Competition, Innovation and Within-Plant Productivity: Evidence from Chilean Plants*

Ana Paula Cusolito
World Bank

Alvaro Garcia-Marin
Universidad de Chile

William F. Maloney
World Bank

April 2017

Abstract

This paper studies the effect of competition on within plant-productivity, and a principal channel through which it may work - innovation. We employ a unique plant-level panel from Chile that allows backing out price-cost markups, constructing direct measures of TFPQ, and tracking innovative activity in detail. We document a strong positive relationship between plant-level markups and within-plant physical productivity, and with innovation investment, with the latter accounting for about half of the productivity impact. These effects are only observed in lagging plants - productivity and innovation of plants at the frontier responds positively, although statistically insignificantly, to competition. Finally, we find that the link between markups and innovation is strong in industries that are more dependent on external financing, suggesting that part of the positive correlation is due to increased markups facilitating self-financing.

JEL: D24, L25, L60

Keywords: Physical productivity, Markups, Innovation, Competition, Chile

*We are grateful to Jan De Loecker, Marcela Eslava, Penny Goldberg, Omar Licandro, José Ignacio López, Chad Syverson, Jo Van Biesebroeck, and participants of the 2016 LACEA Conference and the OECD Summit on "Boosting Productivity in Latin America" for useful comments. Ignacio Nuñez provided superb research assistance. All errors are our own.

1 Introduction

This paper studies the effect of competition on within plant productivity, and in particular, the channel through innovation. It employs a unique 10-year panel data set on Chilean manufacturing plants that permits constructing measures of competition (markups), physical productivity (TFPQ), and investment in a range of innovation activities such as research and development, licenses, patents, and machinery and equipment. To our knowledge, our paper is one of very few papers to establish a tight link between markups and productivity, - and perhaps the first to establish the significance and quantitative importance of the innovation channel. It further explores the variance of the observed effects across plants' distance from the frontier and dependence on external finance.

The impact of competition on productivity is conceptually ambiguous and empirically unresolved (Aghion et al., 2005, see, for instance). The role of competition in reallocating factors of production among firms is long established (Collard-Wexler and De Loecker, 2015; Bloom et al., 2016).[1] Competition is thought to trim the longer left tails that are important to explaining variances in mean productivity (Syverson, 2004; Bloom and Reenen, 2007; Hsieh and Klenow, 2009). In particular, trade liberalizations and pro-market reforms have been documented to improve reallocation (Pavcnik, 2002; Melitz, 2003, among others).

However, much less is known on how competition affects within-plant efficiency, and in particular, through the incentives to innovate.[2] A long literature dating back to Arrow (1962) argues for a positive relationship, although one that becomes theoretically less clear given the market failures attending innovation. Aghion et al. (1997, 2005) on the one hand, argues that increased competition may lead incumbent plants to invest more in efficiency-enhancing activities in order to escape competition. On the other hand, competition reduces rents and hence the incentives to invest in technology and machinery. This view – posed originally by Schumpeter (1942) and incorporated later in a full-fledged model of endogenous growth by Aghion and Howitt (1992) – predicts a negative relation between competition and innovation. To date, there is little conclusive evidence on this channel (see Blundell et al., 1999, and the literature that follows). Aghion et al. (2005) argue that the relationship may in fact follow an inverted U-shape depending on firm distance from

[1] An early example of industry-level productivity improvements is provided by Syverson (2004). He uses the concrete industry to show that after an increase in competition caused by a larger market size, selection of plants got more severe, and this led to aggregate productivity gains. The trade literature also provides a large number of examples documenting aggregate productivity gains occurring through the selection mechanism (see Pavcnik, 2002, among other).

[2] Competition may affect productivity through other mechanisms as well. Matsa (2011) shows that the entrance of Wal-Mart to the U.S. supermarket industry led to improvement in the provision of quality, measured in terms of product availability. Van Reenen (2011) and Bloom and Reenen (2010) show evidence that competition improves plants' productivity, and suggest evidence that these improvements are a consequence of improvements in management practices.

the technological frontier. For firms close to the frontier (leaders), competition increases firms' incentives to innovate as a form of escaping competition (see Aghion et al., 2014, for a detailed exposition); for those further from the frontier (laggards), the reduced rents actually reduce innovation.[3] Bloom and Reenen (2010); Van Reenen (2013) find overall positive effect of competition on the adoption of better managerial technologies, although this result appears somewhat less robust for lagging countries (Maloney and Sarrias, 2014). This paper weighs in precisely on this debate.

Assessing the effect of competition on productivity throughout the innovation channel poses several empirical challenges. First, generating a reliable measure of market forces is challenging and researchers often use indirect proxies such as the Herfindahl index, import penetration or other self-reported measures of number of competitors which may not capture the actual competitive pressure that firms face. We follow De Loecker and Warzynski (2012) in backing out price-cost markups,[4] and other direct measures of price-cost margins for each product produced by plants.

Second, changes in price-cost markups may reflect factors other than competition (Holmes and Schmitz, 2010). For instance, they may change following a cost-saving technological shock unrelated to changes in the competitive environment. We closely follow the strategy of Bartik (1991), to exploit exogenous sources of variation unrelated to plants' own efficiency. In particular, we instrument plants' markups with (4-digit) industry average markups.[5]

A third empirical challenge relates to the measurement of productivity. Revenue-based productivity measures such as TFP mechanically reflect variation in markups and input prices (Katayama et al., 2009; Garcia-Marin and Voigtländer, 2013) and hence confounding price and efficiency changes arising from competition pressure.[6] The Chilean data allows us to construct plant-specific input and output deflators, which we use to compute measures of physical multi-factor productivity (TFPQ).

A fourth challenge is generating detailed information on innovation activities of firms. Often, innovation surveys are easily not undertaken or not easily linked to industrial surveys. For a subset of years and plants, we are able to merge the ENIA data with the Chilean Innovation Survey. This dataset contains detailed information on research and development expenditure and

[3]There is mixed empirical support for e inverted U-shape relationship between competition and innovation. While Aghion et al. (2005) provide supporting for this mechanism using U.K. patent count data, Hashmi (2013) replicate the empirical analysis for the U.S. and finds a *mildly negative* relationship between competition and innovation. Finally, Hashmi and Van Biesebroeck (2016) estimates a structural model for the automobile industry, finding supporting evidence for an inverted U-shape relationship between innovation and competition.

[4]This procedure is flexible with respect to the underlying price-setting model and the functional form of the production function.

[5]The implicit identification assumption of this approach is that not all variation in other plants' markups is caused by technological shocks common to all plants.

[6]Foster et al. (2008) pioneered this literature showing that revenue-based productivity understates the impact of efficiency on survival, because it mixes demand and technology components.

personnel, expenditure in licenses and patents, and investment in machinery and equipment, among other variables. To our knowledge, this matching of firm and innovation surveys is the first that permits quantifying the importance of the innovation channel in the mark-up-productivity relationship. Further, it also sheds light on the counterintuitive empirical findings that process innovation decreases productivity (Janz et al., 2004; Lööf et al., 2003; van Leeuwena and Klomp, 2006; Criscuolo, 2009), showing that the results may be explained by an identification problem when using measures of productivity like revenue TFP, as efficiency improvements from process innovation may not be reflected in total sales when they result in lower prices without corresponding counterbalancing increases in quantities (see Mohen and Hall, 2003; Hall, 2003, for a detailed discussion).

Our estimates suggest that variation in plant-level markups have a substantial effect on physical productivity. Moving a plant from the first to the third quartile of the markup distribution is associated with a 9% difference in physical productivity. Our results are robust to (i) alternative measures of physical productivity, and (ii) alternative measures of competition, such as the traditional Herfindahl-Hirschman Index (HHI), or the more recent Boone (2008) index. Taken together, these findings suggest that our main result is not an artifact of the methodology we use to calculate markups or TFPQ.

Analogous to the productivity findings, we show that plants with higher markups also invest more in research and development activities, purchase more external knowledge embedded in licenses and patents, and invest more in general purpose and innovative machinery. We then use estimates of the direct impact of innovation on TFPQ to assess how much innovation could account for the observed markup induced effects on TFPQ. Back-of-the-envelope calculations suggest that the competition-innovation mechanism could account for up to 50% of the direct relationship between markups and TFPQ.

A potential explanation for the markup-innovation relationship lies in the long gestation periods and lumpiness of innovation investments which make them especially sensitive to financing constraints (Aghion et al., 2012 Bond et al., 2003; Hall and Lerner, 2010; Mulkay et al., 2000). To test this, we show that industries scoring high on Rajan and Zingales' (1998) index of external financing show a stronger relationship between innovation and markups.

Finally, we investigate whether the positive effect of markups on productivity and innovation holds for both leader and laggard plants within each industry. We find that the positive effect of markups on R&D investment and productivity is only observed in laggards. Consistent with Aghion et al. (2005), leaders seem to react positively, although statistically insignificantly, to competition.

The rest of the paper is structured as follows. The section presents the empirical framework we use to calculate the plant-level measures of productivity and markups. Section 3 describes our dataset, and discuss preliminary evidence on the relationship between prices, productivity and markups. Section 4 establishes the positive relation between plant-level markups and physical productivity, and present several robustness checks on this relationship. Section 5 explores whether the innovation-competition mechanism could account for the positive relationship between markups and productivity. Finally, section 6 presents the main conclusions.

2 Empirical Framework

This section presents the empirical strategy for studying the relationship between markups and productivity. First, it discusses the different productivity measures used in the literature -including the commonly used revenue productivity (TFPR) and physical productivity (TFPQ) - and how they may fail to capture differences in productivity across plants. Then, it presents the estimation procedure used for computing plant-level markups and productivity.

