

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

Working Paper

Very simple Markov-perfect industry dynamics

Working Paper, No. 2013-20

Provided in Cooperation with:

Federal Reserve Bank of Chicago

Suggested Citation: Abbring, Jaap H.; Campbell, Jeffrey R.; Tilly, Jan; Yang, Nan (2013) : Very simple Markov-perfect industry dynamics, Working Paper, No. 2013-20, Federal Reserve Bank of Chicago, Chicago, IL

This Version is available at:

<https://hdl.handle.net/10419/96641>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Federal Reserve Bank of Chicago

Very Simple Markov-Perfect Industry Dynamics

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

November 2013

WP 2013-20

Very Simple Markov-Perfect Industry Dynamics

Jaap H. Abbring*

Jeffrey R. Campbell†

Jan Tilly‡

Nan Yang§

November 30, 2013

Abstract

This paper develops an econometric model of industry dynamics for concentrated markets that can be estimated very quickly from market-level panel data on the number of producers and consumers using a nested fixed-point algorithm. We show that the model has an essentially unique symmetric Markov-perfect equilibrium that can be calculated from the fixed points of a finite sequence of low-dimensional contraction mappings. Our nested fixed point procedure extends Rust's (1987) to account for the observable implications of mixed strategies on survival. We illustrate the model's empirical application with ten years of County Business Patterns data from the Motion Picture Theaters industry in 573 Micropolitan Statistical Areas. The results are suggestive of fierce competition between theaters in the market for film exhibition rights.

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

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

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

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

JEL Codes: L13, C25, C73.

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

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

1 Introduction

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

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

Monte Carlo results indicate that the NFXP procedure can accurately estimate sunk costs and profits per customer (as a multiple of the per-period fixed cost of production) using observations on the number of producers and consumers from as few as 100 markets over ten years. We further illustrate the model’s application by estimating its primitives for one industry with many concentrated local markets, Motion Picture Theaters (NAICS 512131). Our data include observations on the number of theaters from 2000 to 2009 serving 573 Micropolitan Statistical Areas (μ SAs). We find that adding a single firm to a monopoly market lowers profits

per customer by almost half. Adding two more firms brings profits per customer to 34 percent of its monopoly value. Since [Davis \(2005, 2006\)](#) found only modest effects of competition on ticket prices, we interpret these results as suggestive of fierce competition between theaters for film exhibition rights. In any case, the maximum likelihood estimation underlying this analysis takes only a few minutes on an ordinary desktop computer.

Both our model and that of [Abbring et al. \(2010\)](#) can be viewed as special cases of the [Ericson and Pakes \(1995\)](#) Markov-perfect industry dynamics framework that restrict firms from varying their investments in productivity improvements. Since computing that model's equilibria has proven to be computationally challenging (see [Doraszelski and Pakes \(2007\)](#) for examples), its estimation has focused on statistically inefficient methods that avoid equilibrium calculation altogether. For example, [Bajari et al. \(2007\)](#) apply the [Hotz and Miller \(1993\)](#) inversion to estimate the structural parameters governing a single agent's dynamic optimization problem after conditioning on the observed distribution of all other agents' choices. Our nested fixed-point algorithm computes maximum likelihood estimates, which are, of course, statistically efficient. Even with parameter estimates in hand, the lack of a fast algorithm for equilibrium computation makes counterfactual analysis of the [Ericson and Pakes \(1995\)](#) model difficult. [Weintraub et al. \(2008\)](#) make equilibrium computation more tractable by assuming that firms ignore current information about competitors' states. Instead, they make dynamic decisions based solely on their own state and knowledge of the long-run average industry state. Such *oblivious equilibria* approximate Markov-perfect equilibria when the number of competitors is large, yet they are much easier to compute and thus can serve as a starting point for empirical analysis, as in [Xu \(2008\)](#). Our model can serve as a similar starting point for the analysis of markets with only few firms.

The remainder of the paper proceeds as follows. The next section presents the model's primitives, and [Section 3](#) discusses equilibrium existence, uniqueness, and computation. [Section 4](#) develops the model's empirical implementation, which includes sampling, likelihood construction, identification, and maximum likelihood estimation using the NFXP procedure. [Section 5](#) demonstrates the light computational demands of the NFXP procedure and explores the estimator's finite sample behavior using Monte Carlo experiments. It also briefly discusses the relative performance of [Su and Judd's \(2012\)](#) mathematical programming with equilibrium

constraints (MPEC) implementation of the estimator. Section 6 illustrates the model’s application with the application to Motion Picture Theaters. Section 7 concludes. The Appendix provides technical details.

