# Very Simple Markov-Perfect Industry Dynamics*

Jaap H. Abbring[†] Jeffrey R. Campbell[‡] Jan Tilly[§]
Nan Yang[¶]

May 28, 2013

**Abstract**

This paper develops an econometric model of industry dynamics for concentrated markets that can be estimated very quickly from market-level panel data on the number of producers and consumers using a nested fixed-point algorithm. Its estimation enables the measurement of economic barriers to entry and the toughness of price competition. The model's econometric error comes from a shock to both potential entrants' sunk costs of entry and incumbents' fixed costs of continuation. We show that the model has an essentially unique symmetric Markov-perfect equilibrium that can be calculated from the fixed points of a finite sequence of low-dimensional contraction mappings. Our nested fixed point procedure extends Rust's (1987) to account for the observable implications of mixed strategies on survival. We illustrate the model's empirical application with ten years of County Business Patterns data from the Motion Picture Theaters in 573 Micropolitan Statistical Areas. Adding four firms to a monopoly market reduces profits per customer by 23 percent.

---

*A replication file will be available at http://research.abbring.org.

[†]CentER, Dept. of Econometrics & OR, Tilburg University. E-mail: jaap@abbring.org.

[‡]Economic Research, Federal Reserve Bank of Chicago, and CentER, Tilburg University. E-mail: jcampbell@frbchi.org.

[§]Department of Economics, University of Pennsylvania. E-mail: jtilly@sas.upenn.edu.

[¶]Business School, National University of Singapore. E-mail: yangnan@nus.edu.sg.

JEL Codes: L13, C23, C25, C73.

Keywords: barriers to entry, demand uncertainty, dynamic oligopoly, firm entry and exit, Markov-perfect equilibrium, nested fixed point estimator, sunk costs, toughness of competition.

The research of Jaap Abbring is financially supported by the Netherlands Organisation for Scientific Research (NWO) through Vici grant 016.125.624. The views expressed herein are those of the authors, and they do not necessarily represent those of the Federal Reserve Bank of Chicago, the Federal Reserve System, or its Board of Governors.

# 1 Introduction

This paper introduces an econometric model of firm entry, competition, and exit in dynamic oligopolistic markets. The model includes market-level demand and cost shocks and sunk entry costs. Because all shocks with dynamic consequences occur at the market level, the model's theoretical analysis and equilibrium computation are straightforward. In particular, we prove that there exists an essentially unique symmetric Markov-perfect equilibrium that can be computed from the fixed points of a finite sequence of low-dimensional contraction mappings. We use these results to develop a nested fixed point (NFXP) algorithm for the model's maximum likelihood estimation. This extends Rust's (1987) algorithm to account for mixed equilibrium survival strategies.

We begin with Abbring, Campbell, and Yang's (2010) model of Markov-perfect duopoly dynamics. They describe it as "simple" because adding a second firm to a market always lowers the equilibrium payoff of a monopolist incumbent. This result allows them to prove that there is an essentially unique "natural" Markov-perfect equilibrium and to develop an algorithm for its fast calculation. At the cost of removing firm-specific shocks to profitability, we extend their equilibrium uniqueness and calculation results to an oligopoly setting. We also add a market-level shock to both potential entrants' sunk costs of entry and incumbents' fixed costs of continuation. This is observed by market participants but not by the econometrician and so serves as the econometric error. Removing firm-specific shocks makes our framework inappropriate for applications that focus on persistent firm heterogeneity, as in Hopenhayn (1992) and Melitz (2003). However, it is well suited for extending Bresnahan and Reiss's (1990; 1991) measurements of the effect of competition on profitability to a dynamic setting.

Monte Carlo results indicate that the model can accurately estimate sunk costs and profits per customer (normalized to the per-period fixed cost of production) using observations on the number of producers and consumers from as few as 25 markets over ten years. We further illustrate the model's application by estimating its primitives for one concentrated industry, Motion Picture Theaters (NAICS 512131). Our data include observations on the number of theaters from 2000 to 2009 serving 573 Micropolitan Statistical Areas ($\mu$SAs). We find that adding a single firm to a monopoly market lowers profits per customer by 10 percent. Adding a third or fourth firm lowers profits only little, while the fifth firm reduces profits by 13 percent more. The maximum likelihood estimation takes only *seconds* on an ordinary personal computer. In this sense, we make Abbring, Campbell, and Yang's (2010) analysis *very* simple.

Both our model and that of Abbring, Campbell, and Yang (2010) can be viewed as special cases of the Ericson and Pakes (1995) Markov-perfect industry dynamics framework in which firms cannot vary their investments in productivity improvements. Since computing that model's equilibria has proven to be computationally challenging (see Doraszelski and Pakes (2007) for examples), that framework's estimation has focused on statistically inefficient methods that avoid equilibrium calculation altogether. For example, Bajari, Benkard, and Levin (2007) apply the Hotz and Miller (1993) inversion to estimate the structural parameters governing a single agent's dynamic optimization problem after conditioning on the observed distribution of all other agents' choices. Our nested fixed-point algorithm computes maximum likelihood estimates which are, of course, statistically efficient. Even with parameter estimates in hand, the lack of a fast algorithm for equilibrium computation makes counterfactual analysis of the Ericson and Pakes (1995) model difficult. Weintraub, Benkard, and Van Roy (2008) make equilibrium computation more tractable by assuming that firms ignore current information about competitors' states. Instead, they make dynamic decisions based solely on their own state and knowledge of the long-run average industry state. Such *oblivious equilibria* approximate Markov-perfect equilibria when the number of competitors is large, yet they are much easier to compute and thus can serve as a starting point for empirical analysis, as in Xu (2008). Our analysis can serve as a similar starting point for the analysis of markets with only few firms.

The remainder of the paper proceeds as follows. The next section presents the model's primitives, and Section 3 discusses equilibrium existence, uniqueness, and computation. Section 4 develops the model's empirical implementation. It subsequently discusses sampling, likelihood construction, identification, and maximum likelihood estimation using the NFXP procedure. Section 5 demonstrates the good computational performance of the NFXP procedure and explores the estimator's finite sample behavior using Monte Carlo experiments. It also briefly discusses the relative performance of Su and Judd's (2012) mathematical programming with equilibrium constraints (MPEC) implementation of the estimator. Section 6 illustrates the model's application with an empirical analysis of entry, competition, and exit in the Motion Picture Theaters (NAICS 512131) industry. Section 7 concludes. The Appendix provides technical details.

## 2 The Model

Consider a market in discrete time indexed by $t \in \mathbb{N}$. In period $t$, firms that have entered in the past serve the market. Each firm has a name $f \in \mathcal{F} \equiv \mathbb{N} \times \mathbb{N}$. The firm's name gives the precise node of the game tree in which the firm has its single opportunity to enter the market. Aside from the timing of their entry opportunities, the firms are identical.

Figure 1 details the actions taken by firms in period $t$ and their consequences for the game's state at the start of period $t+1$. We call this the game's *recursive extensive form.* For expositional purposes, we divide each period into two subperiods, the entry and survival subgames. Play in period $t$ begins on the left with the entry subgame. The period begins with values of $N_t$ and $C_{t-1}$ which were either inherited from period $t-1$ (if $t > 1$) or set by nature at the beginning of play (if $t = 1$). Nature draws a new demand state $C_t$ from the conditional distribution $G_C(\cdot \mid C_{t-1})$ and a real-valued cost state $W_t$ from the marginal distribution $G_W(\cdot)$. We use $\mathcal{C}$ to denote the support of $C_t$, and $W_t$'s support is the real line. All incumbent firms observe $(C_t, W_t)$, and each earns a surplus $\pi(N_t, C_t)$ from serving the market. We assume that

- $\exists \check{\pi} < \infty$ such that $\forall n \in \mathbb{N}$ and $\forall c \in \mathcal{C}$, $\pi(n,c) < \check{\pi}$;
- $\exists \check{n} \in \mathbb{N}$ such that $\forall n > \check{n}$ and $\forall c \in \mathcal{C}$, $\pi(n,c) = 0$; and
- $\forall n \in \mathbb{N}$ and $\forall c \in \mathcal{C}$, $\pi(n,c) \geq \pi(n+1,c)$.

The first assumption is technical and allows us to restrict equilibrium values to the space of bounded functions. We will use the second assumption to bound the number of firms that will participate in the market simultaneously. It is not restrictive in empirical applications to oligopolistic markets. The third assumption requires the addition of a competitor to reduce weakly per-period surplus for all incumbents. Sutton (1991) labelled the rate at which additional competitors lower post-entry surplus the *toughness of competition.*

The period $t$ entry cohort consists of firms with names in $t \times \mathbb{N}$. After incumbents receive their payoffs, these firms make their entry decisions sequentially in the order of their names' second components. We denote firm $f$'s entry decision with $a_E^f \in \{0,1\}$. If $j$ firms with lower names entered the market in this period, then firm $f$ incurs the sunk cost $\varphi(N_t + j)\exp(W_t)$ if it enters the market ($a_E^f = 1$). We assume that

- $\varphi(n+1) \geq \varphi(n) \geq 0$ for all $n \in \mathbb{N}$, so entrants might pay a higher entry cost if the market has more firms. McAfee, Mialon, and Williams (2004) refer to this as an *economic barrier to entry*. If firm $f$ chooses to not enter the market ($a_E^f = 0$), then it earns a payoff of zero and never has an opportunity to enter again. This refusal to enter also ends the entry subgame, so firms remaining in this period's entry cohort that have not yet had an opportunity to enter *never* get to do so. Since the next firm in line faces exactly the same choice as did the firm that refused to enter, this convenient assumption is without substantial loss of generality. Since every period has at least one firm refusing an available entry opportunity, the model is one of free entry.

The total number of firms in the market after the entry stage equals $N_{E,t}$, which sums the incumbents with the actual entrants ($J_t$ in Figure 1). Denote their names with $f_1, \ldots, f_{N_{E,t}}$. In the survival subgame, these firms simultaneously choose probabilities of remaining active, $a_S^{f_1}, \ldots, a_S^{f_{N_{E,t}}} \in [0,1]$ Nature subsequently draws the firms' survival outcomes independently across firms from the chosen Bernoulli distributions. Firms that survive pay a fixed cost $\kappa \exp(W_t)$, with $\kappa > 0$. Firms that exit earn 0 and never again participate in the market. The $N_{t+1}$ surviving firms continue in the next period, $t+1$.

Firms make entry and exit decisions that maximize their flow of payoffs discounted with a factor $0 \leq \rho < 1$, given the entry and exit strategies used by their competitors.

