Very Simple Markov-Perfect Industry Dynamics

Abbring, J.H.; Campbell, J.R.; Tilly, J.; Yang, N.

Publication date:
2014

Citation for published version (APA):
Abbring, J. H., Campbell, J. R., Tilly, J., & Yang, N. (2014). *Very Simple Markov-Perfect Industry Dynamics*. (CentER Discussion Paper; Vol. 2014-008). Econometrics.

CentER Discussion Paper

No. 2014-008

VERY SIMPLE MARKOV-PERFECT INDUSTRY DYNAMICS

By

Jaap H. Abbring, Jeffrey R. Campbell,
Jan Tilly, Nan Yang

15 January, 2014

ISSN 0924-7815
ISSN 2213-9532

Tilburg University

Very Simple Markov-Perfect Industry Dynamics

Jaap H. Abbring*
Jeffrey R. Campbell†
Jan Tilly‡
Nan Yang§

January, 2014

Abstract

This paper develops an econometric model of industry dynamics for concentrated markets that can be estimated very quickly from market-level data on demand shifters and the number of producers. We show that the model has an essentially unique symmetric Markov-perfect equilibrium that can be calculated from the fixed points of low-dimensional contraction mappings. We characterize the model's identification and extend Rust's (1987) nested fixed point estimator to account for the observable implications of mixed strategies on survival. We illustrate the model's application with ten years of County Business Patterns data from Motion Picture Theaters in 573 Micropolitan Statistical Areas.

*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@econ.upenn.edu.

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

JEL Codes: L13, C25, C73.

Keywords: 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. Nan Yang gratefully acknowledges the financial support from the National University of Singapore Start-up Grant R-313-000-100-133. 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

In this paper, we present an econometric model of firm entry, competition, and exit in oligopolistic markets. It features toughness of competition, sunk entry costs, and market-level demand and cost shocks. We allow firms to use mixed strategies and close the model by focusing on symmetric Markov-perfect equilibria. The model's key simplifying assumption is that active producers are homogeneous, as in Bresnahan and Reiss (1990, 1991). Using this and the equilibrium implications of mixed strategies for payoffs, we prove that adding a competitor cannot increase incumbents' equilibrium continuation values. This result in turn ensures not only that a symmetric equilibrium exists but also that it is unique and can be quickly computed from the fixed points of a finite sequence of low-dimensional contraction mappings. Because our algorithm relies on contraction mappings, it is *guaranteed* to calculate the equilibrium.

We use these results to analyze the model's empirical content (given market-level panel data on demand shifters and the number of producers) and to develop a nested fixed point (NFXP) estimation procedure that extends Rust's (1987) algorithm for full information maximum likelihood estimation. One novel aspect of our NFXP procedure is that it accounts for the observable implications of mixing over survival actions. Because equilibrium calculation requires so little time, it can (and in our illustrative application does) calculate a separate equilibrium for each of several hundred heterogeneous markets at every trial value of the parameters. Of course, the estimator is asymptotically efficient. With a Monte Carlo study, we demonstrate that it also performs well in small samples.

To illustrate further the model's applicability and ease of use in an applied setting, we estimate its primitives for one industry with many concentrated local markets, Motion Picture Theaters (NAICS 512131). Our data include observations on the number of theaters from 2000 to 2009 serving 573 Micropolitan Statistical Areas (μSAs). We find that adding a single theater to a monopoly market lowers the producers' surplus per consumer by almost half. Adding two more theaters brings the per consumer surplus to 34 percent of its monopoly value.

Our analysis extends Bresnahan and Reiss's (1990; 1991) approach to the measurement of the effects of entry on profitability to a dynamic setting. We begin with Abbring and Campbell's (2010) model of last-in first-out oligopoly dynamics.

In their model, entry and continuation incur sunk and fixed costs, demand is stochastic, and incumbent firms make continuation decisions sequentially in the order of their entry. They restricted attention to Markov-perfect equilibria in which older firms' first-mover advantages allow them to outlive their younger rivals for sure. They showed that there exists a unique such equilibrium and that— with additional restrictions on the stochastic process for demand— firms use pure entry and survival strategies that depend on demand thresholds (see also Abbring and Campbell, 2013). That result gives a structural interpretation to the dynamic ordered probit estimated by Bresnahan and Reiss (1993).

In this paper, we replace Abbring and Campbell's (2010) (admittedly special) sequential continuation decisions with simultaneous ones. Our equilibrium characterization and computation do not rely on Abbring and Campbell's sequential timing assumptions, nor on their restriction to last-in first-out dynamics, but instead leverage the properties of mixed-strategy equilibria in a novel way. Also, unlike Abbring and Campbell, we provide a full econometric development of our model. To this end, we 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. As in Bresnahan and Reiss (1990, 1991), this market-level cost shock serves as the model's econometric error.

We explore the empirical content of our model and provide conditions for its identification based on minimal and readily available panel data on the number of producers and demand shifters. Our result shows that the structure implied by the firms' use of a mixed strategy identifies the scale of the econometric error's distribution without *a priori* restrictions on how demand or another observable regressor influences payoffs. This distinguishes our identification results from those for single agent models (Magnac and Thesmar, 2002), which require either a known distribution of the econometric error or further structure on the payoffs.

Our model can be viewed as a special case of Ericson and Pakes's (1995) Markov-perfect industry dynamics framework. Ericson and Pakes focused on equilibria in pure strategies. Doraszelski and Satterthwaite (2010) noted that such equilibria may not exist, contrary to an assertion in the original article. Gowrisankaran (1999) ensured equilibrium existence by augmenting the Ericson and Pakes framework with firm-specific privately-observed shocks to the costs of continuation. Further research using this framework (summarized by Doraszelski and Pakes, 2007) has adopted

Gowrisankaran's extended version of Ericson and Pakes's framework. The primitives of this modified framework can be estimated by applying standard methods for dynamic discrete choice models, such as Hotz and Miller's (1993), to the isolated decision problems that each of the model's oligopolists faces after conditioning on its rivals' (observed) equilibrium choices. Bajari et al. (2007) exemplify this approach.

In this paper, we return to Ericson and Pakes's original complete-information approach and show that it can bear fruit. Our model does *not* incorporate firm-specific shocks to the costs of continuation, and its unique symmetric equilibrium strategy generically dictates that firms mix their survival actions in some states that occur with positive probability. This modeling choice renders the many methods developed to estimate the modified Ericson and Pakes framework not applicable to our model without substantial modifications, but it also provides a benefit that more than outweighs this cost: Because firms choosing a mixed strategy always earn the value of the outside option (zero) in expectation, *symmetric equilibrium payoffs equal the maximum of zero and the value of all incumbents choosing certain continuation.* This insight underlies the contraction mappings we use to both calculate the equilibrium and demonstrate its uniqueness.

Our model occupies a particular point on the production possibilities frontier that balances analytical and computational convenience with model complexity. Those who have developed and extended the Ericson and Pakes framework have chosen a very different point that favors model complexity. To illustrate, Doraszelski and Pakes (2007) state

> Indeed by adopting the framework the researcher essentially gives up on analytic elegance in favor of an ability to numerically analyze the more complex situations that might better approximate what we seem to observe in real data sets. Theoretical results from more stylized environments are often used at a later stage as a guide to understanding the economics underlying the phenomena generated by the numerical results.

By its very nature, our model can serve as one of these "more stylized environments." However, it is primarily of use in projects that require many equilibrium calculations, in which "babysitting" a procedure that might not converge would be impractical.

One example of such a project is the analysis of data sets with many heterogeneous markets but little *a priori* information on how firms strategically interact. Some industrial organization economists have structured their studies of such data around models of perfect (Hopenhayn, 1992) or Chamberlinian monopolistic (Melitz, 2003) competition. However, Campbell and Hopenhayn (2005) and Campbell (2010) empirically argue that strategic interaction between small businesses is widespread. Our work can extend models of competitive industry dynamics so that they do not rely on an unrealistic absence of strategic interaction. Because we abstract from the full complexity of Ericson and Pakes's framework, we can easily estimate such extensions— even with a limited number of observations for each distinct market in the data— and analyze them in large-scale computational experiments.

Another example is the construction and analysis of macroeconomic models with non-trivial firm interactions, such as that of Jaimovich (2007). Our model can straightforwardly be incorporated into a dynamic general equilibrium environment to account for variation in firm entry, competition and exit over the business cycle. Because its equilibrium calculation is fast and always convergent, this can be done without unduly complicating the resulting computational analysis. In contrast, difficulties with equilibrium selection and algorithmic convergence make such a standard computational approach to macroeconomic analysis infeasible with an environment that incorporates Ericson and Pakes's framework in its full generality.

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, which includes sampling, likelihood construction, identification, and maximum likelihood estimation using the NFXP procedure. Section 5 demonstrates the light computational demands of the NFXP procedure and explores the estimator's finite sample behavior using Monte Carlo experiments. Section 6 illustrates the model's application with an empirical analysis of Motion Picture Theaters. Section 7 concludes with a perspective on extending this paper's approach to models with commonly observed firm-specific shocks, as in Ericson and Pakes's original framework, along the lines of Abbring et al. (2010). The Appendix provides technical details.

2 The Model

Consider a market in discrete time indexed by $t \in \mathbb{N} \equiv \{1, 2, \ldots\}$. In period t, firms that have entered in the past and not yet exited serve the market. Each firm has a name $f \in \mathcal{F} \equiv \mathbb{N} \times \mathbb{N}$. The first component of a firm's name gives the date in which it has its single opportunity to enter the market, and the second component gives its position in that date's entry queue. 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. If $t = 1$, nature sets N_1, the number of firms serving the market in period 1, and C_1, the initial demand state; if $t > 1$, these are inherited from the previous period. We use $\mathcal{C}$ to denote the support of C_t. Although we consistently refer to C_t as "demand," it can encompass any observable, relevant, and time-varying characteristics of the market.

Each incumbent firm earns a surplus $\pi(N_t, C_t)$ from serving the market, and all firms value future profits and costs with the discount factor $\rho \in [0, 1)$. We assume that

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

Here and throughout; we denote the next period's value of a generic variable Z with Z', random variables with capital Roman letters, and their realizations with the corresponding small Roman letters. 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 each incumbent's surplus. That is, what Sutton (1991) labelled the *toughness of competition* must dominate any exogenous complementarities between firms' activities like those considered by Honoré and De Paula (2010).

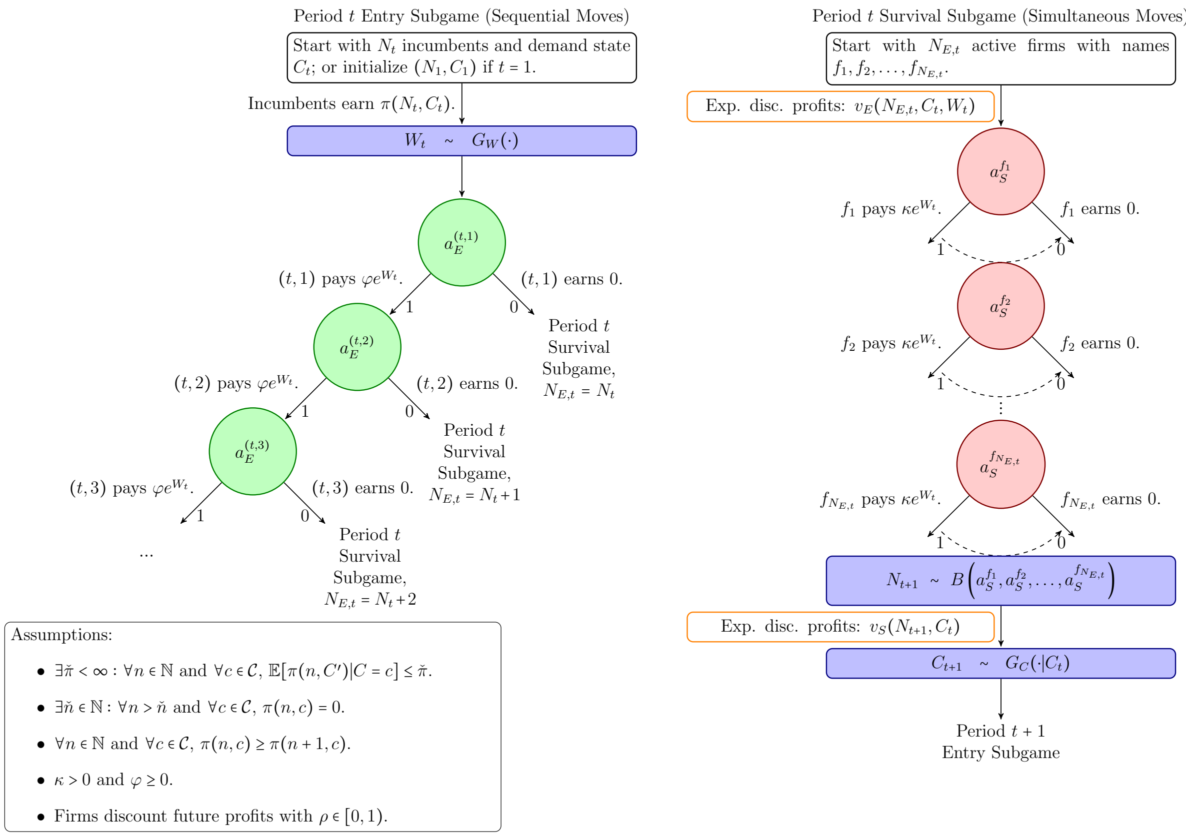

Figure 1: The Model's Recursive Extensive Form

After incumbents earn their surplus, nature draws the current period's real-valued shock to continuation and entry costs from a distribution G_W with positive density everywhere on the real line. Then, entry *per se* can take place. The period t entry cohort consists of firms with names in $\{t\} \times \mathbb{N}$. These firms make their entry decisions sequentially in the order of their names' second components. We denote firm f's entry decision with $a_E^f \in \{0, 1\}$. An entrant pays the sunk cost $\varphi \exp(W_t)$, with $\varphi \geq 0$. A firm choosing not to enter earns a payoff of zero and never has another entry opportunity. Such a 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 does not affect any symmetric equilibrium outcome. 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. Denote their names with $f_1, \ldots, f_{N_{E,t}}$. In the survival subgame, these firms simultaneously choose probabilities of remaining active, $a_S^{f_t}, \ldots, a_S^{f_{N_{E,t}}} \in [0,1]$. Subsequently, all survival outcomes are realized independently across firms according to 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, $t+1$. To end the period, nature draws a new demand state C_{t+1} from the conditional distribution $G_C(\cdot \mid C_t)$.

