A Firm's First Year

Jaap H. Abbring*
Jeffrey R. Campbell†

December, 2005

Abstract

This paper determines the structural shocks that shape a firm's first year by estimating a structural model of firm growth, learning, and survival using monthly sales histories from 305 Texas bars. We find that heterogeneity in firms' pre-entry scale decisions accounts for about 40% of their sales' variance; persistent post-entry shocks account for most of the remainder. We find no evidence of entrepreneurial learning. Variation of the firms' fixed costs consistent with an annual lease cycle explains their exit rates. We use the estimated model to price a new bar's option to exit, which accounts for 124% of its value.

*Department of Economics, Free University, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands, and Tinbergen Institute. Email: jabbring@econ.vu.nl.

†Federal Reserve Bank of Chicago, 230 South LaSalle Street, Chicago IL 60604, USA and NBER. Email: jcampbell@frbchi.org

Keywords: firm exit, option value, fixed costs, Bayesian learning

JEL codes: C34, D83, L11, L83

We gratefully acknowledge the research support of the National Science Foundation through grant #013748 to the NBER. Jaap Abbring's research is supported by a fellowship of the Royal Netherlands Academy of Arts and Sciences and was in parts carried out at The University of Chicago, University College London, and the Institute for Fiscal Studies. We would also like to thank Gadi Barlevy and numerous seminar participants for comments, and the IZA in Bonn for facilitating a joint visit in December 2001. Ronni Pavan and Raghu Suryanarayanan provided superlative research assistance.

1 Introduction

The risk inherent to a new firm's life gives value to the option to exit. In this paper, we quantify this option's contribution to a new firm's value using monthly sales histories from 305 new bars in Texas. Doing so requires identifying the fixed costs that exit avoids and the information that entrepreneurs acquire as their firms age. We accomplish this by estimating a structural model of firm growth and survival. The estimation distinguishes between persistent and transitory post-entry shocks and separates these from firms' different pre-entry scale decisions. Jovanovic's (1982) model of entrepreneurial learning and Hopenhayn's (1992) model of Markovian firm growth are both special cases nested within it. The model's point estimates indicate that firms' pre-entry scale decisions account for about 40% of their sales' variance; persistent post-entry shocks to their profitability account for most of the remainder. All of the growth of surviving firms' average sales reflects selection. The option to exit accounts for 124% of a new firm's value: A firm that must be operated in perpetuity is a liability. Thus, a policy analysis that fails to examine entrepreneurs' exit decisions almost certainly ignores the largest component of a new firm's value.

Twenty percent of the firms in our sample exit in less than one year. Their monthly exit rates are low immediately following entry, but they rise quickly after six months. This survival pattern is consistent with entrepreneurial learning about a firm's unobservable profitability, but it can also reflect fixed costs of continuation that vary over time. A reduced-form analysis of growth and survival cannot distinguish between them, while an identified structural model that nests them, such as ours, can. We find that entrepreneurial learning plays no role in the firm's first year of operation. Instead, fixed costs that are very high in the first month of each new year and relatively low in the remaining 11 months explain the increase in exit rates. This pattern is consistent with many of the costs of annual leases for space and equipment being effectively sunk at the beginning of the year.

The sales of surviving bars in our sample display non-Markovian history dependence: Their initial sales significantly contribute to forecasts of their sales in subsequent years that condition on the firm's recent history. Pakes and Ericson (1998) find a similar result in a long annual panel of Wisconsin retailers, and interpret this as evidence of Bayesian entrepreneurial learning. Our analysis builds on theirs, but it accounts for heterogeneity across firms' pre-entry scale decisions and transitory shocks observed only by entrepreneurs. We find that these can fully explain surviving bars' long history dependence without appealing to entrepreneurial learning.

The paper's structural analysis begins with an extension of Jovanovic's (1982) model of firm selection. A firm's producer surplus is proportional to its sales; and the logarithm of each firm's surplus sums the entrepreneur's choice of scale before entry, post-entry shocks that persistently affect profit, and transitory shocks. After an entrepreneur observes her profit, she must decide either to remain open or to exit and avoid future fixed costs. An entrepreneur only observes her profit and *one* of the transitory shocks, and she bases her continuation decision on her posterior belief regarding the persistent shocks. The scale decision affects the firm's sales and costs proportionally, so it has no relevance for the entrepreneur's continuation decision. Heterogeneity in the scale decisions enters as a random effect in firms' sales and fixed costs, and is thus a potential source of sales' persistence. We follow Jovanovic and assume that the persistent and transitory shocks unobserved by the entrepreneur have normal distributions. Hence, the entrepreneur's posterior beliefs are also normal with a history-dependent mean and an age-dependent variance. The firm exits when the posterior's mean falls below an age-specific threshold.

In our model, the entrepreneur observes her scale decision and one of profit's transitory shocks. With this information, she can better forecast her firm's future than can an observer who relies on the public tax records we use. Thus, the model violates Rust's (1987) assumption of conditional independence of unobservables. This makes existing identification strategies for dynamic discrete choice models, such as Hotz and Miller's (1993), unavailable. Nevertheless, the firm histories we observe identify the model's parameters. We show that differences between all firms' initial sales and the initial sales of those that subsequently survive reveal the distribution of the entrepreneurs' observations as well as the rules they use for their exit decisions. Heuristically, if these samples are very similar we conclude that most observed variation across firms is irrelevant for entrepreneurs' exit decisions.

We augment the Kalman filter to account for sample selection, and we use it to calculate maximum likelihood estimates of the entrepreneurs' optimal exit thresholds and all of the model's identified structural parameters except firms' fixed costs. With these in hand, we can recover the fixed costs using the profit maximization problem's Bellman equation. This two-step approach has both computational and substantial advantages. Because it estimates the entrepreneur's decision rule directly, our first stage avoids the repeated solution of the entrepreneur's dynamic programming problem that is inherent to Rust's (1987) nested fixed point algorithm. Moreover, the procedure accounts for the entrepreneur's superior understanding of her profit maximization problem. For example, time-varying fixed costs associated with annual lease renewal can manifest themselves in the estimated exit thresholds even if we fail to appreciate their importance before estimation.

The estimated structural model characterizes the dependence of entrepreneurs' forward looking exit decisions on fixed costs of continuation and the stochastic process for profit, so it can be used to quantitatively evaluate the effects of policy interventions on an entrepreneur's decisions and her firm's value. For example, we calculate the value of the entrepreneur's option to exit by simulating the model first using optimal exit policies and then constraining the firm not to exit.

This policy analysis places the present paper into a new literature that uses estimated structural models of firms' dynamic decisions for policy analysis. Benkard's (2004) evaluation of competition policies for the wide-bodied commercial aircraft market exemplifies it. He relies on detailed observations of airplane construction costs and product characteristics to identify credibly individual firms' cost and demand functions. Most industries that are less concentrated than aircraft production lack such detailed public information, so this this data-intensive approach becomes infeasible when evaluating policies of broad importance – such as small-business subsidies. Fortunately, many such exercises require only a policy-invariant rule for the evolution of pre-tax *profit*. Our analysis demonstrates how to estimate this from sales and survival histories commonly available from trade directories, statistical agencies' industry surveys, and tax records. In this sense, our approach to policy analysis using limited observations complements Benkard's more data-intensive path.

The remainder of the paper proceeds as follows. The next section describes the data, Section 3 presents the structural model, and Section 4 considers the theory of its empirical content. Section 5 discusses the model's parameter estimates and their implications for the value of a new firm's option to close. The final section offers some concluding remarks.

2 Histories of Firm Growth and Survival

Readily available data sources— such as public tax records, business directories, or economic census records— allow the construction of data documenting the growth and survival of a cohort of entering firms. Gort and Klepper (1982), Dunne, Roberts, and Samuelson (1988), Bahk and Gort (1992), Jovanovic and MacDonald (1994), Holmes and Schmitz (1994), and Pakes and Ericson (1998) have all examined such data sets and characterized their implications for various aspects of firm and industry growth. Our analysis uses a similar data set constructed from a panel of Texas alcohol tax returns. The state of Texas collects a 14% tax on the sale of alcohol for on-premise consumption, and the Texas Alcoholic Beverage Control Board (TABC) makes these returns publicly available. Returns are filed monthly, and a firm must file a separate return for each of its establishments. Information included with each return includes the identity and street address of the establishment's parent firm, its trade name, its own street address, the date its alcohol license was issued, the date it was returned if the establishment no longer operates, and its tax payment for that month.

2.1 Sample Construction

Our observations begin in December, 1993 and end in March, 2001. We divided all sales observations in our sample by the geometric average of tax returns from the establishment's county filed in the same month. This accounts for persistent differences in input prices across counties, inflation over the sample period, and county-level aggregate fluctuations. Campbell and Lapham (2004) measure restaurants' and bars' aggregate responses to large county-specific demand shocks. They are an order of magnitude smaller than typical changes to individual establishments' sizes. For this reason, we henceforth focus on this idiosyncratic variation and suppose that the observations of establishments' survival and scaled sales come from a stationary environment.

We used alcohol tax identification numbers and the establishments' street addresses to group these observations into establishment histories. Our linking process accounts for the fact that a single establishment may have multiple owners over its lifetime.[1] A tax return must be filed for each establishment in every month, even if the establishment sold no alcohol for that month. Therefore, our data contain several reports of zero sales. The data also contain tax returns with a very low tax payment.[2] These apparently reflect unobserved shutdown of the establishment for part of the month or a very small scale of operation. When constructing establishment histories, we equate any tax payment of less than $750 with zero. We consider an establishment to be born in the first month that it pays more than $750 in tax, and we date its exit in the first month that it fails to pay that amount. If an establishment's tax payment temporarily falls below $750, we consider that establishment to have temporarily shut down. We exclude such establishments from our data set altogether.

Alcohol sales is of primary importance for bars, but restaurants substantially profit from

[1] In our sample, there are numerous instances of an establishment being transferred from an individual to a corporate entity with the same address. These appear to be legal reorganizations with few immediate economic implications.

[2] Some of these reflect operation during only part of the first or last month of the establishments' operation. When the given dates of license issuance and return indicate that this is the case, we divide the tax payment by the fraction of the month the establishment operated. Even after this correction, there remain several tax returns with very small but positive tax payments. The smallest positive tax payment in our data is under $1.

both alcohol and food sales. So that our sample is as homogeneous as possible in this regard, we focus only on those establishments that present themselves to the public as bars. To be included in our sample, a firm's trade name must include the word "bar" or one of 10 other words indicative of a drinking place, and it must *not* include the word "restaurant" or one of 20 words indicating the presence of substantial food service.[3] Given the limitations of the data, this minimizes the risk of falsely including restaurants at the expense of falsely excluding bars.