Before continuing to the model's analysis, we review its key assumptions from the perspective of its econometric implementation using data on a panel of markets. In Section 4, we will assume that, for each market, the data contain information on $N_t$, $C_t$, and possibly some time-invariant market characteristics $X$ that shift the market's primitives. The market-level cost shocks $W_t$ are not observed by the econometrician and serve as the model's structural econometric errors. Because they are observed by all firms and affect their payoffs from entry and survival, they ensure that the relation between the market structure $N_t$ and the observed demand state $C_t$ is statistically nondegenerate.

The assumptions on $\{C_t, W_t\}$ make it a first-order Markov chain satisfying Rust's (1987) conditional independence assumption.[1] This ensures that the markets' observed (by the econometrician) initial conditions $(N_1, C_0)$ cannot be

[1] Rust (1987) defines "conditional independence" for a *controlled* Markov process, but his definition immediately specializes to our case of an externally specified process $\{C_t, W_t\}$ if we take the control to be trivial. In terms of our model, Rust's conditional independence assumption more generally allows $W_t$ and $C_t$ to depend on $(C_{t-1}, W_{t-1})$ through a conditional distribution $G_W(\cdot \mid C_t)$. Our analysis easily extends to this case.

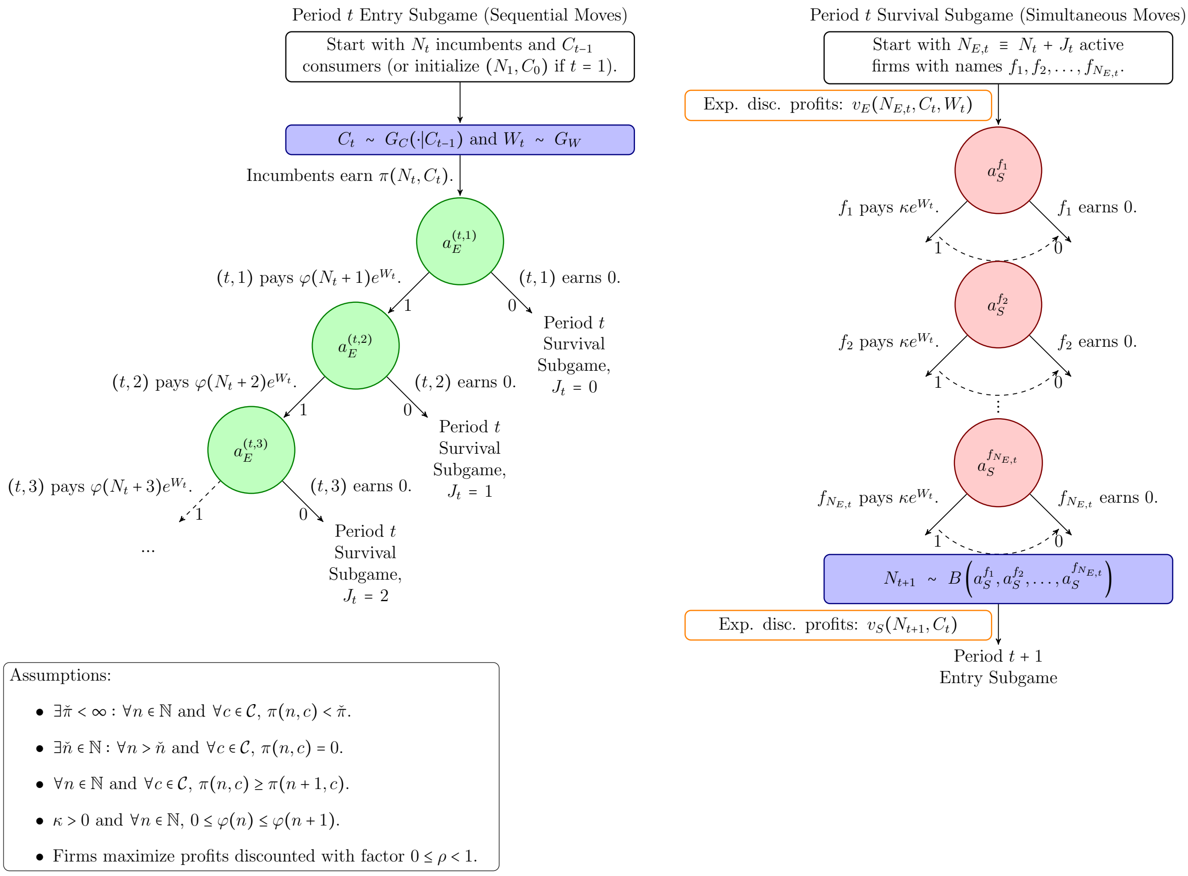


**Figure 1:** The Model's Recursive Extensive Form

informative of the unobserved cost shocks $\{W_t\}$.

# 3 Equilibrium

We assume that firms play a symmetric Markov-perfect equilibrium (Maskin and Tirole, 1988), a subgame-perfect equilibrium in which all firms use the same Markov strategy.

## 3.1 Markov Strategies

A Markov strategy is a strategy that maps *payoff relevant states* into actions. When a potential entrant $(t, j)$ makes its entry decision in period $t$, the payoff-relevant states are the number of firms in the market including all current period's entrants up to $(t, j)$, $M_t^j \equiv N_t + j$, the current state of demand $C_t$, and the cost shock $W_t$. We collect the entrant's payoff relevant state variables into the tuple $(M_t^j, C_t, W_t)$, which takes values in $\mathcal{H} \equiv \mathbb{N} \times \mathcal{C} \times \mathbb{R}$. Similarly, we collect the payoff relevant state variables of a firm $f$ contemplating survival in period $t$ in the $\mathcal{H}$-valued tuple $(N_{E,t}, C_t, W_t)$. Since survival decisions are made simultaneously, this state is the same for all active firms. A Markov strategy is a pair of functions $a_E : \mathcal{H} \to \{0,1\}$ and $a_S : \mathcal{H} \to [0,1]$. Since time itself is not payoff relevant, we drop the subscript $t$ from the payoff relevant states and denote the next period's value of $N$ with $N'$, etcetera.

## 3.2 Symmetric Markov-Perfect Equilibrium

In a symmetric Markov-perfect equilibrium, firms' expected discounted profits (values) at each node of the game are a function of that node's payoff-relevant state variables. Two value functions are particularly useful for the model's equilibrium analysis: the *post-entry* value function, $v_E$, and the *post-survival* value function, $v_S$. The post-entry value $v_E(N_E, C, W)$ equals the expected discounted profits of a firm facing $C$ consumers in a market with $N_E$ firms just after all entry decisions are made. The post-survival value $v_S(N', C)$ equals the expected discounted profits from being active in a market with $N'$ firms just after the survival outcomes are realized. Figure 1 shows the points in the survival subgame where these value functions apply. The post-survival function does not depend on $W$ because that cost shock has no forecasting value and is not directly payoff relevant after survival decisions are made. Since the payoff from leaving the market equals zero, the post-entry value function $v_E$ satisfies

$$v_E(n_E, c, w) = a_S(n_E, c, w)\left(-\kappa \exp(w) + \mathbb{E}_{a_S}\left[v_S(N', c)\middle| N_E = n_E, C = c, W = w\right]\right). \quad (1)$$

Here and throughout, random variables are denoted by capital letters and realizations of random variables by small letters. The expectation $\mathbb{E}_{a_S}$ over $N'$ takes survival of the firm of interest as given. We make its dependence on the equilibrium survival rule explicit with the subscript $a_S$. The post-survival value function $v_S$ satisfies

$$v_S(n', c) = \rho \mathbb{E}_{a_E}\left[\pi(n', C') + v_E(N'_E, C', W')\middle| N' = n', C = c\right]. \quad (2)$$

A strategy $(a_E, a_S)$ forms a symmetric Markov-perfect equilibrium with payoffs $(v_E, v_S)$ if and only if no firm can gain from one-shot deviations (see Sections 4.2 and 13.2 in Fudenberg and Tirole, 1991). Thus, given the pair of payoff functions $(v_E, v_S)$ as obtained in equations (1) and (2), it must hold that

$$a_E(m,c,w) \in \arg\max_{a\in\{0,1\}} a\Big(-\varphi(m)\exp(w) + \mathbb{E}_{a_E}\left[v_E(N_E,c,w)|M=m, C=c, W=w\right]\Big), \quad (3)$$

$$a_S(n_E,c,w) \in \arg\max_{a\in[0,1]} a\Big(-\kappa\exp(w) + \mathbb{E}_{a_S}\left[v_S(N',c)|N_E=n_E, C=c\right]\Big). \quad (4)$$

Before proceeding to the equilibrium analysis, we wish to note and dispense with an uninteresting source of equilibrium multiplicity. If a potential entrant is indifferent between its two choices, we can construct one equilibrium from another by varying only that choice. Similarly, an incumbent monopolist can be indifferent between continuation and exit, and we can construct one equilibrium from another by changing that choice alone. Moreover, the assumption that flow profits *weakly* decline with $n$ leaves open the theoretical (but not so practically relevant) possibility that $n$ oligopolists are each indifferent across *all* possible survival outcomes. To avoid these uninteresting caveats to our results, we focus on equilibria that *default to inactivity.* In such an equilibrium, a potential entrant that is indifferent between entering or not stays out,

$$\mathbb{E}_{a_E}\left[v_E(N_E,c,w)|M=m, C=c, W=w\right] = \varphi(m)\exp(w) \Rightarrow a_E(m,c,w) = 0,$$

and an active firm that is indifferent between all possible outcomes of the survival stage exits,

$$v_S(n,c) = \cdots = v_S(1,c) = \kappa\exp(w) \Rightarrow a_S(n,c,w) = 0.$$

The restriction to equilibria that default to inactivity does *not* restrict the game's strategy space. Hereafter, we require the strategy underlying a "symmetric Markov-perfect equilibrium" to default to inactivity.

## 3.3 Existence, Uniqueness, and Computation

This section presents our analysis of the existence, uniqueness, and computation of the symmetric Markov-perfect equilibrium. The paper's Appendix contains its claims' proofs. We start by noting that the assumption that per-period surplus equals zero if $n > \check{n}$ bounds the long-run number of firms in equilibrium.

**Lemma 1 (Bounded number of firms)** *In a symmetric Markov-perfect equilibrium,* $\forall c \in \mathcal{C}$ *and* $\forall w \in \mathbb{R}$, $a_E(n,c,w) = 0$ *and* $a_S(n,c,w) < 1$ *for all* $n > \check{n}$.