Before continuing to the model's analysis, we review its key assumptions from the perspective of its econometric implementation using market-level panel data. 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 make the relation between the market structure N_t and the observed demand state C_t 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 distribution of (N_t, C_t) conditional on $(N_{t^\star}, C_{t^\star})$ for all $t^\star < t$ depends only on (N_{t-1}, C_{t-1}), so we require only the model's transition rules to calculate the conditional likelihood function.

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.

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

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 state variables are the number of firms in the market including all of the current period's entrants up to and including (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 $\mathcal{H}$-valued vector (M_t^j, C_t, W_t), with $\mathcal{H} \equiv \mathbb{N} \times \mathcal{C} \times \mathbb{R}$. Similarly, we collect the payoff-relevant state variables of a firm f contemplating survival in period t in the $\mathcal{H}$-valued $(N_{E,t}, C_t, W_t)$. 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 and firm names themselves are not payoff relevant, we drop the subscript t and the superscript j from the payoff-relevant states.

3.2 Symmetric Markov-Perfect Equilibrium

In a symmetric Markov-perfect equilibrium, a firm's expected continuation value at a particular node of the game can be written as a function of that node's payoff-relevant state variables. Two of these 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 in a market with demand state C,cost shock W and 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 the same market with N' firms just after the survival outcomes are realized. The post-survival value does not depend on W because that cost shock has no forecasting value and is not directly payoff relevant after survival decisions are made. Figure 1 shows the points in the survival subgame where these value functions apply.

The payoff from leaving the market equals zero, so v_E and v_S satisfy

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

The expectation $\mathbb{E}_{a_S}$ over N' takes survival of the firm of interest as given. Its subscript makes its dependence on a_S explicit. It conditions on the current values of C and W because these influence the survival probability's value. Similarly, we have

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

This expectation operator's subscript indicates its dependence on a_E.

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 a one-shot deviation. Thus, given the pair of payoff functions (v_E, v_S), their corresponding strategy must satisfy

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

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

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. To avoid these uninteresting caveats to our results, we follow Abbring and Campbell (2010) by focusing on equilibria that *default to inactivity*. In such an equilibrium, a potential entrant that is indifferent between entering or not stays out,

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

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

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

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 subsection presents our analysis of equilibrium existence, uniqueness, and computation. For this, three features of the model are disposable: the serial independence of W_t, the additive separability of per-period surplus from the costs of continuation, and the invariance of sunk costs to the number of firms and the current demand state. In the Appendix, Sections A and B generalize the model by removing these assumptions and Section C provides proofs of all of this subsection's results extended to that more general model.

We start by noting that the assumption that per-period surplus equals zero if more than $\check{n}$ firms serve the market bounds the long-run number of firms in equilibrium.

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

Intuitively, in a symmetric equilibrium, firms cannot survive for sure with $n > \check{n}$ firms because this would give them negative payoffs (which they could avoid by exiting instead). To see this, note that if all firms continue for sure, each would incur a positive continuation cost and earn a zero surplus one or more times, and then collect a zero post-entry value in the first future period in which firms leave with positive probability. So, $a_S(n, \cdot, \cdot) < 1$ and $v_E(n, \cdot, \cdot) = 0$. Because no firm would be willing to pay a positive sunk cost to enter a survival subgame with zero expected payoff, $a_E(n, \cdot, \cdot) = 0$.

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 $\{1, 2, \ldots, \check{n}\} \times \mathcal{C}$. Extending an equilibrium strategy over this restricted state space to the full state space is straightforward .

The next step in the equilibrium analysis uses the assumption that *per-period* surplus weakly decreases with the number of competitors to show that the same monotonicity applies to the post survival value functions.

Lemma 2 (Monotone equilibrium payoffs) *In a symmetric Markov-perfect equilibrium,* $\forall c \in \mathcal{C}$, $v_S(n', c)$ *weakly decreases with* n'.

Lemma 2 says that no endogenous complementarity between firms arises in equilibrium. Although this is intuitive, it is not a trivial result. Indeed, Abbring et al. (2010) give a counterexample to the analogous proposition in a model with heterogeneous productivity types.[2]

To appreciate the implications of Lemma 2, consider a one-shot simultaneous-moves survival game played by n_E active firms. In it, each of the n' survivors earns $-\kappa \exp(w) + v_S(n', c)$, where v_S is the post-survival value in a symmetric Markov-perfect equilibrium of our dynamic game, and each exiting firm earns zero. The Nash equilibria of this game are intimately connected to the Markov-perfect equilibria of our model. In particular, (4) implies that a survival rule $a_S(n_E, c, w)$ from a symmetric Markov-perfect equilibrium forms a symmetric Nash equilibrium of the one-shot game, and vice versa.

This one-shot game has many equilibria in the trivial case that $v_S(n_E, c) = \cdots = v_S(1, c) = \kappa \exp(w)$. In this case, our restriction to equilibria that default to inactivity picks $a_S(n_E, c, w) = 0$. In the more interesting case where $v_S(n', c) \neq \kappa \exp(w)$ for at least one $n' \in \{1, \ldots, n_E\}$, Lemma 2 guarantees that the one-shot game has a *unique* symmetric Nash equilibrium. To show this, we distinguish three subcases.

- First, suppose that $v_S(1, c) \leq \kappa \exp(w)$. Lemma 2 implies that $v_S(n', c) \leq \kappa \exp(w)$ for all $n' \in \{1, \ldots, n_E\}$. Therefore, exiting for sure (setting $a_S(n_E, c, w) = 0$) is a weakly dominant strategy and forms one symmetric equilibrium. Furthermore, since $v_S(n', c) \neq \kappa \exp(w)$ for at least one $n' \in \{1, \ldots, n_E\}$, we know that $v_S(n_E, c) < \kappa \exp(w)$. Therefore, exiting for sure is also the unique best response to any positive symmetric continuation probability.

- Next, suppose that $v_S(n_E, c) \geq \kappa \exp(w)$. Lemma 2 implies that $v_S(n', c) \geq \kappa \exp(w)$ for $n' = 1, \ldots, n_E$. Therefore, continuing for sure (setting $a_S(n_E, c, w) = 1$) is a weakly dominant strategy and forms one symmetric equilibrium. Since $v_S(n', c) \neq \kappa \exp(w)$ for at least one $n' \in \{1, \ldots, n_E\}$, we also have $v_S(1, c) > \kappa \exp(w)$. Therefore, continuing for sure is also the unique best response to any continuation probability less than one.

[2] However, Abbring et al. (2010) show that a version of Lemma 2 also holds in a model with heterogeneous productivity if $\check{n} = 2$. See Section 7.

- For the last subcase, suppose that $v_S(1,c) > \kappa\exp(w)$ and $v_S(n_E,c) < \kappa\exp(w)$. No symmetric pure strategy equilibrium exists, because the best response to all other firms continuing for sure is to exit for sure, and vice versa. In a mixed strategy equilibrium, firms must be indifferent between continuation and exit. By the intermediate value theorem, there is some a_S that solves this indifference condition,

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

which establishes existence of a mixed strategy equilibrium. Lemma 2 and the subcase's conditions together guarantee that the left hand side of (5) strictly decreases in a_S. Therefore, there is only one symmetric mixed strategy equilibrium.

For future reference, we state the equilibrium uniqueness result for this nontrivial case with

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 $-\kappa\exp(w) + v_S(c,n')$ with $-\kappa\exp(w) + v_S(n',c) \neq 0$ for at least one $n' \in \{1,\ldots,n_E\}$, 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 it is positive and gives them zero otherwise (whenever 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*

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

Note that, Corollary 2 in combination with Lemma 2 implies that $v_E(n_E,\cdot,\cdot)$ also weakly decreases with n_E.

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. Begin with calculating $v_E(\check{n},\cdot,\cdot)$ and $v_S(\check{n},\cdot)$. From Lemma 1, there will be no entry in a period starting with $\check{n}$ firms, so

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

Insert this into (6) to get

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

The right-hand side defines a contraction mapping on the complete space of bounded functions on $\mathcal{C}\times\mathbb{R}$, with a unique fixed point $v_E(\check{n},\cdot,\cdot)$. Although we are constructing a *candidate* equilibrium, the fixed point's uniqueness implies that this is the only possible equilibrium post-entry value. Inserting this fixed point into (7) immediately yields $v_S(\check{n},\cdot)$. Again, this is the only possible candidate value. The unique entry rule that is consistent with these payoffs and individual optimality that also defaults to inactivity is

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

Here, $\mathbb{1}[x] = 1$ if x is true and 0 otherwise.

With $v_E(\check{n},\cdot,\cdot)$ and $a_E(\check{n},\cdot,\cdot)$ calculated, the construction of the remaining candidate value functions and entry strategies proceeds recursively. For given n, suppose that $v_E(n',\cdot,\cdot)$ and $a_E(n',\cdot,\cdot)$ for $n' = n+1, n+2, ..., \check{n}$ are in hand, and define

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

This is the number of firms that will be active in a period that starts with n firms after all that period's potential entrants have followed the candidate entry strategy. The second equality follows from the monotonicity of v_E (and hence a_E). Corollary

2 and (2) together imply

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

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

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

Since Lemma 2 implies that further entry cannot make an incumbent better off, a necessary condition for (11) is that the firm would enter in the absence of further entry, $v_E(n,c,w) > \varphi\exp(w)$. On the other hand, because later entrants pay the same entry costs, further entry will never take post-survival values below $\varphi\exp(w)$, so $v_E(n,c,w) > \varphi\exp(w)$ is also sufficient for (11). Therefore,

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

is the only possible equilibrium entry rule consistent with $v_E(n,c,w)$.[3]

When this recursion is complete, we have the unique continuation values and entry strategies that are consistent with an equilibrium. To determine a candidate survival rule a_S, we set $a_S(n_E,c,w) = 0$ for all (n_E,c,w) such that $v_S(n_E,c) = \cdots = v_S(1,c) = \kappa\exp(w)$ and find an equilibrium to Corollary 1's one-shot survival game for all other (n_E,c,w). *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.*

[3]In the more general model of the Appendix, we assume that the sunk costs of entry weakly increase with the number of firms already committed to production in the next period. The logic of this paragraph applies straightforwardly to that setting.

4 Empirical Implementation

The previous section shows that there exists a unique symmetric Markov-perfect equilibrium for given primitives π, κ, φ, ρ, G_C, and G_W. Given (N_1, C_1), this equilibrium induces a distribution for the process $\{N_t, C_t\}$. This section studies how data on this process from a panel of markets can be used to estimate the model's primitives.

4.1 Sampling

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

We assume that $(\{N_{r,t}, C_{r,t}; t = 1, \ldots, \check{t}\}, X_r)$ is distributed independently across markets.[4] The initial conditions $(N_{r,1}, C_{r,1}, X_r)$ are drawn from a distribution that we leave unspecified. Conditional on $(N_{r,1}, C_{r,1}, X_r)$, industry dynamics $\{N_{r,t}, C_{r,t}; t = 2, \ldots, \check{t}\}$ follow the transition rules implied by Section 3's unique equilibrium, with primitives $\pi_r(\cdot,\cdot) = \pi(\cdot,\cdot\,|X_r, \theta_P)$, $\kappa_r = \kappa(X_r, \theta_P)$, $\varphi_r = \varphi(X_r, \theta_P)$, and $\rho_r = \rho(X_r, \theta_P)$ for some finite vector θ_P; $G_{C,r}(\cdot \mid \cdot) = G_C(\cdot \mid \cdot\,; X_r, \theta_C)$ for some finite vector θ_C; and $G_{W,r}(\cdot) = G_W(\cdot\,; X_r, \theta_W)$ for some finite vector θ_W.[5]

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 θ for data on market dynamics $\{N_{r,t}, C_{r,t}; t = 2, \ldots, \check{t}; r = 1, \ldots, \check{r}\}$ given the initial conditions $(N_{r,1}, C_{r,1}, X_r; r = 1, \ldots, \check{r})$.[6] Using the model's Markov

[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]These assumptions rule out persistent unobserved heterogeneity in the primitives across markets. Relaxing this and appropriately extending our NFXP procedure is straightforward in principle, but it does require us to provide a model-based solution to the "initial conditions problem" that $(N_{r,1}, C_{r,1}, X_r)$ is not independent of the persistent unobservables.

[6]We neither specify nor estimate the initial conditions' distribution, because we want to be agnostic about their relation to the dynamic model. We could instead assume that the initial conditions are drawn from the model's ergodic distribution. This would allow us to develop a more

structure and conditional independence, this likelihood can be written as $\mathcal{L}(\theta) = \mathcal{L}_C(\theta_C) \cdot \mathcal{L}_N(\theta)$, with

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

the marginal likelihood of θ_C for the demand state dynamics; and

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

the conditional likelihood of θ for the evolution of the market structures. Here, $g_C(\cdot \mid \cdot; X_r, \theta_C)$ is the density of $G_{C,r}$ and $p(n'|n, c; X_r, \theta) = \Pr(N_{r,t+1} = n'|N_{r,t} = n, C_{r,t} = c; X_r, \theta)$ is the equilibrium probability that market r with n firms and in demand state c has n' firms next period.

Note that $\mathcal{L}_C(\theta_C)$ can be computed directly from the demand data, without ever solving the model. To calculate $\mathcal{L}_N(\theta)$ we need to compute the equilibrium transition probabilities $p(\cdot \mid \cdot\,; X_r, \theta)$ for each distinct value of X_r in the sample and substitute these into (13). To this end, for each value of X_r, we first compute the equilibrium post-survival values $v_{S,r}$ corresponding to the primitives implied by X_r and θ. From these, we obtain cost-shock thresholds for entry and *sure* survival, defined by

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

and

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

For $n' > n$, $p\left(n'|n, c; X_r, \theta\right)$ can easily be calculated as the probability that $W_{r,t}$ falls into $[\overline{w}_{E,r}(n'+1, c), \overline{w}_{E,r}(n', c))$. For $n' \leq n$, the computations are complicated by equilibrium mixing of survival decisions. For example, the probability that the number of firms remains unchanged at n sums the probability that $W_{r,t}$ falls into $[\overline{w}_{E,r}(n+1, c), \overline{w}_{S,r}(n, c)]$ with the probability that it instead equals some $w \in (\overline{w}_{S,r}(n, c), \overline{w}_{S,r}(1, c))$ and that all n firms survive when they mix with

efficient estimator, at the price of robustness. Moreover, it would allow us to deal with the initial conditions problems mentioned in Footnote 5.

probability $a_S(n, c, w)$. Similar complications arise for $n' \in \{1, \ldots, n-1\}$, which can only occur if firms follow a non-trivial mixed strategy, and for $n' = 0$, which can occur if either the incumbents all choose certain exit or they are mixing nontrivially and all exit by chance. Accounting for the influence of mixed strategies on $p(n'|n, c; X_r, \theta)$ in these last three cases is tedious but straightforward.

