

Very Simple Markov-Perfect Industry Dynamics: Empirics

Jaap H. Abbring*

Jeffrey R. Campbell[†]

Jan Tilly[‡]

Nan Yang[§]

March 2017

Abstract

This paper develops an econometric model of firm entry, competition, and exit in oligopolistic markets. The model has an essentially unique symmetric Markov-perfect equilibrium, which can be computed very quickly. We show that its primitives are identified from market-level data on the number of active firms and demand shifters, and we implement a nested fixed point procedure for its estimation. Estimates from County Business Patterns data on U.S. local cinema markets point to tough local competition. Sunk costs make the industry's transition following a permanent demand shock last 10 to 15 years.

*CentER, Department of Econometrics & OR, Tilburg University, and CEPR. 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, nested fixed point estimator, sunk costs, toughness of competition, counterfactual policy analysis.

We thank Ben Chabot for sharing his practitioner's knowledge of the Motion Picture Theater industry. The research of Jaap Abbring is financially supported by the Netherlands Organisation for Scientific Research (NWO) through Vici grant 453-11-002. Nan Yang gratefully acknowledges financial support from National University of Singapore Start-up Grant R-313-000-100-133 and an HSS Faculty Research Fellowship. 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

We develop an econometric model of oligopolists' entry, competition, and exit that can be estimated with readily available data on the numbers of active firms and profit shifters for a panel of independent local markets. Our model features toughness of competition, sunk entry costs, and market-level demand and cost shocks. Its firms have complete information and, within each local market, face identical expected payoffs when making entry and survival decisions.

Each local market in our model is a special case of [Abbring, Campbell, Tilly, and Yang's \(2017\)](#) theoretical model in which the state variables are either observed demand shifters or unobserved cost shocks. [Abbring et al.](#)'s analysis implies that there exists an essentially unique symmetric Markov-perfect equilibrium in each local market, which can be quickly computed by solving for the fixed points of a finite sequence of contraction mappings. We leverage this result to analyze our model's identification, develop a rapid and statistically efficient procedure for its estimation, and facilitate large-scale counterfactual experiments based upon its estimated primitives.

We assume that the econometric errors, which are shocks to firms' costs of entry and continuation, enter firms' profits in an additively separable way and satisfy a conditional independence assumption (as in [Rust, 1987](#)). These assumptions are standard in the empirical analysis of dynamic discrete choice models and games and serve two purposes. First, they simplify the identification of the market's state transitions and the firms' entry and exit decision rules by ensuring that these only depend on observed (by the econometrician) state variables and independent transitory shocks. In particular, the conditional independence assumption excludes persistent unobserved heterogeneity across markets. This allows us to give precise conditions for the identification of the primitives that determine these decision rules: the sunk costs of entry, the toughness of competition, and the (transition) distributions of the demand and cost shocks. Second, they further speed up the equilibrium computations, by allowing the unobserved shocks to be integrated out from the contraction mappings before computing their fixed points on lower dimensional spaces. The resulting algorithm can be embedded in procedures that call it many times. We develop a version of [Rust's \(1987\)](#) nested fixed point (NFXP) maximum likelihood estimation procedure that demonstrates this: It computes an

equilibrium for each local market in the data and each trial value of the parameters and performs well both in a Monte Carlo study and in our empirical application.

Our empirical application characterizes entry, competition, and exit in local markets of the U.S. Motion Picture Theaters industry (NAICS 512131). We first show that our model is suitable for this industry's study. We quantify the influence of theater heterogeneity (measured with mid-March employment) on the evolution of the number of theaters serving *local* markets using County Business Patterns data from 2000 through 2009 from 573 Micropolitan Statistical Areas (μ SAs). Using Poisson regressions, we find that our favored measure of heterogeneity—the Herfindahl-Hirschman Index divided by the value it would attain if producers were homogeneous—has *statistically* significant but *economically* negligible effects on the number of theaters in each local market. The average absolute change in the Poisson regressions' forecasts from omitting heterogeneity is 0.036 theaters compared with an average forecast of 2.46 theaters. In this way, the data make a strong case for abstracting from persistent heterogeneity among theaters within a local market. This is fortunate because our model does so and because high concentration in these local industries precludes applying models that gain comparable simplicity from competition between many firms, like [Hopenhayn's \(1992\)](#) and [Weintraub, Benkard, and Van Roy's \(2008\)](#).

Our estimates imply that adding a single theater to a monopoly market nearly halves the producers' surplus per consumer. Adding two more theaters brings the per consumer surplus to 34 percent of its monopoly value. It follows that cinemas compete fiercely in their local markets, despite earlier evidence that local competition has only small effects on ticket prices ([Davis, 2005](#)). We take this as evidence that movie theaters intensely compete for movie screening rights. The estimated model's sunk costs are substantial, so the initial number of incumbents influences the number of producers for 10 to 15 years following a permanent demand shock. Without sunk costs, producers' dynamic considerations practically vanish and transition to the long run is almost instantaneous. The industry that is composed of *all* local markets in our sample adjusts to permanent demand reductions with both decreased entry and increased exit.

The movie industry is no stranger to large and persistent demand shocks: In the early 1950s, the expansion of television halved movie theater attendance (see e.g. [Takahashi, 2015](#)). Current developments like the advance of internet video

streaming may pose a similar threat to cinemas. Netflix, for example, plans to premiere big movies on its video-on-demand service on the same day that they open in cinemas (Kafka, 2013; Harwell, 2015); and Paramount Pictures intends to make some movies available for home viewing only two weeks after their initial theatrical releases (Schwartzel and Fritz, 2015). With this present relevance as motivation, we investigate whether a policy that limits competition for screening rights could undo the impact of such a change on the long-run average number of firms. This would be reminiscent of the 1970 Newspaper Preservation Act, which sought to maintain media variety by allowing local newspapers to collude under “joint operating agreements.” We find that allowing all theaters to split the monopoly surplus would more than offset the effects of a 25 percent permanent reduction in demand on the number of theaters. A policy that restricts joint operating agreements to duopoly markets would still counter the effects of a 17 percent permanent demand decrease. Such large impacts from changing the toughness of local competition on the number of producers illustrate its economic importance.

Methodologically, our analysis extends Bresnahan and Reiss’s (1990; 1991b) approach to the measurement of the effects of entry on profitability to a dynamic setting. Bresnahan and Reiss (1994) proposed using market-level panel data like ours to estimate sunk costs of entry. However, their empirical analysis of local competition among U.S. dentists was not firmly grounded in theory. Abbring and Campbell (2010) provided a theoretical foundation for the ordered choice models employed by Bresnahan and Reiss (1994) and documented the importance of accounting for uncertainty and sunk costs of entry, even in analyses of static competition. However, they stopped short of developing their model into a framework for econometric analysis. We develop such an econometric framework. Our model’s unique equilibrium involves mixing over survival and exit actions in some states, because its incumbent firms simultaneously decide on survival with complete information. As in Abbring et al. (2017), we leverage the properties of mixed strategy equilibrium—notably that firms earn the value of the outside option, zero, whenever they nontrivially randomize over survival and exit—to simplify the equilibrium analysis and computation. We further simplify equilibrium computation using the econometric assumptions on our model—additive separability and conditional independence.

Because firms might mix over survival and exit in equilibrium, standard

identification arguments and NFXP estimation procedures for single-agent dynamic discrete choice models need nontrivial modifications. After all, if firms mix over survival and exit actions, the number of firms that serve a local market does not only depend on the observed state variables and the unobserved cost shock, as in a single agent model, but also on the outcome of equilibrium mixing. Moreover, the equilibrium mixing probabilities depend not only on observed state variables, but also on the unobserved cost shocks.

For identification, we invert probabilities of observing particular market transitions to learn about the equilibrium payoffs that drive firms' decisions. This is reminiscent of [Magnac and Thesmar's \(2002\)](#) approach to the identification of single-agent dynamic discrete choice models, which inverts choice probabilities to identify value contrasts ([Hotz and Miller, 1993](#)). Mixing over survival and exit, however, complicates the inversion. We show that the familiar choice probability inversion arguments can nevertheless be adapted to our model. Moreover, we show that the equilibrium relation between the mixing probabilities and the unobserved cost shocks can be used to identify a scale parameter of the cost shock distribution, whereas this distribution needs to be fixed *a priori* in a single agent setting ([Magnac and Thesmar, 2002](#)).

Our NFXP procedure evaluates the likelihood function for each trial value of the parameters. After calculating an equilibrium for each local market in the data, which the computational result of [Abbring et al. \(2017\)](#) makes straightforward, it constructs the corresponding likelihood. Unlike “two-step” estimators such as [Bajari, Benkard, and Levin's \(2007\)](#) and [Pakes, Ostrovsky, and Berry's \(2007\)](#), our maximum likelihood estimator is statistically efficient, has well-known asymptotic properties, and can easily incorporate local market heterogeneity.

The remainder of the paper proceeds as follows. Section 2 presents the model for a single local market and discusses its equilibrium existence, uniqueness, and computation. Section 3 develops this model's empirical implementation with panel data from independent, but not necessarily identical, local markets. We discuss sampling, identification, likelihood construction, and maximum likelihood estimation using the NFXP procedure. Section 4 applies the resulting empirical framework to the Motion Picture Theaters industry. Section 5 concludes by comparing our approach to one using two-step estimators for games of incomplete information.

2 The Model

For each local market, we specify a special case of [Abbring et al.’s \(2017\)](#) theoretical model. This section presents this model and discusses its equilibrium uniqueness and computation. [Appendix A](#) further details the links with [Abbring et al.’s](#) analysis.

2.1 Primitives

Time is discrete and indexed by $t \in \mathbb{N} \equiv \{1, 2, \dots\}$. In period t , firms that have entered in the past and not yet exited serve the market. Each firm has a name $f \in \mathcal{F} \equiv \mathcal{F}_0 \cup (\mathbb{N} \times \{1, 2, \dots, \check{j}\})$. Initial incumbents have distinct names in \mathcal{F}_0 , while potential entrants’ names are from $\mathbb{N} \times \{1, 2, \dots, \check{j}\}$. The first component of a potential entrant’s name gives the period in which it has its only opportunity to enter the market, and the second component gives its position in that period’s queue of $\check{j} < \infty$ firms. Aside from the timing of their entry opportunities, the firms are identical.

[Figure 1](#) shows the actions taken by firms in period t and their consequences for the game’s state at the start of period $t + 1$. This is the game’s *recursive extensive form*. We divide each period into two subperiods, the entry and survival stages. Period t begins on the left with the entry stage. If $t = 1$, nature sets the number N_1 of firms serving the market in period 1 and the initial demand state C_1 . If $t > 1$, these are inherited from the previous period. We assume that C_t follows a first-order Markov process and denote its support with \mathcal{C} . Throughout the paper, we refer to C_t as “demand,” but it can encompass any observed, relevant, and time-varying characteristics of the market, depending on the empirical context. In [Section 4](#)’s empirical application, C_t is the local market’s residential population.

