

A Structural Empirical Model of Firm Growth, Learning, and Survival

Jaap H. Abbring*
Jeffrey R. Campbell†

April, 2003

Abstract

We present a structural model of firm growth, learning, and survival and consider its identification and estimation. In the model, entrepreneurs have private and possibly error-ridden observations of persistent and transitory shocks to profit. We demonstrate that the model's parameters can be recovered from public observations of sales and survival, and we estimate them using monthly data from new bars in Texas. We find that entrepreneurs observe profit's persistent component without error. In this sense, their information is substantially superior to the public's.

*Department of Economics, Free University, De Boelelaan 1105, 1081 HV Amsterdam, The Netherlands.
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 dynamics, Bayesian learning, identification, estimation

JEL codes: C34, C35, C41, D83, L11

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 seminar participants at University College London 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

We present in this paper a structural model of firm growth, learning, and survival; and we consider its identification and estimation using a panel of Texas bars' monthly sales histories. As in panel data commonly available from company directories or tax returns, these histories record each firm's sales from its birth until the end of the sampling period or its exit. In the model a firm's sales is a linear and increasing function of its profit, but transitory shocks make it an imperfect indicator of profit's persistent component. Entrepreneurs observe some of these shocks, so their forecasts of future profit are necessarily superior to those of the public. Maximum likelihood estimation of the model's parameters measures the extent of entrepreneurs' Bayesian learning while accounting for their private information. We find that the entrepreneurs in our sample observe profit's persistent component without error. In this sense, their information is substantially superior to the public's.

Our analysis begins with an extension of Jovanovic's (1982) model of firm selection. The logarithm of each firm's profit is the sum of three components, one that displays persistence and two others that are transitory. Unlike in Jovanovic's original model, the persistent component possibly changes through time. After the entrepreneur observes her profit, she must decide either to remain open or to close the firm and avoid future fixed costs. An entrepreneur only observes her profit and *one* of its transitory components, and she bases her continuation decision on her posterior belief regarding the persistent component. We follow Jovanovic and assume that the components 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. The data available to the econometrician are these continuation decisions' outcomes and the history of the firm's sales. The econometrician cannot predict the firm's future as well as the entrepreneur, because he observes neither of profit's transitory components.

The characterization of entrepreneurs' private information using only their discrete continuation decisions and public sales data is not straightforward, because the information relevant for their continuation decisions is both *hidden* from the econometrician and *persistent*. The hidden state variable distinguishes our analysis from estimation procedures that use observations of a continuous forward-looking choice to infer a relevant state variable's value, as in Olley and Pakes (1996). Our work also differs from structural dynamic discrete choice estimation that incorporates Rust's (1987) conditional-independence assumption, because the hidden state variable displays persistence even after conditioning on the entrepreneur's

observed choice.

Although continuation decisions do not reveal the relevant state variable, the model's parameters can be recovered from firm histories such as those we observe. We show that differences between all firms' average sales and the average sales of those firms that subsequently survive reveals 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 reflects shocks that the entrepreneur perceives to be irrelevant for the firm's future. This step only requires that the entrepreneur bases her decisions on her private information. It does *not* require the entrepreneur to form rational expectations or to make optimal decisions. The identified distributions and exit rules uniquely determine the parameters relevant for the entrepreneur's signal extraction problem.

Our data contain monthly records of the first year in the lives of 305 new firms that operated a single bar in Texas. To calculate maximum likelihood estimates using this data, we augment the Kalman filter to account for sample selection. The resulting estimates measure the variances of the shocks that contaminate the entrepreneur's and econometrician's observations of profit's persistent component. We find that the entrepreneur observes this component without error.

Pakes and Ericson (1998) precede us in the empirical detection of Bayesian learning by entrepreneurs. One observable implication of Bayesian learning about a fixed parameter is that a firm's initial size is useful for forecasting its size throughout its life. Pakes and Ericson formally develop this idea and nonparametrically test for such Bayesian learning by estimating the relevance of initial size in a firm size regression. With private entrepreneurial information a firm's initial size may improve the public's forecast even without entrepreneurial learning. We apply a test similar in spirit to Pakes and Ericson's to our sample of Texas bars and find that initial size does indeed improve forecasts of firm sizes. However, the estimates of our structural model's parameters indicate that entrepreneurs' private information can account for this finding without reference to entrepreneurial learning.

Our estimates indicate that entrepreneurs face no signal extraction problem. However, a policy maker with access to only public information should use Bayesian updating to learn about a particular firm's profit. Because our estimates directly characterize the informational differences between such a policy maker and an entrepreneur, they can serve as an important input into the evaluation of information-constrained policy interventions such as the taxation of entrepreneurship. We defer the analysis of such policies to future research.

The remainder of the paper proceeds as follows. The next section describes the data and some of its salient characteristics. Section 3 presents the structural model, and Sections 4 and

5 consider its identification and estimation. Section 6 presents and discusses our maximum likelihood estimates, and 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 (1989), 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. To account for inflation over the sample period and persistent differences in input prices across locations, 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. 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

¹In 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 simply legal reorganizations with few immediate economic implications.

²Some 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

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 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 to rely upon. 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.

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.

³The 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.

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, standard deviation, skewness, and kurtosis of the logarithm of sales 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. The exit rates vary greatly over the course of the first year. Exit rates near the end of the firm's first year are relatively high, and no firms in our sample exited following their sixth and thirteenth months.

The lack of exits following the thirteenth month is particularly notable because it continues into the (unreported) fourteenth month. One possible explanation for this sudden and persistent decrease in exit is that the renewal of an annual lease on a commercial location is a watershed event that induces many unprofitable businesses to exit after twelve months. In this case, firms that survived their first year were relatively fit and unlikely to have immediately exited.

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 cross-sectional standard deviation changes little over the first year. The skewness of sales is very close to zero except in the eighth and tenth months, while the kurtosis varies between a high value of 3.53 in the first month and a low value of 3.08 in the fifth month. To better understand the statistical significance of these estimates, we have tabulated 95% bootstrap confidence intervals for all of the statistics in Table 1. These (unreported) confidence intervals reveal that the skewness and kurtosis coefficients are estimated much less precisely than the means and standard deviations. The average width of a confidence interval for a mean is 0.20, while the analogous average for the standard deviations' confidence intervals is 0.16. In contrast, the average widths of the skewness and kurtosis coefficients' confidence intervals are 0.62 and 1.20.⁴

To assess the role of selection in the evolution of these summary statistics, it is helpful to inspect the analogous statistics from those firms that survived the entire sample period. Table 2 reports these statistics for the 244 firms that operated for thirteen or more months. There are no striking differences between the selected sample's second and higher moments and

⁴The bootstrap confidence intervals were constructed from 10,000 random samples of 305 firms drawn with replacement from our data set.

those of the complete sample. However, the means of the selected sample are considerably higher than their analogues from the complete sample. The average log sales in the first month of a firm that survives its first year is eleven percentage points higher than the unconditional average, 0.17 versus 0.06. Thus, selection and the growth of surviving firms each account for approximately one half of the complete sample’s mean growth.

