Productivity, Investment, and Wealth Dynamics under Financial Frictions*

Alvaro Aguirre† Matias Tapia† Lucciano Villacorta†

Abstract

We develop an econometric framework to provide direct microeconomic evidence of the mechanisms underlying macroeconomic models of firm dynamics with financial frictions. Using administrative data, we non-parametrically estimate firm-level productivity and its effect on firms' investment and saving decisions in a unified framework that allows for financial frictions. The productivity process is non-linear, with larger persistence for more productive firms. We uncover heterogeneous responses of investment and wealth accumulation to productivity shocks. Our estimates are consistent with the presence of collateral constraints and the existence of a self-financing channel that can enhance aggregate TFP by up to 27%.

JEL classification: C33, E23, O11, L0
Keywords: Investment, wealth, self-financing, financial frictions, productivity.

*We are grateful to Daniel Ackerberg, Manuel Arellano, Stephane Bonhomme, Richard Blundell, Paco Buera, Andrea Caggese, Manuel Garcia Santana, Jordi Jaumandreu (discussant), Benjamin Moll, Josep Pijoan-Mas, Yongs Shin, Chad Syverson, Alonso Villacorta, Gianluca Violante and attendees at the Conference in honor of Manuel Arellano, STEG Workshop on Firms and Frictions, Midwest Macro Meetings, BIS CCA, the BSE Summer Forum, SED 2024, LACEA, the IPDC, World Congress ES, EEA, Santiago Macro Workshop, SECHI, UCSC, PUC Chile, CEMFI, CBC, UDP, and U Chile. Diego Huerta, Cristian Valencia, and Pablo Barros provided excellent research assistance. This study was developed within the scope of the research agenda conducted by the Central Bank of Chile (CBC). Our views do not represent the views of the CBC or its board members.

† Central Bank of Chile. Contact email: lvillacorta@bcentral.cl

1 Introduction

Over the past decade, a substantial body of literature has employed macro models featuring heterogeneous agents to quantify the impact of financial frictions at the firm level on aggregate productivity, capital, and income [see Buera et al., 2015]. An important insight provided by these models is that the quantitative effects of financial frictions are driven by the joint distribution of firm wealth and productivity and how this distribution evolves over time. If firm productivity and wealth are not well aligned, financial frictions can distort the relationship between investment and productivity, generating misallocation of capital and other production factors across firms with potentially important macro implications. Another pivotal insight from these models is that firms can gradually overcome financial frictions by accumulating wealth in response to persistent productivity shocks. This endogenous "self-financing channel" implies that wealth and productivity will align over time and has the potential to mitigate the aggregate adverse effects of financial frictions.[1] The firm's capability to accumulate wealth and the strength of this self-financing channel depend to a great extent on the persistence of productivity shocks faced by firms and the subsequent response of investment and wealth accumulation to productivity [e.g. Moll, 2014].

Key empirical questions arise: How (mis)aligned are the productivity and wealth distributions in the data? How persistent and volatile are firm productivity shocks? Do firms experiencing positive productivity shocks accumulate wealth rapidly enough to capitalize on these opportunities? Addressing these questions necessitates a thorough empirical examination of firms' productivity patterns and their specific policy functions for investment and wealth accumulation. However, current literature lacks a detailed empirical analysis using microdata, regarding the individual decisions of firms that are pivotal in driving the self-financing channel and its macroeconomic effects. This

[1] Quantitative models, in Banerjee and Moll [2010], and Midrigan and Xu [2014], suggest that self-financing is strong enough to rapidly undo the impact of financial frictions on misallocation. These findings have lent support to the idea that the macro implications of financial frictions are less significant than thought earlier.

paper aims to fill this gap by providing firm-level evidence on the joint dynamics of productivity, investment, and wealth.

Painting the joint distribution of firm wealth and productivity (and its evolution) requires first quantifying firm productivity, which is unobserved. We present a novel methodology to estimate firm-level production functions in the face of financial frictions and use it to nonparametrically characterize the nature of the productivity process and its effect on the firm's decisions. To the best of our knowledge, this is the first paper that estimates, using micro-data, the relevant policy functions for investment and wealth accumulation that emerge in macro models with financial frictions. We use these functions to empirically document the response of investment and wealth accumulation to productivity shocks at the micro level, study how these responses vary along the wealth and productivity distribution, and explore the strength of the self-financing mechanism in reducing misallocation.

To do so, we take advantage of a rich data set of manufacturing firms obtained from a census of administrative records of formal firms in Chile from 2006 to 2016. Besides including data on inputs and output at the firm level, the database provides information on the balance sheet of firms, allowing us to characterize wealth and directly analyze the role of financial frictions. The data set has the advantage of including a panel of firms of different sizes and characteristics, mostly private, in the context of an emerging economy. As we can follow individual firms over time, we can directly observe their wealth accumulation and investment decisions and relate them to the evolution of the estimated productivity process. This data set is an ideal laboratory to study the firm-level dynamics that lie at the foundations of the macroeconomic models used in the literature [see discussion in Diggs and Kaboski, 2022].[2]

Our empirical analysis is guided by the macroeconomic models that study financial frictions and the self-financing channel [e.g. Buera and Shin, 2011, Moll, 2014, Midrigan and Xu, 2014]. However, an im-

[2] "The biggest obstacle in researching financial frictions in developing countries is data availability. Ideally, data would consist of information on the firm ability and wealth over several years." (Diggs and Kaboski, 2022)

portant aspect of our tractable econometric framework is that it uncovers the firms' productivity process and its effects on firms' investment and wealth accumulation decisions without relying on a structural estimation. In contrast to fully specified structural approaches, which require the specification of particular functional forms for preferences and especially the distribution of unobserved productivity, we adopt a semi-parametric approach where we recover productivity from the firm production function and estimate nonlinear firm's policy rules that are compatible with a large class of heterogeneous-agent models with financial frictions where the relevant state variables are firm wealth and productivity. Given that the productivity process is a key determinant of the strength of the self-financing channel, a crucial advantage of our approach is to characterize the firm productivity process flexibly without relying on the linear AR(1) with normally distributed shocks traditionally used in structural models, which as we show, is at odds with the data. Furthermore, our nonlinear policy rules enable a comprehensive characterization of the entire distribution of micro-level investment and wealth accumulation propensities in response to productivity shocks.

We provide four key findings *First*, as opposed to the standard linear AR(1) productivity process with constant persistence embedded in most quantitative macro models, we find a highly nonlinear process revealing heterogeneity in the persistence of productivity shocks, which depends on the previous level of productivity and the magnitude and the sign of new shocks. This has important implications for the strength of self-financing, as higher (lower) persistence provides stronger (weaker) incentives for firms to accumulate wealth after a positive productivity shock (see Moll [2014]). We find that persistence is increasing in previous productivity, ranging from near-unity for high-productivity firms to as low as 0.65 for low-productivity firms. However, persistence can change in the face of extreme events.

Second, we find heterogeneous patterns in the response of firm investment to productivity shocks. We are the first paper to document a nonlinear relationship between investment and productivity at the firm level. We find larger responses in investment to productivity shocks at higher productivity levels. Such a larger response could be explained by our finding that more productive firms display higher productivity persistence. Also, we find that the investment reaction to positive productivity shocks is increasing in the stock of wealth for all productivity levels, providing support for models with collateral constraints.

Third, we are the first paper to document a significant, positive effect of productivity on wealth accumulation, which is heterogeneous across firms and depends on the previous level of wealth and productivity. The elasticity of productivity shocks to savings is high for low-wealth firms and weakens significantly as we move upwards along the wealth distribution, which is in line with the notion of *self-financing*. For high-productivity firms, such elasticity goes from one (which indicates a full transmission to savings) to 0.45 as we move up in the wealth distribution. For low-productivity firms, such elasticity ranges from 0.6 to 0.2.

Fourth, we are the first paper that characterizes and documents the joint distribution of productivity and wealth in the data. Our findings reveal that the correlation between productivity and wealth exhibited fluctuations within the range of 0.3-0.4 during the 2007-2015 period. Notably, the correlation pattern closely aligns with the dynamics of the economic cycle. Additionally, with the estimated parameters of the production function and the firm-level productivity, we estimate the marginal product of capital (MPK) distribution. We observe a substantial dispersion in the MPK across firms, with an average standard deviation of 0.35, indicating high capital misallocation across firms. We then use our estimated empirical model to quantify the strength of the self-financing channel. We find that through self-financing, the joint distribution of productivity and wealth becomes more aligned, and the distribution of MPK becomes more concentrated. Importantly, we find that self-financing has the potential to enhance aggregate productivity by up to twenty-seven percent compared to a scenario without the reallocation stemming from self-financing. However, while our results show that convergence in MPK does occur, it is slow, as differences in MPK persist for more than three decades, even for very productive firms. This suggests that the self-financing channel might operate at a slower pace than previously suggested in the literature [e.g. Banerjee and Moll, 2010, Midrigan and Xu, 2014].

Our approach Our empirical framework shares the spirit of the empirical consumption-household income framework [e.g. Arellano et al., 2017], which exploits panel data to estimate the empirical counterpart of the household policy rules and show how consumption responds to unobserved household income shocks. A crucial econometric difference in our framework lies in estimating the income shocks. In the household framework, income shocks and their effect on consumption are extracted directly from the household income data after removing demographic characteristics that are assumed to be orthogonal to the income shocks. By contrast, to estimate the unobserved firm income shocks (the productivity shocks) and their effect on investment and savings, we need to estimate the production function parameters, where the regressors are endogenous and correlated with unobserved productivity.

The first step of our analysis builds on previous literature that estimates production functions and productivity at the firm level. This literature relies on a proxy variable approach to recover productivity using the firm's input decisions [see Ackerberg et al., 2015, for a review]. For instance, Olley and Pakes [1996] recover productivity by inverting an investment demand function, whereas Levinsohn and Petrin [2003] invert the firm demand function for intermediate inputs. The proxy variable approach uses observed differences in input demands to control for differences in unobserved productivity in a production function regression. We extend this estimation method to allow for financial frictions, as otherwise, proxy methods deliver biased estimates of the production function and the productivity process. Intuitively, financial frictions generate differences in input demands for equally productive firms that the proxy variable method misinterprets as differences in unobserved productivity. Instead, our method compares input demands for firms with similar levels of wealth. Additionally, the proxy method is not well-suited to estimate flexible empirical policy functions, as it does not allow for unobservables besides productivity in the policies. In the presence of other shocks in the policies, productivity cannot solely be represented as a function of observable state variables. [3] This is an empirically restrictive assumption since it rules out the possibility of measurement error and idiosyncratic shocks affecting investment and wealth in the data.

A methodological contribution of this paper is to combine the insights of the self-financing channel with recent developments in non-linear panel models [Hu and Schennach, 2008, Arellano et al., 2017] to jointly identify and estimate the production function, the productivity process, and the investment and wealth accumulation policy functions in a setup with financial frictions. Due to the self-financing channel, more productive firms have higher incentives to invest today and accumulate wealth at the end of the period than less productive firms with the same level of wealth today. We use that conditional correlation between investment in t and wealth in $t+1$ to reveal differences in unobserved persistent productivity across firms and control for it in the production function regression.[4] Once the production function parameters are identified, the productivity process is identified from the time-series dependence structure of the firm net income process. Finally, the policy rules that depend on productivity are identified using non-parametric instrumental variables arguments given the exclusion restrictions provided by our dynamic model. In that sense, we are the first paper to provide the conditions for the identification of the empirical functions underlying the quantitative macro models with financial frictions (see the econometric transparency discussion in Andrews et al. [2020]).

Although the empirical model cannot provide direct policy counterfactuals, our results may be used directly or indirectly to calibrate structural models that are able to do so.

First, our production function and productivity estimates can be used to parametrize the firm's production function and productivity process directly in a structural model. We uncover new empirical

[3] This violates the scalar unobservable assumption in the proxy variable method Ackerberg et al. [2015].

[4] From an instrumental variable perspective, both policy functions can be thought of as noisy measures of unobserved productivity. We show that the wealth accumulation policy provides a valid external instrument for investment, which is used as a proxy variable with noise in a production function regression — that also includes the stock of wealth to control for collateral constraints.

results on the estimates of the firm production function and productivity process as we find significant differences once we control for financial frictions. We show that applying standard methods without controlling for financial frictions underestimates the marginal effect of capital (the constrained input) in the production function due to the negative correlation between capital and financial frictions and underestimates the productivity of constrained firms as they show larger investment gaps with respect to their optimal levels. As a result of the underestimation of the capital parameter and productivity, those methods overestimate the labor parameter to fit the production function. Consistent estimation of these objects is crucial as they are key inputs in structural models that quantitatively study the role of financial frictions and the self-financing channel.

Second, our empirical policy rules can be used as matching targets for other key parameters related to preferences and financial constraints. For example, Catherine et al. [2018] uses reduced-form evidence on the sensitivity of investment to collateral as the target to match in a structural estimation. Our results provide a set of new empirical moments to match like the complete distribution of investment and wealth accumulation sensitivities.

Related literature Our paper contributes to different streams in the literature. Our initial motivation is the macro-finance literature that studies the aggregate effects of financial frictions. We are closer to the set of papers focusing on collateral constraints and self-financing [e.g. Banerjee and Moll, 2010, Buera and Shin, 2011, Buera et al., 2011, Buera and Shin, 2013b, Caggese and Cuñat, 2013, Moll, 2014, Midrigan and Xu, 2014, Buera et al., 2015], as we guide our empirical specification by implications of these models. Our contribution is to provide novel direct evidence and an identification strategy on firms' wealth accumulation and investment decisions, which in these papers are an endogenous outcome of calibrated structural models built under different assumptions. Our estimates may help to discipline these models and provide further insights into their underlying mechanisms. Ours is the first paper to provide direct, firm-level evidence of the self-financing channel.