Each incumbent firm serves the market and earns a surplus $\pi(N_t, C_t)$. We assume that

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

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

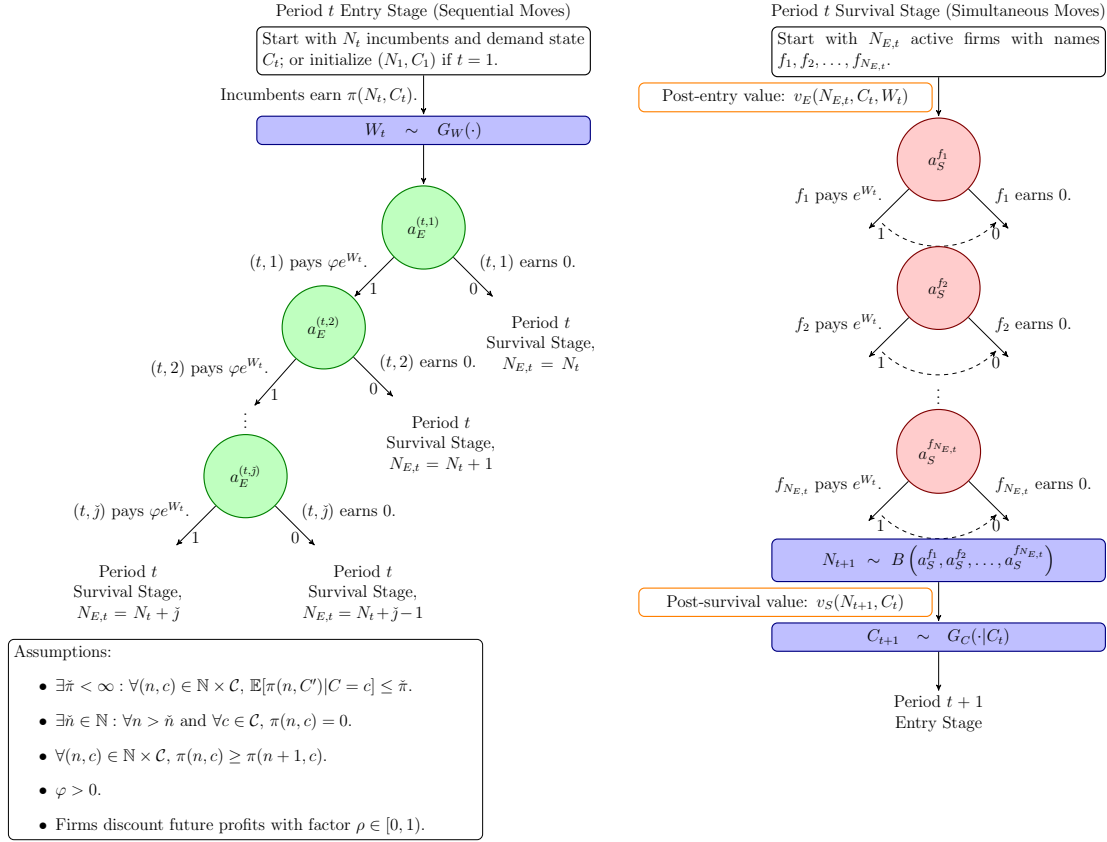


Figure 1: The Model's Recursive Extensive Form

corresponding small Roman letters. The first assumption is technical and allows us to restrict equilibrium values to the space of bounded functions. The second assumption allows us to restrict equilibrium analysis to markets with at most \bar{n} firms. It is not restrictive in empirical applications to oligopolies. The third assumption requires the addition of a competitor to reduce weakly each incumbent's surplus.

After incumbents earn their surpluses, nature draws the current period's shock to continuation and entry costs, W_t , from a distribution G_W with positive density everywhere on the real line. Then, period t 's cohort of potential entrants $\{t\} \times \mathbb{N}$ make entry decisions in the order of the second component of their names. We denote firm f 's entry decision with $a_E^f \in \{0, 1\}$. An entrant ($a_E^f = 1$) pays the sunk cost $\varphi \exp(W_t)$, with $\varphi > 0$. A firm choosing not to enter ($a_E^f = 0$) earns a payoff of zero and never has another entry opportunity. Such a refusal to enter also ends the entry stage, so firms remaining in this period's entry cohort that have not yet had an opportunity to enter *never* get to do so.

We denote the number of firms in the market after the entry stage, the sum of the incumbents and the actual entrants, with $N_{E,t}$. Suppose that the names of these active firms are $f_1, \dots, f_{N_{E,t}}$. In the subsequent survival stage, they simultaneously decide on continuation with probabilities $a_S^{f_1}, \dots, a_S^{f_{N_{E,t}}} \in [0, 1]$. After these decisions, all survival outcomes are realized independently across firms according to the chosen Bernoulli distributions.¹ Firms that survive pay a fixed cost $\exp(W_t)$. A firm can avoid this cost by exiting to earn zero. Firms that have exited cannot reenter the market later. The N_{t+1} surviving firms continue to the next period, $t + 1$. The period ends with nature drawing a new demand state C_{t+1} from the conditional distribution $G_C(\cdot | C_t)$. All firms discount future profits and costs with the discount factor $\rho \in [0, 1)$.

In Section 3, 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 observed demand state C_t and the market structure N_t statistically nondegenerate. [Bresnahan and Reiss \(1991a\)](#) noted that equilibrium multiplicity can arise in static games with econometric errors that have complete support and are at least somewhat independent across both players and outcomes. Our specification of a single shock to all firms’ continuation and entry costs avoids this difficulty. The assumptions on $\{C_t, W_t\}$ make it a first-order Markov process satisfying [Rust’s \(1987\)](#) conditional independence assumption.² This ensures that the distribution of (N_t, C_t) conditional on (N_{t^*}, C_{t^*}) for all $t^* < t$ depends only on (N_{t-1}, C_{t-1}) , so we require only the model’s transition rules to calculate the conditional likelihood function.

¹The assumption that entrants immediately contemplate exit might seem strange, but exit immediately following entry never occurs in equilibrium. Furthermore, this timing assumption removes an unrealistic possibility. If entrants did not make these continuation decisions, then they could effectively commit to continuation. This would allow an entrant to displace an incumbent only by virtue of this commitment power. See [Abbring et al. \(2017\)](#) for further discussion of these timing assumptions.

²[Rust \(1987\)](#) defined “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.

2.2 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. A Markov strategy maps payoff-relevant states into actions. When a potential entrant (t, j) makes its entry decision in period t , the payoff-relevant states are $M_t^j \equiv N_t + j$, the current demand C_t , and the cost shock W_t . Here, M_t^j is the number of firms that would be committed to serve the market in period $t + 1$ if firm (t, j) would decide to enter. We collect these into the vector $(M_t^j, C_t, W_t) \in \mathcal{H} \equiv \mathbb{N} \times \mathcal{C} \times \mathbb{R}$. Similarly, we collect the payoff-relevant state variables of a firm f contemplating survival in period t in the \mathcal{H} -valued $(N_{E,t}, C_t, W_t)$. Since survival decisions are made simultaneously, this state is the same for all active firms. A Markov strategy is a pair of functions $a_E : \mathcal{H} \rightarrow \{0, 1\}$ and $a_S : \mathcal{H} \rightarrow [0, 1]$. The *entry rule* a_E assigns a binary indicator of entry to each possible state. Similarly, a_S gives a *survival probability* for each possible state. Since time and firms' names themselves are not payoff-relevant, we henceforth drop the subscript t and the superscript j from the payoff-relevant states.

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 n_E firms, demand state c , and cost shock w 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 stage where these value functions apply.

A firm's post-survival value equals the expected sum of the profit and post-entry value that accrue to the firm in the next period, discounted to the current period with ρ :

$$v_S(n', c) = \rho \mathbb{E}_{a_E} [\pi(n', C') + v_E(N'_E, C', W') | N' = n', C = c]. \quad (1)$$

Here, \mathbb{E}_{a_E} is an expectation over the next period's demand state C' , cost shock W' , and post-entry number of firms N'_E . This expectation operator's subscript indicates its dependence on a_E .³ In particular, given $N' = n'$, N'_E is a deterministic function of $a_E(\cdot, C', W')$. Because the payoff from leaving the market is zero, a firm's post-entry value in a state (n_E, c, w) equals the probability that it survives, $a_S(n_E, c, w)$, times the expected payoff from surviving:

$$\begin{aligned} v_E(n_E, c, w) & \\ &= a_S(n_E, c, w) \left(\mathbb{E}_{a_S} [v_S(N', c) | N_E = n_E, C = c, W = w] - \exp(w) \right). \end{aligned} \quad (2)$$

The expectation \mathbb{E}_{a_S} over N' takes survival of the firm of interest as given. That is, it takes N' to equal one plus the outcome of $n_E - 1$ independent Bernoulli (survival) trials with success probability $a_S(n_E, c, w)$. 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.

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 from its prescriptions:⁴

$$\begin{aligned} a_E(m, c, w) &\in \operatorname{argmax}_{a \in \{0,1\}} a \left(\mathbb{E}_{a_E} [v_E(N_E, c, w) | M = m, C = c, W = w] - \varphi \exp(w) \right), \\ a_S(n_E, c, w) &\in \operatorname{argmax}_{a \in [0,1]} a \left(\mathbb{E}_{a_S} [v_S(N', c) | N_E = n_E, C = c, W = w] - \exp(w) \right). \end{aligned}$$

Abbring et al. (2017) note that it might be possible to construct one symmetric Markov-perfect equilibrium from another by changing a single firm's entry or continuation decision when that firm is indifferent between its available actions. We follow their approach to eliminating this possibility by restricting attention to equilibria that *default to inactivity*. In such equilibria, a potential entrant

³The assumptions on π ensure that $0 \leq \mathbb{E}_{a_E} [\pi(n', C') | N' = n', C = c] = \mathbb{E}[\pi(n', C') | C = c] \leq \tilde{\pi}$ and that v_E is bounded from above. Moreover, optimal exit behavior ensures that $v_E \geq 0$. Thus, the expectations in (1) and (2) are well defined and v_S is bounded. See Appendix A.

⁴Because the cost shock W can be arbitrarily high, firms' flow payoffs are not bounded from below. Therefore, it is not immediately obvious a strategy profile forms a subgame perfect equilibrium whenever no firm can gain from a one-shot deviation. For example, Theorem 4.2 in **Fudenberg and Tirole (1991)** does not immediately apply to our game. In the online supplement to **Abbring et al. (2017)**, we used the existence of the outside option with a fixed payoff to show that a strategy profile that is one-shot deviation proof does indeed form a subgame-perfect equilibrium.

that is indifferent between entering or not stays out, and an active firm that is indifferent between *all* possible outcomes of the survival stage exits.⁵ The analysis in [Abbring et al. \(2017\)](#) implies that our model has a unique symmetric Markov-perfect equilibrium that defaults to inactivity, with the following properties (see [Appendix A](#)).

1. There will be no entry in markets with \tilde{n} or more active firms, so that we can limit its analysis to states with \tilde{n} or fewer firms.⁶ Intuitively, this follows from the assumption that firms always make negative profits in markets with more than \tilde{n} active firms. If $\check{j} > \tilde{n}$, then at least one potential entrant chooses not to enter every period. In this sense, the model becomes one of free entry. [Abbring et al. \(2017\)](#) impose this free-entry condition, and we follow them.
2. The post-survival value $v_S(n', c)$ weakly decreases with n' . This implies that $a_S(n_E, c, w) = 0$ if $v_S(1, c) \leq \exp(w)$; $a_S(n_E, c, w)$ equals the unique survival probability $a \in (0, 1]$ that makes firms indifferent between exit and survival,

$$0 = -\exp(w) + \sum_{n'=1}^{n_E} \binom{n_E-1}{n'-1} a^{n'-1} (1-a)^{n_E-n'} v_S(n', c), \quad (3)$$

if $v_S(n_E, c) \leq \exp(w) < v_S(1, c)$; and $a_S(n_E, c, w) = 1$ if $v_S(n_E, c) > \exp(w)$.⁷ Moreover, because firms continue and collect the payoff $-\exp(w) + v_S(n_E, c)$ whenever it is positive, and receive a zero payoff otherwise, (2) simplifies to

$$v_E(n_E, c, w) = \max\{0, -\exp(w) + v_S(n_E, c)\}. \quad (4)$$

Thus, the post-entry value in a state (n_E, c, w) can be computed from the post-survival value in state (n_E, c) without knowing the post-survival values that would be obtained after the exit of one or more competitors. This result is key to our recursive procedure for computing the equilibrium values.

3. Equation (4) and the fact that $v_S(n_E, c)$ weakly decreases with n_E imply

⁵The restriction to equilibria that default to inactivity is innocuous in this paper's context. We will assume that W follows a continuous distribution, so that an exact indifference between activity and inactivity occurs with probability zero.

⁶If $N_1 \leq \tilde{n}$, the equilibrium number of active firms never exceeds \tilde{n} ; otherwise, firms leave with positive probability until the number firms is no larger than \tilde{n} .

⁷In (3), we use the convention that $0^0 \equiv 1$, so $a = 1$ if $v_S(n_E, c) = \exp(w)$.

that $v_E(n_E, c)$ weakly decreases with n_E . This ensures that $a_E(m, c, w) = \mathbb{1} [v_E(m, c, w) > \varphi \exp(w)]$, which, with (4), further simplifies to $a_E(m, c, w) = \mathbb{1} [v_S(m, c) > (1 + \varphi) \exp(w)]$. Here, $\mathbb{1} [\cdot] = 1$ if its argument is true and equals 0 otherwise.