2.1 Revenue vs. Physical Productivity

Productivity –the efficiency with which establishments convert inputs into outputs– is typically measured as the log-difference between output and the contribution of inputs. However, detailed data on physical inputs and outputs is generally unavailable; as a result, researchers traditionally rely on revenues to proxy for establishments' physical output. This traditional measure of multi-factor productivity is known in the literature as revenue total factor productivity (TFPR), to differentiate it from the real subject of interest, where inputs and outputs are measured in terms of physical units. This last measure of productivity is known as physical total factor productivity (TFPQ).

As previous research show, using TFPR to proxy for TFPQ is not innocuous.[7] The problem lies in the fact that revenues depend on demand conditions and the nature of competition, among other factors. For instance, if prices fall as plants become more efficient, revenue productivity would be a downward-biased measure of physical productivity.[8] Even more worrisome, under amenable conditions, TFPR and TFPQ are completely unrelated. To illustrate this point, assume a

[7]See Foster, Haltiwanger, and Syverson (2008); Katayama, Lu, and Tybout (2009), among others.

[8]Foster et al. (2008) documents a negative correlation between physical productivity and prices for a sample of U.S. manufacturing establishments producing relatively homogeneous products. In a different context, (Garcia-Marin and Voigtänder, 2016) show that revenue productivity is not helpful for finding export-related efficiency gains, because export entrants tend to reduce their prices when they enter into export markets. In empirical studies, this price bias is commonly tackled by deflating revenues with industry price indexes. However, the bias does not disappear *within* industries, and cross-sectional differences in TFPR will reflect the difference between individual plants' prices and the corresponding industry price index.

Cobb-Douglas production function with (approximately) constant returns to scale, and assume for the moment that output elasticities are correctly measured. Then, using the definition of revenue productivity ($TFPR = P \cdot A$), and the fact that optimal price is equal to marginal cost times (relative) markups (μ), it can be shown that (Katayama et al., 2009; Garcia-Marin and Voigtänder, 2016):

$$\Delta TFPR \;=\; \Delta\mu - \triangle\phi(\mathbf{w}) \tag{1}$$

where $\triangle$ denotes log differences. Equation (1) implies that TFPR does not reflect efficiency differences, unless more efficient plants charge higher markups or face lower input prices. Thus, using TFPR to study the effect of markups on productivity will most likely not yield meaningful results, because TFPR mechanically reflects markup dispersion under reasonable conditions: the identification of the effect of markups on productivity requires sources of exogenous variation.[9]

2.2 Productivity Estimation

To compute TFPR and TFPQ, we estimate separate Cobb-Douglas production functions for each 2-digit manufacturing sector (s), using labor (l), capital (k), and materials (m) as production inputs:[10]

$$q_{it} \;=\; \beta_l^s l_{it} + \beta_k^s k_{it} + \beta_m^s m_{it} + \omega_{it} + \varepsilon_{it} \tag{2}$$

where all lowercase variables are in logs; q_{it} corresponds to output of plant i in year t, l_{it} represents labor, k_{it} denotes capital stock, m_{it} are material inputs, ω_{it} is productivity, and ε_{it} represents measurement error as well as unanticipated shocks to output.

A key aspect in the estimation of (2) relates to the treatment of nominal variables. Our discussion in the previous section suggests that input and output variables ideally should be expressed in terms of quantities, or should be deflated with plant-level prices. In such a case, the estimated productivity term $\widehat{\omega}_{it}$ would correspond to TFPQ. However, detailed information on input and output prices often is unavailable, so that the computation of $\widehat{\omega}_{it}$ relies on the use of revenues and inputs expenditure, deflated with broadly defined industry-level price indexes. In these cases, the productivity term corresponds to TFPR.[11]

To estimate (2), we follow the methodology by Ackerberg et al. (2015, henceforth ACF), who

[9] We explain in more detail our identification strategy for computing the impact of markups on productivity in Section 4.1.

[10] The 2-digit product categories are: Food and Beverages, Textiles, Apparel, Wood, Paper, Chemicals, Plastic, Non-Metallic Manufactures, Basic and Fabricated Metals, and Machinery and Equipment.

[11] In the estimation of TFPR, we use 2-digit industry specific deflators.

extend the framework of Olley and Pakes (1996, henceforth OP) and Levinsohn and Petrin (2003, henceforth LP). This methodology controls for the simultaneity bias that arises because input demand and unobserved productivity are positively correlated.[12] As in Garcia-Marin and Voigtländer (2013), we modify the canonical ACF procedure by specifying an endogenous productivity process that can be affected by export status and plant investment. In addition, we include interactions between export status and investment in the productivity process. This reflects the corrections suggested by De Loecker (2013): if productivity gains from exporting also lead to more investment (and thus a higher capital stock), the standard method would overestimate the capital coefficient in the production function, and thus underestimate productivity (i.e., the residual). Accordingly, the law of motion for productivity is:

$$\omega_{it} \;\; = \;\; g(\omega_{it-1}, d^x_{it-1}, d^i_{it-1}, d^x_{it-1} \times d^i_{it-1}) + \xi_{it} \tag{3}$$

where d^x_{it} is an export dummy and d^i_{it} is a dummy for periods in which a plant invests in physical capital (following De Loecker, 2013).

In the first stage of the ACF routine, we compute a consistent estimate of expected output $\hat{\phi}_t(\cdot)$ using inverse material demand $h_t(\cdot)$ to proxy for unobserved productivity. Expected output is structurally represented by $\phi_t(\cdot) = \beta^s_l l_{it} + \beta^s_k k_{it} + \beta^s_m m_{it} + h_t(m_{it}, l_{it}, k_{it}, \boldsymbol{x}_{it})$.[13] The vector $\boldsymbol{x}_{it}$ contains other variables that affect material demand (time and product dummies, reflecting aggregate shocks and specific demand components). Next, we use the OLS estimate of expected output together with an initial guess for the coefficient vector $\boldsymbol{\beta}^s$ to compute productivity; for any candidate coefficient vector $\tilde{\boldsymbol{\beta}}^s$, productivity is given by $\omega_{it}(\tilde{\boldsymbol{\beta}}^s) = \hat{\phi}_t - \left(\tilde{\beta}^s_l l_{it} + \tilde{\beta}^s_k k_{it} + \tilde{\beta}^s_m m_{it}\right)$. Finally, for any given candidate vector $\tilde{\boldsymbol{\beta}}^s$, the productivity innovation ξ_{it} is recovered as the residual term from non-parametrically estimating Equation (3). The second stage of the ACF routine uses moment conditions on ξ_{it} to iterate over candidate vectors $\tilde{\boldsymbol{\beta}}^s$. In this stage, all coefficients of the production function are identified through GMM using the moment conditions

$$\mathbb{E}\left(\xi_{it}(\boldsymbol{\beta}^s)\mathbf{Z}_{it}\right) = 0 \tag{4}$$

where $\mathbf{Z}_{it}$ is a vector of variables that comprises lags of all the variables in the production function, as well as the current capital stock. Note that current capital stock is a valid instrument, because it is chosen before the productivity innovation is observed.

Given the estimated coefficients for each product category s (the vector $\boldsymbol{\beta}^s$), productivity $\hat{\omega}_{it}$

[12] We follow LP in using material inputs to control for the correlation between input levels and unobserved productivity.

[13] We approximate the function $\hat{\phi}_t(\cdot)$ with a full second-degree polynomial in capital, labor, and materials.

can be calculated at the plant level as , where q_{it} is plant's output (either revenues or physical output), and the term in parentheses represents the estimated contribution of the production factors to total output in plant i. Note that the estimated production function allows for returns to scale, so that the residual $\hat{\omega}_{it}$ is not affected by increasing or decreasing returns.

2.3 Markups Estimation

The methodology for deriving markups follows the production approach recently revisited by De Loecker and Warzynski (2012). This approach computes markups without relying on market-level demand information. The main assumptions are that at least one input is fully flexible and that plants minimize costs for each product j. The first-order condition of a plant's cost minimization problem with respect to the flexible input V can be rearranged to obtain the markup of plant i at time t:

$$\underbrace{\mu_{it}}_{\text{Markup}} \equiv \frac{P_{it}}{MC_{it}} = \underbrace{\left(\frac{\partial Q_{it}(\cdot)}{\partial V_{it}} \frac{V_{it}}{Q_{it}}\right)}_{\text{Output Elasticity}} \Big/ \underbrace{\left(\frac{P_{it}^{V} \cdot V_{it}}{P_{it} \cdot Q_{it}}\right)}_{\text{Expenditure Share}}, \tag{5}$$

where P (P^V) denotes the price of output Q (input V), and MC is marginal cost. According to Equation (5), the markup can be computed by dividing the output elasticity of product j (with respect to the flexible input) by the expenditure share of the flexible input (relative to the sales of product j). Note that, under perfect competition, the output elasticity equals the expenditure share, so that the markup is one (i.e., price equals marginal cost).

In the computation of (5), we use materials (M) as the flexible input to compute the output elasticity, based on our estimates of (2).[14] In our baseline estimation (due to its use of a Cobb-Douglas production function), the output elasticity with respect to material inputs is given by the constant term β_m^s, estimated with physical quantities for inputs and output in (2). The second component needed in (5) – the expenditure share for material inputs – is directly observed in the data. This procedure yields plant-year specific markups μ_{it}.