This paper also connects to two strands of research in corporate finance. One area of literature, starting with Fazzari et al. [1987], tries to identify financially constrained firms through their investment sensitivity to cash flows beyond profitability, typically captured by Tobin's Q or other observable characteristic. A second related area discusses the determinants of firms' cash holding decisions and associates them to firm characteristics such as growth opportunities and risk management.[5] In our framework, the investment and wealth accumulation policy functions are two key outcomes, and we identify unobservable productivity not only to control for profitability but also to estimate non-linear and interaction effects with our measure of collateral. Furthermore, since we follow the structural macro models, we focus on net wealth instead of cash flows.

The paper also connects with the empirical literature that estimates production functions at the firm level using the proxy variable approach [Olley and Pakes, 1996, Levinsohn and Petrin, 2003, Ackerberg et al., 2015, Doraszelski and Jaumandreu, 2013, 2018, Gandhi et al., 2020, Shenoy, 2020]. Our paper differs from these papers in several aspects. First, our paper studies the biases that appear when the proxy method is used to estimate the production function under financial constraints.[6] Second, we use the insights and mechanisms presented in macro models to propose a novel strategy that is robust to financial frictions. In this sense, our paper is the first paper that uses the self-financing channel to identify the firm productivity process and the firm production function. We allow for more flexible policy rules, including transitory shocks, unlike the proxy variable approach. Finally, a key difference is the identification and estimation of the policy functions.

Section 2 presents a model of firm dynamics to motivate the ingredients of the empirical model. Section 3 introduces the empirical

[5] See, for example, Opler et al. [1999], and Almeida et al. [2004]

[6] Shenoy [2020] studies how the proxy variable approach fails when any type of market frictions distort the firm's input choices and proposes a dynamic linear panel data approach (Arellano and Bond [1991], Blundell and Bond [1998]) to estimate the production function. In contrast, our method does not rely on a linear model for productivity, which, as our results suggests, is an assumption not satisfied in our data.

model. Section 4 establishes identification, while Section 5 describes the estimation methods. Section 6 describes the data and the results.

2 A Simple Model with Financial Frictions

We start with a stylized structural model inspired by the macro literature that studies financial frictions and the self-financing channel [see Buera et al., 2015, for a detailed analysis]. The model motivates the empirical policy rules we take to the data, provides the mechanisms to identify the empirical model, and illustrates the biases of the proxy variable approach.

Consider a price-taking firm with initial financial wealth A_{it}, capital K_{it} and productivity Z_{it} that solves the following dynamic problem to maximize the discounted value of distributed profits D_{it} choosing labor L_{it}, investment I_{it} and next period wealth A_{it+1}:

$$
\begin{aligned}
V(A_{it}, K_{it}, Z_{it}) &= \max_{A_{it+1}, I_{it}, L_{it}} D_{it} + \beta E\left[V(A_{it+1}, K_{it+1}, Z_{it+1})|Z_{it}\right], \\
\text{s.t.} \qquad D_{it} + g(A_{it+1}) &= Y_{it} - WL_{it} - (r+\delta)K_{it} + (1+r)A_{it}, \\
Y_{it} &= Z_{it}K_{it}^{\beta_k}L_{it}^{\beta_l} \\
K_{it+1} &= I_{it} + (1-\delta)K_{it}.
\end{aligned}
$$

where Y_{it} is the value added produced by firm i. Investment, which determines the next period's capital, is decided before observing next period's productivity, while labor is decided contemporaneously with productivity.[7] As preferences are linear, $g(\cdot)$ is assumed to be convex, ruling out corner solutions.[8] The firm discounts future flows at β, capital depreciates at δ, and the firm pays interest rate r for its debt, implicitly defined by $K_{it} - A_{it}$.

[7]This assumption implies that it takes a full period for new capital to be delivered and installed.

[8]Assuming linear preferences is not needed in our empirical framework but simplifies the illustrative analysis in this section. Including the convex function g introduces an incentive to smooth assets over time, ruling out solutions in which firms retain either all or none of their earnings. This specification combines ease of analysis with the general qualitative implications of models that introduce concavity in preferences.

The log of productivity z_{it} follows a Markovian process with distribution P_z.

$$P_z\left(z_{it} \mid z_{it-1}, z_{it-2}, \ldots\right) = P_z\left(z_{it} \mid z_{it-1}\right) \tag{1}$$

where $E\left[z_{it} \mid z_{it-1}\right] = \varphi\left(z_{it-1}\right)$ is a continuous function of z_{it-1}. The quantitative macro literature typically assumes normality for P_z and linearity for $\varphi\left(z_{it-1}\right)$.

Financial Constraint Following Buera et al. [2015] we consider the following specification

$$K_{it+1} \leq \kappa(A_{it}, Z_{it}) \tag{2}$$

Equation 2 implies that debt is limited by the repayment capacity of the firm, a combination of its productivity and current wealth. This captures the idea that financial friction depends on the profitability of the firm and its financial status. The first term in 2 is an *asset-based collateral constraint*, as net worth determines the part of the balance sheet that is owned by the firm and can be pledged as collateral. The second term in 2 represents *forward looking constraints* (or *earning-based constraints*), as persistent productivity determines the flow of current and future cash flows, which are the main factors in earning-covenants and earning-based lending.[9] Equation 2 nests standard constraints [Moll, 2014, Buera et al., 2015]. We left the function κ unrestricted and considered policy functions that are compatible with any financial constraint that is a function of wealth and productivity.

Optimality Conditions The FOC with respect to investment can be written as:

$$C_k E(Z_{it+1}|Z_{it})^{\frac{1}{1-\beta_l}} \left(I_{it} + (1-\delta)K_{it}\right)^{\frac{\beta_k}{1-\beta_l}-1} = \beta(r+\delta) + \mu(A_{it}, Z_{it}), \tag{3}$$

[9] Lian and Ma [2020], and Ivashina et al. [2022] provide substantial evidence that both types of financial constraints are prevalent in developed and developing countries, whereas Aguirre [2017], Brooks and Dovis [2020] and Drechsel [2022] show that both constraints are quantitatively important.

where C_k is a constant. The last term on the right-hand side is the wedge caused by financial frictions and is the multiplier of the collateral constraint (2), which is decreasing in A_{it}. Given the wedge, MPKs will not equalize across firms, so the equilibrium allocation of capital, a function of the current distribution of A_{it} and Z_{it}, is not efficient. Equation (3) generates an investment function (in logs) that depends nonlinearly on wealth and productivity.

$$i_{it} = h(z_{it}, k_{it}, a_{it}) \tag{4}$$

Finally, at the end of the period, the firm must decide on wealth accumulation, which is crucial to finance future investments. The FOC is given by:

$$g'(A_{it+1}) = \beta \left(1 + r + E_t \left[\kappa_A \mu(A_{it+1}, Z_{it+1})\right]\right) \tag{5}$$

When firms expect to be constrained tomorrow, an additional dollar of retained earnings allows them to increase investment tomorrow in κ_A dollars. The marginal benefit of wealth is the expected MPK net of borrowing costs, the value of the multiplier. As Z_{it} is persistent, firms with higher productivity anticipate higher future MPKs and consequently have stronger incentives to save compared to low-productivity firms with equivalent current wealth. This dynamic results in the wealth accumulation policy function (in logs)

$$a_{it+1} = g(z_{it}, k_{it}, a_{it}) \tag{6}$$

The extent of financial frictions and the strength of self-financing are reflected in the responses of investment and wealth accumulation to persistent productivity shocks and how these responses vary with available collateral and current productivity.

This simple setup illustrates the **goal of this paper**: to flexibly characterize, using microdata, the firm productivity process in (1) without parametrizing its distribution, and document its impact on firm decisions estimating the firm policy functions in (4) and (6) without relying on approximations and distributional assumptions.

Biases in proxy variable estimators in the presence of financial frictions. In this paper, we follow the industrial organization literature and recover firm productivity without using distributional assumptions by estimating the firm production function. However, the model above provides insights into the biases of the estimates that use standard empirical methods that do not account for financial frictions and how these biases can distort the interpretation of the production function and the productivity process. We illustrate our argument in the context of the influential paper by Olley and Pakes [1996], henceforth OP. However, the same logic applies to Levinsohn and Petrin [2003], as long as financial frictions affect the demand for materials as in Mendoza and Yue [2012]. The analysis also applies to Ackerberg et al. [2015], and Wooldridge [2009]. These biases are problematic, as having consistent estimates of the production function and the productivity process is key for structural macro models that quantitatively study financial frictions.

Here we discuss the intuition behind the biases, and we provide a detailed explanation using the macro model described at the beginning of this section in Online Appendix 1. Consider the log of the value-added production function described above:

$$y_{it} = \beta_l l_{it} + \beta_k k_{it} + z_{it} + \varepsilon_{it}, \tag{7}$$

where ε_{it} is measurement error in value added.[10] The challenge in the estimation of β_l and β_k is that z_{it} is an unobservable variable that is potentially correlated with the regressors k_{it} and l_{it}, creating an endogeneity problem in the OLS regression of y_{it} on k_{it} and l_{it}. The OP approach uses the investment policy function as an auxiliary equation to obtain information on z_{it}. In the absence of collateral constraints, by controlling for investment in the production function, OP eliminates the endogeneity problem and provides consistent estimates of β_l and β_k. Intuitively, OP interprets differences in firm investment in the data as differences in unobserved firm productivity.

However, in the model with financial frictions described above, the investment function does not only depend on productivity and

[10] We focus on a model with perfect competition where output prices are homogeneous across firms as in Olley and Pakes [1996], Levinsohn and Petrin [2003], Ackerberg et al. [2015], and Gandhi et al. [2020]. For production function estimation with monopolist competition, see De Loecker [2011].

initial capital, but also on net worth and its effect on credit access. Therefore, under borrowing constraints, equally productive firms with different wealth might have different investment, capital, and output. In consequence, if the implied heterogeneity in output is driven by heterogeneity in capital due to differences in credit access, the OP approach will wrongly assign such dispersion to productivity, as it will misinterpret differences in investment as solely coming from productivity gaps. As a result, the OP productivity proxy captures an important part of the effect of capital on output, underestimating its true marginal impact, and downward biasing the estimated coefficient for capital. Conversely, as long as financial frictions are relatively less severe for labor, the labor coefficient is upwardly biased. OP interprets a financially constrained firm with low investment as a low-productivity firm that hires "too many" workers and produces "too much" output relative to its proxy-OP productivity. Hence, it will assign a large role to labor in production, overestimating the labor elasticity. Furthermore, these biases in estimated factor elasticities can have significant effects on the measure of returns to scale. Additionally, OP will underestimate the productivity of financially constrained firms, as they have larger investment gaps with respect to their optimal levels. Therefore, the estimated productivity distribution across firms will also be biased. We empirically illustrate these biases and how our methodology corrects them, both with actual as well as from simulated data using a model with financial frictions.

3 General Empirical Framework

This section discusses the empirical model and its stochastic assumptions. The model consists of the production function in (7) and the empirical counterparts of the productivity process in (1) and the policy functions in (4) and (6).[11]

As we want to recover the productivity distribution from the data,

[11] We consider a Cobb-Douglas production function since it is the specification used in the structural models that study the self-financing channel. Our approach can accommodate more flexible production functions as long as productivity is Hicks-neutral.

we consider a very flexible specification for the productivity process using a quantile model:

$$z_{it} = Q_z\left(z_{t-1}, \eta_{it}\right) \tag{8}$$

with innovation $\left(\eta_{it} \mid z_{t-1}\right)$ uniform distributed $U\left(0,1\right)$ and persistence:

$$\rho\left(z_{t-1}, \tau\right) = \frac{\partial Q_z\left(z_{t-1}, \tau\right)}{\partial z} \tag{9}$$

where $\rho\left(z_{t-1}, \tau\right)$ measures the persistence of z_{t-1} when, at time t, the firm productivity is hit by a shock η_{it} that has rank τ. Equation 8 is a direct non-parametric model for $P_z\left(z_{it} \mid z_{it-1}\right)$, that leaves the dependence structure of z unrestricted beyond the Markovian assumption as opposed to the normal distribution usually assumed in macro models. Also, compared to the linear AR(1) traditionally used in the literature, it allows for nonlinear persistence with two main properties. First, for a given shock η_{it} the relationship between z_{it} and z_{it-1} depends on z_{it-1}. Therefore, persistence after a given innovation can vary across firms with different productivity levels z_{it-1}. Second, for a given z_{it-1} the relationship between z_{it} and z_{it-1} depends on η_{it}. This implies that unusually large shocks can change the relationship between current and past productivity, canceling the cumulative effects of past shocks. This modeling approach was introduced by Arellano et al. [2017] to model persistent income shocks to households. We are the first paper to use it to estimate productivity and production functions at the firm level. As it is standard in the literature, shocks η_{it+1} and ε_{it} are not part of the firm's information set when making decisions at t. The assumptions on the stochastic processes underlying both shocks are explained below.

Following the model in Section 2, capital k_{it} is a dynamic but predetermined input, decided in $t-1$ when the firm chose I_{it-1} since there is time to build physical capital, while labor l_{it} is a flexible input. The specification of the empirical policy rules follows the model discussed earlier, but a stochastic shock augments each policy function:

$$i_{it} = h_t\left(z_{it}, k_{it}, a_{it}, v_{it}\right), \tag{10}$$

$$a_{it+1} = g_{t+1}\left(z_{it}, a_{it}, k_{it}, w_{it+1}\right). \tag{11}$$

where h_t and g_{t+1} are the nonlinear reduced-form policy rules of investment and wealth that can be derived in a firm-dynamics model with financial frictions as the one discussed earlier. The terms v_{it} and w_{it+1} capture other unobserved factors besides z_{it} that affect the evolution of investment and wealth. For example, in the context of our earlier stylized model, shocks to collateral constraints could affect the investment policy function. Firms may face temporary idiosyncratic shocks that affect the relationship between debt, productivity, and collateral (i.e., $\kappa(Z_{it}, A_{it}, v_{it})$). In the case of the wealth accumulation policy function, stochastic shocks can come from unexpected fluctuations in the valuation of firms' assets. If these occur in the interim between the distribution of dividends (when equation 5 is solved) and when the firm uses wealth as collateral to borrow (when equation 4 is solved), they will appear as unplanned changes in the valuation of collateral.[12] More generally, the inclusion of shocks can bridge the transition from the insights of the stylized model to an empirical model that deals with actual microdata and issues such as measurement errors that emerge from the use of different datasets.