2.3 Computation

Abbring et al. (2017) provided an algorithm for equilibrium computation which exploits equation (4) to represent equilibrium post-entry values as solutions to a sequence of dynamic programming problems. The relevant Bellman equations are

$$v_E(n, c, w) = \max \{0, -\exp(w) + \rho \mathbb{E}_{a_E} [\pi(n, C') + v_E(N'_E, C', W') \mid N_E = n, C = c]\} \quad (5)$$

for $n = \tilde{n}, \tilde{n} - 1, \dots, 1$. In the case with $n = \tilde{n}$, no additional firms enter in equilibrium. Therefore, only $v_E(\tilde{n}, c, w)$ appears on the right-hand side of (5). With this, Abbring et al.'s (2017) algorithm calculates the only possible equilibrium post-entry value for each of \tilde{n} incumbent oligopolists. This in turn determines the only possible equilibrium entry rule that defaults to inactivity, $a_E(\tilde{n}, c, w) = \mathbb{1} [v_E(\tilde{n}, c, w) > \varphi \exp(w)]$. Proceeding to $n = \tilde{n} - 1$, the right-hand side of (5) includes $v_E(\tilde{n}, c, w)$, $a_E(\tilde{n}, c, w)$, and $v_E(\tilde{n} - 1, c, w)$. The first two of these are known, so Abbring et al.'s (2017) algorithm can use (5) to compute the only possible values of $v_E(\tilde{n} - 1, c, w)$ and $a_E(\tilde{n} - 1, c, w)$. Their algorithm proceeds recursively to calculate all of the post-entry values and entry rules, and it finishes by computing the corresponding post-survival values and equilibrium survival rules.

Using the special structure of this paper's empirical model, we modify Abbring et al.'s (2017) algorithm to make the computation less taxing. Specifically, we recursively compute the post-survival value v_S instead of the post-entry value v_E and thereby remove one dimension from calculated value functions. Figure 2 presents the resulting algorithm in detail. Its recursive portion begins with the calculation of $v_S(\tilde{n}, c)$. This satisfies

$$v_S(\tilde{n}, c) = \rho \mathbb{E} \left[\pi(\tilde{n}, C') + \int_{-\infty}^{\log v_S(\tilde{n}, C')} (-\exp(w) + v_S(\tilde{n}, C')) dG_W(w) \mid C = c \right].$$




Figure 2: Equilibrium Calculation

With $v_S(\check{n}, c)$ in hand, we can represent $a_E(\check{n}, c, w)$ with a cost-shock threshold,

$$\bar{w}_E(\check{n}, c) \equiv \log v_S(\check{n}, c) - \log(1 + \varphi).$$

A firm contemplating entry into a market with $\check{n} - 1$ incumbents does so in equilibrium if and only if $w < \bar{w}_E(\check{n}, c)$. With this completed, the algorithm's loop proceeds through $n = \check{n} - 1, \dots, 1$ calculating $v_S(n, c)$ and the entry threshold

$$\bar{w}_E(n, c) \equiv \log v_S(n, c) - \log(1 + \varphi) \tag{6}$$

recursively. For a generic n , the Bellman operator used in the n 'th pass through the loop is

$$\begin{aligned} T_n(f)(\cdot) = & \rho \mathbb{E} \left[\pi(n, C') + \int_{\bar{w}_E(n+1, C')}^{\log f(C')} (-\exp(w) + f(C')) dG_W(w) \right. \\ & \left. + \sum_{n'=n+1}^{\check{n}} \int_{\bar{w}_E(n'+1, C')}^{\bar{w}_E(n', C')} (-\exp(w) + v_S(n', C')) dG_W(w) \middle| C = \cdot \right]. \end{aligned} \tag{7}$$

When the algorithm's recursive portion is complete, it proceeds to the calculation of the equilibrium survival rule a_S and the post-entry value v_E . By construction, this algorithm's output is identical to that of [Abbring et al.'s \(2017\) Algorithm 1](#), so their Theorem 1 establishes that it is the unique symmetric Markov-perfect equilibrium that defaults to inactivity.

3 Empirical Implementation

The previous section provided an algorithm to compute the unique symmetric Markov-perfect equilibrium for given primitives π , 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 observations of this process from a market panel data can be used to estimate the model's primitives.

3.1 Sampling

Suppose that we have data from \check{r} markets indexed with $r = 1, \dots, \check{r}$. For each market, we observe the number of active firms $N_{r,t}$ and the demand state $C_{r,t}$ in

each period $t = 1, \dots, \check{t}$; 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 $W_{r,t}$.⁸

We assume that $(\{N_{r,t}, C_{r,t}; t = 1, \dots, \check{t}\}, X_r)$ is distributed independently across markets.⁹ The initial conditions $(N_{r,1}, C_{r,1}, X_r)$ are drawn from a distribution that we leave unspecified. Thereafter, industry dynamics follow the transition rules implied by the unique equilibrium of our model, with primitives $\pi_r(\cdot, \cdot) = \pi(\cdot, \cdot | X_r, \theta_P)$, $\varphi_r = \varphi(X_r, \theta_P)$, and $\rho_r = \rho(X_r, \theta_P)$ for some finite vector θ_P ; $G_{C,r}(\cdot | \cdot) = G_C(\cdot | \cdot; X_r, \theta_C)$ for some finite vector θ_C ; and $G_{W,r}(\cdot) = G_W(\cdot; X_r, \theta_W)$ for some finite vector θ_W .¹⁰ We collect the model's structural parameters in a vector $\theta \equiv (\theta_P, \theta_C, \theta_W)$.

3.2 Identification

In this section, we analyze the extent to which we can determine θ when we observe the population $(\{N_t, C_t; t = 1, \dots, \check{t}\}, X)$ underlying our data. Specifically, suppose that we know the distribution of (N', C') conditional on $(N, C, X) = (n, c, x)$ for all $n \in \{0\} \cup \mathbb{N}$, $c \in \mathcal{C}$, and a specific value x of the market characteristics.¹¹ Throughout the remainder of this section, we keep the conditioning information $X = x$ 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 .¹² The

⁸Our estimation does not utilize data on firms' input choices, sales volumes, costs, revenues, or profits. Such data could be used to directly quantify certain model primitives (for example, by equating profits to those from Cournot competition with an estimated demand curve and constant marginal costs) before estimating its other primitives using the procedure outlined in this section.

⁹Our estimation procedure can be extended to allow for observed (to the econometrician) time-varying covariates that are common across markets, such as business cycle indicators, provided that firms can use the model's primitives to forecast their evolution.

¹⁰These 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.

¹¹For 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.

¹²Above, 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.

remaining primitives of interest are the model's sunk cost φ , surplus function π , and the distribution G_W of the econometric error. Our identification argument for these parameters builds upon 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 some knowledge of G_W , we assume that this belongs to the parametric family

$$G_W(w) = \Phi\left(\frac{w + \omega^2/2}{\omega}\right), \quad (8)$$

where Φ is 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 of producers give us no information on the level of profits, we do not attempt to separately identify the location parameter of this distribution.¹³

Analogously to the entry threshold $\bar{w}_E(n, c)$ that we defined in (6), we define a cost-shock threshold for *sure* survival, $\bar{w}_S(n, c) \equiv \log v_S(n, c)$. A firm deciding on continuation in state (n_E, c, w) will survive for sure if $w < \bar{w}_S(n, c)$, exit with positive probability if $w > \bar{w}_S(n, c)$, and exit for sure if $w \geq \bar{w}_S(1, c)$.

We can retrieve $\bar{w}_S(1, c)$ (a monopolist's survival threshold), up to the unknown scale and shift in G_W , from the probability of a monopolist surviving:

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

Similarly, we can recover $\bar{w}_E(n, c)$ (the threshold for entry as the n th active firm) from the probability of at least n firms entering a previously empty market:

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

These and the definitions of $\bar{w}_S(1, c)$ and $\bar{w}_E(1, c)$ can be used to identify the sunk

¹³This 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.

cost of entry up to the scale parameter ω :

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

In turn, this allows us to retrieve

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

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 $c^* \in \mathcal{C}$ and $n^* \in \{2, \dots, \tilde{n}\}$,

$$\bar{w}_S(1, c^*) = \dots = \bar{w}_S(n^* - 1, c^*) > \bar{w}_S(n^*, c^*).$$

This is equivalent to requiring that

$$v_S(1, c^*) = \dots = v_S(n^* - 1, c^*) > v_S(n^*, c^*)$$

for some c^* and n^* . This is a very weak condition, particularly given that we have already established 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^* incumbents simultaneously exiting:

$$\begin{aligned} &\Pr[N' = 0 | N = n^*, C = c^*] \\ &= \Pr[W \geq \bar{w}_S(1, c^*)] + \int_{\bar{w}_S(n^*, c^*)}^{\bar{w}_S(1, c^*)} [1 - a_S(n^*, c^*, w)]^{n^*} dG_W(w) \\ &= \Pr[N' = 0 | N = 1, C = c^*] + \int_{\bar{w}_S(n^*, c^*)}^{\bar{w}_S(1, c^*)} [1 - a_S(n^*, c^*, w)]^{n^*} dG_W(w). \quad (9) \end{aligned}$$

Because the two transition probabilities in (9) 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 $v_S(1, c^*) = \dots = v_S(n^* - 1, c^*)$, we can explicitly solve for the mixing probability $a_S(n^*, c^*, w)$:

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

Rewrite the integral on the right-hand side of (9) by substituting this expression for $a_S(n^*, c^*, w)$, 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

$$\int_{k_{n^*}}^{k_1} \left[1 - \left(\frac{\exp(\omega k_1) - \exp(\omega \varepsilon)}{\exp(\omega k_1) - \exp(\omega k_{n^*})} \right)^{\frac{1}{n^*-1}} \right]^{n^*} \phi(\varepsilon) d\varepsilon, \quad (10)$$

with

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

Because k_1 and k_{n^*} are known,

$$\frac{\exp(\omega k_1) - \exp(\omega \varepsilon)}{\exp(\omega k_1) - \exp(\omega k_{n^*})} = \frac{1 - \exp(-\omega(k_1 - \varepsilon))}{1 - \exp(-\omega(k_1 - k_{n^*}))} \quad (11)$$

is a known function of ω . Moreover, it is straightforward to verify that it is strictly increasing in ω for $\varepsilon \in (k_{n^*}, k_1)$. Hence, the integrand in (10) is a known, strictly decreasing function of ω . Because the domain of integration of the integral in (10) 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 use variation in C that impacts next period's expected post-entry value but *not* next period's expected surplus to identify

ρ .¹⁴ Specifically, 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 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 the relevant Bellman equation. We summarize these results in a theorem.

Theorem 1 *Suppose that ρ is known and that G_W is specified up to scale as in (8). Furthermore, suppose that, for some $c^* \in \mathcal{C}$ and $n^* \in \{2, \dots, \tilde{n}\}$,*

$$\begin{aligned} \Pr[N' = 0|N = 1, C = c^*] &= \dots = \Pr[N' = 0|N = n^* - 1, C = c^*] \\ &< \Pr[N' = 0|N = n^*, C = c^*]. \end{aligned}$$

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

To emphasize that it can be verified in data, we have rewritten the required condition on the sure survival thresholds in terms of known probabilities. The equivalence between the two sets of conditions follows from fact that the integral on the right-hand side of (9) has a positive integrand and so equals zero if and only if its limits of integration equal each other. That is, if and only if $w_S(n^*, c^*) = w_S(1, c^*)$.

We only establish 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

¹⁴Magnac and Thesmar (2002) formalized the use of such exclusion restrictions to identify the discount factor in dynamic discrete choice models. They focused on high level restrictions on a particular value contrast, the “current value.” Abbring and Daljord (2016) explored the identifying power of exclusion restrictions on primitive utility, such as the per period surplus in our model.

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 analysis. First, it is possible to identify the model’s parameters without examining the cross-sectional relationship between N and C that [Bresnahan and Reiss \(1990, 1991b\)](#) used in their estimation. Second, estimation of our model need not follow the NFXP 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. Because our model lacks firm-specific private shocks, this deduction would substantially differ from that pioneered by [Bajari, Benkard, and Levin \(2007\)](#). We choose to estimate our primitives using NFXP maximum likelihood because it avoids the nonparametric estimation of transition probabilities that is required for a [Hotz and Miller](#) style estimator (which does not calculate the decision rules from primitives) and it achieves statistical efficiency. 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), 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 such restrictions on payoffs for estimation with a finite sample.