2.4 Plant-Level Price Indexes

Estimating TFPQ is demanding from a data perspective. First, outputs and inputs need to be observed in terms of physical quantities, and, in general, this information is not available in typical plant or firm-level datasets. Second, even when this information is observed, it requires aggregation of inputs and outputs at the plant level for plants producing multiple outputs or using multiple

[14] In principle, labor could be used as an alternative. However, in the case of Chile's regulated labor market, it would be a strong assumption that labor is a flexible input . A discussion of the evolution of job security and firing cost in Chile can be found in Montenegro and Pagés (2004).

inputs. This may be difficult if inputs and outputs are expressed in terms of different units, or correspond to different products.[15] One way to circumvent these issues is to construct plant-level input and output price indexes, and then deflate plants' sales and input expenditures to obtain inputs and outputs in physical units.[16] We follow this approach and compute Tornqvist price indexes for inputs and outputs.

The advantage of this type of index over other alternatives is that it gives a lower weight to outliers. We start by defining the log-change in plant-level prices Δp_{it} for plant i in period t as:

$$\Delta p_{it} = \sum_{v \in \Phi_v} \phi_{iv}(\ln P_{ivt} - \ln P_{iv,t-1}) \tag{6}$$

where Φ_v denotes the subset of outputs (inputs) sold (used) by plant i, and ϕ_v is the average share of input/ouput v between periods t and $t-1$. Once the price change is obtained, the level for the price index can be computed recursively:

$$\ln P_{it} = \ln P_{i,t-1} + \Delta p_{it} \tag{7}$$

A typical approach to computing price indexes is to normalize the price index in a given year for all plants. This method, however, makes it impossible to identify cross-sectional price differences in the base year (González and Miles-Touya, 2016). To avoid this complication, for the first year of each plant in the sample, we follow the following procedure.[17] First, for each output/input, we compute its log difference with the average industrial price for all plants with outputs/inputs in the same product category. Second, we aggregate outputs/inputs using the shares ϕ_v, but now computed for the initial period. Thus, given the level for the initial plant-level price, and using the recursion (7), the complete series for plant-level price indexes can be recovered.

3 Data

3.1 The Chilean Annual Industrial Survey (ENIA)

The main dataset we use in this paper is the *Encuesta Nacional Industrial Anual* (Annual National Industrial Survey – ENIA) for the years 1996–2007. Data for ENIA are collected annually by the

[15]One special case where these issues are not a problem is for single-product plants producing with a single input. However, this subset of plants represents a small fraction of the universe of plant-year observations –less than 5% of the total.

[16]A shortcoming of this more aggregate approach is that plant-level output price indexes may not account for differences in product scope Hottman, Redding, and Weinstein (2016).

[17]This procedure follows the literature computing Tornqvist TFP indexes, see Aw, Chen, and Roberts (2001).

Chilean National Institute of Statistics (INE), with direct participation of Chilean manufacturing plants. ENIA covers the universe of manufacturing plants with 10 or more workers, and contains detailed information on plant characteristics, such as sales, spending on inputs, employment, wages, investment, and export status. This survey provides information for approximately 4,800 manufacturing plants per year with positive sales and employment information. Out of these, about 20% are exporters, and two-thirds are small (less than 50 workers). Medium-sized plants (50-150 workers) and large plants (more than 150 workers) represent 20 and 12 percent, respectively. In addition to aggregate plant data, ENIA provides rich information for every good produced by each plant, reporting the value of sales, the total variable cost of production, and the number of units produced and sold. We use this information to construct plant-level price indexes (see Section 2.4 for details). Products in ENIA are defined according to the *Clasificador Unico de Productos* (CUP). This ENIA-specific product category is comparable to the 7-digit ISIC code.[18]

3.2 Sample Selection

In this subsection, we explain the procedure we follow to arrive at the dataset used in the main empirical analysis. We follow three steps. First, we exclude plant-product-year observations that have zero values for total employment, demand for raw materials, sales, or product quantities. Second, we exclude plants with missing observations for input and output prices between its first and last year of operation in the sample. This is important, because, as we discuss above, the computation of our main productivity variables requires information for plant-level input and output price indexes. Such indexes relies on price changes between consecutive periods; consequently, they cannot be computed if (i) a plant changes its output or input mix in two consecutive periods, or (ii) if it has missing information for a given year. For an important number of plants – about one-third of the overall sample – we are unable to construct input price indexes for all years. Thus, to preserve representativeness in the results, for the baseline analysis we use physical productivity with output in terms of physical output, but inputs in terms of deflated material expenditure. We show that this does not modify our main conclusions, because the correlation between both variables is very high – almost 90%.[19] Finally, to avoid the possibility that outliers and/or misreported prices affect our results, we exclude observations where the input or output price deviates by more than five times the average. After these adjustments, our sample consists of 31,678 plant-product-year observations.

[18] For example, the wine industry (ISIC 3132) is disaggregated by CUP into 8 different categories, such as "Sparkling wine of fresh grapes", "Cider", "Chicha", and "Mosto".

[19] The reason for this is that in the data input prices are not systematically related to either markups or TFPQ. Consequently, the sign - and, to a first order approximation, the magnitude – of our main results are not affected by the choice of productivity measure.

3.3 Summary Statistics

Before turning to our main results, we present summary statistics for the main variables of interest in Table 1.[20] In this table, all variables are measured in logarithms, and demeaned with respect to the respective 2-digit sector-year averages.[21] Several interesting patterns emerge from this table. First, revenue and physical productivity are highly correlated – although to a lesser extent than for the case of the United States (see Foster et al., 2008). One reason for the discrepancy is that, for the case of United States, Foster et al. (2008) focus on a sample of relatively homogeneous products, while our sample includes the universe of the Chilean manufacturing . Because price differences are less important in relatively homogeneous products, TFPR is expected to be a closer approximation of TFPQ in their sample than in ours.

Second, physical productivity displays a higher dispersion than revenue productivity. The last row of Table 1 reveals that the standard deviations of TFPQ almost doubles the standard deviation of TFPR. One explanation is that, as reported by Foster et al. (2008) for the United States, physical productivity is negatively correlated with output prices (first column, third row in Table 1): relatively more efficient establishments charge lower prices to customers.[22]

Third, revenue productivity and markups show a high correlation within sector-years, with a correlation coefficient of 0.82. Figure 2 plots both variables – both in logarithms and demeaned, as in 1. As the figure shows, most plant-year observations lie close to the 45-degree line, suggesting that TFPR moves almost 1-to-1 with changes in markups. The high correlation between markups and TFPR should not come as a surprise in light of the empirical framework presented in Section 2.1. Indeed, the fact that our sample shows approximately constant returns to scale (see the appendix for the estimated elasticities and the related returns to scale) implies that Equation (1) applies to our data. This decomposition, in combination with the high degree of co-movement between TFPR and markups, implies that variation in input prices most likely explain a small fraction of the variation in TFPR. This is corroborated by a low correlation between input prices and revenue productivity (second column, fourth row of Table 1).

Finally, markups and TFPQ show a positive correlation of about 42%. The fact that this correlation is of the same order of magnitude as the correlation between TFPQ and TFPR is not a

[20] In the online appendix, we show detailed tables for the estimated elasticities for the revenue and physical production functions, as well as markups, for each 2-digit sector.

[21] In the computation of the averages, we exclude outliers by trimming the top and bottom 2 percentiles of each variable's demeaned distribution.

[22] Several authors in the trade context have also documented the negative correlation between physical productivity and prices. Garcia-Marin and Voigtländer (2013) show that export-related efficiency gains are substantially larger when measured in terms of measures that are not affected by prices – such as marginal cost and physical productivity - than when measured in terms of revenue productivity. In different contexts, this is confirmed by Smeets and Warzynski (2013) and Eslava et al. (2013).

coincidence, because, as we discuss above, TFPR mostly reflects variation in markups. In principle, this correlation reflects two forces. On the one hand, more efficient establishments tend to capture a larger market share, which may translate into higher markups under certain demand structures (see Melitz and Ottaviano, 2008). On the other, it may be that larger markups allow establishments to invest more in efficiency-enhancing activities to escape from competition. In the next section, we present our strategy for disentangling both possibilities, and show that, even after controlling for the mechanical connection between markups and TFPQ, exogenous variation in markups are related to efficiency differences in terms of physical productivity.

4 The effect of Markups on Physical Productivity

In this section, we present our main empirical results related to the impact of markup changes on physical productivity. We begin by presenting our empirical strategy and baseline results. We end the section by discussing several robustness checks.

4.1 Empirical Strategy and Baseline Results

Identifying the productivity effect of markups is not straightforward, because of the possibility of reverse causality. As is common in theoretical models with heterogeneous productivity and variable markups (e.g. Eckel and Neary, 2010; Melitz and Ottaviano, 2008), more productive plants tend to exploit their productivity advantage by charging higher markups. Thus, a simple OLS regression between markups and physical productivity will fail to identify whether markups lead to differences in physical productivity. In the following, we attempt to isolate the variation in markups that is not driven by changes in plants' efficiency to assess the impact of competition on productivity.

Our baseline specification establishes, for each plant i operating in industry s at time t, a log-linear relationship between physical productivity, lagged markups and other plant level controls:

$$\ln TFPQ_{ist} \quad = \quad \delta_{st} + \beta_2 \ln(\mu_{is,t-1}) + \gamma_2 X_{ist} + \vartheta_{ist} \tag{8}$$

where $\mu_{is,t-1}$ are lagged markups, X_{ist} are controls varying across time and plants, δ_{st} are (4-digit) sector-year fixed effects, and ϑ_{ist} is a disturbance term. The vector X_{ist} includes log employment (to control for plant size) and initial plant-level TFPQ (to control for within-plant changes in TFPQ). Note that, in (8), log-markups enter the equation lagged one period. There are two reasons for this. First, the underlying theory (Aghion and Howitt, 1992) suggests that competition affects physical efficiency by changing plants' incentives to invest in efficiency-enhancing machinery and equipment. By lagging markups one period, we allow for a delay between the change in markups and the time at which the higher margins translate into greater purchases/development of technology, resulting in an increase in productivity. Second, by using lagged instead of contemporaneous markups, we minimize the possibility of capturing the action of third factors affecting both markups and productivity.