There are also substantial differences between multiple establishment firms and their counterparts that only operate one location. The manager of an incumbent firm's new establishment can use that firm's history and experience to plan its operations and judge its prospects. An entrepreneur starting a single establishment firm has no such information. Accordingly, we exclude any establishment founded by a firm with two or more establishments in Texas from our data set. There are 305 single-establishment firms in our data set that were born in the five years beginning in 1994 and ending in 1999. These single establishment firms comprise our sample.

2.2 Summary Statistics

Table 1 reports summary statistics from our sample of firms. For each age we consider, one to thirteen months old, it reports the number of firms that survived to that age; the mean and standard deviation of sales' logarithm among these survivors; and the fraction of them that did not operate in the following month (the exit rate). Selection during these firms' first year was extensive. Over the course of the year, 20% of the firms exited. Exit rates in the months immediately after entry are quite low, and no firm exited following its sixth month. Thereafter firms' exit rates increase with age until their twelfth month. The exit rate then falls again to zero after the thirteenth month. This lack of exits continues into the (unreported) fourteenth month. The inverted "U" shape of firms' exit rate as a function of age apparently reflects an initial delay in exit followed by a cleansing of less-profitable firms at the end of the first year. In this case, firms that survived their first year were relatively fit and unlikely to have immediately exited.

[3]The words that qualify a firm for inclusion in the data set are "bar", "cantina", "cocktail", "drink", "lounge", "pub", "saloon", "tap", and "tavern." In addition to "restaurant" the words that exclude a firm from our analysis are "bistro", "brasserie", "cafe", "club", "diner", "dining", "food", "grill", "grille", "hotel", 'oyster", "restaurante", "shrimp", "steak", "steakhouse", "sushi", and "trattoria". When selecting this sample, we consider only the trade name listed on the firm's first tax return that reports a tax payment greater than $750.

Unsurprisingly, survivors' average size increases quickly with age. Initial average sales of all bars is slightly greater than the average sales of all license holders. After one year, the survivors' average size is approximately 27% greater than this overall average. The standard deviation of firms' initial sales is 0.87, and this dispersion changes little over the first year.

2.3 Firm Growth

To better understand the importance of entrepreneurs' learning for their exit decisions, Pakes and Ericson (1998) advocate examining the persistence of surviving firms' sales. In particular, they derive nonparametric predictions from two models of learning. In the model of passive learning, entrepreneurs apply Bayesian updating to learn about a time-invariant and firm-specific parameter, as in Jovanovic (1982). This model implies that a firm's initial sales will be useful for forecasting its sales throughout its life. In the active learning model, firms invest to improve their products and processes. Because the outcome of this investment is stochastic and very successful firms optimally choose to invest little and allow their knowledge to depreciate, the observation of a firm's initial sales becomes progressively less relevant for forecasting its future. Pakes and Ericson test these models' contrasting predictions using panels of Wisconsin retail and manufacturing firms. They find that initial sales improve forecasts of retailers' future sales, but manufacturers' sales appear to be Markovian. They conclude that an approach to firm dynamics based on Bayesian learning is promising for retail firms.

We have assessed the properties of our sample of bars by conducting a similar empirical investigation. We estimated density-weighted average derivatives from the regression of a firm's sales on its sales in the previous and first months— all in logarithms— using Powell, Stock, and Stoker's (1989) nonparametric instrumental-variables estimator.[4] These estimates rely on no distributional assumptions beyond standard regularity conditions, so they are appropriate for investigating the importance of a firm's initial sales on its evolution when the structural parameters relevant for the survival decision are unknown.

Table 2 reports these estimates as well as standard errors based on their asymptotic distributions.[5] All of the derivative estimates are positive and statistically significant at the

[4] If $m(x)$ denotes the expected value of y given x and $f(x)$ is the density of x, then the regression's density-weighted average derivatives are defined as $\mathbb{E}[(\partial m(x)/\partial x)f(x)]/\mathbb{E}[f(x)]$, where the expectation is taken with respect to $f(x)$. If the regression function is linear, then these equal the linear regression coefficients. If the regression function depends on an element of x only trivially, then the corresponding density-weighted average derivative equals zero.

[5] To implement this estimation, we follow Powell, Stock, and Stoker's (1988) recommendation and use the

5% level. The derivatives with respect to the previous month's sales are surprisingly similar across months. They are nearly all between 0.80 and 0.95. The derivatives with respect to the firm's sales in its first month are smaller but not negligible. Furthermore, there is no apparent tendency for the firm's initial sales to become less relevant for forecasting as the firm ages. When the dependent variable is the firm's sales in the third month, the derivative with respect to the first month's sales equals 0.154. This is nearly identical to the analogous coefficient when the dependent variable is the thirteenth month's sales, 0.168. The analogous coefficients in the other months vary from a low estimate of 0.021 to a high of 0.190. Apparently, no first-order Markov process can fit surviving firms' observed sales well.[6]

Pakes and Ericson (1998) emphasize that the observable differences between the two models they consider only apply to very old firms if sales depend on transitory shocks observed by only the entrepreneur. Thus, any conclusion about the importance of Bayesian learning based on the application of their methodology to observations of firms' first years is necessarily suspect. Non-Markovian dynamics can also arise from permanent and unobservable differences across entrepreneurs' choices of their firms' intended scales. Even with a very long panel of firm histories, Pakes and Ericson's procedure cannot distinguish between such unobserved heterogeneity and true Bayesian learning. The structural analysis we pursue next overcomes these difficulties.

3 A Structural Model of Firm Growth and Survival

In this section, we present a structural model of firm growth and survival in a monopolistically competitive industry. A firm's life begins before entry with the choice of its scale. New firms' different scales reflect heterogeneity across entrepreneurs' skills. After entry, a persistent shock and two transitory shocks affect demand for the firm's single product. The entrepreneur observes *one* of the transitory shocks and the sum of the other one with the persistent shock. She applies Bayes' rule to optimally infer the persistent shock from her noisy observations. Continuation requires payment of fixed costs that are not necessarily constant over time, and the entrepreneur can irreversibly close the firm to avoid them. She chooses to exit when her optimal inference of the persistent shock falls below a threshold that depends on the firm's age. The pre-entry choice of scale proportionally affects the firm's producer surplus and the

bias-reducing kernel discussed by Bierens (1987). Before estimation, we scaled both explanatory variables by their standard deviations. We used a tenth-order kernel with a bandwidth of 2.

[6]Expanding the set of regressors to include the past three months' sales leaves the estimated effect of the first month's sales intact.

fixed costs required for continuation, so it has no impact on the exit decision. Instead, it plays the role of a random effect across firms' sales. Because the entrepreneur observes both her scale decision and one of the transitory shocks, we cannot perfectly forecast the firm's exit with the observations of sales that we use for estimation.

We assume that firms compete *anonymously*. That is, the behavior of any single firm only possibly depends on the behavior of other firms through some aggregate statistics. There is no direct strategic interaction between any two firms. We also assume that entrepreneurs are not affected by liquidity constraints. That is, they either have deep pockets or unlimited access to credit markets. We next discuss these assumptions' content and empirical plausibility. We then detail the firm's stochastic environment and optimization problem. Following this, we consider the entrepreneur's procedure for optimally assessing her firm's future and deciding upon its survival. With the description of firms' post-entry evolution in place, we then consider the entrepreneur's pre-entry choice of its intended scale. We finish this section by showing how our framework encompasses two prominent models of firm dynamics.

3.1 Imperfect Competition and Firm Dynamics

Bars produce heterogeneous goods and compete with each other in local markets. This compels us to consider imperfect competition as the most likely market structure for our sample of firms.

The theory of competition among a large number of producers offers us two distinct approaches to consider, monopolistic and oligopolistic competition. In models of monopolistic competition such as Dixit and Stiglitz's (1977), Hart's (1985), and Wolinsky's (1986), producers compete anonymously, and strategic interaction is absent. Idiosyncratic shocks to demand and cost have no effect on competitors' optimal actions, so empirical analysis can proceed by considering each producer's choice problem in isolation from those of her rivals. Models of oligopolistic competition, such as Prescott and Visscher's (1977) and Salop's (1979), emphasize strategic interaction. A producer's actions impact the profits of her neighbors in geographic or product space, so shocks that directly influence only one producer's profits can affect her competitors indirectly.[7]

Campbell and Hopenhayn (2005) suggest distinguishing between these possibilities using cross-market comparisons of producer sizes. A robust prediction of anonymous monopolistic competition is that the producer size distribution is invariant to market size if factor prices,

[7]The observation that consumers view similar goods available at different geographic locations as imperfect substitutes does *not* immediately imply that competition is oligopolistically competitive.

consumer demographics, and technology are held constant. For example, a firm's sales in Dixit and Stiglitz's (1977) model equals the product of the constant demand elasticity with the sunk cost of entry. Models of oligopolistic competition generally predict that increasing market size erodes producers' market power. This is true in Salop's (1979) model of competition on the circle, where increasing market size— measured with the density of consumers on the unit circle— decreases the distance between any two firms. This lowers their markups, so firms must sell more to recover their fixed costs. Hence, with oligopolistic competition we expect producers' sizes to be increasing in market size. Campbell and Hopenhayn compare retailers' average sales and employment across large and small markets to determine which of these two approaches is more promising. Their results favor oligopolistic competition.

We have implemented Campbell and Hopenhayn's procedure for our bars' parent four-digit industry, Drinking Places, using exactly the same sample of markets, control variables, measures of market size, and measures of establishment size that they do. Unlike Campbell and Hopenhayn, we fail to find a statistically significant effect of market size on average establishment size in Drinking Places. In most of the specifications we have considered, the estimated coefficient on market size is *negative* and statistically insignificant. This cross-market comparison of establishment sizes does not refute the assumption that bars are monopolistic competitors.

3.2 Access to Financial Markets

Entrepreneurship is risky and may be impeded by liquidity constraints. If an entrepreneur has only limited access to financial markets, her entry and exit decisions will depend on her wealth. If wealth is not controlled for, this would introduce yet another source of non-Markovian dynamics in observed firm histories. It is therefore crucial that we take a clear and empirically sound position on the entrepreneurs' access to financial markets.