4.3 Identification

Before proceeding to use the likelihood function for the model's estimation, we first analyze the extent to which we could determine θ uniquely if we observe the population $(\{N_t, C_t; t = 1, \ldots, \check{t}\}, X)$ underlying our data. Specifically, suppose that we know the distribution of (N', C') conditional on $(N, C, X) = (n, c, x)$ for all $n \in \mathbb{N}_0 \equiv \{0\} \cup \mathbb{N}$, $c \in \mathcal{C}$, and a specific value x of the market characteristics.[7] Throughout the remainder of this section, we keep conditioning on $X = x$ implicit, so the results demonstrate identification of the model's primitives as nonparametric functions of the market characteristics.

To begin, note that the population information directly identifies G_C.[8] The remaining primitives of interest are the model's fixed cost, κ, sunk cost φ, surplus function π, and the distribution G_W of the econometric error. Our identification argument for these parameters follows that of Magnac and Thesmar (2002), who retrieve value functions by applying the inverse cumulative distribution function of the econometric error to observed choice probabilities. Since this strategy requires knowledge of G_W, we assume that this belongs to the parametric family

$$G_W(w) = \Phi\left(\frac{w + \omega^2/2}{\omega}\right), \tag{16}$$

where Φ refers to the cumulative distribution function of a standard normally distributed random variable, with density ϕ. That is, $\exp(W)$ has a log-normal distribution with unit mean and scale parameter ω. Since observations of the number

[7]For x fixed, the hypothetical data scenario that is informative about this distribution involves the number of transitions from (N, C) to (N', C') approaching infinity. Whether such transitions are coming from the same market or many different markets all with characteristics x plays no role in the identification argument.

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

of producers give us no information on the level of profits, we also normalize the mean per-period fixed cost κ to one.[9]

Begin by retrieving $\overline{w}_S(1,c)$, up to the unknown scale and shift in G_W, from the probability of a monopolist surviving:

$$\frac{\overline{w}_S(1,c)+\omega^2/2}{\omega} = \Phi^{-1}\left(\Pr[N' \geq 1 | N = 1, C = c]\right).$$

Similarly, we can recover $\overline{w}_E(n,c)$ from the probability of n firms entering a previously empty market:

$$\frac{\overline{w}_E(n,c)+\omega^2/2}{\omega} = \Phi^{-1}\left(\Pr[N' \geq n | N = 0, C = c]\right).$$

These and the definitions of $\overline{w}_S(1,c)$ and $\overline{w}_E(1,c)$ in (14) and (15) can be used to identify the sunk cost of entry up to the scale parameter ω:

$$\begin{aligned}\frac{\log(\varphi+1)}{\omega} &= \frac{\overline{w}_S(1,c)-\overline{w}_E(1,c)}{\omega} \\ &= \Phi^{-1}\left(\Pr[N' \geq 1 | N = 1, C = c]\right) - \Phi^{-1}\left(\Pr[N' \geq 1 | N = 0, C = c]\right).\end{aligned}$$

In turn, this allows us to retrieve

$$\frac{\overline{w}_S(n,c)+\omega^2/2}{\omega} = \frac{\overline{w}_E(n,c)+\omega^2/2}{\omega} + \frac{\log(\varphi+1)}{\omega}.$$

The argument's next step identifies the scale parameter ω. In a simple probit model, the analogous parameter is *not* identified unless one places an *a priori* restriction on the regressors' coefficients. For the present model, the mixing sometimes employed by exiting oligopolists provides information on the scale of payoffs relative to the econometric error. This information identifies ω without the use of auxiliary restrictions on payoffs.

To proceed, suppose that, for some $c^\star \in \mathcal{C}$ and $n^\star \in \{2,\ldots,\check{n}\}$,

$$\overline{w}_S(1,c^\star) = \cdots = \overline{w}_S(n^\star-1,c^\star) > \overline{w}_S(n^\star,c^\star).$$

[9]This normalization implies that we do not identify cross-market differences in the *scale* of producers' surplus, fixed costs, and sunk costs. Rather, we identify producers' surplus and sunk costs relative to fixed costs for each market.

By (15), this is equivalent to requiring that

$$v_S(1, c^\star) = \cdots = v_S(n^\star - 1, c^\star) > v_S(n^\star, c^\star)$$

for some $c^\star$ and $n^\star$. This is a very weak condition, particularly in light of Lemma 2's result that $v_S(n, \cdot)$ *always* weakly decreases in n. Moreover, it can be verified in data, because we have already determined the sure survival thresholds up to a common scale and location shift.

Now, consider the probability of $n^\star$ incumbents simultaneously exiting:

$$\begin{aligned}
&\Pr[N' = 0|N = n^\star, C = c^\star] \\
&\quad = \Pr[W \geq \overline{w}_S(1, c^\star)] + \int_{\overline{w}_S(n^\star, c^\star)}^{\overline{w}_S(1, c^\star)} [1 - a_S(n^\star, c^\star, w)]^{n^\star} g_W(w) dw \\
&\quad = \Pr[N' = 0|N = 1, C = c^\star] + \int_{\overline{w}_S(n^\star, c^\star)}^{\overline{w}_S(1, c^\star)} [1 - a_S(n^\star, c^\star, w)]^{n^\star} g_W(w) dw. \quad (17)
\end{aligned}$$

Because the two transition probabilities in (17) are known, so is the integral on its right hand side. We will now show that this integral can be written as a known monotone function of ω, so that it identifies ω. Using $v_S(1, c^\star) = \cdots = v_S(n^\star - 1, c^\star)$, we can explicitly solve (5) for the mixing probability $a_S(n^\star, c^\star, w)$:

$$a_S(n^\star, c^\star, w) = \left(\frac{v_S(1, c^\star) - \exp(w)}{v_S(1, c^\star) - v_S(n^\star, c^\star)} \right)^{\frac{1}{n^\star - 1}}.$$

Rewrite the integral on the right hand side of (17) by substituting this expression for $a_S(n^\star, c^\star, w)$, use (15) to replace post-survival values with sure survival thresholds, and change the variable of integration from w to $\varepsilon = (w + \omega^2/2)/\omega$. This gives

$$\int_{k_{n^\star}}^{k_1} \left[1 - \left(\frac{\exp(\omega k_1) - \exp(\omega \varepsilon)}{\exp(\omega k_1) - \exp(\omega k_{n^\star})} \right)^{\frac{1}{n^\star - 1}} \right]^{n^\star} \phi(\varepsilon) d\varepsilon, \quad (18)$$

with

$$k_1 \equiv \frac{\bar{w}_S(1, c^\star) + \omega^2/2}{\omega} \text{ and } k_{n^\star} \equiv \frac{\bar{w}_S(n^\star, c^\star) + \omega^2/2}{\omega}.$$

Because k_1 and $k_{n^\star}$ are known,

$$\frac{\exp(\omega k_1) - \exp(\omega \varepsilon)}{\exp(\omega k_1) - \exp(\omega k_{n^\star})} = \frac{1 - \exp(-\omega (k_1 - \varepsilon))}{1 - \exp(-\omega (k_1 - k_{n^\star}))} \tag{19}$$

is a known function of ω. Moreover, it is straightforward to verify that it is strictly increasing in ω for $\varepsilon \in (k_{n^\star}, k_1)$. Hence, the integrand in (18) is a known, strictly decreasing function of ω. Because the domain of integration of the integral in (18) is also known, this establishes that the integral itself is a known strictly decreasing function of ω, so that ω can be uniquely determined from the integral's known value.

With ω identified, we immediately recover φ, $\bar{w}_S = \log v_S$, $\bar{w}_E$, and v_E (and therewith a_S and a_E). The discount factor and per-period surplus function remain to be identified. For the discount factor, we can follow one of two approaches. First, we can assume that auxiliary information like the average borrowing rate for small businesses identifies ρ. Alternatively, we can follow Magnac and Thesmar (2002) and suppose that there exists an observable component of C that impacts next period's expected continuation values but *not* next period's expected surplus. More precisely, suppose that there exist two values $c_1 \neq c_2$ such that

$$\mathbb{E}_{a_E}[v_E(N'_E, C', W')|N' = n', C = c_1] \neq \mathbb{E}_{a_E}[v_E(N'_E, C', W')|N' = n', C = c_2]$$

but $\mathbb{E}[\pi(n', C')|C = c_1] = \mathbb{E}[\pi(n', C')|C = c_2]$. The former condition can be verified from data because v_E, a_E, G_C and G_W are identified, but the latter is an *a priori* exclusion restriction. Under this assumption, we can use (2) to show that

$$\rho = \frac{v_S(n', c_1) - v_S(n', c_2)}{\mathbb{E}_{a_E}[v_E(N'_E, C', W')|N' = n', C = c_1] - \mathbb{E}_{a_E}[v_E(N'_E, C', W')|N' = n', C = c_2]}.$$

Of course, which of these approaches is most appropriate depends on the application at hand. In either case, given ρ we can recover $\mathbb{E}[\pi(n', C')|C = c]$ from (2).

We summarize these results in a theorem.

Theorem 2 *Suppose that ρ and κ are known and that G_W is specified up to scale as in (16). Furthermore, suppose that, for some $c^\star \in \mathcal{C}$ and $n^\star \in \{2, \dots, \check{n}\}$,*

$$\Pr[N' = 0|N = 1, C = c^\star] = \cdots = \Pr[N' = 0|N = n^\star - 1, C = c^\star] < \Pr[N' = 0|N = n^\star, C = c^\star].$$

Then, the distribution of (N', C') given $(N, C) = (n, c)$ for $n \in \mathbb{N}_0$ and $c \in \mathcal{C}$ uniquely *determines* $\mathbb{E}[\pi(\cdot, C')|C = c]$ *for* $c \in \mathcal{C}$*,* φ*,* G_C*, and* G_W*.*

To emphasize that it can be verified in data, we have rewritten the condition that there exist $c^\star$ and $n^\star$ such that $\overline{w}_S(1, c^\star) = \cdots = \overline{w}_S(n^\star - 1, c^\star) > \overline{w}_S(n^\star, c^\star)$ in terms of known probabilities, using (17) and the fact that $0 < a_S(n^\star, c^\star, w) < 1$ for $w \in (\overline{w}_S(n^\star, c^\star), \overline{w}_S(1, c^\star))$.

Theorem 2 only establishes identification of the expected surplus $\mathbb{E}[\pi(\cdot, C')|C = c]$, not of the surplus function π itself. This makes sense, because entry and exit decisions are taken after a period's surplus is earned and before next period's demand state C' is realized, so that observed market transitions only depend on π through the expected surplus. Nevertheless, in some applications, for example those involving counterfactual specifications of G_C, it may be useful to separately identify π. In these cases, π can be uniquely determined from the expected surplus provided that G_C satisfies a completeness condition of the type now routinely used in nonparametric identification analysis (see e.g. Newey and Powell, 2003).

We take three 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 and Miller (1993), we could instead estimate the equilibrium value functions directly from observed transition probabilities and from these deduce the underlying primitives. Given the absence of idiosyncratic shocks from our model, this deduction would substantially differ from that pioneered by Bajari et al. (2007). We choose to estimate our primitives using maximum likelihood because it achieves statistical efficiency yet it requires no parametric assumptions beyond those used by a Hotz and Miller style estimator. Third, the use of nontrivial mixed strategies can identify the scale of the econometric error without imposing restrictions on players' payoffs. Identifying the analogous parameter in static discrete choice models always requires restricting the non-stochastic portion of payoffs in some way. Although we can impose similar restrictions on $\pi(n, c)$ (for example by requiring linearity in c), we have found that these need not create straightforward restrictions on the equilibrium continuation values useful for extending static identification arguments to this dynamic setting. Nevertheless, we do not doubt the practical usefulness of restrictions on payoffs for estimation with a finite sample.

4.4 Estimation

We have created C++ and Matlab code for computing the full information maximum likelihood estimator of θ. As in Rust (1987), computation proceeds in three steps:

1. Estimate θ_C with $\tilde{\theta}_C \equiv \arg\max_{\theta_C} \mathcal{L}_C(\theta_C)$;
2. estimate (θ_P, θ_W) with $(\tilde{\theta}_P, \tilde{\theta}_W) \equiv \arg\max_{(\theta_P,\theta_W)} \mathcal{L}_N(\theta_P, \tilde{\theta}_C, \theta_W)$; and
3. estimate θ by maximizing the full likelihood function $\hat{\theta} \equiv \arg\max_{\theta} \mathcal{L}(\theta)$, using $\tilde{\theta} \equiv (\tilde{\theta}_P, \tilde{\theta}_C, \tilde{\theta}_W)$ as a starting value.

Note that the partial likelihood estimator $\tilde{\theta}$ computed in the first two steps is consistent, but not efficient. The third step's estimator $\hat{\theta}$ is asymptotically efficient. To compute estimated standard errors, we use the outer-product-of-the-gradient estimator of the (full) information matrix. In particular, we assume that $\check{r}$ is large and $\check{t}$ is small and use the average over markets of the outer products of the market-specific gradients, 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 this 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 that can be used as a sandbox for experimentation and teaching.

