Institute for International Economic Policy Working Paper Series
Elliott School of International Affairs
The George Washington University

Firm Learning and Growth

IIEP-WP-2015-5

Olga A. Timoshenko
George Washington University

February 2015

Institute for International Economics Policy
1957 E St. NW, Suite 502
Voice: (202) 994-5320
Fax: (202) 994-5477
Email: iiep@gwu.edu
Web: www.gwu.edu/~iiep

Title: Firm Learning and Growth

Key Words: Firm dynamics, growth, learning

JEL Codes: F12, F14

Firm Learning and Growth

Costas Arkolakis Theodore Papageorgiou Olga A. Timoshenko *

Yale University and NBER McGill University George Washington University

December 29, 2014

Abstract

We develop a general equilibrium model of firm growth with learning about unobserved demand. Our framework introduces learning (Jovanovic, 1982) into a monopolistically competitive environment with firm productivity heterogeneity, á la Melitz (2003). The model correctly predicts that firm growth rates decrease with age, holding size constant, a fact that models focusing on idiosyncratic productivity shocks have difficulty matching. We calibrate the model using Colombian plant-level data and find that it matches growth and survival patterns well. Unlike the standard Melitz setup the model with learning is no longer efficient, leaving room for welfare improving policies. We illustrate how subsidies to the fixed costs of young firms can be welfare enhancing: they allow young firms to avoid early exit and thus, benefit consumers through access to a larger number of varieties.

*Arkolakis: Yale University and NBER, 28 Hillhouse Av., Room 312 New Haven, CT, 06520, costas.arkolakis@yale.edu. Papageorgiou: McGill University, 855 Sherbrooke St W., Room 519 Montreal, QC, H3A 2T7, theodore.papageorgiou@mcgill.ca. Timoshenko: George Washington University, 2115 G St. NW, Suite 340, Washington, DC, 20052, timoshenko@gwu.edu. We are grateful to Manolis Galenianos, Illenin Kondo, Giuseppe Moscarini, Emanuel Ornelas, Andres Rodriguez-Clare, Kim Ruhl, Erick Sager, James Tybout as well as various seminar participants. This research is supported by NSF grant SES-0921673. This work was supported in part by the facilities and staff of the Yale University Faculty of Arts and Sciences High Performance Computing Center. All remaining errors are ours.

1 Introduction

Young firms are smaller than older ones but grow faster. Even conditional on firm size, age appears to play an important role on firm growth as shown by Evans (1987) and recently discussed at length by Haltiwanger et al. (2011). To stimulate the growth of young firms, policymakers often provide subsidies throughout the firms' initial years of operation (see for instance the recommendations of Commision (2010) and the analysis and recommendations in chapter 5 in the OECD (2013) report). Traditional structural models of firm dynamics that focus on idiosyncratic productivity shocks, in the tradition of Hopenhayn (1992), abstract from the role of age for firm growth and are thus inappropriate to evaluate such policies.

In this paper we develop a general equilibrium model of firm growth to evaluate the welfare impact of subsidies to young firms. To do so, we suitably adapt the standard learning mechanism of Jovanovic (1982) into the monopolistic competition framework of Melitz (2003). In our model firms enter the market small and they are uncertain about the demand their product faces. Over time, as firms observe sales realizations, they may grow large if they have a very successful product or, if not, they shrink and may even exit the market. We calibrate our setup using Colombian plant-level data and find that unlike the canonical Melitz (2003) model, our economy is not efficient. We illustrate how subsidies to the fixed cost of production of young firms can be welfare enhancing: They allow young firms (some with potentially large product appeal) to avoid early exit and, thus, benefit consumers through access to a larger number of varieties.

More specifically, in our setup the demand for a firm's product is uncertain and demand realizations in each period are determined by an unobserved idiosyncratic firm demand component, namely the firm's product appeal in a market, plus a noise component. Each firm decides the quantity produced in the beginning of each period depending on its prior for its product appeal. Once demand is realized, the price adjusts to clear the firm's product market and, using the realized price, the firm revises its posterior for its product appeal. A firm that experiences higher demand than initially expected revises upward its belief, expands production and grows in size. A firm that experiences lower demand than expected cuts back on production, and it may even find it optimal to exit the market.

In our setup two different channels lead to faster growth for younger firms: beliefs updating and selection. First, younger firms face large uncertainty about their product appeal, they revise their beliefs relatively more, compared to firms that are older and better informed. In equilibrium, they are expected to grow faster compared to older firms. This holds even conditional on a firm's size, thus yielding the conditional age and size dependence of firms' growth rates that we observe in the data. Second, the endogenous selection of firms strengthens the above result: since younger firms face larger uncertainty, they are much more likely to exit the market if they face a low demand realization. This selection implies that the measured growth rate of surviving young firms is even higher than the respective growth rate of older firms.

We next investigate the quantitative implications of our setup, by calibrating the model to match moments from a panel of Colombian plant-level data, assuming that in our model each firm owns a unique plant. Guided by our theoretical results, we identify the importance of learning by targeting the impact of age on growth rates, conditional on size. In addition, the model correctly predict moments that are not targeted in the calibration, such as the annual survival rate and the impact of size on growth rates, conditional on age. Given that learning also affects the exit behavior of firms, it is reassuring that our model is also able to match the survival rate.

Armed with the calibrated model, we are able to evaluate policies that subsidize young firms. While the entry and exit decisions are always optimal from the individual firm's point of view, they may not be optimal from an aggregate welfare perspective. For instance, when a firm exits (enters) it does not internalize the loss (addition) of a variety for the consumer. In addition, when a firm enters or exits it does not consider the negative impact its decision has on the profits of incumbent firms through the monopolistic distortion. In the canonical Melitz (2003) model these two effects exactly cancel out leaving no role for welfare improving policies. In our calibrated setup with learning however, in equilibrium, there is an inefficiently low mass of firms and therefore varieties available to consumers.

We show that a subsidy to the fixed operating costs increases both the equilibrium mass of firms and welfare. We experiment with a range of different subsidies and find that, for instance, a 40% subsidy to the fixed cost of production applied to young firms leads to up to a 20% increase in the mass of operating firms in equilibrium. The subsidy is financed through lump-sum taxation. The welfare gains are of similar order of magnitude to the cost of the subsidy: a subsidy equal to 0.36% of GDP leads to a 0.23% increase in aggregate consumption (consumption multiplier of 0.64). A lower subsidy leads to an even higher consumption multiplier: for instance, a subsidy equal to 0.10% of GDP leads to a 0.08% increase in aggregate consumption (consumption multiplier of approximately 0.8).

Our paper contributes to a new, but growing literature on the quantitative implications of learning on firm growth. Eaton et al. (2012) consider a model where firms actively learn their product appeal by forming new matches with buyers. Their learning framework is flexible enough to match a number of patterns in the buyer-seller exporting data. Abbring and Campbell (2003) also develop a structural model of firm learning, which they estimate with data on Texan bars. Both these frameworks are much richer in the specification of the learning mechanism, but they use a partial equilibrium model and as such it cannot be used for macro policy evaluation, a key focus of our analysis. Ruhl and Willis (2014) modify a standard model with idiosyncratic productivity shocks by specifying a foreign demand function that increases with firm age in the market. Abstracting from general equilibrium effects, they illustrate that this addition allows the model to better match the entry and growth patterns of exporters. Albornoz et al. (2012) consider a two-period model where a firm is uncertain about its demand in the first period and uncertainty is resolved by incurring a fixed cost. Finally, closer to our approach, Timoshenko (2014) develops a general equilibrium model of learning in the context of multi-product firms and shows that such a framework can predict well the age dependence of product switching among exporters.