Intuitively, the post-survival payoff to one of more than $\check{n}$ firms must be negative because the flow payoff can become positive only when some other firm leaves. Since *all* firms must earn zero expected payoffs if the common survival strategy gives a positive probability to exit, any positive expected profits earned after other firms' departures are balanced by the payment of the (stochastic) fixed cost $\kappa \exp(W)$ when more than $\check{n}$ firms continue. Thus survival with $\check{n}$ or more rivals incurs a cost — the current value of $\kappa \exp(W)$ — with no benefit. Consequently, no firm would pay a positive sunk cost to enter the market ($a_E(n, c, w) = 0$) and all incumbent firms choosing sure continuation is inconsistent with individual payoff maximization ($a_S(n, c, w) < 1$).

In equilibrium, the market can only have more than $\check{n}$ active firms if $N_1 > \check{n}$. Because these firms exit with positive probability until there are $\check{n}$ or fewer of them, $N_t$ must eventually enter $\{0, 1, \ldots, \check{n}\}$ permanently. Consequently, the equilibrium analysis hereafter focuses on the restrictions of $a_E$, $v_E$, and $a_S$ to $\{1, 2, \ldots, \check{n}\} \times \mathcal{C} \times \mathbb{R} \subset \mathcal{H}$ and of $v_S$ to $\{1, 2, \ldots, \check{n}\} \times \mathcal{C}$. With an equilibrium strategy over this restricted state space in hand, it is straightforward to extend it to the full state space.

The next step in the equilibrium analysis extends the assumption that flow payoffs decrease with the number of competitors to the post entry and survival value functions.

**Lemma 2 (Monotone equilibrium payoffs)** *In a symmetric Markov-perfect equilibrium,* $\forall c \in \mathcal{C}$ *and* $\forall w \in \mathbb{R}$, $v_E(n, c, w)$ *and* $v_S(n, c)$ *weakly decrease with* $n$.

The monotonicity assumption on $\pi$ rules out exogenously specified complementarities between firms in the market. Lemma 2 says that endogenous complementarity also does not arise in equilibrium. Although this is intuitive, it is not a trivial result. Indeed, Abbring, Campbell, and Yang (2010) give a counterexample to the analogous proposition in a model with heterogeneous productivity types. In it, two high-productivity firms mutually benefit each other by jointly deterring the entry of two low-quality potential rivals. That counterexample shows that the lack of post-entry heterogeneity is crucial for obtaining Lemma 2.[2]

Consider a one-shot simultaneous-moves survival game played by $n_E$ active firms. In it, each of the $n'$ survivors earns $-\kappa \exp(w) + v_S(n', c)$, with $v_S$ referring

[2] Abbring, Campbell, and Yang (2010) show that a version of Lemma 2 also holds in a model with heterogeneous productivity if $\check{n} = 2$, so proving payoff monotonicity does not (strictly speaking) require post-entry homogeneity.

to the post-survival value in a symmetric Markov-perfect equilibrium of our model, and those that choose exit earn zero. The Nash equilibria of this game are intimately connected to the Markov-perfect equilibria of our model. In particular, (3) and (4) imply that a strategy that forms a symmetric Nash equilibrium of the one-shot game equals the survival rule $a_S$, evaluated at the state $(n_E, c, w)$, in a symmetric Markov-perfect equilibrium of our model, and vice versa.

Now note that Lemma 2 guarantees that the one-shot game has a *unique* symmetric Nash equilibrium (that defaults to inactivity). If $v_S(1,c) < \kappa \exp(w)$, then Lemma 2 guarantees that $v_S(n,c) < \kappa \exp(w)$ for all $n > 1$. Therefore, exiting for sure (setting $a_s(n,c,w) = 0$) is a dominant strategy. On the other hand, if $v_S(n_E,c) > \kappa \exp(w)$, Lemma 2 requires $v_S(n',c) > \kappa \exp(w)$ for $n' = 1, \ldots, n_E - 1$. Again, we have a dominant strategy, $a_S(n,c,w) = 1$. Finally, if neither of these conditions holds, then no pure-strategy symmetric equilibrium exists, but there is an equilibrium in a mixed strategy. In this equilibrium, $a_S(n,c,w)$ has to make firms indifferent between continuation and exit:

$$\sum_{n'=1}^{n} \binom{n-1}{n'-1} a_S(n,c,w)^{n'-1} \left[1 - a_S(n,c,w)\right]^{n-n'} \left[-\kappa \exp(w) + v_S(n',c)\right] = 0. \tag{5}$$

The left hand side of (5) is the expected payoff of a firm that continues and that assumes that all its $n-1$ rivals continue with probability $a_S(n,c,w)$. Lemma 2 guarantees that this payoff is decreasing in $a_S(n,c,w)$, so that there is only one mixed strategy equilibrium. For future reference, we state this result as

**Corollary 1** *Let $v_S$ be the post-survival value function associated with a symmetric Markov-perfect equilibrium. Consider the one-shot survival game in which $n_E$ firms simultaneously choose between survival and exit (as in the survival subgame of Figure 1), each of the $n'$ survivors earns $-\kappa \exp(w) + v_S(c, n')$, and each exiting firm earns zero. This game has a unique symmetric Nash equilibrium, possibly in mixed strategies.*

It follows that the survival rule in a symmetric Markov-perfect equilibrium is unique and takes values equal to the symmetric Nash equilibrium strategies of the one-shot game. Because this rule gives firms the individual payoff from joint continuation if positive and zero otherwise (because the equilibrium strategy puts positive probability on exit), we also have

**Corollary 2** *If $v_E$ and $v_S$ are the post-entry and post-survival value functions* *associated with a symmetric Markov-perfect equilibrium, then*

$$v_E(n_E, c, w) = \max\{0, -\kappa \exp(w) + v_S(n_E, c)\}. \tag{6}$$

With Corollaries 1 and 2 in hand, we proceed to demonstrate equilibrium existence constructively. Our equilibrium uniqueness result and algorithm for equilibrium calculation follow from the construction as byproducts. The construction of a candidate equilibrium begins by calculating $v_E(\check{n}, \cdot, \cdot)$ and $v_S(\check{n}, \cdot)$. From Lemma 1, there will be no entry in the next period, so

$$v_S(\check{n}, c) = \rho\mathbb{E}\big[\pi(\check{n}, C') + v_E(\check{n}, C', W')\big|C = c\big]. \tag{7}$$

Using Corollary 2 to replace $v_E(\check{n}, C', W')$ yields

$$v_S(\check{n}, c) = \rho\mathbb{E}\big[\pi(\check{n}, C') + \max\{0, -\kappa \exp(W') + v_S(\check{n}, C')\}\big|C = c\big]. \tag{8}$$

The right-hand side defines a contraction mapping on the complete space of bounded functions on $\mathcal{C}$, with a unique fixed point $v_S(\check{n}, \cdot)$. Although we are constructing a *candidate* equilibrium, the fixed point's uniqueness implies that this is the only possible equilibrium post-survival value. Applying Corollary 2 to this immediately yields $v_E(\check{n}, \cdot, \cdot)$. Again, this is the only possible candidate value. The unique entry rule that is consistent with these payoffs and individual optimality is

$$a_E(\check{n}, c, w) = \mathbb{1}\left[v_E(\check{n}, c, w) > \varphi(\check{n}) \exp(w)\right].$$

Here, $\mathbb{1}(\cdot) = 1$ if $\cdot$ is true and $\mathbb{1}(\cdot) = 0$ otherwise.

With $v_E(\check{n}, \cdot, \cdot)$ and $a_E(\check{n}, \cdot, \cdot)$ in hand, the construction of the remaining candidate value functions and entry strategies proceeds recursively. To this end, define

$$\mu(n, c, w) \equiv n + \sum_{m=n+1}^{\check{n}} a_E(m, c, w). \tag{9}$$

This is the number of firms that will be active after potential entrants follow the candidate entry strategies. For given $n$, suppose that $v_E(m, \cdot, \cdot)$ and $a_E(m, \cdot, \cdot)$ for $m = n+1, n+2, ..., \check{n}$ are in hand. Then, the equilibrium optimality conditions

(3) and (4) along with Corollary 2 imply

$$v_S(n,c) = \rho\mathbb{E}\Big[\pi(n,C') + \sum_{m=n+1}^{\check{n}} \mathbb{1}\{\mu(n,C',W') = m\} v_E(m,C',W') + \mathbb{1}\{\mu(n,C',W') = n\} \max\{0, -\kappa\exp(W') + v_S(n,C')\}\Big| C = c\Big]. \tag{10}$$

Given the values of $v_E(m,\cdot,\cdot)$ for $m = n+1,\ldots,\check{n}$, the right hand side defines a contraction mapping with $v_S(n,\cdot)$ as its unique fixed point. Corollary 2 again yields the unique $v_E(n,\cdot,\cdot)$. Finally, by (3), a firm in state $(n,c,w)$ enters if

$$\mathbb{E}_{a_E}\left[v_E(N_E,c,w)|M = n, C = c, W = w\right] > \varphi(m)\exp(w). \tag{11}$$

Because, by Lemma 2, further entry can only make an incumbent worse off, a necessary condition for (11) is that the firm would enter in the absence of further entry, $v_E(n,c,w) > \varphi(n)\exp(w)$. On the other hand, because later entrants pay (weakly) higher entry costs, further entry will never take post-survival values below the firm's entry cost $\varphi(n)\exp(w)$, and it also suffices for (11) that $v_E(n,c,w) > \varphi(n)\exp(w)$. It follows that

$$a_E(n,c,w) = \mathbb{1}\left[v_E(n,c,w) > \varphi(n)\exp(w)\right].$$

When this recursion is complete, we have the unique continuation values and entry strategies that are consistent with an equilibrium. To find a candidate survival strategy $a_S(n,c,w)$, we find an equilibrium to the one-shot survival game described above. If the candidate is actually an equilibrium, then Corollary 1 guarantees that these survival strategies are unique. This is indeed the case.

**Theorem 1 (Equilibrium existence and uniqueness)** *There exists a unique symmetric Markov-perfect equilibrium that defaults to inactivity, with equilibrium payoffs* $(v_E, v_S)$ *and equilibrium strategy* $(a_E, a_S)$.

This theorem concludes the model's theoretical analysis. We assumed that the costs of entry and continuation are additively separable from per-period flow profits and that $\{C_t, W_t\}$ satisfies Rust's (1987) conditional independence assumption. However, the proof of Theorem 1 requires neither of these conditions.

# 4 Empirical Implementation

The previous section shows that there exists a unique symmetric Markov-perfect equilibrium for given primitives $\pi$, $\kappa$, $\varphi$, $\rho$, $G_C$, and $G_W$. With some distribution of the initial condition $(N_1, C_0)$, this equilibrium implies an equilibrium distribution of the process $\{N_t, C_t\}$. This section studies how data on this process for a panel of markets can be used to infer these markets' primitives.