5 Monte Carlo Experiments

In this section we investigate the statistical properties and computational performance of our estimation procedure with Monte Carlo experiments. For these, we set the maximum number of firms entering any market to $\check{n} = 5$, let the cost shocks be log-normally distributed with unit mean and scale parameter ω, normalize the mean per period fixed costs κ to one, fix the discount factor ρ at $\frac{1}{1.05}$, and interpret C_t as the number of consumers in the market. The statistical process governing C_t has support on 200 grid points that are equally spaced on the logarithmic scale with distance d: $c_{[1]}$, $c_{[2]} = c_{[1]} \exp(d), \ldots, c_{[200]} = c_{[1]} \exp(200d)$. So that the growth of C_t is approximately normally distributed with mean μ and variance σ^2, we follow

Tauchen (1986) and specify the probability of transitioning to $c_{[i]}$ from $c_{[j]}$ for any $i = 2, ..., 199$ and $j = 1, ..., 200$ with

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

The probabilities of transitioning to the grid's end points equal

$$\Pr[C' = c_{[1]}|C = c_{[j]}] = \Phi\left(\frac{\log c_{[1]} - \log c_{[j]} + \frac{d}{2} - \mu}{\sigma}\right)$$

and

$$\Pr[C' = c_{[200]}|C = c_{[j]}] = 1 - \Phi\left(\frac{\log c_{[200]} - \log c_{[j]} - \frac{d}{2} - \mu}{\sigma}\right)$$

respectively. We set $\mu = 0$ and $\sigma = 0.02$.

Each Monte Carlo experiment consists of 1,000 synthetic samples. We use four different sample sizes, each of them with ten time periods and between 100 and 1,000 ex ante identical markets. We compute the equilibrium and generate each sample market by simulating the evolution of (N, C), beginning with a draw from the model's ergodic distribution. We then use each sample to estimate the model's parameters with the three step procedure presented in Section 4. (Since this specification excludes variation in market characteristics, a single equilibrium calculation can support the likelihood function calculations for all of a sample's observations.) The starting parameter vector used for the likelihood function's maximization equals a vector of ones multiplied by one random variable uniformly distributed on $[1, 10]$. Dubé et al. (2012) caution that a nested fixed point algorithm can falsely converge when the tolerance criterion for the inner loop (which calculates the equilibrium) is set too loosely relative to that of the outer loop (which maximizes the likelihood function). We fix the convergence tolerance for the value function iteration at a value that is multiple orders of magnitude smaller than that for the likelihood maximization to avoid this potential pitfall.[10]

We first simulate data from a model where the surplus function is parameterized as $\pi(c, n) = (c/n)k$ for $n \leq \check{n}$ and some fixed $k > 0$. This means that per consumer surplus is constant in the number of active firms n for $n \leq \check{n}$. We set the true values of

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

Table 1: Monte Carlo Results with Constant Surplus Per Consumer

	$\check{r}$ = 100	$\check{r}$ = 250	$\check{r}$ = 500	$\check{r}$ = 1,000
Averages of Estimates				
k	1.501	1.500	1.501	1.499
φ	10.255	10.113	10.075	10.029
ω	0.995	0.999	0.999	1.000
$\mu \times 10^2$	-0.002	-0.000	0.001	-0.000
$\sigma \times 10^2$	2.000	1.999	1.999	1.998
Averages of Estimated Standard Errors				
k	0.049	0.031	0.022	0.015
φ	2.924	1.790	1.254	0.879
ω	0.070	0.044	0.031	0.022
$\mu \times 10^2$	0.068	0.043	0.030	0.021
$\sigma \times 10^2$	0.049	0.031	0.022	0.015
Monte Carlo Estimates of 95% Confidence Interval Coverage				
k	0.954	0.937	0.945	0.958
φ	0.924	0.944	0.948	0.956
ω	0.950	0.936	0.946	0.946
μ	0.948	0.948	0.951	0.963
σ	0.943	0.933	0.939	0.943

Note: Results of a Monte Carlo experiment using the three step NFXP estimator to estimate the model with one profit parameter (k) using 1,000 synthetic samples. The true value of k equals 1.5 and the true value of φ equals 10. The true value of the standard deviation of the log costs (ω) equals 1. Demand is discretized into 200 states. The demand process is governed by the drift parameter μ, which is set to zero, and the innovation standard deviation σ, which equals 0.02. The bottom-most panel displays the fraction of samples for which the estimated 95 percent confidence interval contained the parameter's true value.

k, φ, and ω, to 1.5, 10, and 1 respectively. Table 1 reports the corresponding Monte Carlo experiments' results. Its first panel gives the averages of the 1,000 estimates for each parameter, and it shows that the NFXP estimator is essentially without bias, even for the sample with only 100 markets. The second panel reports the averages of the estimated standard errors. For the sample with 100 markets, the average estimated standard error for the estimate of the sunk cost is 2.924. Therefore, we would expect a 95 percent confidence interval to approximately correspond to $(4, 16)$. This is possibly too wide for empirical usefulness, but the other estimates' standard errors are relatively small. As expected, increasing the sample size decreases the standard errors approximately at the rate $\sqrt{\check{r}}$. So for $\check{r} = 500$ the standard error on $\hat{\varphi}$ is only 1.254. The table's final panel reports the Monte Carlo estimates of 95

Table 2: Monte Carlo Results with Decreasing Surplus Per Consumer

	ř = 100	**ř = 250**	**ř = 500**	**ř = 1,000**
Averages of Estimates				
$k(1)$	1.804	1.800	1.801	1.800
$k(2)$	1.397	1.399	1.400	1.400
$k(3)$	1.198	1.200	1.200	1.198
$k(4)$	1.000	1.000	1.000	0.999
$k(5)$	0.897	0.898	0.899	0.898
φ	9.939	10.014	9.951	9.977
ω	0.983	0.994	0.996	0.999
$\mu \times 10^2$	-0.004	-0.001	-0.002	-0.001
$\sigma \times 10^2$	1.995	1.997	1.998	1.998
Averages of Estimated Standard Errors				
$k(1)$	0.087	0.053	0.037	0.026
$k(2)$	0.094	0.059	0.041	0.029
$k(3)$	0.079	0.049	0.035	0.024
$k(4)$	0.078	0.049	0.034	0.024
$k(5)$	0.100	0.061	0.043	0.030
φ	3.401	2.093	1.451	1.024
ω	0.085	0.053	0.037	0.026
$\mu \times 10^2$	0.068	0.043	0.030	0.021
$\sigma \times 10^2$	0.049	0.031	0.022	0.015
Monte Carlo Estimates of 95% Confidence Interval Coverage				
$k(1)$	0.948	0.943	0.946	0.954
$k(2)$	0.938	0.952	0.957	0.960
$k(3)$	0.941	0.943	0.956	0.938
$k(4)$	0.958	0.960	0.957	0.950
$k(5)$	0.946	0.942	0.946	0.937
φ	0.882	0.918	0.934	0.941
ω	0.926	0.940	0.953	0.944
μ	0.946	0.941	0.945	0.953
σ	0.949	0.957	0.949	0.953

Note: Results of a Monte Carlo experiment using the three step NFXP estimator to estimate the model with five profit parameters $(k(1), k(2), ..., k(5))$ using 1,000 synthetic samples. The true value of $(k(1), k(2), ..., k(5))$ equals $(1.8, 1.4, 1.2, 1.0, 0.9)$ and the true value of φ equals 10. The true value of the standard deviation of the log costs (ω) equals 1. Demand is discretized into 200 states. The demand process is governed by the drift parameter μ, which is set to zero, and the innovation standard deviation σ, which equals 0.02. The bottom-most panel displays the fraction of samples for which the estimated 95 percent confidence interval contained the parameter's true value.

percent confidence intervals' coverage probabilities. With the exception of that for φ with $\check{r} = 100$, these are all within 1.5 probability points of their common nominal value. Apparently, the estimated standard errors provide accurate inference.

For our second set of simulations we parameterize the flow surplus function as $\pi(c,n) = (c/n)k(n)$, 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.[11] Table 2 reports the results of the corresponding Monte Carlo experiments. Again, all parameter estimates are essentially without bias, the estimated standard errors are small enough to be empirically useful, and the 95 percent confidence intervals have coverage probabilities close to their common nominal value. To check whether the estimator is able to distinguish a model with a decreasing per consumer surplus from a model with a constant surplus, we compute a likelihood ratio test for each sample. We can reject the null hypothesis $k(1) = \ldots = k(5)$ at the 95 percent confidence level in all of our synthetic samples regardless of the sample size. Overall, we conclude that the NFXP procedure has the potential to be empirically useful. Verifying that potential with observations not created by the model is the subject of the next section.

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 is the case. Even in the largest of our synthetic samples, the average computation of the maximum likelihood estimator takes about two minutes using the C++ code.

Su and Judd's (2012) results suggest that we might be able to improve on the 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. We implemented the MPEC estimator of our model in C++ using analytical gradients of both the objective function and the constraints. The MPEC estimator always

[11]These values generate a realistic distribution of firms per market. No firm is active in about 5 percent of the markets, a monopolist serves about 30 percent of the markets, and five firms serve about 5 percent of the markets. In contrast, the previous specification with constant per consumer surplus generates a distribution with an additional local mode at five firms (the value of $\check{n}$), because competition does not get "tougher" when the number of firms increases.

Table 3: Computational Performance

	$\check{r}$ = 100	$\check{r}$ = 250	$\check{r}$ = 500	$\check{r}$ = 1,000
one entry cost parameter, one profit parameter				
time per run (in seconds)	30.46	38.27	50.17	73.97
one entry cost parameter, five profit parameters				
time per run (in seconds)	65.79	74.90	90.50	123.73

Note: Average computational performance of the NFXP estimators in the synthetic samples. The estimator is implemented in C++ and Knitro and runs as a single thread on an Intel Core i7-860 with 2.8GHz.

yielded the same estimates as our NFXP procedure, but we found it to be more than ten times slower than the NFXP implementation.[12] MPEC's relatively poor performance arises from the computation of the objective function's gradients with respect to the nuisance parameters, which requires repeatedly retrieving information from very large and relatively dense matrices. These computational challenges might not be insurmountable, but our NFXP estimator seems to balance the costs of programmer time and execution time well.

6 Empirical Illustration

The Monte Carlo results suggest that we can use observations from a few hundred markets over ten time periods to estimate the model's parameters accurately enough for differentiating between economically distinct hypotheses about the magnitude of sunk costs and the effects of additional competition on producers' surplus. In this section we take the model beyond "data" of its own making and estimate its parameters with observations from the Motion Picture Theaters industry. Although we intend this application to be illustrative, we have tried to make our results more useful by including in X_r a measure of the geographic diversity of a market's consumers. Davis (2006) finds that theater location substantially influences

[12] Su and Judd (2012) emphasize the usefulness of passing "sparsity patterns" to the optimizer, which indicate which derivatives of the constraints with respect to the nuisance parameters are identically zero. We followed their advise, and we also initialized the nuisance parameters at their true values. For the parameters of interest, we used the same starting values as in the NFXP estimation.

consumers' decisions about whether and where to attend film screenings. Indeed, one's probability of attending a given theater declines substantially when the travel distance moves from between zero and five miles to between five and ten miles.[13] Davis (2002) finds that the concomitant low cross-price elasticities from such spatial preferences impact firms' pricing behavior. Using observations from a New Haven area theater that experimented with a temporary price cut, Davis (2002) found that rivals five to seven miles away responded with lower prices but those ten to twelve miles away did not. This spatial differentiation of Motion Picture Theaters makes market definitions based on readily available geographic data plausibly applicable to this industry. Since we place a measure of how far apart consumers are from each other into X_r, our estimates quantify how the spatial structure of demand impacts the level of profits and the toughness of competition.

6.1 The Data

This analysis equates a market with a Micropolitan Statistical Area (μSA) as defined by the Office of Management and Budget. Each one is based around an urban core of at least 10,000 but less than 50,000 inhabitants.[14] We dropped the μSA "The Villages, FL", because its population growth far exceeds that of any other μSA. The remaining 573 μSAs account for about ten percent of the United States population. We measured the diversity of the geographic preferences of each μSA's residents using the locations and populations of its constituent year 2000 Census tracts. For this, we supposed that each census tract is a circle with an area equal to that of the tract itself, that population is uniformly distributed over the area enclosed by the circle, that all travel within a tract must pass through its center, and that travel between tracts follows straightline roads that connect their centers.[15] We then measured *geographic preference diversity* with the average distance between two randomly-chosen residents of the μSA. Likewise, we can measure the average distance between two randomly chosen individuals from two distinct μSAs. By

[13]See the logit model estimates reported in Table 5 of Davis (2006).

[14]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 μSAs.

[15]For these calculations, we used the tract population and geographic location information from the National Census Tracts Gazetteer File for the 2000 Decennial Census. See `http://www.census.gov/geo/maps-data/data/gazetteer2000.html` for its documentation. We used each tract's latitude and longitude in this file as its center.

Table 4: Summary Statistics for μSAs

	Quantile				
	10	25	50	75	90
Population	23.51	32.57	42.67	62.32	87.71
Median Household Income	32.77	37.10	41.29	45.83	51.10
Geographic Preference Diversity	9.24	11.17	13.37	16.81	21.23
Geographic Market Isolation	23.94	28.77	37.61	51.83	72.70

Note: All variables are measured as of 2000 for the 573 μSAs in our sample. Population is expressed in thousands of people, Median household income is expressed in thousands of dollars per year, and the remaining variables are expressed in miles. Please see the text for further details.

construction, μSAs are geographically isolated from larger Metropolitan Statistical Areas, so we measure a given μSA's *geographic market isolation* as the shortest such distance to another μSA.

For the 573 μSAs, Table 4 displays the five standard quantiles for population, median household income, geographic preference diversity, and geographic market isolation. Population varies by about a factor of four from the 10th to the 90th percentiles. For the United States as a whole, median household income equalled $\$42,148$ in 2000. This is very close to the median value across the μSA's, $\$41,288$. About 80 percent of the μSAs have median household incomes within $\$10,000$ of this central tendency. The median geographic preference diversity is 13.37 miles. Perhaps unsurprisingly, this variable is highly skewed to the right. The 10th percentile is 9.24 miles, while the 90th percentile is 21.23 miles. Given the evidence from Davis (2002, 2006) regarding urban consumers' transportation costs for attending movies, it is plausible that the least geographically diverse μSAs in our sample might form a single geographic market. On the other hand, those with the most geographic preference diversity might actually be collections of two or more "markets" with relatively low elasticities of substitution across them. In any case, the measures of geographic isolation indicate that the elasticities of substitution across locations within a μSA should be much larger than those across μSAs. Its median value across μSAs equals 37.61 miles. Indeed, there are only eight μSAs where this distance is less than twenty miles. We conclude that the μSAs are isolated enough from each other so that substitution between them can be ignored.