Models that focus on idiosyncratic productivity shocks (Hopenhayn (1992), Luttmer (2007), Arkolakis (2011)) cannot explain the dependence of growth rates on age, conditional on size.[1] The reason for this failure is that growth is based on an underlying Markov process. This assumption implies that all firms of the same size have the same growth profile, which is independent of their age.[2] Financial constraints together with idiosyncratic productivity advances, as in Cooley and Quadrini (2001), can explain age dependence (even conditional on size) if the entry of new firms is at high productivity levels. Our approach does not require idiosyncratic productivity advances to generate that age dependence.[3]

[1]Hopenhayn (1992) notes that "since size is a sufficient statistic for [productivity], age has no extra predictive role" (page 1141).

[2]A similar reasoning applies to the Klette and Kortum (2004) approach. However, even if the productivity process is not a simple Markov process, but depends on a finite number of past realizations as well, the two approaches give distinct predictions regarding the variance of sales as a function of age. The learning explanation proposed here implies that the variance of sales declines with age, as large firms become better informed regarding the demand they face. However, productivity models have no such implication. In general, in a learning model, the dependence of firm actions on past realizations does not erode away as the firm ages. This implication is used in Ericson and Pakes (1998) to distinguish learning from a model of productivity advances, in which the state dependence has ergodic characteristics.

[3]In a multi-market setup, however, one can distinguish between our setup and the models of Jovanovic (1982) and Cooley and Quadrini (2001). Since a firm can have different demand realizations across markets, our framework is consistent with differential behavior of the same firm across the various markets in which it operates. For instance, it is possible that a firm expands production in a given market, while choosing to exit another one. If the financial constraints are important or if the firm is learning about its underlying cost structure, then its behavior is restricted to be the same across markets. Furthermore, the evidence in Eaton et al. (2008) and Albornoz et al. (2012) suggest that the age of a firm in a market is inversely related to firm growth and probability of exit in that market, as implied by our model.

The rest of the paper is organized as follows: Section 2 introduces a model that suitably adapts the standard learning mechanism of Jovanovic (1982) into the monopolistic competition framework of Melitz (2003). Section 3 calibrates the model using a panel of Colombian plant-level data, whereas Section 4 investigates the welfare implications of a subsidy of the fixed operating costs of young firms. Section 5 concludes.

2 The Model

This section describes a model which introduces a demand-learning mechanism into a dynamic setup. The learning mechanism is similar to that of Jovanovic (1982) and firms must learn about their unobserved demand level which is subject to transient preference shocks. Firms operate in a monopolistically competitive environment as in Melitz (2003).

2.1 Environment

Time is discrete and denoted by t. The economy is populated by a continuum of consumers of mass L. Each consumer derives utility from the consumption of a composite good, C_t, according to the utility function

$$U = E \sum_{t=0}^{+\infty} \beta^t \ln \left(C_t \right),$$

where β is the discount factor. The composite good consists of the consumption of a continuum of differentiated varieties $c_t(\omega)$, aggregated using a constant elasticity of substitution (CES) aggregator with elasticity of substitution $\sigma > 1$

$$C_t = \left(\int_{\omega \in \Omega_t} \left(e^{a_t(\omega)} \right)^{\frac{1}{\sigma}} c_t \left(\omega \right)^{\frac{\sigma-1}{\sigma}} d\omega \right)^{\frac{\sigma}{\sigma-1}},$$

where Ω_t is the continuum of differentiated varieties available at time t; $a_t(\omega)$ is the demand shock at time t for variety $\omega \in \Omega_t$. Consumers are endowed with one unit of labor that they inelastically supply to the market to receives a wage, w_t, in return. They also own an equal share of domestic firms.

Each good is supplied by a monopolistically competitive firm. These firms maximize the expected present discounted value of profits and have the same discount factor as consumers, β. The realization of the demand variable $a_t(\omega)$ for each firm is determined by a time invariant component, $\theta(\omega)$, and a shock, $\varepsilon_t(\omega)$, and is given every period by

$$a_t(\omega) = \theta(\omega) + \varepsilon_t(\omega), \; \varepsilon_t(\omega) \sim N\left(0, \sigma_\varepsilon^2\right) \text{ i.i.d.}$$

The time-invariant demand parameter, $\theta(\omega)$, can be interpreted as the true underlying demand for the firm's product, its product appeal, and it is unobserved by the firm. This is subject to transient preference shocks, $\varepsilon_t(\omega)$, which are independent over time and across firms.

Following Chaney (2008), we assume that every period there is an exogenous mass of potential entrants J. Entrants draw their unobserved demand parameter, $\theta(\omega)$, from a normal distribution with mean $\bar{\theta}$ and variance σ_θ^2. There is no sunk cost of entry.

Firms exit the market either exogenously with probability δ or endogenously if the sum of their discounted expected profits is less than zero. If the firm stays in the market, it decides the quantity to be produced, $q_t(\omega)$. Output is linear in labor, $l_t(\omega)$, which is the only factor of production, and there is a fixed cost of production, f, measured in the units of labor. Firms are heterogeneous in their productivity level, $e^{z(\omega)}$, which is drawn upon entry and is time-invariant (in Appendix A we relax this assumption and allow productivity to change over time). Unlike the demand parameter, $\theta(\omega)$, each firm's productivity level is observed by the firms.

The sequence of actions is the following: At the beginning of each period, each incumbent firm decides whether to stay in the market and what quantity to produce or to endogenously exit. Next, potential entrants draw their productivity, $e^{z(\omega)}$, and their (unobserved) demand parameter, $\theta(\omega)$, and decide whether to sell or exit. Those that decide to sell pay the fixed cost of production. Subsequently, the firm decides the quantity produced and production takes place. Next, the demand shock, $a_t(\omega)$, is realized and the price of each good, $p_t(\omega)$, adjusts so that the good's market clears. The firm observes the price, infers the underlying demand realization and updates its belief for its product appeal. Finally, firms exit the market exogenously with probability δ.

2.2 Consumer Demand

Each consumer chooses the consumption levels, $c_t(\omega)$, to minimize the cost of acquiring C_t taking into account the prices of varieties, $p_t(\omega)$, as well as his income. The resulting demand equation for each variety is given by

$$q_t(\omega) = e^{a_t(\omega)} \frac{(p_t(\omega))^{-\sigma}}{P_t^{1-\sigma}} Y_t, \tag{1}$$

where Y_t is the aggregate spending level and is given by $Y_t = w_t L + \Pi_t$, while Π_t is total profits of firms. P_t is the aggregate price index associated with the consumption of the composite good C_t and is given by

$$P_t^{1-\sigma} = \int_{\omega \in \Omega_t} e^{a_t(\omega)} (p_t(\omega))^{1-\sigma} \, d\omega.$$

To economize on notation, in what follows we drop the variety index ω and also, since we will be focusing on a stationary equilibrium, the t subscript. Future variables are denoted with a prime.