3.3 Likelihood

We now focus on inferring the structural parameters θ from the conditional likelihood $\mathcal{L}(\theta)$ of θ for data on market dynamics $\{N_{r,t}, C_{r,t}; t = 2, \dots, \check{t}; r = 1, \dots, \check{r}\}$ given the initial conditions $(N_{r,1}, C_{r,1}, X_r; r = 1, \dots, \check{r})$.¹⁵ Using the model’s Markov structure

¹⁵We 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

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

$$\mathcal{L}_C(\theta_C) \equiv \prod_{r=1}^{\tilde{r}} \prod_{t=1}^{\tilde{t}-1} g_C(C_{r,t+1} | C_{r,t}; X_r, \theta_C),$$

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

$$\mathcal{L}_N(\theta) \equiv \prod_{r=1}^{\tilde{r}} \prod_{t=1}^{\tilde{t}-1} p(N_{r,t+1} | N_{r,t}, C_{r,t}; X_r, \theta),$$

the conditional likelihood of θ for the evolution of the market structures.¹⁶ Here, $g_C(\cdot | \cdot; X_r, \theta_C)$ is the density of $G_{C,r}$ and $p(n'|n, c; X_r, \theta) \equiv \Pr(N_{r,t+1} = n' | N_{r,t} = n, C_{r,t} = c; X_r, \theta)$ 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 $p(\cdot | \cdot; X_r, \theta)$ for each distinct value of X_r in the sample. To this end, 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, $\bar{w}_{E,r}(n, c) \equiv \log v_{S,r}(n, c) - \log(1 + \varphi_r)$ and $\bar{w}_{S,r}(n, c) \equiv \log v_{S,r}(n, c)$.

For $n' > n$, $p(n'|n, c; X_r, \theta)$ can easily be calculated as the probability that $W_{r,t}$ falls into $[\bar{w}_{E,r}(n' + 1, c), \bar{w}_{E,r}(n', c))$. For $n' \leq n$, the computations are complicated by the equilibrium mixing of survival decisions. For example, the number of firms can remain unchanged either because survival is a dominant action or because firms choose to exit with positive probability but by chance they all survive. Therefore, the probability that $n' = n > 0$ sums the probability that $W_{r,t}$ falls into $[\bar{w}_{E,r}(n + 1, c), \bar{w}_{S,r}(n, c))$ (so that survival is a dominant action but no entry occurs) with the probability that it instead equals some $w \in [\bar{w}_{S,r}(n, c), \bar{w}_{S,r}(1, c))$ (so that incumbents mix exit and survival) and that all n firms survive when they

efficient estimator, at the price of robustness. Moreover, it would allow us to deal with the initial conditions problems mentioned in Footnote 10.

¹⁶As in [Ericson and Pakes \(1995\)](#), firms begin to earn profits in the period *after* their entry decisions. Since $N_{r,t+1}$ is determined before the realization of $C_{r,t+1}$, its conditional distribution depends only on $C_{r,t}$.

mix with probability $a_S(n, c, w)$.

Similar complications arise for $n' \in \{1, \dots, n - 1\}$, which can only occur if firms nontrivially mix exit and survival, 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 $p(n'|n, c; X_r, \theta)$ is tedious but straightforward. The interested reader should consult Appendix B.

3.4 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 $\tilde{\theta}_C \equiv \operatorname{argmax}_{\theta_C} \mathcal{L}_C(\theta_C)$;
2. estimate (θ_P, θ_W) with $(\tilde{\theta}_P, \tilde{\theta}_W) \equiv \operatorname{argmax}_{(\theta_P, \theta_W)} \mathcal{L}_N(\theta_P, \tilde{\theta}_C, \theta_W)$; and
3. estimate θ by maximizing the full likelihood function $\hat{\theta} \equiv \operatorname{argmax}_{\theta} \mathcal{L}(\theta)$, using $\tilde{\theta} \equiv (\tilde{\theta}_P, \tilde{\theta}_C, \tilde{\theta}_W)$ as a starting value for the chosen optimization routine.

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 \tilde{r} is large and \tilde{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. We use a standard non-linear, gradient-based optimizer to perform the optimization, and we compute all gradients analytically. The Matlab code provides a more user friendly implementation of the NFXP procedure with neither covariates nor analytical gradients that can be used as a sandbox for experimentation and teaching.

In Appendix C, we report the results of Monte Carlo experiments that estimate the model's parameters given data generated by the model itself. Using these same "data," we also estimate our model using Su and Judd's (2012) mathematical programming with equilibrium constraints (MPEC) procedure. Those experiments lead us to four conclusions. First, the NFXP estimator can distinguish between economically meaningful hypotheses with observations from as few as 250 markets

over 10 years. Second, asymptotic distribution theory gives a good guide to standard confidence intervals' coverage probabilities with such small samples. Third, the NFXP estimator takes very little time to compute. The average estimation time across 1,000 experiments was only 21 seconds. Thus, our estimator's computational burden is not substantially greater than that of static models of long-run industry structure, a point emphasized by [Pakes, Ostrovsky, and Berry \(2007\)](#) regarding their two-step estimators. Fourth and finally, the MPEC procedure always calculated estimates practically indistinguishable from those of our NFXP estimator but was 40 times slower. Thus, it appears that our estimator passes the initial quality assurance test of being accurate and relatively easy to compute when applied to simulated data.

4 Motion Picture Theaters

In this section we apply the model and its estimator to an empirical analysis of the Motion Picture Theaters industry. [Davis \(2006\)](#) showed that theater locations substantially influence consumers' decisions about whether and where to attend film screenings. Indeed, one's probability of attending a given theater declines considerably when the travel distance moves from between zero and five miles to between five and ten miles.¹⁷ [Davis \(2002\)](#) found 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](#) established 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 has two implications for our empirical analysis. First, theaters that are sufficiently far apart plausibly operate in distinct markets. Consequently, in [Section 4.1](#), we define such markets using readily available geographic data. Second, variation across markets in the within-market spatial structure of demand is likely to come with variation in their profits and toughness of competition. To capture this, we include in X_r a measure of the diversity of consumers' geographic preferences.

The empirical analysis proceeds in four steps. First, we describe the data we use for the estimation. Second, we conduct an empirical examination of whether

¹⁷See the logit model estimates reported in [Table 5 of Davis \(2006\)](#).

the model’s abstraction from persistent heterogeneity across producers makes it unsuitable for application to these data. We conclude that it does not. Third, we present estimates of the model’s parameters and discuss their implications for the toughness of competition between theater owners for screening rights. Fourth, we present several simulations of the model that highlight the long-lasting impact of initial conditions, illustrate the importance of sunk costs in determining the length of “short-run” transitions to the long run, and quantify how lenient antitrust policy can offset permanent negative demand shocks.

4.1 The Data

Our analysis equates a market with a Micropolitan Statistical Area (μ SA) 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.¹⁸ We dropped the μ SA “The Villages, FL,” because its population growth far exceeds that of any other μ SA. The remaining 573 μ SAs account for about ten percent of the United States population. We measured the diversity of the geographic preferences of each μ SA’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.¹⁹ We then measured *geographic preference diversity* with the average distance between two randomly-chosen residents of the μ SA. Likewise, we can measure the average distance between two randomly chosen individuals from two distinct μ SAs. By construction, μ SAs are geographically isolated from larger Metropolitan Statistical Areas, so we measure a given μ SA’s *geographic market isolation* as the shortest such distance to another μ SA.

For the 573 μ SAs, Table 1 displays the five standard quantiles for population,

¹⁸We 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, which includes information on 574 μ SAs.

¹⁹For 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 1: Summary Statistics for μ SAs

	Quantile				
	10	25	50	75	90
Population	23.51	32.57	42.67	62.32	87.71
Median Household Income	27.40	30.40	33.52	38.04	42.40
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 μ SAs 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.

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 \$41,990 in 2000. This is considerably higher than the median value across the μ SAs, \$33,520. More than 80 percent of the μ SAs 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 μ SAs 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 μ SA should be much larger than those across μ SAs. Its median value across μ SAs equals 37.61 miles. Indeed, there are only eight μ SAs where this distance is less than twenty miles. We conclude that the μ SAs are isolated enough from each other so that consumer 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, except for drive-in theaters. Our estimation uses annual counts of the number of *theaters* (not firms) in each μ SA from the County Business Patterns (CBP), beginning in 2000 and ending in 2009. The top panel of [Table 2](#) reports the

Table 2: Frequencies and Transition Rates from the County Business Patterns

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

% of Transitions Given N_{t-1}					
$\downarrow N_{t-1}/N_t \rightarrow$	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 μ SA from 2000 to 2009 from the County Business Patterns for the 573 μ SAs in our sample. The bottom panel displays the conditional probability of transitioning from N_{t-1} movie theaters in a μ SA at time $t - 1$ (row) to N_t theaters at time t (column).

frequencies of the number of theaters across all of the μ SA-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 2's bottom panel 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 μ SA 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.

4.2 Heterogeneity and Industry Dynamics

Since the pioneering empirical work of [Dunne, Roberts, and Samuelson \(1988\)](#), a rich literature has arisen that documents high producer turnover within narrowly-defined industries and measures its role in aggregate productivity growth. Our analysis focuses on a very different (and somewhat older) question: How does changing market size change the number of producers and the profits they earn? This question does not obviously require consideration of producer heterogeneity, and indeed both [Bresnahan and Reiss \(1990\)](#) and [Campbell and Hopenhayn \(2005\)](#) examined it in the context of models with homogeneous firms. However, just because it is convenient to treat all producers symmetrically does not mean it is appropriate to do so. Persistent differences between firms might substantially impact their entry and continuation decisions and thereby bring the results of applying our model into question. At the same time, the simple observation that producers differ on some observable dimensions does not invalidate an approach like ours that treats them symmetrically. Our model only requires that producers' *expected* profits are identical, because expected profits govern entry and exit decisions. This leaves a great deal of room for incorporating observable transitory heterogeneity across producers. For example, the theaters in our data might always have very different sizes (measured with e.g. sales) if the (randomly-chosen) theater with the highest-quality film (the “blockbuster”) attracts all quality-sensitive consumers and the remaining quality-insensitive consumers split themselves equally across theaters. As long as blockbusters do not systematically go to one particular theater, such within-period heterogeneity is irrelevant for producers' entry and continuation decisions.

Whether abstracting from producer heterogeneity within a local market is a helpful simplification or a fatal error is ultimately an empirical question. The common observation of high producer turnover at the national level might seem to settle that question, since our model trivially predicts that entry and exit never occur simultaneously. However, only turnover *within* μ SAs is relevant for the question of

whether the model can be usefully applied to our sample of local markets.²⁰ We seek to settle this question for the theaters in our data set by examining directly how producer heterogeneity influences the evolution of the number of producers. Our homogeneity assumption requires measures of producer heterogeneity to have insubstantial effects in a forecasting model of the number of firms that accounts for time-invariant market characteristics, the realization of demand (population in our case), and the lagged values of the number of producers. To investigate this, we have estimated Poisson regression models with the number of firms in year t as the dependent variable. The independent variables include the logarithm of population in year $t - 1$ and its square, dummy variables for the number of theaters in year $t - 1$, calendar-year dummies, census-division dummies, linear and squared terms in median income measured in 2000, and the logarithms of geographic preference diversity and geographic isolation.