Our assumption that entrepreneurs are not liquidity-contrained allows us to ignore entrepreneur's wealth in our empirical analysis. Hurst and Lusardi (2004) provide ample support for this assumption, while refuting earlier evidence in favor of liquidity constraints. They show that the relation between wealth and entrepreneurial activity, which is often cited as evidence of binding liquidity constraints, is highly nonlinear. In particular, in the Panel Study of Income Dynamics (PSID) wealthy households are more likely to start a business, but only at the very top-end of the wealth distribution. Then, they provides various pieces of evidence against earlier interpretations of this relation in terms of liquidity contraints, and argue for an alternative intrepretation based on common determinants of wealth and entrepreneurship. Moore (2004) confirms Hurst and Lusardi's findings using data from the Survey of Consumer Finances. Similar results have recently been found for other countries (*e.g.* Hochguertel, 2004).

Finally, the Panel Study of Entrepreneurial Dynamics lends independent support for our assumption. Less than ten percent of the former owners of failed small businesses cited financial considerations as the most important cause of failure. An even smaller percentage listed financial considerations as a secondary cause. We conclude that we can ignore liquidity contraints in our structural model and empirical analysis.

We now proceed to present our model of firm dynamics under monopolistic competition and unlimited access to financial markets. We first discuss the post-entry evolution of a firm of a given scale, and we then detail the forward-looking scale choice before entry.

3.3 The Stochastic Environment

Consider the life of a single firm, which begins production at time $t = 1$. The firm is a monopolistic competitor that produces a single good at a single location. In period t, consumers demand $Q_t = \mathrm{e}^{R+X_t+W_t}P_t^{-\varepsilon}$ units of the firm's good, where $\varepsilon > 1$ is the absolute value of the demand elasticity and P_t is its price. The time-invariant demand-shifter R reflects a pre-entry choice of the intended scale of operation of the firm. Examples of decisions affecting R for our sample of bars are the centrality of the bar's location and the installation of attractions like a music stage or a mechanical bull. The composite random variable X_t+W_t shifts the demand curve through time, and we explain the properties of its two components below. Throughout, we adopt conventional notation and reserve capital letters for random variables and small letters for their realizations.

An affine cost function, $\vartheta Q_t + \mathrm{e}^R\kappa_t$, describes the firm's technology, with both ϑ and κ_t strictly greater than zero. The fixed costs can vary with the age of the firm and are higher for firms with larger intended scale. In our application to Texas' bars, time variation in the fixed cost κ_t might reflect periodic renewal of leases for equipment or space. The effect of R on fixed costs reflects the greater expenditure associated with maintaining a larger firm.

It is straightforward to show that the entrepreneur's profit-maximizing price choice is constant, $P_t = \left(\frac{\varepsilon}{\varepsilon-1}\right)\vartheta$. The resulting sales and profits are

$$\mathrm{e}^{S_t} = P_tQ_t = \left(\frac{\varepsilon}{\varepsilon-1}\right)^{1-\varepsilon}\mathrm{e}^{R+X_t+W_t}\vartheta^{1-\varepsilon} \tag{1}$$

and

$$(P_t - \vartheta)\, Q_t - \mathrm{e}^R \kappa_t = \left(\frac{1}{\varepsilon - 1}\right)\left(\frac{\varepsilon}{\varepsilon - 1}\right)^{-\varepsilon} \mathrm{e}^{R+X_t+W_t} \vartheta^{1-\varepsilon} - \mathrm{e}^R \kappa_t, \tag{2}$$

We choose the unit of account to set $\vartheta = (\varepsilon - 1)/\varepsilon$, so that log sales are simply $S_t = R + X_t + W_t$ and the firm's profit equals $\varepsilon^{-1}\mathrm{e}^{S_t} - \mathrm{e}^R \kappa_t$.

The random variable $X_t = A_t + U_t$. Of these, A_t is the persistent shock to the firm's demand and U_t is transitory. The initial value of A_t is drawn from a normal distribution with mean μ_1 and variance σ_1^2. Thereafter, a first-order autoregression governs A_t's evolution,

$$A_t = \mu + \rho A_{t-1} + Z_t \qquad \text{with } Z_t \sim \mathcal{N}(0, \sigma^2) \tag{3}$$

In (3), $\rho > 0$ and the disturbances Z_t are independent over time

The two series of transitory shocks, $\{W_t\}$ and $\{U_t\}$, are independent over time and from each other and $\{Z_t\}$. We assume that they are normally distributed with mean zero and variances γ^2 and η^2. The entrepreneur cares about these transitory shocks individually because she observes only one of them, W_t. She also observes X_t, but U_t and A_t are hidden from her. The variance of W_t contributes to the model's econometric error term. Nontrivial Bayesian learning about A_t arises from the variance of U_t.

For concreteness, we have assumed that the pre-entry choice of scale and all post-entry shocks affect the firm's demand and not its marginal cost. We could alternatively have assumed that these variables affect marginal cost as well. If we only have data on sales and survival, these models are observationally equivalent. In this paper, we restrict attention to the analysis of such data. We do not address the separate identification of idiosyncratic demand and cost variation, which would require the observation of firms' prices.

3.4 Bayesian learning and selection

At the end of each period, the entrepreneur must decide whether or not to close the firm and exit. Exit is an irreversible decision, and its payoff equals zero. The entrepreneur is risk-neutral and discounts the firm's future profits with the constant factor $\delta < 1$.

The normality of Z_t and U_t imply that the entrepreneur can use the Kalman filter to calculate an optimal inference of A_t given the relevant information at hand, $(X_1, \ldots, X_t) \equiv \bar{X}_t$.[8] Denote this optimal forecast and its mean squared error with

$$\hat{A}_t \equiv \mathbb{E}[A_t | \bar{X}_t] \quad \text{and} \quad \Sigma_t \equiv \mathbb{E}\left[\left(A_t - \hat{A}_t\right)^2\right].$$

[8]Throughout this paper, we denote the vector $(a_1, a_2, \ldots, a_t)$ with $\bar{a}_t$.

The Kalman filter calculates $\hat{A}_t$ and Σ_t recursively using

$$\hat{A}_t = \mu + \rho \hat{A}_{t-1} + \lambda_t (X_t - \mu - \rho \hat{A}_{t-1}), \tag{4}$$

and $\Sigma_t = \eta^2 \lambda_t$, where $\lambda_1 \equiv \sigma_1^2 / \left(\sigma_1^2 + \eta^2\right)$ and $\lambda_t \equiv \left(\rho^2 \Sigma_{t-1} + \sigma^2\right) / \left(\rho^2 \Sigma_{t-1} + \sigma^2 + \eta^2\right)$ for $t > 1$. The coefficient λ_t is the Kalman gain, and it measures the informativeness of the entrepreneur's observation of X_t. The firm's sales does not directly reveal X_t, so the entrepreneur's estimate of A_t is necessarily more accurate than that of an outside observer.

Although the entire history of X_t is, in principle, relevant for the entrepreneur's exit decision, $\hat{A}_t$, the firm's age, and R are sufficient for characterizing the distribution of future profits. Define $v_t\left(\hat{A}_t, S_t, R\right)$ to be the value of a firm of age t and intended scale R to an entrepreneur who estimates A_t to be $\hat{A}_t$ and observes current sales to be S_t. The Bellman equation that this value function satisfies is

$$v_t\left(\hat{A}_t, S_t, R\right) = \varepsilon^{-1} \mathrm{e}^{S_t} - \mathrm{e}^R \kappa_t + \delta \max\{0, \mathbb{E}_t[v_{t+1}(\hat{A}_{t+1}, S_{t+1}, R)]\}. \tag{5}$$

The entrepreneur calculates the expectation $\mathbb{E}_t$ in (5) using the joint distribution of $\hat{A}_{t+1}$ and S_{t+1} conditional on the available information at time t, for which $(t, \hat{A}_t, R)$ is sufficient.

The entrepreneur's optimal exit policy is simple. The expectation in the right-hand side of (5) is a function of $(t, \hat{A}_t, R)$ and it is continuous and increasing in $\hat{A}_t$. Therefore, there exists a threshold value α_t such that the entrepreneur chooses to exit if and only if $\hat{A}_t \leq \alpha_t$. The firm's scale, R, has no impact on the exit decision, because it only scales up the firm's surplus and fixed costs. Instead, it plays the role of an unobservable (to us) random effect on the sales observations.

3.5 Pre-Entry Choice of Intended Scale

With the firm's exit policy and value function in place, we are now prepared to consider the entrepreneur's pre-entry choice of intended scale. For a randomly selected entrepreneur, the cost of entry with scale R is $F\mathrm{e}^{\tau R}$, where $\tau > 1$. The assumption that the entry cost is convex in e^R reflects limits to entrepreneurs' abilities, as in Lucas (1978). The positive random variable F embodies heterogeneity across entrepreneurs in those abilities. With this cost function, doubling the scale more than doubles the costs of entry. The expected value of a new firm is a linear function of e^R, so equating the marginal cost and benefits of increasing the firm's scale for an entrepreneur with ability F yields $R = \left(\ln(\upsilon/\tau) - \ln F\right) / (\tau - 1)$, where υ is the slope of the expected value of new firm with respect to e^R. We assume that

$\ln F$ has a normal distribution with mean $\ln(\upsilon/\tau)$ and variance $(\tau-1)^2\nu^2$ so that R is also normally distributed, with mean zero and variance ν^2.

3.6 Existing Models of Firm Dynamics

As we noted in the introduction, the model encompasses versions of both Jovanovic's (1982) and Hopenhayn's (1992) models of firm dynamics. Although both models assume a perfectly competitive market structure, we see the change to a simple version of monopolistic competition as inconsequential for our purposes. The simplest versions of both models assume that fixed costs are constant over time; $\kappa_t = \kappa$ for all t. Jovanovic's model of learning about time-invariant random profitability sets $\rho = 1$ and $\sigma^2 = 0$, while we recover Hopenhayn's model with $\rho < 1$ and $\eta^2 = 0$. Our versions of both models add a transitory econometric error term, W_t, and unobservable heterogeneity, R, to make them suitable for estimation.

4 The Model's Empirical Content

The TABC panel contains observations of each firm's sales (S_t) and an indicator for its survival to age t, which we denote with N_t. Because we do not observe S_t after the firm's exit, we set $S_t = 0$ if $N_t = 0$. The sales observations cover the first $T = 12$ months in the firm's life. The TABC panel data set also records N_{T+1}. That is, the firm's survival at the end of the sample period is known. In this section, we describe our sequential procedure for inferring the structural model's parameters from the joint distribution of these observations.