The Motion Picture Theaters industry (NAICS code 512131) consists of all establishments that primarily display first-run and second-run motion pictures,

Table 5: Frequencies and Transition Rates from the County Business Patterns

% of μSA-Year Observations by Number of Movie Theaters				
0	1	2	3	≥ 4
19.3	50.6	19.4	5.8	4.9

$\downarrow N_{t-1}/N_t \rightarrow$	% of Transitions Given N_{t-1}: 0	1	2	3	≥ 4
0	88.9	9.3	1.6	0.2	0.0
1	4.3	89.2	5.9	0.4	0.1
2	0.7	13.1	77.9	7.2	1.1
3	0.0	3.7	22.1	62.2	11.9
≥ 4	0.4	0.8	2.4	13.0	83.5

Note: The top panel gives the distribution of the number of movie theaters per μSA from 2000 to 2009 from the County Business Patterns for the 573 μSAs in our sample. The bottom panel displays the conditional probability of transitioning from N_{t-1} movie theaters in a μSA at time $t-1$ (row) to N_t theaters at time t (column).

except for drive-in theaters. Our estimation uses annual counts of the number of *theaters* (not firms) in each μSA from the County Business Patterns (CBP), beginning in 2000 and ending in 2009. The top panel of Table 5 reports the frequencies of the number of theaters across all of the μSA-year observations. No theaters serve the market in about twenty percent of the observations, a single theater serves about half of them, and about thirty percent of our observations have more than one theater. The maximum number of theaters observed is nine, but only 4.9 percent of the observations have four or more. Each row of Table 5 reports the observed frequencies of the number of theaters conditional on its previous year's value. Regardless of the initial number of theaters, the most common outcome is for it to remain unchanged. Nevertheless, the number of theaters changes in about 15 percent of the observed annual transitions.

In addition to this panel of producer counts, our estimation requires repeated measurements of the demand indicator C and cross-sectional measurements of time-invariant market characteristics X. The time-invariant market characteristics we employ are median income, dummy variables indicating membership in the nine U.S. Census Divisions, and an indicator for geographic preference diversity above its median value. For C, we use annual population for each μSA as published by the Census Bureau. For our sample from 2000 to 2009, the mean and standard deviation of the annual population growth rate equal 0.34 percent and 1.11 percent. The Census Bureau estimates these for non-census years using the most recent decennial census as a baseline, so they have very large adjustments between 2009 and 2010. The mean and standard deviation of population growth between these two years equals 1.5 percent and 3.1 percent. Since the measured changes between 2009 and 2010 disproportionately arise from differences in measurement methodology rather than true population changes, we end our estimation sample in 2009.

6.2 Estimates

Our model's estimation uses the demand-process specification from our Monte Carlo exercise. We restrict $C_{r,t}$ to a grid of 200 points equally spaced on a logarithmic scale. Its minimum value equals the minimum population observed in our data, $11,011$, divided by 1.25. Analogously, its maximum value equals the maximum population, $197,912$ multiplied by 1.25. For estimation, we replace each observation of μSA population with the closest grid point. The maximum number of movie theaters sustainable, $\check{n}$, is fixed at the maximum number of theaters observed in the data, nine. We give the cost shocks' logarithms a normal distribution with standard deviation ω and mean $-\omega^2/2$, normalize κ to one, and fix the discount factor ρ at $\frac{1}{1.05}$. The specification for the producers' surplus function is augmented with the effects of market characteristics $X_r \equiv (X_r^{(1)}, X_r^{(2)})$:

$$\pi_r(n,c) = \exp(\beta' X_r^{(1)}) \frac{c}{n} k(n; X_r^{(2)}). \tag{20}$$

The market characteristics $X_r^{(1)}$ affect the surplus log-linearly and include the logarithm of median income, in deviation from the logarithm of the average median income across our 573 μSAs, and dummies for all Census Divisions excluding New England. The remaining characteristics, $X_r^{(2)}$, interact with k and thereby affect the toughness of competition in a general way.[16] We both estimate the model without heterogeneity in k across markets (trivial $X_r^{(2)}$) and with k depending on whether a market's geographic preference diversity is above or below its median value, 13.4 miles ($X_r^{(2)}$ equal to an indicator for diversity exceeding 13.4 miles). We

[16]Our estimation code constrains the flow surplus to weakly decrease with the number of firms, as required by the monotonicity assumption in Section 2. This constraint does not bind at the maximum-likelihood estimates.

set $k(4) = k(5) = \cdots = k(9)$ to accommodate the paucity of observations with four or more theaters.

Table 6 reports the estimated parameters for two specifications, one that ignores geographic preference diversity and another that takes geographic diversity into account. In the first specification, the full-information maximum likelihood estimates of the demand process drift and innovation standard deviation are very close to the unconditional sample mean and standard deviation of population growth; 0.34 and 1.21 percent versus 0.34 and 1.11 percent. The coefficients in β are jointly and (with the exceptions of those multiplying three division dummies) individually significant. The mean sunk cost of entry, φ, is over fifty times the mean fixed cost of continuation. However, one should *not* interpret this as a measure of the typical sunk cost paid because entry only occurs when the realization of the cost shock is low. To calculate more informative measures of fixed and sunk costs, we simulated the estimated model for the New England Census Division (which includes Connecticut, Maine, Massachusetts, New Hampshire, Rhode Island and Vermont). In the simulation, the average fixed cost of continuation and sunk cost of entry *paid* were 0.78 and 1.53. The estimates of all these parameters from the specification that accounts for geographic preference diversity are similar to these baseline estimates.

Table 7 reports transformations of the estimates from Table 6 with a more straightforward economic interpretation. The first row reports $1/(k(1) \times 10^3)$, the population in thousands that sets a monopolist's current profit (the surplus earned minus the fixed continuation cost incurred in a period) to zero in a New England market with average median income when the fixed continuation cost equals one. The baseline specification's estimate of this is 16,790 people. We expect that concentrating customers' locations increases a monopolist's profit by making it easier to simultaneously satisfy their geographic preferences. The estimates from the model that accounts for geographic preference diversity support this prior. It takes $18,730$ people to support a monopolist in a μSA with geographic preference diversity above the median and $16,540$ to support a monopolist in a μSA with preference diversity below the median. A Wald test indicates that this difference is significant at the five percent level.

The remaining rows of Table 7 report estimates of $k(n+1)/k(n)$, the share of the surplus per consumer left after the addition of a competitor. These indicate

Table 6: Parameter Estimates

	All μSAs	Geographic Preference Diversity: Diversity > 13.4 miles	Geographic Preference Diversity: Diversity ≤ 13.4 miles
$k(1) \times 10^5$	5.96	5.34	6.05
	(0.79)	(0.73)	(0.85)
$k(2) \times 10^5$	3.22	3.20	2.90
	(0.47)	(0.47)	(0.48)
$k(3) \times 10^5$	2.65	2.68	2.26
	(0.41)	(0.42)	(0.43)
$k(4) \times 10^5$	2.03	2.12	1.50
	(0.32)	(0.35)	(0.35)
φ	51.24	49.09	
	(9.82)	(9.76)	
Median Income	0.88	0.85	
	(0.17)	(0.17)	
Mid Atlantic	-0.64	-0.60	
	(0.15)	(0.15)	
East North Central	-0.49	-0.45	
	(0.14)	(0.14)	
West North Central	0.05	0.09	
	(0.16)	(0.16)	
South Atlantic	-0.73	-0.71	
	(0.15)	(0.15)	
East South Central	-0.52	-0.49	
	(0.16)	(0.16)	
West South	-0.33	-0.29	
	(0.15)	(0.15)	
Mountain	-0.12	-0.07	
	(0.15)	(0.15)	
Pacific	-0.11	-0.08	
	(0.15)	(0.15)	
ω	1.75	1.74	
	(0.07)	(0.08)	
$\mu \times 10^2$	0.34	0.34	
	(0.02)	(0.02)	
$\sigma \times 10^2$	1.21	1.21	
	(0.01)	(0.01)	
$-\mathcal{L}$	9199.09	9192.20	
Number of Markets	573	287	286

Note: Standard errors are reported in parentheses. The data include 573 μSAs from 2000 to 2009, and $\check{n}$ equals nine, which is the maximum of the number of active firms observed in the data. The values of $k(5), \ldots, k(9)$ identically equal $k(4)$. The first column reports estimates of a specification in which k does not vary between markets. The second and third columns report estimates of a specification in which k differs between markets with geographic preference diversity below and above its median value of 13.4 miles ($X_r^{(2)}$). Both specifications include the logarithm of median income (in deviation from the logarithm of the average median income across μSAs) and Census Division dummies (excluding New England) as market characteristics in $X_r^{(1)}$. Please see the text for further details.

Table 7: Estimates of the Toughness of Competition

		Geographic Preference Diversity	
	All μSAs	Diversity > 13.4 miles	Diversity ≤ 13.4 miles
$1/\left(k(1) \times 10^3\right)$	16.79	18.73	16.54
	(2.23)	(2.55)	(2.32)
$k(2)/k(1)$	0.54	0.60	0.48
	(0.14)	(0.14)	(0.20)
$k(3)/k(2)$	0.82	0.84	0.78
	(0.06)	(0.06)	(0.10)
$k(4)/k(3)$	0.77	0.79	0.66
	(0.08)	(0.08)	(0.21)
Number of Markets	573	287	286

Note: This Table is based on the model's estimates as reported in Table 6. Standard errors are reported in parentheses. The ratio $1/(k(1) \times 10^3)$ can be interpreted as the population (in thousands of people) that sets a monopolist's current profit to zero in a New England market with average median income when the fixed cost equals one. The ratio $k(n+1)/k(n)$ is an indicator of the toughness of competition. Please see the text for further details.

very tough competition. In the first specification, duopolists' producers' surplus per consumer equals 54 percent of a monopolist's. A third competitor decreases this surplus to 82 percent of the duopolists'. In markets with four or more competitors, the producers' surplus per consumer further decreases to 77 percent of that in a market with three firms. Altogether, adding three or more theaters to a market with a single incumbent brings the surplus per customer down to 34 percent of its monopoly value.

The theoretical literature on spatial differentiation overwhelmingly points to heterogeneity of consumers' locations as a source of market power. This leads us to expect producers' surplus to fall less rapidly with additional competition in the high-diversity markets. The estimates from the second specification support this conjecture. The ratios of the producers' surplus $k(n+1)/k(n)$ are indeed higher in high diversity than in low diversity markets. A Wald test indicates that these differences are jointly statistically significant at the one percent level. This suggests that entering theaters can lessen the toughness of competition with their location choices.

These estimates of tough competition are difficult to reconcile with other evidence from this industry. Davis (2005) provides evidence on competition for customers from regressions of theaters' admissions prices against indicators of the presence of other theaters at various distances using data from large (relative to

μSAs) U.S. cities in the 1990s. Based on both across-market and within-market-over-time variation, he concludes that

> ... the magnitude of the price-reducing effect of local competition appears to be economically modest.

Prior research on the vertical relationships between theater owners and their upstream suppliers, film distributors, has emphasized formal and informal arrangements to manage the *popcorn conflict* over the final ticket price: Popcorn and other concession sales are complements with theater attendance, and theater owners keep all surplus from concession sales while splitting surplus from ticket sales with the film distributor. Therefore, theater owners prefer lower ticket prices than do distributors. The motion picture industry operates under a relatively unique legal regime, under which the producers of films are legally barred from directly influencing box-office pricing or vertically integrating with motion picture theaters. Nevertheless, repeated interactions between distributors and theater owners might give distributors indirect and extralegal control over box-office prices.[17] Supporting the view that film distributors constrain theaters' pricing choices, Davis (2006) finds that

> ... the average theater owner would prefer to actually lower admissions prices, if she could attract the same set of films.

Accordingly, we find it implausible that the estimates reflect fierce competition for customers.

We can think of two alternative causes for our estimates of very tough competition that are consistent with the lack of competitive pressure on admissions prices. In the first, adding a theater increases competition in the market for film exhibition rights. Indeed, Gil and LaFontaine (2012) find some evidence that owners of Spanish theaters with higher local market shares get better deals from film distributors. In the second, incumbent firms successfully deter entry and appropriate expansion opportunities from market growth to themselves. The apparent unwillingness of new firms to enter a market really reflects an incumbent simultaneously deterring entry by rival firms and delaying expansion of its own

[17] Orbach and Einav (2007) review the legal environment in which theater owners negotiate with film distributors and set admissions prices.

theaters so as not to cannibalize its existing operations. Rigorously assessing the first possibility requires heretofore unavailable U.S. data on contracts between film distributors and theater owners in a variety of competitive environments. In principle, we can test the second possibility using records of theaters' firm affiliations in the payroll tax records that underlie the CBP data. These are available within U.S. Census Research Data Centers, and our research agenda includes their examination.

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. μSAs. That this relatively complete development and application of a dynamic oligopoly model was feasible validates our title's assertion that our model's Markov-perfect industry dynamics are "very simple."

We have emphasized that our model features homogeneous firms, but accounting for observable heterogeneity with *serially uncorrelated* shocks is straightforward and merely requires reinterpreting $\pi(n, c)$ as expected surplus. Adding *persistent* heterogeneity while retaining analytic simplicity is more challenging but still feasible. Abbring et al. (2010) extend the present paper's model with persistent and publicly-observed firm-specific shocks to profitability. These satisfy Ericson and Pakes's (1995) assumption that idiosyncratic shocks can only improve a firm's type. In the case with $\check{n} = 2$, they prove an analogue of Lemma 2: Adding a second firm to a market never raises an incumbent's equilibrium payoff. For this reason, they describe that model's equilibrium dynamics as "simple." This result allows them to prove an analogue to Theorem 1— that there is an essentially unique natural (Cabral, 1993) Markov-perfect equilibrium— and to extend this paper's algorithm for its fast calculation. If instead $\check{n} > 2$, adding a firm can *raise* incumbents' continuation values *even though adding competitors weakly reduces per-period producers' surplus.*

Nevertheless, Abbring et al. show that an appropriate extension of this paper's algorithm can calculate *every* natural Markov-perfect equilibrium in which firms coordinate on joint survival when this Pareto dominates simultaneous exit. Thus, although we have developed our analysis in a very simple environment, many of its insights apply to richer frameworks.

Appendix

The model in the main text embodies structure on the stochastic processes for the state variables and their effects on profits that simplifies its maximum likelihood estimation but contributes nothing to its theoretical analysis. This appendix presents a *general model* without this structure that encompasses the *special model* of the main text, and it proves the appropriate generalizations of Lemmas 1 and 2, Corollaries 1 and 2, and Theorem 1.