In one of the forecasting models, we also include a particular measure of heterogeneity. If we denote the size of the j th active producer in market r in year t with $Q_{j,r,t}$, then this can be written as

$$H_{r,t} \equiv \frac{1}{N_{r,t}} \sum_{j=1}^{N_{r,t}} Q_{j,r,t}^2 / \left(\frac{1}{N_{r,t}} \sum_{j=1}^{N_{r,t}} Q_{j,r,t} \right)^2 .$$

This is the *uncentered* second moment of producer size divided by its first moment squared. It can be easily shown that $H_{r,t}$ equals the Herfindahl-Hirschman Index (HHI) multiplied by $N_{r,t}$. Since the HHI obtains its minimum value of $1/N_{r,t}$ when producers have equal sizes, $H_{r,t}$ equals the multiplicative “correction” one must apply to $1/N_{r,t}$ to get the true HHI. Its minimum value (obtained with identical

²⁰We estimate that the entry and exit rates of theaters between the 2002 and 2007 Economic Censuses were *at least* 9.2 percent and 11.3 percent. Therefore, a more thorough investigation would almost certainly find that the Motion Picture Theaters industry is no exception to the empirical rule of high producer turnover at the national level. We calculated these lower bounds using reports of the number of establishments in both years’ Economic Censuses, 4979 and 4879, and the number of establishments in the 2002 Census that did not operate for all of 2002. We counted these as exits. Since exits could also occur in any of the inter-census years, this is a lower bound. The 11.3 percent exit rate equals $100 \times 2 \times 555 / (4979 + 4879)$ percent. To get the entry rate, we subtracted the number of establishments that were active for all of 2002 from the number of establishments in the 2007 Economic Census. The result, 455, is a lower bound on the number of entrants between the two censuses. The 9.2 percent entry rate equals $100 \times 2 \times 455 / (4979 + 4879)$ percent. The data underlying these calculations can be found at http://factfinder.census.gov/bkmk/table/1.0/en/ECN/2007_US/51SSZ1.




Figure 3: Histogram of the Absolute Differences between Model Forecasts

Note: The figure plots the fraction of market-year observations for which the absolute difference of the two models' predicted (expected) values of the number of theaters falls into bins with width equal to 0.01 theater. Both models are Poisson regressions that include calendar-year dummies, census-division dummies, linear and squared terms in the previous year's population, linear and squared terms in median income measured in 2000, the logarithms of geographic preference diversity and geographic isolation (as defined in the text), and indicators for the number of theaters serving the market in the previous year. One model also includes the logarithm of $H_{r,t-1}$, heterogeneity's contribution to the HHI in the previous year. Both models are estimated with the sample of 1,572 observations from 2001 through 2009 with $N_{r,t-1} \geq 2$.

producers) is one. By including this measure in our forecasting model, we give the data the opportunity to indicate whether or not producer heterogeneity substantially impacts the evolution of $N_{r,t}$. If our structural model were literally true, then we could exclude $H_{r,t-1}$ from the forecasting model without cost. On the other hand, we expect $H_{r,t-1}$ to substantially improve forecasts of $N_{r,t}$ generated from other models of industry dynamics. For example, in models of increasing market dominance, such as that in Cabral (2002), heterogeneity increases over time on average, which increases the incentive for lagging producers to exit. Similarly, the introduction of a big-box retailer to a market increases producer heterogeneity and induces smaller competitors to exit, so a high value of $H_{r,t-1}$ should predict a reduction in $N_{r,t}$.²¹

In our data, we measure each theater's size with the midpoint of the employment

²¹See Haltiwanger, Jarmin, and Krizan (2010) for evidence of this effect of big-box retailers on their smaller competitors.

size category to which it belongs in the year’s mid-March pay period. The County Business Patterns always reports this for each theater without identifying the theater itself. We include the logarithm of $H_{r,t-1}$ in our forecasting model, so its estimated coefficient can be interpreted as an elasticity. So that we do not bias the forecasting model towards finding that heterogeneity is unimportant, we include only observations for which $N_{r,t-1} \geq 2$. Over our ten year sample, there are 1,572 such observations. Over them, the mean value of $H_{r,t-1}$ is 1.31, and its standard deviation is 0.30. The estimated coefficient on the logarithm of $H_{r,t-1}$ equals -0.096 , and its standard error is 0.036. Therefore, $H_{r,t-1}$ has a statistically-significant and negative effect on $N_{r,t}$. However, this effect is economically small. We can see this in two ways. First, we note that the logarithm of $H_{r,t-1}$ has a standard deviation of 0.21 in our sample. Therefore, a one-standard deviation increase in this measure of heterogeneity decreases the predicted number of firms by only about 2 percent. Second, we estimated the same forecasting model but excluding the logarithm of $H_{r,t-1}$ and calculated forecasts using it and our complete model. Figure 3 gives the histogram of the absolute differences between the two models’ forecasts. The median absolute difference between the two models’ predictions is 0.030 theaters, and the mean absolute difference is 0.036 theaters. In comparison, the mean forecasted values for both models equal 2.46 theaters. Since there is no economically significant effect of producer heterogeneity on the evolution of $N_{r,t}$, we conclude that abstracting from producer heterogeneity is a helpful simplification rather than a fatal error.

4.3 Estimates

The NFXP procedure requires us to specify the demand process G_C , the distribution G_W of the cost shocks, and the per period surplus function π as functions of a finite vector of parameters. For the demand process, we follow [Tauchen \(1986\)](#). We restrict $C_{r,t}$ to a grid of 200 points equally spaced on a logarithmic scale with distance d : $c_{[1]}, c_{[2]} = c_{[1]} \exp(d), \dots, c_{[200]} = c_{[1]} \exp(199d)$. So that the growth of C_t is approximately normally distributed with mean μ_C and variance σ_C^2 , we specify the probability of transitioning to $c_{[i]}$ from $c_{[j]}$ for any $i = 2, \dots, 199$ and $j = 1, \dots, 200$ with

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

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

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

and

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

respectively. The lower bound of the demand grid equals the minimum population observed in our data, 11,011, divided by 1.25. Analogously, the upper bound equals the maximum population, 197,912, multiplied by 1.25. For estimation, we replace each observation of μ SA population with the closest grid point.

The maximum number of movie theaters sustainable, \tilde{n} , is fixed at the maximum number of theaters observed in the data, nine. We specify the distribution G_W of the cost shocks to be normal with standard deviation ω and mean $-\omega^2/2$, so that $\exp(W_{r,t})$ has a log-normal distribution with unit mean and scale parameter ω . We fix the discount factor ρ at $\frac{1}{1.05}$. The specification for the producers' surplus function is

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

For this, we split the market characteristics in X_r into two sub-vectors, $X_r \equiv (X_r^{(1)}, X_r^{(2)})$. Those in $X_r^{(1)}$ affect the surplus log-linearly and include the logarithm of median income (expressed as a deviation from the logarithm of the average median income across our 573 μ SAs) and dummies for all Census Divisions excluding New England. The remaining characteristics in $X_r^{(2)}$ interact with k and thereby affect the toughness of competition in a general way.²² 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 set $k(4) = k(5) = \dots = k(9)$ to accommodate the paucity of observations with four or more theaters.

²²The monotonicity assumption in Section 2 requires the flow surplus to weakly decrease with the number of firms. This assumption is satisfied at the maximum-likelihood estimates.

Table 3 reports the estimated parameters for two specifications, one that ignores geographic preference diversity and another that takes it into account. The entire estimation of both specifications required about thirty minutes using two Intel Xeon E5-2699 v3 CPUs (released by Intel in 2014) on a single machine with C++ code.

In the first specification, the full-information maximum likelihood estimates of the demand process’s drift and innovation standard deviation, 0.34 and 1.21 percent, are very close to the unconditional sample mean and standard deviation of population growth, 0.34 and 1.11 percent. The coefficients in β are jointly and (with the exceptions of those multiplying two 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. In the simulation, the average fixed cost of continuation and sunk cost of entry *paid* were 0.47 and 0.92. The estimates of all these parameters from the specification that accounts for geographic preference diversity are similar to these baseline estimates.

Table 4 reports transformations of the estimates from Table 3 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 26,360 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 29,420 people to support a monopolist in a μ SA with geographic preference diversity above the median and 25,980 to support a monopolist in a μ SA 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 4 report estimates of $k(n+1)/k(n)$, the share of the surplus per consumer left after the addition of a competitor. These indicate very tough competition. In the first specification, duopolists’ producers’ surplus per consumer equals 54 percent of a monopolist’s. A third competitor decreases this

Table 3: Parameter Estimates

	Geographic Preference Diversity		
	All μ SAs	Diversity > 13.4 miles	Diversity \leq 13.4 miles
$k(1) \times 10^5$	0.38 (0.05)	0.34 (0.05)	0.38 (0.05)
$k(2) \times 10^5$	0.21 (0.03)	0.20 (0.03)	0.18 (0.03)
$k(3) \times 10^5$	0.17 (0.03)	0.17 (0.03)	0.14 (0.03)
$k(4) \times 10^5$	0.13 (0.02)	0.14 (0.02)	0.10 (0.02)
φ	51.23 (9.82)		49.06 (9.75)
Median Income	1.75 (0.07)		1.74 (0.08)
Mid Atlantic	0.88 (0.17)		0.85 (0.17)
East North Central	-0.64 (0.15)		-0.60 (0.15)
West North Central	-0.49 (0.14)		-0.45 (0.14)
South Atlantic	0.05 (0.16)		0.09 (0.16)
East South Central	-0.73 (0.15)		-0.71 (0.15)
West South	-0.52 (0.16)		-0.49 (0.16)
Mountain	-0.33 (0.15)		-0.29 (0.15)
Pacific	-0.12 (0.15)		-0.07 (0.15)
ω	1.75 (0.15)		1.74 (0.15)
$\mu_C \times 10^2$	0.34 (0.00)		0.34 (0.00)
$\sigma_C \times 10^2$	1.21 (0.00)		1.21 (0.00)
$-\log \mathcal{L}$	9199.13		9192.24
Number of Markets	573	287	286

Note: Standard errors are reported in parentheses. The data include 573 μ SAs from 2000 to 2009, and \tilde{n} equals nine, which is the maximum of the number of active cinemas observed in the data. The values of $k(5), \dots, k(9)$ identically equal $k(4)$. The first column reports estimates of a specification in which $k(n)$ does not vary between markets. The second and third columns report estimates of a specification in which $k(n)$ 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 μ SAs) and Census Division dummies (excluding New England) as market characteristics in $X_r^{(1)}$. Please see the text for further details.

Table 4: Estimates of the Toughness of Competition

	Geographic Preference Diversity		
	All μ SAs	Diversity > 13.4 miles	Diversity \leq 13.4 miles
$1/(k(1) \times 10^3)$	26.35 (3.50)	29.42 (4.00)	25.98 (3.64)
$k(2)/k(1)$	0.54 (0.14)	0.60 (0.14)	0.48 (0.20)
$k(3)/k(2)$	0.82 (0.06)	0.84 (0.06)	0.78 (0.10)
$k(4)/k(3)$	0.77 (0.08)	0.79 (0.08)	0.66 (0.21)
Number of Markets	573	287	286

Note: This table is based on the model's estimates as reported in Table 3. 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.

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 contrast with other evidence from this industry. [Davis \(2005\)](#) provided 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 μ SAs) U.S. cities in the 1990s. Based on both across-market and within-market-over-time variation, he concluded 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.²³ Supporting the view that film distributors constrain theaters' pricing choices, [Davis \(2006\)](#) found 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 our estimates reflect fierce competition for customers. Instead, we believe that adding theaters to a market increases competition for film exhibition rights. Indeed, [Gil and LaFontaine \(2012\)](#) presented some evidence that owners of Spanish theaters with higher local market shares get better deals from film distributors.


4.4 Dynamic Implications

The estimates in [Table 3](#) show that movie theaters face substantial sunk entry costs and uncertainty about future profits. Consequently, nontrivial dynamic considerations govern their entry and exit. This section explores the implied dynamics of local movie theater markets by simulating the estimated model and counterfactual versions of it. This is straightforward because its equilibrium is unique and easy to calculate.

We first consider the dynamics of a single local market. We find that a market's initial number of active firms and initial demand state (population) have long-lasting

²³[Orbach and Einav \(2007\)](#) reviewed the legal environment in which theater owners negotiate with film distributors and set admissions prices.