We also assume that h_t and g_{t+1} are monotonic in v_{it} and w_{it+1}, respectively. The specification in (10) nests several nonlinear empirical investment functions studied in the literature [e.g. Olley and Pakes, 1996, Cooper and Haltiwanger, 2006, Gala et al., 2020]. The two major innovations of our framework are (i) the inclusion of wealth a as an additional state variable in (10) to control for the existence of collateral constraints and (ii) the explicit modeling of wealth dynamics in (11) and its relationship with productivity (the self-financing channel). An additional important difference to Gala et al. [2020] and Cooper and Haltiwanger [2006] is that we explicitly include z_{it} as a state variable. In contrast, those papers replace z_{it} for value-added,

[12] The inclusion of v_{it} in the investment function represents a departure from the unobservable scalar assumption required by the proxy variable. It is important to recall that under the proxy variable approach, the investment function is of no interest by itself, as it is only an auxiliary equation to recover the production function.

an endogenous variable.[13] Finally, as it is standard in the macro literature, we model the labor input as a non-dynamic input:

$$l_{it} = n_t\left(z_{it}, a_{it}, k_{it}, w_{l,it}\right), \tag{12}$$

where equation (12) is the empirical labor decision. An extension from the stylized model in Section 2 is that our empirical specification allows for potential effects of financial frictions over hiring decisions, as represented by the inclusion of a_{it} in the policy function. Once again, the term $w_{l,it}$ represents an i.i.d. shock that is independent of the state variables a_{it}, k_{it}, and z_{it}. This shock can capture exogenous transitory shocks to wages in the model in Section 2 or optimization errors as discussed in Ackerberg et al. [2015]. The inclusion of $w_{l,it}$ and v_{it} introduce stochastic fluctuations in l_{it} and k_{it} besides z_{it} and state variables. This feature mitigates the functional dependence issue inherent in the proxy variable approach, as noted in Ackerberg et al. [2015].

Timing Firms begin period t with a stock of productive capital k and wealth a_t. During t, firms learn about their productivity from equation (8) and decide labor according to (12) and produce output from (7). Since productivity follows a Markov process, firms forecast next period's productivity $E\left[z_{it+1} \mid z_{it}\right]$, based on z_{it} and decide k_{t+1} according to (10) to prepare for production in the subsequent period. Finally, at the end of period t, firms decide whether to allocate their earnings towards increasing wealth (such as by holding cash, investing in financial assets, or debt repayment) and be less financially constrained in $t+1$ or distribute dividends, as outlined in equation (11).

To complete the model description, we formally make the following assumptions, using the notation $x_i^t = (x_{i1}, \ldots, x_{it})$ for any variable x_{it}.

Assumption 1. *(Conditional Independence). For all $t \geq 1$:*

[13] As Gala et al. [2020] argue in footnote 10, including z_{it} instead of y_{it} requires the estimation of the production function, which adds several econometric problems, most significantly, endogeneity. One of the contributions of our paper is to address this issue and consistently estimate the production function and the correct investment equation as a function of unobserved productivity.

(i) ***Output Shock:*** *ε_{it+s} for all $s \geq 0$ is independent over time and independent of $a_i^{t-1}, z_i^{t-1}, i_i^{t-1}, k_i^{t-1}, l_i^t$, y_i^{t-1} and η_{it+s}. Also ε_{i1} is independent of z_{i1}, a_{i1} and k_{i1}, and $E[\varepsilon_{it}] = 0$.*

(ii) ***Productivity Shock:*** *η_{it+s} for all $s \geq 0$ is independent over time and independent of $a_i^{t-1}, z_i^{t-1}, i_i^{t-1}, k_i^{t-1}, l_i^{t-1}$, and y_i^{t-1}.*

(iii) ***Policy Functions Shocks:*** *v_{it} and w_{it+1} are mutually independent, independent over time, and also independent of z_{i1}, a_{i1}, k_{i1} $(\varepsilon_{is}, \eta_{is})$ for all s and of v_{is} and w_{is+1} for all $s \neq t$.*

Assumption 2. *(First Order Markovian). For all $t \geq 1$:*
(i) a_i^{t+1} is independent of $\left(a_i^{t-1}, k_i^{t-1}, z_i^{t-1}\right)$ conditional on (a_{it}, k_{it}, z_{it})
(ii) i_i^t is independent of $\left(a_i^{t-1}, k_i^{t-1}, z_i^{t-1}\right)$ conditional on (a_{it}, k_{it}, z_{it})

Parts (i) and (ii) of Assumption 1 state that current and future productivity and production shocks, which are independent of past productivity and production shocks are also independent of the current and past wealth and capital stocks, investment, and labor decisions. The initial wealth stock a_{i1}, initial capital stock k_{i1}, and initial productivity z_{i1} are arbitrarily dependent. Allowing for a correlation between a_{i1}, k_{i1}, and z_{i1} is important, as wealth and capital accumulation upon entry in the sample may be correlated with past persistent productivity shocks. Part (iii) requires investment and wealth shock to be mutually independent, independent over time, and independent of production components. Assumption 1 implies that ε_{it}, v_{it} and w_{it+1} are independent of the state variables $(k_{it}, a_{it,} z_{it})$ and mutually independent conditional on $(l_{it}, k_{it}, a_{it,} z_{it})$. Hence, Assumption 1 provides the exclusion restrictions necessary for identification, while Assumption 2 is a first-order Markov condition on wealth and capital dynamics. Assumption 2 is a standard assumption both in macro models and in the empirical literature that estimates production functions.

Investment and Wealth Accumulation Propensities The nonlinear functions h_t and g_{t+1} allow for heterogeneous effects of productivity shocks on investment and wealth accumulation, depending on the collateral and productivity of the firm. Our objects of interest are the following average derivative effects with respect to z_{it}:

$$\Phi_{it}^{h} = \Phi^{h}\left(a_{it}, k_{it}, z_{it}\right) = E_{v_{it}}\left[\frac{\partial h_t\left(z_{it}, k_{it}, a_{it}, v_{it}\right)}{\partial z}\right] \tag{13}$$

$$\Phi_{it+1}^{g} = \Phi^{g}\left(a_{it}, k_{it}, z_{it}\right) = E_{w_{it+1}}\left[\frac{\partial g_{t+1}\left(z_{it}, k_{it}, a_{it}, w_{it+1}\right)}{\partial z}\right] \tag{14}$$

where the expectations are taken with respect to the idiosyncratic shocks in the policies. Φ_{it}^{h} and Φ_{it+1}^{g} measure the average propensities of investment and wealth accumulation in response to productivity shocks. These are key objects to understand the nature of financial frictions and how firms adjust to them.

Estimating these propensities can shed light on the specific nature of the credit constraints. For example, if firms face collateral constraints, Φ_{it}^{h} should be increasing in a_{it} for a given z_{it}, as the investment of wealthier firms can respond more to productivity shocks. Similarly, the existence of earning-based [Drechsel, 2022] or forward-looking constraint [Buera et al., 2015] would imply that Φ_{it}^{h} is increasing in z_{it} for a given level of wealth, as more productive firms can leverage up more in function of their future flows. Of course, both types of constraints can be active, reflecting the heterogeneous nature of contracts and credit relationships firms face. As mentioned earlier, these propensities also provide evidence of firms' response to financial frictions. For instance, if the self-financing channel exists as in Moll [2014], Φ_{it+1}^{g} is always positive and (weakly) decreasing in a_{it} for a given productivity and increasing in z_{it} for a given value of current wealth.

4 Identification

Given the goal of this paper, it is crucial to show that the nonlinear model we aim to estimate can be identified from the micro-data we have at hand without distributional assumptions. Recently, Hu and Schennach [2008] and Arellano et al. [2017] have established conditions under which nonlinear dynamic models with latent variables are non-parametrically identified under conditional independence restrictions. We build on these papers and use the insights of the self-financing channel to provide non-parametric identification of the empirical model introduced in Section 3. In particular, the goal of this section is to show that β_k, β_l, $Q_z(z_{t-1}, \eta_{it})$, h_t, g_{t+1} are identified from data on (y_{it}, k_{it}, l_{it}, i_{it}, a_{it}, a_{it+1}) given that (z_{it}, w_{it+1}, v_{it}, ε_{it}) are not observed by the econometrician and z_{it} is correlated with (l_{it}, a_{it}, k_{it}).

The sketch of identification is sequential. First, we identify the production function parameters β_k and β_l. Once we have identified the production function, we define the firm net income process $\tilde{y}_{it} = y_{it} - \beta_k k_{it} - \beta_l l_{it}$ and use its autocorrelation structure to establish the identification of the productivity process. Finally, we identify the policy functions that depend on the latent productivity we identified from the firm net income.

4.1 Intuition in a linear model

To build intuition for identification, let's consider the case where the policy functions and the productivity process are normally distributed: $i_{it} = h_z z_{it} + h_a a_{it} + v_{it}$, $a_{it+1} = g_z z_{it} + g_a a_{it} + w_{it+1}$ and $z_{it} = \rho_z z_{it-1} + \eta_{it}$.[14]

Production function Similar to the proxy approach, we can invert the investment function: $z_{it} = \pi_1 i_{it} + \pi_2 a_{it} + \pi_4 v_{it}$ where $\pi_1 = 1/h_z$, $\pi_2 = -h_a/h_z$ and $\pi_4 = -1/h_z$ and replaced into the production function:

$$y_{it} = \beta_l l_{it} + \beta_k k_{it} + \pi_1 i_{it} + \pi_2 a_{it} + \tilde{\varepsilon}_{it} \tag{15}$$

where $\tilde{\varepsilon}_{it} = \varepsilon_{it} + \pi_4 v_{it}$. In the absence of investment shocks (i.e., $v_{it} = 0$), a simple OLS regression between y_{it} on l_{it}, k_{it}, i_{it} and a_{it} identifies β_l and β_k, as in the proxy variable approach. The main difference to the proxy variable approach is that the regression in (15) controls for a_{it}. Hence, rather than looking for unconditional differences in investment across firms to control for differences in productivity, we are considering differences in investment across firms with the same collateral constraints.

[14]For simplicity of exposition, we consider the case where $h_k = g_k = 0$. When $h_k \neq 0$ we recover β_l from (15) and β_k in a second step (see Section 5.1)

In the more general case with investment shocks (i.e., $v_{it} \neq 0$), z_{it} can not be expressed only as a function of observables and parameters. Therefore, even after controlling for a_{it}, we cannot disentangle variation in investment coming from z_{it} from variation in other shocks.[15]

The self-financing channel is key for identification. According to the model in Section 2, for a given level of state variables a_{it} and k_{it}, more productive firms should increase investment more and simultaneously accumulate more wealth to reduce the constraint in the future. Therefore, the covariance between i_{it} and a_{it+1}, conditioned on current wealth a_{it}, allows us to isolate the variation in i_{it} due to variation in z_{it} from the variation in i_{it} due to variation in v_{it}. Hence, a_{it+1} can be used as an instrument for investment in equation (15) given conditional independence (*Assumption* 1) - wealth does not have a direct effect in the production function- and the relevance condition implied by the self-financing channel $g_z \neq 0$. A regression between $E\left[y_{it} \mid a_{it+1}, l_{it}, k_{it}, a_{it}\right]$, and $\left[l_{it}, k_{it}, E\left[i_{it} \mid a_{it+1}, k_{it}, l_{it}, a_{it}\right], a_{it}\right]$ identifies $\{\beta_l, \beta_k\}$.

The identification sketch we develop here provides a novel and simple estimation procedure by doing an IV regression within the proxy method where the theoretical insights of macro models justify the external instrument.

Productivity Process and Policy Functions Once we have identified β_k and β_l, we can write the firm net-income process $\tilde{y}_{it} = y_{it} - \beta_k k_{it} - \beta_l l_{it}$ as an additive model with two independent latent variables (given *Assumption* 1).

$$\tilde{y}_{it} = z_{it} + \varepsilon_{it} \tag{16}$$

Then, we can use the autocorrelation structure of $(\tilde{y}_{i1,\cdots,}\tilde{y}_{iT})$ to construct valid instruments in a similar fashion to Arellano and Bond [1991]. For example, we can use equation (16) to express the productivity model in terms of the observed net income process $\tilde{y}_{it} = \rho_z \tilde{y}_{it-1} + \eta_{it} + \varepsilon_{it} - \rho_z \varepsilon_{it-1}$ and use three waves of the net-income

[15] This violates the unobservable scalar assumption required by the proxy approach and, therefore, the production function model can not be consistently estimated using OLS since $E\left(i_{it}\tilde{\varepsilon}_{it}\right) \neq 0$

process $\{\tilde{y}_{it-2}, \tilde{y}_{it-1}, \tilde{y}_{it}\}$ to identify ρ_z from an IV regression using $\tilde{y}_{it-2}$ as an instrument for $\tilde{y}_{it-1}$ (given that $\tilde{y}_{it-1}$ and ε_{it-1} are correlated). Then, the variance of the productivity shock and the variance of the measurement error in income are identified from $E(\tilde{y}_{it}\tilde{y}_{it-1}) = \rho_z E(\tilde{y}_{it-1}\tilde{y}_{it-1}) - \rho_z \sigma_\varepsilon^2$ and $E(\tilde{y}_{it}\tilde{y}_{it}) = \rho_z^2 E(\tilde{y}_{it-1}\tilde{y}_{it-1}) + \sigma_\eta^2 + (1-\rho_z^2)\sigma_\varepsilon^2$.

For the policy rules, we can use equation (16) to express the model (that depends on the unobserved z_{it}) in terms of the observed net income process $a_{it+1} = g_z\tilde{y}_{it} + g_a a_{it} + g_k k_{it} + w_{it+1} - g_z\varepsilon_{it}$, and identify the linear policy functions from an IV regression using $\tilde{y}_{it-1}$ as an instrument for $\tilde{y}_{it}$ (given that $\tilde{y}_{it}$ and ε_{it} are correlated) and controlling for a_{it} and k_{it}.