2 The Model

Consider a market in discrete time indexed by $t \in \mathbb{N} \equiv \{1, 2, \dots\}$. In period t , firms that have entered in the past serve the market. Each firm has a name $f \in \mathcal{F} \equiv \mathbb{N} \times \mathbb{N}$. The firm’s name gives the precise node of the game tree in which the firm has its single opportunity to enter the market. Aside from the timing of their entry opportunities, the firms are identical.

Figure 1 details the actions taken by firms in period t and their consequences for the game’s state at the start of period $t+1$. We call this the game’s *recursive extensive form*. For expositional purposes, we divide each period into two subperiods, the entry and survival subgames. Play in period t begins on the left with the entry subgame. If $t = 1$, nature sets the number N_1 of firms serving the market in period 1 and an initial demand state C_0 ; if $t > 1$, the number of incumbent firms N_t and the demand state C_{t-1} are inherited from period $t-1$. Nature draws a new demand state C_t from the conditional distribution $G_C(\cdot | C_{t-1})$ and a real-valued cost state W_t from the marginal distribution $G_W(\cdot)$. We use \mathcal{C} to denote the support of C_t ; W_t ’s support is the real line. All incumbent firms observe (C_t, W_t) , and each earns a surplus $\pi(N_t, C_t)$ from serving the market. We assume that

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

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

The period t entry cohort consists of firms with names in $t \times \mathbb{N}$. After incumbents receive their payoffs, these firms make their entry decisions sequentially in the order of their names' second components. We denote firm f 's entry decision with $a_E^f \in \{0, 1\}$. Firm f incurs the sunk cost $\varphi \exp(W_t)$ if it enters the market (i.e. $a_E^f = 1$). Otherwise (i.e. if $a_E^f = 0$), it earns a payoff of zero and never has an opportunity to enter again. This refusal to enter also ends the entry subgame, so firms remaining in this period's entry cohort that have not yet had an opportunity to enter *never* get to do so. Since the next firm in line faces exactly the same choice as did the firm that refused to enter, this convenient assumption is without economic content. Since every period has at least one firm refusing an available entry opportunity, the model is one of free entry.

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

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

The assumptions on $\{C_t, W_t\}$ make it a first-order Markov chain satisfying Rust's (1987) conditional independence assumption.¹ This ensures that the markets'