2.3 Firm Growth

As in many other panels of firm histories, entrepreneurs’ exit decisions clearly contribute to the evolution of our sample. To better understand the importance of entrepreneurial learning for this selection and other forward-looking decisions, Pakes and Ericson (1998) advocate examining the persistence of surviving firms’ sales. In particular, they derive robust 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, a firm’s initial sales becomes progressively less relevant for forecasting its future as it ages. Pakes and Ericson test these models’ contrasting predictions using panels of Wisconsin retail and manufacturing firms. They find that initial sales improves forecasts of retailers’ future sales, but manufacturers’ future sales appear to be Markovian. From this, 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. For each month of the life cycle that we consider, we have nonparametrically characterized the regression of the logarithm of a firm’s sales on the logarithms of its sales in the previous and first months. For this we estimated density-weighted average derivatives of the regression functions using Powell, Stock, and Stoker’s (1989) nonparametric instrumental-variables estimator.⁵ 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.

⁵To implement this estimation, we follow Powell, Stock, and Stoker’s (1998) recommendation and use 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.

Table 3 reports these estimates as well as standard errors based on their Gaussian asymptotic distributions. All of the derivative estimates are positive and statistically significant at 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. Overall, the estimates indicate that no low-order Markov process can fit surviving firms' observed sales well. Expanding the set of regressors to include the past three months' sales attenuates the estimated effect of the first month's sales, but does not alter this conclusion.

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 our sample of new firms is necessarily suspect. The structural analysis we will pursue next overcomes this difficulty by using entrepreneurs' observed continuation decisions to measure the variation in sales due to those transitory shocks that the entrepreneur observes.

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. We assume that potential entrants are *ex ante* identical and that firms only compete *anonymously* *ex post*. 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 first discuss this assumption's content and empirical plausibility. We then detail the firm's stochastic environment and optimization problem. Finally, we consider the entrepreneur's procedure for optimally assessing her firm's future and deciding upon its survival.

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. These models inherit the simplicity of perfect competition. The actions of any single producer are irrelevant for any competitor's profits. This irrelevance immediately implies that idiosyncratic shocks to demand and cost have no effect on competitors' profits or optimal actions. Therefore, the empirical analysis of a monopolistic competition model can proceed by considering each producer's choice problem in isolation from those of her rivals after conditioning on the appropriate aggregate variables.

In contrast, 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. The presence of these indirect effects complicates these models' empirical analysis.⁶

Campbell and Hopenhayn (2002) suggest a simple procedure for distinguishing between these two approaches to imperfect competition based on cross-market comparisons of the producer size distribution. A robust prediction of anonymous monopolistic competition is that the producer size distribution is invariant to market size if factor prices, demographics, and technology are held constant.⁷ This prediction follows directly from the failure of entry to erode monopolistic competitors' market power. In contrast, models of oligopolistic competition generally predict that producers in large and competitive markets recover their fixed costs by selling more at a lower markup than their counterparts in small and less competitive markets. Thus, a comparison of average producer size across large and small markets can

⁶The observation that consumers view similar goods available at different geographic locations as imperfect substitutes does *not* immediately imply that competition is oligopolistically competitive. Campbell and Lapham (2001) illustrate this with a model of monopolistic competition and cross-border shopping. Although competition is anonymous in that model, goods produced in the same location are better substitutes than are goods produced in different locations.

⁷Campbell (2002) demonstrates that this familiar feature of symmetric monopolistic competition models with free entry and a single technology is robust to introducing product placement and technology choice decisions of arbitrary complexity.

determine which of these two approaches to competition among large numbers of producers is more empirically promising. 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 in the majority of industries that Campbell and Hopenhayn consider, 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. The remainder of our analysis rests accordingly on the assumption of monopolistic competition.

3.2 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. Idiosyncratic shocks cause the demand curve for the firm’s product to vary over time. In period t , it is $Q_t = e^{X_t+W_t} P_t^{-\varepsilon}$, where $\varepsilon > 1$ is the absolute value of the firm’s demand elasticity and P_t is its price. The composite random variable $X_t + W_t$ shifts the demand curve’s location 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, $\vartheta \cdot Q_t + \kappa_t$, describes the firm’s technology, with both ϑ and κ_t strictly greater than zero. In our application to Texas’ bars, time variation in the fixed cost κ_t may reflect the fact noted in Section 2 that leases on bars typically have fixed terms. 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

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

and

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

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

The random variables W_t and X_t represent the transitory and persistent components of the firm's profit. The transitory component is independent through time and continuously distributed with density f_{W_t} in period t . We assume that this random variable has mean zero, but currently we leave its distribution otherwise unspecified. The persistent component follows a Gaussian linear state-space model with parameters that (potentially) depend on the firm's age. That is

$$A_t = \mu_t + \rho_t A_{t-1} + Z_t \quad \text{with } Z_t \sim \mathcal{N}(0, \sigma_t^2) \quad (3)$$

and

$$X_t = A_t + U_t \quad \text{with } U_t \sim \mathcal{N}(0, \eta_t^2). \quad (4)$$

In (3), $\rho_1 \equiv 0$ and $\rho_t > 0$ for all $t > 1$, so that A_t and X_t are positively autocorrelated. The disturbance processes $\{Z_t\}$ and $\{U_t\}$ are independent over time and mutually independent of each other. The intercept μ_t reflects systematic firm growth as the firm ages, such as that arising from passive learning about the production process.

For concreteness, we have assumed that variation in the firm's profit reflects only demand shocks and that the firm's marginal cost is constant. We could alternatively have assumed that demand is constant and marginal cost is stochastic, or that both are stochastic. If we only have data on sales and survival, all of these models are observationally equivalent. In this paper, we restrict attention to the analysis of such data and we do not address the separate identification of idiosyncratic demand and cost shocks. If we also observed firms' price choices, the assumption of a constant percentage markup would allow us to identify separate shocks to demand and marginal cost.

3.3 Bayesian learning and selection

We assume that the entrepreneur observes W_t and X_t separately at the beginning of period t , but she observes neither A_t nor U_t . 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 information at hand, $(X_1, \dots, X_t) \equiv \bar{X}_t$. Denote this optimal forecast and its mean squared error with

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

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

$$\hat{A}_t = \mu_t + \rho_t \hat{A}_{t-1} + \lambda_t (X_t - \mu_t - \rho_t \hat{A}_{t-1}), \quad (5)$$

and

$$\Sigma_t = \eta_t^2 \lambda_t, \quad (6)$$

where

$$\lambda_t \equiv \frac{\rho_t^2 \Sigma_{t-1} + \sigma_t^2}{\rho_t^2 \Sigma_{t-1} + \sigma_t^2 + \eta_t^2}. \quad (7)$$

The initialization of (7) uses $\Sigma_0 = 0$. The coefficient λ_t is the Kalman gain, and it measures the informativeness of the entrepreneurs' observation of X_t . The firm's sales reflects both X_t and W_t , so the entrepreneur's inference of A_t is necessarily superior to that of an econometrician using only the current and past realizations of S_t .