Figure 4: Initial Conditions and the Evolution of the Expected Number of Firms




Note: These panels report the evolution of the expected number of active firms implied by the model. The model estimates are taken from the second specification reported in Table 3 and correspond to a high diversity market in New England with median income equal to \$34,417, the average median income across μ SAs. The left panel sets the initial demand state to 32,558, the first quartile of the 2009 population distribution across μ SAs in our data. The right panel sets it to 64,119, the third quartile. The initial values for the number of active firms are marked as such in both graphs.

effects on its expected number of active firms. Figure 4 illustrates this for a New England market with average median income, high geographic preference diversity, and parameters equal to the estimates from the second specification in Table 3. Its left panel sets initial demand in this market equal to the first quartile of the 2009 population distribution across μ SAs (32,570) and plots the evolution of the expected number of firms from various initial numbers of firms. Its right panel sets initial demand to the third quartile of that same distribution (65,119). In both panels, the dependence on the initial number of firms vanishes only after 10 to 15 years. Sunk costs are key to this gradual adjustment; if we were to set them to zero in our model, the expected number of firms would lose its relation to the initial number of firms right away.²⁴

Figure 4 also shows that the average number of firms serving the industry after

²⁴Because incumbent firms can commit to serving the market in the next period before entry takes place, their survival decisions are nontrivially dynamic even in the absence of sunk costs. Nevertheless, the equilibrium without sunk costs closely mimics repeated static Nash play of a one-shot entry game with no incumbency.

Figure 5: Expected Number of Active Firms in the Long Run



Note: These panels show the log of the expected number of active firms after thirty years as a function of the log initial population. The black curves show this relationship for the estimated model. They are based on the second specification reported in Table 3 and correspond to a market in New England with high geographic preference diversity and median income equal to \$34,417 (the average median income across μ SAs). In the left panel, the orange curve corresponds to a counterfactual model without sunk costs. In this counterfactual, fixed costs are raised to $1 + \varphi(1 - \rho)$. In the right panel, the red curve shows the relationship for the Netflix counterfactual, i.e. a 25 percent reduction in consumers' propensity to go to the movies. The green and blue curves correspond to the counterfactuals that add JOAs to the Netflix counterfactual for all markets and duopolies, respectively.

30 years depends on the market's initial demand state. If the market starts at the first quartile of the population distribution, it will have on average 1.32 active firms after 30 years (left panel); if it starts at the third population quartile, it will end up with 2.20 firms (right panel). By construction the demand process is stationary, so this dependence vanishes if given enough time. However, this process requires well over 1,000 years. (This is not surprising, since we designed G_C to approximate the short-run behavior of a random walk.) Therefore, the usual mathematically-convenient definition of the “long run”—the model's ergodic distribution—is not practically relevant for this application. Instead, we define the “long-run” with the distribution of the number of active firms 30 years from the present. Based on this definition and on the results in Figure 4, we conclude that the transition to the long run requires between 10 and 15 years.


Figure 5 further explores how the long-run average number of firms depends on initial demand. In its left panel, we plot the log of the market's initial population

against the log of its expected number of active firms 30 years later. The black and orange curves show this function for our estimated model and for a counterfactual variant without sunk costs.²⁵ Their slopes give the percentage change in the number of active firms for a one percent change in initial population. In both cases, for large enough initial demand (and thus a high enough expected number of firms), it takes more than a 1 percent increase in initial demand for a 1 percent increase in the expected number of firms. This reflects our estimates of tough competition. At very high initial demand states, the expected number of firms is close to the maximum number of firms, even large increases in demand cannot entice further entry, and the slopes taper off to zero. At low enough initial demand states, the expected number of firms is close to zero and its elasticity with respect to demand is larger. Comparing both curves shows that sunk costs magnify this effect; this is consistent with [Abbring and Campbell's \(2010\)](#) finding that the corresponding option values shift the population thresholds at which firms exit down.

Next, we consider the market's dynamic responses to a permanent 25 percent fall in its consumers' propensity to patronize its theaters. Since this demand reduction could follow the screening of new movies by an internet streaming platform, we will refer to it as the "Netflix shock." Figure 6 plots the implied differences between the counterfactual outcomes and Figure 4's baseline outcomes as a share of those baseline outcomes for the same initial conditions as in Figure 4. Like the baseline outcomes themselves, the short-run responses to the Netflix shock depend strongly on the initial conditions, with only their dependence on initial demand persisting in the long run. With low initial demand, the market adjusts more quickly to the Netflix shock if it starts with a large number of firms. Similarly, with high initial demand, it adjusts more quickly if it starts with few firms. Eventually, the market sheds around 20 percent of its theaters. This long-run loss is larger if initial demand is low. Figure 5's right panel confirms this: The gap between the log long-run expected number of active firms in the baseline (black) and in the Netflix counterfactual (red) decreases with initial demand. One possible explanation for this is that the reduction of competition dampens the negative effects of the Netflix shock on the profits of surviving incumbents in high demand markets, which are

²⁵In the counterfactual, the sunk costs are annualized and added to the fixed costs (see the note to Figure 5). Increasing the fixed costs further (so that the implied number of firms is closer to that of the estimated model) shifts the orange line down without substantially changing its shape.

Figure 6: Relative Changes in the Expected Number of Firms after a Netflix Shock



Note: These panels report responses to the Netflix shock, a permanent 25 percent reduction in demand, for various initial conditions. Specifically, they plot the differences between the counterfactual expected numbers of firms following the Netflix shock and Figure 4's baseline expected numbers of firms as a share of those same baseline numbers. The model estimates are taken from the second specification reported in Table 3, and the simulations correspond to a market in New England with high geographic preference diversity and median income equal to 34,417 (the average median income across μ SAs). The left panel sets the initial demand state to 32,558, the first quartile of the 2009 population distribution across μ SAs. The right panel sets it to 64,119, the third quartile. The initial values for the number of active firms are marked as such in both graphs.

more likely to support two or more firms.

Our estimates contain no information on how consumers value variety, so we cannot evaluate the social optimality of the number of active theaters or its adjustment. However, our model does allow us to calculate the *positive* responses of the industry to a policy intervention.²⁶ To demonstrate this capability, we consider one possible policy response to the Netflix shock: an antitrust exemption that allows competing theaters to sign joint operating agreements (JOAs) that centralize the acquisition of screening rights. Such a policy would be reminiscent of the 1970 Newspaper Preservation Act. This allowed newspapers, which have long been in decline, to centralize the choices of advertising rates. We consider two variants of this intervention. One allows agreements in all markets and the other restricts them

²⁶If external information on the benefits of variety was available, appending it to this empirical analysis would be straightforward.

to duopolies.²⁷ The latter would make sense if the benefits of variety are exhausted with two competitors.


Figure 5's right panel plots the long run log expected number of firms if the Netflix shock is compensated with an all markets JOA (green) and if it is compensated with a duopoly JOA (blue). Neither JOA has much impact if initial demand is very low and markets are unlikely to support more than one theater. The baseline JOA policy is increasingly effective as initial demand, and thus the expected number of firms, increases. For large enough initial demand, it more than offsets the Netflix shock. On the other hand, the duopoly JOA only just compensates for the Netflix shock at intermediate levels of demand.

So far, we have focused on the dynamics of a single local market under various initial conditions. We finish this section by exploring the (counterfactual) evolution of our sample's actual markets following a Netflix shock with and without a JOA at the end of the sample period. To this end, we compute each market's equilibrium outcome paths starting from its actual 2009 demand state and market structure under various combinations of a Netflix shock and a JOA intervention in 2010. Figure 7 summarizes the results by partitioning the markets into four subsamples, one for each quartile of their 2009 population distribution, and plotting the average expected number of firms in each such subsample against the years elapsed since 2010. As our findings for the long run effects on a single market suggest, a duopoly JOA falls short of compensating the 25 percent Netflix shock in all years following it and all four subsamples. The baseline JOA more than offsets the Netflix shocks in all but the smallest markets, which benefit little from agreements among three or more theaters.

To further quantify the JOAs' effects, we have computed the sizes of the negative demand shocks that are exactly offset by each JOA in terms of the resulting average (across markets) long-run (30 year) expected number of firms. The baseline JOA exactly offsets a 36 percent permanent reduction in demand. This makes sense, given that it more than compensates for the 25 percent Netflix shock in at least 75 percent of the markets. In contrast, a duopoly JOA only offsets a 17 percent demand reduction.

²⁷We operationalize the JOAs for all markets by setting $k(n)$ to the estimated value of the monopolist's surplus $k(1)$ for all $n \geq 2$. We operationalize the JOAs for duopoly markets by setting $k(2)$ equal to the estimated value of $k(1)$ while keeping $k(n)$ unchanged and equal to their estimated values for all $n \geq 3$.

Figure 7: Expected Number of Firms following a Netflix Shock (Averaged over Markets by Demand Quartile)



Note: These panels report the evolution of the expected number of active firms implied by the estimated model and three counterfactual versions of it from 2010, averaged over the sampled markets by 2009 demand quartile. All three counterfactual models involve a Netflix shock, a permanent 25 percent reduction in demand, in 2010. The second and third counterfactual models add, respectively, an all markets and a duopoly JOA from 2010. The model estimates are taken from the second specification reported in Table 3.

5 Conclusion

This paper’s dynamic oligopoly model is a version of [Ericson and Pakes’s \(1995\)](#) framework in which firms face identical expected payoffs when making entry and survival decisions. Like [Ericson and Pakes](#), we focus on a game with complete information. Unlike them, we allow for mixed strategies to ensure the existence of a symmetric Markov-perfect equilibrium and leverage the implications of mixing to simplify the equilibrium analysis and computation.

Applied research following [Ericson and Pakes](#) (as summarized by [Doraszelski and Pakes, 2007](#)) has generally ensured the existence of an equilibrium in pure strategies by augmenting [Ericson and Pakes’s \(1995\)](#) model with firm-specific privately-observed shocks. In such augmented models, computing an equilibrium can be challenging, and there is no guarantee that one in hand is unique. Therefore, methods for estimating these models—such as [Bajari, Benkard, and Levin’s \(2007\)](#) and [Pakes, Ostrovsky, and Berry’s \(2007\)](#)—have avoided equilibrium calculation altogether. Instead, they assume that producers in all sample markets employ the same equilibrium strategy, use this to estimate firms’ expected behavior from their observed choices, and identify the structural parameters of interest from the implied individual choice problems. This “two-step” approach cleverly solves the problem of estimation when there could be equilibrium multiplicity. However, unlike the maximum likelihood estimator of our model, the two-step estimators are generally not efficient. Moreover, their asymptotic distributions may be hard to compute; in fact, [Bajari, Benkard, and Levin](#) (p. 1349) “believe it will typically be easiest to use subsampling or the bootstrap to estimate standard errors.” Finally, users of the two-step approach have had to coarsen covariates before estimation so that they could estimate transition probabilities by pooling data across markets (e.g. [Dunne, Klimek, Roberts, and Xu, 2013](#)).

Even with two-step estimates of the structural parameters of the augmented [Ericson and Pakes](#) framework in hand, equilibrium multiplicity seriously impedes its practical application. After all, structural parameters are rarely of interest *per se* but rather serve as intermediate inputs into the analysis of environmental changes and policy interventions. In contrast, conducting counterfactual policy analysis with our model, which has a unique equilibrium, is straightforward. For example, we repeatedly calculated the model’s unique equilibrium (separately for each of the 573

markets in our sample) to determine the size of a permanent demand shock that a joint operating agreement exactly offsets. Not every important question concerning dynamic oligopolies can be cast within an environment with a unique equilibrium; and further methodological developments for policy analysis with multiple equilibria (such as calculating bounds on counterfactuals as in [Eizenberg, 2014](#), and [Reguant, 2015](#)) are needed. Nevertheless, our analysis has surmounted these theoretical and computational obstacles for an empirical framework capable of answering some of the oldest questions in dynamic industrial organization.

References

- ABBRING, J. H. AND J. R. CAMPBELL (2010): “Last-In First-Out Oligopoly Dynamics,” *Econometrica*, 78, 1491–1527.
- ABBRING, J. H., J. R. CAMPBELL, J. TILLY, AND N. YANG (2017): “Very Simple Markov-Perfect Industry Dynamics: Theory,” *Econometrica*, forthcoming.
- ABBRING, J. H. AND Ø. DALJORD (2016): “Identifying the Discount Factor in Dynamic Discrete Choice Models,” Discussion Paper 11133, Centre for Economic Policy Research, London.
- BAJARI, P. L., C. L. BENKARD, AND J. LEVIN (2007): “Estimating Dynamic Models of Imperfect Competition,” *Econometrica*, 75, 1331–1370.
- BRESNAHAN, T. F. AND P. C. REISS (1990): “Entry in Monopoly Markets,” *Review of Economic Studies*, 57, 531–553.
- (1991a): “Empirical models of discrete games,” *Journal of Econometrics*, 48, 57–81.
- (1991b): “Entry and Competition in Concentrated Markets,” *Journal of Political Economy*, 99, 977–1009.
- (1994): “Measuring the Importance of Sunk Costs,” *Annales d’Économie et de Statistique*, 34, 181–217.
- CABRAL, L. M. B. (2002): “Increasing Dominance with No Efficiency Effect,” *Journal of Economic Theory*, 102, 471–479.
- CAMPBELL, J. R. AND H. A. HOPENHAYN (2005): “Market Size Matters,” *Journal of Industrial Economics*, 53, 1–25.
- DAVIS, P. (2002): “Estimating Multi-Way Error Components Models with Unbalanced Data Structures,” *Journal of Econometrics*, 106, 67–95.
- (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.