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

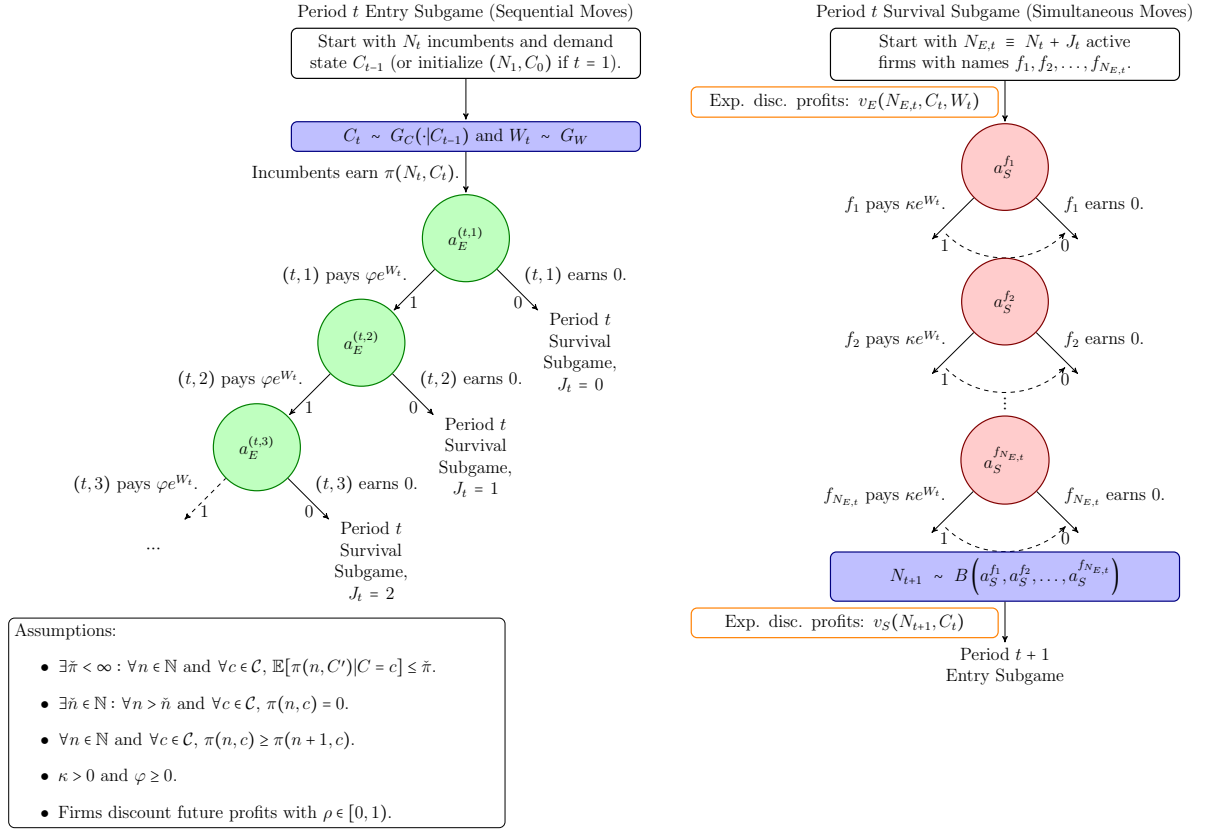


Figure 1: The Model's Recursive Extensive Form

observed (by the econometrician) initial conditions (N_1, C_0) cannot be informative of the unobserved cost shocks $\{W_t\}$.

3 Equilibrium

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

3.1 Markov Strategies

A Markov strategy is a strategy that maps *payoff relevant states* into actions. When a potential entrant (t, j) makes its entry decision in period t , the payoff-relevant states are the number of firms in the market including all of the current period's entrants up to and including (t, j) , $M_t^j \equiv N_t + j$, the current state of demand C_t ,

and the cost shock W_t . We collect the entrant's payoff relevant state variables into (M_t^j, C_t, W_t) , which takes values in $\mathcal{H} \equiv \mathbb{N} \times \mathcal{C} \times \mathbb{R}$. Similarly, we collect the payoff relevant state variables of a firm f contemplating survival in period t in the \mathcal{H} -valued $(N_{E,t}, C_t, W_t)$. 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 $a_S : \mathcal{H} \rightarrow [0, 1]$. Since time itself is not payoff relevant, we drop the subscript t from the payoff relevant states and denote the next period's value of a generic variable Z with Z' .

3.2 Symmetric Markov-Perfect Equilibrium

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

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

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

Here and throughout, capital and small letters denote random variables and their realizations. The expectation \mathbb{E}_{a_S} over N' takes survival of the firm of interest as given. Its subscript makes its dependence on a_S explicit. Similarly, we have

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

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

A strategy (a_E, a_S) forms a symmetric Markov-perfect equilibrium with payoffs (v_E, v_S) if and only if no firm can gain from a one-shot deviation (see Sections 4.2 and 13.2 in [Fudenberg and Tirole, 1991](#)). Thus, given the pair of payoff functions (v_E, v_S) , their corresponding strategy must satisfy

$$a_E(m^j, c, w) \in \arg \max_{a \in \{0,1\}} a \left(-\varphi \exp(w) \right. \quad (3)$$

$$\left. + \mathbb{E}_{a_E} [v_E(N_E, c, w) | M^j = m^j, C = c, W = w] \right),$$

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

Before proceeding to the equilibrium analysis, we wish to note and dispense with an uninteresting source of equilibrium multiplicity. If a potential entrant is indifferent between its two choices, we can construct one equilibrium from another by varying only that choice. Similarly, an incumbent monopolist can be indifferent between continuation and exit, and we can construct one equilibrium from another by changing that choice alone. To avoid these uninteresting caveats to our results, we follow [Abbring and Campbell \(2010\)](#) by focusing on equilibria that *default to inactivity*. In such an equilibrium, a potential entrant that is indifferent between entering or not stays out,

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

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

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

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

3.3 Existence, Uniqueness, and Computation

This section presents our analysis of equilibrium existence, uniqueness, and computation. For this, three features of the model are disposable: the serial

independence of W_t , the additive separability of per-period flow profits from the costs of continuation, and the invariance of sunk costs to the number of firms and the current demand state. In the Appendix, Sections A and B generalize the model by removing these assumptions and Section C provides proofs of all of this section's appropriately extended results.

