsiveness of growth measures to productivity and market appeal changes is substantially lower than that predicted by a frictionless model.*

This finding can be rationalized by the presence of frictions limiting the effect of demand and TFP shocks. In that case, the gap between predicted and estimated elasticities could be simply due to the fact that our theoretical framework does not account for them and, therefore, predicts “full” response of employment and investment to changes in productivity and market appeal. In hypothesizing the existence of such frictions we join a vast literature

that has discussed adjustment costs in factor demand and their role in preventing efficient allocation of inputs.¹⁷

Following a large literature we have estimated the production function abstracting from adjustment costs, with the exception of the one period lag in building capital. One may think that ignoring adjustment costs when estimating the production function makes our estimates inconsistent. We argue that this is not the case. To begin with, it must be stressed that only costs that modify the amount of output obtained for given inputs are problematic for our approach. Many instances of adjustment costs, such as firing costs on labor or bureaucratic and administrative costs to modify the scale of operation, do distort input choices away from the unconstrained optimum but not the inputs-output relationship: these types of adjustment costs do not affect our estimates. For instance, a firm might keep workers whose marginal product is below the marginal costs if firing is costly. Still, conditional on employing such workers, the firing cost does not modify the technical relationship between inputs and output, the only determinant of the estimates of the production function. Frictions that enter the production function pose more serious challenges to our estimation procedure. Cooper and Haltiwanger (2006) estimate a general capital adjustment cost function, allowing for the possibility that new investments disrupt production. For example, adjusting the capital stock may require the firm to temporarily shut down operations to install the new machinery or to retrain workers to use the technology. Ignoring such costs would lead to biased estimates, as we observe lower output when the firm is increasing its capital stock. There are, however, several reasons to believe that our production function estimates are robust to the presence of disruption costs. First, our measures of inputs (utilized capital and worked hours) account for this type of disruption, unlike those typically used in the literature. In fact, whereas the installed capital and the number of workers do not fluctuate as a result of a temporary plant shutdown, this will result in a lower utilization of installed capital and in fewer hours worked. Therefore, even without explicitly introducing frictions in the production function, our input measures protect us from the bias they may introduce. Second, as we argue in Appendix B.6, any fixed costs of changing the capital stock drop out when first differencing, as the control function approach forces us to only use the observations in which firms are investing. Finally, we note that at lower frequencies the size of variation in production inputs should be large

¹⁷Recent contributions include studies of investment adjustment costs (Collard-Wexler et al. 2011), financial constraints (Banerjee and Moll 2012, Midrigan and Xu 2010), and employment protection legislation (Petrin and Sivadasan 2011). Hamermesh and Pfann (1996) provides a comprehensive survey of the earlier literature.

enough to swamp the disruption cost. The change in output over a few years' arc will result from the cumulated investments over those years; whereas only current disruption costs will be reflected in current output. As discussed in Section 4.2, our TFP estimates change only marginally when estimating the production function using longer lags.

The comparison of our estimates of the elasticities to demand and productivity shocks with the model predictions delivers a second interesting result. Not only are measured elasticities much lower than those that the model predicts, but the gap is significantly larger for the response to productivity than to market appeal. For example, the predicted elasticity of real output to TFP changes is 2.8, whereas we estimate it to be .98, while our estimate of the elasticity to market appeal is .27, much closer to the predicted value of 0.44. The response of price to shocks is very similar for demand (0.11 vs. 0.13), while much smaller in absolute value for TFP shocks (-0.56 vs. -0.15).

Finding 2. *Deviations between actual and predicted responses are much larger for TFP than for market appeal shocks.*

Our second finding is, to the best of our knowledge, completely novel. It implies that adjustment costs affect asymmetrically firms responses to demand and TFP. In particular, frictions are stronger when adjusting to changes in productivity. Detecting that frictions are not independent of the nature of the shock necessarily requires a model allowing for multiple forcing variables and our paper is one of the first taking this approach.¹⁸

To provide evidence that the wedge between the model's prediction and our estimates is indeed caused by adjustment frictions, we consider a general implication of adjustment costs: they should induce lagged response to changes in TFP and market appeal.¹⁹ If adjusting prices or inputs takes time, we should find that current output growth is a function not only of contemporaneous but also of lagged shocks. Note that, even in the presence of adjustment costs, the production and demand functions are static: output produced depends on current inputs, and quantity sold depends on price. Introducing dynamic effects does not affect our general identification strategy but it does complicate the control function approach. Since

¹⁸The low response of investment to shocks could be explained without introducing adjustment costs if the persistence of the processes were low. However in Section 3 we showed that, if anything, TFP shocks seem to be more persistent than demand shocks. If the lack of response were driven by low persistence, we should have found the opposite of that stated in Finding 2.

¹⁹Although it is generally true that the presence of adjustment costs induces lagged response, the dynamic pattern of adjustment itself depends on the form of the adjustment cost function. For example, convex adjustment costs imply smooth adjustment, while fixed costs lead to bunching. In both cases, adjustment depends not only on current but also on past shocks. The literature has been inconclusive on the shape of the adjustment cost function. Contributing to this debate is beyond the scope of this paper.

the firm chooses investment based also on the lagged values of the shocks, we need to increase the number of controls. We use the forecast for next year’s investment, the expected change in technical capacity and two lags of the demand appeal shocks as additional controls.²⁰

We investigate the importance of lagged shocks in Table 8. We consider two lags of both TFP and demand appeal shocks. Past TFP shocks have a sizeable effect on the growth rate of output (Column 1): .15 at lag 1 and .036 at lag 2. Slow adjustments implies that real output should keep growing after impact and that the price should keep falling. Column (2) shows that the pattern for price is consistent with this prediction. A positive shock to TFP leads to price cuts in the current year (-.16), as well as in the next two (-.04 and -.02, respectively). The lagged effects are even stronger on “quasi-fixed” inputs, that is in the (end of the year) number of employees and in the investment rate (columns 3 and 4). Even at lag two the elasticity is similar to the contemporaneous one. This indicates that the time required to update the productive capacity to a TFP shock is substantial.

The dynamic of demand shocks follows a rather different pattern. Lagged demand shocks have a small negative effect on output at lag 1 and no effect at lag 2. The negative effect at lag one might seem counterintuitive, but can be understood by considering the fact that the price still continues to increase in response to higher market appeal one period after the shock occurred. This pattern is consistent with a sluggish price adjustment: after a positive demand shock, firms do not immediately increase prices to the new equilibrium level. As a result, the immediate increase in output is larger than the “long run” one. As prices are further increased, output falls. Quasi-fixed inputs follow the same pattern observed for TFP shocks, with positive responses at all lags. This is not at odds with the output results. In fact, in unreported regressions we found that variable inputs display a negative elasticity at lag one, as does output. Still, the firm upgrades the productive capacity slowly, consistently with adjustment costs. In terms of asymmetry, lagged effects are stronger for TFP shocks, confirming that adjustment costs are more important for them. Even if we take into account dynamic effects, the cumulated response is far from that predicted by the frictionless model, and the larger deviation for TFP shocks still persist.

²⁰In Appendix B.7 we argue that this is a valid control function for the case at hand. We recompute the coefficients of the production function and the corresponding TFP levels for this modified setting and use these estimates for the regressions in Table 8. The resulting coefficients are similar to those obtained in the basic specification.

6 Is adjusting to TFP shocks more difficult?

In our model, optimal factor demand is derived assuming that firms equalize the marginal value product of each input to its marginal cost, efficiently allocating resources across production units. This condition is, however, not borne out in the data: there are large differences between the responses to shocks we measure and those predicted on the basis of the model. Therefore, our results suggest that factors are not efficiently allocated. This is in line with a growing empirical literature that measures factor misallocation through the dispersion in the marginal value product of inputs (see for example Hsieh and Klenow (2009) for China and India or Yang (2011) for Indonesia).

We also document a new fact about misallocation, and shed light on the type of frictions that may be responsible for it. In fact, we show that the extent of misallocation, measured as the magnitude of the deviation from the elasticities predicted by the baseline model, depends on the nature of the shock: it is larger for TFP than for demand shocks. The frictions commonly considered by the literature do not display this property. For example, the need to pay bribes (Hsieh and Klenow 2009) or the presence of firing costs (Hopenhayn and Rogerson 1993) are often cited as obstacles for firms' growth. However, they would have the same impact whether a firm wanted to grow because it became more productive or because of an increase in its market appeal.

Drawing from the expanding literature on firm organization and managerial practices,²¹ we propose a friction that could cause asymmetric effects of the type we described above. We proceed from the premise that responding to a TFP shock requires more reorganization and restructuring activity than reacting to a market appeal one. When demand increases, the firm is enjoying more recognition by customers and needs to cater to a larger residual demand. This can be done by simply scaling up operations, moving along a given cost function. A TFP shock on the other hand entails a shift in the production technology itself. Although our model does not point to any specific source for TFP growth, some classic examples of productivity improvement have the distinctive feature of being transformative events that require substantial reorganization of work routines within the firm. For example, access to broadband connection has a direct impact on productivity, as workers can now access the web at higher speed (the measured TFP shock). At the same time, in order to fully exploit the opportunities offered by such a shock the firm might require some reorganization of business operations, a different skill mix, different types of capital inputs

²¹See Bloom and Van Reenen (2010) for a recent survey.

etc. These are more complex tasks than proportionally changing all inputs.²² The need for (costly) reorganization can therefore generate a low response to TFP shocks.