## 4.1 Sampling

Suppose that we have data on $R \geq 1$ markets $r = 1, \ldots, R$. For each market $r$, we observe the number of active firms $N_{r,t}$ and the demand state $C_{r,t}$ in each period $t = 1, \ldots, T$; for some $T \geq 2$.[3] We also observe some time-invariant characteristics of each market $r$, which we store in a row vector $X_r$. However, we have no data on its cost shocks $W_{r,t}$.

We assume that $(\{N_{r,t}, C_{r,t}; t = 1, \ldots, T\}, X_r)$ is distributed independently across markets $r$.[4] The initial conditions $(N_{r,1}, C_{r,1}, X_r)$ are drawn from some distribution that is common across markets $r$. Conditional on $(N_{r,1}, C_{r,1}, X_r)$, industry dynamics $\{N_{r,t}, C_{r,t}; t = 2, \ldots, T\}$ follow the transition rules implied by Section 2's unique equilibrium, with primitives $\pi_r$, $\kappa_r$, $\varphi_r$, $\rho_r$, $G_{C,r}$, and $G_{W,r}$. The primitives may vary across markets $r$, but with $X_r$ only.[5] We make this explicit by parameterizing $\pi_r(\cdot,\cdot) = \pi(\cdot,\cdot \mid X_r, \theta_P)$, $\kappa_r = \kappa(X_r, \theta_P)$, $\varphi_r(\cdot) = \varphi(\cdot \mid X_r, \theta_P)$, and $\rho_r = \rho(X_r, \theta_P)$ for some common parameter $\theta_P$; $G_{C,r}(\cdot \mid \cdot) = G_C(\cdot \mid \cdot\,; X_r, \theta_C)$ for some parameter $\theta_C$; and $G_{W,r}(\cdot) = G_W(\cdot\,; X_r, \theta_W)$ for some parameter $\theta_W$.

Our estimation procedure is designed for *finite* parameters $\theta_P$, $\theta_C$, and $\theta_W$. For example, in Section 6's empirical illustration, the demand state $C_{r,t}$ will be the population of market $r$ at time $t$ and we will use the specification $\pi_r(c, n) = (c/n)k(n)\exp(X_r\beta)$ of market $r$'s flow surplus, with $k$ the average surplus per

[3]In a typical application, like our empirical illustration in Section 6, $T$ would be small and $R$ would be large. However, the model can be estimated with data on a sufficiently long time series for a single market; that is, with $T$ large and $R = 1$.

[4]Our estimation procedure can be straightforwardly extended to allow for observed (to the econometrician) time-varying covariates that are common across markets, such as business cycle indicators, provided that we make appropriate assumptions on their evolution.

[5]This rules out unobserved (to the econometrician) heterogeneity in the markets' primitives. It is straightforward to introduce e.g. a finite number of unobserved market types and extend the NFXP procedure to estimate a version of the model with such finite unobserved heterogeneity, under the assumption that $\{C_t, W_t\}$ satisfies Rust's (1987) conditional independence assumption conditional on the unobserved heterogeneity. This extension requires a solution to the usual initial conditions problem that, typically, $(N_{r,1}, C_{r,1}, X_r)$ would not be independent of the unobserved heterogeneity (see Footnote 6).

consumer served at $X_r = 0$. Then, $\theta_P$ will include the finite number of values of $k$ and the parameter vector $\beta$. Section 4.3's identification discussion, however, does not rely on parametric restrictions.

## 4.2 Likelihood

We focus on inferring the structural parameters $\theta \equiv (\theta_P, \theta_C, \theta_W)$ from the conditional likelihood $\mathcal{L}(\theta)$ of $\theta$ for data on market dynamics $\{N_{r,t}, C_{r,t}; t = 2, \ldots, T; r = 1, \ldots, R\}$ given the initial conditions $(N_{r,1}, C_{r,1}, X_r; r = 1, \ldots, R)$.[6] Using the model's Markov structure and conditional independence, this likelihood can be written as $\mathcal{L}(\theta) = \mathcal{L}_C(\theta_C) \cdot \mathcal{L}_N(\theta)$, with

$$\mathcal{L}_C(\theta_C) \equiv \prod_{r=1}^{R} \prod_{t=2}^{T} g_C\left(C_{r,t} \mid C_{r,t-1}; X_r, \theta_C\right), \tag{12}$$

the marginal likelihood of $\theta_C$ for the demand state dynamics; and

$$\mathcal{L}_N(\theta) \equiv \prod_{r=1}^{R} \prod_{t=2}^{T} p\left(N_{r,t} \mid N_{r,t-1}, C_{r,t-1}; X_r, \theta\right), \tag{13}$$

the conditional likelihood of $\theta$ for the evolution of the market structures. Here, $g_C(\cdot \mid \cdot; X_r, \theta_C)$ is the density of $G_{C,r}$ with respect to its dominating measure and $p(n'|n, c; X_r, \theta) = \Pr(N_{r,t} = n'|N_{r,t-1} = n, C_{r,t-1} = c; X_r, \theta)$ is the equilibrium probability that market $r$ with $n$ firms and in demand state $c$ has $n'$ firms next period. Note that $\mathcal{L}_C$ can be computed directly from the demand data, without ever solving the model. The remainder of this section focuses on the construction of $\mathcal{L}_N$.

First, we recursively compute the fixed point $v_{S,r}$ of (10) and use this to derive the equilibrium entry rule $a_{E,r}$ and exit rule $a_{S,r}$. Next, we use $a_{E,r}(m+1, c, w)$, $m \geq n$, and $a_{S,r}(m, c, w)$ to construct the equilibrium $\Pr(N_{r,t} = n'|N_{r,t-1} = n, C_{r,t-1} = c, W_{r,t-1} = w; X_r, \theta)$. Finally, we integrate this probability over $w$ with respect to the distribution $G_{W,r}$ of $W_{r,t-1}|X_r, \theta_W$. This gives $p(n'|n, c; X_r, \theta)$. Repeating these calculations for all markets $r$ and substituting into (13) gives $\mathcal{L}_N(\theta)$.

[6] We do not specify nor estimate the distribution of the initial conditions. In particular, we ignore information about $\theta$ in the initial conditions, because we want to be agnostic about their relation to the dynamic model. Alternatively, we could, for example, assume that the initial conditions and covariates are drawn from their ergodic distribution in the dynamic model, which is fully determined by $\theta$. This would allow us to develop a more efficient estimator, at the price of robustness. Moreover, in an extension with unobserved heterogeneity of markets, it would allow us to deal with the initial conditions problems alluded to in Footnote 5.

In practice, we can simplify these calculations considerably by exploiting that the entry and exit rules are monotone in the unobserved cost state. In particular, this monotonicity allows us to (partially) characterize these rules in terms of thresholds on the unobserved cost state. This brings two key simplications. First, it allows us to rewrite (10) so that it no longer involves the full strategy, but only the corresponding thresholds. This simplifies the computation of the fixed points $v_{S,r}$, and the corresponding thresholds, in the first step. Second, it allows for simple expressions of $p(n'|n,c;X_r,\theta)$ in terms of these thresholds.

First, consider the (partial) characterization of the equilibrium strategy in terms of thresholds on the cost state. From Section 3, it follows that market $r$'s entry rule satisfies

$$a_{E,r}(m,c,w) = \mathbb{1}\left[v_{E,r}(m,c,w) > \varphi_r(m)\exp(w)\right].$$

Lemma 2 and (6) imply that $v_{E,r}(m,c,w)$ is weakly decreasing with $w$, so that we can rewrite $a_{E,r}$ as a cost threshold rule:

$$a_{E,r}(m,c,w) = \mathbb{1}\left[w < \overline{w}_{E.r}(m,c)\right],$$

with

$$\overline{w}_{E,r}(m,c) \equiv \log v_{S,r}(m,c) - \log\left[\kappa_r + \varphi_r(m)\right], \tag{14}$$

the value of $w$ that sets $v_{E,r}(m,c,w) = \varphi_r(m)\exp(w)$. Moreover, because $v_{E,r}(m,c,w)$ is weakly decreasing with $m$, $\overline{w}_{E,r}(m,c)$ decreases with $m$. Consequently, a market with $n$ firms and demand state $c$ in period $t-1$ will have exactly $m > n$ firms in period $t$ if $\overline{w}_{E,r}(m+1,c) \leq W_{r,t-1} < \overline{w}_{E,r}(m,c)$. Similarly,

$$\overline{w}_{S,r}(n,c) \equiv \log v_{S,r}(n,c) - \log\kappa_r \tag{15}$$

is the value of $w$ at which $v_{S,r}(n,c) = \kappa_r\exp(w)$. An incumbent firm in market $r$ in demand state $c$ that faces $n-1$ rivals when deciding on survival in period $t-1$ will survive for sure if $W_{r,t-1} < \overline{w}_{S,r}(n,c)$ and exit with positive probability if $W_{r,t-1} > \overline{w}_{S,r}(n,c)$. Therefore, we will refer to $\overline{w}_{S,r}(n,c)$ as a *sure* survival threshold. Because entrants incur a sunk cost, entry implies sure survival: If $W_{r,t-1} < \overline{w}_{E,r}(n,c)$, then $W_{r,t-1} < \overline{w}_{S,r}(n,c)$. Consequently, there cannot be entry and exit within the same period.[7]

Now consider the computation of $\mathcal{L}_N$. First, note that we can rewrite (10) in terms of the entry and sure survival thresholds:

$$
\begin{aligned}
&v_{S,r}(n,c) = \\
&\rho \int_{\mathcal{C}} \Bigg[ \pi_r(n,c') + \int_{\overline{w}_{E,r}(n+1,c')}^{\overline{w}_{S,r}(n,c')} \left[ -\kappa_r \exp(w') + v_{S,r}(n,c') \right] dG_{W,r}(w') \\
&\qquad + \sum_{m=n+1}^{\check{n}} \int_{\overline{w}_{E,r}(m+1,c')}^{\overline{w}_{E,r}(m,c')} \left[ -\kappa_r \exp(w') + v_{S,r}(m,c') \right] dG_{W,r}(w') \Bigg] dG_{C,r}(c'|c).
\end{aligned} \tag{16}
$$

A key advantage of (16) over (10) is that it does not involve the full strategy $(a_{E,r}, a_{S,r})$ on $\mathcal{H}$, but only the corresponding thresholds, which are defined on the smaller space $\{1,\ldots,\check{n}\}\times\mathcal{C}$. Recursively computing the fixed point $v_{S,r}$ of (16), using (14) and (15) to compute $\overline{w}_{E,r}$ and $\overline{w}_{S,r}$ along the way, is straightforward.

With $v_{S,r}$, $\overline{w}_{E,r}$, and $\overline{w}_{S,r}$ in hand, it is straightforward to compute $p(n'|n,c;X_r,\theta)$. There are four cases to consider.