4.2 Nonparametric Identification

We generalize the ideas sketched in the linear version to provide identification of the more general model in section 3, where the policy functions and the productivity process are modeled non- parametrically. As in the linear case, the sketch of identification is sequential. To establish nonparametric identification, we use the following statistical conditions that we connect with the insights of the macro model discussed in Section 2.

Let $X_{it} = (a_{it,} k_{it}, l_{it})$ be the covariates of the model and let $f(a \mid b)$ be a generic notation for the conditional density $f_{A|B}(a \mid b)$.

Condition 1. *Almost surely in covariate values X_t: (i) the joint density $f(y_t, i_t, a_{t+1}, z_t \mid X_t)$ is bounded, as well as all its joint and marginal densities; (ii) the characteristic function of ε_{it} has no zeros on the real line; (iii) for all $z_{1t} \neq z_{2t}$, $Pr[f(i_{it} \mid z_{1t}, X_t) \neq f(i_{it} \mid z_{2t}, X_t)] > 0$; (iii) $f(a_{t+1} \mid z_t, X_t)$ is complete in z_{it}. (iv) for $\tilde{y}_{it} = y_{it} - \beta_l l_{it} - \beta_k k_{it}$, $f(\tilde{y}_{it} \mid \tilde{y}_{it-1})$, $f(z_{it} \mid \tilde{y}_{it-1})$, $f\left(z_{it} \mid \tilde{y}_i^T\right)$ are complete and the distribution of $f\left(z_{it} \mid a_i^t, k_i^t, \tilde{y}_i^T\right)$ is complete in $\left(a_i^{t-1}, k_i^{t-1}, \tilde{y}_i^T\right)$.*

Condition 1-(i) requires bounded densities. Condition 1-(ii) is a technical assumption previously used in the literature.[16] The normal

[16] This condition is used for the i.i.d shock of the household income in Arellano et al. [2017] and for the i.i.d shock in the firm production function in Hu et al. [2020].

distribution and many other standard distributions satisfy this condition. Condition 1-(iii) requires that $f(i_{it} \mid z_{it}, X_{it})$ be non-identical at different values of z_{it}. This condition is weaker than the assumption in Olley and Pakes [1996] and Ackerberg et al. [2015], where the realization of investment has to be monotonic in z_{it}. Here, we require that two firms with the same current wealth and capital level but different productivity levels have different investment probabilities. In the presence of financial frictions, the question arises whether the high z_t, which has stronger investment incentives compared to the low z_t firm, can obtain better access to credit, given that both firms have the same current k_t and a_t and hence face the same collateral constraint. The macroeconomic model outlined in Section 2, which incorporates forward-looking or earnings-based constraints where lenders possess insights into firm profitability, satisfies this condition.[17] With forward-looking constraints, the high z_t firm consistently appears more financially attractive than the low z_t firm, even if both firms possess the same current wealth a_t. Lian and Ma [2020] and Ivashina et al. [2022] provide substantial evidence that earnings-based lending practices are widespread in developed and developing countries. Additionally, models featuring only asset-based constraints meet this criterion, as long as the firms can borrow as much as they want, paying a premium in the interest rate that depends on the collateral as in Cavalcanti et al. [2021].

Condition 1-(iv) is a completeness condition commonly assumed in the literature on nonparametric instrumental variables [Newey and Powell, 2003].[18] Intuitively, we need enough variation in the density $f(a_{it+1} \mid z_{it}, a_{it}, k_{it})$ for different values of z_{it}. This requires a statistical dependence between wealth accumulation a_{it+1} and productivity z_{it} conditioned on the observed state variables. This requirement is met by the self-financing channel in the model described in Section 2, where conditional on the same level of current wealth, highly productive firms should accumulate more wealth to relax the friction in

[17] Lenders may not possess complete information regarding z_{it}. However, as long as lenders have some imperfect signals of future firm cash flows, the high z_t firm is more likely to secure a loan than the low z_t.

[18] The distribution of $\tilde{y}_{it} \mid \tilde{y}_{it-1}$ is complete if $E[\phi(\tilde{y}_{it}) \mid \tilde{y}_{it-1}] = 0$ implies that $\phi(\tilde{y}_{it}) = 0$ for all ϕ in some space.

the future than less productive firms. In instrumental variable terminology, this is a relevance condition that ensures that a_{it+1} is a valid instrument for z_{it}. For example, in the linear case, the completeness condition collapses to $g_z \neq 0$. Condition 1-(v) requires that z_{it} and z_{it-1} are statistically dependent, which is ensured by the Markovian assumption.

These conditions lead to the following theorem, which sequentially combines the results in Hu and Schennach [2008] and Arellano et al. [2017].

Theorem 1. *If Assumption 1, Assumption 2 and condition 1 (i)-(v) hold, then* β_k, β_l, $Q_z(z_{t-1}, \eta_{it})$, h_t, g_{t+1} *are identified from data on* y_{it}, k_{it}, l_{it}, i_{it}, a_{it} *for* $T \geq 4$.

Below we discuss the sketch of the sequential identification and leave the details for Online Appendix 2.

Production Function From *Assumption* 1, ε_{it}, v_{it}, and w_{it+1} are independent conditional on $(l_{it}, k_{it}, a_{it,} z_{it})$, which can be interpreted as the exclusion restrictions in a nonlinear IV setting. Using this conditional independence assumption, we can write the conditional distribution of the observed variables $f(y_t, i_t \mid a_{t+1}, X_t)$, which is a data object, in terms of some elements of the model that we aim to identify:

$$f(y_t, i_t \mid a_{t+1}, X_t) = \int f(y_t \mid z_t, k_t, l_t)\, f(i_t \mid z_t, X_t)\, f(z_t \mid a_{t+1}, X_t)\, dz_t \tag{17}$$

We notice that equation (17) can be framed into the setup studied in Hu and Schennach [2008]. Given *condition* 1(i)-(iv), Theorem 1 of Hu and Schennach [2008] can be applied to our setting to show that the distribution of the production function $f(y_t \mid z_t, k_t, l_t)$ is identified from the data, which leads to the identification of the production function parameters [see Hu et al., 2020]). A novelty of our approach is that our model with financial frictions provides a second policy rule (the self-financing channel) that connects the latent productivity with an observed variable a_{it+1} that is not directly linked to the production function regression (i.e., a_{it+1} is not an input in the production function regression), so we can use it as an instrument. In online Appendix 2, we show that the production function is identified through a "non-parametric" version of the proxy-IV discussed in the linear version, where the investment function is used as a proxy variable with noise for productivity, and the wealth accumulation function is used as an instrument.

Productivity Process Once we have identified β_k, β_l, and given that the productivity is Hicks-neutral, we can write the firm net-income process $\tilde{y}_{it} = y_{it} - \beta_k k_{it} - \beta_l l_{it}$ as an additive model with two independent latent variables (given *Assumption* 1).[19]

$$\tilde{y}_{it} = z_{it} + \varepsilon_{it} \tag{18}$$

To non-parametrically identify the productivity process, we rely on the fact that the net-income process in (18) has a Hidden-Markov structure (by *Assumption* 1) where $\{\tilde{y}_{it-2}, \tilde{y}_{it-1}, \tilde{y}_{it}\}$ are independent given z_{it-1}. The additivity of the net-income process and *condition* 1-(v) allow us to identify the joint distribution of $(\varepsilon_{i2,\cdots,}\varepsilon_{iT-1})$ and the joint distribution of $(z_{i2,\cdots,}z_{iT-1})$ from the autocorrelation structure of $(\tilde{y}_{i1,\cdots,}\tilde{y}_{iT})$ for $T \geq 3$ and identify $Q_z\,(z_{t-1}, \eta_{it})$ for $T \geq 4$ extending the arguments of the linear case.[20]

Policy Functions Once $\left(z_{i1} \mid \tilde{y}_i^T\right)$ is identified, we use *Assumption* 1 and *Assumption* 2 to construct the following IV moment restriction, which allows us to relate the conditional distribution of observable variables $f\left(a_1, k_1 \mid \tilde{y}^T\right)$, $f\left(a_{t+1} \mid a^t, k^t, \tilde{y}^T\right)$, and $f\left(i_t \mid a^t, k^t, \tilde{y}^T\right)$ which are data objects, to the distribution of the policy rules we want to identify (see proposition 3).

$$f\left(a_1, k_1 \mid \tilde{y}^T\right) = E\left[f\left(a_1, k_1 \mid z_1\right) \mid \tilde{y}_i^T = \tilde{y}^T\right] \tag{19}$$

$$f\left(a_{t+1} \mid a^t, k^t, \tilde{y}^T\right) = E\left[f\left(a_{t+1} \mid z_t, a_t, k_t\right) \mid a_i^t = a^t, k_i^t = k^t, \tilde{y}_i^T = \tilde{y}^T\right] \tag{20}$$

[19] For identification and estimation of production functions with non-neutral productivity see Doraszelski and Jaumandreu [2018] and Villacorta [2018].

[20] Equation (18) has a similar structure to the household income process model with non-linear Markovian persistent shocks studied in Arellano et al. [2017].

$$f\left(k_{t+1} \mid a^t, k^t, \tilde{y}^T\right) = E\left[f\left(k_{t+1} \mid z_t, a_t, k_t\right) \mid a_i^t = a^t, k_i^t = k^t, \tilde{y}_i^T = \tilde{y}^T\right] \tag{21}$$

where the expectation in (19) is taken with respect to the density of z_{i1} given $\tilde{y}_i^T$ for fixed values of a_1 and k_1 and the expectation in (20) and (21) are taken with respect to the density of z_{it} given $\tilde{y}_i^T$, k_i^t, and a_i^t for a fixed value of a_{t+1} and k_{t+1}, respectively. Equation (19) is analogous to a nonlinear IV problem where z_{i1} is the endogenous regressor and $\tilde{y}_i^T$ is the vector of instruments. The difference with a standard nonlinear IV is that the "endogenous regressor" in the moment condition in (19) is a latent variable. However, this is not a problem since we have identified $\left(z_{i1} \mid \tilde{y}_i^T\right)$ using the production function. Provided that the distribution of $\left(z_{i1} \mid \tilde{y}_i^T\right)$ is complete (*condition* 1(v)), the unknown density $f\left(a_1, k_1 \mid z_1\right)$ is identified from (19). Similarly, equations (20) and (21) can be interpreted as nonlinear IV restrictions where a_{it} and k_{it} are the controls (they are arguments in the wealth function and investment functions), and the vector $\tilde{y}_i^T$ contains the excluded instruments. Given *condition* 1(v) and *Assumption* (2), the distributions $f\left(a_{t+1} \mid z_t, a_t, k_t\right)$ and $f\left(k_{t+1} \mid z_t, a_t, k_t\right)$ for $t > 2$ are identified recursively from equations (20) and (21). The identification of $f\left(a_{t+1} \mid z_t, a_t, k_t\right)$ and $f\left(k_{t+1} \mid z_t, a_t, k_t\right)$ allows us to recover the policy functions $g_{t+1}\left(\right)$ and $h_t\left(\right)$. In the same fashion of the linear case, we are using the autocorrelation structure of $\tilde{y}_i^T$, conditioned on current values of a_{it} and k_{it}, to construct instruments (lagged and lead values of the firm's net income process) to identify the policies.

5 Empirical Strategy

In this section, we discuss two approaches to estimate different versions of the model. First, we consider a parsimonious model where at least one of the policies is a quasi-linear function in productivity. For this model, we propose a novel procedure that consists of an IV regression within the proxy variable framework, following the identification strategy presented in section 4.1. Second, we consider a more flexible model that allows for unrestricted nonlinear effects of productivity, and we propose a flexible estimation method well suited for nonlinear models with latent variables as the one identified in section 4.2.

5.1 Parsimonious policy functions

Proxy-IV The identification of β_l and β_k using the IV-proxy method strategy requires that at least one of the two policy functions is polynomial of degree one in z_{it} and separable in z_{it} and the shock. The other policy function and the distribution of the shocks are left unrestricted.[21] Consider the following investment function:

$$i_{it} = h\left(z_{it}, k_{it}, a_{it}, v_{it}\right) = h_1\left(k_{it}, a_{it}\right) + h_2\left(k_{it}, a_{it}\right) z_{it} + v_{it}, \qquad (22)$$

It is important to notice that model (22) is flexible enough to capture heterogeneous effects of productivity on investment depending on the level of collateral. Meanwhile, the wealth accumulation policy function is left unrestricted. As in the proxy variable approach, we can invert equation (22) in terms of productivity:

$$z_{it} = \pi_1\left(k_{it}, a_{it}\right) + \pi_2\left(k_{it}, a_{it}\right) i_{it} + \omega_{it} \qquad (23)$$

where $\pi_1\left(k_{it}, a_{it}\right) = -h_1\left(k_{it}, a_{it}\right)/h_2\left(k_{it}, a_{it}\right)$, $\pi_2\left(k_{it}, a_{it}\right) = 1/h_2\left(k_{it}, a_{it}\right)$ and $\omega_{it} = -v_{it}/h_2\left(k_{it}, a_{it}\right)$. Replacing (23) in the production function:

$$y_{it} = \beta_l l_{it} + \phi\left(k_{it}, a_{it}\right) + \pi_2\left(k_{it}, a_{it}\right) i_{it} + \omega_{it+1} + \varepsilon_{it}, \qquad (24)$$

where $\phi(k_{it}, a_{it}) = \beta_k k_{it} + \pi_1\left(k_{it}, a_{it}\right)$. As we emphasized in section 4, an OLS regression of (24) does not provide a consistent estimator of β_l since $E\left(\omega_{it} \mid i_{it}\right) \neq 0$. However, given *Assumption 1*, a_{it+1} can be used as an instrument for i_{it} in equation (24). Therefore, we propose the following two-stage procedure:

First Stage: Estimate (24) with an IV estimator using $\pi_2\left(k_{it}, a_{it}\right) a_{it+1}$ as the instrument for $\pi_2\left(k_{it}, a_{it}\right) i_{it}$. The IV regression delivers a consistent estimator of β_l, $\phi\left(k_{it}, a_{it}\right)$ and $\pi_2\left(k_{it}, a_{it}\right) a_{it+1}$.