Although the entire history of X_t is in principle relevant for the entrepreneur's exit decision, \hat{A}_t and the firm's age are sufficient for characterizing the distribution of future profits given the information in \bar{X}_t . Define $v_t(\hat{A}_t, S_t)$ to be the value of a firm of age t 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

$$v_t(\hat{A}_t, S_t) = \varepsilon^{-1} e^{S_t} - \kappa_t + \delta \max\{0, \mathbb{E}_t[v_{t+1}(\hat{A}_{t+1}, S_{t+1})]\}. \quad (8)$$

The entrepreneur calculates the expectation \mathbb{E}_t in (8) 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) is sufficient.

The owner's optimal exit policy is simple. The expectation in the right-hand side of (8) is a function of (t, \hat{A}_t) and it is continuous and increasing in \hat{A}_t . Therefore, there exists a threshold value α_t such that the owner chooses to exit if and only if $\hat{A}_t \leq \alpha_t$. Define N_t to be an indicator variable that equals one if the firm has survived through that date and equals zero otherwise. Given the exit rule, this is generated recursively by

$$N_t = \begin{cases} 1 & \text{if } N_{t-1} = 1 \text{ and } \hat{A}_{t-1} > \alpha_{t-1} \\ 0 & \text{otherwise} \end{cases}. \quad (9)$$

This specification of our model assumes that the market is in a stationary equilibrium, as in Hopenhayn (1992), so that aggregate disturbances change neither the firm's demand nor its cost structure. This assumption is undesirable if aggregate shocks are important for the industry under consideration. However, it is straightforward to generalize the model to allow the evolution of the sales process and the chosen exit policy to depend on aggregate statistics such as the number of producers per customer.

4 Identification

Suppose an econometrician observes a set of firms from our model for T periods. On each date, the econometrician records both S_t and N_t . For completeness, we define $S_t = 0$ if $N_t = 0$. At the end of period T , the econometrician also records N_{T+1} . That is, the firms' survival at the end of the sample period is known. In this section, we investigate which of the parameters of our structural model can be recovered from the joint distribution of these observations.

In our model, the exit thresholds $(\alpha_1, \dots, \alpha_T) \equiv \bar{\alpha}_T$ represent the entrepreneur's optimal exit decisions. Unlike Rust's (1987) nested fixed-point algorithm, our approach to identification and estimation does not rely on an explicit solution to the entrepreneur's dynamic programming problem. Rather, we treat these thresholds as parameters to be estimated jointly with parameters that determine the evolution of S_t ; $\bar{\mu}_T$, $\bar{\rho}_T$, $\bar{\sigma}_T^2$ and $\bar{\eta}_T^2$, and the unknown densities $(f_{W_1}, \dots, f_{W_T})$. The statistical problem we face is to identify and estimate the true values of these parameters and functions using the realizations of $(\bar{S}_T, \bar{N}_{T+1})$ from a large cross-section of firms. We do not identify or estimate the parameters that influence the observable data only through the choice of $\bar{\alpha}_T$. These are δ , ε , and $\bar{\kappa}_T$.

We begin by placing our model into a less restrictive encompassing model in which neither the latent variable A_t nor the entrepreneur's Bayesian learning play an explicit role in the evolution of X_t , S_t , and N_t . Rust (1994) refers to such a model that directly characterizes dynamic decisions without reference to the parameters of technology or preferences as a *reduced-form* model, and we adopt this terminology here. We first show that the parameters of the reduced-form model are identified. Then, we prove that these parameters can be used to recover all of our structural model's parameters except μ_1 , α_1 , σ_1^2 , η_1^2 , σ_2 , and ρ_2 . The identification of these remaining parameters requires one additional restriction. This sequential approach to identification broadens the applicability of our results beyond the specific structural model we consider.

4.1 A Reduced-Form Model of Survival and Growth

Our structural model can be viewed as a special case of a less restrictive model of firm growth and survival in which (3) and (4) are replaced by a sequence of age-specific regression functions,

$$X_t = \beta_t (\bar{X}_{t-1}) + V_t \tag{10}$$

for $t = 1, \dots, T$. The expectational error V_t is normally distributed with mean zero and variance $\omega_t^2(\bar{X}_{t-1})$ conditional on \bar{X}_{t-1} . β_t and ω_t are arbitrary (measurable) functions of \bar{X}_{t-1} . The firm has no history at $t = 1$, so β_1 and ω_1 are simply constant parameters. The entrepreneur's exit rule also makes no reference to A_t or Bayesian learning. Instead, the entrepreneur closes the firm if and only if X_t falls below a history-dependent threshold, $\tau_t(\bar{X}_{t-1})$. Again, τ_t is an arbitrary (measurable) function of \bar{X}_{t-1} and τ_1 is simply a constant parameter. The recursive rule determining N_t in this case is

$$N_t = \begin{cases} 1 & \text{if } N_{t-1} = 1 \text{ and } X_{t-1} > \tau_{t-1}(\bar{X}_{t-2}) \\ 0 & \text{otherwise} \end{cases} . \quad (11)$$

The model given by (10), (11), and

$$S_t = N_t(X_t + W_t) \quad (12)$$

encompasses the structural model. For given values of the structural model's parameters, there are unique choices of β_t , ω_t^2 and τ_t that yield the same distribution of $(\bar{S}_T, \bar{N}_{T+1}, \bar{X}_T)$. In particular, equations (3) and (4) imply that $\beta_t(\bar{X}_{t-1}) = \mu_t + \rho_t \hat{A}_{t-1}$, so that we can build the regression functions β_t recursively using (5), (6), and (7). The regression error V_t is simply $X_t - \beta_t(\bar{X}_{t-1}) = \rho_t(A_{t-1} - \hat{A}_{t-1}) + Z_t + U_t$, so that $\omega_t^2 = \rho_t^2 \Sigma_{t-1} + \sigma_t^2 + \eta_t^2$. Finally, the exit threshold that is consistent with the structural model's exit rule is

$$\tau_t(\bar{X}_{t-1}) = \frac{\alpha_t - \beta_t(\bar{X}_{t-1})}{\lambda_t} + \beta_t(\bar{X}_{t-1}) . \quad (13)$$

Because its parameters display general history dependence, the reduced form model's identification demonstrates that structural models of firm growth and survival other than our own also have empirical content. In particular, the heteroskedasticity of the errors in the state-evolution equation (10) allows the volatility of firms' profits to depend on the resolution of a profit-relevant event, such as acquiring a stable clientele. Such modifications to our relatively simple structural model may be relevant for considering other aspects of the life cycle in other industries, and we defer their exploration to future research.