A Primitives

In the general model, firms' profits in period t depend on the $\mathcal{Y}$-valued vector Y_t. Figure 2 gives the model's recursive extensive form. Period t starts in state (N_t, Y_t), and each incumbent earns profits $\tilde{\pi}(N_t, Y_t)$.[18] As in the special model, all players have names giving the date of their entry opportunity and their position in that date's entry queue. In the entry subgame of period t, firm (t, j) pays the sunk cost $\tilde{\varphi}(M_t^j, Y_t)$ upon entry, where again $M_t^j \equiv N_t + j$. As before, a potential entrant's payoff from choosing inactivity equals zero. Progressing to the period t survival subgame, an active firm choosing survival incurs no cost *during period* t. The expected profits from operating in period $t+1$ subsume the special model's costs of continuation. At the end of the period, nature draws Y_{t+1} from from the Markov transition distribution $\tilde{G}(\cdot\,|Y_t)$.

The restrictions we place on the payoffs are

A1. $\mathbb{E}\left[\tilde{\pi}(n, Y')|Y = y\right]$ exists and $\exists \breve{\pi} < \infty : \forall n \in \mathbb{N}$ and $\forall y \in \mathcal{Y}$, $\mathbb{E}\left[\tilde{\pi}(n, Y')|Y = y\right] < \breve{\pi}$;

A2. $\exists \breve{n} \in \mathbb{N} : \forall n > \breve{n}$ and $\forall y \in \mathcal{Y}$, $\tilde{\pi}(n, y) < 0$;

A3. $\forall n \in \mathbb{N}$ and $\forall y \in \mathcal{Y}$, $\tilde{\pi}(n, y) \geq \tilde{\pi}(n+1, y)$; and

A4. $\forall m \in \mathbb{N}$ and $\forall y \in \mathcal{Y}$, $0 \leq \tilde{\varphi}(m, y) \leq \tilde{\varphi}(m+1, y)$.

[18] Here and throughout this appendix, we place a tilde over any primitive, strategy, or value function with a similar name in the special model.

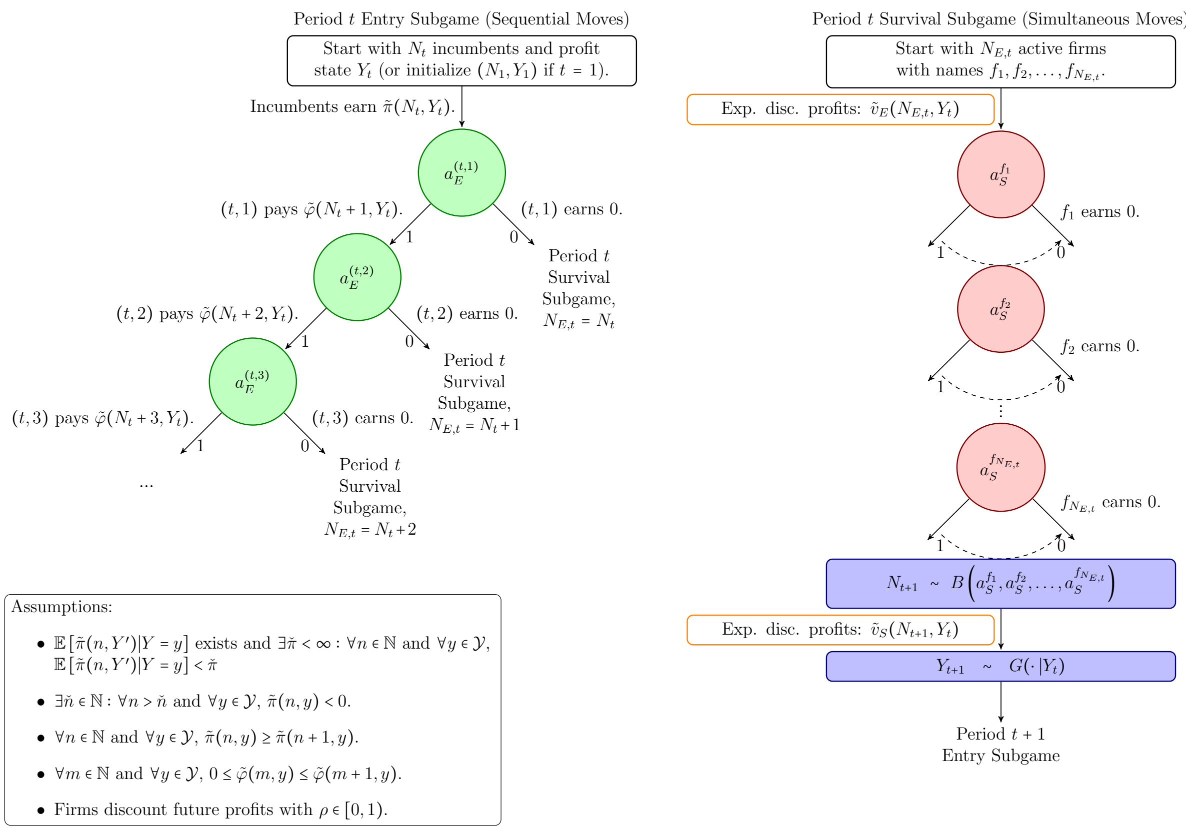

Figure 2: The General Model's Recursive Extensive Form

To cast the special model within this more general framework, set

$$
\begin{aligned}
Y_t &\equiv (C_t, W_t, W_{t-1}), \\
\tilde{\pi}(n; c, w, w_{-1}) &\equiv \pi(n, c) - \rho^{-1}\kappa \exp(w_{-1}), \\
\tilde{\varphi}(m; c, w, w_{-1}) &\equiv \varphi \exp(w), \text{ and} \\
\tilde{G}(c, w, w_{-1}|C_{t-1}, W_{t-1}, W_{t-2}) &\equiv \begin{cases} G_C(c|C_{t-1})G_W(w) & \text{if } W_{t-1} \leq w_{-1} \\ 0 & \text{otherwise.} \end{cases}
\end{aligned}
$$

Note that in the special model,

$$
\mathbb{E}[\tilde{\pi}(n, Y')|Y = (c, w, w_{-1})] = \mathbb{E}[\pi(n, C')|C = c] - \rho^{-1}\kappa \exp(w).
$$

Thus, the assumptions that $\mathbb{E}[\pi(n, C')|C = c] \leq \check{\pi}$, $\pi(n, c) \geq 0$, and $\kappa > 0$ ensure that $\mathbb{E}[\tilde{\pi}(n, Y')|Y = y]$ exists and is bounded above by $\check{\pi}$, so Assumption A1 holds.

This allows us to restrict equilibrium post-entry values to the space of bounded non-negative functions. The special model's assumptions that $\kappa > 0$ and that there exists $\check{n}$ such that $\pi(n, c) = 0$ for all $n > \check{n}$ imply Assumption A2, which is key to bounding the number of firms in equilibrium. The assumption that $\pi(n, c) \geq \pi(n+1, c)$ in the special model directly gives Assumption A3. Finally, Assumption A4 generalizes the special model's assumption of a constant $\varphi \geq 0$. Since $\tilde{\varphi}(m, \cdot)$ may increase in m, the general model can have an *economic barrier to entry* (McAfee et al., 2004).[19]

B Equilibrium

For both a potential entrant and a firm contemplating survival, the payoff-relevant variables are the number of firms committed to play that period's survival subgame and the current state of demand. A Markovian strategy is a pair of functions, $\tilde{a}_E : \mathbb{N} \times \mathcal{Y} \to \{0, 1\}$ and $\tilde{a}_S : \mathbb{N} \times \mathcal{Y} \to [0, 1]$. An equilibrium strategy and its associated continuation values $\tilde{v}_E$ and $\tilde{v}_S$ satisfy conditions analogous to the special model's equations (1), (2), (3) and (4):

$$\tilde{v}_E(n_E, y) = \max_{a \in [0,1]} a \mathbb{E}_{\tilde{a}_S} \left[\tilde{v}_S(N', y) | N_E = n_E, Y = y \right], \tag{21}$$

$$\tilde{v}_S(n', y) = \rho \mathbb{E}_{\tilde{a}_E} \left[\tilde{\pi}(n', Y') + \tilde{v}_E(N'_E, Y') | N' = n', Y = y \right], \tag{22}$$

$$\tilde{a}_E(m, y) \in \arg\max_{a \in \{0,1\}} a \Big(-\tilde{\varphi}(m, y) + \mathbb{E}_{\tilde{a}_E} \left[\tilde{v}_E(N_E, y) | M = m, Y = y \right] \Big), \tag{23}$$

and

$$\tilde{a}_S(n_E, y) \in \arg\max_{a \in [0,1]} a \Big(\mathbb{E}_{\tilde{a}_S} \left[\tilde{v}_S(N', y) | N_E = n_E, Y = y \right] \Big). \tag{24}$$

The expectation operators condition on the deciding firm choosing to be active in the next period and on all other firms using the strategy in the operator's subscript. Note that in the special model $\tilde{v}_E$, $\tilde{a}_E$ and $\tilde{a}_S$ equal v_E, a_E and a_S; but $\tilde{v}_S(n'; c, w, w_{-1}) = v_S(n', c) - \kappa \exp(w)$ because the general model incorporates the continuation costs into $\tilde{\pi}(n, y)$. As in the text, we restrict attention to equilibria in strategies that default to inactivity. In such an equilibrium, a potential entrant that is indifferent

[19] Although adding a barrier to entry to the model's theoretical analysis is straightforward, our identification proof does rely on the special model's constant specification for $\tilde{\varphi}(m, \cdot)$. Thus, the identification of barriers to entry remains an important open area of inquiry.

between entering or not stays out,

$$\mathbb{E}_{\tilde{a}_E}\left[\tilde{v}_E(N_E, y) | M = m, Y = y\right] = \tilde{\varphi}(m, y) \Rightarrow \tilde{a}_E(m, y) = 0,$$

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

$$\tilde{v}_S(n_E, y) = \cdots = \tilde{v}_S(1, y) = 0 \Rightarrow \tilde{a}_S(n_E, y) = 0.$$

The general model's equilibrium characterization begins with the appropriate analogues to Lemmas 1 and 2 and Corollaries 1 and 2. Section C contains their proofs.

Lemma 1⋆ (Bounded number of firms in the general model) *In a symmetric Markov-perfect equilibrium, $\forall y \in \mathcal{Y}$ and $\forall n > \check{n}$, $\tilde{a}_E(n, y) = 0$ and $\tilde{a}_S(n, y) < 1$.*

Lemma 2⋆ (Monotone equilibrium payoffs in the general model) *In a symmetric Markov-perfect equilibrium, $\forall y \in \mathcal{Y}$, $\tilde{v}_S(n', y)$ weakly decreases with n'.*

Corollary 1⋆ *Let $\tilde{v}_S$ 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 2), each of the n' survivors earns $\tilde{v}_S(n', y)$, with $\tilde{v}_S(n', y) \neq 0$ for at least one $n' \in \{1, \ldots, n_E\}$, and each exiting firm earns zero. This game has a unique symmetric Nash equilibrium, possibly in mixed strategies.*

Corollary 2⋆ *If $\tilde{v}_E$ and $\tilde{v}_S$ are the post-entry and post-survival value functions associated with a symmetric Markov-perfect equilibrium, then*

$$\tilde{v}_E(n_E, y) = \max\{0, \tilde{v}_S(n_E, y)\}.$$

Just as with the special model, we demonstrate equilibrium existence and uniqueness by constructing an equilibrium. We present the algorithm for this equilibrium calculation in Algorithm 1. It begins by initializing the number of firms under consideration n to $\check{n}$ and both the candidate equilibrium entry rule $\tilde{\alpha}_E : \mathbb{N} \times \mathcal{Y} \to \{0, 1\}$ and a dummy function $f^\star : \mathcal{Y} \to [0, \frac{\rho\check{\pi}}{1-\rho}]$ to zero. We will denote the candidate equilibrium post-entry and post-survival value functions by

$\tilde{\nu}_E$ and $\tilde{\nu}_S$, respectively. The algorithm then enters its main loop, which begins by constructing the number of firms that will be active if we begin with n firms and potential entrants follow the rule in the current value of $\tilde{\alpha}_E$. For a given $y \in \mathcal{Y}$, this is

$$\tilde{\mu}(n, y) \equiv n + \sum_{n'=n+1}^{\infty} \prod_{m=n+1}^{n'} \tilde{\alpha}_E(m, y).$$

In the first pass through the loop, $\tilde{\mu}(\check{n}, y) = \check{n}$. Next, the algorithm uses Bellman equation iteration (on the dummy function $f^\star$) to solve the general model's analogue to the dynamic programming problem in (10). The relevant Bellman operator is

$$\begin{aligned} T_n(f^\star)(y) = \max\{0, \rho\mathbb{E}[\tilde{\pi}(n, Y') + \mathbb{1}[\tilde{\mu}(n, Y') = n] f^\star(Y') \\ + \mathbb{1}[\tilde{\mu}(n, Y') > n]\tilde{\nu}_E(\tilde{\mu}(n, Y'), Y') \,|\, Y = y]\}. \end{aligned} \tag{25}$$

The next two steps assign the fixed point of T_n stored in $f^\star(\cdot)$ to $\tilde{\nu}_E(n, \cdot)$ and use (22) to construct $\tilde{\nu}_S(n, \cdot)$. In the loop's final step, the initial value of $\tilde{\alpha}_E(n, y)$ is replaced with $\mathbb{1}[\tilde{\nu}_E(n, y) > \tilde{\varphi}(n, y)]$. If the current value of n exceeds one, it is decremented and the algorithm returns to the top of the main loop. If instead n equals one, then the algorithm proceeds to its final task, setting the candidate equilibrium survival rule to zero in the trivial case that $\tilde{\nu}_S(n, \cdot) = \cdots = \tilde{\nu}_S(1, \cdot) = 0$, and to the highest probability consistent with equilibrium in the survival subgame otherwise.[20] Algorithm 1 only computes candidate post-entry and post-survival values and a candidate survival rule on $\{1, \ldots, \check{n}\} \times \mathcal{Y}$. As noted in the main text, it is straightforward to extend them to the full state space $\mathbb{N} \times \mathcal{Y}$. Because this extension is not of much applied interest, we keep it implicit and simply refer to Algorithm 1 as computing a candidate equilibrium.