[21] This model is more flexible than the simple example discussed in section 4.1 to give intuition on the identification approach. However, the identification approach in section 4.1 still applies to this model.

For instance, in the linear case where $g_2(k_{it}, a_{it}) = 1$, a_{it+1} will be the instrument for i_t.

Second Stage: Combining equation (23) with the Markovian model for a linear productivity process $z_{it} = \rho_z z_{it-1} + \eta_{it}$:

$$z_{it} = \rho_z \pi_1 (k_{it-1}, a_{it-1}) + \rho_z \pi_2 (k_{it-1}, a_{it-1}) i_{it-1} + \rho_z \omega_{it-1} + \eta_{it}, \quad (25)$$

Replacing equation (25) into the production function:

$$y_{it} - \beta_l l_{it} = \beta_k k_{it} + \rho_z \pi_1 (k_{it-1}, a_{it-1}) + \rho_z \pi_2 (k_{it-1}, a_{it-1}) i_{it-1} + \rho_z \omega_{it-1} + \eta_{it} + \varepsilon_{it}, \quad (26)$$

using *assumption 1* we can define the following moment condition from equation (26)

$$E (\omega_{it-1} + \eta_{it} + \varepsilon_{it} \mid k_{it}, k_{it-1}, a_{it-1}, a_t) = 0, \quad (27)$$

The moment condition in (27) allows us to identify β_k. We refer to this novel estimator as *Proxy-IV*. Once β_l and β_k are estimated we can estimate the productivity process and the policy functions following the IV strategy discussed in section 4 for the simple cases where productivity and the policies are linear functions.

5.2 Flexible policy functions

To estimate a more flexible model that allows for nonlinear persistence in productvity and nonlinear interactions between z_{it} and observed state variables in the policies, we bring the following nonlinear specifications to the data:

(i) For productivity, we implement the following quantile specification:

$$z_{it} = Q_z (z_{t-1}, \eta_{it}) = \sum_{r=1}^{R} \alpha_r^Q(\tau) \phi_r^Q (z_{it-1})$$

where τ represents the τth conditional quantile of z_{it} given z_{it-1}, ϕ_r^Q is a dictionary of functions and α_r^Q the parameters associated which are quantile-specific, allowing the effect of z_{it-1} on z_{it} to change with the shocks. The quantile model is a direct non-parametric model for the conditional distribution of productivity, as it does not assume normality or impose separability in the productivity process, leaving the dependence structure of z_{it} unrestricted beyond the Markovian assumption. We also implement a more parsimonious model that is nonlinear in past productivity but separable in the shock.

(ii) For the policy functions, we use these nonlinear specifications:

$$i_{it} = \sum_{r=1}^{R} \alpha_r^h \phi_r^h \left(z_{it}, k_{it}, a_{it}, \delta_t^h \right) + v_{it}$$

$$a_{it+1} = \sum_{r=1}^{R} \alpha_r^g \phi_r^g \left(z_{it}, k_{it}, a_{it}, \delta_t^g \right) + w_{it+1}$$

$$l_{it} = \sum_{r=1}^{R} \alpha_r^n \phi_r^n \left(z_{it}, k_{it}, a_{it} \right) + w_{l,it+1}$$

where ϕ_r^h, ϕ_r^g and ϕ_r^n are dictionaries of functions and α_r^h, α_r^g and α_r^n are the associated parameters. Note that ϕ_r^h, ϕ_r^g and ϕ_r^n are anonymous functions without an economic interpretation, as they are just building blocks of flexible models. The objects of interest will be summary measures of the derivative effects constructed from the model, like the propensities discussed in Section 3. We follow the proxy variable literature and model the functions as high-order polynomials to allow for flexible interactions between productivity and observed state variables. In our baseline specification of the nonlinear model, we specify stationary policy functions with additive errors that are normally distributed to have a more parsimonious model to take to the data, but, as we showed in Section 4, the model is non-parametrically identified with time-varying functions, non-additive errors and without parametric distributions.

Stochastic EM Estimation Algorithm (SEM) We adapt the stochastic EM algorithm in Arellano and Bonhomme [2016], and Arellano et al. [2017] to our firm's framework to estimate the nonlinear model. See details in Online Appendix 3.

6 Data and Empirical Results

6.1 Data

We use administrative records generated by Chile's tax collection agency (*Servicio de Impuestos Internos* - SII). The records cover all firms that operate in the formal sector. Each firm is assigned a unique identifier by SII, so they can be tracked across time preserving anonymity. We use the information contained in income tax form F22, submitted annually by firms. The data set contains information on *firms* (as opposed to *plants*) of all sizes and sectors, although we focus on the manufacturing sector. Firms are defined as productive units that generate revenue, utilize production factors and operate under a unique tax ID. Form F22 has firm-level information on annual sales, expenditures on intermediate materials, a proxy for the capital stock ("immobile assets") and the firm's wage bill, as well as the firm's economic sector. We combine this information with tax form 1887, which reports monthly information on workers employed by the firm, and therefore allows us to calculate a measure of annual employment adjusted by the number of months per worker. Crucially, form F22 provides information on the firm's balance sheets. In particular, we can build a measure of net worth, defined as the difference between reported total assets and total liabilities[22]. This allows us to combine the information on the production side traditionally used in the literature on production functions with information on the firm's self-reported wealth and its evolution across time.

To clean up the raw data, we follow several steps. First, we drop observations with zero or missing information for the capital proxy, sales, expenditures on intermediate inputs, employment, or net worth. Second, we focus on firms that have at least five workers. Third, we build a measure of annual investment by using the annual change in

[22]In particular, we use code 123 of form F22, "Total del Pasivo", for total liabilities. This variable is the combination of all the liabilities of the firm, as the tax form does not provide a decomposition between financial liabilities, credit from suppliers, etc. Similarly, total assets come from code 122, "Total de Activos", which combines all assets, including financial instruments as well as our capital proxy, "Activo Inmovilizado", code 647. Net worth is calculated simply as the difference between both. This means that our measure of physical capital (code 647) is equal to net worth (code 122 - code 123) plus total liabilities (code 123) net of non-capital assets (code 122 - code 647).

the capital stock and assuming a 10% depreciation rate.[23] The final dataset has 4867 firms in the manufacturing sector between 2005 and 2016. As discussed earlier, the data set provides a panel of firms of different sizes and characteristics in the context of an emerging economy. Although we do not have information on whether firms are publicly traded, the relatively small coverage of the Chilean stock market (768 firms across all sectors) implies that almost all of our firms must be private. Having information on balance sheets is an advantage relative to most databases used in the literature on production function estimations, either from surveys or administrative records, which typically provide detailed information on the production side but do not account for assets or wealth. Moreover, we can directly observe wealth accumulation and investment decisions at the individual level, as well as the dynamics of output, inputs, and the estimated productivity process. The combination of financial statements and information on the production side is not unique to our dataset. Long and detailed information is available for a large number of countries in datasets such as Compustat, Amadeus, and Orbis. However, relative to those sources, our dataset has the advantage of including a heterogeneous set of firms operating in an emerging economy. In that sense, this might be a better setup to study the effects of financial frictions that are likely to be less relevant in the developed world, in particular for relatively large firms.[24]

Table 1 presents some descriptive statistics of the data. As expected, there is a large degree of heterogeneity between firms. Sales for firms in the 90th percentile are 40 times larger than those in the 10th percentile, while differences in capital or investment are even larger. While the average firm has 91 workers, the median firm has only 20, and firms in the 10th percentile have 7. There is also a large variation in net worth, both in levels as a ratio to capital. This high-

[23]As an alternative, we also use the information on tax form F29, which has monthly data on investment in machinery and equipment. The behavior of both investment series is very similar.

[24]Other datasets, such as the Enterprise Surveys conducted by the World Bank, are similar to ours in that they also include firms of all sizes in developing countries, although, by their nature, they are less suited to follow a specific firm across several consecutive years, as we do here.

lights that the data contains a diverse set of firms, some of them quite small and with very low levels of wealth. While our data still has omissions (as it can not account for firms in the informal sector), it seems fit to provide a rich characterization of the behavior of heterogeneous firms over time, and the potential role of financial frictions in the context of a developing country and enriches the evidence previously available in the literature, in the spirit of the discussion in Kaboski [2021] and Diggs and Kaboski [2022].[25]

	Mean	p10	p50	p90
Value Added (million CLP)	1647.4	39.7	188.0	1536
Capital (million CLP)	2393.9	7.90	90.5	1197.9
Number of Employees	91.73	7	20	150
Investment (million CLP)	549.7	0.7	16.1	270.7
Net Worth (million CLP)	868.0	5.1	37.2	365.0
Capital to Output ratio	2.19	0.06	0.46	2.43
Net Worth to Capital Ratio	4.76	0.05	0.41	3.79

Table 1: Sample Descriptive Statistics

6.2 Empirical Results

We now use the data to implement the methodology discussed in Section 5. We begin by estimating firm-level production functions, accounting for the presence of financial frictions, and then use those estimates to study the properties of the productivity process. In the second part of the section, we present a novel empirical characterization of investment and wealth accumulation policy functions, highlighting the role of non-linearities and providing an empirical analysis of the self-financing channel.

[25] "Perhaps the biggest obstacle in researching financial frictions in developing countries is data availability. Ideally, data would consist of information on the firm ability and wealth over several years. Additionally, data may not include representative coverage of all firms. To have a full understanding of the firm side of an economy, it is necessary to include businesses across sectors and wealth, privately and publicly owned, and formal and informal firms." (Diggs and Kaboski, 2022)

In line with the earlier discussion, our estimates for the production function and the underlying productivity process, robust to financial frictions, differ significantly from those of the proxy variable approach. We also show that the productivity process of firms is highly nonlinear, with larger persistence for highly productive firms. However, extremely large productivity shocks can change productivity. These findings are in stark contrast with the linear AR(1) productivity process typically assumed in the literature. Regarding policy functions, we find large heterogeneity in the sensitivity of investment and wealth accumulation to productivity shocks. We present novel evidence on the nonlinear relationship between investment and productivity, with larger investment responses to productivity shocks in more productive firms. Our results suggest that both collateral constraints and earning-based constraints are present. Finally, we also show novel evidence of the existence of self-financing, with a very large savings propensity in low-wealth, productive firms. However, the impact of self-financing appears to be limited, as convergence in the MPK between constrained and unconstrained firms is slow.

6.2.1 Production Functions

We start by comparing our two novel estimators that control for financial frictions (Proxy-IV and SEM) with OP -the proxy variable approach in Olley and Pakes [1996]- which uses investment as an auxiliary equation to recover productivity, and LP -the proxy variable approach of Levinsohn and Petrin [2003]-, which uses intermediate inputs to recover productivity.

As discussed earlier, we expect OP to underestimate the capital elasticity and overestimate the labor elasticity. By a similar argument, we expect the same type of bias in other methodologies relying on a proxy variable approach, such as LP. In contrast, by controlling for wealth in the policy functions, our estimators can discriminate between productivity and the effects of collateral constraints. In addition, by relying on the co-movements between wealth accumulation and investment decisions, after controlling for the current stock of net wealth, our estimators can disentangle productivity shocks from transitory shocks that can temporarily affect investment and saving decisions.

Table 2 presents the results of the full estimation of the production function parameters (β_l, β_k) using the four methodologies. The general pattern is consistent with the presence of financial constraints and with our earlier theoretical predictions.

The estimate of β_l is 0.67 for OP and, as expected, decreases significantly to 0.44 in Proxy-IV and 0.46 in SEM. This estimated labor elasticity is in line with the labor share computed from aggregate data in Chile. Conversely, the opposite pattern holds for capital: the estimate of β_k is 0.35 for OP and increases to 0.42 for Proxy-IV and 0.43 for SEM. Similar biases appear in LP, which suggest that financial frictions are also present in demand for intermediate goods as in Mendoza and Yue [2012] and Bigio and La'o [2020].[26]

It is worth noting that the estimates of the production function parameters are very similar in proxy-IV and SEM. Even though we show that the model is non-parametrically identified from data, in order to devise tractable estimation methods, we impose some restrictions on the empirical model. Having robust results with both estimators suggests that neither the parametric assumption on the policy shocks in SEM nor the quasi-linear policy in IV affects the estimation of the production function parameters.[27]

These differences in input elasticities have relevant implications for the degree of returns to scale at the firm level, a crucial parameter to understand aggregate dynamics. In particular, OP results are consistent with constant returns to scale, while Proxy-IV and SEM both imply decreasing returns to scale with a span of control around 0.87. This figure lies on the upper end of the range used in the related literature.[28] This lower span of control relative to OP implies a larger

[26] Our proxy variables estimates of the production function are similar to ones in Gandhi et al. [2020]. Their proxy variable estimates with the Chilean data are 0.77 for β_l and 0.33 for β_k.

[27] For the proxy-IV, we assume that one of the policies is a quasi-linear function of degree 1 in productivity and leave the other policy function and the distribution of productivity and shocks completely unrestricted. In SEM, we allow all the policies to be nonlinear in all state variables, including productivity, but we parameterize the distribution of the policy shocks.

[28] For instance, Buera and Shin [2013a] use 0.79 while Midrigan and Xu [2014] use 0.85.

entrepreneurial income share that can be retained by firms, allowing for faster wealth accumulation to overcome financial constraints.

To complement our results, we simulate data from an extended version of our stylized model to confirm the biases of the proxy variable approach and the robustness of our proposed estimators. In line with the empirical estimates, we set $\beta_k = 0.43$ and $\beta_l = 0.44$ in the calibrated model.[29] Table 3 presents the estimates for simulated data. As expected, OP delivers biased estimates, whereas Proxy-IV and SEM recover the true underlying parameters.[30] Therefore, data generated from a quantitative model, explicitly including financial frictions, supports our insights regarding the biases of traditional methodologies, as well as our novel estimators.