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

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

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

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

The next step in the equilibrium analysis uses the assumption that *flow* payoffs weakly decrease with the number of competitors to show that the same monotonicity applies to the post entry and survival value functions.

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

The monotonicity assumption on π rules out exogenously specified complementarities between firms in the market. Lemma 2 says that no endogenous complementarity arises in equilibrium. Although this is intuitive, it is not a trivial

result. Indeed, [Abbring et al. \(2010\)](#) give a counterexample to the analogous proposition in a model with heterogeneous productivity types. In it, two high-productivity firms benefit each other by jointly deterring the entry of two low-quality potential rivals. That counterexample shows that Lemma 2 cannot be easily extended to the case with post-entry heterogeneity.²

Consider a one-shot simultaneous-moves survival game played by n_E active firms. In it, each of the n' survivors earns $-\kappa \exp(w) + v_S(n', c)$, with v_S the post-survival value in a symmetric Markov-perfect equilibrium of our dynamic game, and each exiting firm earns zero. The Nash equilibria of this game are intimately connected to the Markov-perfect equilibria of our model. In particular, (3) and (4) imply that a survival rule $a_S(n_E, c, w)$ from a symmetric Markov-perfect equilibrium forms a symmetric Nash equilibrium of the one-shot game, and vice versa.

First note that this one-shot game has many equilibria in the trivial case that $v_S(n_E, c) = \dots = -v_S(1, c) = \kappa \exp(w)$. In this case, our restriction to equilibria that default to inactivity picks the unique one in which $a_S(n_E, c, w) = 0$.

In the more interesting case where $v_S(n', c) \neq \kappa \exp(w)$ for at least one $n' \in \{1, \dots, n_E\}$, Lemma 2 guarantees that the one-shot game has a *unique* symmetric Nash equilibrium. To show this, we distinguish three subcases.

- First, suppose that $v_S(1, c) \leq \kappa \exp(w)$. If $v_S(1, c) = \kappa \exp(w)$, then Lemma 2 implies that $v_S(n', c) \leq \kappa \exp(w)$ for all $n' \in \{1, \dots, n_E - 1\}$. Furthermore, since $v_S(n', c) \neq \kappa \exp(w)$ for at least one $n' \in \{1, \dots, n_E\}$, we know that $v_S(n_E, c) < \kappa \exp(w)$. Since exiting for sure (setting $a_s(n_E, c, w) = 0$) is a weakly dominant strategy, it forms one symmetric equilibrium. Exiting for sure is also the unique best response to any positive symmetric continuation value, so it is the only symmetric equilibrium.
- Next, suppose that $v_S(n_E, c) \geq \kappa \exp(w)$. Lemma 2 implies that $v_S(n', c) \geq \kappa \exp(w)$ for $n' = 2, \dots, n_E - 1$. Since $v_S(n', c) \neq \kappa \exp(w)$ for at least one $n' \in \{1, \dots, n_E\}$, we also have and $v_S(1, c) > \kappa \exp(w)$. Therefore, continuing for sure (setting $a_S(n, c, w) = 1$) is a strictly dominant strategy and forms the unique Nash equilibrium.

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

- For the last subcase, suppose that $v_S(1, c) > 0$ and $v_S(n_E, c) < 0$. No pure strategy equilibrium exists, because the best response to all other firms continuing for sure is to exit for sure, and vice versa. However, the intermediate value theorem guarantees that that there is at least one survival probability a_S satisfying symmetric mixed strategy equilibria exist. no pure-strategy symmetric equilibrium exists, but there is an equilibrium in a mixed strategy. In it, $a_S(n_E, c, w)$ leaves firms indifferent between continuation and exit:

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

Lemma 2 guarantees that the left hand side of (5) weakly decreases in a_S , and the subcase's conditions further require it to be strictly decreasing for all $a_S \in [0, 1]$. Therefore, there is only one mixed strategy equilibrium.