In the procedure's first stage, we treat the exit thresholds $(\alpha_1, \ldots, \alpha_T) \equiv \bar{\alpha}_T$ as parameters to be estimated jointly with the parameters that determine the evolution of S_t. The procedure's second stage recovers the sequence of fixed costs κ_t up to scale using the first-stage estimates, an assumed value of the discount factor δ, and the Bellman equation. Unlike Rust's (1987) nested fixed-point algorithm, our approach to inference separates the estimation of most of the model's parameters from the repeated solution of the entrepreneurs' dynamic programming problem. This separation brings with it a distinct advantage: We estimate many of the model's structural parameters without assuming that the entrepreneurs' dynamic program is stationary or otherwise specifying firms' production possibilities beyond the sample period.

We use the structural model to estimate the content of entrepreneurs' information without actually having that information, so it is particularly important to understand the sources of the model's identification. We demonstrate next that a sample of infinitely many firm histories identifies the exit thresholds and the parameters governing the evolution of S_t, and we discuss the features of the data that are particularly influential in this inference. We then consider the first-stage estimation of these parameters and policies using maximum likelihood. Finally, we show how to recover the sequence of fixed costs in the procedure's second stage. With these in hand, we can calculate the value of the firm's option to exit and conduct other policy experiments.

4.1 First-Stage Identification

The structural model decomposes the fluctuations of S_t into four components, the firm's scale (R), the transitory shock observed by the entrepreneur (W_t), and the transitory and persistent shocks unobserved by the entrepreneur (A_t and U_t). In our procedure's first stage, we estimate the parameters that characterize these shocks' distributions— ν^2, γ^2, ρ, σ_1^2, σ^2, η^2, μ_1, and μ— and the entrepreneurs' optimal exit thresholds $\bar{\alpha}_T$. The proof of the following proposition demonstrates that observations of firm histories identify these parameters and policies.

Proposition. *If $T \geq 3$, then the joint distribution of $\left(\bar{S}_T, \bar{N}_{T+1}\right)$ uniquely determines ν^2, γ^2, ρ, σ_1^2, σ^2, η^2, μ_1, μ, and $\bar{\alpha}_T$.*

The appendix contains the proposition's proof. Here we describe the three key insights that it requires.

Blundell and Preston (1998) identify a consumer's permanent income process using the covariance of his current income (a noisy proxy) with his consumption (a forward-looking choice). The proof adapts their approach to the identification of the parameters that govern X_t. Consider, for example, the identification of $\mathbb{V}\left[X_1\right] = \sigma_1^2 + \eta^2$ from the joint distribution of S_1 and N_2. The observed value of S_1 is a noisy proxy for the model's true state variable, $\hat{A}_1 = \mu_1 + \lambda_1\left(X_1 - \mu_1\right)$. Its covariance with the forward looking choice N_2 is

$$\mathbb{E}[S_1 N_2] - \mathbb{E}[S_1]\mathbb{E}[N_2] = \sqrt{\mathbb{V}\left[X_1\right]}\phi\left(\Phi^{-1}\left(1 - \mathbb{E}[N_2]\right)\right), \tag{6}$$

where $\Phi(\cdot)$ and $\phi(\cdot)$ are the *c.d.f.* and *p.d.f.* of a standard normal random variable. This immediately identifies $\mathbb{V}\left[X_1\right]$. Intuitively, if S_1 is very useful for predicting N_2 in a linear-in-probabilities model, then much of its variation must arise from the factor that influences survival, X_1. Otherwise, we conclude that $R + W_1$ dominates its variation. The expression in (6) translates this intuition into a quantitative measure.[9]

[9] Olley and Pakes (1996), Pavcnik (2002), and Levinsohn and Petrin (2003) follow a related identification

The model contains two distinct sources of persistence in S_t, autocorrelation of A_t and variation of R across firms. Because the model is nonlinear, standard dynamic panel data techniques for distinguishing between them, such as Blundell and Bond's (1998), are inapplicable. We instead exploit the different implications of random effects and autocorrelation for the impact of selection on the evolution of S_t. Consider the regression of S_2 on S_1. In the absence of selection, this regression would be linear. Because X_1 and R contribute to its error, the regression conditional upon survival will also include a Heckman (1979) selection correction. That is

$$\mathbb{E}[S_2 - \mathbb{E}[S_2]|S_1, N_2 = 1] = \varrho(S_1 - \mathbb{E}[S_1]) + \varsigma m(S_1). \tag{7}$$

In (7), $m(S_1)$ is the relevant inverse Mills' ratio, which is known. Hence, we can obtain the coefficients ϱ and ς from this regression function.

The key to using (7) to separately identify $\text{cov}[X_2, X_1]$ from ν^2 is that these two sources of persistence have very different implications for the Heckman selection correction. Given S_1, learning that the firm survived increases the best estimate of X_1 and *reduces* the best estimate of R. Hence if persistence in S_t arises mostly from variation in R, then the coefficient multiplying the Heckman selection correction is negative. If instead the variance of R is small relative to the variability and persistence of X_t, then selection increases the expectation of S_2. In addition to this qualitative information, ϱ and ς determine $\text{cov}[X_2, X_1]$ and ν^2 uniquely. We demonstrate this in the proposition's proof.

The final key insight required by the proof is that of Pakes and Ericson (1998). As we noted in Section 2, their procedure applied directly to the original data cannot distinguish between true entrepreneurial learning, transitory shocks observed by entrepreneurs, and permanent differences in firms' scales. To reliably apply their procedure to a sample generated by our model, we must instead consider the joint distribution of $\bar{X}_T$. From this we can calculate the linear regression of X_3 on X_2 and X_1,

$$X_3 - \mathbb{E}[X_3] = \beta_{32}(X_2 - \mathbb{E}[X_2]) + \beta_{31}(X_1 - \mathbb{E}[X_1]) + V_3. \tag{8}$$

In (8), V_3 is a composite error term. If η^2 equals zero, then there is no scope for entrepreneurial learning, $\beta_{31} = 0$, and $\beta_{32} = \rho$. Otherwise, $\beta_{31} > 0$ and $\beta_{32} < \rho$. It is straightforward to write β_{31} and β_{32} as explicit functions of the model's structural parameters. These expressions show that these coefficients directly reveal ρ and η^2.

strategy. They use nonlinear functions of a continuous decision (investment or materials use) to replace a firm's unobserved productivity in the estimation of a production function. We cannot apply this strategy, because the forward-looking decision we employ is discrete.

4.2 First-Stage Estimation

With the identification of the parameters' and policies' established, we now proceed to consider their estimation using maximum likelihood. This poses no conceptual problems, but computation of the likelihood is nontrivial because the model involves repeated selection on the basis of a persistent latent state variable. In the following, we denote the vector of parameters and policies estimated in the first-stage with ϖ. Our method of calculating the likelihood function follows the non-Gaussian state-space approach of Kitagawa (1987), which utilizes the likelihood's prediction-error decomposition. If we condition on the realization of R, then this is

$$f_{\bar{Y}_{T+1}}\left(\overline{Y}_{T+1}|R;\varpi\right) = f_{Y_1}\left(Y_1|R;\varpi\right)\prod_{t=2}^{T+1} f_{Y_t}\left(Y_t|\overline{Y}_{t-1},R;\varpi\right), \tag{9}$$

where $Y_t \equiv (S_t, N_t)$ unless $t = T+1$, in which case it equals N_{T+1}.

The first term in (9), the density of Y_1, is known. We calculate the prediction-error decomposition's remaining terms recursively. We initialize the recursion with $f_{\widehat{A}_1}\left(\widehat{a}_1|Y_1,R;\varpi\right)$, which is normal with known mean and variance. To calculate the likelihood function's remaining terms, suppose that $f_{\widehat{A}_{t-1}}\left(\hat{a}_{t-1}|\bar{Y}_{t-1},R;\varpi\right)$ is known. We consider three separate cases. If the firm exits production following period $t-1$, then $N_t = 0$ and the relevant term in the likelihood function is the probability of exit conditional on the observed history.

$$f_{Y_t}\left((0,0)|\bar{Y}_{t-1},R;\varpi\right) = \int_{-\infty}^{\alpha_{t-1}} f_{\widehat{A}_{t-1}}\left(\hat{a}_{t-1}|\bar{Y}_{t-1},R;\varpi\right)d\hat{a}_{t-1} \tag{10}$$

Following exit, the evolution of S_t is trivial and the remaining terms in the prediction error decomposition identically equal one. In the second case, the firm continues production following period $t-1$, but $t = T+1$ so the data do not contain the realized value of S_t. In this case of right censoring, the final term in the prediction error decomposition is

$$f_{Y_{T+1}}\left(1|\bar{Y}_T,R;\varpi\right) = \int_{\alpha_T}^{\infty} f_{\widehat{A}_T}\left(\hat{a}_T|\bar{Y}_T,R;\varpi\right)d\hat{a}_T \tag{11}$$

In the final case, the firm produces in period $t < T+1$, so $N_t = 1$. For this case, the term of interest in the likelihood function can be written as

$$f_{Y_t}\left((S_t,1)|\bar{Y}_{t-1},R;\varpi\right) = \int_{\alpha_{t-1}}^{\infty} \frac{1}{\sqrt{\rho^2\Sigma_{t-1}+\sigma^2+\eta^2+\gamma^2}}\,\phi\left(\frac{S_t-\mu-\rho\widehat{a}_{t-1}-R}{\sqrt{\rho^2\Sigma_{t-1}+\sigma^2+\eta^2+\gamma^2}}\right) \times f_{\widehat{A}_{t-1}}\left(\hat{a}_{t-1}|\bar{Y}_{t-1},R;\varpi\right)d\hat{a}_{t-1} \tag{12}$$

Equation (12) follows from Bayes' rule and the definition of $\widehat{A}_{t-1}$.

This final case is the only one in which we wish to continue the recursion. To do so, we must calculate $f_{\widehat{A}_t}\left(\widehat{a}_t|\bar{Y}_t, R; \varpi\right)$, which is

$$f_{\widehat{A}_t}\left(\hat{a}_t|\bar{Y}_t, R; \varpi\right) = \frac{\int_{\alpha_{t-1}}^{\infty} f_{\widehat{A}_t, S_t}\left(\hat{a}_t, S_t|\hat{a}_{t-1}, R; \varpi\right) f_{\widehat{A}_{t-1}}\left(\hat{a}_{t-1}|\bar{Y}_{t-1}, R; \varpi\right) d\hat{a}_{t-1}}{f_{Y_t}\left(Y_t|\bar{Y}_{t-1}, R; \varpi\right)}. \tag{13}$$

The distribution of $\left(\widehat{A}_t, S_t\right)$ conditional on $\widehat{A}_{t-1}$ and R is bivariate normal, so both terms in the integrand of (13) are known. Therefore, the recursion can continue.