Belief Updating

The firm, by observing the market-clearing price for a given quantity decision, q, is able to infer the demand realization, a, which is informative about the unobserved demand parameter, θ. Using Bayes' rule, the firm updates its beliefs regarding θ: the posterior belief of a firm that has observed n signals with mean $\bar{a}$, is given by a normal distribution with mean

$$\mu_n = \frac{\sigma_\varepsilon^2}{\sigma_\varepsilon^2 + n\sigma_\theta^2}\bar{\theta} + \frac{\sigma_\theta^2}{\sigma_\varepsilon^2 + n\sigma_\theta^2} n\bar{a} \ ,$$

and variance

$$\nu_n^2 = \frac{\sigma_\theta^2 \sigma_\varepsilon^2}{\sigma_\varepsilon^2 + n\sigma_\theta^2}.$$

Thus, the belief of the firm regarding the realization of its demand shock, a, follows $N\left(\mu_n, \nu_n^2 + \sigma_\varepsilon^2\right)$. An entrant knows the distribution from which the demand parameter is drawn, and that distribution serves as the prior. Notice that in the long run, upon observing infinitely many signals, the posterior belief converges to a degenerate distribution centered at $\bar{a}$ and $\lim_{n\to\infty} \bar{a} = \lim_{n\to\infty} \frac{\sum_{i=0}^{n} a_i}{n} = E(a) = \theta$. Thus, in the long run, firms learn their product appeal level, while in the short run their knowledge can be summarized by the two state variables: the number of observed shocks n, and the mean of the observed shocks $\bar{a}$.

2.3 Firm Optimization

In describing the firm's problem, we focus on the stationary equilibrium, which exists as long as $\delta > 0$. For simplicity, we suppress the time notation and denote next period variables with a prime. Each firm has to make two decisions every period: whether to stay in a market or exit and how much to produce, if it chooses to stay. When the firm decides whether to produce or exit, it takes into account not only that period's profits, but also the value of learning by observing demand realizations.

More specifically, learning allows the firm to make more informed decisions in the future about how much to produce and whether to exit or not. When deciding how much to produce however, learning considerations are not a factor, since the demand realization is independent of how much the firm produces. Therefore its quantity choice is a static one.[4] We begin by analyzing the static quantity decision and then consider the entry and exit decision.

Quantity Decision

The static per-period profit maximization problem is given by

$$\max_{q} E_{a|n,\bar{a}}\left[p\left(a\right)q\right] - w\frac{q}{e^{z}} - wf, \tag{2}$$

subject to the consumer inverse demand

$$p\left(a;q\right) = \left(\frac{e^{a}Y}{q}\right)^{\frac{1}{\sigma}} P^{\frac{\sigma-1}{\sigma}}. \tag{3}$$

Substituting the inverse demand constraint (3) into the static maximization problem (2) and taking the first order conditions leads to the optimal quantity choice given by

$$q\left(z,b\right) = \left(\frac{\sigma-1}{\sigma}\right)^{\sigma}\left(\frac{be^{z}}{w}\right)^{\sigma}\frac{Y}{P^{1-\sigma}}\ , \tag{4}$$

where we define

$$b \equiv E_{a|\bar{a},n}(e^{\frac{a}{\sigma}}) = \exp\left\{\frac{\mu_n}{\sigma} + \frac{1}{2}\left(\frac{\nu_n^2+\sigma_\varepsilon^2}{\sigma^2}\right)\right\},$$

i.e. the firm's belief regarding (e^{a}). Notice from equation (4) that firms which are more productive (have higher z) and firms which have higher beliefs regarding θ (have higher b) sell more. Using equation (4), Proposition 1 below established the main implication of the model: the age dependence of growth rates.

[4] See Bergemann and Välimäki (2000) and Eaton et al. (2012) for models of active learning and experimentation.

Proposition 1 Assume there is no endogenous exit. For a given z, conditional on firm size, q_t, the expected growth rate of young firms is higher than the growth rate of older firms.

Proof: See Appendix B.

Note that with the endogenous exit the observed growth rate is even larger for young firms near the exit threshold than what is implied by the proposition. The reason is that the firms that observe low sales realizations may prefer to exit the market, especially if there is large uncertainty for their true demand, rather than stay in and record low or negative growth rates.

The above result formalizes the intuition of Evans (1987) who argued that learning implies that firm growth rates decrease with age, holding size constant. In the calibration of the model presented in the next section, motivated by the above result, we use the age effect on firm growth rates, conditional size to identify the importance of learning.

Substituting the firm's quantity choice, equation (4), into the consumer inverse demand, equation (3), gives the market clearing price

$$p\left(a,z,b\right)=\frac{\sigma}{\sigma-1}\frac{w}{e^{z}}\frac{\left(e^{a}\right)^{\frac{1}{\sigma}}}{b}.$$

The price depends on the firm's productivity, z and belief, b, as well as the realization of the demand shock, a. Taking expectation with respect to a, the expected market clearing price is given by

$$Ep=\frac{\sigma}{\sigma-1}\frac{w}{e^{z}}\ .$$

Thus, in expectation, price is a constant mark-up over marginal cost.

Substituting the optimal quantity and price into equation (2), the expected per-period firm profits are given by

$$E\pi\left(z,b\right)=\frac{\left(\sigma-1\right)^{\sigma-1}}{\sigma^{\sigma}}b^{\sigma}\left(\frac{e^{z}}{w}\right)^{\sigma-1}\frac{Y}{P^{1-\sigma}}-wf. \tag{5}$$

Notice that, since $\sigma > 1$, profits are a convex function of firm beliefs.

Exit and Entry Decision

As discussed above, a firm decides whether to continue paying the fixed cost and stay in the market or exit, taking the maximized expected profits, $E\pi(z, b)$, as given. Thus, the value of a firm of age n, with productivity z and beliefs b, is given by

$$V(z, b, n) = \max\left\{E\pi(z, b) + \beta(1-\delta)E_{b'|b,n}V(z, b', n+1), 0\right\}, \tag{6}$$

where the value of exit is normalized to zero.

We now consider the entry decision. Potential entrants pay the fixed per-period cost and start to sell in the initial period if the value of entry is greater than zero, or otherwise immediately exit. All potential entrants share the same prior about the demand level, but vary in terms of their productivity draw, e^z. Thus the entry condition determines a productivity threshold value $\underline{z}$, such that an entrant that draws productivity $z < \underline{z}$, exits immediately. This threshold is implicitly given by

$$V(\underline{z}, b^e, 0) = 0, \tag{7}$$

where b^e is the potential entrant's belief and is given by

$$b^e = \exp\left\{\frac{\bar{\theta}}{\sigma} + \frac{1}{2}\left(\frac{\sigma_\theta^2 + \sigma_\varepsilon^2}{\sigma^2}\right)\right\}.$$

2.4 Equilibrium

Denote the equilibrium mass of operating firms at every period by M. Each potential entrant draws its productivity e^z from the following Pareto distribution

$$f(e^z) = \frac{\xi}{(e^z)^{\xi+1}}, \tag{8}$$

where $\xi > 0$. They also draw their unobserved, demand parameter θ from $N\left(\bar{\theta}, \sigma_\theta^2\right)$. As discussed above, based on the productivity realization, each potential entrant decides whether to enter the market and begin selling or exit immediately.

To study the equilibrium, we normalize the wage rate, w, to 1. The stationary equilibrium is described by the entry productivity threshold $\underline{z}$, the aggregate price index P, the aggregate expenditure level Y, the mass of operating firms M, the probability mass function $m(z, b, n)$, firms' optimal quantity choice q, firms' optimal entry and exit decisions, consumers' optimal consumption choice c such that

1. Consumers are optimizing: given equilibrium values, c satisfies demand equation (1).

2. Firms are optimizing: given equilibrium values, q, $\underline{z}$ and the exit thresholds solve the firm's optimization problem in equations (2) and (6).

3. Goods market clears: $Y = L + \Pi$.

4. The probability mass function of active firms, $m(z, b, n)$, delivers the aggregate mass of firms, M, i.e. $M = \sum_n \int \int m(z, b, n)\, dzdb$.