**Case I: $\mathbf{n'>n}$.** If the number of firms increases from $n$ in period $t-1$ to $n'>n$ in period $t$, then it must be profitable for the $n'$th firm to enter, but not for the $(n'+1)$th: $\overline{w}_{E,r}(n'+1,c)\leq W_{r,t-1}<\overline{w}_{E,r}(n',c)$. The probability of this event is

$$
p(n'|n,c;X_r,\theta) = G_{W,r}\left[\overline{w}_{E,r}(n',c)\right] - G_{W,r}\left[\overline{w}_{E,r}(n'+1,c)\right]. \tag{17}
$$

**Case II: $\mathbf{0<n'<n}$.** If the number of firms decreases from $n$ in period $t-1$ to $n'$ in period $t$, with $0<n'<n$, then $W_{r,t-1}$ must take a value $w$ such that firms exit with probability $a_{S,r}(n,c,w)\in(0,1)$. Thus, this value $w$ must be high enough so that $n$ firms cannot survive profitably, $w\geq\overline{w}_{S,r}(n,c)$, but low enough for a single firm to survive profitably, $w<\overline{w}_{S,r}(1,c)$. Given such value $w$ and $(N_{r,t-1}=n, C_{r,t-1}=c, W_{r,t-1}=w; X_r,\theta)$, $N_{r,t}$ is binomially distributed with success probability $a_{S,r}(n,c,w)$ and population size $n$. Hence, the probability of observing a transition from $n$ to $n'$ with $0<n'<n$ equals

$$
\begin{aligned}
&p(n'|n,c;X_r,\theta) \\
&\quad = \int_{\overline{w}_{S,r}(n,c)}^{\overline{w}_{S,r}(1,c)} \binom{n}{n'} a_{S,r}(n,c,w)^{n'} \left[1-a_{S,r}(n,c,w)\right]^{n-n'} g_{W,r}(w)dw,
\end{aligned} \tag{18}
$$

[7] The model can be applied to data with simultaneous entry and exit by specifying data periods as aggregates of model periods, specifying data markets as aggregates of model markets, or both.

where $g_{W,r}$ is the Lebesgue density of $G_{W,r}$. The integrand in (18) involves the mixing probabilities $a_{S,r}(n,c,w)$. We avoid computing these mixing probabilities by substituting for $w$ in (18) the value $\omega_r(a;n,c)$ that sets $a_{S,r}[n,c,\omega_r(a;n,c)] = a$ for given survival probability $a \in (0,1)$. This gives

$$p(n'|n,c;X_r,\theta) = \int_0^1 \binom{n}{n'} a^{n'} (1-a)^{n-n'} \frac{d\omega_r(a;n,c)}{da} g_{W,r}[\omega_r(a;n,c)]\, da. \tag{19}$$

A simple rearrangement of (5) gives an explicit expression for $\omega_r(a;n,c)$:

$$\omega_r(a;n,c) \equiv -\log \kappa_r + \log \sum_{n'=1}^{n} \binom{n-1}{n'-1} a^{n'-1} (1-a)^{n-n'} v_{S,r}(n',c).$$

Using this, and the explicit expression for $d\omega_r(a;n,c)/da$ that can be derived from it, gives an explicit expression for the integrand in (19). We compute the integral with Gauss-Legendre quadrature.

**Case III: $\mathbf{n' = 0}$.** If all firms exit in period $t-1$, then either it is not profitable for even a single firm to continue, $W_{r,t-1} \geq \overline{w}_{S,r}(1,c)$, or it is profitable for a single firm but not for all firms to continue, $\overline{w}_{S,r}(n,c) \leq W_{r,t-1} < \overline{w}_{S,r}(1,c)$, firms exit with probability $a_S(n,c) \in (0,1)$ as in Case II, and by chance none of the $n$ firms survives. The probability of these events is

$$\begin{aligned} &p(0|n,c;X_r,\theta) \\ &\quad = 1 - G_{W,r}[\overline{w}_{E,r}(1,c)] + \int_{\overline{w}_{S,r}(n,c)}^{\overline{w}_{S,r}(1,c)} [1 - a_{S,r}(n,c,w)]^n g_{W,r}(w) dw. \end{aligned} \tag{20}$$

As in Case II, the integral in the right hand side of (20) can be computed by substituting $w = \omega_r(a;n,c)$, which gives

$$\int_0^1 (1-a)^n \frac{d\omega_r(a;n,c)}{da} g_{W,r}[\omega_r(a;n,c)]\, da,$$

and applying Gauss-Legendre quadrature.

**Case IV: $\mathbf{n' = n}$.** If there is entry nor exit in period $t-1$, then either no firm finds it profitable to enter and all $n$ incumbents find it profitable to stay, $\overline{w}_{E,r}(n+1,c) \leq W_{r,t-1} < \overline{w}_{S,r}(n,c)$, or the $n$ incumbents mix as in Cases II and

III, but by chance end up all staying. The probability of these events is

$$p(n|n,c;X_r,\theta) = G_{W,r}\left[\overline{w}_{S,r}(n,c)\right] - G_{W,r}\left[\overline{w}_{E,r}(N+1,c)\right] + \int_{\overline{w}_{S,r}(n,c)}^{\overline{w}_{S,r}(1,c)} a_{S,r}(n,c,w)^n g_{W,r}(w)dw. \tag{21}$$

The integral in (21) can be computed as in Cases II and III.

This completes our construction of $p(n'|n,c;X_r,\theta)$. Substituting in (13) gives $\mathcal{L}_N(\theta)$.

### 4.3 Identification

This section inverts the construction of the likelihood function to consider identification: How can the model's primitives be recovered if we are given the joint distribution of $N', C'$ conditional on $N$, $C$ (and implicitly $X$)? Rust (1994) established the non-identifiability of the discount rate in a decision theoretic model of dynamic discrete choice. Since his fundamental insight holds good in our model, we will assume that auxiliary information that identifies $\rho$, such as the average borrowing rate for small businesses, is in hand. Furthermore, the density of $C'$ given $C$ can be read directly off of the given joint distribution.[8]

The remaining primitives of interest are the model's fixed cost, $\kappa$, sunk cost function $\varphi$, surplus function $\pi$, and the distribution of the econometric error, $G_W$. The observations on the number of producers in $N'$ gives us no information about the level of producer surplus or the magnitude of its variation with the econometric error. Just like discrete choice models, our model requires an a-priori assumption on these quantitites. For this, we set the per-period fixed cost $\kappa$ to one and assume that $W$ has a standard normal distribution.

If $\overline{w}_S(n,c)$ were known, we could recover $\pi(n,c)$ by first using the thresholds' definitions to recover the post-survival values,

$$v_S(n,c) = \exp(\overline{w}_S(n,c)), \tag{22}$$

(where we have imposed $\kappa$'s assumed value of one) and then finding the values of $\pi(n,c)$ for which these value functions satisfy their Bellman equations constructed with the assumed value of $\rho$ and the known transition probabilities for $C'$ and $N'$. To recover these sure survival thresholds, we begin with that for

[8]Above, we specified this density as a function of a vector of parameters, $\theta_C$. Such a parametric restriction might be of use when estimating using a finite sample, but it is not necessary for identification.

a monopolist. Since $\Pr[N' = 0|N = 1, C = c] = 1 - G_W(\overline{w}_S(1,c))$, we have

$$\overline{w}_S(1,c) = G_W^{-1}\left(1 - \Pr[N' = 0|N = 1, C = c]\right).$$

The identification of $\overline{w}_S(n,c)$ for $n > 1$ is complicated by the mixed survival strategies firms follow when $W$ is between it and $\overline{w}_S(1,c)$. For $n = 2$, we observe

$$\Pr(N' = 0|N = 2, C = c) = \\ \Pr[W > \overline{w}_S(1,c)] + \int_{\overline{w}_S(2,c)}^{\overline{w}_S(1,c)} (1 - a_S(2,c,w))^2 g_W(w)dw$$

The first term equals $\Pr[N' = 0|N = 1, C = c]$ and so is known. The second term depends on the known $\overline{w}_S(1,c)$ and the unknown $\overline{w}_S(2,c)$ and $a_S(2,c,w)$. Using the indifference condition for the mixed strategy, we can replace the unknown strategy with a simple function of $w$, $v_S(1,c)$, and $v_S(2,c)$. Substituting the sure survival thresholds for these two value functions with (22) yields

$$\int_{\overline{w}_S(2,c)}^{\overline{w}_S(1,c)} \left(1 - \frac{\exp(\overline{w}_S(1,c)) - \exp(w)}{\exp(\overline{w}_S(1,c)) - \exp(\overline{w}_S(2,c))}\right)^2 g_W(w)dw.$$

This decreases with $\overline{w}_S(2,c)$, so the probability of two dupolists simultaneously exiting identifies it. Proceeding to $n > 2$, a recursive argument shows that the probability that $k \leq n$ oligopolists simultaneously exit identifies $\overline{w}_S(n,c)$. As with $n = 2$, the key to the argument is the recognition that the indifference condition defining the equilibrium strategy gives that strategy as a function of the sure survival thresholds already in hand and the unknown threshold being recovered.

Of course, the survival thresholds contain no information on $\varphi(n)$. To get this, we recover entry thresholds from the probabilities of the number of firms growing. Select an $n' \in \{1, 2, \dots, \check{n}\}$ and $n < n'$. Since $\Pr[N' \geq n'|N = n, C = c] = G_W(\overline{w}_E(n',c))$, we have

$$\overline{w}_E(n',c) = G_W^{-1}\left(\Pr[N' \geq n'|N = n, C = c]\right).$$

From the definition of this threshold, we get

$$\varphi(n',c) = v_S(n',c)\exp(-\overline{w}_E(n',c)) - 1$$

We take two lessons away from this identification argument. First, it is possible to identify the model's parameters without examining the cross-sectional relationship between $N$ and $C$ that Bresnahan and Reiss (1990, 1991) use in their estimation. Second, estimation of our model need not follow the nested fixed point approach that we adopt. In the spirit of Hotz, Miller, Sanders, and Smith (1994), we could instead estimate the equilibrium value functions directly from observed transition probabilities and from these deduce the underlying primitives.

## 4.4 Estimation

We have created C++ and Matlab code for computing a full information maximum likelihood estimator of $\theta$. Following Rust (1994), we compute the estimator in three steps:

1. Estimate $\theta_C$ by computing the marginal likelihood estimator $\tilde{\theta}_C \equiv \arg\max_{\theta_C} \mathcal{L}_C(\theta_C)$;

2. estimate $(\theta_P, \theta_W)$ by computing the conditional likelihood estimator $(\tilde{\theta}_P, \tilde{\theta}_W) \equiv \arg\max_{(\theta_P, \theta_W)} \mathcal{L}_N(\theta_P, \tilde{\theta}_C, \theta_W)$; and

3. estimate $\theta$ by computing the full information maximum likelihood estimator $\hat{\theta} \equiv \arg\max_\theta \mathcal{L}(\theta)$, using $\tilde{\theta}$ as a starting value.