4.2 Identification of the Reduced Form Model

Identification of the reduced form model requires deconvolution of the observed sales process $\{S_t\}$ into the two components observed by the entrepreneur, $\{X_t\}$ and $\{W_t\}$. To achieve identification, we use sample selection to partially solve the deconvolution problem. Because

the entrepreneur's continuation decisions only depend on $\{X_t\}$, we can disentangle $\{X_t\}$ and $\{W_t\}$ by comparing the unconditional distribution of S_t with its distribution conditional on $N_{t+1} = 1$. Our solution is partial in the sense that we identify $\beta_t(\bar{X}_{t-1})$, $\omega_t^2(\bar{X}_{t-1})$, and $\tau_t(\bar{X}_{t-1})$ only for those values of \bar{X}_{t-1} that imply $N_t = 1$. The following proposition formally states this result.⁸

Proposition 1. (Identifiability of the reduced-form model) *The joint distribution of $(\bar{S}_T, \bar{N}_{T+1})$ uniquely determines the density f_{W_t} and the functions β_t , ω_t^2 , and τ_t for all \bar{X}_{t-1} such that $N_t = 1$ for $t = 1, \dots, T$.*

Proof. The proof proceeds recursively. It is helpful to define \vec{S}_t and \vec{N}_{t+1} to be the random vectors (S_t, \dots, S_T) and $(N_{t+1}, \dots, N_{T+1})$. 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 the econometrician knows this joint density for a particular age t . For $t = 1$, $f_{\vec{S}_1, \vec{N}_2, \bar{X}_0}(\cdot)$ is simply the joint density of the econometrician's data. We first show that β_t , ω_t^2 , τ_t and f_{W_t} can be recovered from $f_{\vec{S}_t, \vec{N}_{t+1}}(\cdot | \bar{X}_{t-1})$. 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 $\beta_t(\bar{X}_{t-1})$, $\omega_t(\bar{X}_{t-1})$ and $\tau_t(\bar{X}_{t-1})$ for some value of \bar{X}_{t-1} such that $N_t = 1$. This requires only the knowledge of the expected value of S_t conditional on this history, the conditional probability of survival, and the expected value of S_t given this history *and* survival to period $t + 1$. These are

$$\mathbb{E}[S_t | \bar{X}_{t-1}] = \beta_t(\bar{X}_{t-1}), \quad (14)$$

$$\mathbb{E}[N_{t+1} | \bar{X}_{t-1}] = 1 - \Phi\left(\frac{\tau_t(\bar{X}_{t-1}) - \beta_t(\bar{X}_{t-1})}{\omega_t(\bar{X}_{t-1})}\right), \quad (15)$$

$$\begin{aligned} \mathbb{E}[S_t | \bar{X}_{t-1}, N_{t+1} = 1] &= \beta_t(\bar{X}_{t-1}) + \omega_t(\bar{X}_{t-1}) \phi\left(\frac{\tau_t(\bar{X}_{t-1}) - \beta_t(\bar{X}_{t-1})}{\omega_t(\bar{X}_{t-1})}\right) \\ &/ \left[1 - \Phi\left(\frac{\tau_t(\bar{X}_{t-1}) - \beta_t(\bar{X}_{t-1})}{\omega_t(\bar{X}_{t-1})}\right)\right] \end{aligned} \quad (16)$$

where $\Phi(\cdot)$ and $\phi(\cdot)$ and the *c.d.f.* and *p.d.f.* of a standard normal random variable. Because the joint distribution of (S_t, N_{t+1}) conditional on the given history is assumed to be known, these three equations immediately yield $\beta_t(\bar{X}_{t-1})$, $\omega_t(\bar{X}_{t-1})$ and $\tau_t(\bar{X}_{t-1})$. To obtain f_{W_t} ,

⁸Because we have not even assumed continuity of f_{W_1}, \dots, f_{W_T} , β_2, \dots, β_T , $\omega_2, \dots, \omega_T$ and τ_2, \dots, τ_T , these functions can only possibly be identified up to almost-everywhere equivalence. We will not explicitly qualify our results this way.

note that the probability density of S_t conditional on the given history can be written as

$$f_{S_t}(s|\bar{X}_{t-1}) = \int_{-\infty}^{\infty} \frac{1}{\omega_t(\bar{X}_{t-1})} \phi\left(\frac{x - \beta_t(\bar{X}_{t-1})}{\omega_t(\bar{X}_{t-1})}\right) f_{W_t}(s - x) dx. \quad (17)$$

All components of (17) except f_{W_t} are known, so applying a standard deconvolution argument establishes that f_{W_t} is identified.⁹

We will now show that $f_{\bar{S}_{t+1}, \bar{N}_{t+2}, \bar{X}_t}(\cdot)$ 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 $(\bar{X}_{t-1}, \vec{S}_{t+1}, \vec{N}_{t+1})$ displays independence. Therefore, the probability density of S_t conditional on these variables is

$$f_{S_t}(s|\bar{X}_{t-1}, \vec{S}_{t+1}, \vec{N}_{t+1}) = \int_{-\infty}^{\infty} f_{X_t}(x|\bar{X}_{t-1}, \vec{S}_{t+1}, \vec{N}_{t+1}) f_{W_t}(s - x) dx. \quad (18)$$

As in (17), \bar{X}_{t-1} is assumed to imply that $N_t = 1$. The left-hand side of (18) and f_{W_t} are known, so deconvolution yields $f_{X_t}(\cdot|\bar{X}_{t-1}, \vec{S}_{t+1}, \vec{N}_{t+1})$. For any value of X_t such that $\bar{X}_t \equiv (X_t, \bar{X}_{t-1})$ implies that $N_{t+1} = 1$, we can immediately recover the joint distribution of interest by multiplying this conditional distribution by the known joint distribution of the conditioning variables. Thus, the recursion may continue. \square

Equation (16) directly parallels Heckman's (1979) decomposition of the regression function of a selected sample into the unconditional regression function and the expectation of the error conditional upon selection. As in his model, the choice probability in (15) immediately identifies the inverse of Mill's ratio in (16), but without imposing further parametric restrictions on $\beta_t(\bar{X}_{t-1})$ and $\omega_t(\bar{X}_{t-1})$, the resulting equation does not uniquely determine the model's parameters. However, because we observe the unselected sample, the identification of $\beta_t(\bar{X}_{t-1})$ is immediate using (14). With this additional information, the identification of $\omega_t(\bar{X}_{t-1})$ is straightforward.