The appropriate generalization of Theorem 1 to the general framework states that the candidate equilibrium strategies and payoffs arising from Algorithm 1 correspond to the unique Markov-perfect Nash equilibrium that defaults to inactivity.

Theorem 1* (Equilibrium existence and uniqueness in the general model) *There exists a unique symmetric Markov-perfect equilibrium that defaults to inactivity. The equilibrium strategy and corresponding equilibrium payoffs are those computed by Algorithm 1.*

[20] Because these survival rule calculations do not build on each other, they can be parallelized.

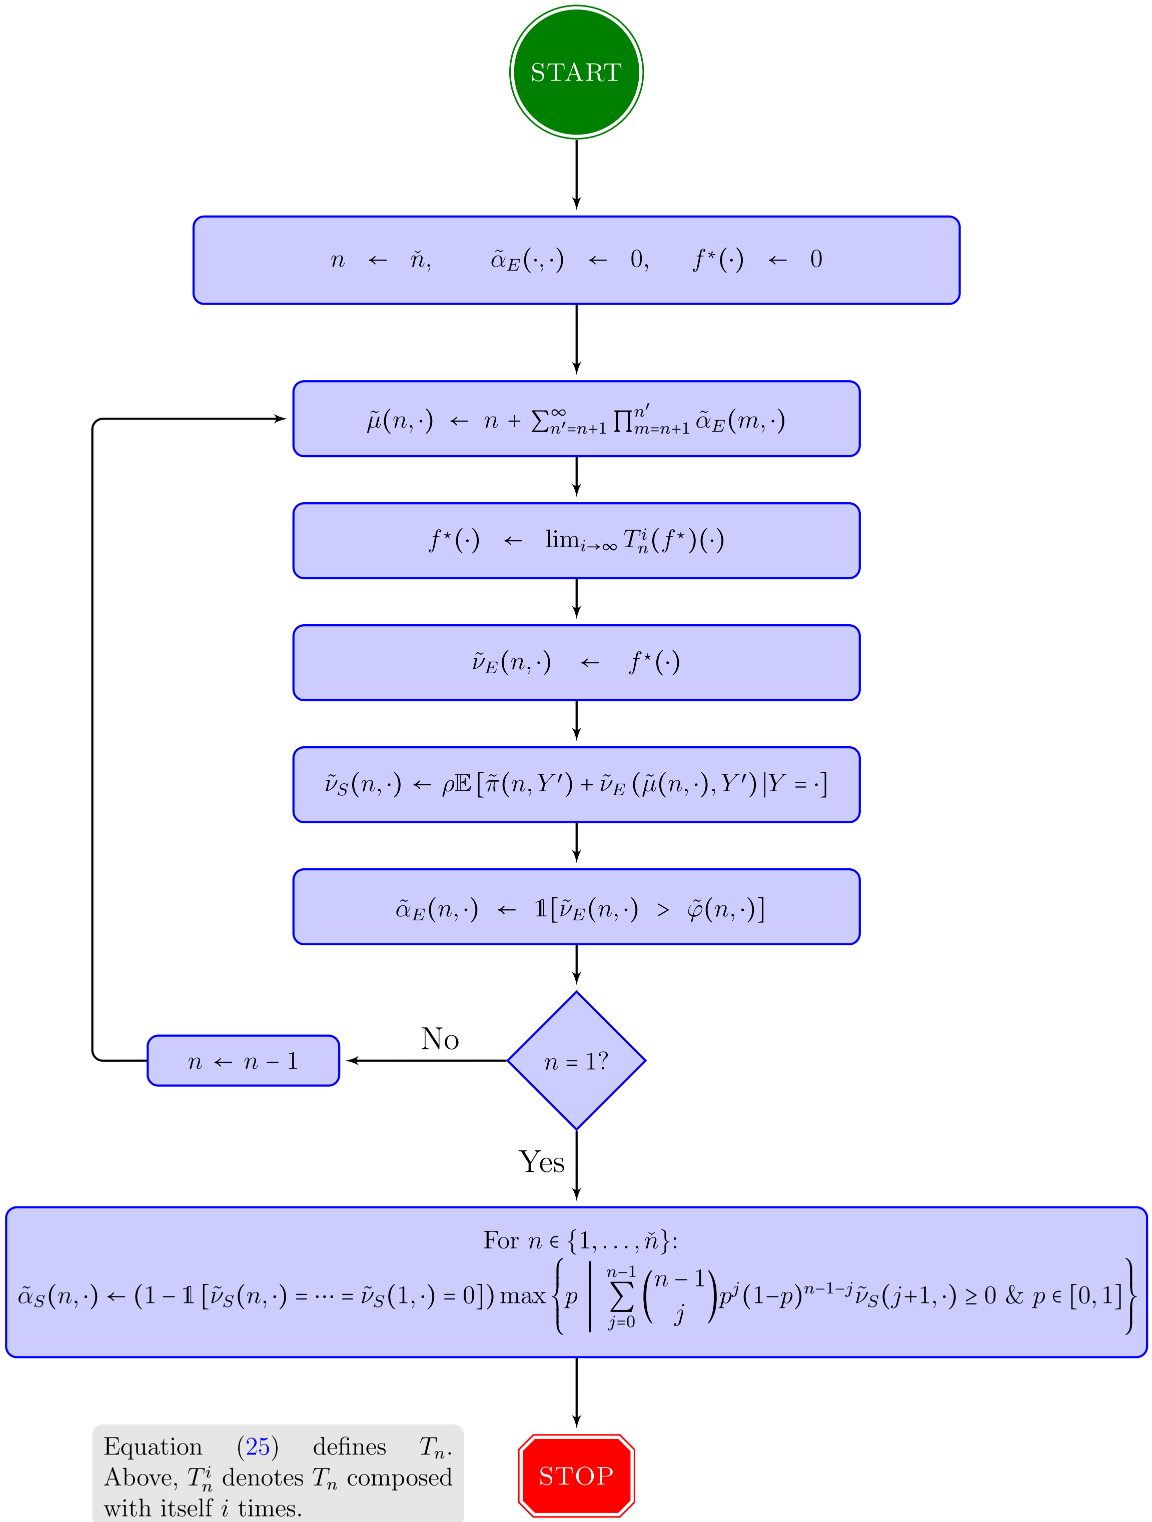

Algorithm 1: Equilibrium Calculation for the General Model

C Proofs

This final section contains the formal proofs of the numbered results.

Proof of Lemma 1*. Consider a period-$t^\star$ survival subgame with $N_{E,t^\star} = n > \check{n}$ firms and profit state $Y_{t^\star} = y$. Define the random time τ as the first period weakly after $t^\star$ in which the firms choose exit with a positive probability, with $\tau \equiv \infty$ if they never do:

$$\tau \equiv \min\{\{t \geq t^\star : \tilde{a}_S(N_{E,t}, Y_t) < 1\} \cup \{\infty\}\}.$$

If indeed $\tilde{a}_S(n, y) < 1$ as asserted, then $\tau = t^\star$. Suppose to the contrary that $\tilde{a}_S(n, y) = 1$, so $\tau > t^\star$. By definition, exit can occur only in or after period τ, so we know that $N_t = N_{E,t-1} \geq n$ for $t \in \{t^\star + 1, \ldots, \tau\}$. Since $n > \check{n}$, this together with Assumption A2 implies that $\tilde{\pi}(N_t, Y_t) < 0$ for $t \in \{t^\star + 1, \ldots, \tau\}$. If $\tau = \infty$, then the incumbent firms receive an infinite sequence of strictly negative payoffs. If instead $\tau < \infty$, then the incumbent firms receive a finite sequence of strictly negative payoffs followed by a zero expected continuation value from playing the period-τ survival subgame ($\tilde{v}_E(N_\tau, Y_\tau) = 0$). Therefore, the expectation of the discounted sum of payoffs, $\tilde{v}_S(n, y)$, must be strictly negative. Since any incumbent firm can raise its payoff to zero by choosing certain exit, the supposition that $\tilde{a}_S(n, y) = 1$ must be incorrect.

Next, consider the decision of the first potential entrant, firm $(t^\star, 1)$, in a period-$t^\star$ entry subgame that starts with $N_{t^\star} = n - 1 > \check{n} - 1$ incumbents and profit state $Y_{t^\star} = y$. (This value of n need not equal that used above to initialize an arbitrary period $t^\star$ survival subgame.) Note that this firm pays $\tilde{\varphi}(n, y)$ upon entry. In return, it receives a continuation value of zero (as proven above). If $\tilde{\varphi}(n, y) > 0$, then it maximizes its payoff by staying out of the industry and earning zero. If instead $\tilde{\varphi}(n, y) = 0$, then the assumption that the equilibrium strategy defaults to inactivity dictates the same action. In either case, $\tilde{a}_E(n, y) = 0$ as asserted. ■

Proof of Lemma 2*. Consider a subgame beginning at the start of period $t^\star$ with $N_{t^\star} = n'$ and $Y_{t^\star} = y'$. We refer to this as the *original* subgame and denote its equilibrium outcomes for $t \geq t^\star$ simply with N_t and $N_{E,t}$. Now consider a second period-$t^\star$ subgame with $Y_{t^\star} = y'$ and one additional firm. We refer to this as the *perturbed* game and denote its initial number of firms with $N^+_{t^\star} = n' + 1$ and its outcomes for $t \geq t^\star$ with N_t^+ and $N_{E,t}^+$. We assume that the original and perturbed subgames games share a realization of $\{Y_t\}_{t=t^\star+1}^{\infty}$.

Define the random times τ and τ^+ with

$$\begin{aligned} \tau &\equiv \min\{\{t \geq t^\star : \tilde{a}_S(N_{E,t}, Y_t) < 1\} \cup \{\infty\}\} \text{ and} \\ \tau^+ &\equiv \min\{\{t \geq t^\star : \tilde{v}_S(N_{E,t}^+, Y_t) \leq 0\} \cup \{\infty\}\}. \end{aligned}$$

With these definitions the probability of survival in the original subgame equals one until period τ, when it falls to a number less then one. Similarly (but not identically), the probability of survival in the perturbed subgame equals one at least until τ^+. By these definitions, N_t and N_t^+ are weakly increasing with t for all $t^\star \leq t \leq \tau$ and $t^\star \leq t \leq \tau^+$, respectively. We also know that $N_t \leq N_t^+$ for all $t^\star \leq t \leq \tau^+$, because otherwise the two subgames would have potential entrants in the same payoff-relevant states making different entry decisions. That would violate either the assumption that the equilibrium strategy is Markovian or the assumption that the equilibrium entry rule defaults to inactivity.

We first wish to show that $\tau^+ \leq \tau$ always. Suppose to the contrary that there exist realizations of $\{Y_t\}_{t=t^\star}^{\infty}$, $\{N_{E,t}\}_{t=t^\star}^{\infty}$, $\{N_t\}_{t=t^\star}^{\infty}$, $\{N_{E,t}^+\}_{t=t^\star}^{\infty}$, and $\{N_t^+\}_{t=t^\star}^{\infty}$ such that $\tau < \tau^+$. Consider the original and perturbed period-τ survival subgames under such realizations. The definition of τ requires that $\tilde{a}_S(N_{E,\tau}, Y_\tau) < 1$. Suppose that one of the original subgame's $N_{E,\tau}$ incumbents deviates from this strategy and chooses certain survival. For $t > \tau$, let $\mathcal{N}_t$ denote the number of firms that arises if (i) the deviating incumbent survives until period t and (ii) all other players follow the equilibrium strategy. Since $\tilde{a}_S(N_{E,\tau}, Y_\tau) < 1$, we know that no entry occurred in period τ of the original subgame.[21] Therefore, $N_{E,\tau} = N_\tau$. Since some of the $N_{E,\tau} - 1$ non-deviating active firms might not survive, we have $\mathcal{N}_{\tau+1} \leq N_{E,\tau}$. Combining these with $N_\tau \leq N_\tau^+$ and $N_\tau^+ \leq N_{\tau+1}^+$ yields $\mathcal{N}_{\tau+1} \leq N_{\tau+1}^+$. Note that (i) no exit can occur in the perturbed subgame before τ^+, (ii) exit in the post-deviation original subgame only reduces $\mathcal{N}_t$, and (iii) entry in the post-deviation original subgame cannot leave $\mathcal{N}_t > N_t^+$ for $t \in \{\tau+1, \ldots, \tau^+\}$. Therefore, we know that $\mathcal{N}_t \leq N_t^+$ for $t = \tau+1, \ldots, \tau^+$. Assumption A3 then implies that $\tilde{\pi}(\mathcal{N}_t, Y_t) \geq \tilde{\pi}(N_t^+, Y_t)$ for $t = \tau+1, \ldots, \tau^+$.

By choosing survival until the realization of τ^+ and exiting at that date, the deviating incumbent in the original period-τ survival subgame earns weakly greater

[21] Otherwise, either an entrant pays a positive sunk cost for a zero continuation value or the proposed equilibrium strategy does not default to inactivity.

payoffs than an incumbent in the perturbed period-τ survival subgame. Because $\tau^+ > \tau$, $\tilde{v}_S(N^+_{E,\tau}, Y_\tau) > 0$ in the perturbed survival subgame. Therefore, we conclude that the deviating incumbent in the original subgame receives a payoff strictly greater than that earned from following the equilibrium strategy, zero. This is inconsistent with the supposition of equilibrium, so we conclude that $\tau^+ \leq \tau$.

Now return to the original and perturbed subgames from the beginning of period $t^\star$. Since $\tau^+ \leq \tau$, period-$t^\star$ incumbents in both subgames survive until τ^+; and because $N_t \leq N^+_t$ for $t \in \{t^\star + 1, \ldots, \tau^+\}$, the original game's incumbents earn a weakly higher payoff until τ^+. Furthermore, the original game's incumbents must have a weakly higher period-τ^+ post-survival value because (i) $a_S(N_{\tau^+}, Y_{\tau^+}) = 1$ implies that $v_S(N_{\tau^+}, Y_{\tau^+}) \geq 0$ and (ii) $v(N^+_{\tau^+}, Y_{\tau^+}) \leq 0$ from the definition of τ^+. These two conclusions together imply that an incumbent at the start of the original period-$t^\star$ subgame, with n' firms and profit state y', has weakly higher expected discounted payoffs than one of the $n' + 1$ incumbents in the perturbed subgame. Because this ranking is preserved under integration over y' with respect to $d\tilde{G}(y'|y)$, it follows that $\tilde{v}_S(n', y) \geq \tilde{v}_S(n' + 1, y)$. ■

Proof of Corollaries 1⋆ and 2⋆. Corollary 1⋆'s one-shot survival game falls into one of three mutually-exclusive cases.