Iterating on (12) and (13) until the individual either exits or is right censored produces the likelihood function at any given choice of ϖ for a fixed value of R. We obtain the unconditional likelihood function by calculating the expectation of this density with respect to R. In practice, evaluating the likelihood function requires approximating these integrals. We do so using Gaussian quadrature procedures. Maximization of the resulting likelihood function is straightforward; and the corresponding maximum-likelihood estimate, $\hat{\varpi}$, has the usual distributional properties.

4.3 Second-Stage Estimation

With $\hat{\varpi}$ in hand, we can proceed to exploit the entrepreneur's Bellman equation to estimate the remaining identified structural parameters. First, we fix the discount factor δ at a value consistent with a 5% annual rate of interest. Rust's (1994) nonidentification result for Markovian dynamic discrete choice models shows that this assumption is unavoidable. Next, note that the Bellman equation (5) and the linearity of the producer's surplus, $\varepsilon^{-1}\mathrm{e}^{S_t}$, and the fixed costs, $\mathrm{e}^R\kappa_t$, in e^R together imply that exit decisions are governed by the scaled Bellman equation

$$g_t\left(\hat{A}_t, X_t + W_t\right) = \mathrm{e}^{X_t+W_t} - \varepsilon\kappa_t + \delta\max\{0, \mathbb{E}_t[g_{t+1}(\hat{A}_{t+1}, X_{t+1} + W_{t+1})]\}. \tag{14}$$

In (14), $g_t\left(\hat{A}_t, X_t + W_t\right) = \varepsilon v_t\left(\hat{A}_t, S_t, R\right)/\mathrm{e}^R$.[10] Given an infinite sequence $\{\varepsilon\kappa_t\}$ and the parameters estimated in the first stage, we can iterate equation (14) to recover the scaled value function g_t. Obviously, multiplying ε by a positive constant and dividing the sequence $\{\kappa_t\}$ by that same constant yields an observationally equivalent model. Therefore, we will focus on estimating $\{\varepsilon\kappa_t\}$. This only recovers fixed costs up to scale, but we use this to

[10] This alternative representation of the exit decision process clarifies the contribution of R to the data-generating process. It shows that R influences only the firm's sales and not its exit decisions.

express a firm's fixed costs as a fraction of its expected first-period surplus, $\varepsilon\kappa_t/\mathbb{E}[e^{X_1+W_1}]$. Note that this ratio does not depend on its intended scale R.[11]

The Bellman equation (14) suggests a simple procedure for estimating $\{\varepsilon\kappa_t\}$ from the first-stage estimates of the model's other parameters and of the exit thresholds: Choose the scaled fixed costs so that the corresponding optimal exit thresholds mimic the estimated thresholds. A practical problem is that we do not have estimates of α_t for $t > T$. We circumvent this in our application by assuming that $\varepsilon\kappa_t$ follows an annual cycle. That is, $\varepsilon\kappa_{t+12} = \varepsilon\kappa_t$. This specification of the scaled fixed costs captures the annual lease cycles mentioned in Section 3. The first T optimal exit thresholds can then be expressed as a known smooth function of ϖ and $\varepsilon\bar{\kappa}_{12}$, $\bar{\alpha}_T(\varpi, \varepsilon\bar{\kappa}_{12})$. We estimate the vector of fixed costs with the minimizer of the distance between these optimal thresholds and their first-stage estimates, $\hat{\bar{\alpha}}_T$,

$$\widehat{\varepsilon\bar{\kappa}}_{12} = \arg\min_{\varepsilon\bar{\kappa}_{12}} \left(\hat{\bar{\alpha}}_T - \bar{\alpha}_T(\hat{\varpi}, \varepsilon\bar{\kappa}_{12})\right)' \hat{\mathbb{V}}(\hat{\bar{\alpha}}_T)^{-1} \left(\hat{\bar{\alpha}}_T - \bar{\alpha}_T(\hat{\varpi}, \varepsilon\bar{\kappa}_{12})\right),$$

where $\hat{\mathbb{V}}(\hat{\bar{\alpha}}_T)$ is a consistent estimator for the variance-covariance matrix of $\hat{\bar{\alpha}}_T$. This estimator of the scaled fixed costs has well-known statistical properties.

Once we have estimated $\{\varepsilon\kappa_t\}$, it is straightforward to estimate the expected value of a new firm relative to its expected first-period surplus. Recall from Subsection 3.5 that υe^R is the expected value of a new firm with intended scale R. The Bellman equation (14), the first stage parametes, and $\{\varepsilon\kappa_t\}$ together yield $\varepsilon\upsilon = \mathbb{E}\left[g_1\left(\hat{A}_1, X_1 + W_1\right)\right]$. Using this, we can estimate the expected value of a new firm relative to its first-period producer's surplus, $\varepsilon\upsilon/\mathbb{E}[e^{X_1+W_1}]$, by simply plugging in our estimates of these parameters. We can also easily compute the value of a new firm if it must be operated in perpetuity. The difference between these two gives the value of the option to close the firm.

5 Empirical Results

In this section, we report the results from estimating the model's parameters by maximum likelihood using the sample of 305 new bars described in Section 2. We first demonstrate that the maximum likelihood estimate of η^2 is zero. This implies that entrepreneurs observe A_t without error and do no not learn. We then estimate a version of our model in which η^2 is constrained to equal zero. We find that post-entry idiosyncratic shocks have a half life of 47

[11]Alternatively, we could have used the ratio $\varepsilon\kappa_t/\mathbb{E}[e^{X_t+W_t}|N_t = 1]$. This measure of period t's fixed costs relative to the expected producer surplus would reflect both the evolution of fixed costs and the typical growth of surviving firms, so it is not as easily interpreted as the expression we prefer.

months and account for approximately 60% of sales' variance. The remaining 40% of sales' variance is due to heterogeneity in pre-entry scale decisions. We conclude with a presentation of the second-stage estimates of the model's fixed costs and the policy experiments that yield the value of the firm's option to exit.

5.1 First-Stage Estimates

We first discuss the estimation of η^2, which determines the speed and extent of entrepreneurial learning. Figure 1 summarizes our results. For values of η^2 between zero and 0.45, it plots the value of the likelihood function after enveloping out the other elements of ϖ. The result is clear: The likelihood function attains its maximum at $\eta^2 = 0$. That is, our maximum likelihood estimates imply that entrepreneurs observe the persistent component of profit without error. We have conducted the same exercise for several alternative specifications of our model, and they have all yielded this result. Thus, we find no evidence that entrepreneurial learning contributes to the dynamics of the firms in this sample.

Table 3 reports the first-stage maximum likelihood estimates of a version of our model in which η^2 is constrained to equal zero. Below each estimate is its asymptotic standard error, which we calculated using the outer product estimate of the information matrix. Recall from Section 2 that no firm in our sample exited following its sixth month. In light of our identification proof's reliance on selection, this led us to suspect that our data are uninformative regarding the sixth month's exit threshold. To determine whether or not this was the case, we fixed α_6 at a variety of large and negative values and chose the other elements of ϖ to maximize the likelihood function. The resulting parameter estimates are nearly invariant to our choice of α_6. Accordingly, we present here estimates based on setting α_6 to -2.5.

Consider first the estimates of μ_1, σ_1, γ, and ν; which jointly determine the distribution of firms' initial sizes. The estimate of μ_1, 0.069, is virtually identical to the unconditional mean of firms' initial size reported in Table 1. However, the standard error attached to this estimate is relatively large, 0.055. The estimates of σ_1, γ, and ν are 0.714, 0.167, and 0.581. These imply that the standard deviation of S_1 equals 0.94, which is above the estimated standard deviation of 0.87. At these parameter estimates, the pre-entry choice of location quality accounts for 38% of the variance of S_1. The persistent post-entry shock accounts for most of S_1's remaining variance. We conclude from this that entrepreneurs gain significant information about their firm's prospects immediately after entry. However, firms' initial sales also embody substantial heterogeneity of pre-entry scale decisions.

The parameters that govern the evolution of A_t are ρ, σ, and μ. The estimate of ρ is 0.985, indicating that post-entry shocks have persistent effects on firms' sizes. This is estimated with great precision— its standard error equals 0.0002, so we can reject the hypothesis that it equals one and Z_t permanently changes firm sizes. This implies that the relatively small firms that compose our sample do not obey Gibrat's law even after correcting for selection. An increase in Z_t that increases current sales by one percent increases sales twelve and twenty-four months later by 0.84% and 0.70%.

The estimate of the intercept μ implies an estimate of the mean of A_t's ergodic distribution of 0.093, which is only slightly above our estimate of μ_1. The difference between them is insignificant at any conventional level. This implies that mean log sales would not grow much in the absence of exits. Selection explains nearly all the growth in the data. We have also estimated a version of the model with age-dependent intercepts, but do not reject the assumption that the intercepts are constant.[12] Thus, we do not find evidence of systematic firm growth, such as that arising from the gradual acquisition of a clientele, as the firm ages.

The estimated standard deviation of Z_t is 0.161. This implies that the standard deviation of A_t's ergodic distribution is 0.94. This is substantially and significantly above the estimate of σ_1, so the dispersion of an entering cohort's sizes *increases* with its age in the absence of exit. In fact, the dispersion of firm sizes neither increases nor decreases with the cohort's age, so selection offsets this underlying increase in variation.

The final estimates to report are those describing the entrepreneurs' exit policy, $\bar{\alpha}_{12}$. The estimate of the initial exit threshold is -1.65. The implied probability that a firm exits following its first month is 0.008. The estimated exit thresholds tend to rise with the firm's age. The estimate of α_2 is -1.53, while that for α_{12} is -0.54. Entrepreneurs apparently become more selective when deciding on continuation of their operations during the course of the firms' first year. There is no Bayesian learning in this specification of the model, and we have estimated very little drift in average sales. Therefore, accounting for these changes in the entrepreneurs' exit policies requires parallel changes in their fixed costs. We measure these changes below with the procedure's second stage.

To help gauge how well the estimated model fits the data, Table 4 reports the implied population values of the summary statistics that we considered in Section 2.[13] Like the raw exit rates in Table 1, the exit rates implied by our model are low in the first half of the year, and high in the remainder. Substantial increases in mean sales again reflect the importance of selection effects. Consistently with the data and despite the importance of selection, sales' standard deviation does not decrease substantially as the cohort ages. Quantitatively, the

[12] The Appendix shows that such extensions of the model with age-dependent parameters are identified.