We evaluate the viability of this explanation by assessing whether those characteristics of firms which that can be meaningfully related to their capacity to restructure after shocks contribute to explain the under-response to TFP changes. We construct two proxies for firms' ability to restructure and check whether firms that score low in this metric are also less responsive to shocks, particularly to TFP ones.

Our first proxy exploits the information in the INVIND survey which provides a direct measure of reorganization hurdles internal to the firm. Each year the the interviewees are asked to compare actual investments with planned ones²³ and, where the two differ by more than 5%, to identify the causes of the discrepancy. One of the causes listed is "reasons related to the internal organization of the firm". Around 55% of firms not fulfilling their investment plans list problems with internal reorganization among the causes. We assume that firms selecting this option are facing higher costs of organizational restructuring.

The second measure is based on the identity of the controlling shareholder. A recent literature has documented a large degree of heterogeneity in firm's managerial practices and organizational structures (Bloom and Van Reenen 2007, Bloom, Sadun and Van Reenen 2012) and linked their quality to the ownership structure. In particular, family or government controlled firms tend to be managed less efficiently than widely-held or institutionally controlled firms because they are more inclined, when selecting their management, to favoring personal acquaintanceship such as family membership. (Lippi and Schivardi 2010). We use direct information on ownership structure and proxy management quality, with an indicator equal to one if the firm is controlled by a family or the government (around half of of the sample), and zero if controlled by a financial institution, a conglomerate or a foreign entity.

To test our conjecture, we allow the sensitivity to shocks to differ according to the firm categorization by running the following regression:

$$\Delta y_{it} = b_0 + b_1 \Delta TFP_{it} + b_2 D_R \Delta TFP_{it} + b_3 \Delta \xi_{it} + b_4 D_R \Delta \xi_{it} + b_5 D_R + b_6 X_{it} + e_{it} \quad (19)$$

where D_R is the dummy for high restructuring costs. In Table 9 we report the results when classifying firms according to the self-reported measure of organizational hurdles. We find that, following a TFP shock, firms reporting internal organization problems adjust output,

²²In fact, the literature on ICT adoption (Caroli and Van Reenen 2001, Bresnahan, Brynjolfsson and Hitt 2002) has shown that ICT affects firm performance only if the firm also reorganizes.

²³The survey asks firms each year to report planned investments. Therefore, information on unfulfilled plans is derived from an objective forecast on record from the past year. Of course, this implies that the question can only be answered by managers in firms that appear in the survey in consecutive years.

inputs and prices less than the other firms. All differences are statistically significant, with the exception of investments. There is no difference, however, in how firms with and without internal problems respond to demand shocks. Moreover, the indicator variable for organizational bottlenecks in itself is not significant for any of the growth measures we analyzed except output. This indicates that facing organizational problems does not affect growth in itself, directly, but only through its interplay with responses to TFP shocks. Table 10 repeats the exercise, distinguishing between family/government firms and other types of ownership. Again, we find that firms controlled by a family or the government adjust output, prices and employment less than other firms. On the other hand, ownership structure is not a factor in the responses to demand shocks. We obtain similar results if we use a finer partition, creating a dummy for each of the five types of ownership structure present in the data. Overall, this evidence is consistent with the idea that TFP shocks require some degree of restructuring to fully take advantage of them, while accommodating demand shocks does not entail specific reorganization activity and that there is substantial cross-firm heterogeneity in firms' restructuring capacity.

7 Conclusions

In this paper we use a unique dataset on a panel of Italian manufacturing firms which contains information on firm level prices. We identify and distinguish the roles of idiosyncratic productivity and demand in firm growth, and go on to show that, though mostly neglected in the literature, heterogeneity in demand is an important determinant of growth. We also show that the measured effect of idiosyncratic variables is lower than that which would be predicted by a frictionless theoretical framework. Finally we show that the size of this gap is different for technology and demand shocks, suggesting a role for frictions that can generate such asymmetry. We propose an example of this type of friction and present supportive evidence for it.

The main conclusion of this study is that the barriers to the efficient allocation of resources are not exclusively of a regulatory nature, as the literature on this subject has typically assumed (Hopenhayn and Rogerson 1993, Restuccia and Rogerson 2008, Hsieh and Klenow 2009). This conclusion has important implications for the debate on how to reduce productivity losses from misallocation. On the one hand, regulation and corruption call directly government policies into question. On the other hand the type of firm idiosyncratic friction we stress –managerial practices and capabilities and the propensity to restructure– are much less under the direct policy influence. Recent literature has shown

them to depend on a plurality of factors, such as corporate governance and control, managerial and entrepreneurial abilities, work attitudes, competition in the product markets, etc. (Bloom and Van Reenen 2010). Improving our understanding of the determinants of managerial practices and capabilities is of paramount importance.

One limitation of our approach is that demand and supply shocks are considered exogenous. In reality, firms can affect both, for example by investing in R&D and advertising. Endogenizing productivity and market appeal is an exciting avenue to explore in future research.

References

- Akerberg, D. A., Caves, K. and Frazer, G. (2006), Structural identification of production functions. Mimeo, UCLA.
- Akerberg, D., Benkard, L., Berry, S., and Pakes, A. (2007), Econometric tools for analyzing market outcomes, *in* J. Heckman and E. Leamer, eds, ‘Handbook of Econometrics, Volume 6’, Amsterdam: North Holland, pp. 4171–4276.
- Banerjee, A. and Moll, B. (2012), ‘Why does misallocation persist?’, *American Economic Journal: Macroeconomics* **2**, 189–206.
- Bartelsman, E., Haltiwanger, J. and Scarpetta, S. (Forthcoming), ‘Cross-country differences in productivity: The role of allocation and selection’, *American Economic Review* .
- Berry, S. (1994), ‘Estimating discrete-choice models of product differentiation’, *RAND Journal of Economics* **25**, 242–262.
- Berry, S., Levinsohn, J. and Pakes, A. (1995), ‘Automobile prices in market equilibrium’, *Econometrica* **63**, 841–890.
- Bloom, N. (2009), ‘The impact of uncertainty shocks’, *Econometrica* **77**, 623–685.
- Bloom, N., Sadun, R. and Van Reenen, J. (2012), ‘Americans do I.T. better: US multinationals and the productivity miracle’, *American Economic Review* **102**, 167–201.
- Bloom, N. and Van Reenen, J. (2007), ‘Measuring and explaining management practices across firms and countries’, *The Quarterly Journal of Economics* **122**, 1351–1408.
- Bloom, N. and Van Reenen, J. (2010), Human resource management and productivity, *in* O. Ashenfelter and D. Card, eds, ‘Handbook of Labor Economics, Volume 4’, Elsevier.
- Brandolini, A. and Cipollone, P. (2001), Multifactor productivity and labour quality in Italy, 1981-2000. Temi di discussione (Economic working papers), Bank of Italy.
- Bresnahan, T., Brynjolfsson, E. and Hitt, L. (2002), ‘Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence’, *Quarterly Journal of Economics* **117**, 339–376.
- Caroli, E. and Van Reenen, J. (2001), ‘Skill-biased organizational change? Evidence from a panel of British and French establishments’, *Quarterly Journal of Economics* **116**, 1449–1492.
- Collard-Wexler, A., De Loecker, J. and Asker, J. (2011), Productivity volatility and the misallocation of resources in developing economies. Mimeo, NYU.
- Cooper, R. W. and Haltiwanger, J. (2006), ‘On the nature of capital adjustment costs’, *The Review of Economic Studies* **73**, 611–633.

- De Loecker, J. (2011), ‘Product differentiation, multi-product firms and estimating the impact of trade liberalization on productivity’, *Econometrica* **79**, 1407–1451.
- De Loecker, J., Goldberg, P., Khandelwal, A. and Pavcnik, N. (2012), Prices, markups and trade reform. Mimeo, Princeton University.
- Doraszelski, U. and Jaumandreu, J. (2012), Measuring the bias of technical change. Mimeo, The Wharton School.
- Dunne, T., Roberts, M. and Samuelson, L. (1988), ‘Patterns of entry and exit in US manufacturing’, *RAND Journal of Economics* **19**, 495–515.
- Dunne, T., Roberts, M. and Samuelson, L. (1989), ‘The growth and failure of US manufacturing plants’, *Quarterly Journal of Economics* **104**, 671–698.
- Evans, D. (1987a), ‘The relationship between firm growth, size, and age: Estimates for 100 manufacturing industries’, *Journal of Industrial Economics* **35**, 567–581.
- Evans, D. (1987b), ‘Testing alternative theories of firm growth’, *Journal of Political Economy* **95**, 657–674.
- Fan, X., Roberts, M., Xu, D. Y. and Zhang, S. (2012), Demand, cost, and profitability across Chinese exporting firm. Mimeo, Penn State University.
- Foster, L., Haltiwanger, J. C. and Syverson, C. (2012), The slow growth of new plants: Learning about demand? NBER w17853.
- Foster, L., Haltiwanger, J. and Syverson, C. (2008), ‘Reallocation, firm turnover, and efficiency: Selection on productivity or profitability?’, *American Economic Review* **98**, 394–425.
- Gandhi, A., Navarro, S. and Rivers, D. (2011), Does value added overstate productivity dispersion? Identification and estimation of the gross output production function. Mimeo, University of Wisconsin-Madison.
- Gourio, F. and Rudanko, L. (2012), Customer capital. Mimeo, Boston University.
- Hamermesh, D. and Pfann, G. (1996), ‘Adjustment costs in factor demand’, *Journal of Economic Literature* **34**, 1264–1292.
- Hopenhayn, H. (1992), ‘Entry, exit, and firm dynamics in long run equilibrium’, *Econometrica* **60**, 1127–1150.
- Hopenhayn, H. and Rogerson, R. (1993), ‘Job turnover and policy evaluation: A general equilibrium analysis’, *Journal of Political Economy* **101**, 915–938.
- Hsieh, C. and Klenow, P. (2009), ‘Misallocation and manufacturing TFP in China and India’, *The Quarterly Journal of Economics* **124**, 1403.