As we discuss above, plant-level markups are likely to be correlated with the error term in (8). Therefore, estimating this equation by OLS would most likely yield inconsistent estimates. To alleviate endogeneity concerns, we instrument plant-level markups with the average markup charged by all other plants operating in the same (4-digit) industry. More concretely, for each plant i operating in industry s at time t, we define our instrument as:

$$\overline{\mu}_{-ist} \;=\; \sum_{j\neq -i} \frac{1}{(I-1)}\mu_{jst} \tag{9}$$

Note that (9) excludes plant i's markups from the computation of the average industry-level markup.[23] Thus, the instrument captures markup variation originated by factors affecting plant i's competitors.

In the first stage, we predict plant-level log markups based on the average markup of other plants, sector-year fixed effects, and controls:

$$\ln(\mu_{is,t-1}) = \alpha_{st} + \beta_1 \ln(\overline{\mu}_{-is,t-1}) + \gamma_1 X_{ist} + \varepsilon_{ist} \tag{10}$$

Equation (10) exploits within-industries-year variation, using sector-year fixed effects. Correspondingly, all standard errors are clustered at the 4-digit sector-year level.

Column 1 in Table 2 presents our first-stage results. As this column shows, there is strong evidence supporting the use of industry-level markups as instruments for plant-level markups. The first stage F-statistic is significantly above the critical value of 16.4 for 10% maximal IV bias, and the coefficient on industry-level markups is significant at the 1% level. Notice that the coefficient on industry level markups is well below 1, implying that industry-level markup shocks are only partially passed through to plant's markups. Indeed, the magnitude of the first-stage coefficient implies that only 44% of the variation in the average markup of other plants in the same industry are transmitted to plant i's markup.

Next, we proceed with the second stage, where we regress log TFPQ on predicted markups

[23] Note the resemblance of the instrument to the family of instruments based on Bartik (1991). Unlike that family, in our case the cross-sectional variation of the instrument originates from the exclusion of firm i from the calculation of the average, while in Bartik (1991) the cross-sectional variation comes from differences in the initial employment share.

($\ln\widehat{(\mu_{is,t-1})}$) and other controls:

$$\ln TFPQ_{ist} = \delta_{st} + \beta_2 \ln\widehat{(\mu_{is,t-1})} + \gamma_2 X_{ist} + \vartheta_{ist} \tag{11}$$

Column 2 in Table 2 reports the second-stage results for TFPQ. The estimated coefficient is positive and statistically significant even at the 1% level (we report weak-IV robust Anderson-Rubin p-values in square brackets, based on Andrews and Stock (2005). Interestingly, Column 3 in Table 2 reveals that the 2SLS coefficient is almost three times the OLS coefficient. This provides support to our strategy of using exogenous markup variation to assess the TFPQ-markups relationship. In quantitative terms, the coefficient implies that physical productivity increases by about 2% after an industry-driven markup increase of 10%. In other words, our estimates suggest that moving a plant from the first to the third quartile of the markup distribution leads to a 9% difference in physical productivity.

Finally, Column 4 shows reduced form estimates from regressing physical productivity directly on industry-level markups. Again, the coefficient is positive and significant at the 1% level, delivering a similar pattern as the 2SLS estimates. In sum, the results presented in Table 2 provide compelling evidence that markups lead to efficiency differences.

4.2 Robustness and Additional Results

We perform a number of robustness tests, adding covariates and considering a series of extensions. In this subsection, we discuss the most important of them. In some cases, we summarize the results without providing detailed tables; many of these are provided in the appendix and the others are available on request.

We first show that our results are very similar when accounting for plant-level input price differences in physical productivity. As we discussed in Section 3, our baseline measure of physical productivity does not exploit differences in plant-level input prices, because input prices are not available for about one-third of the sample. Because both measures of physical productivity display a high correlation – about 0.90 – we expect results to be very similar with both productivity measures. In the first column of Table 3, we show the 2SLS estimate for lagged log-markups. As in our baseline case, the coefficient is positive and significant at the 5% level. However, the coefficient is about 35% lower than in Table 2 (0.124 vs. 0.189). To evaluate whether the difference in coefficients is due to the difference in the sample used in the two tables, we re-estimate Equation (11) with our benchmark TFPQ measure (which does not consider differences in input prices) but with the sample of Table 3. This exercise yields a 2SLS of 0.108, with a p-value of about 7%. Thus, we conclude that the choice of our baseline physical productivity measure does not drive our results - i.e., that the differences in coefficient and statistical significance in Tables 2 and 3 are only to a minor extent due to the different TFPQ definition.

Next, we evaluate whether the presence of multi-product plants affects our results. Note that the methodology we follow for deriving plant-level markups (Equation 5) depends on the estimated output elasticity of material inputs. For multi-product plants, the relevant materials' elasticity may differ from the value assigned to the plant, especially if the plant produces products from different product categories. Note that, in the sample of single product plants, this is not a concern (because there is a single relevant elasticity). In the third column of Table 3, we show our results for this sample of plants. Note that the sample size is significantly reduced – the regression is run with only one-third of the observations considered by our baseline case. However, the coefficient stays positive and highly significant, and the magnitude is close to our benchmark estimates.

Then, we investigate to what extent our measure of competition drives our results. We address this concern by analyzing the sensitivity of our results to different measures of competition. First, recall from the date section that ENIA provides information for the average variable cost and unit value of each plant-product. We divide unit value and average variable cost to construct a plant-product level measure of margin, and then average the margin of all outputs produced by the plant to obtain a plant-level measure of average variable margin. Column 2 in Table 3 shows that that the 2SLS coefficient stays positive and statistically significant, as in our baseline case. The point estimate for the case of average variable margin is significantly larger than our baseline estimates – almost three times larger. However, the economic impact in both cases is closer than what is suggested by a simple comparison of point estimates, because the measure of average variable margin displays about half of the variation of markups.

Finally, we construct alternative measures of competition. We begin by computing the popular Herfindahl-Hirschman index (HHI). The index reflects the concentration of market shares in the industry, and is defined as the sum of the squares of the market shares of the firms in the (4-digit) industry. Next, we compute the recently developed Boone (2008) index. The underlying idea behind this index is that, in more competitive industries, there should be a closer relationship between profits and efficiency. In practice, the index is computed as the elasticity of plants' profits to marginal or average costs. An important caveat of the HHI and Boone indexes is that they can only be computed at the industry-year level. Consequently, when using these indicators to proxy for competition, we cannot use the full set of industry-year fixed effects, and instead use industry and year fixed effects separately. When using these indexes to proxy for competition, we find a positive effect of competition on physical productivity, although the effect is significant only at the 10% level in both cases.

5 Markups and Investment in Efficiency-Enhancing Technology

Why do higher markups improve plants' efficiency? In this section, we address this question through the optic of Shumpeterian theory, and analyze whether competition – measured in terms of plants' markups – affects plants' incentives to invest in technology and research and development activities. In theory, more stringent competition would reduce establishments' incentives to innovate, because the rents related to a successful innovation are reduced in markets that are more competitive. In the following, we seek evidence of this mechanism in our sample of Chilean plants.

5.1 Measures of Innovation and Technological Investment in the Chilean Data

To investigate the effect of competition on technological investment, we merge our baseline data with an auxiliary dataset, the Chilean Technological Innovation Survey (EIT). This dataset is a nationally representative survey of Chilean establishments, conducted by the Chilean National Statistical Agency. There are two important differences between ENIA and EIT. First, the innovation survey is administered every two or three years, while ENIA is conducted on an annual basis. Second, unlike ENIA, EIT samples establishments randomly every time the survey is administered.[24] As a consequence, most establishments in EIT are surveyed only once, which impedes efforts to exploit panel variation.

In the results that follow, we use information from the third, fourth, fifth and sixth waves of EIT. Although EIT is administered every two or three3 years, the survey asks plants to report disaggregated R&D expenditure information for the years in between two consecutive surveys. Thus, in practice the innovation information is available for all years covered by ENIA between 2000-2007, with the exception of 2002.[25] The combined ENIA-EIT dataset consists of 3,900 plant-year observations, which corresponds to about 20% of the ENIA sample for the years for which EIT is available.

Before turning to our main analysis, we investigate whether plants surveyed in EIT are systematically different than the rest of the plants surveyed in ENIA. Because EIT only includes a portion of the establishments in ENIA, this exercise is important to understand how representative are the results obtained with the EIT-ENIA sample in terms of the full dataset. For this purpose, we run a simple regression, where each variable of interest is run against a dummy variable that takes the value one if the plant is included in EIT (and zero otherwise), with sector-year fixed effects included. The coefficient accompanying the EIT dummy variable is interpreted as the percentage-

[24] However, all entities representing more than 2% of manufacturing added value enter compulsory in the innovation survey.