Our analysis is also related to Blundell and Preston's (1998) identification of a consumer's permanent income process using the covariance of his current income (a noisy proxy) with his consumption (a forward-looking choice). To see this, rewrite equation (16) using (14) and (15) as

$$\mathbb{E}[S_t N_{t+1} | \bar{X}_{t-1}] - \mathbb{E}[S_t | \bar{X}_{t-1}] \mathbb{E}[N_t | \bar{X}_{t-1}] = \omega_t(\bar{X}_{t-1}) \phi(\Phi^{-1}(1 - \mathbb{E}[N_t | \bar{X}_{t-1}]))$$

This equation immediately yields $\omega(\bar{X}_t)$ from the observed covariance of S_t with N_t . As the covariance approaches zero for a given survival probability, so does the inferred value

⁹See Feller (1971) for an introduction to deconvolution.

of $\omega(\bar{X}_t)$. Thus, if current size is very helpful at forecasting survival in a linear probability model, we infer that much of the observed variance of S_t reflects persistent shocks. Otherwise, we infer that transitory shocks without dynamic implications for the entrepreneur dominate variation in S_t .¹⁰

4.3 Identification of the Structural Model

Proposition 1 shows that identification of \bar{X}_T 's distribution and the firm's exit policy relies on only on the assumption that $\{X_t\}$ follows some (possibly non-linear, higher-order, and heteroskedastic) auto-regression with normal disturbances and that the entrepreneur uses *some* history-dependent threshold-based exit rule. We now turn to the implications of Proposition 1 for our particular structural model. If we assume that entrepreneurial learning is absent from our framework, then the identification of the structural model's remaining parameters is immediate. The following proposition asserts that the addition of entrepreneurial learning imposes no fundamental obstacle to identification.

Proposition 2. *Given the functions β_t , ω_t , τ_t for $t = 1, \dots, T$, the following parameters of the structural model are uniquely determined: $\mu_1, \mu_3, \dots, \mu_T$; $\sigma_3^2, \dots, \sigma_T^2$; ρ_3, \dots, ρ_T ; $\eta_2^2, \dots, \eta_T^2$; and $\alpha_2, \dots, \alpha_T$. Furthermore, these functions uniquely determine the values of $\mu_2 + \rho_2\mu_1$, $\sigma_1^2 + \eta_1^2$, $\lambda_1\rho_2$, $\rho_2^2\Sigma_1 + \sigma_2^2$, and $(\alpha_1 - \mu_1)/\sigma_1^2$.*

Proposition 2 can be easily proven using the choices of β_t , ω_t , and τ_t that are consistent with the structural model. The proof is available from the authors upon request. The intuition behind this result is straightforward: if entrepreneurial learning is unimportant, then $\beta_t(\bar{X}_{t-1})$ will depend only on X_{t-1} and $\tau_t(\bar{X}_{t-1})$ will be a constant. With entrepreneurial learning, both of these functions depend in general on the entire history \bar{X}_{t-1} . Therefore, it is possible to use these functions' history dependence to infer the presence and magnitude of entrepreneurial learning. Because the lag structure plays a crucial role here, Proposition 2 does not deliver full identification of the structural model. In particular, there may be multiple ways of decomposing X_1 into the persistent component unobserved by the entrepreneur, A_1 , and the noise that masks it, U_1 . However, this lack of identification can be addressed with a single additional restriction. In all of our empirical specifications, we will impose the

¹⁰This argument depends critically on the assumption that V_t has a normal distribution. Because this assumption arises repeatedly in theoretical treatments of entrepreneurial learning, such as that of Jovanovic (1982), we view our structural model as a natural benchmark. The identification of a variant of the reduced form model without the assumption of normality is a subject of our current research.

assumption that $\rho_2 = \rho_3 = \dots = \rho_T$, which is more than sufficient to guarantee identification of all the structural model's parameters.¹¹

4.4 Discussion

Before proceeding to the estimation of our model, we wish to consider briefly our identification result from the perspective of the extensive literature on the identification of discrete choice models. A key aspect of our analysis is that entrepreneurs base their decisions on information that is unavailable to the econometrician. This is embodied in our model by the assumption that the econometrician only observes a noisy signal, S_t , of the relevant state variable, X_t . As such, our discrete-choice problem involves an error that is independent of the *latent* variable that determines the entrepreneur's choice. Therefore, we cannot directly apply non-parametric identification results for static binary-choice models involving an error that is independent of the observed covariates, such as some of Manski's (1988) results. The discrepancy between private and public information in our model furthermore implies a violation of Rust's (1987) conditional-independence assumption. Taber (2000) and Magnac and Thesmar (2002) have studied the identifiability of dynamic discrete-choice models without this assumption. Our identification result differs from theirs, because we observe no state variable.

5 Maximum Likelihood Estimation

The proof of Proposition 1 is constructive, so it immediately suggests a semi-parametric estimator for our model built on the empirical analogues of (14), (15), and (16). Forming these moment conditions requires repeated deconvolution applied to the joint density of $(\bar{S}_T, \bar{N}_{T+1})$. Rather than pursuing this semi-parametric approach, we consider instead parametric maximum likelihood estimation. To this end, we assume that W_t is normally distributed with mean zero and variance γ_t^2 . Maximum likelihood estimation 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.

¹¹The structural model also implies several cross-equation restrictions that allow one to test the hypotheses that (i) the persistence in X_t reflects the evolution of a single latent factor and (ii) entrepreneur's exit decisions are based on a rational assessment of this latent factor. We will explore these cross-equation restrictions in future work.

In the following, we suppose that the structural model's parameters are themselves known functions of a vector of "primitive" parameters, ϖ . Our method of calculating the likelihood function follows the non-Gaussian state-space model approach of Kitagawa (1987), which utilizes the likelihood's prediction-error decomposition

$$f_{\bar{Y}_{T+1}}(\bar{Y}_{T+1}; \varpi) = f_{Y_1}(Y_1; \varpi) \prod_{t=2}^{T+1} f_{Y_t}(Y_t | \bar{Y}_{t-1}; \varpi), \quad (19)$$

where $Y_t \equiv (S_t, N_t)$ unless $t = T + 1$, in which case it equals N_{T+1} . In (19), we have explicitly included the vector of primitive parameters to be estimated as an argument.

Consider the first term in (19), the density of Y_1 . Because $N_1 = 1$ always, this is simply

$$f_{Y_1}((S_1, 1); \varpi) = \frac{1}{\sqrt{\omega_1^2 + \gamma_1^2}} \phi\left(\frac{S_1 - \mu_1}{\sqrt{\omega_1^2 + \gamma_1^2}}\right), \quad (20)$$

where $\omega_t^2 = \rho_t^2 \Sigma_{t-1} + \sigma_t^2 + \eta_t^2$ is the variance defined in the analysis of the reduced form model.

Our recursive procedure for computing the likelihood's remaining terms is based on initializing and updating the density of \hat{A}_t given \bar{Y}_t . The distribution of \hat{A}_1 given S_1 is normal with mean $\mu_1 + \lambda_1 \omega_1^2 (S_1 - \mu_1) / (\omega_1^2 + \gamma_1^2)$ and variance $\lambda_1^2 \omega_1^2 \gamma_1^2 / (\omega_1^2 + \gamma_1^2)$. Therefore, $f_{\hat{A}_1}(\hat{a}_1 | Y_1; \varpi)$ is known.