- Jovanovic, B. (1982), ‘Selection and the evolution of industry’, *Econometrica* **50**, 649–670.
- Klette, T. and Griliches, Z. (1996), ‘The inconsistency of common scale estimators when output prices are unobserved and endogenous’, *Journal of Applied Econometrics* **11**, 343–361.
- Levinsohn, J. and Petrin, A. (2003), ‘Estimating production functions using intermediate inputs to control for unobservables’, *Review of Economic Studies* **70**, 317–341.
- Lippi, F. and Schivardi, F. (2010), Corporate control and executive selection. CEPR D.P. 8301.
- Melitz, M. (2000), Estimating firm-level productivity in differentiated product industries. Mimeo, Harvard University.
- Midrigan, V. and Xu, D. (2010), Finance and misallocation: Evidence from plant-level data. Mimeo, NYU.
- Olley, S. and Pakes, A. (1996), ‘The dynamics of productivity in the telecommunications equipment industry’, *Econometrica* **64**, 1263–1298.
- Petrin, A. and Sivadasan, J. (2011), ‘Estimating lost output from allocative inefficiency, with an application to Chile and firing costs’, *Review of Economics and Statistics* . Forthcoming.
- Restuccia, D. and Rogerson, R. (2008), ‘Policy distortions and aggregate productivity with heterogeneous establishments’, *Review of Economic Dynamics* **11**, 707–720.
- Schivardi, F. and Torrini, R. (2008), ‘Identifying the effects of firing restrictions through size-contingent differences in regulation’, *Labour Economics* **15**, 482–511.
- Syverson, C. (2011), ‘What determines productivity?’, *Journal of Economic Literature* **49**, 326–365. Forthcoming.
- Trajtenberg, M. (1989), ‘The welfare analysis of product innovation with an application to computer tomography scanners’, *Journal of Political Economy* **97**, 444–479.
- Yang, M.-J. (2011), Micro-level misallocation and selection: Estimation and aggregate implications. Mimeo, U.C. Berkeley.

Table 1: Summary statistics for main variables, by sector

	All	Textile and leather	Paper	Chemicals	Minerals	Metals	Machinery	Vehicles
Panel A: Levels								
Sales	126,572 (595,936)	54,055 (109,611)	114,224 (254,860)	169,027 (313,142)	71,758 (119,066)	117,691 (341,440)	106,777 (245,165)	484,263 (2,117,926)
Output	126,514 (572,606)	54,370 (110,007)	110,263 (234,334)	173,626 (319,267)	73,187 (121,902)	119,889 (342,851)	108,486 (246,671)	461,690 (2,019,429)
Workers	525 (2,455)	314 (559)	445 (823)	510 (973)	331 (479)	335 (904)	564 (1,271)	1,953 (8,857)
Panel B: Growth rates								
Δ Sales	.020 (.19)	-.005 (.17)	.027 (.13)	.020 (.14)	.016 (.18)	.021 (.17)	.036 (.19)	.035 (.38)
Δ Output	.023 (.22)	-.007 (.20)	.035 (.16)	.029 (.20)	.023 (.19)	.034 (.20)	.030 (.23)	.043 (.31)
Δ Interm. inputs	.024 (.32)	-.012 (.31)	.039 (.25)	.027 (.32)	.027 (.25)	.031 (.32)	.038 (.34)	.049 (.44)
Δ hours worked	-.004 (.13)	-.017 (.14)	-.005 (.09)	.001 (.12)	-.008 (.12)	.004 (.13)	.001 (.14)	-.003 (.15)
Δ utilized capital	.038 (.20)	.015 (.20)	.052 (.19)	.041 (.21)	.040 (.21)	.053 (.18)	.043 (.19)	.044 (.25)
Δ prices	.021 (.06)	.023 (.05)	.016 (.08)	.021 (.06)	.026 (.05)	.027 (.08)	.017 (.06)	.016 (.04)
Obs.	12,102	2,718	705	1,664	1,192	1,885	3,156	782

Notes: Figures reported are sample averages; standard deviations are in parentheses. Sales and Output are expressed in thousands of 2007 euros. Δ Sales and Δ Output are computed net of growth in firm level prices.

Table 2: Estimates of σ , by sector

	(1) All firms	(2) Single product firms	(3) Single plant firms	(4) Non exporting firms	(5) OLS firms	(6) IV
Textile and leather	4.5 (3.3)	4.7 (3.3)	4.6 (3.4)	8 (3.7)	.27 (.08)	6.1 (.54)
Paper	5.1 (3.4)	4.7 (3.3)	5.5 (3)	5.6 (4)	.39 (.06)	4.6 (1.03)
Chemicals	4.7 (3.3)	5.7 (4.3)	5.4 (3.7)	5.6 (3.7)	.40 (.06)	5.1 (.73)
Minerals	5.4 (3.2)	3.5 (3.1)	5.4 (3.3)	6 (4.1)	-.04 (.10)	5.5 (.91)
Metals	5.5 (3.5)	6.4 (3.5)	5.3 (3.6)	7 (0)	.28 (.05)	4.9 (.61)
Machinery	5 (3.2)	5.1 (3.1)	5.1 (3.2)	7.4 (4)	.39 (.09)	5.7 (.46)
Vehicles	6 (3.6)	8.4 (2.7)	5.7 (3.4)	8.2 (3.2)	.63 (.28)	7.2 (1.58)

Notes: In the first four columns we report the average of the self-reported price elasticities and their standard deviations (in parentheses). In the last two columns, we report the coefficient of the price variable in an OLS and IV specification, respectively, and the standard error of the estimate (in parentheses). Single product firms are defined as those claiming to derive at least 80% of their revenues from a single product line. Single plant firms are those reporting that all their employees work in the same macro-region. The IV column uses unexpected variation in ΔTFP as instrument (see Appendix B.2).

Table 3: Production function estimation: OP and OLS results

	(1) Txt+leather	(2) Paper	(3) Chemicals	(4) Minerals	(5) Metals	(6) Machinery	(7) Vehicles
Panel A: Output deflated with firm prices							
Δk	0.14*** (0.027)	0.09** (0.042)	0.11*** (0.023)	0.12*** (0.033)	0.08*** (0.028)	0.11*** (0.023)	0.17** (0.066)
Δl	0.17*** (0.025)	0.31*** (0.055)	0.23*** (0.030)	0.24*** (0.045)	0.24*** (0.031)	0.17*** (0.029)	0.33*** (0.070)
Δm	0.49*** (0.023)	0.37*** (0.045)	0.58*** (0.027)	0.38*** (0.032)	0.52*** (0.023)	0.52*** (0.019)	0.38*** (0.053)
$\alpha + \beta + \gamma$	0.8	0.77	0.92	0.74	0.85	0.8	0.88
Obs.	1,800	443	1,076	815	1,354	2,071	418
R ²	0.67	0.55	0.71	0.59	0.65	0.72	0.63
Panel B: Output deflated with sectoral prices							
Δk	0.11*** (0.023)	0.06 (0.038)	0.08*** (0.020)	0.10*** (0.030)	0.07*** (0.024)	0.08*** (0.018)	0.13** (0.062)
Δl	0.13*** (0.022)	0.20*** (0.050)	0.17*** (0.025)	0.23*** (0.039)	0.17*** (0.027)	0.15*** (0.023)	0.31*** (0.064)
Δm	0.43*** (0.020)	0.36*** (0.041)	0.55*** (0.025)	0.34*** (0.029)	0.47*** (0.021)	0.50*** (0.017)	0.36*** (0.050)
$\tilde{\alpha} + \tilde{\beta} + \tilde{\gamma}$	0.67	0.62	0.80	0.67	0.71	0.73	0.80
$\frac{\sigma(\tilde{\alpha} + \tilde{\beta} + \tilde{\gamma})}{\sigma^{-1}}$	0.86	0.77	1.01	.82	.86	.91	.96
Obs.	1,806	446	1,083	816	1,356	2,076	419
R ²	0.77	0.72	0.82	0.70	0.76	0.79	0.67

Notes: The dependent variable is the growth rate of output, deflated with firm level prices. Δk is the log difference of the stock of capital used in production, taking capital utilization into account, Δl is the log difference of the number of hours worked and Δm is the log difference of intermediates. All regressions include the control function and year dummies. Robust standard errors are reported in parenthesis. Robust standard errors are reported in parenthesis. Significance levels: *: 10%, **: 5%, *** : 1%

Table 4: Descriptive statistics: Δ TFP and $\Delta\xi$

	<i>N</i>	<i>Mean</i>	<i>Std.dev.</i>	Percentiles				
				<i>5th</i>	<i>25th</i>	<i>50th</i>	<i>75th</i>	<i>95th</i>
Panel A: ΔTFP								
Olley and Pakes	12,102	.008	.14	-.16	-.04	.008	.06	.16
Δ TFP- Δ^3	12,110	.004	.15	-.16	-.04	.004	.05	.16
Panel B: $\Delta\xi$								
Sectoral avg.	12,102	.014	.32	-.46	-.12	.02	.15	.47
Class avg.	10,307	.010	.34	-.48	-.12	.02	.15	.47
Individual reported	1,089	-.012	.46	-.65	-.13	.02	.15	.49

Notes: *Olley and Pakes* refers to estimates of TFP recovered using Olley and Pakes (1996); Δ TFP- Δ^3 reports estimates obtained through the same methodology but using 3-year differences in the estimation of the production function (rather than first differences). *Sectoral* and *class* rows refer to estimates of $\Delta\xi$ obtained using self-reported elasticities averaging firm responses at the sector and class level respectively. The row *individual* reports estimates of $\Delta\xi$ which rely only on the firms that replied directly to the question on price elasticity in the 1996 wave of INVIND.