[25] We do not use earlier waves of the survey because they do not allow us to establish a one-to-one correspondence with the questions from the most recent versions of EIT.

point difference between plants in EIT and plants not included in EIT. The results of this exercise are reported in Table 4.[26] As the table shows, plants in EIT are systematically different in many dimensions. First, plants in EIT are significantly larger; they hire more workers (Column 1), and have higher sales (Column 2) than the establishments not included in EIT. Second, they are more likely to be exporters – about 20% more likely, as shown in Column 3. Third, they have 7% higher physical productivity, as shown in Column 4. However, markups do not seem to be systematically different for plants surveyed in EIT (Column 5). In sum, the subsample we use to analyze the effect of competition on efficiency-enhancing investment displays important differences in key dimensions. Thus, in the analysis that follows, we are careful to point out how and when the sample differences could be driving our results.

5.2 Markups and Technological Investment

In this subsection, we study whether or not higher markups lead plants to increase their technological investment. To answer this question, we apply the same 2SLS strategy presented in Section 4 for a series of indicators of technological investment.[27] In particular, we consider the following variables from EIT: overall and in-house R&D expenditure, investment in machinery and equipment intended for innovative activities, and expenditure on patents and licenses. We complement these variables with information on overall investment in machinery and equipment from ENIA.

We begin the analysis by corroborating whether our main result on the positive relationship between markup and TFPQ holds in the reduced EIT-ENIA sample. We present the result of this exercise in Table 5. As in our baseline case, the first stage shows a high F-statistic, and the 2SLS coefficient is positive and highly significant.[28] Next, we analyze the effect of markups on innovative activities and general investment in machinery and equipment. In Columns 2–6, each dependent variable x is expressed according to the formula $\ln(1 + x)$. Note that this formulation explicitly accounts for zeros; therefore, the estimated coefficients reflect both extensive margins (whether the plant invested in R&D) and intensive margins (how much they invested). The bottom panel of Table 5 focuses on the extensive margin only, expressing all dependent variables as categorical variables that take the value one only for strictly positive observations. Thus, the coefficients in this panel can be directly interpreted as changes in the probability that a plant engaged

[26] As in the rest of the paper, we cluster standard errors at the industry-year level.

[27] There is only one difference: in the light of previous results on the nature of R&D expenditure, we control for the logarithm of plants' sales instead of log employment.

[28] Note that the magnitude of the estimated coefficient is almost three times the coefficient of our base-case estimates in Table 2. One reason for this difference could be the relatively higher participation of larger plants in the reduced sample. The larger coefficient in the EIT sample suggests that large plants benefit more than small plants from high markups in terms of TFPQ.

in technological investment after doubling plant markups.

As can be seen, regardless of the variable under consideration, the impact of markups is positively related to the probability and amount of investment in efficiency-enhancing technologies and equipment. In the upper panel, all coefficients but one are statistically significant at the 1% level; the exception (investment in innovative machinery) is statistically significant at the 5% level. The estimated magnitudes in this table are relatively large. Taken at their face value, these coefficients imply that moving a plant from the 25th to the 75th percentile of the markup distribution would result in differences in technological investment ranging from 105%, for investment in innovative machinery and equipment, to over 200% for overall R&D investment.[29]

The bottom panel of Table 5 suggests that part of the overall effect of markups on the amount invested in technological activities operates through an increase in the probability of engaging in technological activities. The coefficients in this panel are somewhat weaker than in the upper panel, where intensive and extensive margins are considered. For instance, investment in machinery and equipment is non-significant for the case of general-purpose machinery. However, as in the upper panel, the implied economic effect of the coefficients is relatively large. Indeed, our estimates suggest that doubling plant level markups would increase a plant's probability of making a technological investment by between 14 percent – for the case of innovative machinery – to about 30% for the case of overall R&D investment.

5.3 Economic significance

So far, we have documented two main findings: (i) exogenously driven markup differences are related to differences in physical productivity, and (ii) they are related to differences in innovative effort and technological investment. This suggest that at least part of the positive relationship between markups and TFPQ could be explained by the increased incentives to invest in efficiency-enhancing activities. In this section, we estimate the quantitative impact of technological investment activities on physical productivity. Then, we use these estimates to provide back-of-the-envelope calculations on how much of the direct relationship between markup and TFPQ can be accounted for by the technological investment channel.

We begin by presenting estimates for the impact of efficiency-enhancing investment on physical productivity, in Table 6. The upper panel presents OLS estimates; each column in this table assesses the impact of each technological investment variable on physical productivity. Correspondingly, in all columns, we regress the logarithm of TFPQ on each variable of technological investment, and control for plant's employment, initial physical productivity, and sector-year fixed

[29]These magnitudes are equivalent to one-fifth of a standard deviation, in the case of investment in innovative machinery and equipment, and one-half of standard deviation, in the case of overall R&D investment.

effects. As Table 6 shows, for all variables under consideration, least squares estimates show a positive and statistically significant relationship between physical productivity and technological innovation.[30] Then, we move to 2SLS estimates to address the possible endogeneity of technological investment (bottom panel of Table 6). Our IV strategy resembles the approach we follow for the case of the relationship between markups and TFPQ. In particular, we consider two sets of instruments for plants' technological investment. Our first instrument corresponds to the investment undertaken by other plants in the industry. In addition, we consider a series of indicators on the main obstacles that plants face for developing innovations.[31]

Before proceeding, a word of caution is due for the interpretation of these coefficients. Although in all specifications the instruments are statistically significant, the first stage F-statistic points toward the presence of weak instruments in all specifications. Thus, as in previous 2SLS regressions, to assess the statistical significance of the coefficients, we rely on weak-IV robust Anderson-Rubin p-values (see Andrews and Stock, 2005). As in the case of least squares estimation, all coefficients are positive and statistically significant. The magnitude of our 2SLS estimates is about twice the size of the least squares coefficient presented in the upper panel. Taken together, these findings suggest that all components of technological investment are positively related to physical productivity.

Next, we illustrate the importance of the technological innovation channel, using our estimates in Tables 5 and 6. The numbers we provide next should be considered as exploratory evidence. A more definitive evaluation would require a fully specified model, estimated to explicitly account for the cross-correlation of the error term of the different equations. In spite of this caveat, we believe the back-of-the-envelope calculation that follows is important, because it provides a first approximation on the overall importance of the innovation mechanism.

In order to provide comparable estimates with our previous results and discussion, we repeat the exercise of moving a plant from the 25th to the 75th percentile of the markup distribution. Recall that the discussion of results in Section 4 revealed a direct economic impact of markups of 9% in terms of TFPQ. However, this is not a valid benchmark to evaluate the importance of the innovation channel, because our estimates for technological investment are based on the reduced EIT-ENIA sample. In this sample, the direct impact of markups on TFPQ is about two times larger

[30] In the appendix, we provide results with these variables entering jointly the regressions. When we include overall R&D expenditure – which comprises in-house R&D, expenditure in licenses and patents, and investment in innovative equipment – together with plants' general investment in machinery and equipment, both coefficients remain highly significant, and their magnitudes only fall in about 10% with respect to the case with separate regressions. In addition, when we include all disaggregated innovative variable together with general investment in machinery and equipment, all variables stay positive and statistically significant.

[31] We use as instrument the following obstacles: `"lack of incentives as an obstacle for innovation"`, and `"innovation too easy to be imitated"`.

– 29% in terms of TFPQ. Next, we move to evaluating the indirect impact that occurs through technological innovation. For this purpose, we first compute the additional investment of each technological component related to a markup change that is equivalent to moving a plant from the 25th to the 75th percentile of the markup distribution. Finally, we multiply the resulting numbers by the impact of each component on physical productivity according to our estimates in Table 6. This exercise give us the additional TFPQ that would result through incremental investment in each of the technological components. To compute the overall impact due to the technological channel, we sum the individual effect occurring in each component. The result of this exercise reveals that markup variation leads to a 14% variation in physical productivity. This is equivalent to about half of the direct effect of markups on TFPQ (14% out of the 29% overall effect). Thus, the technological investment accounts for a large part of the effect of markups on TFPQ.

5.4 Additional Results

5.4.1 Heterogeneity: Leaders and laggards

Tables 2 and 5 provide strong support to the idea that competition hurts productivity by reducing incentives to invest in efficiency-enhancing technology. In this subsection, we explore whether this result holds for different types of establishments. A recent literature following Aghion et al. (2005) suggests that the relationship between competition and innovation depends on plants' productivity. These authors argue that the negative effect of competition would predominantly be observed in plants that are far from the technological frontier (laggards). In contrast, competition would have the opposite effect in highly productive plants. For them, competition provides a means to escape from innovation. Therefore, in more competitive environments, frontier plants are predicted to increase their innovative effort, while the negative effect of competition on innovation would tend to be concentrated in lagging establishments.

To investigate whether the effect of markups on productivity and innovation depends on plants' distance to the technological frontier, we construct a variable that measures the distance to the technological frontier as:

$$\ln(TFPQ_{ist}^{GAP}) \quad = \quad \max_{v \in I}\{\ln(TFPQ_{vst})\} - \ln(TFPQ_{ist}) \tag{12}$$

where the $\max$ operator runs over all plants in the same sector and year as plant i. Then, we modify our baseline specifications for TFPQ and overall R&D expenditure by adding the TFPQ-gap variable and its interaction with lagged log-markups. According to the underlying theory, we expect the interacted term to be positive – reflecting a more negative impact of competition for plants farther away from the technological frontier – while we expect the baseline coefficient on log-markups to be negative or non-significant – reflecting a non-negative impact of competition on frontier plants.