Note that the partial likelihood estimator $\tilde{\theta}$ computed in the first two steps is consistent, but not efficient. Under the usual regularity conditions, the third step's estimator $\hat{\theta}$ has the standard properties of full information maximum likelihood, including asymptotic efficiency. Standard errors are computed using the outer-product-of-the-gradient estimator of the (full) information matrix. In particular, we assume that $R$ is large and $T$ is small and use the average of the outer products of the market-specific gradients over markets, evaluated at $\hat{\theta}$.

The C++ code provides a full implementation of this three-step NFXP procedure for specifications with and without covariates. It uses Knitro for the optimization, with analytical gradients. We use the C++ code for the Monte Carlo experiments in Section 5 and the empirical illustration in Section 6.

The Matlab code provides a more user friendly implementation of the NFXP procedure. It only covers specifications without covariates. It also uses Knitro, but with numerical derivatives. The Matlab code can be used as a sandbox for testing variants of the specifications and procedures and for teaching.

## 5 Monte Carlo Experiments

In this section we investigate the statistical properties and computational performance of our estimation procedure. We set the maximum number of firms sustainable in each market to $\check{n} = 5$. Following the discussion of identification in Section 4.3, we let the cost shocks be standard normally distributed and normalize $\kappa$ to one. We fix the discount factor $\rho$ at $\frac{1}{1.05}$. Each Monte Carlo experiment consists of 1,000 repetitions. We consider six different sample sizes, each of them with ten time periods and between 25 and 1,000 markets. The statistical process governing the demand state is assumed to be known and has support on 200 grid points that are equally spaced on the logarithmic scale. We compute the equilibrium and simulate the evolution of $N$, beginning with a draw of $(N, C)$ from the model's ergodic distribution. We construct the likelihood function as laid out in Section 4. We initialize each estimation procedure by setting all the parameter values to the same randomly drawn number from a uniform distribution on the interval $[1, 10]$. Dubé, Fox, and Su (2012) caution that a nested fixed point algorithm may falsely converge when the tolerance criterion for the inner is set too loosely relative to that of the outer loop. We fix the the tolerance criterion for the value function iteration at a value that is multiple orders of magnitude smaller than that for the outer loop to avoid this potential pitfall.[9]

We first simulate data from a model where the surplus function is parameterized as $\pi_r(c, n) = (c/n)k$, where $k$ is set to 1.5, which means that per consumer surplus is constant in the number of active firms. The sunk cost of entry, $\varphi$, is fixed at 10. The top panel of Figure 2 shows the resulting distribution of the number of firms per market. Table 1 reports the corresponding Monte Carlo experiments' results. The averages of the point estimates are on target even for the smallest sample with only 25 markets. Also note that the means of the asymptotic standard errors almost equal the standard deviations of the estimates across the simulated data sets.

For our second set of simulations we parameterize the flow surplus function as $\pi_r(c, n) = (c/n)k(n)$, where $(k(1), k(2), k(3), k(4), k(5))$ is set to $(1.8, 1.4, 1.2, 1.0, 0.9)$. This specification has the average surplus per consumer decrease in the number of active firms. The bottom panel of Figure 2 displays the distribution of firms per market implied by this model. Table 2 reports the results of the corresponding Monte Carlo experiments. Again, all parameter

[9]We set the tolerance value to $10^{-10}$ for the inner loop and to $10^{-6}$ for the outer loop.

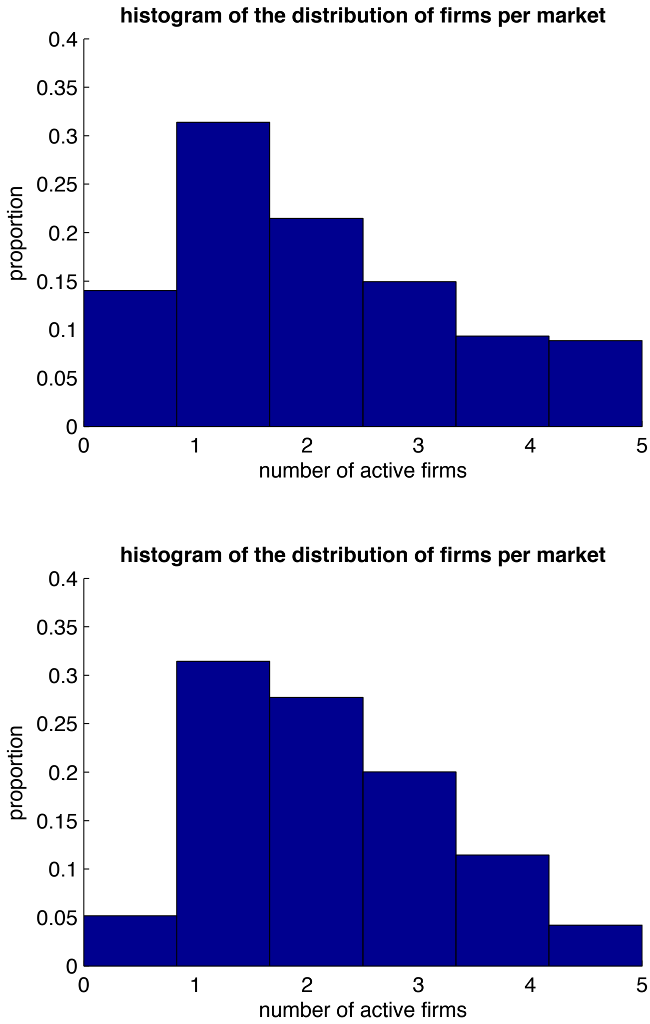


**Figure 2:** Distribution of the number of active firms per market in the specification with a single profit parameter (top panel) and in the specification with five profit parameters (bottom panel).

estimates are essentially without bias. As the number of markets gets larger, all differences between the profit parameterss $k(1), ..., k(5)$ become statistically significant.

Since our equilibrium computation algorithm finds fixed points to relatively low dimensional contraction mappings, one would expect the estimation procedure to be relatively fast. Table 3 shows that this in fact the case. Even in the largest of our Monte Carlo samples, the average computation of the maximum likelihood estimator does not take longer than twenty seconds in a C++ implementation.

Su and Judd's (2012) results suggest that we might be able to improve on the

| | N=25 | N=50 | N=100 | N=250 | N=500 | N=1000 |
|---|---|---|---|---|---|---|
| *means* | | | | | | |
| $\mathbb{E}\hat{k}$ = 1.5 | 1.5081 | 1.5018 | 1.5004 | 1.4991 | 1.4993 | 1.4988 |
| $\mathbb{E}\hat{\varphi}$ = 10.0 | 10.3795 | 10.2146 | 10.0672 | 10.0330 | 10.0130 | 10.0022 |
| $\sqrt{V(\hat{k})}$ | 0.0973 | 0.0684 | 0.0489 | 0.0302 | 0.0222 | 0.0153 |
| $\sqrt{V(\hat{\varphi})}$ | 2.1358 | 1.4623 | 0.9890 | 0.6112 | 0.4443 | 0.3149 |
| *standard errors* | | | | | | |
| $\mathbb{E}SE(\hat{k})$ | 0.0979 | 0.0683 | 0.0480 | 0.0302 | 0.0213 | 0.0151 |
| $\mathbb{E}SE(\hat{\varphi})$ | 2.1204 | 1.4393 | 0.9906 | 0.6208 | 0.4372 | 0.3085 |
| $\sqrt{V(\hat{S}E(k))}$ | 0.0117 | 0.0056 | 0.0027 | 0.0011 | 0.0005 | 0.0003 |
| $\sqrt{V(\hat{S}E(\varphi))}$ | 0.6290 | 0.2866 | 0.1334 | 0.0512 | 0.0262 | 0.0131 |
| *root mean squared errors* | | | | | | |
| $RMSE(\hat{k})$ | 0.0976 | 0.0684 | 0.0489 | 0.0302 | 0.0222 | 0.0153 |
| $RMSE(\hat{\varphi})$ | 2.1682 | 1.4772 | 0.9907 | 0.6118 | 0.4442 | 0.3148 |

**Table 1:** Results of a Monte Carlo experiments using the NFXP estimator with 1,000 repetitions estimating the model with one profit parameter, $k$, and one entry cost parameter, $\varphi$. Demand is discretized into 200 states.

already rapid performance of our estimation procedure by using a mathematical programming with equilibrium constraints (MPEC) procedure in lieu of a nested fixed point algorithm. The MPEC estimator treats the value functions as a vector of nuisance parameters to be estimated subject to the equilibrium constraint implied by the sequence of Bellman equations and thereby omits the inner loop. Our implementation of the MPEC estimator uses analytical gradients of both the objective function and the constraints, and we also passed sparsity patterns to the optimizer. The MPEC estimator always yielded the same estimates as our NFXP procedure, but we found it to be more than ten times slower than the NFXP implementation.[10] MPEC's relatively poor performance reflects the computation of the objective function's gradients with respect to the nuisance parameters, which requires computing and repeatedly retrieving information from very large and relatively dense matrices.[11]

[10]We used the same starting values for both the NFXP and the MPEC algorithm. In addition, we initialized the value functions for the MPEC such that the equilibrium constraint is satisfied with equality.