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

Corollary 1 *Let v_S be the post-survival value function associated with a symmetric Markov-perfect equilibrium. Consider the one-shot survival game in which n_E firms simultaneously choose between survival and exit (as in the survival subgame of Figure 1), each of the n' survivors earns $-\kappa \exp(w) + v_S(c, n')$, with $-\kappa \exp(w) + v_S(n', c) \neq 0$ for at least one $n' \in \{1, \dots, n_E\}$, and each exiting firm earns zero. This game has a unique symmetric Nash equilibrium, possibly in mixed strategies.*

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

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

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

With Corollaries 1 and 2 in hand, we proceed to demonstrate equilibrium existence constructively. Our equilibrium uniqueness result and algorithm for

equilibrium calculation follow from the construction as byproducts. Begin with calculating $v_E(\check{n}, \cdot, \cdot)$ and $v_S(\check{n}, \cdot)$. From Lemma 1, there will be no entry in the next period, so

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

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

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

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

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

Here, $\mathbb{1}[x]$ gives x 's truth value.

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

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

This is the number of firms that will be active after potential entrants follow the candidate entry strategies. For given n , suppose that $v_E(m, \cdot, \cdot)$ and $a_E(m, \cdot, \cdot)$ for $m = n + 1, n + 2, \dots, \check{n}$ are in hand. Then, the equilibrium optimality conditions (3) and (4) along with Corollary 2 imply

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

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

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

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

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

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

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

Theorem 1 (Equilibrium existence and uniqueness) *There exists a unique symmetric Markov-perfect equilibrium that defaults to inactivity.*

4 Empirical Implementation

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

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

4.1 Sampling

Suppose that we have data on $\check{r} \geq 1$ markets $r = 1, \dots, \check{r}$. For each market r , we observe the number of active firms $N_{r,t}$ and the demand state $C_{r,t}$ in each period $t = 1, \dots, \check{t}$; for some $\check{t} \geq 2$.⁴ We also observe some time-invariant characteristics of each market r , which we store in a row vector X_r . However, we have no data on the cost shocks $W_{r,t}$.

We assume that $(\{N_{r,t}, C_{r,t}; t = 1, \dots, \check{t}\}, X_r)$ is distributed independently across markets r .⁵ The initial conditions $(N_{r,1}, C_{r,1}, X_r)$ are drawn from some distribution that is common across markets r . Conditional on $(N_{r,1}, C_{r,1}, X_r)$, industry dynamics $\{N_{r,t}, C_{r,t}; t = 2, \dots, \check{t}\}$ follow the transition rules implied by Section 2’s unique equilibrium, with primitives $\pi_r, \kappa_r, \varphi_r, \rho_r, G_{C,r}$, and $G_{W,r}$. The primitives may vary across markets r , but with X_r only. We make this explicit by parameterizing $\pi_r(\cdot, \cdot) = \pi(\cdot, \cdot | X_r, \theta_P)$, $\kappa_r = \kappa(X_r, \theta_P)$, $\varphi_r(\cdot) = \varphi(\cdot | X_r, \theta_P)$, 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 .⁶

4.2 Likelihood

We focus on inferring the structural parameters $\theta \equiv (\theta_P, \theta_C, \theta_W)$ from the conditional likelihood $\mathcal{L}(\theta)$ of θ for data on market dynamics $\{N_{r,t}, C_{r,t}; t = 2, \dots, \check{t}; r = 1, \dots, \check{r}\}$ given the initial conditions $(N_{r,1}, C_{r,1}, X_r; r = 1, \dots, \check{r})$.⁷ Using the model’s Markov structure and conditional independence, this likelihood can be written as $\mathcal{L}(\theta) =$

⁴In a typical application, like our empirical illustration in Section 6, \check{t} would be small and \check{r} would be large. However, the model can be estimated with data on a sufficiently long time series for a single market; that is, with \check{t} large and $\check{r} = 1$.

⁵Our estimation procedure can be straightforwardly extended to allow for observed (to the econometrician) time-varying covariates that are common across markets, such as business cycle indicators, provided that we make appropriate assumptions on their evolution.

⁶These assumptions rule out persistent unobserved (to the econometrician) heterogeneity in the markets’ primitives. Extending our NFXP procedure to account for unobserved heterogeneity 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 the persistent unobservables. See Footnote 7 below for further discussion of this point.