[13] We calculated these summary statistics using $100,000$ simulated firm histories from the estimated model.

exit rates increase faster in the model than in the data. The estimated model's exit rates in the ninth and eleventh months are particularly high compared to their empirical counterparts. Mean sales in the first 9 months are matched well by the model. The large selection effects in the ninth and eleventh months, however, cause the model to overstate the growth of mean sales between months 9 and 10 and months 11 and 12, and to generate a small decrease in sales' standard deviation in these months that is absent from the data.[14]

Similarly, Table 5 provides the population version of the nonparametric regression of surviving firms' log sales on their log sales in the previous and first months. Recall from Section 2 that its empirical counterpart shows that sales in the first month help predicting current sales conditional on sales in the previous month. The model generates a similar pattern of non-Markovian dynamics, even though the estimated model contains no entrepreneurial learning–the source of non-Markovian dynamics emphasized by Pakes and Ericson (1998). Instead, unobserved heterogeneity in the scale of operations and transitory shocks observed by the entrepreneurs are sufficient to generate the observed non-Markovian dynamics in our short panel of bars.

5.2 Second-Stage Estimates and the Value of the Option to Exit

Table 6 reports estimates of the fixed costs as a fraction of the average first-period producer surplus, $\epsilon\kappa_t/\mathbb{E}[e^{X_1+W_1}]$. Below each estimated fixed cost is an estimate of its asymptotic standard error.[15] The most remarkable result is that the fixed costs in the first period are some 10 times larger than the first period's producer surplus. A claim on a surplus that is on average above the per-period fixed costs in the remaining 11 months of the year compensates for this somewhat. Note that the relative fixed costs are estimated to be particularly low in the seventh month, reflecting the observed low exit rate after six months. We do not estimate the fixed costs very precisely. The standard error on the first period's cost is 1.80. The other estimates' standard errors are comparable to this when expressed as a fraction of the corresponding fixed costs. However, all of the estimated fixed costs are statistically different from zero at conventional sizes. Overall, the sum of the estimated fixed costs equals

[14]We could pursue a better match of the data by allowing the variances of the structural shocks to vary over time. The analysis in the Appendix shows that such an extended model is identified. However, we think the simplicity of our present setup outweighs a slightly better match with the data.

[15]Because we set $\eta^2 = 0$, the optimal exit thresholds follow an annual cycle. In this special case, we can recover $\widehat{\varepsilon\overline{\kappa}}_{12}$ as the solution to a set of linear equations based on the assumed optimality of $\widehat{\overline{\alpha}}_{12}$. We calculate these equations' coefficients and intercepts by simulating a large number of firm histories that begin with $\hat{A}_t = \hat{\alpha}_t$. The reported standard errors account for the resulting simulation error.

14.59 times the expected first-period producer surplus. The standard error for this sum is 2.26.

Our estimates indicate that fixed costs are high at the start of each new year. We interpret this as the effect of an annual lease cycle. Signing a lease for space or equipment that is difficult to break effectively sinks its proscribed costs. This explains why exit rates are high near the end of the year, when next year's lease costs can still be avoided, and low at the beginning of the year, when the returns on the investment in the lease have not yet materialized.

With estimates of the model's structural parameters in hand, we finally compute the value of the option to close a new firm. In the model, a newly created firm is worth 29.7 times its expected first-period surplus. The value of the entrepreneur's exit option is an important component of this. Consider a hypothetical new firm that faces the same sales process as the bars in our sample but has no option to exit. Our estimates imply that such a firm is a liability: Its estimated value is -7.0 times its expected first-period surplus. The fixed costs' standard errors naturally affect these estimates. The 95% confidence interval for a new firm's value is $(22.8, 36.7)$, and the corresponding interval for its value if operated in perpetuity is $(-20.8, 6.7)$. Together, these estimates imply that the exit option accounts for 124% of the value of a new firm. The corresponding 95% confidence interval is $(0.75, 1.73)$. The width of this interval is not small, but it does imply that the option to exit accounts for most of the value of a new firm.

6 Conclusion

Risk and selection dominate a firm's first year. Although new bars differ considerably in their intended scales, persistent post-entry shocks account for the majority of their sales' cross-sectional variation. Immediately after entry, only very unfavorable shocks induce an entrepreneur to exit. As the firm's first anniversary approaches, the entrepreneur raises her standards for continuation. Firms' exit rates correspondingly increase over the course of the first year. In principle, this behavior is consistent with Bayesian learning about persistent shocks: The entrepreneur delays her exit decision until her posterior beliefs about profitability are sufficiently tight. However, our estimates indicate that entrepreneurs observe their bars' persistent shocks without error. Instead, the increased selectivity reflects a substantial fixed cost incurred at the beginning of each year.

We conclude that a new firm's value arises mostly from the option to operate it only after favorable shocks. Hence, an analysis of small business policy that ignores entrepreneurs' exit decisions almost certainly fails to capture its primary effect on firms' values. Our estimates directly characterize these decisions as well as the informational differences between entrepreneurs and policy makers with access only to public information. Hence, they can support the quantitative evaluation of information-constrained policy interventions. One relevant example is the taxation of entrepreneurship.

The owner of an existing firm that expands by creating a new establishment faces a problem similar to that of an entering firm. In both cases, the new venture carries substantial risk. However, the owner of an expanding firm has his previous experience as a guide. Presumably, the new establishment imperfectly copies the firm's older businesses. For the manufacturing sector; Dunne, Roberts, and Samuelson (1988) documented that establishments founded by expanding firms are larger and more likely to survive than other entrants without previous industry experience. The measurement of how these new establishments' first years depend on their parent firms' histories is a natural application for this paper's framework. A related application is the measurement of the sunk costs of exporting, which Roberts and Tybout (1997) characterized for Columbian manufacturing. Our future research will investigate these important aspects of firm growth.

Appendix

The time-invariance of the model's parameters plays a limited role in the proposition's proof. To make this explicit, we consider a slightly more general model in which the parameters vary over time. In obvious notation, the parameters of this model are ν^2, $\bar{\gamma}_T^2$, $\bar{\rho}_T$, $\bar{\sigma}_T^2$, $\bar{\eta}_T^2$, $\bar{\mu}_T$, and $\bar{\alpha}_T$. In this appendix, we prove

Proposition. *The distribution of* $\left(\bar{S}_T, \bar{N}_{T+1}\right)$ *uniquely determines* ν^2; $\bar{\gamma}_T^2$; $\lambda_1\rho_2, \rho_3, \ldots, \rho_T$; $\sigma_1^2 + \eta_1^2, \rho_2^2\Sigma_1 + \sigma_2^2, \sigma_3^2, \ldots, \sigma_T^2$; $\eta_2^2, \ldots, \eta_T^2$; $\mu_1, \mu_2 + \rho_2\mu_1, \mu_3, \ldots, \mu_T$; *and* $\left(\alpha_1 - \mu_1\right)/\sigma_1^2, \alpha_2, \ldots, \alpha_T$.

This proposition implies that all of the model's parameters except μ_1, α_1, σ_1^2, η_1^2, σ_2, and ρ_2 are identified. Complete identification requires one additional restriction, such as $\rho_2 = \rho_3$. For $T \geq 3$, the model of Section 3 imposes this and additional restrictions. Hence the proposition stated in Section 4 follows as a corollary.

The proposition's proof requires some new notation. Recall that $S_t = N_t(X_t + W_t + R)$. Write the transition for X_t conditional on $\bar{X}_{t-1}$ as $X_t = \beta_t\left(\bar{X}_{t-1}\right) + V_t$, with $\beta_t(\bar{X}_{t-1}) \equiv \mu_t + \rho_t\hat{A}_{t-1}$ and $V_t \equiv \rho_t(A_{t-1} - \hat{A}_{t-1}) + Z_t + U_t$. Note that the regression functions β_t can be computed recursively using the Kalman filter and that $\beta_1 = \mu_1$. The disturbance V_t is normally distributed with mean zero and variance $\omega_t^2 \equiv \rho_t^2\Sigma_{t-1} + \sigma_t^2 + \eta_t^2$ and is independent of $\bar{X}_{t-1}$; the processes $\{V_t\}$ and $\{W_t\}$ and mutually independent and independent over time. The recursive specification for the survival process is

$$N_t = \begin{cases} 1 & \text{if } N_{t-1} = 1 \text{ and } X_{t-1} > \tau_{t-1}\left(\bar{X}_{t-2}\right) \\ 0 & \text{otherwise} \end{cases},$$

with $\tau_t\left(\bar{X}_{t-1}\right) \equiv \lambda_t^{-1}\left[\alpha_t - \beta_t\left(\bar{X}_{t-1}\right)\right] + \beta_t\left(\bar{X}_{t-1}\right)$. Note that β_t and τ_t are both affine functions. They are therefore uniquely determined by their restrictions to any set in $\mathbb{R}^{t-1}$ that is not contained in a linear subspace of $\mathbb{R}^{t-1}$.

The proof proceeds in four steps. We first prove

Lemma 1 (Identification of the distribution of intended scales). ν^2 *is identified from the distribution of* $(\bar{S}_2, \bar{N}_2)$.

Next, define scaled sales to be $S_t^* \equiv N_t(S_t - R)$. We exploit Lemma 1 to prove

Lemma 2 (Identification of the scaled data distribution). *The distribution of* $(\bar{S}_T^*, \bar{N}_{T+1})$ *is identified from the distribution of* $(\bar{S}_T, \bar{N}_{T+1})$.

Lemma 2 ensures that we can apply

Lemma 3 (Identification of $\bar{\gamma}_T^2$, $\bar{\omega}_T^2$, $\bar{\beta}_T$, and $\bar{\tau}_T$). *The distribution of* $\left(\bar{S}_T^*, \bar{N}_{T+1}\right)$ *uniquely determines the parameters* $\bar{\omega}_T^2$ *and* $\bar{\gamma}_T^2$ *and the functions* $\bar{\beta}_T$ *and* $\bar{\tau}_T$.

With these lemmas in place, Lemma 4 delivers the proposition's conclusion.

Lemma 4 (Identification of the structural parameters from $\bar{\omega}_T$, $\bar{\beta}_T$, and $\bar{\tau}_T$). *Given* $\bar{\omega}_T$, $\bar{\beta}_T$, *and* $\bar{\tau}_T$, *the following parameters of the structural model are uniquely determined:* $\lambda_1\rho_2, \rho_3, \ldots, \rho_T$; $\sigma_1^2 + \eta_1^2, \rho_2^2\Sigma_1 + \sigma_2^2, \sigma_3^2, \ldots, \sigma_T^2$; $\eta_2^2, \ldots, \eta_T^2$; $\mu_1, \mu_2 + \rho_2\mu_1, \mu_3, \ldots, \mu_T$; *and* $\left(\alpha_1 - \mu_1\right)/\sigma_1^2, \alpha_2, \ldots, \alpha_T$.