The computation of the stationary equilibrium is based on the method developed in Timoshenko (2014) and the setup is briefly outlined in Appendix C.

3 Calibration

For the calibration we assume that the data has been generated from the steady state of our model and match theoretical moments to their empirical counterparts. In the next section we use the calibrated model to investigate the welfare implications of a subsidy on young firms' per-period fixed operating costs.

We use the Colombian plant-level data collected in the DANE survey (see Roberts (1996)). In our calibration we treat a plant in the data as a firm in the model. The data cover all the manufacturing plants in Colombia with 10 or more employees for the period of 1983-1991.[5] For our purposes this database is particularly informative since it reports the real output of the plant, the age of the plant (i.e. the year of a plant's start-up which we use to infer the plant's age), and the decision of the plant to discontinue operations. In Appendix D we discuss additional details of the data and the construction of the data moments used in the calibration.

We first calibrate some of the parameters independently. In particular, we set the elasticity of substitution across goods, σ, to 7.49 following Broda and Weinstein (2006). Similarly we set the discount factor, β, to 0.9606, which implies a quarterly discount rate of 1%. In order to calibrate the exogenous exit probability, δ, we use the model's prediction that only the smallest firms exit endogenously, due to an adverse demand realization. Given this, we set δ to 0.025 to generate an exit rate of 2.56% among the largest 5% of plants in the Colombian data.

The remaining parameters are jointly calibrated. In particular, in order to identify the importance of learning we exploit the result in Proposition 1 of the previous section and run a regression of firm growth rates on age and size. If we compare two firms of the same size, one young and one old, in our setup the younger firm is more uncertain of its true demand. As a result, it is more likely to grow as new information arrives, compared to the older firm. On the contrary, a model of productivity evolution, as in Hopenhayn (1992), cannot generate this age dependence, conditional on size (see also Ericson and Pakes (1998) and discussion therein).

To be precise, we calibrate the following four parameters: the per period fixed cost, f, the standard deviation of shocks to demands, σ_ε, the shape parameter of the distribution of the productivity draws, ξ, and the standard deviation of the distribution of the unobserved demand parameters, σ_θ.[6] In order to calibrate the

[5] Earlier data from 1977 to 1982 cover all manufacturing plants in Colombia but the data on real production are not as reliable for that time period. To avoid measurement bias caused by the change in the coverage we use only the latest part of the survey, 1983-1991.

[6] Note that $\bar{\theta}$ in this setup is not separately identified from f, so we set $\bar{\theta} = 0$.

Table 1: Parameter Values

Parameter	Value
ξ	13.09
σ_ϵ	1.46
σ_θ	0.997
f	253,181

Sources:Model Baseline Calibration

parameters, we use the following four moments: the mean of log sales; the sales-weighted growth rate of entrants; the entrants' share of sales; and, following the discussion above, the age coefficient in a regression of firms growth rates on age and size.[7]

Although a rigorous identification argument is impossible due to the complexity of the setup, we give an informal argument of how each parameter is identified from the data. The per period fixed cost, f, is pinned down by the mean log sales, since increasing the fixed cost pushes less productive firms out of the market thereby increasing the mean sales. The shape of the distribution of the initial productivity draws, ξ, is pinned down by entrants' share of sales: A lower value for ξ, all else equal, leads to a higher initial productivity dispersion and therefore more entrants starting off relatively large. Finally, the parameters related to the learning process, σ_ε and σ_θ are pinned down by the sales-weighted growth rate of entrants and the age coefficient. The calibrated parameters for this baseline calibration are presented in Table 1 and the targeted moments are presented in Table 2.

As shown in Table 2, the fit for all four moments is close to exact. In addition, Table 2 reports the size coefficient of the regression of growth of sales on age and size. This moment was not targeted in the calibration and, nevertheless, the predicted coefficient is very close to the empirical one. Moreover, the calibrated model does well in predicting the annual survival rate of firms, which was also not targeted in the calibration. Given that learning affects both the growth and the exit behavior

[7]The simplex search method is used to search over the parameter space $(\xi, f, \sigma_\epsilon, \sigma_\theta)$. To compute simulated moments, 40,000 firms are simulated.

Table 2: Data and Simulation Moments

Targeted Moments:	**Data**	**Model**
Mean of log-sales	14.99	14.99
Share of sales from entrants	2.66%	2.66%
Sales-Weighted growth rate of entrants	23%	23%
Age Coefficient	-0.035	-0.035
Other Moments:		
Size Coefficient	-0.022	-0.022
Annual Survival Rate	91%	91%

Sources: DANE survey data (see text for details) and Model Baseline Calibration

of firms, it's reassuring that our model is able to match that moment well.

Finally, as noted above, in Appendix A we repeat the calibration, but allow productivity to evolve over time. In particular, we assume it follows an AR(1) process. The results suggest that learning is the most important determinant of young firms growth and exit behavior.

4 Counterfactual Experiments

We next perform counterfactual experiments to evaluate the extent to which government policies can affect the aggregate welfare in an economy. Our focus lies on analyzing the effect of the fixed cost subsidies or taxes imposed on different groups of young firms and lump-sum taxed or rebated to the consumers, respectively. This particular configuration, while stylized, is meant to capture various forms of government or venture capital interventions which seek to provide support to firms in the first years of their existence. Subsidies to young firms, for example, are often proposed as a means to boost growth and employment (See the recommendations of Commision (2010) and in chapter 5 in the OECD (2013) report). The complications that arise in solving general equilibrium models of learning and structurally estimating them have not permitted so far the evaluation of these policies.

We evaluate the effect of fixed cost subsidies/taxes on firm entry-exit and ag-

Figure 1: **Fixed cost subsidy/tax to different firms and aggregate welfare: Melitz (2003) Model**

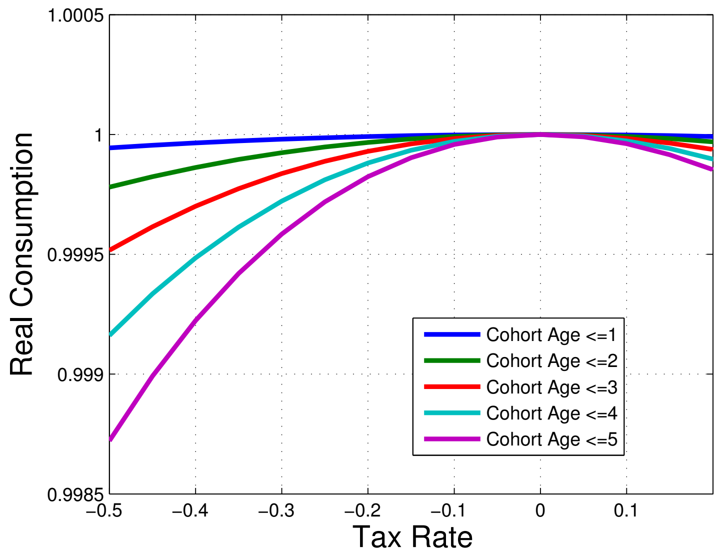

Source: Model calculations.
Notes: Tax applied to firms $\leq$ number of ages.

gregate welfare. In addition, we consider the effect of these policies implemented on different cohorts of firms, such as only entrants, firms which have been in the market for at most two years, at most three years, etc. Since youngest firms face the largest uncertainty about their product appeal we expect that the subsidies will have a more pronounced effect on them. Additionally, the effect of the policy will depend on the size of the targeted group.