Table 5: Sales and output growth

	Sales		Price	Output	
	(1) Nominal	(2) Real	(3)	(4) Nominal	(5) Real
ΔTFP	0.597*** (0.018)	0.736*** (0.021)	-.154*** (0.004)	0.807*** (0.019)	0.981*** (0.022)
$\Delta \xi$	0.408*** (0.006)	0.265*** (0.008)	.132*** (0.002)	0.356*** (0.006)	0.222*** (0.007)
Observations	10,611	10,607	10,712	10,648	10,650
R^2	0.67	0.47	0.76	0.59	0.51

Notes: All dependent variables and the demand and TFP shocks are in delta logs. ΔTFP is calculated using Olley and Pakes (1996) control function approach. $\Delta \xi$ is computed using self-reported sectoral price elasticities from the INVIND survey 1996. The columns labeled “Real” use output and sales deflated using individual firm prices, rather than a sectoral deflator. All specifications include region and industry-year fixed effects. Both dependent and independent variables are trimmed to drop outliers above the 99th or below the 1st percentile. Standard errors are calculated from 500 bootstrap simulations. Robust standard errors are reported in parenthesis. Significance levels: *: 10%, **: 5%, *** : 1%

Table 6: Inputs growth

Panel A: Variable inputs

	(1) Hours worked	(2) Intermediate inputs	(3) Utilized capital
Δ TFP	0.013 (0.013)	0.239*** (0.037)	0.006 (0.019)
$\Delta\xi$	0.103*** (0.005)	0.373*** (0.010)	0.110*** (0.007)
Observations	10,569	10,646	10,573
R-squared	0.12	0.28	0.09

Panel B: Quasi-fixed inputs

	(1) Employment	(2) Hires	(3) Separations	(4) Investment rate
Δ TFP	0.062*** (0.010)	0.066*** (0.012)	-0.005 (0.012)	0.092*** (0.019)
$\Delta\xi$	0.074*** (0.004)	0.068*** (0.004)	-0.015*** (0.004)	0.039*** (0.006)
Observations	10,552	10,651	10,650	8,448
R-squared	0.11	0.10	0.04	0.10

Notes: All dependent variables and the demand and TFP shocks are in delta logs. Δ TFP is calculated using the Olley and Pakes (1996) control function approach. $\Delta\xi$ is computed using self-reported sectoral price elasticities from the INVIND survey 1996. All specifications include region and industry-year fixed effects. Both dependent and independent variables are trimmed to drop outliers above the 99th or below the 1st percentile. Standard errors are calculated from 500 bootstrap simulations. Robust standard errors are reported in parenthesis. Significance levels: *: 10%, **: 5%, *** : 1%

Table 7: Responses to changes in TFP and market appeal: Model's prediction vs. empirical estimates

	$\Delta(p+q)$		Δq		Δp		Δx	
	<i>Predicted</i>	<i>Actual</i>	<i>Predicted</i>	<i>Actual</i>	<i>Predicted</i>	<i>Actual</i>	<i>Predicted</i>	<i>Actual</i>
ΔTFP	2.2	0.81	2.8	0.98	-0.56	-0.15	2.2	0.09
$\Delta\xi$	0.56	0.41	0.44	0.27	0.11	0.13	0.56	0.20

Notes: $p+q$ indicates nominal output, q is real output deflated with firm level prices, x represents inputs (labor, capital and intermediates). The columns labeled *Predicted* report the elasticity implied by the theoretical model, conditional on the estimates of the parameters of the demand and production functions. The columns labeled *Actual* display the reduced form estimate of the same elasticity. For these latter, we report the estimate obtained using output for TFP (Table 5, columns (4) and (5)) and sales for market appeal (Table 5, columns (1) and (2)). For inputs, we use the simple average of the elasticities of variable inputs in Panel A of Table 6.

Table 8: Lagged effects

	(1) Output	(2) Price	(3) Employment	(4) Investment rate
ΔTFP_t	0.988*** (0.030)	-0.160*** (0.006)	0.076*** (0.014)	0.115*** (0.026)
ΔTFP_{t-1}	0.153*** (0.022)	-0.041*** (0.004)	0.110*** (0.014)	0.092*** (0.026)
ΔTFP_{t-2}	0.037* (0.021)	-0.021*** (0.004)	0.062*** (0.013)	0.103*** (0.026)
$\Delta \xi_t$	0.240*** (0.009)	0.133*** (0.002)	0.075*** (0.005)	0.041*** (0.009)
$\Delta \xi_{t-1}$	-0.027*** (0.008)	0.010*** (0.002)	0.024*** (0.004)	0.024*** (0.008)
$\Delta \xi_{t-2}$	0.002 (0.007)	-0.001 (0.001)	0.023*** (0.004)	0.035*** (0.008)
Observations	5,418	5,429	5,371	4,378
R-squared	0.52	0.79	0.16	0.11

Notes: All dependent variables (except investment rate) and the demand and TFP shocks are in delta logs. ΔTFP is calculated using the Olley and Pakes (1996) control function approach. Since lags are included in the specification, the control function is augmented to include the forecast for the next year's investment, the expected change in technical capacity and two lags of the demand appeal shocks as additional controls; more details are in Appendix B.7. $\Delta \xi$ is computed using self-reported sectoral price elasticities from the INVIND survey 1996. Output is deflated using firm level prices. Employment is measured as the number of workers employed at the firm at the end of the year. All specifications include region and industry-year fixed effects. Both dependent and independent variables are trimmed to drop outliers above the 99th or below the 1st percentile. Robust standard errors are reported in parenthesis. Standard errors are calculated from 500 bootstrap simulations. Robust standard errors are reported in parenthesis. Significance levels: *, 10%, **, 5%, ***, 1%

Table 9: Evidence of organizational hurdles: Self-reported

	(1) Output	(2) Price	(3) Employment	(4) Investment rate
ΔTFP	1.041*** (0.035)	-0.167*** (0.007)	0.097*** (0.018)	0.140*** (0.035)
$\Delta TFP \times$ Organizational hurdles	-0.120** (0.047)	0.020** (0.009)	-0.046** (0.023)	-0.066 (0.045)
$\Delta \xi$	0.225*** (0.010)	0.131*** (0.003)	0.075*** (0.006)	0.038*** (0.012)
$\Delta \xi \times$ Organizational hurdles	-0.004 (0.011)	0.001 (0.003)	0.002 (0.007)	0.004 (0.016)
Organizational hurdles	0.006** (0.002)	-0.001 (0.001)	0.002 (0.002)	-0.002 (0.004)
Observations	8,035	8,070	7,960	6,414
R-squared	0.51	0.77	0.13	0.10

Notes: All dependent variables (except investment rate) and the demand and TFP shocks are in delta logs. ΔTFP is calculated using the Olley and Pakes (1996) control function approach. $\Delta \xi$ is computed using self-reported sectoral price elasticities from the INVIND survey 1996. *Organizational hurdles* is an indicator variable for firms that face higher costs in reshaping their internal organization. It takes value 1 for firms reporting that they have not met their investment plans for the past year due to “problems with the internal organization of the firm”. Output is deflated using firm level prices. Employment is measured as the number of workers employed at the firm at the end of the year. All specifications include region and industry-year fixed effects. Both dependent and independent variables are trimmed to drop outliers above the 99th or below the 1st percentile. Robust standard errors are reported in parenthesis. Standard errors are calculated from 500 bootstrap simulations. Robust standard errors are reported in parenthesis. Significance levels: *: 10%, **: 5%, *** : 1%

Table 10: Evidence of organizational hurdles: Family firms

	(1) Output	(2) Price	(3) Employment	(4) Investment rate
ΔTFP	1.051*** (0.024)	-0.164*** (0.005)	0.078*** (0.013)	0.085*** (0.020)
$\Delta TFP \times \text{Family}$	-0.145*** (0.036)	0.022*** (0.008)	-0.036* (0.019)	-0.013 (0.030)
$\Delta \xi$	0.221*** (0.007)	0.133*** (0.002)	0.075*** (0.004)	0.031*** (0.006)
$\Delta \xi \times \text{Family}$	0.004 (0.011)	-0.002 (0.003)	-0.001 (0.006)	0.006 (0.010)
Family	-0.004* (0.002)	0.001** (0.000)	0.005*** (0.001)	-0.002 (0.002)
	10,650	10,683	10,522	8,428
	0.52	0.76	0.12	0.05