To recursively calculate the likelihood function's remaining terms, suppose that for a given $t \geq 2$, the conditional density $f_{\hat{A}_{t-1}}(\hat{a}_{t-1} | \bar{Y}_{t-1}; \varpi)$ 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.

$$f_{Y_t}((0, 0) | \bar{Y}_{t-1}; \varpi) = \int_{-\infty}^{\alpha_{t-1}} f_{\hat{A}_{t-1}}(\hat{a}_{t-1} | \bar{Y}_{t-1}; \varpi) d\hat{a}_{t-1} \quad (21)$$

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

$$f_{Y_{T+1}}(1 | \bar{Y}_T; \varpi) = \int_{\alpha_T}^{\infty} f_{\hat{A}_T}(\hat{a}_T | \bar{Y}_T; \varpi) d\hat{a}_T \quad (22)$$

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

$$f_{Y_t}((S_t, 1) | \bar{Y}_{t-1}; \varpi) = \int_{\alpha_{t-1}}^{\infty} \frac{1}{\sqrt{\omega_t^2 + \gamma_t^2}} \phi\left(\frac{S_t - \mu_t - \rho_t \hat{a}_{t-1}}{\sqrt{\omega_t^2 + \gamma_t^2}}\right) f_{\hat{A}_{t-1}}(\hat{a}_{t-1} | \bar{Y}_{t-1}; \varpi) d\hat{a}_{t-1} \quad (23)$$

Equation (23) 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 $f_{\widehat{A}_t}(\widehat{a}_t|\bar{Y}_t; \varpi)$, which is

$$f_{\widehat{A}_t}(\widehat{a}_t|\bar{Y}_t; \varpi) = \frac{\int_{\alpha_{t-1}}^{\infty} f_{\widehat{A}_t, S_t}(\widehat{a}_t, S_t|\widehat{a}_{t-1}; \varpi) f_{\widehat{A}_{t-1}}(\widehat{a}_{t-1}|\bar{Y}_{t-1}; \varpi) d\widehat{a}_{t-1}}{f_{Y_t}(Y_t|\bar{Y}_{t-1}; \varpi)}. \quad (24)$$

The distribution of (\widehat{A}_t, S_t) conditional on \widehat{A}_{t-1} is bivariate normal, so both terms in the integrand of (24) are known. Therefore, the recursion can continue.

Iterating on (23) and (24) until the individual either exits or is right censored produces the likelihood function at any given choice of ϖ . In practice, evaluating the likelihood function requires approximating the integrals in (21), (22), (23), and (24). We do so using a Gaussian quadrature procedure.

6 Empirical Results

We now turn to maximum likelihood estimation of our model's parameters using the TABC panel. We begin with a relatively simple model based on Hopenhayn's (1992) industry dynamics model. In this model, the entrepreneur faces no signal extraction problem because $\eta_t^2 = 0$ and W_t has an age-invariant normal distribution with variance γ^2 . Furthermore, the evolution of A_t does not directly depend on the firm's age because μ_t , ρ_t , and σ_t^2 equal the constants μ , ρ , and σ^2 for $t \geq 2$. If the fixed production cost does not vary with age, then the optimal exit policy in such a stationary environment is itself stationary. We impose this by constraining α_t for all t to equal the constant α . We refer to this below as the model's stationary specification. After considering these estimates, we turn to a nonstationary specification in which μ_t and α_t are both allowed to be arbitrary functions of t . Finally, we further generalize the nonstationary specification by introducing entrepreneurial learning. To do so, we assume that η_t^2 equals a possibly nonzero constant, η^2 , for all t . This section then concludes with a brief summary of our results.

6.1 Stationary Specification

Table 4 reports the maximum likelihood estimates of our model's stationary specification using the sample of 305 new bars described in Section 2. Below each estimate is its asymptotic standard error, which we calculated using the outer product estimate of the information

matrix. First, consider the estimates of μ_1 , σ_1 , and γ , which jointly determine the distribution of firms' initial sizes. The estimate of μ_1 , 0.109, is somewhat higher than the unconditional mean of firms' initial size reported in Table 1, 0.06. However, the standard error attached to this estimate is relatively large, 0.055. The estimates of σ_1 and γ_1 , 0.766 and 0.294, together imply that the standard deviation of S_1 equals 0.82, which differs little from the estimated standard deviation of 0.87. At these parameter estimates, the transitory shock W_1 accounts for approximately 13% of the variance of S_1 . The remainder is attributable to the persistent shock X_1 . The estimate of α equals -1.23 , implying that the probability of a firm exiting after just one month of operation is 0.04. This is substantially higher than the exit rate we observe in the TABC sample.

Perhaps the most striking feature of Table 4 is the estimated value of ρ , which is exactly one. That is, the persistent component of profitability in the model follows a random walk. This is estimated with great precision — its standard error equals 0.006. With this estimate of ρ the intercept μ becomes the drift in the random walk of A_t . This estimate is negative, -0.017 and significantly different from zero. The innovation to this random walk has a standard deviation of σ , which we estimate to equal 0.187. Together, the estimated negative drift and substantial innovation variance imply that the expected growth rate of the persistent component's *level*, $\mathbb{E}[e^{\Delta A_t} - 1]$ is slightly above zero. Given the estimates of σ and γ , we can calculate the variance of the change in S_t , $\sigma^2 + 2\gamma^2$. The associated standard deviation equals 0.46. Variation in W_t accounts for just over 80% of this variance.

To help gauge how well the estimated model fits the data, Table 5 reports the population values of the summary statistics we considered above in Section 2.¹² As we noted above, the exit rate for new firms in the model is much higher than its empirical analogue. Similarly, the model's exit rates for firms nine to twelve months old are lower than those we estimate. Overall, the model with a stationary exit threshold fits the observed exit rates poorly. As in the data, selection drives growth in the average size of survivors, and it has little impact on the standard deviation of their sizes. As expected, selection increases the cohort's skewness as it ages. The cohort's kurtosis initially falls with selection and then gradually increases. Recall from Section 2 that these latter two moments are very imprecisely estimated, so the differences between the model's predictions of skewness and kurtosis and their sample counterparts does not indicate a poor fit between the model and the data.

¹²We calculated these summary statistics using 100000 simulated firm histories from the estimated model.