Before we start, it is important to stress that in the standard static Melitz (2003) model without learning no form of government intervention will improve welfare since the decentralized equilibrium outcome is Pareto optimal. This is not ex-ante obvious: with firm entry and exit, firms do not internalize the benefits to consumers of an additional variety. In addition, their entry has adverse effects on the profits of incumbent firms through the monopolistic distortion.[8] These effects exactly cancel out in the basic static setup with CES demand and, as a result, efficiency holds

[8]See also discussion in Grossman and Helpman (1991), page 82.

Figure 2: **Equilibrium Mass of Firms and Demand Heterogeneity**

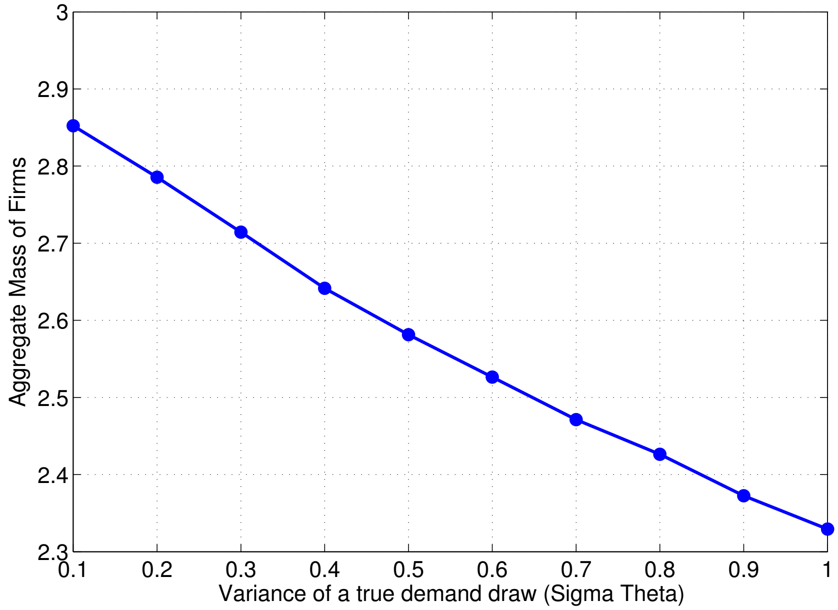

Source: Authors calculations on model baseline calibration.

even in the presence of firm heterogeneity, as pointed out by Dhingra and Morrow (2012). To illustrate this efficient outcome in the Melitz model we use our parameter values to calculate the implications of a fixed cost tax or subsidy to young firms in the canonical Melitz (2003) model. As shown in Figure 1, for any level of tax or subsidy and for any targeted group, welfare is reduced by this policy intervention. In fact, larger subsidies (indicated with a more negative tax rate) and larger taxes on the fixed cost imply a lower real consumption in the equilibrium. In addition, if more firms are subsidized, real consumption falls as well: real consumption is lower if both the first and the second cohort are subsidized, versus only the first cohort only etc.

These results are no longer true in the presence of learning, which brings about new room for policies affecting firm growth and aggregate welfare. We first notice that, as illustrated in Figure 2, when the variance of the demand draw σ_θ, increases, we move away from the outcomes of the static Melitz model (which is obtained here

Figure 3: **Fixed cost subsidy/Tax and mass of firms of different ages in the equilibrium.**

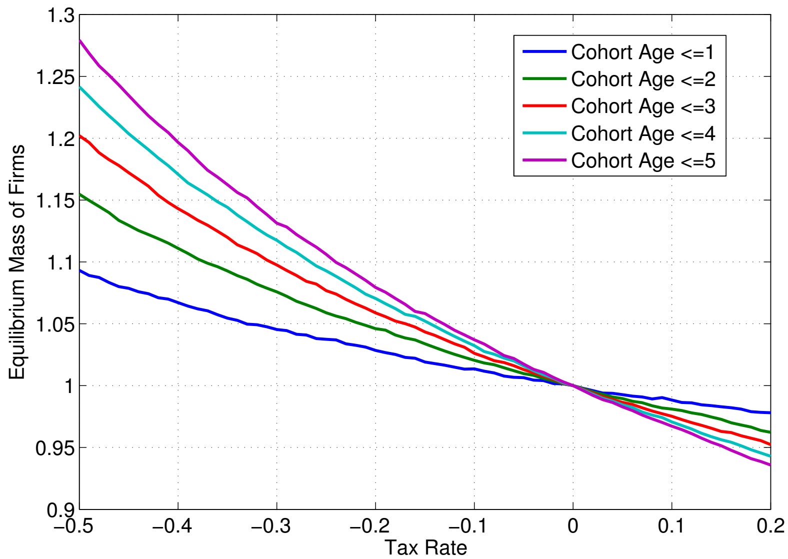

Source: Authors calculations on model baseline calibration.
Notes: Tax applied to firms $\leq$ number of ages.

when $\sigma_\theta = 0$) and fewer firms operate in equilibrium. This will be the driving force of the inefficiency in our setup since the exit of young firms reduces the number of varieties consumed by the consumers. Thus, policies which can avert this early exit may increase consumer welfare.

In our learning model subsidizing young firms increases the mass of operating firms in equilibrium. As shown in Figure 3, a 40% subsidy to the fixed costs of production leads to up a 20% increase in the mass of operating firms in equilibrium, and thus to an overall increase in the number of varieties available to consumers. The mass of operating firms increases when the subsidy base expands from subsidizing only entrants to subsidizing the fixed cost of all firms up to 5 years of age. As shown in Figure 4, which considers the case of a subsidy to entrants only, it is not only the mass of very young firms that goes up with the subsidy: even the mass of firms of age 10 has increased with a subsidy on entrants. By subsidizing entry, a higher

Figure 4: **Fixed cost Tax on Entrants**

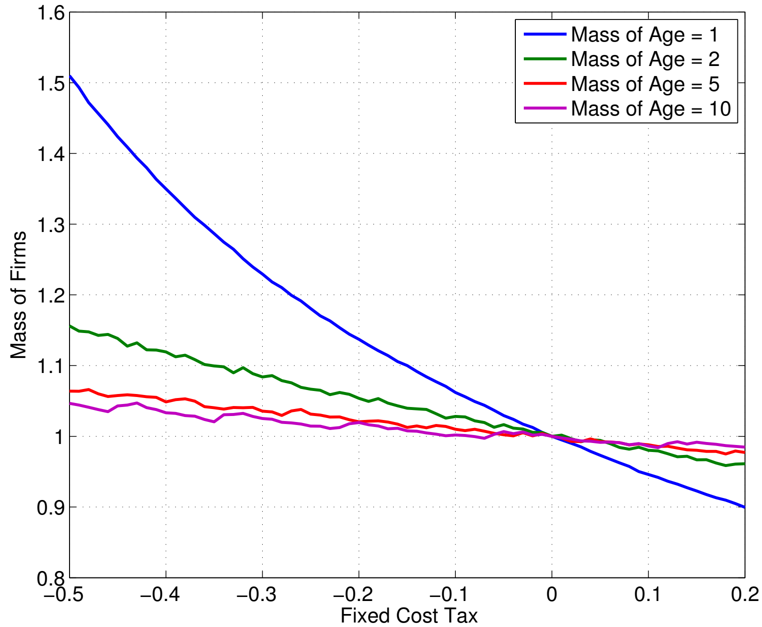

Source: Model baseline calibration.

fraction of young firms survive in the long-run.