Notes: All dependent variables (except investment rate) and the demand and TFP shocks are in delta logs. ΔTFP is calculated using the Olley and Pakes (1996) control function approach. $\Delta \xi$ is computed using self-reported sectoral price elasticities from the INVIND survey 1996. *Family* is an indicator variable for firms that are controlled by an individual/family or by the government. Output is deflated using firm level prices. Employment is measured as the number of workers employed at the firm at the end of the year. All specifications include region and industry-year fixed effects. Both dependent and independent variables are trimmed to drop outliers above the 99th or below the 1st percentile. Robust standard errors are reported in parenthesis. Standard errors are calculated from 500 bootstrap simulations. Robust standard errors are reported in parenthesis. Significance levels: *: 10%, **: 5%, *** : 1%

Figure 1: Distribution of price changes in 1997, by sector


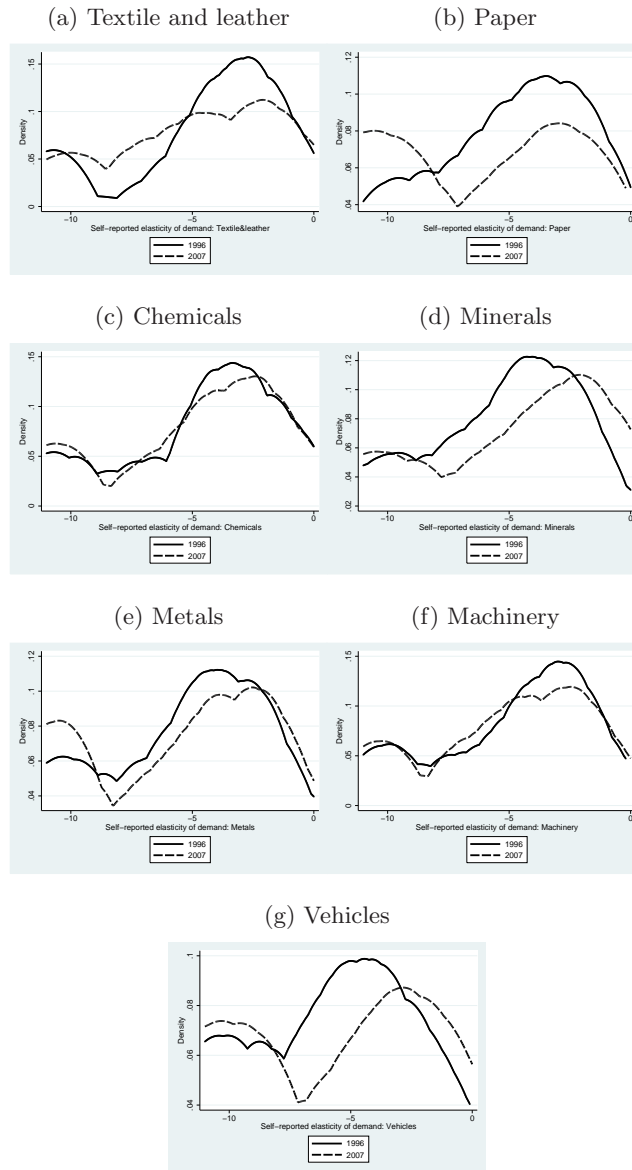


Figure 2: Distribution of self-reported elasticities in 1996 and 2007, by sector



Appendix

A Model details and extensions

A.1 Equilibrium with full capacity utilization

We derive the optimal firm choices when the capital stock is binding. When inequality (9) holds, the firm is characterized by a high level of demand and/or productivity shocks, so that it wants to expand its production accordingly, but the capital stock in place is below the optimal unconstrained level. In this case, the firm uses all its capital stock, which becomes a parameter, and the production function is Cobb-Douglas in just L and M :

$$Q_{it} = (\Omega_{it} \bar{K}_{it}^\alpha) L_{it}^\beta M_{it}^\gamma \quad (\text{A-1})$$

It is then straightforward to derive the optimal quantities when the capital constraint binds:

$$q_{it}^* = \bar{c}_q + \frac{\sigma}{\lambda}(\omega_{it} + \alpha \bar{k}_{it}) + \frac{(\beta + \gamma)}{\lambda} \xi_{it} \quad (\text{A-2})$$

$$p_{it}^* = \bar{c}_p - \frac{1}{\lambda}(\omega_{it} + \alpha \bar{k}_{it}) + \frac{(1 - \beta - \gamma)}{\lambda} \xi_{it} \quad (\text{A-3})$$

$$x_{it}^* = \bar{c}_x + \frac{(\sigma - 1)}{\lambda}(\omega_{it} + \alpha \bar{k}_{it}) + \frac{1}{\lambda} \xi_{it} \quad (\text{A-4})$$

where $\lambda \equiv \beta + \gamma + (1 - \beta - \gamma)\sigma$ and $x = l, m$. When the firm hits the capital constraint, the degree of returns to scale decreases from $\alpha + \beta + \gamma$ to $\beta + \gamma$, as the capital stock is now a fixed input. As a consequence, all the endogenous variables become less responsive to shocks, given that $\lambda > \theta$ (recall that the elasticity to productivity and demand in the unconstrained case is $\frac{\sigma-1}{\theta}$ and $\frac{1}{\theta}$ respectively). Given the elasticity of output and inputs to shocks changes for constrained firms, we exclude these observations from the analysis in Section 6 of the paper.

A.2 The dynamic problem

As discussed in the text, the firm cannot alter its level of capital in place within the period, but it chooses the degree of capital utilization. Capital in place can be increased through investments which will deliver capital in the next period. A higher level of investment decreases the likelihood that the firm will face capital constraints in the following year. The problem is described by the three state variables $\{\bar{K}_{it}, \Omega_{it}, \Xi_{it}\}$. Let $\mathcal{I}_{it} \equiv \mathcal{I}(\bar{K}_{it}, \Omega_{it}, \Xi_{it})$ be the indicator function for the case in which the capital stock is not constrained, i.e., when equation (9) is satisfied. To ease the notation, we have dropped the dependence on the

state variables. Similarly, define Π_{it} and $\bar{\Pi}_{it}$ as the static profits for the constrained and unconstrained case respectively. Simple algebra shows that

$$\Pi_{it} = C_{\Pi} \Omega_{it}^{\frac{\sigma-1}{\theta}} \Xi_{it}^{\frac{1}{\theta}} \quad (\text{A-5})$$

$$\bar{\Pi}_{it} = C_{\bar{\Pi}} \Omega_{it}^{\frac{\sigma-1}{\lambda}} \Xi_{it}^{\frac{1}{\lambda}} \bar{K}_{it}^{\frac{\alpha(\sigma-1)}{\lambda}} - p_K \bar{K}_{it} \quad (\text{A-6})$$

where $c_{\Pi}, c_{\bar{\Pi}}$ are constants. Note that in the unconstrained case the within-period profits do not depend on \bar{K}_{it} .

The recursive formulation of the dynamic problem is as follows:

$$\begin{aligned} V(\bar{K}_{it}, \Omega_{it}, \Xi_{it}) = \max_{\bar{K}_{it+1}} \{ & \Pi_{it} \mathcal{I}_{it} + \bar{\Pi}_{it} (1 - \mathcal{I}_{it}) - p_I (\bar{K}_{it+1} - (1 - \delta) \bar{K}_{it}) + \\ & \psi E(V(\bar{K}_{it+1}, \Omega_{it+1}, \Xi_{it+1}) | \Omega_{it}, \Xi_{it}) \} \\ & \text{subject to} \\ & \bar{K}_{it+1} = I_{it} + (1 - \delta) \bar{K}_{it} \\ & \Omega_{it+1} \sim F_{\Omega}(\cdot | \Omega_{it}) \\ & \Xi_{it+1} \sim F_{\Xi}(\cdot | \Xi_{it}) \end{aligned} \quad (\text{A-7})$$

where ψ is the discount factor and p_I is the price of investment. Under the assumptions on the stochastic processes followed by the shocks, standard dynamic programming considerations ensure that, if $I_t > 0$, the policy function for investment $g(\bar{K}_t, \Omega_{it}, \Xi_{it})$ is increasing in Ξ_{it}, Ω_{it} for each level of \bar{K}_{it} . This implies that we can invert it and express the productivity shock as function of $\{I_{it}, \Xi_{it}, \bar{K}_{it}\}$

A.3 Extension to multi-product firms

33% of the firms in our sample report deriving all of their revenues from a single product line; for 51% of the firms the share is at least 80%. This implies that half of the firms in our sample obtains revenues from selling a variety of products. Since we are assuming the same elasticity for all firms in the same sector, the fact that a firm sells different products in the same sector (e.g. shirts and sweaters within textiles) has no consequence for us. Here we show that even the presence of generic multiproduct firms does not cause problems for our identification of the demand and productivity shocks. To keep the notation simple, we eliminate both the firm and time subscript and focus on the product subscript. Each firm produces G goods. We assume that the number of goods produced by each firm is constant over time. We want to show under what assumptions our procedure based on aggregate firm level data still recovers meaningful indicators of demand and productivity shocks. For

demand, it is natural to assume that each product has its own demand schedule. Then we can either assume that the demand shock is common to all goods, $\xi_g = \xi$, or that each good has its own shock ξ_g . For production, one can assume that the firm has a unique production line on which all goods are produced or that each good is produced with a separate production function. In the latter case, the productivity shock can be common to all production functions, $\omega_g = \omega$, or each production function can have its own productivity shock ω_g . In what follows, we use capital boldface to indicate aggregate quantities: $\mathbf{X} \equiv \sum_g X_g$ and small case boldface for its log: $\mathbf{x} \equiv \log\left(\sum_g X_g\right)$.