Table 7 shows the results for leading and lagging plants. For these results, we trim the upper and bottom percentiles of the productivity distribution to avoid the possibility that extreme productivity outliers will affect the frontiers' definition. We present 2SLS results for TFPQ (Columns 1-2) and overall R&D expenditure (Columns 3-4). For each case, we first present the baseline coefficient with no interaction terms for comparison, and then the specification including the productivity gap and its interaction with markups. We begin by confirming that our results for the case of TFPQ hold in this sample. This is not very surprising, since the sample is almost identical to the sample used in Table 6. Next, Column 2 adds the interaction between markups and TFPQ. Note that the first stage F-statistic drops drastically when adding the interactive terms; however, it stays well above the Stock-Yogo value for 10% maximal IV bias. Regarding the estimated coefficients, the table reveals that the negative effect of competition on TFPQ concentrates in laggards. Indeed, the point estimates suggest a positive impact of competition for the 5-10% of plant-years that are closer to the technological frontier. However, these impacts are not statistically significant. Indeed, even for the 50% of plant-years that are closer to the technological frontier, the effect of markups on TFPQ is not statistically significant. In contrast, for the remaining 50% of observations with relatively low productivity, the negative effect of competition on productivity is statistically significant. Even more noticeable, for the bottom part of the productivity distribution, the negative effect of competition is significantly larger than the average effect reported in Column 1.

Then, we move to the case of overall R&D expenditure. Here, we focus on analyzing the interactive terms (Column 4). As in the case of TFPQ, the estimated coefficients reveal a non-significant effect of competition for plants close to the technological frontier – even though the point estimates suggest a positive effect. In contrast, competition has a statistically negative effect on the R&D expenditure of lagging plants, and the negative effect appears more pronounced than in the case of TFPQ. In fact, the effect of competition is negative even for establishments located in the upper 10th percentile of the productivity distribution.

In sum, results in Table 7 provide plant-level supporting evidence for the mechanism proposed in Aghion et al. (2005). Plants close to the frontier do not seem to be affected by competition; in contrast, laggards react to competition by reducing their innovative effort, and, ultimately, their productivity. These results have important policy implications. Pro-competitive policies, by reducing innovation incentives in lagging plants, may effectively increase the future market power of leading businesses. By making it more likely that leaders will invest in efficiency-enhancing technology, and less likely that laggards will do so, competition may contribute to widening the productivity gap between leaders and laggards, which may ultimately translate into an increase of the market power of leaders.

5.4.2 Why markups do matter for innovation?

We have determined that markups matter for innovation and productivity, but we have not yet discussed why competition could lead plants to reduce their investment in efficiency-enhancing activities. In this section, we argue, following a large theoretical and empirical literature, that the existence of a 'funding gap' for R&D activities may explain the negative effect of competition on innovation (see Hall, 2002; Hall and Lerner, 2010, for reviews). The argument goes as follows. The non-rival nature of newly created knowledge makes it difficult for researchers to fully appropriate the return of R&D activities. Therefore, private external investors may be unwilling to finance this type of activity, making businesses dependent on internal resources to fund innovative activities. Thus, increasing plants' markups would partially solve the financing problem, by providing incremental sources of internal funds for investing in efficiency-enhancing activities.

If the funding gap were at least partially behind explain our results, then we would expect a stronger relationship between markups and technological investment in industries more dependent on external funds. To implement this test, we use the well-known Rajan and Zingales (1998) index to split the sample according to the degree of dependence of external financing of each industry. Table 8 shows the results of this exercise. The upper panel shows results for 4-digit industries with above-median financing needs, while the panel repeats the exercise for industries for industries with below-median financing needs. Reassuringly, all estimates are positive as in the baseline case with the full sample (Table 6), and most of the coefficients are statistically significant. In terms of the importance of external financing, a simple comparison on the coefficients reveals that all components of technological investment but one – general investment in machinery and equipment – are larger in industries more dependent on external financing (upper panel). Standard errors are relatively large because of the reduced size of the subsamples, resulting in non-significant difference for most coefficients. However, it is worth to be mentioned that the one component of technological investment that is significantly different – in-house R&D expenditure – is statistically higher in industries with high dependence on external financing even at the 1% level.

6 Concluding Remarks

A large literature has studied the effect of competition on productivity. The theory – originally developed in Schumpeter (1942), and more recently revisited by Aghion and Howitt (1992) – points toward a negative effect of competition on plants' productivity. Competition, by reducing the rents related to successful innovation, weakens businesses' incentives to engage in innovative activities. So far, empirical evidence on this mechanism is mixed, and focuses on particular links of the full competition-innovation-productivity relationship, studying either the effect of competition on innovation, or the direct effect of competition on productivity. In this paper, we provide evidence for the full competition-innovation-productivity mechanism, using a rich dataset of Chilean plants. We show that competition negatively affects plants' productivity. Then, to empirically assess the innovation mechanism, we analyze the effect of competition on innovation and efficiency-enhancing activities. Next, we estimate the return to these activities in terms of productivity. By studying each element of the competition-innovation-productivity relationship, we are able to compute back-of-the envelope calculations on the empirical relevance of the innovation mechanism for the reduced-form relationship between competition and productivity. Our findings suggest that the innovation mechanism is quantitatively important, explaining up to 50% of the competition-productivity relationship.

To gauge the quantitative importance of our findings in terms of aggregate productivity, we first note that our findings abstract from the positive cross-sectional effects of competition on productivity; extensive evidence shows that increased competition improves aggregate productivity, inducing reallocation of resources from the least to the most productive businesses. Estimates from the trade literature suggest that the productivity gains from the reallocation mechanism are quite high; Pavcnik (2002) show that, after trade liberalization, about two-thirds of the aggregate productivity gains can be related to reallocation of resources. Our evidence suggests that, in the absence of the negative effect of competition on innovation, within-plant efficiency gains related to competition would be substantially larger.

Together, our findings induce subtlety into the discussion of competition and productivity policies for developing countries which, almost by definition, have a larger share of their firms further from the technological frontier. Simply increasing competition without supportive policies for innovation may, indeed, trim the low productivity left tails. But without supportive policies for innovation, it will also dampen innovative activity by potentially more viable firms.

It also raises some questions of the mechanisms through which policies, such as trade liberalization affect productivity. For India, De Loecker et al. (2016), find that competition through tariff reductions disproportionately lower input costs, leading to a rise in markups, at least over the medium term. Combined with our results here, this may suggest that trade liberalization increases innovation in productivity through increasing markups and hence the resources available for investment.

Finally, our results suggest that the negative effect of competition concentrates on businesses far from the technological frontier. This has important implication for pro-competitive policies. By reducing lagging plants' incentives to invest in efficiency-enhancing activities, pro-competitive policies may have a negative effect on competition in the medium term. The lack of investment by laggards distances them farther from leaders, undermining their competitive position.

References

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

Aghion, P., U. Akcigit, and P. Howitt (2014, December). What Do We Learn From Schumpeterian Growth Theory? In *Handbook of Economic Growth*, Volume 2 of *Handbook of Economic Growth*, Chapter 0, pp. 515–563. Elsevier.

Aghion, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt (2005). Competition and Innovation: an Inverted-U Relationship. *The Quarterly Journal of Economics 120*(2), 701–728.

Aghion, P., C. Harris, and J. Vickers (1997). Competition and growth with step-by-step innovation: An example. *European Economic Review 41*(3-5), 771–782.

Aghion, P. and P. Howitt (1992, March). A Model of Growth through Creative Destruction. *Econometrica 60*(2), 323–351.

Andrews, D. W. K. and J. H. Stock (2005). Inference with Weak Instruments. NBER Working Paper 313.

Arrow, K. (1962). Economic welfare and the allocation of resources for invention. In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, pp. 609–626. National Bureau of Economic Research, Inc.

Aw, B. Y., X. Chen, and M. J. Roberts (2001). Firm-level evidence on productivity differentials and turnover in Taiwanese manufacturing. *Journal of Development Economics 66*(1), 51–86.

Bartik, T. J. (1991, November). *Who Benefits from State and Local Economic Development Policies?* Number wbsle in Books from Upjohn Press. W.E. Upjohn Institute for Employment Research.

Bloom, N., M. Draca, and J. V. Reenen (2016). Trade Induced Technical Change? The Impact of Chinese Imports on Innovation, IT and Productivity. *Review of Economic Studies 83*(1), 87–117.

Bloom, N. and J. V. Reenen (2007). Measuring and Explaining Management Practices Across Firms and Countries. *The Quarterly Journal of Economics 122*(4), 1351–1408.

Bloom, N. and J. V. Reenen (2010, Winter). Why Do Management Practices Differ across Firms and Countries? *Journal of Economic Perspectives 24*(1), 203–224.

Blundell, R., R. Griffith, and J. van Reenen (1999). Market Share, Market Value and Innovation in a Panel of British Manufacturing Firms. *Review of Economic Studies 66*(3), 529–554.

Bond, S., J. A. Elston, J. Mairesse, and B. Mulkay (2003, February). Financial Factors and Investment in Belgium, France, Germany, and the United Kingdom: A Comparison Using Company Panel Data. *The Review of Economics and Statistics 85*(1), 153–165.

Boone, J. (2008, 08). A New Way to Measure Competition. *Economic Journal 118*(531), 1245–1261.

Collard-Wexler, A. and J. De Loecker (2015). Reallocation and Technology: Evidence from the US Steel Industry. *American Economic Review 105*(1), 131–171.

Criscuolo, C. (2009). *Innovation and productivity: Estimating the core model across 18 countries*. Innovation in Firms: A Microeconomic Perspective. Paris, France: Organization for Economic Cooperation and Development.

De Loecker, J. (2013). Detecting Learning by Exporting. *American Economic Journal: Microeconomics 5*(3), 1–21.