6.2 Nonstationary Specification

The discrepancy between the very low frequency of exit among brand-new entrants and the relatively high rate predicted by the estimated stationary specification suggests that the model's fit could be improved by allowing either the evolution of A_t or the exit policy to depend on t . The nonstationary specification allows for both of these sources of age dependence. Table 6 contains the maximum likelihood estimates of its parameters, and Table 7 presents the estimated specification's summary statistics. The estimates of those parameters that are common to the two specifications, μ_1 , σ_1 , σ , ρ , and γ , are very close to the analogous estimates in Table 4. In particular, the estimate of ρ , 0.992, is statistically indistinguishable from one. Next, consider the estimates of the age-specific intercepts, μ_2, \dots, μ_{12} . These are measured relatively imprecisely, but the tendency for them to be negative is clear: only three of the eleven estimates are positive. Overall, these estimates reveal no substantial tendency for a firm's growth to accelerate or decelerate as it ages.

The final parameters to report are those describing the entrepreneurs' exit policy, $\bar{\alpha}_{12}$. Unlike with the estimated intercepts, these are clearly not constant. The estimate of the initial exit threshold is -1.94 . This is estimated imprecisely, as the relatively little exit following the first month leads us to expect. Nevertheless, it is substantially below the estimated exit threshold from the stationary specification. The implied probability that a firm exits following its first month, 0.008, is much closer to that we observe. As the firm ages, the estimated exit thresholds tend to rise. Although this increase is not uniform over time, its effects are clear. The estimate of α_2 is -1.53 , while that for α_{12} is -0.97 . These are both relatively precisely estimated. The only remaining exit threshold that we estimate imprecisely is α_6 . This is unsurprising, given the fact that no firm in our sample exits immediately following its sixth month of operation. Overall, allowing α_t to vary with the firm's age significantly improves the model's fit with the data.¹³ A comparison of the model's summary statistics with those of the data reinforces this impression. As in the data, the exit rate from the model is initially low and then gradually rises through the cohort's first year. Similarly, the growth of surviving firms' average size matches that from the data reasonably well.

¹³To formally gauge this improvement, we have calculated a likelihood ratio statistic of the hypothesis that the values of μ_t are unconstrained and $\alpha_t = \alpha$ for all t relative to the alternative hypothesis embodied in the nonstationary specification. The likelihood ratio test statistic equals 236.09, which exceeds all conventional critical values for a χ^2 random variable with 11 degrees of freedom.

6.3 Maximum Likelihood Estimation of η^2

We now turn to 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 0.07, it plots the value of the likelihood function after enveloping out the nonstationary specification's other parameters. 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.

6.4 Summary of Results

Four results summarize our estimates. First, profit's persistent component follows a random walk. Thus, a version of Gibrat's law characterizes our young and small firms. Second, the variance of firms' growth rates is dominated by transitory shocks. In both specifications, these account for approximately 80% of the observed variance. Third entrepreneurs' exit thresholds increase as the firm ages. Because we observe no firms exiting after their thirteenth and fourteenth months, we believe that the annual timing of lease renewals plays a large role in the increasing exit thresholds we estimate. Finally, entrepreneurs' observe profit's persistent component without error. That is, entrepreneurial learning appears to be irrelevant for the firms in our sample.

7 Conclusion

Our model of firm growth and survival with private entrepreneurial information has empirical content. In particular, the extent of entrepreneurial learning is econometrically identified, even though entrepreneurs' observations are hidden from the econometrician. Our maximum likelihood estimates indicate that entrepreneurs' observe profit's persistent component without error, so the theoretical identification result has empirical relevance for our sample of new bars in Texas.

The model we develop in this paper is structural— its unknown parameters either are primitive to the entrepreneur's decision problem or represent the entrepreneur's solution to it. Because we estimate the entrepreneur's decision rule directly, our procedure avoids the repeated solution of the decision maker's dynamic programming problem inherent to Rust's

(1987) nested fixed point algorithm. An additional advantage of estimating entrepreneurs' exit policies directly is that the procedure thereby accounts for the entrepreneur's superior understanding of her profit maximization problem. The nonstationarity of the estimated exit thresholds illustrates this. Time-varying fixed costs associated with annual lease renewal can manifest themselves in the estimated exit thresholds even if the econometrician fails to appreciate their importance before estimation.

One obvious omission from this paper's model is the influence of observable covariates on either the stochastic process for firm size or the entrepreneur's exit thresholds. Incorporating observable covariates is a straightforward extension of our analysis, but it may be of considerable relevance for the model's further application. A structural model of the growth and survival of new establishments owned by incumbent firms should incorporate observable characteristics of the parent firms as covariates. Examples of such characteristics in the TABC data set are firm size, geographic location, and use of a franchised identity. The effects of these variables on establishment growth and survival are of independent relevance for theories of contracts and firm formation. We plan to pursue these applications in future research.

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.F. Bewley, editor, *Advances in Econometrics. Fifth World Congress*, volume I, Cambridge University Press, Cambridge.
- Blundell, R. and I. Preston (1998), “Consumption inequality and income uncertainty”, *Quarterly Journal of Economics*, 113, 603–640.
- Campbell, J.R. (2002), “A robust prediction of monopolistic competition”, Mimeo, Federal Reserve Bank of Chicago and NBER.
- Campbell, J.R. and H.A. Hopenhayn (2002), “Market size matters”, Working Paper 9113, National Bureau of Economic Research.
- Campbell, J.R. and B. Lapham (2001), “Real exchange rate fluctuations and the dynamics of retail trade industries on the U.S.-Canada border”, Working Paper 8558, National Bureau of Economic Research.
- 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 (1989), “Plant turnover and gross employment flows in the U.S. manufacturing sector”, *Journal of Labor Economics*, 7, 48–71.
- Ericson, R. and A. Pakes (1995), “Markov-perfect industry dynamics: A framework for empirical work”, *Review of Economic Studies*, 62, 53–82.
- Feller, W. (1971), *An Introduction to Probability Theory and Its Applications*, volume II, Wiley, New York, 2nd edition.
- Fishman, A. and R. Rob (2000), “Consumer inertia, firm growth and industry dynamics”, Technical report, University of Haifa and University of Pennsylvania.
- 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.
- 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.
- 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.
- Magnac, T. and D. Thesmar (2002), “Identifying dynamic discrete choice processes”, *Econometrica*, 70, 801–816.
- Manski, C.F. (1988), “Identification of binary response models”, *Journal of the American Statistical Association*, 83, 729–738.
- 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.
- 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.
- Rust, J. (1987), “Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher”, *Econometrica*, 55, 999–1033.
- Salop, S.C. (1979), “Monopolistic competition with outside goods”, *Bell Journal of Economics*, 10, 141–156.
- Taber, C. (2000), “Semiparametric identification and heterogeneity in discrete choice dynamic programming models”, *Journal of Econometrics*, 96, 201–229.
- 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	Mean	Standard Deviation	Skewness	Kurtosis	Exit Rate ⁽ⁱⁱⁱ⁾
1	305	0.06	0.87	-0.003	3.53	0.003
2	304	0.08	0.89	-0.066	3.39	0.016
3	299	0.10	0.86	0.085	3.17	0.007
4	297	0.11	0.85	0.114	3.24	0.020
5	291	0.12	0.85	0.071	3.08	0.017
6	286	0.11	0.87	0.034	3.28	0.000
7	286	0.11	0.89	0.058	3.33	0.028
8	278	0.17	0.84	0.199	3.34	0.004
9	277	0.17	0.86	0.101	3.49	0.032
10	268	0.20	0.86	0.218	3.35	0.019
11	263	0.20	0.89	-0.006	3.23	0.042
12	252	0.23	0.89	0.037	3.23	0.032
13	244	0.27	0.87	0.052	3.30	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) 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: Summary Statistics from the First 13 Months in the Lives of Surviving Bars⁽ⁱ⁾