	OP	LP	Proxy-IV	SEM
β_l	0.67	0.81	0.44	0.46
	0.008	*0.007*	*0.01*	*0.003*
β_k	0.35	0.33	0.42	0.43
	0.05	*0.04*	*0.01*	*0.007*
σ_ϵ	0.68	0.62	0.22	0.20
Observations	13516	13516	13516	13516
Firms	4867	4867	4867	4867

Table 2: Production Function Estimates from Microdata

Note: The table shows the Production function estimates from administrative data for Chile, using alternative methodologies: OP - Olley and Pakes [1996]-, LP- Levinsohn and Petrin [2003]-, and two estimators that control for financial friction, Proxy-IV and SEM.

[29] See Online Appendix 4 for calibration details.

[30] As the model does not include intermediate inputs as required by the LP estimator, we only use the OP, Proxy-IV, and SEM estimators.

	OP	Proxy-IV	SEM
β_l	0.505	0.443	0.442
β_k	0.397	0.424	0.431

Table 3: Production Function Estimates Using Simulated Data

Note: Production function estimates from simulated data using alternative methodologies: OP - Olley and Pakes [1996]-, and two estimators that control for financial friction, Proxy-IV and SEM. The model used to generate data is described in Online Appendix 4.

6.2.2 Productivity Process: Distribution

Figure 1 depicts the productivity distribution across firms for the proxy variable approach and our more general model (SEM) that controls for financial frictions. In OP, the standard deviation of productivity is 0.18, significantly lower than 0.40 under SEM (see Table 4). We also find that the gap between ours and OP's productivity estimates, i.e., the fraction by which true productivity is underestimated, is increasing in the firm's productivity. For instance, the coefficient of a linear regression between $z_{it}^{SEM} - z_{it}^{OP}$ and z_{it}^{SEM} is 0.7.

The fact that OP dampens productivity differences between firms is once again consistent with the presence of financial frictions: OP underestimates the productivity of constrained, high-productivity firms. Conversely, the productivity of unconstrained but low-productivity firms, which can invest comparatively more, is overestimated. Hence, by ignoring firm wealth, OP estimates a more compressed distribution relative to methods that are robust to frictions. From figure 1 we can see that using the non-parametric quantile model (SEM-Quantile) delivers a very similar distribution than the one estimated using the nonlinear model with normal errors, so assuming normality for productivity seems to be a reasonable assumption.

6.2.3 Productivity Process: Persistence

As discussed earlier, the persistence of productivity is key for self-financing, as it relates to the incentives for wealth accumulation (Moll, 2014). If positive productivity shocks are not expected to last, incentives for wealth accumulation are weaker.

To highlight the importance of controlling for financial frictions when estimating productivity persistence, we first show the estimates of a linear AR(1) model for productivity using the proxy variable approach and our proposed estimator. Table 4 presents the results for productivity persistence when we fit a linear model. The first row displays the autocorrelation of the estimated productivity, ρ_z. The estimated value of ρ_z raises from 0.53 under OP to 0.87 in proxy-IV and 0.85 under SEM, respectively. This implies that OP could underestimate the incentives for self-financing.

Non-linearities. In most quantitative macro models, productivity is assumed to follow an AR(1) process like the one in Table 4. As discussed in Section 3, one of the contributions of this paper is to uncover firm productivity without relying on either linearity or distributional assumptions. Figure 2 shows that the productivity process appears to be highly non-linear. Therefore, the traditional assumption of linearity might be at odds with the data.[31] We estimate two different models to disentangle the role played by past and new productivity shocks on the nonlinear persistence and the role of a normal parametric assumption. The left panel of Figure 2 depicts the estimated persistence for different levels of initial productivity (horizontal axis) in a model that is non-linear in past productivity but separable in new shocks $z_{it} = \varphi\left(z_{it-1}\right) + \eta_{it}$ with η_{it} normally distributed. This model allows persistence to be heterogeneous across firms but does not allow new shocks to change the current persistence. The micro-data reveals high heterogeneity in productivity persistence with a positive monotonic relationship in past productivity. That is, firms at the lower end of the productivity distribution have smaller persistence (around 0.65), whereas ex-ante very productive firms display a very high per-

[31] Also, production function estimators that relies on linearity of the productivity as in Arellano and Bond [1991], Blundell and Bond [1998] or Shenoy [2020] cannot be applied in our framework.

sistence (near unity). Therefore, highly productive firms will very likely remain at the upper end of the productivity distribution in the future. This novel result has important implications for the study of financial frictions and bodes well for self-financing, as it suggests that highly productive firms have both the ability and the incentives to build up collateral in order to converge toward their optimal capital. We analyze this notion more formally in the next section when we embed the non-linear productivity process in the estimation of the policy functions of firms and evaluate whether wealth and investment decisions change with productivity.

The right panel (panel b) of Figure 2 displays the estimated persistence of the more flexible model $z_{it} = Q_z\left(z_{t-1}, \eta_{it}\right)$. As discussed in Section 3, this model allows the persistence to change with past productivity and new shocks and does not restrict the conditional distribution of productivity. Thus, for a given value of a new productivity shock, the relationship between z_{it} and z_{it-1} depends on z_{t-1}, and for a given value of z_{it-1} the relationship between z_{it} and z_{it-1} may change in the face of extremely large (negative or positive) shocks. The 3-d graph displays the estimated persistence for different values of past productivity and new shocks. On the two horizontal axes, we report the percentile of past productivity and the percentile of the innovation (the shock) of the quantile process. A value at the lower end of the innovation distribution represents a very low positive productivity shock, whereas a value at the upper end represents a very large positive shock. As before, we uncover a huge heterogeneity in persistence across firms that aligns with the results of the parsimonious model with normal errors. For the most common types of events (shocks) of a size close to the median of the shock distribution (the middle section of the right horizontal axis), the relationship between past productivity (left horizontal axis) and persistence (vertical axis) is positive and qualitatively consistent with the result in the left panel, implying that for *median shocks*, persistence is higher for ex-ante highly productive firms. However, persistence is high if firms experience shocks that align with their previous productivity level. For firms starting with high productivity levels and experiencing substantial positive shocks, persistence approaches unity. Likewise, firms with initially low productivity and encountering low productivity shocks also exhibit near-unity persistence. However, persistence can change abruptly in the face of extreme events.[32] For example, productivity persistence in very productive firms drops from almost one to 0.6 in the wake of an extremely adverse shock. A similar thing happens at the bottom of the productivity distribution after an extremely favorable shock. This means that large, infrequent shocks, besides having a direct effect on impact, can also alter the existing relationship between past and current productivity, canceling the cumulative effect of past shocks and permanently altering the trajectory of productivity. Therefore, in the aftermath of an unusually large shock, self-finance incentives can change drastically.

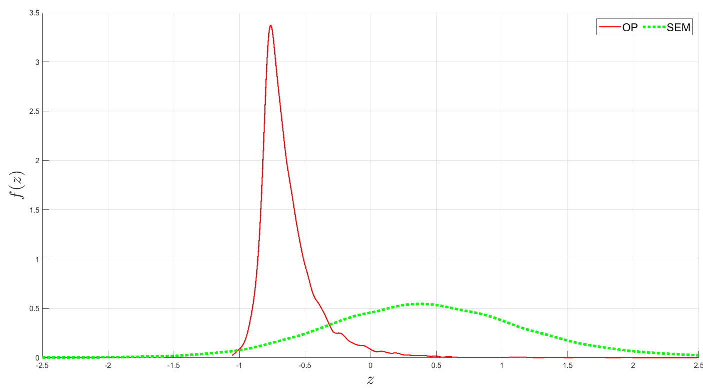

FIGURE 1: Estimated distribution of productivities

Notes: The figure shows the estimated distribution of firm-level productivities using administrative microdata for Chile, under alternative methodologies: OP - Olley and Pakes [1996] and the SEM algorithm using Normal shocks.

6.2.4 Policy Functions

We now present the estimated policy functions, one of the main goals of our empirical exercise. Given our interest in understanding the role

[32]This is consistent with similar results for household income shocks (Arellano et al. [2017])

	OP	Proxy-IV	SEM
ρ_z	0.53	0.87	0.85
	0.01	*0.01*	*0.01*
σ_η	0.18	0.30	0.39
Observations	13516	13516	13516
Firms	4867	4867	4867
R^2	0.37	-	0.70

Table 4: Estimated Parameters of the Productivity Process

Note: The table shows the estimated parameters for the firm-level productivity process from administrative microdata for Chile, using alternative methodologies: OP - Olley and Pakes [1996]-, and the two estimators that control for financial frictions, Proxy-IV and SEM.

of financial frictions and the self-financing channel, we pay special interest to the estimation of policy functions and the analysis of the economic forces that underlie them.

6.2.5 Investment Policy Function

The estimated propensities from the investment model inform us about how firms' behavior in response to the same productivity shock varies with different levels of the state variables. Panels (a)-(c) of Figure 3 display the estimated average derivative effect of productivity on investment $\hat{\Phi}_t^h(a,k,z)$ (the investment propensity). These panels visually represent how the investment response to productivity changes along the $\frac{A}{K}$ distribution for three different values of z. Assessing the propensities for different wealth-to-capital ratios $\frac{A}{K}$ under a constant capital level elucidates the varying investment responses to productivity shocks contingent on the firm's financial position. It is worth noting that lower $\frac{A}{K}$ values correspond to highly leveraged firms by

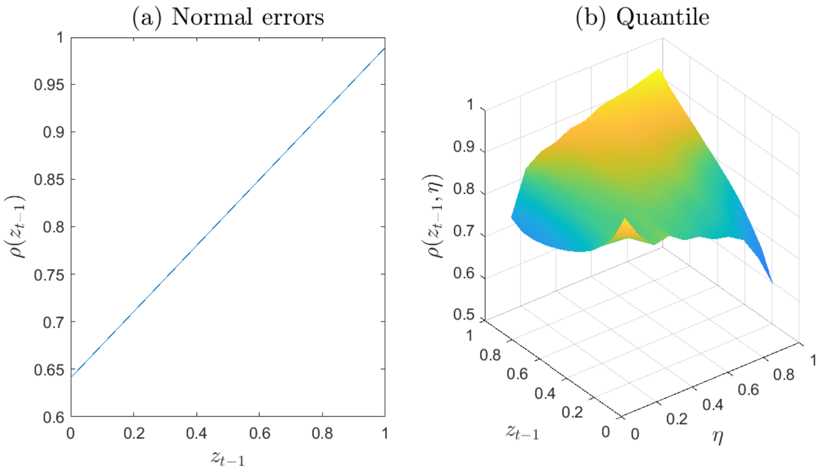

FIGURE 2: Estimated nonlinear persistence of productivity

Notes: The figure shows the estimated nonlinear persistence of firm-level productivity using administrative microdata for Chile. The first plot displays the persistence of estimated productivity using the model with separable errors along the distribution of past productivity, whereas the second plot displays the persistence of estimated productivity using the quantile model where the size and the sign of the shock might affect the persistence depending on the past value of productivity.

definition.

The estimated investment responses to productivity exhibit significant heterogeneity, spanning a range from approximately 0.05 to almost 0.7. Notably, propensities are observed to be lowest in firms characterized by both low productivity levels and meager wealth. Arguably, these firms are less able to adjust investment after a positive and persistent productivity shock as they might be collateral-constrained and can not rely too much on current and future earnings. However, investment propensities increase as we move along both the wealth and productivity distribution.

The main message from figure 3 is that investment propensities are increasing in wealth regardless of productivity. For instance, the investment propensity rises from 0.05 to 0.10 for low-productivity firms, from 0.28 to 0.38 for medium-productivity firms, and from 0.52 to

0.67 for high-productivity firms as we move along the wealth distribution. Interestingly, the biggest change (in levels) occurs for the most productive firms. This result might reflect that, even with forward-looking, productive firms with low wealth are still more financially constrained in relative terms and benefit the most from an additional unit of wealth.

Another important message is that the sensitivity of investment to productivity shocks increases with z. This is, for given levels of wealth and capital, investment responses are larger for ex-ante, more productive firms. This is consistent with the characteristics of the estimated non-linear productivity process described in the previous subsection and the fact that persistence increases with productivity. Moreover, the higher propensity for high-productivity firms also appears to be consistent with the empirical implications of models of financial constraints in which firm productivity can affect firm lending contracts and borrowing opportunities, in which firms can use their future cash flows as collateral and expand their investment.[33] For instance, the investment propensity of a high-productivity firm located at the bottom of the $\frac{A}{K}$ distribution is high (around 0.5). Despite these firms possessing few assets to be pledged as collateral (relative to their level of debt), they can strongly react to a positive productivity shock, a result that is at odds with a model in which only assets can be used as collateral.[34]

To have a taste of how propensities behave using the actual combinations of state variables in the data, we compute the propensity of each of the firms in our sample and plot it against the wealth-to-capital ratio $\frac{A}{K}$ in Figure 4 panels (a)-(c). We use our estimated productivity variable to cluster firms in three different "productivity groups": (i) low-productivity firms with productivity below the 50 percentile of

[33] As in models with earning-based constraints as Lian and Ma [2020], Drechsel [2022] and Camara and Sangiacomo [2022], or forward-looking constraints as in Buera et al. [2015].

[34] The higher investment propensity for highly productive firms may also reflect a form of conditional convergence, as their current capital might be further away from their optimal capital relative to low-productivity firms. However, in the absence of forward-looking or earning-based constraints, under financial frictions, the investment of a low $\frac{A}{K}$ firm might not adjust, even if it is very productive.

the productivity distribution, (ii) median-productivity firms with productivity between the 50 and 75 percentile, and (iii) high-productivity firms with productivity above the 75 percentile. The data replicates the patterns suggested by the estimated policy functions. Investment propensities are increasing in $\frac{A}{K}$ and in z. As we can see, there is a positive relationship between the investment propensity and $\frac{A}{K}$ for all productivity levels, although the marginal impact of $\frac{A}{K}$ is decreasing in $\frac{A}{K}$. Moreover, propensities are more prominent for more productive firms.