- $\tilde{v}_S(1, y) \leq 0$. Lemma 2⋆ implies that $\tilde{v}_S(n', y) \leq 0$ for all $n' > 1$. Therefore, exiting for sure (setting $\tilde{a}_S(n_E, y) = 0$) is a weakly dominant strategy and forms one symmetric equilibrium. Furthermore, since $\tilde{v}_S(n', y) \neq 0$ for at least one $n' \in \{1, \ldots, n_E\}$, we know that $\tilde{v}_S(n_E, y) < 0$. Therefore, exiting for sure is also the unique best response to any positive symmetric continuation probability.

- $\tilde{v}_S(n_E, y) \geq 0$. Lemma 2⋆ implies that $\tilde{v}_S(n', y) \geq 0$ for all $n' \in \{1, \ldots, n_E\}$. Therefore, continuing for sure (setting $\tilde{a}_S(n_E, y) = 1$) is a weakly dominant strategy and forms a symmetric equilibrium. Furthermore, since $\tilde{v}_S(n', y) \neq 0$ for at least one $n' \in \{1, \ldots, n_E\}$, we know that $\tilde{v}_S(1, y) > 0$. Therefore, continuing for sure is also the unique best response to any symmetric continuation probability less than one.

- $\tilde{v}_S(1, y) > 0$ and $\tilde{v}_S(n_E, y) < 0$. No symmetric pure strategy equilibrium exists, because the best response to all other firms continuing for sure is to exit for sure, and vice versa. In a mixed strategy equilibrium, firms must be indifferent between continuation and exit. By the intermediate value theorem, there is some a_S that solves

$$\sum_{n'=1}^{n_E} \binom{n_E - 1}{n' - 1} a_S^{n'-1} (1 - a_S)^{n_E - n'} \tilde{v}_S(n', y) = 0. \tag{26}$$

Lemma 2⋆ and the subcase's conditions guarantee that the left hand side of (26) strictly decreases in a_S. So, there is only one symmetric mixed strategy equilibrium.

This establishes the equilibrium uniqueness asserted by Corollary 1⋆. To establish Corollary 2⋆, simply note that the equilibrium payoff to the survival subgame equals $\tilde{v}_S(n_E, y)$ if this is positive and equals zero otherwise. ■

Proof of Theorem 1⋆. The proof is divided into three parts. First, we show that the candidate continuation values satisfy the monotonicity requirements of Lemma 2⋆. Second, we use this to demonstrate that the candidate continuation values, survival rule, and entry rule satisfy (21)-(24). This establishes that the candidate strategy indeed forms an equilibrium. Third, we demonstrate that an equilibrium's existence implies its uniqueness.

Let $\tilde{\mu}$ be the value of the transition rule at the conclusion of Algorithm 1. (Because $\tilde{\mu}(n', \cdot)$ is set at the beginning of the loop iteration with $n = n'$ and is never altered again, it has the same value as that was used to construct the Bellman operator $T_{n'}$.) Fix $n \in \{1, 2, \ldots, \check{n}\}$ and suppose that for all n' such that $n+1 \leq n' \leq \check{n}$, we know that $\tilde{\nu}_E(n', \cdot) \geq \tilde{\nu}_E(n' + 1, \cdot)$. (This condition is trivially true for $n = \check{n}$.) Consider evaluating T_n at the value of $f^\star$ in memory *after* the completion of the

$n+1$-indexed dynamic programming problem. For all $y \in \mathcal{Y}$, this gives

$$
\begin{aligned}
T_n(f^\star)(y) = &\max\{0, \rho\mathbb{E}[\tilde{\pi}(n, Y') + f^\star(Y') \\
&+ \mathbb{1}[\tilde{\mu}(n, Y') > n](\tilde{\nu}_E(\tilde{\mu}(n, Y'), Y') - f^\star(Y'))|Y = y]\} \\
\geq &\max\{0, \rho\mathbb{E}[\tilde{\pi}(n+1, Y') + f^\star(Y') && (27)\\
&+ \mathbb{1}[\tilde{\mu}(n, Y') > n](\tilde{\nu}_E(\tilde{\mu}(n, Y'), Y') - f^\star(Y'))|Y = y]\} \\
= &\max\{0, \rho\mathbb{E}[\tilde{\pi}(n+1, Y') + f^\star(Y') && (28)\\
&+ \mathbb{1}[\tilde{\mu}(n, Y') > n+1](\tilde{\nu}_E(\tilde{\mu}(n, Y'), Y') - f^\star(Y'))|Y = y]\} \\
= &\max\{0, \rho\mathbb{E}[\tilde{\pi}(n+1, Y') + f^\star(Y') && (29)\\
&+ \mathbb{1}[\tilde{\mu}(n+1, Y') > n+1](\tilde{\nu}_E(\tilde{\mu}(n+1, Y'), Y') - f^\star(Y'))|Y = y]\} \\
= &\tilde{\nu}_E(n+1, y).
\end{aligned}
$$

The inequality in (27) follows from Assumption A3, and the equality in (28) from the equivalence of $\tilde{\nu}_E(n+1, Y')$ with $f^\star(Y')$. Since $\tilde{\nu}_E(n', Y)$ weakly decreases in n' for $n' > n$, so does $\tilde{\alpha}_E(n', Y')$. Therefore, $\tilde{\mu}(n, Y') = \tilde{\mu}(n+1, Y')$ whenever $\tilde{\mu}(n, Y') > n+1$. This gives us (29). The final equality follows again from the equivalence of $\tilde{\nu}_E(n+1, Y')$ with $f^\star(Y')$. The operator T_n is a monotone contraction mapping, so $T_n(f^\star)(n, \cdot) \geq \tilde{\nu}_E(n+1, \cdot)$ implies that its fixed point, $\tilde{\nu}_E(n, \cdot)$, weakly exceeds $\tilde{\nu}_E(n+1, \cdot)$. This is the desired monotonicity result for the proof's first part.

For the second part, begin with (24). For states (n, y) such that $\tilde{\nu}_S(n, y) = \cdots = \tilde{\nu}_S(1, y) = 0$, (24) imposes only the trivial requirement that $\tilde{\alpha}_S(n, y) \in [0, 1]$. Algorithm 1's selection of $\tilde{\alpha}_S(n, y) = 0$ satisfies this requirement. For all other states (n, y), since $\tilde{\nu}_S(n, y)$ is weakly decreasing in n, Algorithm 1 sets $\tilde{\alpha}_S(n, y)$ to one if $\tilde{\nu}_S(n, y) \geq 0$, to zero if $\tilde{\nu}_S(1, y) \leq 0$, and to the highest symmetric survival probability that leaves all firms indifferent between continuation and exit if both $\tilde{\nu}_S(n, y) < 0$ and $\tilde{\nu}_S(1, y) > 0$. In all three cases, $\tilde{\alpha}_S(n, y)$ forms a Nash equilibrium to the one-shot survival game with survival payoffs $\tilde{\nu}_S(n', y)$ for $n' = 1, \ldots, n$, so it satisfies (24). Equation (21) requires $\tilde{\nu}_E(n, y)$ to equal the expected payoff to this game, $\max\{0, \tilde{\nu}_S(n, y)\}$. To establish this equality, recall that Algorithm 1 sets $\tilde{\nu}_E(n, y)$ to the fixed point of (25). Therefore, we have

$$\tilde{\nu}_E(n, y) = \max\{0, \rho\mathbb{E}[\tilde{\pi}(n, Y') + \tilde{\nu}_E(\tilde{\mu}(n, Y'), Y') \,|\, Y = y]\},$$

where $\tilde{\mu}(n, Y') = n + \sum_{n'=n+1}^{\infty} \prod_{m=n+1}^{n'} \tilde{\alpha}_E(m, Y')$ is the number of firms active at the start of the next period's survival game when all potential entrants use the candidate equilibrium entry rule. Algorithm 1 assigns this maximum operator's second argument to $\tilde{\nu}_S(n, y)$. This directly ensures that the continuation values satisfy (22), and it sets $\tilde{\nu}_E(n, y)$ to $\max\{0, \tilde{\nu}_S(n, y)\}$, as required to satisfy (21). The only condition remaining unverified is (23). Since $\tilde{\nu}_E(n, y)$ weakly decreases in n, $\tilde{\alpha}_E(n, y)$ weakly decreases in n. Therefore, $\mathbb{E}_{\tilde{\alpha}_E}\left[\tilde{\nu}_E(N_E, y)|M = n, Y = y\right] > \tilde{\varphi}(n, y)$ if and only if $\tilde{\nu}_E(n, y) > \tilde{\varphi}(n, y)$. Since $\tilde{\alpha}_E(n, y)$ is an indicator for this inquality's truth, it prescribes entry if and only if entry gives a positive expected payoff. Thus, the candidate strategy and continuation values satisfy (23). We conclude that Algorithm 1's candidate strategy indeed forms an equilibrium.

The remainder of this proof demonstrates equilibrium uniqueness. Corollary $1^\star$ and the requirement that the strategy defaults to inactivity together imply that there is a unique equilibrium survival rule corresponding to every equilibrium post-survival value function. Lemma $2^\star$ and the requirement that the strategy defaults to inactivity imply that an equilibrium entry rule prescribes entry if and only if the continuation value from entering and immediately proceeding to the survival subgame is strictly positive. Therefore, each pair of equilibrium continuation value functions $\tilde{v}_S$ and $\tilde{v}_E$ has exactly one corresponding equilibrium strategy. Corollary $2^\star$ requires any equilibrium continuation values to be fixed-points to the Bellman operators used in Algorithm 1. Since these operators are contractions, the continuation values constructed by Algorithm 1 are the only possible equilibrium continuation values. Their corresponding strategy forms the unique symmetric Markov-perfect Nash equilibrium that defaults to inactivity. ■

References

Abbring, J. H. and J. R. Campbell (2010): "Last-In First-Out Oligopoly Dynamics," *Econometrica*, 78, 1491–1527. 1, 2, 9

——— (2013): "Last-In First-Out Oligopoly Dynamics: Corrigendum," Mimeo, Tilburg University and Federal Reserve Bank of Chicago. 2

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. 4, 11, 36, 37

Bajari, P. L., C. L. Benkard, and J. Levin (2007): "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 75, 1331–1370. 3, 21

Bresnahan, T. F. and P. C. Reiss (1990): "Entry in Monopoly Markets," *Review of Economic Studies*, 57, 531–553. 1, 2, 21

——— (1991): "Entry and Competition in Concentrated Markets," *Journal of Political Economy*, 99, 977–1009. 1, 2, 21

——— (1993): "Measuring the Importance of Sunk Costs," *Annales d'Economie et de Statistique*, 31, 181–217. 2

Cabral, L. M. (1993): "Experience Advantages and Entry Dynamics," *Journal of Economic Theory*, 59, 403–416. 36

Campbell, J. R. (2010): "Competition in large markets," *Journal of Applied Econometrics*, 26, 1113–1136. 4

Campbell, J. R. and H. A. Hopenhayn (2005): "Market Size Matters," *Journal of Industrial Economics*, 53, 1–25. 4

Davis, P. (2002): "Estimating multi-way error components models with unbalanced data structures," *Journal of Econometrics*, 106, 67–95. 28, 29

——— (2005): "The Effect of Local Competition on Admission Prices in the U.S. Motion Picture Exhibition Market," *The Journal of Law and Economics*, 48, 677–707. 34

——— (2006): "Spatial Competition in Retail Markets: Movie Theaters," *RAND Journal of Economics*, 37, 964–982. 27, 28, 29, 35

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, Amsterdam: Elsevier Science, vol. 3, chap. 4. 2, 3

Doraszelski, U. and M. Satterthwaite (2010): "Computable Markov-Perfect Industry Dynamics," *RAND Journal of Economics*, 41, 215–243. 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, 2231–2267. 23

Ericson, R. and A. Pakes (1995): "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *Review of Economic Studies*, 62, 53–82. 2, 3, 4, 36

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

Gil, R. and F. LaFontaine (2012): "Using Revenue Sharing to Implement Flexible Prices: Evidence from Movie Exhibition Contracts," *Journal of Industrial Economics*, 60, 187–219. 35

Gowrisankaran, G. (1999): "A dynamic model of endogenous horizontal mergers," *RAND Journal of Economics*, 30, 56–83. 2, 3

Honoré, B. E. and Á. De Paula (2010): "Interdependent Durations," *Review of Economic Studies*, 77, 1138–1163. 5

Hopenhayn, H. A. (1992): "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica*, 60, 1127–1150. 4

Hotz, V. J. and R. A. Miller (1993): "Conditional Choice Probabilities and the Estimation of Dynamic Models," *Review of Economic Studies*, 60, 497–529. 3, 21

Jaimovich, N. (2007): "Firm dynamics and markup variations: Implications for sunspot equilibria and endogenous economic fluctuations," *Journal of Economic Theory*, 137, 300–325. 4

Magnac, T. and D. Thesmar (2002): "Identifying Dynamic Discrete Decision Processes," *Econometrica*, 70, 801–816. 2, 17, 20

Maskin, E. and J. Tirole (1988): "A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs," *Econometrica*, 56, 549–569. 7

McAfee, R. P., H. M. Mialon, and M. A. Williams (2004): "What Is a Barrier to Entry?" *American Economic Review*, 94, 461–465. 40

Melitz, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71, 1695–1725. 4

Newey, W. K. and J. L. Powell (2003): "Instrumental Variable Estimation of Nonparametric Models," *Econometrica*, 71, 1565–1578. 21

Orbach, B. Y. and L. Einav (2007): "Uniform prices for differentiated goods: The case of the movie-theater industry," *International Review of Law and Economics*, 27, 129–153. 35

Rust, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55, 999–1033. 0, 1, 7, 22

Su, C.-L. and K. L. Judd (2012): "Constrained Optimization Approaches to Estimation of Structural Models," *Econometrica*, 80, 2213–2230. 26, 27

Sutton, J. (1991): *Sunk Costs and Market Structure*, MIT Press. 5

Tauchen, G. (1986): "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions," *Economics Letters*, 20, 177–181. 23