Age ⁽ⁱⁱ⁾	Mean	Standard Deviation	Skewness	Kurtosis
1	0.17	0.85	0.089	3.38
2	0.21	0.86	-0.033	3.53
3	0.20	0.86	0.083	3.13
4	0.23	0.85	0.079	3.25
5	0.24	0.83	0.006	3.16
6	0.21	0.87	-0.065	3.51
7	0.24	0.84	0.166	3.35
8	0.24	0.83	0.204	3.40
9	0.25	0.84	0.205	3.41
10	0.26	0.85	0.194	3.11
11	0.24	0.88	0.074	3.24
12	0.25	0.89	0.015	3.30
13	0.27	0.87	0.052	3.30

Notes: (i) The reported statistics are for the sample of firms that operated for at least 13 months. See the text for further details regarding its construction. (ii) Age is measured in months and equals one for a firm filing its first tax return.

Table 3: Regression Estimates of Sales on Previous and First Months' Sales⁽ⁱ⁾

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

Parameter ⁽ⁱⁱ⁾	μ_1	σ_1	ρ	μ	σ	α	γ
Estimate	0.109	0.766	1.000	-0.017	0.187	-1.230	0.294
Standard Error ⁽ⁱⁱⁱ⁾	(0.055)	(0.037)	(0.006)	(0.005)	(0.004)	(0.007)	(0.002)

Notes: (i) This specification assumes that entrepreneurs face no signal extraction problem ($\eta^2 = 0$), that the parameters determining the evolution of X_t and W_t after the firm's birth do not depend on the firm's age, and that the entrepreneur's exit threshold also does not vary with age. (ii) The parameters μ and σ refer to μ_t and σ_t for $t \geq 2$. All other parameters without a time subscript are assumed to be invariant to the firm's age. (iii) 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 5: Summary Statistics from the Stationary Specification⁽ⁱ⁾

Age ⁽ⁱⁱ⁾	Firms ⁽ⁱⁱⁱ⁾	Mean	Standard Deviation	Skewness	Kurtosis	Exit Rate ^(iv)
1	100.0	0.11	0.82	0.00	3.00	0.041
2	95.9	0.16	0.78	0.22	2.84	0.014
3	94.6	0.17	0.79	0.25	2.87	0.012
4	93.4	0.17	0.80	0.28	2.88	0.012
5	92.3	0.18	0.81	0.30	2.89	0.013
6	91.1	0.18	0.82	0.33	2.91	0.013
7	89.9	0.18	0.83	0.34	2.91	0.013
8	88.8	0.19	0.84	0.36	2.93	0.013
9	87.6	0.19	0.85	0.37	2.93	0.013
10	86.4	0.20	0.86	0.39	2.94	0.013
11	85.3	0.20	0.87	0.40	2.95	0.014
12	84.1	0.20	0.88	0.42	3.00	0.014

Notes: (i) The reported statistics are calculated from a synthetic sample of 100,000 firm histories drawn from the estimated stationary specification. (ii) Age is measured in months and equals one for a new 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 6: Maximum Likelihood Estimates: Nonstationary Specification⁽ⁱ⁾

Parameter ⁽ⁱⁱ⁾	σ_1	σ	ρ	γ
Estimate	0.832	0.190	0.992	0.278
Standard Error ⁽ⁱⁱⁱ⁾	(0.040)	(0.004)	(0.005)	(0.003)

Parameter	μ_1	μ_2	μ_3	μ_4	μ_5	μ_6
Estimate	0.063	0.026	-0.014	-0.001	-0.024	-0.030
Standard Error ⁽ⁱⁱⁱ⁾	(0.067)	(0.048)	(0.045)	(0.047)	(0.048)	(0.041)

Parameter	μ_7	μ_8	μ_9	μ_{10}	μ_{11}	μ_{12}
Estimate	0.004	0.017	-0.008	-0.013	-0.029	-0.017
Standard Error ⁽ⁱⁱⁱ⁾	(0.041)	(0.043)	(0.051)	(0.054)	(0.046)	(0.035)

Parameter	α_1	α_2	α_3	α_4	α_5	α_6
Estimate	-1.94	-1.53	-1.44	-1.19	-1.27	-1.55
Standard Error ⁽ⁱⁱⁱ⁾	(0.17)	(0.03)	(0.04)	(0.03)	(0.03)	(0.31)

Parameter	α_7	α_8	α_9	α_{10}	α_{11}	α_{12}
Estimate	-1.32	-1.45	-1.16	-1.10	-0.92	-0.97
Standard Error ⁽ⁱⁱⁱ⁾	(0.03)	(0.04)	(0.03)	(0.04)	(0.02)	(0.02)

Notes: (i) This specification assumes that entrepreneurs face no signal extraction problem ($\eta^2 = 0$) and that the evolution of X_t and W_t depends directly on age only because μ_t may be a nontrivial function of t . (ii) The parameter σ refers to σ_t for $t \geq 2$. The other parameters without a time subscript are assumed to be invariant to the firm's age. (iii) 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 7: Summary Statistics from the Nonstationary Specification⁽ⁱ⁾

Age ⁽ⁱⁱ⁾	Firms ⁽ⁱⁱⁱ⁾	Mean	Standard Deviation	Skewness	Kurtosis	Exit Rate ^(iv)
1	100.0	0.07	0.88	0.00	3.00	0.007
2	99.2	0.11	0.87	0.09	2.87	0.021
3	97.1	0.13	0.85	0.19	2.85	0.014
4	95.7	0.16	0.85	0.24	2.85	0.035
5	92.4	0.19	0.83	0.32	2.87	0.010
6	91.5	0.17	0.83	0.34	2.90	0.001
7	91.4	0.17	0.85	0.31	2.89	0.013
8	90.3	0.21	0.85	0.34	2.91	0.003
9	90.0	0.20	0.86	0.32	2.90	0.027
10	87.6	0.23	0.85	0.38	2.93	0.022
11	85.7	0.24	0.84	0.41	2.95	0.051
12	81.3	0.29	0.82	0.47	3.05	0.017

Notes: (i) The reported statistics are calculated from a synthetic sample of 100,000 firm histories drawn from the estimated nonstationary specification. (ii) Age is measured in months and equals one for a new 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 .

Figure 1: Maximum Likelihood Estimation of η^2