Figure 5 shows aggregate welfare as a function of the subsidy or tax applied to the younger firms in the economy. We see that real consumption increases with a subsidy to fixed cost until a subsidy level of about 40%. On the contrary, real consumption falls when an additional tax is imposed on fixed costs. Again, welfare gains are larger when the subsidy base expands from subsidizing only entrants to subsidizing the fixed cost of all firms up to 5 years of age, as long as the subsidy levels stay below 40%.

As mentioned above, the subsidy is financed through lump-sum taxation of consumers. It's worth noting that the welfare gains are of similar order of magnitude to the cost of the subsidy. For instance, a 35% subsidy on the fixed cost of all firms up to two years of age, costs 0.36% of GDP (see Figure 6) and leads to a 0.23%

Figure 5: **Fixed cost subsidy/tax to different firms and aggregate welfare: Learning Model**

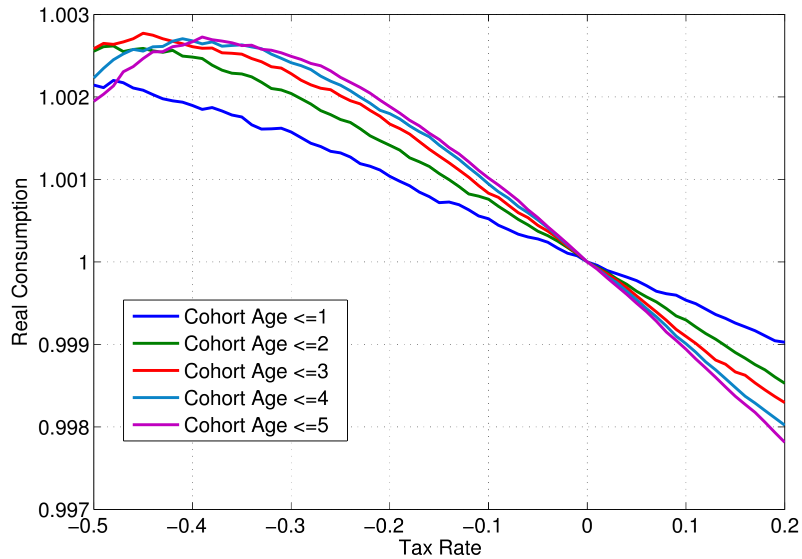

Source: Model baseline calibration.
Notes: Tax applied to firms ≤ number of ages.

increase in aggregate consumption and therefore welfare implying a consumption multiplier of 0.64. A lower subsidy leads to an even higher consumption multiplier: a subsidy for the first two years equal to 0.10% of GDP leads to a 0.08% increase in aggregate consumption (consumption multiplier of approximately 0.8). In standard neoclassical models the consumption multiplier as a result of a government tax/policy is typically negative due to a strong wealth effect that decreases both consumption and leisure (see e.g. Baxter and King (1993) and Ramey (2011) and references therein). In our model there is no leisure and the negative impact of the lower wealth on consumption is compensated by the large increase in varieties as a result of the policy, resulting in an increase in real consumption.

Our results indicate that learning implies an inefficiently low mass of operating firms, due to both higher exit as well as lower entry of firms. A large uncertainty for the firm product appeal results in a premature exit, even by firms that may have

Figure 6: **Tax Collections to GDP Ratio**

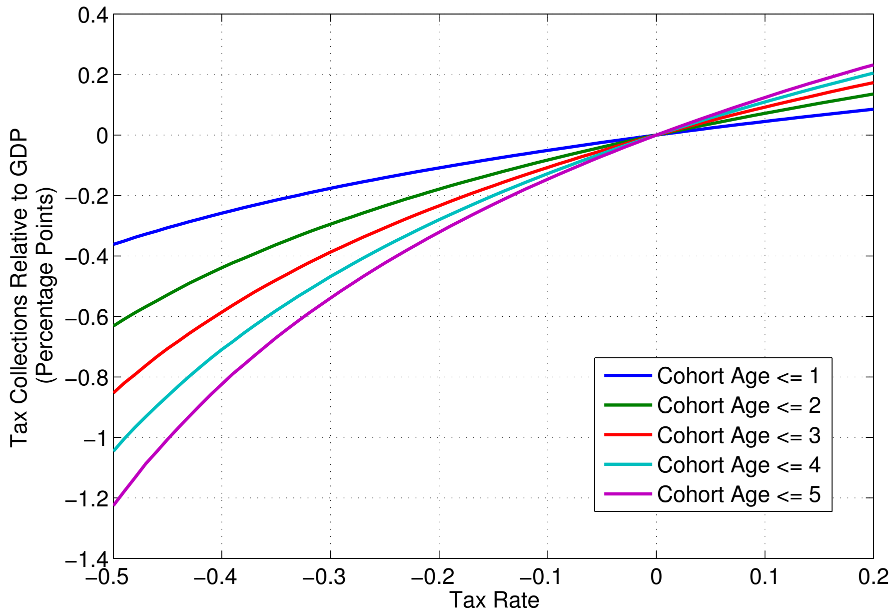

Source: Authors calculations on model baseline calibration.

large true product appeal. While this decision is optimal from the individual firm's point of view, the firm does not internalize the loss of a variety for the consumer.[9] The same argument applies for the firm entry decision. A fixed cost subsidy improves efficiency: lower exit and higher entry both lead to a larger mass of operating firms.

5 Conclusion

In this paper we develop a framework to evaluate the importance of learning for firm growth by suitably adapting the standard learning mechanism of Jovanovic (1982) into the monopolistic competition framework of Melitz (2003). Our setup captures the dependence of growth rates on age, even conditional on firm size. We calibrate our setup using this moment, as well as other moments from a panel of Colombian

[9]This externality in the setup with learning, as shown in Figure 5 outweights the effect of exit on the monopolistic distortion.

plant-level data and illustrate how targeted subsidies to the fixed costs of operations of young firms can lead to positive welfare gains. In our baseline setup we estimate gains from these targeted subsidies which are similar order of magnitude to the cost of the subsidy.

References

Abbring, J. H. and J. R. Campbell (2003): "A Structural Empirical Model of Firm Growth, Learning, and Survival," *NBER working paper*, 9712.

Albornoz, F., G. Corcos, E. Ornelas, and H. F. C. Pardo (2012): "Sequential Exporting," *Journal of International Economics*, 88, 17–31.

Arkolakis, C. (2011): "A Unified Theory of Firm Selection and Growth," *NBER working paper*, 17553.

Baxter, M. and R. G. King (1993): "Fiscal Policy in General Equilibrium," *American Economic Review*, 83, 315–334.

Bergemann, D. and J. Välimäki (2000): "Experimentation in markets," *The Review of Economic Studies*, 67, 213–234.

Broda, C. and D. Weinstein (2006): "Globalization and the Gains from Variety," *Quarterly Journal of Economics*, 121, 541–585.

Chaney, T. (2008): "Distorted Gravity: The Intensive and Extensive Margins of International Trade," *American Economic Review*, 98, 1707–1721.

Commision, E. (2010): "Communication from the Commision to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions," *European Commission*.

Cooley, T. F. and V. Quadrini (2001): "Financial Markets and Firm Dynamics," *American Economic Review*, 91, 91.

Dhingra, S. and J. Morrow (2012): "The Impact of Integration on Productivity and Welfare Distortions Under Monopolistic Competition," *mimeo, LSE*.

Eaton, J., M. Eslava, C. J. Krizan, M. Kugler, and J. Tybout (2012): "A Search and Learning Model of Export Dynamics," Manuscript, Pennsylvania State University.