The typical assumption in the literature is that each good has its own production function but that both the demand and the productivity shocks are common across all goods (Foster et al. 2008, De Loecker 2011):

$$\begin{aligned} Q_g &= \Omega K_g^\alpha L_g^\beta M_g^\gamma \\ Q_g &= P_g^{-\sigma} \Xi \end{aligned}$$

In this case, the solution for the single product case applies to each single good, so that:

$$\begin{aligned} q_g^* &= q^* = c_q + \frac{\sigma}{\theta} \omega + \frac{(\alpha + \beta + \gamma)}{\theta} \xi \\ p_g^* &= p^* = c_p - \frac{1}{\theta} \omega + \frac{(1 - \alpha - \beta - \gamma)}{\theta} \xi \end{aligned}$$

So the firm produces exactly the same amount and applies the same price for all goods. As a consequence, the average price change is just the price change of the common price, Δp . Moreover, using the fact that $\sum_g P_g Q_g = G * P^* Q^*$ and the assumption that G is fixed over time, the change in the log of revenues is given by:

$$\Delta \ln \left(\sum_g P_g Q_g \right) = \Delta p^* + \Delta q^*$$

Therefore, when deflating the change in the log of the revenues with the average change in prices Δp^* , we obtain the correct measure of real output. Now use the production function and the fact that each product accounts for an equal share of total output:

$$\mathbf{Q} \equiv \sum_g Q_g = \Omega \sum_g K_g^\alpha L_g^\beta M_g^\gamma = G^{1-\alpha-\beta-\gamma} \Omega \mathbf{K}^\alpha \mathbf{L}^\beta \mathbf{M}^\gamma$$

where $\mathbf{K} = \sum_g K_g = G * K$ and similarly for L and M , so that

$$\Delta \mathbf{q} = \Delta \omega + \alpha \Delta \mathbf{k} + \beta \Delta \mathbf{l} + \gamma \Delta \mathbf{m}$$

This shows that, under the stated assumptions, the estimation procedure correctly recovers TFP even in the presence of multi-product firms. Similarly, using the fact that $\ln\left(\sum_g Q_g\right) = \ln(G * Q)$ and $\ln\left(\sum_g P_g^{-\sigma}\right) = \ln(G * P_g^{-\sigma})$, we can quickly verify that from firm-level aggregate sales and average price changes we recover the demand shock:

$$\Delta\xi = \Delta\mathbf{q} - \sigma\Delta p.$$

Things are slightly more complicated when the firm sells G products, each with its own demand shock:

$$Q_g = P_g^{-\sigma}\Xi_g \tag{A-8}$$

As long as all goods are produced with only one production function,

$$\mathbf{Q} \equiv \sum_g Q_g = \Omega K^\alpha L^\beta M^\gamma \tag{A-9}$$

we can still recover TFP and a meaningful aggregate measure of demand shocks. The assumption of one production function for all goods is consistent with a production technology in which the firm has one production line on which it can produce alternatively the different goods it sells. Define $C(\mathbf{Q})$ as the cost function to produce the aggregate quantity \mathbf{Q} . Simple algebra shows that $C(\mathbf{Q}) = c_x \left(\frac{\mathbf{Q}}{\Omega}\right)^{1/(\alpha+\beta+\gamma)}$ where c_x is a constant that depends on input prices and the production function coefficients. The firm problem is

$$\max_{\{Q_g\}_{g=1}^G} \sum_g \left(Q_g^{\frac{\sigma-1}{\sigma}} \Xi_g^{\frac{1}{\sigma}} \right) - c_x \left(\frac{\mathbf{Q}}{\Omega} \right)^{1/(\alpha+\beta+\gamma)} \tag{A-10}$$

subject to (A-8) and (A-9). The FOCs are:

$$\frac{\sigma-1}{\sigma} Q_g^{-\frac{1}{\sigma}} \Xi_g^{\frac{1}{\sigma}} = \frac{c_x}{\alpha+\beta+\gamma} \left(\frac{\mathbf{Q}}{\Omega} \right)^{\frac{1-(\alpha+\beta+\gamma)}{\alpha+\beta+\gamma}} \Omega^{-1} \tag{A-11}$$

from which we obtain:

$$\frac{Q_g}{Q_h} = \frac{\Xi_g}{\Xi_h}.$$

Summing over h , total output is:

$$\mathbf{Q} = \frac{Q_g}{\Xi_g} \sum_{h=1}^G \Xi_h \equiv \frac{Q_g}{\Xi_g} \Xi$$

Substitute in (A-11), solve for price and quantities and take logs:

$$q_g^* = C_q + \frac{\sigma}{\theta}\omega + \xi_g - \frac{(1-(\alpha+\beta+\gamma))\sigma}{\theta}\xi \tag{A-12}$$

$$p_g^* = C_p - \frac{1}{\theta}\omega + \frac{(1-(\alpha+\beta+\gamma))}{\theta}\xi \tag{A-13}$$

So the price of the single good does not depend on the idiosyncratic demand shock (apart from the effect through Ξ , common to all goods). The average price is therefore equal to the individual price, so that the average price change is the correct aggregator of the change in individual prices. It is immediate to verify that log total output is:

$$\mathbf{q}^* = c_q + g + \frac{\sigma}{\theta}\omega + \frac{\alpha + \beta + \gamma}{\theta}\boldsymbol{\xi} \quad (\text{A-14})$$

This shows that our estimation procedure recovers the correct measure of TFP and an aggregate indicator of demand appeal shocks, $\boldsymbol{\xi} = \ln\left(\sum_g \Xi_g\right)$.

Consider finally the case in which each product has its own demand and production function shocks. Given the CES demand assumption, the firm can be seen as a collection of products. In this case, revenues are equal to:

$$\mathbf{P} * \mathbf{Q} \equiv \sum_g P_g * Q_g \propto \sum_g \Omega_g^{\frac{\sigma-1}{\theta}} \Xi_g^{\frac{1}{\theta}} \quad (\text{A-15})$$

Unfortunately, from this expression we cannot factor out any combination of ω_g and ξ_g . When we subtract the average price change from the change in revenues, therefore, we do not recover exactly real output.

B Robustness

B.1 Validation of INVIND price variable

The changes in firm-level prices are the key piece of information for this study; without such variable we would not be able to separately identify productivity and market appeal. Like all the information contained in the INVIND survey, price changes are self reported by the interviewee and one may worry about the accuracy of these statements.

To assess the reliability of the price variable, we compare a price index based on price changes reported by respondent to the INVIND survey with the official producer price index (PPI) constructed by the National Statistical Office (ISTAT). The INVIND price index is constructed as the average of the price changes reported by individual firms weighted using sampling weights provided by the Bank of Italy. ISTAT releases estimates of the inflation in producer prices monthly. We use the estimates released in the month of March since INVIND interviews take place in that same month. Both the INVIND price index and the official PPI are normalized to 1 for the year 2009.

In Figure A-1 we plot the time series of the two price indexes for each ATECO sector used in the analysis. The two series are highly correlated and generally close in levels.

The only exception is the electrical machinery sector where the INVIND price index shows prices falling over the sample period, whereas the ISTAT PPI indicates positive price growth. Whereas the correlation could be simply driven by inflation, which would lead both series to trend up, the fact that the actual level of the two indexes are quite similar provides strong evidence that the INVIND price variable picks up more than just noise.

B.2 Estimates of price elasticity of demand

The INVIND survey question eliciting elasticity actually refers to the elasticity of revenues to a 10% increase in price. We apply the following transformation to obtain the familiar price elasticity of quantity. Define $\varepsilon_{10\%}^R$ as the number provided by the interviewee; we can quickly show that the elasticity can be obtained as: $\varepsilon_{1\%}^Q = \frac{\varepsilon_{10\%}^R}{10} - 1$. Also note that the phrasing of the question induces censoring in the self-reported revenue elasticities. In fact, for every firm with elasticity above 10, a 10% increase in price would cause the maximum revenue loss reportable, 100%. This implies that all firms with revenue elasticity greater than 10 (or price elasticity greater than 11) will be bunched at 10 (11). Inspecting the data, the share of firms bunched at 10 does not seem alarmingly large. Furthermore, we have checked the robustness of the results estimating a Tobit model that accounts for censoring. The results (available upon request) are not significantly affected.

We now discuss alternative estimates of price elasticity reported in Table 2. The first potential problem we want to address relates the presence of multiproduct firms. Though we have presented in appendix A.3 a theoretical argument for the robustness of our procedure to the presence of multiproduct firms, in the second column of Table 2 we present a robustness check to dispel any residual concern. Exploiting a question of the survey, we compute sectoral averages of self-reported elasticities for the subset of firms reporting that they have earned at least 80% of their revenues from a single product line. Figures for single product firms are similar to the sample averages. Next, we check whether using firm level, rather than plant level, information has any effect on the elasticities by reporting estimates based on the group of single plant firms.²⁴ Once again, the estimates do not change much.

Exporting firms face residual demand curves with different slopes in the domestic and foreign markets, leaving some doubt on the which figures they are reporting in the survey.

²⁴We define as single plant firms those that report to have 100% of their workforce employed in a single macro-region of Italy. The INVIND survey includes an explicit question on the number of plants but the question was introduced long after 1996 and we therefore elected not to use it. However, all the firms in the sample reporting that the entirety of their labor works in a single region also say that they are single plant in the wave of the survey asking that question.