6.2.6 Wealth Accumulation Policy Function

Panels (d)-(f) of Figure 3 display the estimated average derivative effect of productivity on wealth accumulation (the nonlinear propensity) $\hat{\Phi}^{g}_{t+1}(a, k, z)$. In almost all cases, the average derivative effect of productivity on savings decreases as wealth grows, consistent with the notion that self-financing is more important for firms with low wealth.

Regarding non-linearities, for a given level of capital, propensities are largest in firms that are highly productive but hold little wealth. In fact, the savings propensity to productivity shocks in firms on the upper end of the productivity distribution and the lower end of the wealth distribution is close to 1. This is, unexpected earnings shocks for highly productive but severely constrained firms are almost entirely saved, as the value of alleviating the constraint is comparatively large. As discussed earlier, this effect is reinforced by the larger persistence of productivity for highly-productive firms, which provides more incentives to wealth accumulation for productive firms as the theoretical mechanism in Moll [2014]. Consistent with the self-financing channel, the propensity decreases as we move along the wealth distribution since high-wealth firms are less constrained and have fewer incentives to save.

The savings propensity is also heterogeneous in productivity, as it is significantly lower for low-productivity firms, which are probably less constrained and have fewer incentives to save. However, at low wealth levels, even low-productivity firms save a considerable share of the earnings associated with a productivity shock when wealth is low (the propensity is 0.6). This propensity decreases to 0.2 as wealth increases.

We see similar patterns when we characterize saving propensities using the actual combination of all state variables that we see in the data in figure 4 panels (d)-(f). Propensities are positive for all firms in the data and are increasing in productivity and decreasing in wealth. Again, the propensity is higher for high-productivity firms with low levels of wealth. As we discussed above, even for high-productivity firms that can also rely on future earnings, the magnitude of the increment in the investment propensity as wealth increases is higher for high-productivity firms (see figure 4-(c)), so these firms have strong motives to save and accumulate wealth (see figure 4-(f)).

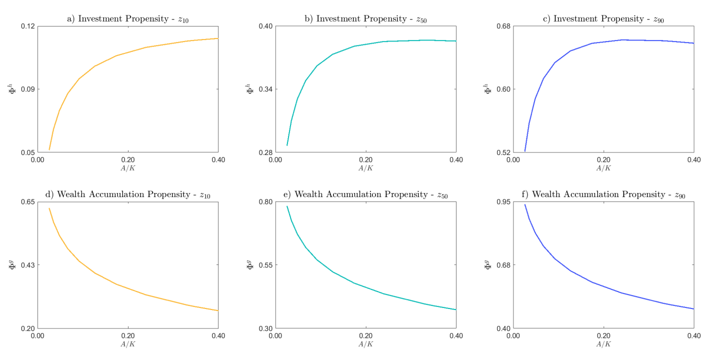

FIGURE 3: Nonlinear model: Investment and Wealth accumulation propensities

Notes: The figure exhibits the estimated derivative effect of productivity in the investment policy function (panels a, b and c) and the estimated derivative effect of productivity in the wealth accumulation policy function (panels d, e and f). The estimated model is highly non-linear, so the figure displays the marginal effect for three different values of productivity (percentiles 10th, 50th and 90th) and the wealth to capital ratio for the median value of capital.

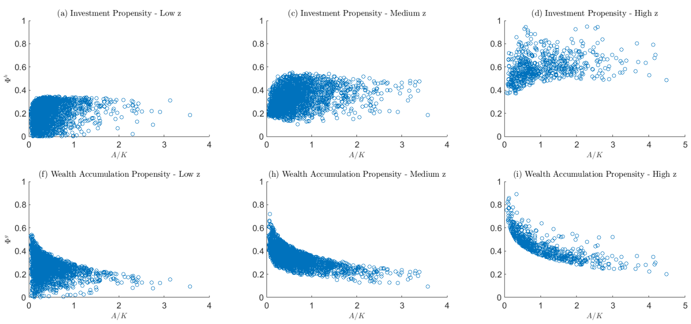

FIGURE 4: Investment and Wealth accumulation propensities in response to productivity

Notes: The figure exhibits how the investment and wealth accumulation propensity varies along the distribution of $\frac{A}{K}$ in the microdata for different productivity values. Each point represents the propensity of each particular firm evaluated at its actual value of a, k, and z. Figures (a), (b), and (c) are the investment propensities for low-, median- and high-productivity firms. Figures (d), (e), and (f) is the wealth accumulation propensities for low-, median- and high-productivity firms.

6.2.7 How does the distribution of wealth and productivity look in the data?

We can use our productivity estimates to graph the joint distribution of firm wealth and productivity in the dataset. Panel-(a) of Figure 5 features a shaded contour plot illustrating the joint distribution in 2014, the year with the largest number of firms in the data. We also present the same figure for all years in our panel data in Appendix 5.

In the figure, regions with higher density are depicted in yellow. The joint distribution has an elliptical shape rather than a circular one, indicating a positive correlation between wealth and productivity.

Figure 6 illustrates the evolution of the productivity-wealth correlation from 2007 to 2015. Throughout this period, the correlation fluctuates between 0.3 and 0.4. Interestingly, the wealth-productivity correlation appears to be procyclical, as it closely follows GDP growth (also depicted in 6). It peaks during the commodity boom, experiences a dip during the financial crisis period, rebounds in the aftermath of recovery, and subsequently declines as output stalls during 2014-2015.

We can also use our productivity estimates and the estimated production function parameters to compute the marginal product of capital (MPK) at the firm-level. Panel-(b) of Figure 5 displays the distribution of MPKs for the year 2014. There is a large degree of dispersion among firms, with values spanning from 0.02 to 3 and a standard deviation of 0.35. The mean of the distribution is 0.37, whereas the median is 0.25. The firm in the 10th percentile of the distribution exhibits an MPK of 0.08, whereas the firm in the 90th percentile exhibits an MPK of 0.8. Similar patterns are evident across other years, as detailed in Appendix 6.

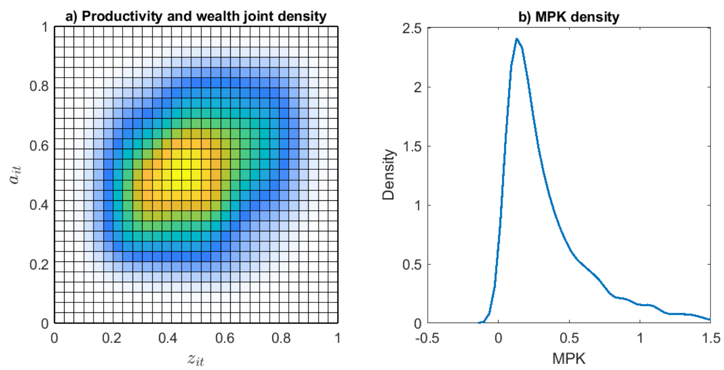

FIGURE 5: Joint distribution of productivity and wealth and MPK distribution

Notes: Nonparametric estimates of densities based on data and estimated firm-level productivity, using a Gaussian kernel. Panel a) shows the joint distribution of firm level productivity and wealth for 2014. Panel b) illustrates the distribution of the estimated MPKs in the dataset for the same year.

6.2.8 Quantitative Implications of Self-financing

For a more direct assessment of the implications of our estimated policy functions for the self-financing channel, we leverage our data and estimates to simulate the evolution of each firm's capital and wealth over time. For this, we use the estimated policy functions, assuming for simplicity that firm-level productivity remains constant over time. We initiate the model at the firm-level values of all observable state

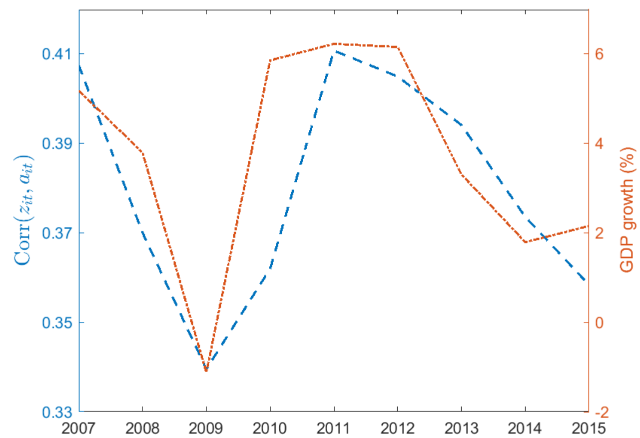

Figure 6: Evolution of productivity and wealth correlation

Notes: The figure depicts the evolution of the correlation between productivity and wealth in the data throughout the period from 2007 to 2015 (blue line). Figure also depicts Chile's per capita GDP growth for the same period (orange line).

variables in the year 2014. For the time-invariant firm-level productivity, we use the average productivity over time for each firm.[35]

Utilizing the productivity data for each firm, we can calculate the optimal capital level for each firm by equating the MPK across firms to the optimal MPK. We assume that the optimal MPK corresponds to the minimum observed (estimated) MPK. We then employ the empirical policy functions to assess the extent and speed at which the self-financing channel allows firms to converge to their optimal capital levels.

Figure 7 illustrates the joint distribution of wealth and produc-

[35]In Appendix 7, we repeat the exercise initializing the model with the 2014 firm productivity levels and show results are almost identical. Additionally, we extend our analysis by including all firms in our ten-year panel dataset. In this scenario, we initialize the firm-level state variables using the values observed in the first year each firm appears in the sample. Again, we use two approaches to define firm-level productivity: average productivity over time for each firm and productivity in the first year in which the firm is observed in the data. Again, results are very similar both qualitatively and quantitatively, underscoring the robustness of our findings.

tivity for the starting period and the simulated distribution 50 years ahead. Consistent with the implications of self-financing, the figure clearly shows that the distributions of wealth and productivity become more aligned over time, with a much stronger positive correlation than the one observed in the initial period. Panel-(a) of Figure 8 illustrates the reduction in the dispersion of MPKs over time. Meanwhile, Panel-(b) of Figure 8 shows that the standard deviation steadily decreases over time, as self-financing dampens the extent of misallocation.

Finally, we calculate various metrics to assess the aggregate efficiency gains resulting from the reallocation of resources generated by self-financing.

Our baseline measure compares the sum of individual firm-level outputs, a representation of aggregate output, across two economies. In the first economy, firm-level capital evolves as dictated by the estimated empirical policy functions. In contrast, the second economy presents a counterfactual scenario where firm-level capital grows at the same rate across all firms, such that the aggregate levels of capital are identical to the one in the first economy (i.e., each firm grows at the growth rate of the aggregate capital in the first economy). Consequently, while both economies have the same aggregate levels of capital at every moment in time, their distributions across firms are different. Notably, reallocation occurs only in the first economy, as the relative inputs across firms remain constant in the second economy, as determined by the initial distribution.

Panel-(d) of Figure 8 presents the ratio in the aggregate production level between the two economies. Reallocation through self-financing enhances aggregate productivity by up to twenty-seven percent compared to the counterfactual without reallocation. It is worth noting that this exercise is analogous to calculating gains in aggregate Total Factor Productivity (TFP) as the residual of an aggregate production function after aggregating individual outputs (the evolution of $\frac{TFP_t}{TFP_0}$). Therefore, the initial allocation of capital represents a loss of twenty-seven percent in TFP relative to the optimal allocation.

Panel-(b) of Figure 9 displays an alternative measure for calculating aggregate productivity gains. It depicts the weighted average of individual productivities, where weights are given by each firm's capital relative to total capital, following the measure proposed by Moll

[2014]. This measure suggests that aggregate productivity gains from reallocation derived from self-financing can be as large as 40%. It is important to recall that our production function estimates show that firms in our sample operate with decreasing returns to scale. Therefore, the theoretical result presented by Moll [2014], which expresses aggregate productivity as a weighted average of firm-level productivities, does not directly apply in our context. This is because the theoretical framework in Moll [2014] assumes constant returns to scale and linear policy functions. Thus, we consider this measure only as an approximation to aggregate TFP.

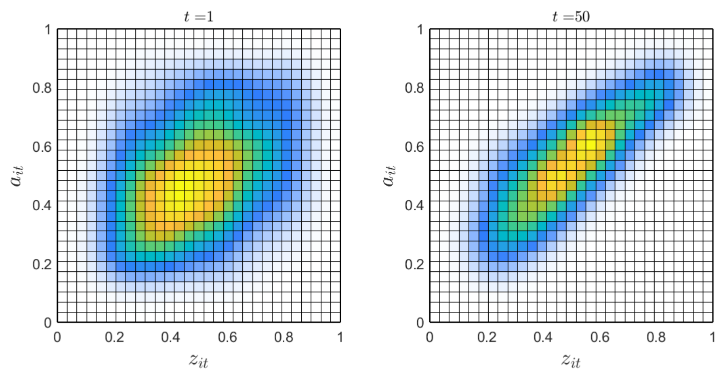

FIGURE 7: Simulated joint distribution of productivity and wealth

Notes: The figure illustrates the estimated joint distribution of firm-level productivity and wealth for the first and 50th years of the simulation, respectively. The state variables' values for the first period correspond to those observed in the data in 2014, except for the values for firm-level productivity, which are constant across the simulated years and equal to the within-firm average productivity in the time series. This is the reason why the joint distribution in the first panel is different from the one in the first panel of Figure 5.

6.2.9 Quantitative Implications: Speed of convergence in MPKs

Our previous results highlight the empirical strength of the self-financing channel in mitigating misallocation and narrowing the disparities in MPK across firms. However, the length of the period over which convergence occur remains a critical question. Figure 10 depicts the speed of convergence in the MPKs of constrained and unconstrained firms, in the spirit of the exercise in Banerjee and Moll [2010].