- (2006): “Spatial Competition in Retail Markets: Movie Theaters,” *RAND Journal of Economics*, 37, 964–982.
- 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.
- 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.
- DUNNE, T., S. KLIMEK, M. J. ROBERTS, AND D. Y. XU (2013): “Entry, Exit, and the Determinants of Market Structure,” *RAND Journal of Economics*, 44, 462–487.
- DUNNE, T., M. J. ROBERTS, AND L. SAMUELSON (1988): “Patterns of Firm Entry and Exit in U.S. Manufacturing Industries,” *RAND Journal of Economics*, 19, 495–515.
- EIZENBERG, A. (2014): “Upstream Innovation and Product Variety in the U.S. Home PC Market,” *Review of Economic Studies*, 81, 1003–1045.
- ERICSON, R. AND A. PAKES (1995): “Markov-Perfect Industry Dynamics: A Framework for Empirical Work,” *Review of Economic Studies*, 62, 53–82.
- FUDENBERG, D. AND J. TIROLE (1991): *Game Theory*, Cambridge: 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.
- HALTIWANGER, J., R. JARMIN, AND C. J. KRIZAN (2010): “Mom-and-Pop meet Big-Box: Complements or Substitutes?” *Journal of Urban Economics*, 67, 116–134.
- HARWELL, D. (2015): “Netflix Plan: Bring Big-Screen Premieres to Subscribers’ Small Screens,” *Washington Post*,

https://www.washingtonpost.com/business/economy/netflix-plan-bring-big-screen-premiers-to-subscribers-small-screens/2015/10/14/cf197154-71d0-11e5-8248-98e0f5a2e830_story.html; accessed on December 4, 2015.

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

HOTZ, V. J. AND R. A. MILLER (1993): “Conditional Choice Probabilities and the Estimation of Dynamic Models,” *Review of Economic Studies*, 60, 497–529.

ISKHAKOV, F., J. LEE, J. RUST, B. SCHJERNING, AND K. SEO (2016): “Comment on ‘Constrained Optimization Approaches to Estimation of Structural Models’,” *Econometrica*, 84, 365–370.

KAFKA, P. (2013): “Netflix Flirts with a New Idea: “Big” Movies at Your House, the Same Day They’re in Theaters,” *AllThingsD.com*, <http://allthingsd.com/20131028/netflix-flirts-with-a-new-idea-big-movies-at-your-house-the-same-day-theyre-in-theaters/>; accessed on May 31, 2015.

MAGNAC, T. AND D. THESMAR (2002): “Identifying Dynamic Discrete Decision Processes,” *Econometrica*, 70, 801–816.

MASKIN, E. AND J. TIROLE (1988): “A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs,” *Econometrica*, 56, 549–569.

MCAFEE, R. P., H. M. MIALON, AND M. A. WILLIAMS (2004): “What Is a Barrier to Entry?” *American Economic Review*, 94, 461–465.

NEWHEY, W. K. AND J. L. POWELL (2003): “Instrumental Variable Estimation of Nonparametric Models,” *Econometrica*, 71, 1565–1578.

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.

PAKES, A., M. OSTROVSKY, AND S. BERRY (2007): “Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples),” *RAND Journal of Economics*, 38, 373–399.

- REGUANT, M. (2015): “A Bounds Approach to Counterfactual Analysis,” Mimeo, Department of Economics, Northwestern University, Evanston, IL.
- RUST, J. (1987): “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55, 999–1033.
- SCHWARTZEL, E. AND B. FRITZ (2015): “Paramount to Speed Up Home Release of Movies,” *Wall Street Journal*, <http://www.wsj.com/articles/paramount-to-break-hollywoods-home-video-window-1436377631>; accessed on July 10, 2015.
- SU, C.-L. AND K. L. JUDD (2012): “Constrained Optimization Approaches to Estimation of Structural Models,” *Econometrica*, 80, 2213–2230.
- TAKAHASHI, Y. (2015): “Estimating a War of Attrition: The Case of the U.S. Movie Theater Industry,” *American Economic Review*, 105, 2204–2241.
- TAUCHEN, G. (1986): “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions,” *Economics Letters*, 20, 177–181.
- WEINTRAUB, G. Y., C. L. BENKARD, AND B. VAN ROY (2008): “Markov Perfect Industry Dynamics with Many Firms,” *Econometrica*, 76, 1375–1411.

Appendices

A General Model

In this appendix, we show that our model of a local market is a special case of the model studied by [Abbring, Campbell, Tilly, and Yang \(2017\)](#) and translate their main theoretical results to our model. We will refer to the two models as the *special* model and the *general* model, respectively.

In the general model, firms' profits in period t depend on the \mathcal{Y} -valued vector Y_t . Figure 8 gives the general model's recursive extensive form. Period t starts in state (N_t, Y_t) , and each incumbent earns profits $\tilde{\pi}(N_t, Y_t)$. 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 stage 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 stage, 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 the Markov transition distribution $\tilde{G}(\cdot | Y_t)$.

[Abbring et al.](#) restrict the general model's payoffs as follows:

- A1. $\exists \tilde{\pi} < \infty$ such that $\forall (n, y) \in \mathbb{N} \times \mathcal{Y}$, $-\infty < \mathbb{E}[\tilde{\pi}(n, Y') | Y = y] < \tilde{\pi}$;
- A2. $\exists \tilde{n} \in \mathbb{N} : \forall n > \tilde{n}$ and $\forall y \in \mathcal{Y}$, $\tilde{\pi}(n, y) < 0$;
- A3. $\forall (n, y) \in \mathbb{N} \times \mathcal{Y}$, $\tilde{\pi}(n, y) \geq \tilde{\pi}(n + 1, y)$; and
- A4. $\forall (m, y) \in \mathbb{N} \times \mathcal{Y}$, $0 < \tilde{\varphi}(m, y) \leq \tilde{\varphi}(m + 1, y)$.

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

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




Figure 8: The General Model's Recursive Extensive Form

This reduces the general model to the main text's special model, except that the continuation cost $\exp(w)$ to a firm in state (c, w, w_{-1}) is not incurred when survival outcomes are realized but subsumed, with an appropriate adjustment for discounting, in the next period's expected profits:

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

Because the general model's continuation cost, $\rho^{-1} \exp(w)$, cannot be avoided by surviving firms and has a present value of $\exp(w)$, it is equivalent to a cost $\exp(w)$ that is due right after survival, as in the special model.

The special model's assumptions that $\pi(n, c) \geq 0$ and $\mathbb{E}[\pi(n, C') | C = c] \leq \tilde{\pi}$, ensure that [A1](#) holds. The assumption that there exists \tilde{n} such that $\pi(n, c) = 0$ for all $n > \tilde{n}$ implies [A2](#). The assumption that $\pi(n, c) \geq \pi(n + 1, c)$ in the special model directly gives [A3](#). Finally, [A4](#) generalizes the special model's assumption of

a constant $\varphi > 0$. Since $\tilde{\varphi}(m, \cdot)$ may increase in m , the general model can have an *economic barrier to entry* (McAfee, Mialon, and Williams, 2004). 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.

Because our empirical model is a special case of Abbring et al.'s general model, their theoretical results apply to it. In equilibrium, no firm will enter a market with \tilde{n} or more incumbents and incumbents will leave a market served by more than \tilde{n} firms with positive probability (Lemma 1). Moreover, the special model has a unique symmetric Markov-perfect equilibrium that defaults to inactivity (Theorem 1). In this equilibrium, the general model's post-entry value $\tilde{v}_E(n_E; c, w, w_{-1})$ equals the special model's value $v_E(n_E, c, w)$. However, because of the (innocuous) difference in the timing of the continuation cost, the general and special model's post-survival values differ by the continuation cost: $\tilde{v}_S(n'; c, w, w_{-1}) = v_S(n', c) - \exp(w)$.

Because $\tilde{v}_S(n'; c, w, w_{-1})$ is weakly decreasing in n' (Lemma 2), so is $v_S(n', c)$. For given v_S , there is a unique equilibrium survival rule a_S that defaults to inactivity (Corollary 1). In particular, $a_S(n_E, c, w) = 0$ if $\tilde{v}_S(1; c, w, w_{-1}) = v_S(1, c) - \exp(w) \leq 0$ (in the subcase that $v_S(1, c) = \dots = v_S(n_E, c) = \exp(w)$, this follows from the restriction that rules default to inactivity); $a_S(n_E, c, w)$ equals the unique survival probability $a \in (0, 1]$ that makes firms indifferent between exit and survival,

$$\begin{aligned} 0 &= \sum_{n'=1}^{n_E} \binom{n_E-1}{n'-1} a^{n'-1} (1-a)^{n_E-n'} \tilde{v}_S(n', (c, w, w_{-1})) \\ &= -\exp(w) + \sum_{n'=1}^{n_E} \binom{n_E-1}{n'-1} a^{n'-1} (1-a)^{n_E-n'} v_S(n', c), \end{aligned}$$

if $v_S(n_E, c) - \exp(w) = \tilde{v}_S(n_E, (c, w, w_{-1})) \leq 0 < \tilde{v}_S(1, (c, w, w_{-1})) = v_S(1, c) - \exp(w)$; and $a_S(n_E, c, w) = 1$ if $\tilde{v}_S(n_E, (c, w, w_{-1})) = v_S(n_E, c) - \exp(w) > 0$.

Monotonicity of v_S also implies that $v_E(n_E, c, w) = \max\{0, \tilde{v}_S(n_E, (c, w, w_{-1}))\} = \max\{0, -\exp(w) + v_S(n_E, c)\}$ (Corollary 2), so that $v_E(n_E, c, w)$ is weakly decreasing in n_E . In turn, this ensures that $\mathbb{E}_{a_E}[v_E(N_E, c, w) | M = m, C = c, W = w] > \tilde{\varphi}(m, (c, w, w_{-1})) = \varphi \exp(w)$ if and only if $v_E(m, c, w) > \varphi \exp(w)$, so that $a_E(m, c, w) = \mathbb{1}[v_E(m, c, w) > \varphi \exp(w)]$ (Abbring et al.'s Section 3.2).

Finally, because \tilde{v}_E is bounded (Abbring et al.'s Section 3.1), both v_E and v_S are bounded.

B Likelihood

In this appendix, we construct a key building block of Section 3.3's likelihood, $p(n'|n, c; X_r, \theta)$, which gives the probability that in a market r , the number of firms evolves from n to n' when the demand state is given by c . Suppose that we have obtained cost-shock thresholds for entry ($\bar{w}_{E,r}$) and sure survival ($\bar{w}_{S,r}$) by solving the model for given parameters. There are four cases to consider.