Since INVIND contains information on the amount of revenue generated through exports by each firm, we can concentrate our attention on the small number of firms in our sample that do not export. The fourth column displays demand elasticities for non exporting firms; these firms appear to face a more elastic demand. This suggests that the “best” firms (i.e. those with some degree of market power) are more likely to become exporters. Overall, though, elasticities estimated using the subsample of non-exporting firms are in line with our baseline figures, with the noticeable exception of Textile and leather, for which the elasticity almost doubles. In the last two columns, we report elasticity values from direct estimation of the demand function in equation (12). The OLS estimates are much lower than the self-reported ones; this is expected since the endogeneity of price should bias the elasticity towards zero. To obtain a number that should be more comparable to self-reported elasticities, we follow an instrumental variables approach analogous to that in Foster et al. (2008) by using the estimates of TFP as cost shifters to instrument for price changes. However, the construction of the instrument is different in our case and requires further explanation.

We start by obtaining estimates of the market appeal shock ($\Delta\xi$) following the steps described in section 4 and using the self reported measure of price elasticity. The estimated $\Delta\xi$'s are necessary to run the Olley and Pakes (1996) procedure and obtain $\Delta\omega$ and ϵ from equation 14; the sum of the two being our measure of Δ TFP. The figures showed in column 6 of Table 2 are obtained instrumenting prices with the ϵ component.

Instrumenting prices using ϵ allows using cost variation that is completely unexpected by the firm. Using the whole Δ TFP, as done in Foster et al. (2008), involves including in the instrument also the persistent part of the TFP process ($\Delta\omega$) to which firms can react, at least in expectation. In that sense the exclusion restriction is more likely to hold for our instrument. At the same time, our IV estimates of price elasticities are not independent from the self reported ones, which are still used in the first step of the procedure. This means that our IV estimates of elasticities cannot completely validate the self-reported measure we use in our preferred specifications. However, the fact that the results obtained with these two approaches are similar is surely a comforting sign.

B.3 Production function estimation details

In estimating the production function we face the usual problem of the endogeneity of inputs. We address it using the control function approach first introduced by Olley and Pakes

(1996). Whereas they assume scalar unobservability, we introduce a vector of unobserved components, a demand and a productivity shifter. It follows that the policy function for investments depends not only on the initial capital stock and on productivity, but also on demand, as shown in equation (11). We therefore need to include controls for demand as well. Given that demand can be estimated independently from production, we address the issue by including our estimate of $\Delta\xi$ in the control function, as suggested by Akerberg et al. (2007). As discussed in Section A.2, this gives us a valid control function. By log linearizing equation (11) and taking first differences, we can express the change in log of TFP as a function of the change in the log of the demand shock, the capital stock in place and investment. To improve the fit, we also include the interactions of the changes, up to the third degree polynomial.

Ideally, one should use the contemporaneous and lagged values of the variables rather than their first difference, as $\Delta h(x) \neq h(\Delta x)$. Unfortunately, we can only compute the first difference of the demand shock. As a further check, given that we do observe the levels of both k and i , we have experimented with a specification in which the polynomial is in the change in the demand shock and in the current and lagged levels of k and i . The estimated coefficients, reported in the Appendix Table A-1, are virtually unchanged. This suggests that a polynomial in the first differences is flexible enough to proxy for the unobserved productivity shock.

A final distinctive characteristic of our setting is that we allow firms to choose the degree of capital utilization. Thus, effective capital is not a predetermined variable, but it is chosen after observing $\{\omega_{it}, \xi_{it}\}$, like labor and intermediates. This implies that, provided a valid control function is used, we can estimate all inputs' coefficients in a single regression, without the need of the second stage as in most applications of the Olley and Pakes (1996) procedure.²⁵ Finally, the control function approach has been criticized by Akerberg, Caves and Frazer (2006) on the basis of the fact that the input demand functions should be perfectly collinear with the control function h . We assume that, due to strikes, power shortages, machines breakdowns and delivery lags there are variations in k_{it} , l_{it} and m_{it} independent from ω_{it} and ξ_{it} .²⁶

²⁵We ignore the problem of selection also stressed by Olley and Pakes (1996): in our data we cannot distinguish exit from simple nonresponse to the questionnaire.

²⁶Alternatively, one can assume the DGP postulated by Akerberg et al. (2006), where labor, intermediates and capital are set prior to the investment decision, and Ω changes between the two points in time.

B.4 Measurement error in TFP

Throughout the analysis, we follow Foster et al. (2008) and measure the productivity shocks using the whole residual of the production function: $\Delta TFP = \Delta\omega + \Delta\epsilon$. This means that TFP shocks contain a noise component (the ϵ) which could bias the results downwards. Since the low response of growth indicators to TFP shocks is one of our key findings, we want to make sure that it cannot be completely explained by the presence of measurement error in ΔTFP . In Table A-2 we recompute the elasticity of growth in output, price, and quasi-fixed inputs to productivity and demand shocks and measure the former using $\Delta\omega$. As expected, the elasticities to TFP shocks increase (in absolute value), but they are nowhere near to what the model would predict. Furthermore, it is still the case that the gap between measured and predicted elasticities is more severe for TFP than for market appeal shocks. In other words, the two main findings stated in section 5 are not a mere product of measurement error in our measure of TFP.

B.5 Measuring serial correlation in ξ and TFP

Assume that ξ and TFP are AR(1) processes, that is

$$\xi_t = \rho^\xi \xi_{t-1} + \epsilon_t^\xi \tag{A-16}$$

$$TFP_t = \rho^{TFP} TFP_{t-1} + \epsilon_t^{TFP} \tag{A-17}$$

with ϵ_t^{TFP} and ϵ_t^ξ iid across time and potentially correlated contemporaneously. Take first differences of (A-16):

$$\Delta\xi_t = \rho^\xi \Delta\xi_{t-1} + \Delta\epsilon_t^\xi \tag{A-18}$$

Clearly, $\Delta\xi_{t-1}$ is correlated with $\Delta\epsilon_t^\xi$ via ϵ_{t-1}^ξ so that an OLS regression of $\Delta\xi_t$ on its lag would give inconsistent estimates. However, both $\Delta\xi_{t-2}$ and ΔTFP_{t-2} are valid instruments, as uncorrelated with ϵ_{t-1}^ξ but correlated with $\Delta\xi_{t-1}$ via ξ_{t-2} . The same reasoning applies to ΔTFP .

B.6 Frictions and production function estimates

Although we have argued in the main text that input usage should account for temporary shutdowns, we could still worry that output might be reduced even when capital utilization and hours are not, for example if part of the inputs are used not for production but to put new capital in place. In this (quite special) case, our estimates could be biased. Cooper

and Haltiwanger (2006) estimate that disruption of production caused by new investment reduces profits by nearly 20 percent, while Bloom (2009) obtains a substantially smaller number. Disruption costs in the form of reduced output given inputs would bias our estimates of α , β , and γ . In fact, instead of observing the output Q_{it} that should be delivered by a given combination of production factors as in equation 4 we would observe a lower level of output (\widetilde{Q}_{it}), scaled down by the disruption costs (λ)

$$\widetilde{Q}_{it} = [\Omega_{it} K_{it}^{\alpha} L_{it}^{\beta} M_{it}^{\gamma}] \cdot \lambda(I)$$

where λ is lower than one when investment is strictly positive. This would lead us to underestimate the coefficients of the production function, as we observe lower output when the firm is increasing its capital stock. Note however that our estimates are robust to any cost that is proportional to output and that is paid whenever the firm invests, which is the case considered in Cooper and Haltiwanger (2006). Since we follow Olley and Pakes (1996) in estimating the production function, we only consider periods with positive investment, so that this cost is always paid in the observations used to estimate production. As such, it drops out once we take the first difference of the logs. A more serious problem arises when forgone output depends on the size of the investment, rather than being a fixed proportion. We are not aware of any estimation of such specification of the adjustment cost function.