[11]These computational challenges are not insurmountable and it may well be possible to speed up the MPEC estimator. However, the code involves adjustments and careful handling of very large matrices with up to $32 \times 10^6$ elements for the specification we consider in the Monte Carlo simulation. None of this is required for the NFXP estimator, where all objects are relatively low dimensional.

| | N=25 | N=50 | N=100 | N=250 | N=500 | N=1000 |
|---|---|---|---|---|---|---|
| *means* | | | | | | |
| $\mathbb{E}\hat{k}(1)$ = 1.8 | 1.8880 | 1.8309 | 1.8130 | 1.8048 | 1.8037 | 1.8013 |
| $\mathbb{E}\hat{k}(2)$ = 1.4 | 1.4161 | 1.4068 | 1.3998 | 1.3998 | 1.4005 | 1.4004 |
| $\mathbb{E}\hat{k}(3)$ = 1.2 | 1.1969 | 1.1960 | 1.1986 | 1.1996 | 1.1998 | 1.1998 |
| $\mathbb{E}\hat{k}(4)$ = 1.0 | 0.9960 | 0.9983 | 0.9979 | 0.9984 | 0.9998 | 0.9998 |
| $\mathbb{E}\hat{k}(5)$ = 0.9 | 0.8353 | 0.8680 | 0.8858 | 0.8969 | 0.8965 | 0.8981 |
| $\mathbb{E}\hat{\varphi}$ = 10.0 | 10.3614 | 10.1936 | 10.0921 | 10.0455 | 10.0309 | 10.0121 |
| $\sqrt{V(\hat{k}(1))}$ | 0.3217 | 0.1443 | 0.0909 | 0.0555 | 0.0378 | 0.0270 |
| $\sqrt{V(\hat{k}(2))}$ | 0.1927 | 0.1353 | 0.0956 | 0.0591 | 0.0417 | 0.0292 |
| $\sqrt{V(\hat{k}(3))}$ | 0.1521 | 0.1098 | 0.0752 | 0.0466 | 0.0333 | 0.0238 |
| $\sqrt{V(\hat{k}(4))}$ | 0.1588 | 0.1075 | 0.0730 | 0.0477 | 0.0325 | 0.0229 |
| $\sqrt{V(\hat{k}(5))}$ | 0.2441 | 0.1569 | 0.0897 | 0.0555 | 0.0394 | 0.0272 |
| $\sqrt{V(\hat{\varphi})}$ | 2.7556 | 1.9171 | 1.2677 | 0.7794 | 0.5354 | 0.3719 |
| *standard errors* | | | | | | |
| $\mathbb{E}SE(\hat{k}(1))$ | 0.2928 | 0.1338 | 0.0867 | 0.0527 | 0.0369 | 0.0259 |
| $\mathbb{E}SE(\hat{k}(2))$ | 0.2488 | 0.1422 | 0.0947 | 0.0584 | 0.0410 | 0.0289 |
| $\mathbb{E}SE(\hat{k}(3))$ | 0.2108 | 0.1158 | 0.0783 | 0.0485 | 0.0341 | 0.0240 |
| $\mathbb{E}SE(\hat{k}(4))$ | 0.2332 | 0.1145 | 0.0747 | 0.0460 | 0.0323 | 0.0228 |
| $\mathbb{E}SE(\hat{k}(5))$ | 0.3411 | 0.1583 | 0.0914 | 0.0552 | 0.0388 | 0.0272 |
| $\mathbb{E}SE(\hat{\varphi})$ | 2.9538 | 1.8297 | 1.2272 | 0.7558 | 0.5301 | 0.3727 |
| $\sqrt{V(\hat{SE}(k(1)))}$ | 0.5456 | 0.0649 | 0.0159 | 0.0054 | 0.0026 | 0.0012 |
| $\sqrt{V(\hat{SE}(k(2)))}$ | 0.1834 | 0.0396 | 0.0121 | 0.0045 | 0.0023 | 0.0011 |
| $\sqrt{V(\hat{SE}(k(3)))}$ | 0.1897 | 0.0297 | 0.0121 | 0.0033 | 0.0016 | 0.0008 |
| $\sqrt{V(\hat{SE}(k(4)))}$ | 0.2777 | 0.0603 | 0.0099 | 0.0031 | 0.0016 | 0.0008 |
| $\sqrt{V(\hat{SE}(k(5)))}$ | 0.4763 | 0.1601 | 0.0231 | 0.0060 | 0.0030 | 0.0014 |
| $\sqrt{V(\hat{SE}(\varphi))}$ | 1.4053 | 0.5067 | 0.2139 | 0.0783 | 0.0372 | 0.0180 |
| *root mean squared errors* | | | | | | |
| $RMSE(\hat{k}(1))$ | 0.3334 | 0.1475 | 0.0918 | 0.0557 | 0.0379 | 0.0270 |
| $RMSE(\hat{k}(2))$ | 0.1933 | 0.1354 | 0.0956 | 0.0591 | 0.0417 | 0.0292 |
| $RMSE(\hat{k}(3))$ | 0.1521 | 0.1098 | 0.0752 | 0.0466 | 0.0333 | 0.0238 |
| $RMSE(\hat{k}(4))$ | 0.1588 | 0.1074 | 0.0729 | 0.0477 | 0.0325 | 0.0229 |
| $RMSE(\hat{k}(5))$ | 0.2524 | 0.1600 | 0.0907 | 0.0555 | 0.0396 | 0.0272 |
| $RMSE(\hat{\varphi})$ | 2.7777 | 1.9259 | 1.2704 | 0.7803 | 0.5360 | 0.3719 |

**Table 2:** Results of a Monte Carlo experiments using the NFXP estimator with 1,000 repetitions estimating the model with five profit parameters $k(1), k(2), ..., k(5)$ and one entry cost parameter $\varphi$. Demand is discretized into 200 states.

|  | N=25 | N=50 | N=100 | N=250 | N=500 | N=1000 |
|---|---|---|---|---|---|---|
| *one entry cost parameter, five profit parameters* |  |  |  |  |  |  |
| time per run (in seconds) | 1.89 | 2.08 | 2.47 | 3.57 | 5.39 | 9.08 |
| iterations | 13.54 | 13.43 | 13.44 | 13.50 | 13.50 | 13.60 |
| *five profit parameters, one entry cost parameter* |  |  |  |  |  |  |
| time per run (in seconds) | 4.82 | 5.13 | 5.94 | 8.35 | 11.91 | 18.83 |
| iterations | 27.60 | 27.27 | 27.58 | 27.96 | 27.46 | 26.66 |

**Table 3:** Average computational performance of the NFXP estimators in the Monte Carlo samples. The estimators is implemented in C++ and Knitro and runs as a single thread on a 3GhZ Intel Xeon CPU.

## 6 Empirical Illustration

This section illustrates the empirical application of our model and its estimator with observations from the Motion Picture Theaters industry (NAICS code: 512131). The observations of producer counts come from the County Business Patterns dataset from 2000 through 2009. We restrict our analysis to Micropolitan Statistical Areas ($\mu$SAs), a set of geographic entities defined by the Office of Management and Budget consisting of 574 areas in the United States based around an urban core of at least 10,000 but less than 50,000 inhabitants.[12] The 574 $\mu$SAs account for about ten percent of the United States population.[13] By definition, $\mu$SAs are approximately geographically isolated from other urban centers and thus represent a convenient market definition for the study of industry dynamics. The Census Bureau publishes annual $\mu$SA population estimates, and we use these as the demand indicator, $C$. For each market we include a set of covariates taken from the 2010 wave of the American Community Survey. These include the proportion of the population aged 65 and older, the proportion of the population employed, the proportion of the population employed in agriculture and related industries, median household income, and the median value of owner-occupied housing. Given the short panel, we treat these market characteristics as time-invariant. Table 4 reports their descriptive statistics.

In the first step, we need to estimate the empirical analog of the transition density of the demand process, $g_C\left(C_{r,t} \mid C_{r,t-1}; X_r, \theta_C\right)$. We discretize the demand state into 200 grid points and denote the grid by $\{C_{[1]}, C_{[2]}, ..., C_{[200]}\}$. The grid points are chosen such that they are equally spaced on the logarithmic scale with distance $d$. We then follow Tauchen (1986) and specify the probability of transitioning from $C_{[j]}$ to $C_{[i]}$ for any $i = 2, ..., 199$ and $j = 1, ..., J$, by

$$\Pr(C_{[i]}|C_{[j]}) = \Phi\left(\frac{\log C_{[i]} - \log C_{[j]} + \frac{d}{2} - \mu}{\sigma}\right) - \Phi\left(\frac{\log C_{[i]} - \log C_{[j]} - \frac{d}{2} - \mu}{\sigma}\right),$$

where $\Phi$ refers to the standard normal cumulative distribution function. The probabilities of transitioning to one of the end points of the grid

[12]The classification of Micropolitan Statistical Areas has been subject to several revisions in the past years. We use the release of the "Annual Estimates of the Population of Metropolitan and Micropolitan Statistical Areas from April 1, 2000 to July 1, 2009" from the US Census Bureau as baseline for our analysis, which includes information on 574 $\mu$SAs.

[13]For the purpose of this paper, we dropped the $\mu$SA "The Villages, FL", because its population almost doubled during the time horizon of this study making it very different from the other markets that we consider.

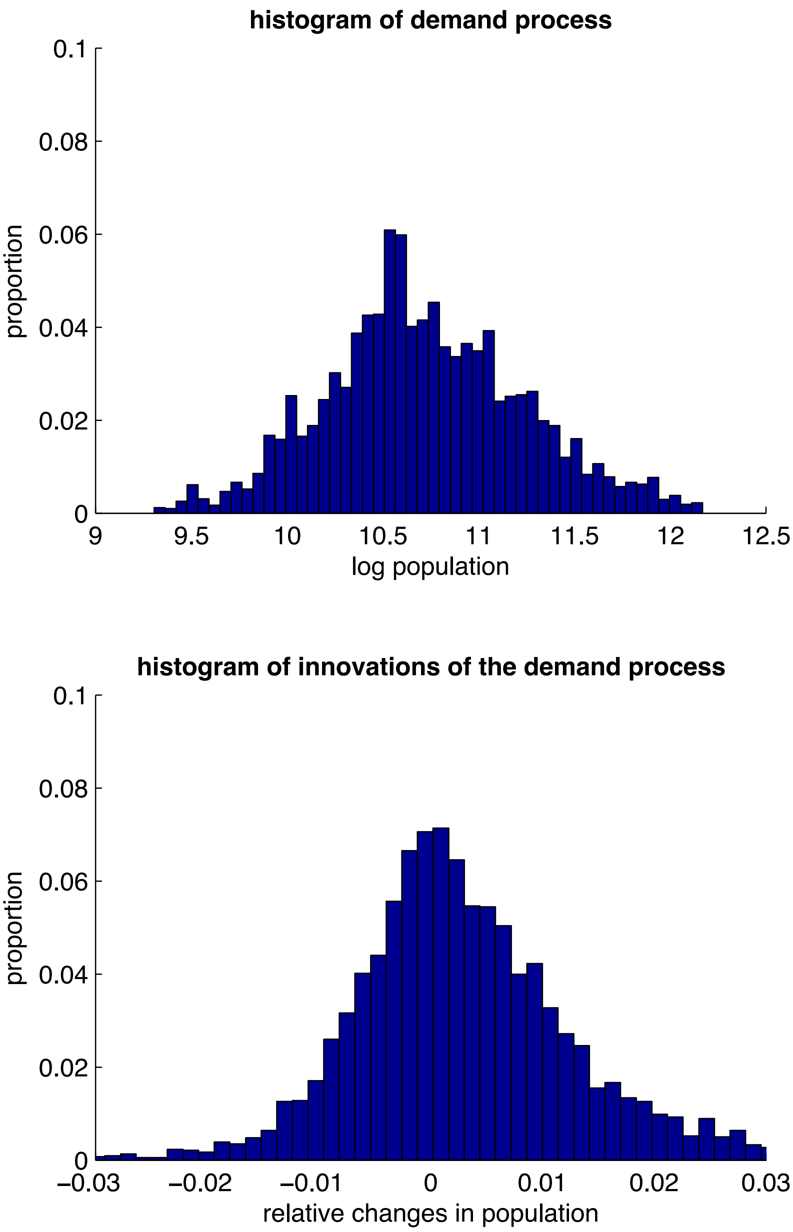


**Figure 3:** Distribution of log population and log population differences for 573 Micropolitan Statistical Areas from 2000 through to 2009.