We now present the lemmas' proofs.

Proof of Lemma 1. Write the regression of X_1 on S_1 as $X_1 = \iota + \zeta S_1 + E_1$, where $\zeta \equiv \omega_1^2 / \left(\omega_1^2 + \gamma_1^2 + \nu^2\right)$, $\iota \equiv \beta_1 (1-\zeta)$, and $E_1 \equiv (1-\zeta)(X_1 - \beta_1) - \zeta (W_1 + R)$. By construction, $E_1 \perp\!\!\!\perp S_1$, $\mathbb{E}[E_1] = 0$, and $\mathbb{V}[E_1] = \chi^2 \equiv \omega_1^2 (1-\zeta)$. Hence, we can write

$$\Pr(N_2 = 1 | S_1) = \Phi\left(\frac{\iota + \zeta S_1 - \tau_1}{\chi}\right).$$

Because $\Pr(N_2 = 1|S_1)$ is data, this identifies $\chi/\zeta = \gamma_1^2 + \nu^2$ and $(\iota - \tau_1)/\chi$. Because $\mathbb{V}[S_1] = \omega_1^2 + \gamma_1^2 + \nu^2$ and $\mathbb{E}[S_1] = \beta_1$ are data, we can also identify ω_1^2, ζ, and χ^2.

Next, consider the autoregression of S_2 on S_1 without accounting for selection, $\mathbb{E}[S_2|S_1] = \xi + \psi S_1$. Here, $\psi \equiv \left(\beta_{21}\omega_1^2 + \nu^2\right) / \left(\omega_1^2 + \nu^2 + \gamma_1^2\right)$ and $\xi \equiv \beta_2(\beta_1) - \psi\beta_1$, with $\beta_{21} \equiv d\beta_2(x)/dx = \rho_2\lambda_1$. The residual from this regression is $E_2 \equiv (\beta_{21} - \psi)(X_1 - \beta_1) + (1-\psi)R + V_2 + W_2 - \psi W_1$. We have that

$$\mathbb{E}[E_1 E_2] = (1-\zeta)(\beta_{21} - \psi)\omega_1^2 + \zeta\psi\left(\gamma_1^2 + \nu^2\right) - \zeta\nu^2.$$

Because ζ/χ, $(\iota - \tau_1)/\chi$, and χ are identified and (see Heckman, 1979)

$$\mathbb{E}[S_2|S_1, N_2 = 1] = \xi + \psi S_1 + \frac{\mathbb{E}[E_1 E_2]}{\chi} \frac{\phi\left(\frac{\iota + \zeta S_1 - \tau_1}{\chi}\right)}{\Phi\left(\frac{\iota + \zeta S_1 - \tau_1}{\chi}\right)}$$

both ψ and $\mathbb{E}[E_1 E_2]$ are identified. Using that ζ, ω_1^2, and $\gamma_1^2 + \nu^2$ are already known, this identifies ν^2. □

Proof of Lemma 2. Denote the densities of $\bar{S}_t$ and $\bar{S}_t^*$ on $\{N_t = 1\}$ by $f_{\bar{S}_t}(\cdot|N_t = 1)$ and $f_{\bar{S}_t^*}(\cdot|N_t = 1)$, respectively. Because R is independent of $(\bar{S}_t^*, \bar{N}_{T+1})$, we have that

$$f_{\bar{S}_t}\left(s_1, \ldots, s_t | \bar{N}_{T+1}\right) = \int_{-\infty}^{\infty} f_{\bar{S}_t^*}\left(s_1 - r, \ldots, s_t - r | \bar{N}_{T+1}\right) \phi\left(\frac{r}{\nu}\right) dr \tag{15}$$

on $\{N_t = 1\}$. The left-hand side of (15) is data and ν is identified by Lemma 1. Thus, a standard deconvolution argument establishes that $f_{\bar{S}_t^*}(\cdot|\bar{N}_{T+1})$ is identified on $\{N_t = 1\}$.[16] This immediately identifies the distributions of $\bar{S}_T^*$ on $\{N_{T+1} = 1\}$ and $\{N_T = 1, N_{T+1} = 0\}$. Using that $S_{t+1}^* = \cdots = S_T^* = 0$ on $\{N_{t+1} = 0\}$, we can in addition identify the distribution of $\bar{S}_T^*$ on $\{N_t = 1, N_{t+1} = 0\}$ for $t = 1, \ldots, T-1$. Because the distribution of $\bar{N}_{T+1}$ is data, this identifies the distribution of $(\bar{S}_T^*, N_{T+1})$. □

Proof of Lemma 3. The proof proceeds recursively. It is helpful to define $\vec{S}_t^*$ and $\vec{N}_{t+1}$ to be the random vectors $(S_t^*, \ldots, S_T^*)$ and $(N_{t+1}, \ldots, N_{T+1})$. Denote their joint density with $\bar{X}_{t-1}$ by $f_{\vec{S}_t^*, \vec{N}_{t+1}, \bar{X}_{t-1}}(\cdot)$. We begin with the assumption that this density for a particular age t is known. For $t = 1$, $f_{\vec{S}_1^*, \vec{N}_2, \bar{X}_0}(\cdot)$ is simply the density of the data. We first show that β_t, ω_t^2, τ_t and f_{W_t} can be recovered from $f_{\vec{S}_t^*, \vec{N}_{t+1}}\left(\cdot|\bar{X}_{t-1}\right)$. We then demonstrate that these parameters in turn identify $f_{\vec{S}_{t+1}^*, \vec{N}_{t+2}, \bar{X}_t}(\cdot)$ if $t < T$, allowing the recursion to continue.

We begin with the identification of β_t, ω_t^2 and τ_t. This requires only the knowledge of the expected value of S_t^* conditional on $\bar{X}_{t-1}$, the conditional probability of survival, and the expected value of S_t^* given this

[16] See Feller (1971) for an introduction to deconvolution. Formally, the right-hand side of (15) is a t-variate convolution of $f_{\bar{S}_t^*}(\cdot|N_t = 1)$ and the singular t-variate density of $(R, \ldots, R)$.

history *and* survival to period $t+1$. On $\{N_t = 1\}$, these are

$$\begin{aligned}
\mathbb{E}[S_t^* | \bar{X}_{t-1}] &= \beta_t(\bar{X}_{t-1}), \\
\mathbb{E}[N_{t+1} | \bar{X}_{t-1}] &= 1 - \Phi\left(\frac{\tau_t(\bar{X}_{t-1}) - \beta_t(\bar{X}_{t-1})}{\omega_t}\right), \text{ and} \\
\mathbb{E}[S_t^* | \bar{X}_{t-1}, N_{t+1} = 1] &= \beta_t(\bar{X}_{t-1}) + \omega_t \frac{\phi\left(\frac{\tau_t(\bar{X}_{t-1}) - \beta_t(\bar{X}_{t-1})}{\omega_t}\right)}{1 - \Phi\left(\frac{\tau_t(\bar{X}_{t-1}) - \beta_t(\bar{X}_{t-1})}{\omega_t}\right)},
\end{aligned}$$

where $\Phi(\cdot)$ and $\phi(\cdot)$ are the cumulative distribution function and density function of a standard normal random variable. Because the distribution of (S_t^*, N_{t+1}) conditional on $\bar{X}_{t-1}$ is assumed to be known, these three equations immediately yield ω_t^2 and the restrictions of β_t and τ_t to the support of $\bar{X}_{t-1}$ on $\{N_t = 1\}$. Because the support of $\bar{X}_{t-1}$ on $\{N_t = 1\}$ is not contained in any linear subspace of $\mathbb{R}^{t-1}$, this identifies β_t and τ_t.

To obtain γ_t^2, note that the probability density of S_t^* conditional on $\bar{X}_{t-1}$ can be written as

$$f_{S_t^*}\left(s | \bar{X}_{t-1}\right) = \int_{-\infty}^{\infty} \frac{1}{\omega_t} \phi\left(\frac{x - \beta_t\left(\bar{X}_{t-1}\right)}{\omega_t}\right) \phi\left(\frac{s - x}{\gamma_t}\right) dx \tag{16}$$

on $\{N_t = 1\}$. Applying a standard deconvolution argument establishes that γ_t^2 is identified.

We will now show that $f_{\vec{S}_{t+1}^*, \vec{N}_{t+2}, \bar{X}_t}(\cdot)$ is identified. The independence of W_t from $\vec{W}_{t+1}$ and $\bar{X}_T$ implies that the joint distribution of W_t and X_t conditional upon $\left(\bar{X}_{t-1}, \vec{S}_{t+1}^*, \vec{N}_{t+1}\right)$ displays independence. Therefore, the probability density of S_t^* conditional on these variables is

$$f_{S_t^*}\left(s | \bar{X}_{t-1}, \vec{S}_{t+1}^*, \vec{N}_{t+1}\right) = \int_{-\infty}^{\infty} f_{X_t}\left(x | \bar{X}_{t-1}, \vec{S}_{t+1}^*, \vec{N}_{t+1}\right) \phi\left(\frac{s - x}{\gamma_t}\right) dx \tag{17}$$

on $\{N_t = 1\}$. The left-hand side of (17) and f_{W_t} are known, so deconvolution yields $f_{X_t}\left(\cdot | \bar{X}_{t-1}, \vec{S}_{t+1}^*, \vec{N}_{t+1}\right)$ on $\{N_t = 1\}$. We can immediately recover the distribution of $\left(\bar{X}_t, \vec{S}_{t+1}^*, \vec{N}_{t+1}\right)$ on $\{N_{t+1} = 1\} \subset \{N_t = 1\}$ by multiplying this conditional distribution by the known joint distribution of the conditioning variables. Using that both $\vec{S}_{t+1}^*$ and $\vec{N}_{t+1}$ are identically zero on $\{N_{t+1} = 0\}$, we can then construct the distribution of $\left(\bar{X}_t, \vec{S}_{t+1}^*, \vec{N}_{t+1}\right)$. Thus, the recursion may continue. □

Proof of Lemma 4. This proof is straightforward. It is available from the authors upon request. □

References

Bahk, B. and M. Gort (1993). Decomposing learning by doing in new plants. *Journal of Political Economy 101*, 561–583.

Bierens, H. J. (1987). Kernel estimators of regression functions. In T. Bewley (Ed.), *Advances in Econometrics. Fifth World Congress*, Volume I. Cambridge: Cambridge University Press.