Eaton, J., M. Eslava, M. Kugler, and J. Tybout (2008): "The Margins of Entry Into Export Markets: Evidence from Colombia," in *Globalization and the Organization of Firms and Markets*, ed. by E. Helpman, D. Marina, and T. Verdier, Massachusetts: Harvard University Press.

Ericson, R. and A. Pakes (1998): "Empirical Implications of Alternative Models of Firm Dynamics," *Journal of Economic Theory*, 79, 1–45.

Evans, D. S. (1987): "Tests of Alternative Theories of Firm Growth," *The Journal of Political Economy*, 95, 657–674.

Grossman, G. and E. Helpman (1991): *Innovation and Growth in the Global Economy*, Cambridge, Massachussets: MIT press.

Haltiwanger, J., R. Jarmin, and J. Miranda (2011): "Who Creates Jobs? Small Vs. Large Vs. Young," *NBER working paper*, 16300.

Hopenhayn, H. A. (1992): "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica*, 60, 1127–1150.

Jovanovic, B. (1982): "Selection and the Evolution of Industry," *Econometrica*, 50, 649–670.

Klette, J. and S. Kortum (2004): "Innovating Firms and Aggregate Innovation," *Journal of Political Economy*, 112, 986–1018.

Luttmer, E. G. J. (2007): "Selection, Growth, and the Size Distribution of Firms," *Quarterly Journal of Economics*, 122, 1103–1144.

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

OECD (2013): "OECD Science, Technology and Industry Scoreboard 2013," *OECD Publishing*.

Ramey, V. A. (2011): "Can Government Purchases Stimulate the Economy?" *Journal of Economic Literature*, 49, 673–685.

Roberts, M. J. (1996): *Colombia, 1977-85: Producer Turnover, Margins, and Trade Exposure*, Oxford, UK: Oxford University Press, 227–259.

Ruhl, K. J. and J. Willis (2014): "New Exporter Dynamics," Manuscript New York University.

Timoshenko, O. (2014): "Product Switching in a Model of Learning," *Journal of International Economics*, forthcoming.

A Learning Model with Productivity Evolution

In this appendix we modify the learning model to allow for firm productivity, e^{z_t}, to change over time. We also calibrate this extended model with both learning and productivity evolution.

In particular, conditional on the firm's survival, productivity $z_t(\omega)$ changes over time according to

$$z_t(\omega) = \rho z_{t-1}(\omega) + u_t(\omega), \; u_t(\omega) \sim N\left(0, \sigma_u^2\right) \text{ i.i.d.}$$

Firms observe this period's productivity when making their quantity decision, but do not know next period's productivity realizations. In addition, today's quantity choice does not affect the evolution of productivity. Thus, the firm's optimal quantity decision is still given by equation (4). The firm's value function (dropping the ω and the t) is now given by

$$V(z, b, n) = \max\left\{E\pi(z, b) + \beta(1-\delta)E_{z'|z, b'|b,n}V(z', b', n+1), 0\right\}.$$

Since productivity now changes over time, firms take this into account both when making their entry decision, as well as when considering whether to remain in a market or exit. In particular, there is an option value of higher productivity draws in the future, especially given the convexity of e^z.

We now turn to the calibration. As in the calibration of the baseline model we calibrate some of the parameters independently. In particular, β and δ take the values of 0.9606 and 0.025 respectively, while the autocorrelation coefficient, ρ, is set to 0.999.

The remaining parameters are jointly calibrated. In particular, we calibrate the following five parameters: the per period fixed cost, f, the standard deviation of shocks to demands, σ_ε, the shape parameter of the distribution of the initial productivity draws, ξ, the standard deviation of the distribution of the unobserved demand parameters, σ_θ and the standard deviation of shocks to productivity, σ_u. We use the following five moments: mean log sales, the standard deviation of log sales, the sales-weighted growth rate of entrants, the entrants' share of sales and the annual survival rate.

As in the baseline model, it's impossible to provide a rigorous identification argument, but here's an informal argument of how each parameter is identified from the data. The per period fixed cost, f, is pinned down by mean log sales, since increasing the fixed cost pushes less productive firms out of the market increasing mean sales. The standard deviation of log sales informs us about the standard deviation of the distribution of the unobserved demand, σ_θ, since higher dispersion in that distribution translates into a more disperse sales distribution. The shape of the distribution of the initial productivity draws, ξ, is pinned down by entrants' share of sales: A lower value for ξ, all else equal, leads to a higher initial productivity dispersion and therefore more entrants starting off relatively large. Finally, the sales-weighted growth rate of entrants and the annual survival rate identify the standard deviation of shocks to demand, σ_ε and the standard deviation of shocks to productivity, σ_u. Intuitively, a lower σ_ε implies faster learning and therefore, all else equal, low demand firms exit sooner, while surviving firms grow faster. Similarly, a higher dispersion of productivity shocks, σ_u, leads to both a higher exit rate, as more firms are hit by large negative shocks, but also higher growth rates. Therefore the growth and survival moments allow us to identify these two moments.

The calibrated parameters are presented in Table 3 and the targeted moments are presented in Table 4. Two comments are in order: First the fit of the model is quite good, with all moments quite close to their targets. In addition as shown in Table 4, the model also matches the size coefficient of the regression of growth of sales on age and size. This moment was not targeted in the calibration and nonetheless the fit is close to exact.

Table 3: Parameter Values for the Learning Model with Productivity Evolution.

Parameter	**Value**
ξ	10.80
σ_ϵ	1.32
σ_θ	0.95
σ_u	0.02
f	248,572

Sources:Model Calibration with Productivity Evolution

Table 4: Data and Simulation Moments for the Learning Model with Productivity Evolution.

Targeted Moments:	**Data**	**Model**
Mean of log-sales	14.99	15.09
Share of sales from entrants	2.66%	2.65%
Sales-Weighted growth rate of entrants	23%	21%
Age Coefficient	-0.035	-0.035
Annual Survival Rate	91%	90%
Other Moments:		
Size Coefficient	-0.022	-0.023

Sources:DANE survey data (see text for details) and Model Calibration with Productivity Evolution

Second, the parameters indicate that most of the action is driven by learning rather than productivity. Indeed the standard deviation of productivity shocks, σ_u, is very small, especially compared to the two parameters capturing the importance of learning, σ_ε and σ_θ.

This becomes even more apparent when we use the calibrated model with productivity growth and shut down productivity evolution (set $z_t = z_0$ for all t). The resulting moments are presented in Table 5. All moments do not change much suggesting that learning is quite important in accounting for the growth and exit behavior of young firms.

Table 5: Counterfactual No Productivity Growth (i.e. $z_{t+1} = z_t$).