De Loecker, J., P. K. Goldberg, A. Khandelwal, and N. Pavcnik (2016). Prices, Markups and Trade Reform. *Econometrica 84*(2), 445–510.

De Loecker, J. and F. Warzynski (2012). Markups and Firm-Level Export Status. *American Economic Review 102*(6), 2437–71.

Eckel, C. and J. P. Neary (2010). Multi-Product Firms and Flexible Manufacturing in the Global Economy. *Review of Economic Studies 77*(1), 188–217.

Eslava, M., J. Haltiwanger, A. Kugler, and M. Kugler (2013). Trade and market selection: Evidence from manufacturing plants in colombia. *Review of Economic Dynamics 16*(1), 135–158.

Foster, L., J. Haltiwanger, and C. Syverson (2008). Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability? *American Economic Review 98*(1), 394–425.

Garcia-Marin, A. and N. Voigtänder (2016). Exporting and Plant-Level Efficicnecy Gains: It's in the Measure. Technical report, UCLA and Universidad de Chile.

Garcia-Marin, A. and N. Voigtländer (2013). Exporting and Plant-Level Efficiency Gains: It's in the Measure. NBER working paper #19033.

González, X. and D. Miles-Touya (2016). Estimating Production Functions when Prices are Partially Observed. Technical report, Manuscript, Universidad de Vigo.

Hall, B. H. (2002, Spring). The Financing of Research and Development. *Oxford Review of Economic Policy 18*(1), 35–51.

Hall, B. H. (2003). Innovation and Productivity. *Nordic Economic Policy Review 2*, 167–204.

Hall, B. H. and J. Lerner (2010). *The Financing of R&D and Innovation*, Volume 1 of *Handbook of the Economics of Innovation*, Chapter 0, pp. 609–639. Elsevier.

Hashmi, A. R. (2013). Competition and Innovation: The Inverted-U Relationship Revisited. *The Review of Economics and Statistics 95*(5), 1653–1668.

Hashmi, A. R. and J. Van Biesebroeck (2016). The Relationship between Market Structure and Innovation in Industry Equilibrium: A Case Study of the Global Automobile Industry. *The Review of Economics and Statistics 98*(1), 192–208.

Holmes, T. J. and J. A. Schmitz (2010). Competition and Productivity: A Review of Evidence. *Annual Review of Economics 2*(1), 619–642.

Hottman, C., S. J. Redding, and D. E. Weinstein (2016). Quantifying the Sources of Firm Heterogeneity. *Quarterly Journal of Economics*. forthcoming.

Hsieh, C.-T. and P. J. Klenow (2009). Misallocation and Manufacturing TFP in China and India. *The Quarterly Journal of Economics 124*(4), 1403–1448.

Janz, N., H. Lööf, and B. Peters (2004). Firm Level Innovation and Productivity: Is there a Common Story Across Countries? *Problems and Perspectives in Management 2*(2), 184–204.

Katayama, H., S. Lu, and J. R. Tybout (2009). Firm-level productivity studies: Illusions and a solution. *International Journal of Industrial Organization 27*(3), 403–413.

Levinsohn, J. and A. Petrin (2003). Estimating Production Functions Using Inputs To Control For Unobservables. *Review of Economic Studies 70*(2), 317–341.

Lööf, H., A. Heshmati, R. Asplund, and S.-O. Nåås (2003). Innovation and Performance in Manufacturing

Industries: A Comparison of the Nordic Countries. *International Journal of Management Research 2*(1), 5–36.

Maloney, W. F. and M. Sarrias (2014). Convergence to the managerial frontier. Policy Research Working Paper Series 6822, The World Bank.

Matsa, D. A. (2011). Competition and Product Quality in the Supermarket Industry. *The Quarterly Journal of Economics 126*(3), 1539–1591.

Melitz, M. J. (2003). The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity. *Econometrica 71*(6), 1695–1725.

Melitz, M. J. and G. I. P. Ottaviano (2008). Market Size, Trade, and Productivity. *Review of Economic Studies 75*(1), 295–316.

Mohen, P. and B. H. Hall (2003). Innovation and Productivity: An Update. *Eurasian Business Review 3*(1), 47–65.

Montenegro, C. E. and C. Pagés (2004). Who Benefits from Labor Market Regulations? Chile, 1960-1998. In *Law and Employment: Lessons from Latin America and the Caribbean*, NBER Chapters, pp. 401–434. National Bureau of Economic Research, Inc.

Mulkay, B., B. H. Hall, and J. Mairesse (2000, December). Firm Level Investment and R&D in France and the United States: A Comparison. NBER Working Papers 8038, National Bureau of Economic Research, Inc.

Olley, G. S. and A. Pakes (1996). The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica 64*(6), 1263–97.

Pavcnik, N. (2002). Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants. *Review of Economic Studies 69*(1), 245–276.

Rajan, R. G. and L. Zingales (1998). Financial Dependence and Growth. *American Economic Review 88*(3), 559–86.

Schumpeter, J. A. (1942). *Capitalism, socialism, and democracy*. New York: Harper and Row.

Smeets, V. and F. Warzynski (2013). Estimating productivity with Multi-Product Firms, Pricing Heterogeneity and the Role of International Trade. *Journal of International Economics 90*(2), 237–244.

Syverson, C. (2004, December). Market Structure and Productivity: A Concrete Example. *Journal of Political Economy 112*(6), 1181–1222.

van Leeuwena, G. and L. Klomp (2006). On the Contribution of Innovation to Multi-Factor Productivity Growth. *Economics of Innovation and New Technologies 15*(4/5), 67–390.

Van Reenen, J. (2011, May). Does competition raise productivity through improving management quality? *International Journal of Industrial Organization 29*(3), 306–316.

Van Reenen, J. (2013). *Advances in Economics and Econometrics: Proceedings of the Tenth World Congress of the Econometric Society*, Volume 3, Chapter Productivity and Management Practices. Cambridge University Press.

FIGURES

Figure 1: Productivity Distribution

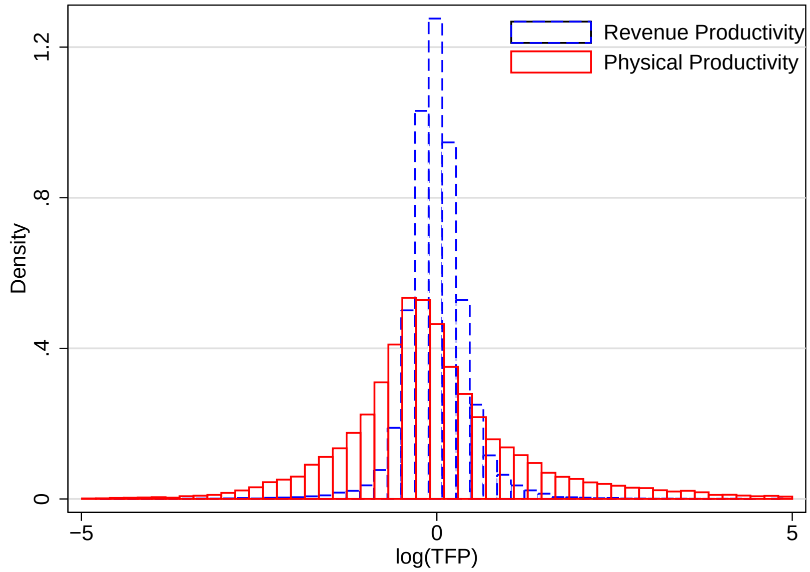

Notes: This figure shows the distribution of physical productivity and revenue productivity ("TFPR", blue-dashed bars) over a sample of 46,058 plant-year observations over 1996-2007. All variables are measured in logarithms, and are demeaned with respect to the respective (2-digit) sector-year averages.

TABLES

Table 1: Correlations and Standard Deviation

Correlations	Output Price	Input Price	TFPQ	TFPR	Markup
Output Price	1.0000	—	—	—	—
Input Price	0.1021	1.0000	—	—	—
Physical Productivity (TFPQ)	$\boxed{-0.8269}$	-0.1011	1.0000	—	—
Revenue Productivity (TFPR)	0.036	-0.035	$\boxed{0.5026}$	1.0000	—
Markup	0.0242	-0.0139	$\boxed{0.4243}$	$\boxed{0.8219}$	1.0000
Standard Deviation	0.6975	0.6164	0.8192	0.4482	0.4950

Notes: Note: This table shows correlations and standard deviations for plant-level variables for a sample of 25,396 plant-year observations (6,115 plants) over 1996-2007. All variables are measured in logarithms, and are demeaned with respect to the respective sector-year averages.

Table 2: Markups and productivity: IV results

	OLS	First Stage	2SLS	Red. Form
	(1)	(2)	(3)	(4)
Dep. Variable	$\ln(TFPQ_{ist})$	$\ln(\mu_{is,t-1})$	$\ln(TFPQ_{ist})$	$\ln(TFPQ_{ist})$
$\ln(\mu_{is,t-1})$	.0501***	—	.189***	—
	(.0102)		[.000]	
$\log(\mu_{-is,t-1})$	—	.440***	—	.0834***
		(.0242)		(.0233)
First Stage F-Stat	—	331.0	—	—
Industry-year FE	✓	✓	✓	✓
Observations	25,404	25,404	25,404	25,404

Notes: This table examines the effect markups on TFPQ. The OLS regression between of TFPQ on markups are reported in column 1. Column 2 reports first-stage results, together with the (cluster-robust) Kleibergen-Paap rK Wald F-statistic. The corresponding Stock-Yogo value for 10% (15%) maximal IV bias is 16.4 (8.96). Second stage results (column 3) report the p-values [in square brackets] for the Anderson-Rubin (Chisquare) test of statistical significance (heteroskedasticity-robust). This test is robust to weak instruments (see Andrews and Stock, 2005, for a detailed review). All regressions are run at the plant-year level, control for the logarithm of employment and for initial plant-level physical productivity, and include industry-year (at the 2-digit level) fixed-effects. Standard errors are clustered at the industry-year level. Key: ** significant at 1%; ** 5%; * 10%.