Blundell, R. and S. Bond (1998). Initial conditions and moment restrictions in dynamic panel data models. *Journal of Econometrics 87*, 115–143.

Blundell, R. and I. Preston (1998). Consumption inequality and income uncertainty. *Quarterly Journal of Economics 113*, 603–640.

Campbell, J. R. and H. A. Hopenhayn (2005). Market size matters. *Journal of Industrial Economics 53*(1), 1–25.

Campbell, J. R. and B. Lapham (2004). Real exchange rate fluctuations and the dynamics of retail trade industries on the U.S.-Canada border. *American Economic Review 94*, 1194–1206.

Davis, S. J. and J. Haltiwanger (1992). Gross job creation, and employment reallocation. *Quarterly Journal of Economics 107*, 820–863.

Dixit, A. K. and J. E. Stiglitz (1977). Monopolistic competition and optimum product diversity. *American Economic Review 67*, 297–308.

Dunne, T., M. J. Roberts, and L. Samuelson (1988). Patterns of firm entry and exit in U.S. manufacturing industries. *Rand Journal of Economics 19*, 495–515.

Feller, W. (1971). *An Introduction to Probability Theory and Its Applications* (2nd ed.), Volume II. New York: Wiley.

Gort, M. and S. Klepper (1982). Time paths in the diffusion of product innovations. *Economic Journal 92*, 630–653.

Hart, O. (1985). Monopolistic competition in the spirit of Chamberlin: A general model. *Review of Economic Studies 52*, 529–546.

Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica 47*, 153–161.

Hochguertel, S. (2004). The dynamics of self-employment and household wealth: New evidence from panel data. Mimeo, Free University, Amsterdam.

Holmes, T. J. and J. A. Schmitz, Jr. (1995). On the turnover of business firms and business managers. *Journal of Political Economy 103*, 1005–1038.

Hopenhayn, H. A. (1992). Entry, exit, and firm dynamics in long run equilibrium. *Econometrica 60*, 1127–1150.

Hurst, E. and A. Lusardi (2004). Liquidity constraints, household wealth, and entrepreneurship. *Journal of Political Economy 112*(2), 319–347.

Jovanovic, B. (1982). Selection and the evolution of industry. *Econometrica 50*, 649–670.

Jovanovic, B. and G. M. MacDonald (1994). The life cycle of a competitive industry. *Journal of Political Economy 102*, 322–347.

Kitagawa, G. (1987). Non-Gaussian state-space modeling of nonstationary time series. *Journal of the American Statistical Association 82*, 1032–1041.

Levinsohn, J. and A. Petrin (2003). Estimating production functions using inputs to control for unobservables. *Review of Economic Studies 70*, 317–341.

Lucas, Jr., R. E. (1978). On the size distribution of firms. *Bell Journal of Economics 9*, 508–523.

Moore, K. (2004). Do liquidity constraints matter for new entrepreneurs? Finance and Economics Discussion Series 2004-42, Federal Reserve Board, Washington, D.C.

Olley, G. S. and A. Pakes (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica 64*, 1263–1298.

Pakes, A. and R. Ericson (1998). Empirical implications of alternative models of firm dynamics. *Journal of Economic Theory 79*, 1–45.

Pavcnik, N. (2002). Trade liberalization, exit, and productivity improvement: Evidence from Chilean plants. *Review of Economic Studies 69*, 245–276.

Powell, J. L., J. H. Stock, and T. M. Stoker (1989). Semiparametric estimation of index coefficients. *Econometrica 57*, 1403–1430.

Prescott, E. C. and M. Visscher (1977). Sequential location among firms with foresight. *Bell Journal of Economics 8*, 378–393.

Roberts, M. J. and J. R. Tybout (1997). The decision to export in Columbia: An empirical model of entry with sunk costs. *American Economic Review 87*, 545–564.

Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica 55*, 999–1033.

Rust, J. (1994). Structural estimation of Markov decision processes. In R. Engle and D. McFadden (Eds.), *Handbook of Econometrics*, Volume 4, pp. 3081–3143. Amsterdam: North-Holland.

Salop, S. C. (1979). Monopolistic competition with outside goods. *Bell Journal of Economics 10*, 141–156.

Wolinsky, A. (1986). True monopolistic competition as a result of imperfect information. *Quarterly Journal of Economics 101*, 493–512.

Table 1: Summary Statistics from the First 13 Months in the Lives of New Texas Bars[(i)]

Age[(ii)]	Firms	Sales' Mean[(iii)]	Sales' Standard Deviation	Exit Rate[(iv)]
1	305	0.06	0.87	0.003
2	304	0.08	0.89	0.016
3	299	0.10	0.86	0.007
4	297	0.11	0.85	0.020
5	291	0.12	0.85	0.017
6	286	0.11	0.87	0.000
7	286	0.11	0.89	0.028
8	278	0.17	0.84	0.004
9	277	0.17	0.86	0.032
10	268	0.20	0.86	0.019
11	263	0.20	0.89	0.042
12	252	0.23	0.89	0.032
13	244	0.27	0.87	0.000

Notes: (i) See the text for details regarding the sample's construction. (ii) Age is measured in months and equals one for a firm filing its first tax return. (iii) Sales are measured in logarithms and are relative to the sales of *all* establishments selling alcoholic beverages for on-premise consumption, whether or not we classify them as bars. (iv) The exit rate is defined as the number of firms operating in month t that do not operate in month $t+1$ divided by the number of firms operating in month t.

Table 2: Regression Estimates of Sales on Previous and First Months' Sales(i)

	Logarithm of Sales in	
Age(ii)	Previous Month	First Month
3	0.809	0.154
	(0.010)	(0.010)
4	0.778	0.190
	(0.008)	(0.008)
5	0.942	0.021
	(0.008)	(0.008)
6	0.859	0.088
	(0.008)	(0.008)
7	0.880	0.086
	(0.010)	(0.009)
8	0.851	0.132
	(0.009)	(0.008)
9	0.905	0.073
	(0.005)	(0.005)
10	0.916	0.058
	(0.006)	(0.006)
11	0.894	0.062
	(0.006)	(0.006)
12	0.830	0.122
	(0.009)	(0.007)
13	0.822	0.168
	(0.008)	(0.007)

Notes: (i) For the third to thirteenth months, this table reports Powell, Stock, and Stoker's (1989) instrumental variable density-weighted average derivative estimates for single-index regression models of log sales on the logarithms of sales in the previous and first months. Standard errors are reported in parentheses below each estimate. For each month, the estimation was conducted using the sample of firms that survived to that month. (ii) Age is measured in months and equals one for a firm filing its first tax return.

Table 3: Maximum Likelihood Estimates[(i)]

Parameter	σ_1	σ	ρ	γ	ν	μ_1	μ
Estimate	0.714	0.161	0.985	0.167	0.581	0.069	0.0014
Standard Error[(ii)]	(0.076)	(0.005)	(0.000)	(0.004)	(0.000)	(0.055)	(0.0030)

Parameter	α_1	α_2	α_3	α_4	α_5	α_6
Estimate	-1.65	-1.53	-1.44	-1.38	-1.34	-2.50
Standard Error[(ii)]	(0.068)	(0.038)	(0.049)	(0.029)	(0.033)	$(\cdot)$

Parameter	α_7	α_8	α_9	α_{10}	α_{11}	α_{12}
Estimate	-1.22	-1.26	-0.83	-0.79	-0.64	-0.54
Standard Error[(ii)]	(0.032)	(0.063)	(0.040)	(0.039)	(0.040)	(0.025)

Notes: (i) This specification assumes that entrepreneurs face no signal extraction problem ($\eta^2 = 0$). (ii) The standard errors are calculated using an estimate of the information matrix based on the outer product of the scores. See the text for further details.

Table 4: Model Summary Statistics[(i)]

Age[(ii)]	Firms[(iii)]	Sales' Mean	Sales' Standard Deviation	Exit Rate[(iv)]
1	100.0	0.07	0.94	0.008
2	99.2	0.08	0.93	0.007
3	98.5	0.10	0.93	0.007
4	97.8	0.11	0.92	0.007
5	97.1	0.12	0.92	0.006
6	96.5	0.13	0.92	0.000
7	96.5	0.13	0.93	0.016
8	95.0	0.15	0.92	0.006
9	94.5	0.16	0.92	0.069
10	88.0	0.25	0.89	0.026
11	85.6	0.28	0.88	0.053
12	81.1	0.33	0.86	0.049

Notes: (i) The reported statistics are calculated from a synthetic sample of $100,000$ firm histories drawn from the estimated model. (ii) Age is measured in months and equals one for a newly born firm. (iii) The number of firms is measured in thousands. (iv) The exit rate is defined as the number of firms operating in month t that do not operate in month $t+1$ divided by the number of firms operating in month t.

Table 5: Model Population Regression of Sales on Previous and First Months' Sales[(i)]

Age[(ii)]	Logarithm of Sales in Previous Month	Logarithm of Sales in First Month
3	0.663	0.325
4	0.705	0.238
5	0.765	0.191
6	0.797	0.150
7	0.818	0.142
8	0.842	0.119
9	0.867	0.123
10	0.839	0.114
11	0.872	0.092
12	0.875	0.075

Notes: (i) For the third to twelfth months, this table reports the density-weighted average derivatives from the regression of log sales on the logarithms of sales in the previous and first months for surviving firms. These are calculated by applying Powell, Stock, and Stoker's (1989) estimator to a synthetic sample of $100,000$ firm histories drawn from the estimated model. (ii) Age is measured in months and equals one for a newly born firm.

Table 6: Estimated Fixed Costs as Fractions of First Month's Producer's Surplus[(i)]

Age[(ii)]	1	2	3	4	5	6
$\mathbb{E}[e^R]\epsilon\kappa/\mathbb{E}[e^{S_1}]$	10.23	0.19	0.24	0.30	0.33	0.44
	(1.80)	(0.06)	(0.05)	(0.06)	(0.06)	(0.05)

Age[(ii)]	7	8	9	10	11	12
$\mathbb{E}[e^R]\epsilon\kappa/\mathbb{E}[e^{S_1}]$	0.06	0.28	0.21	0.68	0.55	1.09
	(0.00)[(iii)]	(0.03)	(0.02)	(0.19)	(0.14)	(0.40)

Notes: (i) This table reports the minimum-distance estimates of each month's scaled fixed costs as described in Subsection 4.3. (ii) Age is measured in months and equals one for a firm filing its first tax return. (iii) The reported standard error is less than 0.01.

Figure 1: Maximum Likelihood Estimation of η^2