The exercises uses the data and our estimates of firm productivity and the production function parameters to calculate the initial levels of MPK for two firms that have the same level of initial productivity but different levels of initial wealth and capital. We then use the estimated policy functions to simulate the evolution of their capital, labor, and wealth across time, assuming, as in the previous section, that productivity is constant and there are no additional shocks.

Results are presented in Figure 10. For each row, the graphs plot the evolution of the MPK for a firm that starts on the lower end of the wealth distribution (10th percentile) vis-a-vis firms with the same level of productivity z but larger levels of initial wealth (50th percentile in the first column, 75th percentile in the second, 90th in the third). We report the convergence in MPKs between a constrained and unconstrained firm for three different productivity scenarios. The first row depicts firms in the 10th percentile of the productivity distribution, while the 50th and 90th productivity deciles are presented in the second and third rows.

Consistent with our previous results, low-wealth, constrained firms increase their capital stock across time, such that the MPK converges towards that of firms with similar firm productivity z but higher levels of initial wealth a_0. Convergence, however, is relatively slow, and marginal productivity gaps persist for decades. For example, across all three productivity levels, the marginal product of capital in a firm with initial wealth in the 10th percentile of the wealth distribution is close to three times larger than in a firm in the 90th wealth percentile. While this gap closes steadily across the years, marginal products in low-wealth firms are still at least twice as large as those of high-wealth firms after one decade. The speed of convergence in our data is much slower than in Banerjee and Moll [2010], where, for a similar initial gap, differences in marginal product between constrained and unconstrained firms vanish in less than a decade. For example, even for firms in the 90th percentile of the productivity distribution, convergence in the marginal product of capital between firms in the 10th and 90th wealth percentiles takes almost 30 years, although half of the initial gap disappears after ten years.

Therefore, our results indicate that while the self-financing channel plays an important role in reducing productivity gaps and the extent of misallocation in this context, it seems unable to offset the persistence of significant differentials in MPKs over the medium term.

7 Conclusions

We provide an empirical analysis of wealth accumulation and investment dynamics in firms that operate under financial frictions and how these decisions relate to the unobservable firm's productivity process. We present a novel framework, robust to financial frictions, to jointly model and estimate the firm wealth accumulation dynamics, its investment decisions, and the unobservable productivity process.

Our results, using Chilean manufacturing data, show that the productivity process seems to be largely non-linear, with larger persistence for more productive firms, while persistence can change significantly in the face of extreme events. This differs significantly from the assumptions for the productivity process used in the structural macro literature.

We use our setup to provide a detailed analysis of the firm´s policy functions. We find that the behavior of firms is consistent with the predictions of macro models with financial frictions, although there are significant non-linearities. Results suggest that both collateral and earnings-based constraints are present in the data. We also find support for the existence of an active self-financing channel, although its ability to overcome misallocation appears to be limited. This novel microeconomic evidence, as well as the methodology used to generate it, can provide support and guidance to the quantitative macroeconomic models that have studied the role of financial frictions and their aggregate implications.

References

Daniel A Ackerberg, Kevin Caves, and Garth Frazer. Identification properties of recent production function estimators. *Econometrica*, 83(6):2411–2451, 2015.

Alvaro Aguirre. Contracting institutions and economic growth. *Review of Economic Dynamics*, 24:192–217, 2017.

Heitor Almeida, Murillo Campello, and Michael Weisbach. The cash flow sensitivity of cash. *Journal of Finance*, 59:1777–1804, 2004.

Isaiah Andrews, Matthew Gentzkow, and Jesse M Shapiro. Transparency in structural research. *Journal of Business & Economic Statistics*, 38(4):711–722, 2020.

Manuel Arellano and Stephen Bond. Some tests of specification for panel data: Monte carlo evidence and an application to employment equations. *The Review of Economic Studies*, 58(2):277–297, 1991.

Manuel Arellano and Stéphane Bonhomme. Nonlinear panel data estimation via quantile regressions. *Econometrics Journal*, 19(3): C61–C94, 2016.

Manuel Arellano, Richard Blundell, and Stéphane Bonhomme. Earnings and consumption dynamics: a nonlinear panel data framework. *Econometrica*, 85(3):693–734, 2017.

Abhijit Banerjee and Benjamin Moll. Why does misallocation persist? *American Economic Journal: Macroeconomics*, 2:189–206, 2010.

Saki Bigio and Jennifer La'o. Distortions in production networks. *The Quarterly Journal of Economics*, 135(4):2187–2253, 2020.

Richard Blundell and Stephen Bond. Initial conditions and moment restrictions in dynamic panel data models. *Journal of econometrics*, 87(1):115–143, 1998.

Wyatt Brooks and Alessandro Dovis. Credit market frictions and trade liberalizations. *Journal of Monetary Economics*, 111:32–47, 2020.

Francisco J Buera and Yongseok Shin. Self-insurance vs. self-financing: A welfare analysis of the persistence of shocks. *Journal of Economic Theory*, 146(3):845–862, 2011.

Francisco J Buera and Yongseok Shin. Financial frictions and the persistence of history: A quantitative exploration. *Journal of Political Economy*, 121(2):221–272, 2013a.

Francisco J Buera and Yongseok Shin. Financial frictions and the persistence of history: A quantitative exploration. *Journal of Political Economy*, 121(2):221–272, 2013b.

Francisco J Buera, Joseph P Kaboski, and Yongseok Shin. Finance and development: A tale of two sectors. *The American Economic Review*, 101(5):1964–2002, 2011.

Francisco J Buera, Joseph P Kaboski, and Yongseok Shin. Entrepreneurship and financial frictions: A macro-development perspective. *Annual Review of Economics*, 2015.

Andrea Caggese and Vicente Cuñat. Financing constraints, firm dynamics, export decisions, and aggregate productivity. *Review of Economic Dynamics*, 16(1):177–193, 2013.

Santiago Camara and Maximo Sangiacomo. Borrowing constraints in emerging markets. *arXiv preprint arXiv:2211.10864*, 2022.

Sylvain Catherine, Thomas Chaney, Zongbo Huang, David Alexandre Sraer, and David Thesmar. Quantifying reduced-form evidence on collateral constraints. *forthcoming Journal of Finance*, 2018.

Tiago V Cavalcanti, Joseph P Kaboski, Bruno S Martins, and Cezar Santos. Dispersion in financing costs and development. Technical report, National Bureau of Economic Research, 2021.

Russell W Cooper and John C Haltiwanger. On the nature of capital adjustment costs. *The Review of Economic Studies*, 73(3):611–633, 2006.

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

Savita Diggs and Joseph Kaboski. Smoothing financial frictions for structural change. *Policy Brief, STEG Pathfinding Papers*, 5, 2022.

Ulrich Doraszelski and Jordi Jaumandreu. R&d and productivity: Estimating endogenous productivity. *The Review of Economic Studies*, 80(4):1338–1383, 2013.

Ulrich Doraszelski and Jordi Jaumandreu. Measuring the bias of technological change. *Journal of Political Economy*, 126(3):1027–1084, 2018.

Thomas Drechsel. Earnings-based borrowing constraints and macroeconomic fluctuations. *DP16975*, 2022.

Steven Fazzari, R Glenn Hubbard, and Bruce C Petersen. Financing constraints and corporate investment. Technical report, National Bureau of Economic Research, 1987.

Vito D Gala, Joao F Gomes, and Tong Liu. Investment without q. *Journal of Monetary Economics*, 116:266–282, 2020.

Amit Gandhi, Salvador Navarro, and David A Rivers. On the identification of gross output production functions. *Journal of Political Economy*, 128(8):2973–3016, 2020.

Yingyao Hu and Susanne M Schennach. Instrumental variable treatment of nonclassical measurement error models. *Econometrica*, 76 (1):195–216, 2008.

Yingyao Hu, Guofang Huang, and Yuya Sasaki. Estimating production functions with robustness against errors in the proxy variables. *Journal of Econometrics*, 215(2):375–398, 2020.

Victoria Ivashina, Luc Laeven, and Enrique Moral-Benito. Loan types and the bank lending channel. *Journal of Monetary Economics*, 126: 171–187, 2022.

Joseph P Kaboski. Financial frictions, financial market development, and macroeconomic development. *Unpublished Working Paper*, 2021.

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

Chen Lian and Yueran Ma. Anatomy of corporate borrowing constraints. *The Quarterly Journal of Economics*, 136(1):229–291, 2020.

Enrique G Mendoza and Vivian Z Yue. A general equilibrium model of sovereign default and business cycles. *The Quarterly Journal of Economics*, 127(2):889–946, 2012.

Virgiliu Midrigan and Daniel Yi Xu. Finance and misallocation: Evidence from plant-level data. *The American Economic Review*, 104 (2):422–458, 2014.

Benjamin Moll. Productivity losses from financial frictions: Can self-financing undo capital misallocation? *American Economic Review*, 104(10):3186–3221, 2014.

Whitney K Newey and James L Powell. Instrumental variable estimation of nonparametric models. *Econometrica*, 71(5):1565–1578, 2003.

Steven Olley and Ariel Pakes. The dynamics of productivity in the telecomunications equipment industry. *Econometrica*, 64:1263–97, 1996.

Tim Opler, Lee Pinkowitz, Rene Stultz, and Rohan Williamson. The determinants and implications of corporate cash holdings. *Journal of Financial Economics*, 52:3–46, 1999.

Ajay Shenoy. Estimating the production function under input market frictions. *Review of Economics and Statistics*, pages 1–45, 2020.

Lucciano Villacorta. Estimating country heterogeneity in capital-labor substitution using panel data. Technical report, Mimeo, 2018.

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

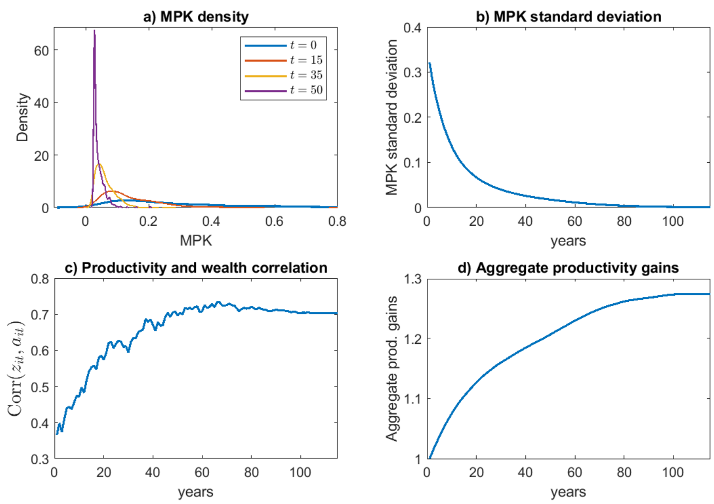

Figure 8: Simulated quantitative implications of self-financing

Notes: The figure shows several results derived from simulations of the evolution of each firm's capital and wealth over time. Simulations were conducted using the estimated policy functions and time-invariant firm-level productivity and labor. For productivity we use the within-average productivity in the time series for each firm, and for labor firm's initial values were kept constant. First period values for capital, wealth and labor correspond to those observed in the data in 2014. Wealth accumulation policy function were shocked by i.i.d. normal disturbances. Panel a) shows the density of MPKs in the initial year and its evolution 15, 35 and 50 years ahead. Panel b) shows the evolution of the of the MPKs' standard deviation across the years. Panel c) illustrates the evolution of the correlation between firm-level productivity and wealth through the years. Panel d) depicts the evolution of our baseline measure of aggregate efficiency gains: the ratio between aggregate output in the economy with reallocation and aggregate output in the economy without reallocation.

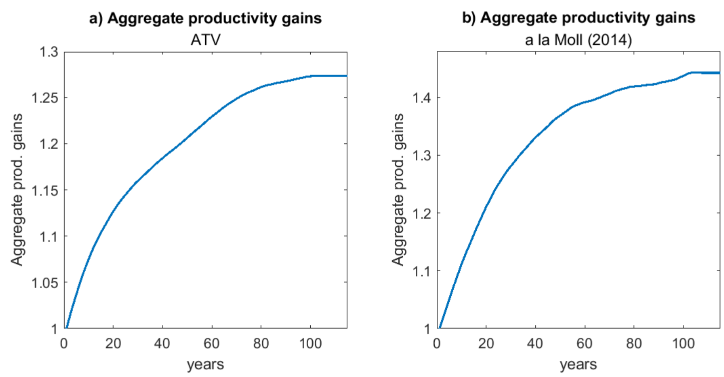

Figure 9: Simulated aggregate productivity gains

Notes: The figure compares the evolution of two measures of aggregate efficiency gains using the simulated data. Panel a) depicts the evolution of our baseline measure of aggregate efficiency gains: the ratio between aggregate output in the economy with reallocation and aggregate output in the economy without reallocation. Panel b) exhibits the evolution of an alternative measure, which is the weighted average of individual productivities with weights given by each firm's capital relative to the total capital, following the measure proposed by Moll [2014].

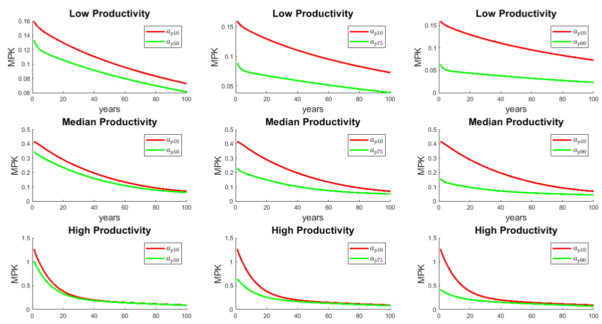

FIGURE 10: Convergence in the marginal product of capital across firms

Notes: The figure exhibits the simulated evolution of the marginal product of capital for firms with different levels of initial productivity and wealth. Low-wealth firms (10th percentile) are depicted in red, while high-wealth firms (50th percentile in column 1, 75th in column 2, and 90th in column 3) are depicted in green. The first row presents firms in the 10th percentile of the productivity distribution, while the second and third rows present figures in the 50th and 90th productivity deciles. The simulation uses the estimated production function and investment and wealth accumulation policy functions, holding firm productivity constant.