Case I: $n' > n$. If the number of firms increases from n to $n' > n$, then it must be profitable for $n' - n$ firms to enter, but not for $n' - n + 1$, i.e.

$$\bar{w}_{E,r}(n' + 1, c) \leq W_{r,t} < \bar{w}_{E,r}(n', c).$$

The probability of this event is

$$p(n'|n, c; X_r, \theta) = G_{W,r}(\bar{w}_{E,r}(n', c)) - G_{W,r}(\bar{w}_{E,r}(n' + 1, c)). \quad (12)$$

Case II: $0 < n' < n$. If the number of firms decreases from n to $n' > 0$, then $W_{r,t}$ must take a value w such that firms exit with probability $a_{S,r}(n, c, w) \in (0, 1)$. That is, w must be high enough so that n firms cannot survive profitably, $w \geq \bar{w}_{S,r}(n, c)$, but low enough for (at least) a monopolist to survive profitably, $w < \bar{w}_{S,r}(1, c)$. Given such a w , N' is binomially distributed with success probability $a_{S,r}(n, c, w)$ and population size n . Hence, the probability of observing a transition from n to n' with $0 < n' < n$ equals

$$\begin{aligned} p(n'|n, c; X_r, \theta) & \\ &= \int_{\bar{w}_{S,r}(n,c)}^{\bar{w}_{S,r}(1,c)} \binom{n}{n'} a_{S,r}(n, c, w)^{n'} (1 - a_{S,r}(n, c, w))^{n-n'} g_{W,r}(w) dw, \end{aligned} \quad (13)$$

where $g_{W,r}$ is the density of $G_{W,r}$. The integrand in (13) involves the mixing probabilities $a_{S,r}(n, c, w)$ that are implicitly defined in (3). We avoid computing these mixing probabilities directly by solving for the roots of (3) and perform a

change of variable instead. We substitute for w the value $\omega_r(a; n, c)$ that sets $a_{S,r}(n, c, \omega_r(a; n, c)) = a$ for a given survival probability $a \in (0, 1]$. We then integrate over a instead of over w . This gives

$$p(n'|n, c; X_r, \theta) = \int_0^1 \binom{n}{n'} a^{n'} (1-a)^{n-n'} \frac{d\omega_r(a; n, c)}{da} g_{W,r}(\omega_r(a; n, c)) da. \quad (14)$$

A simple rearrangement of (3) gives an explicit expression for $\omega_r(a; n, c)$:

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

Using this and its derivative $d\omega_r(a; n, c)/da$, we obtain an analytic expression for the integrand in (14). We then compute the integral with Gauss-Legendre quadrature.

Case III: $n' = 0, n > 0$. If all firms exit, then either it is not profitable for even a single firm to continue, $W_{r,t} \geq \bar{w}_{S,r}(1, c)$; or it is profitable for some firms but not for all firms to continue, $\bar{w}_{S,r}(n, c) \leq W_{r,t} < \bar{w}_{S,r}(1, c)$, firms mix over exit and survival as in Case II, and by chance none of the n firms survives. The probability of these events' union is

$$p(0|n, c; X_r, \theta) = 1 - G_{W,r}(\bar{w}_{S,r}(1, c)) + \int_{\bar{w}_{S,r}(n, c)}^{\bar{w}_{S,r}(1, c)} (1 - a_{S,r}(n, c, w))^n g_{W,r}(w) dw. \quad (15)$$

As in Case II, the integral in the right-hand side of (15) can be computed by substituting $w = \omega_r(a; n, c)$, which gives

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

and applying Gauss-Legendre quadrature.

Case IV: $n' = n = 0$. In this case, entry into an empty market is not profitable.

$$p(0|n = 0, c; X_r, \theta) = 1 - G_{W,r}(\bar{w}_{E,r}(1, c)) \quad (16)$$

Case V: $\mathbf{n}' = \mathbf{n}, \mathbf{n} > \mathbf{0}$. If there is neither entry nor exit, then either no firm finds it profitable to enter and all n incumbents find it profitable to stay, $\bar{w}_{E,r}(n+1, c) \leq W_{r,t} < \bar{w}_{S,r}(n, c)$; or the n incumbents mix as in Cases II and III, but by chance end up all staying. The probability of these events is

$$\begin{aligned}
 p(n|n, c; X_r, \theta) &= G_{W,r}(\bar{w}_{S,r}(n, c)) - G_{W,r}(\bar{w}_{E,r}(n+1, c)) \\
 &\quad + \int_{\bar{w}_{S,r}(n, c)}^{\bar{w}_{S,r}(1, c)} a_{S,r}(n, c, w)^n g_{W,r}(w) dw.
 \end{aligned} \tag{17}$$

The integral in (17) can be computed as in Cases II and III. This completes our construction of $p(n'|n, c; X_r, \theta)$.

C Monte Carlo Experiments

In this appendix, 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 as in (8), 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. We use the discretization procedure described in Section 4. We set $\mu_C = 0$ and $\sigma_C = 0.02$.

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 3. (Since this specification excludes variation in market characteristics, a single equilibrium calculation can and does support the likelihood function calculations for all of a sample's observations.) We use the mean and standard deviation of the log innovations of the generated C as starting values for the first step's likelihood function maximization. The starting parameter vector used for the second step equals a vector of ones multiplied by one random variable uniformly distributed

on $[1, 10]$. [Dubé, Fox, and Su \(2012\)](#) cautioned that an NFXP 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 two orders of magnitude smaller than that for the likelihood maximization to avoid this potential pitfall.²⁸

Table 5: Monte Carlo Results with Constant Surplus Per Consumer

	$\check{r} = 100$	$\check{r} = 250$	$\check{r} = 500$	$\check{r} = 1,000$
<i>Averages of Estimates</i>				
k	1.502	1.501	1.500	1.500
φ	10.219	10.064	10.054	10.017
ω	0.993	0.997	0.999	1.000
$\mu_C \times 10^2$	0.003	0.001	0.000	0.000
$\sigma_C \times 10^2$	1.999	2.001	2.001	2.001
<i>Averages of Estimated Standard Errors</i>				
k	0.050	0.031	0.022	0.015
φ	2.930	1.789	1.258	0.883
ω	0.070	0.044	0.031	0.022
$\mu_C \times 10^2$	0.068	0.043	0.030	0.021
$\sigma_C \times 10^2$	0.049	0.031	0.022	0.015
<i>Monte Carlo Estimates of 95% Confidence Interval Coverage</i>				
k	0.954	0.948	0.959	0.960
φ	0.930	0.932	0.944	0.938
ω	0.943	0.952	0.942	0.947
μ_C	0.948	0.944	0.951	0.949
σ_C	0.950	0.952	0.948	0.947

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_C , which is set to zero, and the innovation standard deviation σ_C , 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.

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 k , φ , and ω , to 1.5, 10, and 1, respectively. The lower and upper bounds of the demand grid equal 0.5 and 5. Table 5 reports the corresponding Monte Carlo

²⁸We set the tolerance value to 10^{-10} for the inner loop and to 10^{-8} for the outer loop.

Table 6: Monte Carlo Results with Decreasing Surplus Per Consumer

	$\bar{r} = 100$	$\bar{r} = 250$	$\bar{r} = 500$	$\bar{r} = 1,000$
<i>Averages of Estimates</i>				
$k(1)$	1.813	1.806	1.802	1.802
$k(2)$	1.398	1.399	1.400	1.400
$k(3)$	1.196	1.198	1.198	1.200
$k(4)$	0.997	0.998	0.998	0.999
$k(5)$	0.892	0.899	0.899	0.898
φ	10.004	9.958	9.976	10.026
ω	0.986	0.993	0.997	1.000
$\mu_C \times 10^2$	0.003	0.002	0.000	-0.000
$\sigma_C \times 10^2$	2.000	2.000	2.000	2.001
<i>Averages of Estimated Standard Errors</i>				
$k(1)$	0.095	0.057	0.040	0.028
$k(2)$	0.094	0.058	0.041	0.029
$k(3)$	0.079	0.049	0.034	0.024
$k(4)$	0.075	0.047	0.033	0.023
$k(5)$	0.088	0.054	0.038	0.027
φ	3.433	2.082	1.457	1.031
ω	0.086	0.053	0.037	0.026
$\mu_C \times 10^2$	0.068	0.043	0.030	0.021
$\sigma_C \times 10^2$	0.049	0.031	0.022	0.015
<i>Monte Carlo Estimates of 95% Confidence Interval Coverage</i>				
$k(1)$	0.955	0.954	0.961	0.946
$k(2)$	0.945	0.942	0.948	0.939
$k(3)$	0.945	0.941	0.949	0.956
$k(4)$	0.951	0.945	0.943	0.950
$k(5)$	0.940	0.940	0.943	0.935
φ	0.884	0.901	0.926	0.944
ω	0.932	0.929	0.931	0.945
μ_C	0.941	0.951	0.965	0.956
σ_C	0.948	0.950	0.950	0.930

Note: Results of a Monte Carlo experiment using the three step NFXP estimator to estimate the model with five profit parameters ($k(1), k(2), \dots, k(5)$) using 1,000 synthetic samples. The true value of ($k(1), k(2), \dots, k(5)$) 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_C , which is set to zero, and the innovation standard deviation σ_C , 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.

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 2.958. 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 $\sqrt{\tilde{r}}$. So for $\tilde{r} = 500$ the standard error on φ is only 1.254. The table's final panel reports the Monte Carlo estimates of 95 percent confidence intervals' coverage probabilities. These are all within 2.0 probability points of their common nominal value. Apparently, the estimated standard errors provide accurate inference.

For our second set of simulations we use the same parameterization as before except that we define 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.²⁹ Table 6 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 $k(1) = \dots = k(5)$ 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.

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 7 shows that this in fact is the case. On average, the computation of the three step maximum likelihood estimator takes between 15 and

²⁹These 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 \tilde{n}), because competition does not get "tougher" when the number of firms increases.

Table 7: Computational Performance

	$\tilde{r} = 100$	$\tilde{r} = 250$	$\tilde{r} = 500$	$\tilde{r} = 1,000$
<i>one entry cost parameter, one profit parameter</i>				
time per run	15.09	13.76	14.99	14.48
<i>one entry cost parameter, five profit parameters</i>				
time per run	20.81	18.12	19.97	21.05

Note: Average computational performance of the NFXP estimators in the synthetic samples. The estimator is implemented in C++ and runs as a single thread on an Intel Xeon CPU E5-2699 v3 with 2.30GHz.

a little over 20 seconds depending on the specification. This is much faster than the estimation time reported above for our application. The speed increase is entirely attributable to the absence of covariates, which requires computing an equilibrium for each market and each trial value of the parameters. One unexpected feature of Table 7 is that computation time can *decrease* as the number of markets grows. We speculate that this happens because larger sample sizes smooth the objective function and thereby reduce the number of likelihood function evaluations required for optimization.

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 NFXP algorithm. The MPEC estimator treats the value functions as a vector of nuisance parameters to be estimated subject to the equilibrium constraints implied by the sequence of Bellman equations. Thereby, MPEC omits the NFXP procedure’s inner loop. Su and Judd compared the MPEC and NFXP procedures in a simulation study of Rust’s (1987) bus engine renewal problem and concluded that MPEC is much faster. However, Iskhakov, Lee, Rust, Schjerning, and Seo (2016) pointed out that a more efficient version of the NFXP procedure, as originally proposed by Rust, is as fast as the MPEC procedure. Because results for a single-agent renewal problem do not necessarily carry over to our game, it is instructive to compare our NFXP results to those from an application of the MPEC procedure.

We only use the MPEC method in the second step of the three-step procedure, since the first step is independent of the estimation procedure used in the second

step and the third step essentially takes no time. Thus, if there are substantial differences between MPEC and the NFXP, these differences will be most visible in the second step of the three-step-procedure.

We implemented the MPEC estimator of our model in C++. We provided analytical derivatives for the objective function and the constraint Jacobian, and we explicitly accounted for the sparsity pattern of the constraint Jacobian. The Hessian was approximated using the Quasi-Newton BFGS method. The objective function was optimized using the commercial optimizer Knitro.

In contrast to [Su and Judd](#), we found the MPEC estimator to be considerably slower than the NFXP estimator. We obtained this result under very favorable starting values that fall within 10 percent of the truth. For these starting values, the MPEC estimator always calculated estimates practically indistinguishable from those of our NFXP estimator but was 40 times slower. We can only speculate as to why the MPEC estimator performs relatively poorly compared to the NFXP estimator. [Su and Judd](#) emphasized 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. In their application to [Rust’s](#) bus-engine replacement problem the constraint Jacobian is relatively sparse. Only 7 percent of its entries are non-zero. In our application, the constraint Jacobian is about 60 percent dense.³⁰ This is partly driven by the fact that the transition probability matrix for the demand state is fully dense. MPEC’s relatively poor performance in our application also might arise from the computation of the objective function’s gradients with respect to the nuisance parameters, which requires repeatedly retrieving information from 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.

³⁰For the specification that corresponds to the Monte Carlo simulation reported in [Table 6](#), there are 603,000 nonzeros in the Jacobian (out of 1,007,000 entries).