Table 3: Markups and productivity: Robustness Checks

Specification	(1) Input & Output Prices	(2) Reported AV Margin	(3) Single-product plants
$\log(\mu_{ij,t-1})$	.124**	.725***	.213**
	[.0494]	[.000]	[.0143]
First Stage F-Stat	202.0	80.71	299.0
Industry-year FE	✓	✓	✓
Observations	16,955	25,120	8,352

Notes: This table examines the effect markups on TFPQ. The OLS regression between of TFPQ on markups are reported in column 1. Column 2 reports first-stage results, together with the (cluster-robust) Kleibergen-Paap rK Wald F-statistic. The corresponding Stock-Yogo value for 10% (15%) maximal IV bias is 16.4 (8.96). Second stage results (column 3) report the p-values [in square brackets] for the Anderson-Rubin (Chisquare) test of statistical significance (heteroskedasticity-robust). This test is robust to weak instruments (see Andrews and Stock, 2005, for a detailed review). All regressions are run at the plant-year level, control for the logarithm of employment and for initial plant-level physical productivity, and include industry-year (at the 2-digit level) fixed-effects. Standard errors are clustered at the industry-year level. Key: ** significant at 1%; ** 5%; * 10%.

Table 4: Sample differences: ENIA vs EIT

	(1)	(2)	(3)	(4)	(5)
	Plant Size		Export Prob.	Productivity	Markup
Dependent Variable:	ln(workers)	ln(sales)	D(Exp=1)	ln(TFPQ)	ln(markup)
EIT dummy	1.269***	.203***	.766***	-.00854	.0668**
	(.137)	(.0241)	(.0848)	(.00854)	(.0278)
Sector-year FE	✓	✓	✓	✓	✓]
R^2	.067	.122	.059	.057	.264
Observations	25,404	25,404	25,404	25,404	25,404

Notes: The table reports the percentage-point difference of the dependent variable between plants in the survey of technological innovation (EIT) compared to plants in the manufacturing survey (ENIA). All regressions control for sector-year effects at the 2-digit level. Standard errors are clustered at the sector-year level. Key: ** significant at 1%; ** 5%; * 10%.

Table 5: Markups and Investment in Efficiency-Enhancing Activities

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent		R&D Expenditure		Patents	Machinery & Equipment	
Variable	$\ln(TFPQ_{ist})$	Overall	In-House	Licenses	Innovative	General
Panel A. Overall Effect						
log($\mu_{ij,t-1}$)	.570***	4.038***	3.122***	2.658***	2.062**	3.000***
weak-IV robust p-value:	[.000]	[.0002]	[.0019]	[.0018]	[.0319]	[.0023]
First Stage F-Stat	256.8	271.1	271.1	271.1	271.1	271.1
Industry-year FE	✓	✓	✓	✓	✓	✓
Observations	3,428	3,428	3,428	3,428	3,428	3,428
Panel B. Extensive Margin						
log($\mu_{ij,t-1}$)	—	.296***	.251***	.212***	.143*	.140
weak-IV robust p-value:	—	[.0002]	[.0046]	[.0055]	[.0810]	[.115]
First Stage F-Stat	—	271.1	271.1	271.1	271.1	271.1
Industry-year FE	—	✓	✓	✓	✓	✓
Observations	—	3,428	3,428	3,428	3,428	3,428

Notes: This table examines the effect of markups on investment in machinery, equipment and technology. Dependent variable ’x’ in the upper panel for columns 2-6 are log(1+x) to include zeros. The bottom panel uses dummies for positive values of ’x’ as dependent variables. All regressions controls for the initial physical productivity, size and for industry-year fixed effects. The first-stage statistic corresponds to the (cluster-robust) Kleibergen-Paap rK Wald F-statistic. The corresponding Stock-Yogo value for 10% (15%) maximal IV bias is 16.4 (8.96). P-values [in square brackets] are for the Anderson-Rubin (Chi-square) test of statistical significance (heteroskedasticity-robust). Standard errors are clustered at the industry-year level. Key: ** significant at 1%; ** 5%; * 10%.

Table 6: Effect of Technological Investment on TFPQ – 2SLS

	(1)	(2)	(3)	(4)	(5)
	R&D Expenditure		Patents &	Machinery & Equipment	
	Overall	In-house	Licenses	Innovative	General
Panel A. OLS Estimates					
Coefficient	.00850***	.00640***	.00551**	.00662***	.00812***
weak-IV robust p-value:	(.00203)	(.00241)	(.00215)	(.00205)	(.00252)
R-squared	.6115	.6101	.6099	.6104	.6107
Industry-year FE	✓	✓	✓	✓	✓
Observations	3,427	3,427	3,427	3,427	3,427
Panel B. 2SLS Estimates					
Coefficient	.0138***	.0102***	.0128***	.0303***	.0659***
weak-IV robust p-value:	[.0033]	[.0019]	[.0030]	[.0028]	[.0091]
First Stage F-Stat	10.60	14.14	9.359	2.603	5.581
Industry-year FE	✓	✓	✓	✓	✓
Observations	3,427	3,427	3,427	3,427	3,427

Notes: This table analyzes the effect of different components of technological investment on TFPQ. Each column runs a regression of TFPQ on each component of technological investment, initial physical productivity, size and industry-year fixed effects. Panel A presents OLS estimates; panel B present 2SLS estimates. In panel B, we instrument plants' technological investment with the sum of the technological investment made by other plants in the industry, an cateorical variables for innovation obstacles. The first-stage statistic corresponds to the (cluster-robust) Kleibergen-Paap rK Wald F-statistic. The corresponding Stock-Yogo value for 10% (15%) maximal IV bias is 16.4 (8.96). P-values [in square brackets] are for the Anderson-Rubin (Chi-square) test of statistical significance (heteroskedasticity-robust). Standard errors are clustered at the industry-year level. Key: ** significant at 1%; ** 5%; * 10%.

Table 7: Heterogenous Effects: Leaders and Laggards

Dependent Variable	Physical Productivity		R&D Expenditure	
	(1)	(2)	(3)	(4)
$log(\mu_{is,t-1})$	.576***	-.511	3.060***	-4.170
	(.0804)	(.373)	(1.037)	(3.480)
$log(\mu_{is,t-1}) \times TFPQ^{GAP}_{is,t-1}$	—	.142**	—	1.554*
		(.0717)		(.809)
$TFPQ^{GAP}_{is,t-1}$	—	-.719***	—	.543**
		(.0362)		(.255)
First Stage F-Statistic	276.5	26.57	281.3	28.33
Industry-year FE	✓	✓	✓	✓
Observations	3,344	3,344	3,344	3,344

Notes: This table analyzes the presence of heterogeneous effects depending on plants' distance to the technological frontier. We define frontier at the industry-year level, as the difference between the physical productivity of the most efficient plant and plant's TFPQ. R&D expenditure corresponds to overall investment in technological activities, and includes in-house R&D expenditure, spending in patent and licenses, and investment in innovative machinery and equipment, among other factors. All regressions controls for the initial physical productivity, size and for industry-year fixed effects. The first-stage statistic corresponds to the (cluster-robust) Kleibergen-Paap rK Wald F-statistic. The corresponding Stock-Yogo value for 10% (15%) maximal IV bias is 16.4 (8.96). Standard errors are clustered at the industry-year level. Key: ** significant at 1%; ** 5%; * 10%.

Table 8: Markups effect on Technological Investment by External Financing Dependence – 2SLS

	(1)	(2)	(3)	(4)	(5)
	R&D Expenditure		Patents &	Machinery & Equipment	
	Overall	In-house	Licenses	Innovative	General
Panel A. Above Median External Financing					
log $\mu_{is,t-1}$ (predicted)	4.809**	7.370***	4.364***	2.696	2.757**
weak-IV robust p-value:	[.0103]	[.0000]	[.0055]	[.1530]	[.0376]
First Stage F-Stat	52.44	52.44	52.44	52.44	52.44
Industry-year FE	✓	✓	✓	✓	✓
Observations	1,589	1,589	1,589	1,589	1,589
Panel B. Below Median External Financing					
log $\mu_{is,t-1}$ (predicted)	3.808***	.747	1.911*	1.864*	3.128**
weak-IV robust p-value:	[.0032]	[.4900]	[.0564]	[.0852]	[.0239]
First Stage F-Stat	283.5	283.5	283.5	283.5	283.5
Industry-year FE	✓	✓	✓	✓	✓
Observations	1,815	1,815	1,815	1,815	1,815

Notes: This Table replicates Table 5, but splitting the sample for industries with high and low external financing needs. Industries' dependence on external financing are defined according to Rajan and Zingales (1998) index. Dependent variable 'x' in all are log(1+x) to include zeros. All regressions controls for the initial physical productivity, size and for industry-year fixed effects. The first-stage statistic corresponds to the (cluster-robust) Kleibergen-Paap rK Wald F-statistic. The corresponding Stock-Yogo value for 10% (15%) maximal IV bias is 16.4 (8.96). P-values [in square brackets] are for the Anderson-Rubin (Chi-square) test of statistical significance (heteroskedasticity-robust). Standard errors are clustered at the industry-year level. Key: ** significant at 1%; ** 5%; * 10%.