B.7 Control function with lagged shocks

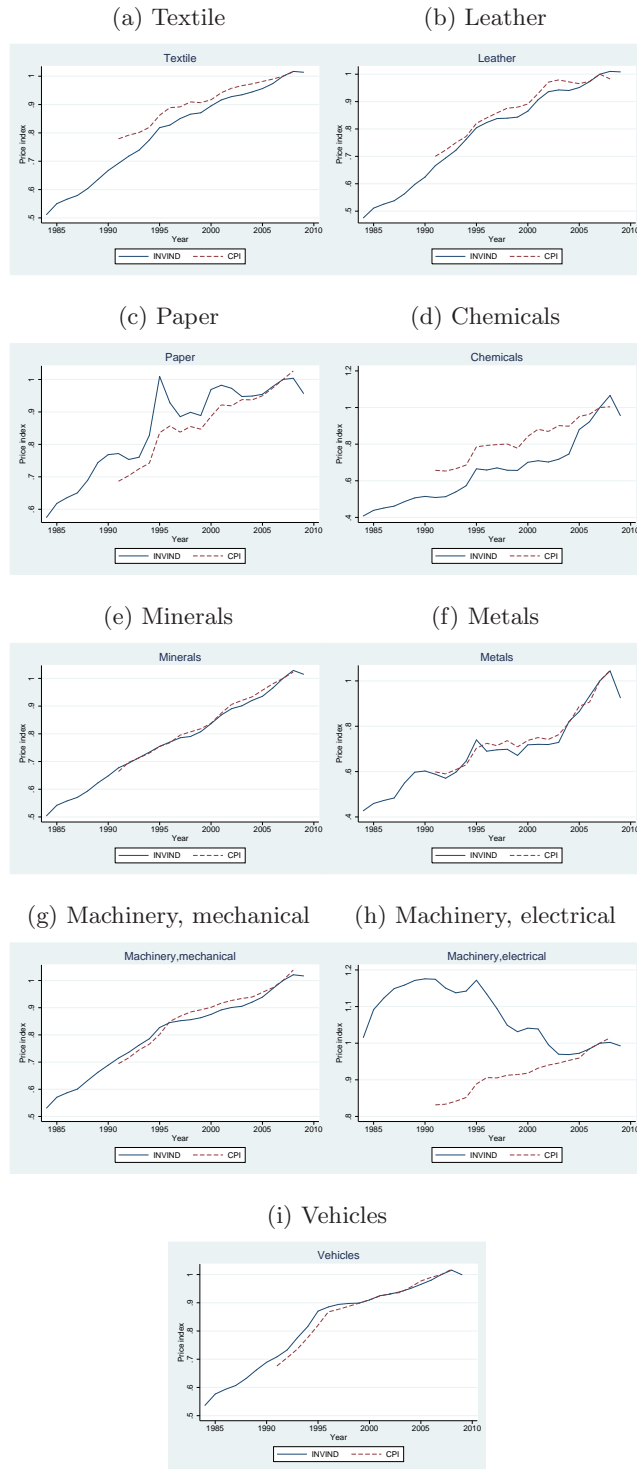
Next, we discuss the validity of the control function when using the lagged values of the shocks as regressors, as in the specifications in Table 8. If the firm responds to lagged shocks, then the investment equation also depends on such shocks. This means that we need to increase the number of controls in the control function to ensure invertibility. We use the forecast for the next year's investment, the expected change in technical capacity and two lags of the demand appeal shocks as additional controls. Therefore; we have 7 state variables (the capital stock in place and the current, lag one and lag two values of each shock) and 7 controls (the investment rate, the installed capital stock, the current, lag one and lag two values of the demand shock, and the forecast for the next year's investment and the expected change in technical capacity). For these to be a valid set of controls, we need to assume that not only current investment but also the forecast for the next year's investment and the expected change in technical capacity are monotonic functions of the shocks: higher demand or TFP today implies higher expected investment next year and

higher expected change in technical capacity. Given the monotonicity that characterizes the model, these are natural assumptions.

B.8 Elasticities to shock, sector by sector

Table A-3 reports sector-by-sector elasticities to shocks of the main growth measures analyzed in the paper. It emerges that the results of the pooled data are not driven by specific sectors. The qualitative patterns are the same for each of the sectors included in the analysis.

Figure A-1: Comparison of price index based on INVIND self-reported price changes and official Producer Price Index computed by the Italian Statistical Office (ISTAT)



Notes: The INVIND price index is computed averaging firm-reported price changes using sample weights provided by the Bank of Italy. The ISTAT PPI refers to the figure released in the month of March. Both indexes are normalized to 1 in 2009.

Table A-1: Additional estimates of the production function

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Txt+leather	Paper	Chemicals	Minerals	Metals	Machinery	Vehicles
Panel A: OP controls in levels for i, k							
Δk	0.14*** (0.026)	0.10** (0.045)	0.10*** (0.024)	0.12*** (0.032)	0.08*** (0.028)	0.12*** (0.024)	0.18** (0.074)
Δl	0.16*** (0.025)	0.30*** (0.059)	0.22*** (0.030)	0.21*** (0.045)	0.22*** (0.031)	0.17*** (0.030)	0.33*** (0.076)
Δm	0.48*** (0.022)	0.37*** (0.044)	0.58*** (0.028)	0.36*** (0.031)	0.51*** (0.023)	0.52*** (0.019)	0.36*** (0.062)
Obs.	1,805	443	1,083	815	1,354	2,072	419
R^2	0.70	0.60	0.73	0.63	0.67	0.73	0.66
Panel B: Using three year differences							
Δk	0.13*** (0.038)		0.11*** (0.041)	0.17*** (0.044)	0.17*** (0.041)	0.08*** (0.031)	
Δl	0.23*** (0.025)		0.31*** (0.041)	0.25*** (0.054)	0.27*** (0.033)	0.20*** (0.024)	
Δm	0.58*** (0.022)		0.62*** (0.030)	0.40*** (0.031)	0.64*** (0.028)	0.63*** (0.017)	
Observations	1,043		547	473	738	1,098	
R-squared	0.77		0.77	0.71	0.76	0.82	

Notes: Panel A reports the estimates of the production function when using the current and lagged levels of i, k in the control function instead of their first difference. Panel B reports the estimates of the production function when using 3 year differences instead of first differences. Robust standard errors are reported in parenthesis. Significance levels: *: 10%, **: 5%, *** : 1%

Table A-2: Elasticities to demand and TFP shocks, robustness to measurement error

	(1) Output	(2) Price	(3) Employment	(4) Investment
$\Delta\omega$	1.478*** (0.112)	-0.309*** (0.024)	0.364*** (0.053)	0.473*** (0.073)
$\Delta\xi$	0.230*** (0.010)	0.128*** (0.003)	0.080*** (0.004)	0.040*** (0.006)
Observations	7,371	7,451	7,337	7,431
R-squared	0.26	0.69	0.12	0.06

Notes: This table replicates some of the pooled estimates displayed in Tables 5 and 6 using the $\Delta\omega$ in equation 14 as the measure of productivity rather than the $\Delta TFP = \Delta\omega + \varepsilon$ we used in the main specification. All dependent variables and the demand and productivity shocks are in delta logs. $\Delta\omega$ is calculated using Olley and Pakes (1996) control function approach. $\Delta\xi$ is computed using self-reported sectoral price elasticities from the INVIND survey 1996. All specifications include region and industry-year fixed effects. Both dependent and independent variables are trimmed to drop outliers above the 99th or below the 1st percentile. Standard errors are calculated from 500 bootstrap simulations. Robust standard errors are reported in parenthesis. Significance levels: *: 10%, **: 5%, *** : 1%

Table A-3: Elasticities to demand and TFP shocks, by sector

Panel A: Output							
	Txt+leather	Paper	Chemicals	Minerals	Metals	Machinery	Vehicles
Δ TFP	0.925*** (0.050)	0.993*** (0.058)	1.034*** (0.064)	0.957*** (0.050)	0.928*** (0.055)	0.980*** (0.046)	1.021*** (0.064)
$\Delta\xi$	0.296*** (0.017)	0.103*** (0.012)	0.227*** (0.016)	0.151*** (0.015)	0.173*** (0.011)	0.292*** (0.014)	0.218*** (0.030)
Observations	1,926	659	1,585	1,114	1,733	2,944	695
R-squared	0.55	0.56	0.43	0.60	0.45	0.53	0.60
Panel B: Price							
	Txt+leather	Paper	Chemicals	Minerals	Metals	Machinery	Vehicles
Δ TFP	-0.159*** (0.011)	-0.172*** (0.014)	-0.169*** (0.013)	-0.138*** (0.012)	-0.141*** (0.011)	-0.151*** (0.009)	-0.100*** (0.013)
$\Delta\xi$	0.140*** (0.005)	0.162*** (0.006)	0.155*** (0.005)	0.134*** (0.005)	0.142*** (0.003)	0.109*** (0.004)	0.087*** (0.007)
Observations	1,939	646	1,588	1,118	1,736	2,984	709
R-squared	0.73	0.92	0.82	0.79	0.84	0.66	0.62

Table A-3: Elasticities to demand and TFP shocks, by sector (continued)

Panel C: Employment

	Txt+leather	Paper	Chemicals	Minerals	Metals	Machinery	Vehicles
ΔTFP	0.076*** (0.023)	0.084** (0.036)	0.053* (0.031)	0.051 (0.032)	0.040 (0.029)	0.054*** (0.016)	0.030 (0.032)
$\Delta \xi$	0.113*** (0.009)	0.044*** (0.010)	0.069*** (0.011)	0.067*** (0.012)	0.045*** (0.007)	0.086*** (0.007)	0.097*** (0.012)
Observations	1,910	658	1,578	1,102	1,697	2,925	689
R-squared	0.17	0.09	0.10	0.11	0.07	0.13	0.19

Panel D: Investment

	Txt+leather	Paper	Chemicals	Minerals	Metals	Machinery	Vehicles
ΔTFP	0.094*** (0.033)	0.129*** (0.050)	0.028 (0.050)	0.047 (0.033)	0.112*** (0.037)	0.056** (0.028)	0.096 (0.062)
$\Delta \xi$	0.042*** (0.013)	0.018 (0.018)	0.021 (0.014)	0.031** (0.012)	0.037*** (0.012)	0.046*** (0.011)	0.002 (0.029)
Observations	1,590	503	1,232	878	1,465	2,313	482
R-squared	0.05	0.09	0.03	0.06	0.06	0.04	0.05

Notes: This table replicates sector by sector some of the pooled estimates displayed in Tables 5 and 6. All dependent variables and the demand and TFP shocks are in delta logs. ΔTFP is calculated using Olley and Pakes (1996) control function approach. $\Delta \xi$ is computed using self-reported sectoral price elasticities from the INVIND survey 1996. All specifications include region and year fixed effects. Both dependent and independent variables are trimmed to drop outliers above the 99th or below the 1st percentile for each sector. Standard errors are calculated from 500 bootstrap simulations. Robust standard errors are reported in parenthesis. Significance levels: *: 10%, **: 5%, *** : 1%