| | mean | median | st. dev. | max | min |
|---|---|---|---|---|---|
| *establishments* | | | | | |
| number of movie theaters | 1.2837 | 1 | 1.0898 | 9 | 0 |
| *covariates* | | | | | |
| population | 52019 | 43629 | 29461 | 197912 | 11011 |
| employed | 25722 | 21068 | 15508 | 107531 | 4381 |
| 65 and older (%) | 15.38 | 15.21 | 3.39 | 36.71 | 6.16 |
| employed in agriculture (%) | 2.25 | 1.63 | 2.09 | 18.15 | 0.07 |
| median income in USD | 42117 | 41288 | 8450 | 103643 | 22881 |
| median house value in USD | 126221 | 103300 | 71649 | 614600 | 31400 |

**Table 4:** Number of movie theaters for 573 Micropolitan Statistical Area ($\mu$SAs) from the County Business Patterns from 2000 to 2009 and descriptive statistics of the covariates from the 2010 American Community Survey. Population in agriculture refers to the share of population employed in agriculture, forestry, fishing, hunting, and mining.

are given by $\Pr(C_{[1]}|C_{[j]}) = \Phi\left(\frac{\log C_{[1]} - \log C_{[j]} + \frac{d}{2} - \mu}{\sigma}\right)$ and $\Pr(C_{[200]}|C_{[j]}) = 1 - \Phi\left(\frac{\log C_{[200]} - \log C_{[j]} - \frac{d}{2} - \mu}{\sigma}\right)$ respectively. We then estimate the parameters $\mu$ and $\sigma$ with maximum likelihood. The standard deviation, $\sigma$, of the underlying normal distribution is equal to about 1.2%. The drift parameter, $\mu$, is equal to about 0.3%.

With the first step estimates in hand, we proceed by estimating four specifications of our model. The maximum number of movie theaters sustainable is fixed at ten, which is above the maximum number of theaters observed in our data, nine (see Table 4). We maximize the likelihood following the procedure described in Section 4. The covariates from the 2010 American Community Survey enter market $r$'s flow surplus through $\pi_r(c, n) = (c/n)k(n)\exp(X_r\beta)$, where $X_r$ is a vector of covariates for market $r$ and $\beta$ is a vector of coefficients. All covariates are demeaned and included in logs. The per consumer surplus $k(n)$ is constrained to be weakly decreasing with the number of firms which is stronger than required by the assumptions in Section 2. Table 5 reports the results.

The first specification that we estimate, reported in the first column, holds $k(n)$ constant and includes no covariates. The sunk cost of entry is 7.5 times the fixed cost. The surplus from serving 10,000 people is only 47% of the fixed cost. The specification reported in the second column includes a vector of covariates. Although three of the five covariates have statistically significant coefficients, our estimates for $k$ and $\varphi$ do not meaningfully change. For the third specification we

| | constant profits | | varying profits | |
|---|---|---|---|---|
| $k(1)$ | 0.4694 | 0.5014 | 0.4763 | 0.5138 |
| | (0.0061) | (0.0071) | (0.0070) | (0.0083) |
| $k(2)$ | | | 0.4454 | 0.4671 |
| | | | (0.0081) | (0.0093) |
| $k(3)$ | | | 0.4454 | 0.4671 |
| | | | (0.0113) | (0.0126) |
| $k(4)$ | | | 0.4413 | 0.4533 |
| | | | (0.0145) | (0.0165) |
| $k(5, ..., 10)$ | | | 0.4100 | 0.4066 |
| | | | (0.0189) | (0.0210) |
| $\varphi$ | 7.5150 | 7.4387 | 6.9138 | 6.6211 |
| | (0.2566) | (0.2625) | (0.2591) | (0.2587) |
| Population aged 65 and older | | -0.0554 | | -0.0457 |
| | | (0.0543) | | (0.0556) |
| Population Employed | | 0.7575 | | 0.8094 |
| | | (0.1485) | | (0.1518) |
| Population Employed in Agriculture | | 0.2124 | | 0.2039 |
| | | (0.0159) | | (0.0162) |
| log(Median Income) | | 0.1384 | | 0.1802 |
| | | (0.0997) | | (0.1013) |
| log(Median Housing) | | 0.0801 | | 0.1203 |
| | | (0.0348) | | (0.0373) |
| $-\mathcal{L}$ | 3561.93 | 3427.18 | 3553.19 | 3411.51 |

**Table 5:** Coefficient estimates for the NFXP estimator for 573 $\mu$SAs from 2000 to 2009. The first step is estimated using the Tauchen approximation. Standard errors are in parentheses. Standard errors are not adjusted for the first step estimation.

allow producers' surplus per customer to decline with the number of producers. The surplus per customer is freely parameterized for the first four firms in a market and held constant from the fifth entrant onwards. We reject the null that per consumer surplus is constant in the number of active firms. The specification reported in the fourth column adds covariates to this. By taking log ratios of the estimates for $k(n)$, we find that adding a single firm to a monopoly market lowers profits per customer by 10 percent. Adding a third or fourth firm lowers profits only little, while the fifth firm reduces profits by an additional 13 percent.

# 7 Conclusion

We have demonstrated uniqueness of our model's symmetric Markov-perfect equilibrium, provided an algorithm for its fast calculation, shown that its parameters can be identified from observations on the joint evolution of demand and the number of active firms, provided a nested fixed-point algorithm for its maximum-likelihood estimation, evaluated the estimator's statistical properties and computational burden with Monte Carlo experiments, and applied all of these tools to estimate the toughness of competition between Motion Picture Theaters in U.S $\mu$SAs. That this relatively complete development and application of a dynamic oligopoly model was feasible validates our title's assertion that our model's dynamics are "Very Simple."

We anticipate three applications of our model and its maximum-likelihood estimator. First, they can be used to estimate the impact of observable cross-market heterogeneity on the primitive determinants of industry dynamics. For example, one might speculate that differences in zoning create an economic barrier to entry that protects incumbent firms and allow measurable aspects of zoning to influence the sunk costs of entry by including them in $x$. Second, our model is simple enough for inclusion as a moving part in general equilibrium models with entry, exit, and endogenous markups, such as Jaimovich's (2007). Third, the estimated model can serve as a point of departure for an analysis with a more computationally and theoretically demanding model. For example, in a trade context a potential entrant might choose between two imperfectly integrated markets. By estimating our model first, one can gain familiarity with the industry's dynamics and obtain starting values for homotopy-based estimation and equilibrium calculation.

## References


Abbring, J. H., J. R. Campbell, and N. Yang (2010): "Simple Markov-Perfect Industry Dynamics," Working Paper 2010–21, Federal Reserve Bank of Chicago, Chicago. 1, 2, 8

Bajari, P. L., C. L. Benkard, and J. Levin (2007): "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 75(5), 1331–1370. 2

Bresnahan, T. F., and P. C. Reiss (1990): "Entry in Monopoly Markets," *Review of Economic Studies*, 57(4), 531–553. 1, 19

——— (1991): "Entry and Competition in Concentrated Markets," *Journal of Political Economy*, 99(5), 977–1009. 1, 19

Doraszelski, U., and A. Pakes (2007): "A Framework for Applied Dynamic Analysis in IO," in *Handbook of Industrial Organization*, ed. by M. Armstrong, and R. H. Porter, vol. 3, chap. 4. Elsevier Science, Amsterdam. 2

Dubé, J.-P., J. T. Fox, and C.-L. Su (2012): "Improving the Numerical Performance of Static and Dynamic Aggregate Discrete Choice Random Coefficients Demand Estimation," *Econometrica*, 80(5), 2231–2267. 20

Ericson, R., and A. Pakes (1995): "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *Review of Economic Studies*, 62, 53–82. 2

Fudenberg, D., and J. Tirole (1991): *Game Theory*. MIT Press, Cambridge, MA. 6

Hopenhayn, H. A. (1992): "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica*, 60(5), 1127–1150. 1

Hotz, V. J., and R. A. Miller (1993): "Conditional Choice Probabilities and the Estimation of Dynamic Models," *Review of Economic Studies*, 60(3), 497–529. 2

Hotz, V. J., R. A. Miller, S. G. Sanders, and J. A. Smith (1994): "A Simulation Estimator for Dynamic Models of Discrete Choice," *Review of Economic Studies*, 61(2), 265–289. 19

Jaimovich, N. (2007): "Firm Dynamics and Markup Variations: Implications for Sunspot Equilibria and Endogenous Economic Fluctuations," *Journal of Economic Theory*, 137(1), 300–325. 29

Maskin, E., and J. Tirole (1988): "A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs," *Econometrica*, 56(3), 549–569. 5

McAfee, R. P., H. M. Mialon, and M. A. Williams (2004): "What is a Barrier to Entry?," *American Economic Review*, 94(2), 461–465. 4

Melitz, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71(6), 1695–1725. 1

Rust, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55, 999–1033. 0, 1, 4, 11, 12

——— (1988): "Maximum Likelihood Estimation of Discrete Control Processes," *SIAM Journal on Control and Optimization*, 26, 1006–1024. 42

——— (1994): "Structural Estimation of Markov Decision Processes," in *Handbook of Econometrics*, ed. by R. Engle, and D. McFadden, vol. 4, pp. 3081–3143. North-Holland, Amsterdam. 17, 19

Stokey, N. L., and R. E. Lucas, Jr. (1989): *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge, MA. 39, 40

Su, C.-L., and K. L. Judd (2012): "Constrained Optimization Approaches to Estimation of Structural Models," *Econometrica*, 80(5), 2213–2230. 2, 21

Sutton, J. (1991): *Sunk Costs and Market Structure*. MIT Press. 3

Tauchen, G. (1986): "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions," *Economics letters*, 20(2), 177–181. 25

Weintraub, G. Y., C. L. Benkard, and B. Van Roy (2008): "Markov Perfect Industry Dynamics with Many Firms," *Econometrica*, 76(6), 1375–1411. 2

Xu, Y. (2008): "A Structural Empirical Model of R&D, Firm Heterogeneity and Industry Evolution," Mimeo, Department of Economics, New York University, New York, NY. 2