Targeted Moments:	**Baseline Model**	**No Productivity Growth**
Mean of log-sales	15.09	15.08
Share of sales from entrants	2.65%	2.50%
Sales-Weighted growth rate of entrants	21%	19%
Age Coefficient	-0.035	-0.032
Annual Survival Rate	90%	92%
Size Coefficient	-0.023	-0.017

Sources:Model Calibration with Productivity Evolution

B Proof of Proposition 1

The expected growth rate of a firm with current size q_t is given by

$$E_t\left(\frac{q_{t+1}}{q_t}\right) = \frac{E_t\left(q_{t+1}\right)}{q_t}. \tag{9}$$

Using equation (4), we can substitute in for the firm's quantity choice each period to obtain

$$\frac{E_t\left(q_{t+1}\right)}{q_t} = \frac{E_t\left(\left(\frac{\sigma-1}{\sigma}\right)^{\sigma}\left(\frac{b_{t+1}e^{z}}{w}\right)^{\sigma}\frac{Y}{P^{1-\sigma}}\right)}{\left(\frac{\sigma-1}{\sigma}\right)^{\sigma}\left(\frac{b_t e^{z}}{w}\right)^{\sigma}\frac{Y}{P^{1-\sigma}}}$$

$$= \frac{\left(\frac{\sigma-1}{\sigma}\right)^{\sigma}\left(\frac{e^{z}}{w}\right)^{\sigma}E_t\left(b_{t+1}^{\sigma}\right)\frac{Y}{P^{1-\sigma}}}{\left(\frac{\sigma-1}{\sigma}\right)^{\sigma}\left(\frac{b_t e^{z}}{w}\right)^{\sigma}\frac{Y}{P^{1-\sigma}}} = \frac{E_t\left(b_{t+1}^{\sigma}\right)}{b_t^{\sigma}}.$$

We know that b_{t+1}^{σ} is log normally distributed, with mean

$$m_n = \log\left(b_t^{\sigma}\right) - \frac{v_n^2 - v_{n+1}^2}{2\sigma}$$

and variance given by

$$s_n^2 = \frac{\lambda^2\left(v_n^2 + \sigma_\varepsilon^2\right)}{\left(1 + (n+1)\lambda\right)^2}$$

where

$$\lambda = \frac{\sigma_\theta^2}{\sigma_\varepsilon^2}$$

and

$$v_n^2 = \frac{\sigma_\theta^2 \sigma_\varepsilon^2}{\sigma_\varepsilon^2 + n\sigma_\theta^2} = \frac{\lambda \sigma_\varepsilon^2}{1+n\lambda}$$

and n is the firm's age (number of observations).

Thus the expected growth rate is given by

$$E_t\left(\frac{q_{t+1}}{q_t}\right) = \frac{E_t\left(b_{t+1}^\sigma\right)}{b_t^\sigma} = \frac{\exp\left(m_n + \frac{s_n^2}{2}\right)}{b_t^\sigma} = \frac{b_t^\sigma \exp\left(\frac{1}{2}\frac{\lambda^2\left(v_n^2+\sigma_\varepsilon^2\right)}{(1+(n+1)\lambda)^2} - \frac{v_n^2 - v_{n+1}^2}{2\sigma}\right)}{b_t^\sigma} =$$

$$\exp\left(\frac{1}{2}\frac{\lambda^2\left(v_n^2+\sigma_\varepsilon^2\right)}{\left(1+(n+1)\lambda\right)^2} - \frac{v_n^2 - v_{n+1}^2}{2\sigma}\right).$$

Straightforward calculations show that we can rewrite the above growth rate as

$$\exp\left(\frac{\lambda^2\sigma_\varepsilon^2(\sigma-1)}{2\sigma(1+n\lambda)(1+(n+1)\lambda)}\right).$$

The derivative of the above growth rate with age (n) gives

$$\frac{\partial \exp\left(\frac{\lambda^2\sigma_\varepsilon^2(\sigma-1)}{2\sigma(1+n\lambda)(1+(n+1)\lambda)}\right)}{\partial n}$$

$$\begin{aligned}
&= \exp\left(\frac{\lambda^2\sigma_\varepsilon^2(\sigma-1)}{2\sigma(1+n\lambda)(1+(n+1)\lambda)}\right)\frac{\partial\left(\frac{\lambda^2\sigma_\varepsilon^2(\sigma-1)}{2\sigma(1+n\lambda)(1+(n+1)\lambda)}\right)}{\partial n} \\
&= -\exp\left(\frac{\lambda^2\sigma_\varepsilon^2(\sigma-1)}{2\sigma(1+n\lambda)(1+(n+1)\lambda)}\right)\frac{2\sigma\lambda^2\sigma_\varepsilon^2(\sigma-1)\left(2\lambda+2n\lambda^2+\lambda^2\right)}{\left(2\sigma(1+n\lambda)(1+(n+1)\lambda)\right)^2} < 0
\end{aligned}$$

since $\sigma > 1$ and all other parameters are positive.

C Solving the Stationary Equilibrium

The stationary equilibrium objects that need to be solved for are $e^{\underline{z}}, M, P, Y, w$.

Change of variables

$$
\begin{aligned}
u^{\sigma-1} &= \frac{(e^{z})^{\sigma-1}P^{\sigma-1}Y}{w^{\sigma-1}} \\
(u^{*})^{\sigma-1} &= \frac{(e^{\underline{z}})^{\sigma-1}P^{\sigma-1}Y}{w^{\sigma-1}}
\end{aligned}
$$

u^{*} is a solution to the firm's entry problem

$$V(u^{*}, b^{e}, 0) = 0$$

Next

$$M = \frac{L}{\tilde{r} - \tilde{\pi}},$$

where $\tilde{r}$ is the mean revenue of firms and $\tilde{\pi}$ is the mean profit level of firms.

$e^{\underline{z}}$ is a solution to

$$
\begin{aligned}
M &= J\left(\frac{e^{z_{min}}}{e^{\underline{z}}}\right)^{\xi} \times Mass \\
e^{\underline{z}} &= \left(\frac{J \times Mass}{M}\right)^{\frac{1}{\xi}} e^{z_{min}},
\end{aligned}
$$

where, $Mass$ is the mass of firms as determined by $m(z, b, n)$. Next

$$
\begin{aligned}
Y &= L + M\tilde{\pi}, \\
P &= \frac{u^{*}}{Y^{\frac{1}{\sigma-1}} e^{\underline{z}}}.
\end{aligned}
$$

Table 6: OLS Regression: the Age-Size Dependence of Growth Rates

$\log(RP_{i,t})$	-0.022***
	(0.002)
$\log(Age_{i,t})$	-0.035***
	(0.002)
Constant	0.428***
	(0.022)
Sample	$Age_{i,t} \leq 20$
No. Obs.	36,882
R^2	0.02

*** statistically significant at 1% level.

D Data

The measure of a plant's sales is taken to be the real value of production (variable RP in the dataset). The real value of production is reported in thousands of pesos. Thus, all values are multiplied by 1000.

The first calibrating moment - the mean of the logarithm of sales - is constructed by taking a cross sectional mean of the logarithm of plants' sales for a given year. The reported value is the mean across annual observations between 1983 and 1991.

For the second calibrating moment - the share of sales from entrants - an entrant is defined as a plant which is observed selling a positive amount in a given period, and is not observed in the sample in the previous period. The reported value is the mean across annual observations between 1983 and 1991.

The third calibrating moment - the sales-weighted growth rate of entrants - is constructed by taking the sales-weighted mean of the cumulative growth rate of sales of surviving entrants. The reported value is the mean across annual observations between 1983 and 1991.

The fourth calibrating moment - the age coefficient - is taken to be the value of the coefficient β_2 in the regression below

$$\log\left(\frac{RP_{i,t+1}}{RP_{i,t}}\right) = \alpha + \beta_1 \log(RP_{i,t}) + \beta_2 \log(Age_{i,t}) + \epsilon_{i,t},$$

where $RP_{i,t}$ is the real value of production for plant i at time t. Variable Age measures the number of *consecutive* years a plant is observed in the sample. For example, for a plant that is observed in years 1984, 1986, and 1987, the plant's age in 1984 is 1, in 1986 is 1, in 1987 is 2. The age of a plant in 1983 (the start of our sample period) is determined by the difference between 1983 and the plants start year (variable X6 in the dataset). The regression is run on the subsample of plants with $Age_{i,t} \leq 20$. Results are reported in Table 6.