⁷We 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 and covariates are drawn from their ergodic distribution in the dynamic model, which is fully determined by θ . This would allow us to develop a more efficient estimator, at the price of robustness. Moreover, it would allow us to deal with the initial conditions problems alluded to in Footnote 6 in an extension with unobserved heterogeneity across markets.

$\mathcal{L}_C(\theta_C) \cdot \mathcal{L}_N(\theta)$, with

$$\mathcal{L}_C(\theta_C) \equiv \prod_{r=1}^{\tilde{r}} \prod_{t=2}^{\tilde{t}} g_C(C_{r,t} | C_{r,t-1}; X_r, \theta_C), \quad (12)$$

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

$$\mathcal{L}_N(\theta) \equiv \prod_{r=1}^{\tilde{r}} \prod_{t=2}^{\tilde{t}} p(N_{r,t} | N_{r,t-1}, C_{r,t}; X_r, \theta), \quad (13)$$

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

Note that $\mathcal{L}_C(\theta_C)$ 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 and substitute these into (13). To this end, for each value of X_r , we first compute the equilibrium post-survival values $v_{S,r}$ corresponding to the primitives implied by X_r and θ . From these, we obtain cost-shock thresholds for entry and *sure* survival, defined by

$$\bar{w}_{E,r}(n, c) \equiv \log v_{S,r}(n, c) - \log(\kappa_r + \varphi_r), \quad (14)$$

and

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

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 equilibrium mixing of survival decisions. For example, the probability that the number of firms remains unchanged at n sums the probability that the cost shock falls into $[\bar{w}_{E,r}(n + 1, c), \bar{w}_{S,r}(n, c)]$ with the probability that it instead equals some $W_{r,t} \in (\bar{w}_{S,r}(1, c), \bar{w}_{S,r}(n, c)]$ and that all the n firms outcomes from mixed strategy $a_S(n, c, W_{r,t})$ dictate survival. Similar complications arise for $n' \in \{1, \dots, n - 1\}$, which requires firms to follow a non-trivial mixed strategy, and for $n' = 0$, which

can occur if either the incumbents all choose certain exit or they are mixing nontrivially and all exit by chance. Accounting for the influence of mixed strategies on $p(n'|n, c; X_r, \theta)$ in these last three cases is tedious but straightforward.

4.3 Identification

Before proceeding to use the likelihood function for the model's estimation, we first analyze to what extent we could determine θ uniquely if we would observe not just our sample of \check{r} market histories $(\{N_{r,t}, C_{r,t}; t = 1, \dots, \check{t}\}, X_r)$ but the population $(\{N_t, C_t; t = 1, \dots, \check{t}\}, X)$ from which it is drawn. Specifically, suppose that we know the distribution of (N', C') conditional on $(N, C, X) = (n, c, x)$ for all $n \in \mathbb{N}_0 \equiv \{0\} \cup \mathbb{N}$, $c \in \mathcal{C}$ and $x \in \mathcal{X}$.⁸ Throughout the remainder of this section, we keep conditioning on $X = x$ implicit, so the results demonstrate identification of the model's primitives as nonparametric functions of x .

To begin, note that there is little hope that we can uniquely determine ρ from this distribution: [Rust \(1994\)](#) established the non-identifiability of the discount rate in a decision theoretic model of dynamic discrete choice and his fundamental insight holds good in our model. Therefore, we will assume that auxiliary information that identifies ρ , such as the average borrowing rate for small businesses, is in hand. Next, note that $G_C(\cdot|c)$ can be identified directly with the given distribution of $C'|C = c$.⁹

The remaining primitives of interest are the model's fixed cost, κ , sunk cost φ , surplus function π , and the distribution G_W of the econometric error. Our identification argument for these parameters follows that of [Hotz et al. \(1994\)](#), who retrieve value functions by applying the inverse cumulative distribution function of the econometric error to observed choice probabilities. Since this strategy requires knowledge of G_W , we assume that this belongs to the parametric family

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

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

⁸We ignore information about θ in the initial conditions (N_1, C_0, X) , for reasons given in Footnote 7.

⁹Above, we specified this distribution as a function of a vector of parameters, θ_C . Such a parametric restriction might be of use when estimating using a finite sample, but it is not necessary for identification.

of profits, we also normalize the mean per-period fixed cost κ to one.

The identification argument begins by retrieving $\bar{w}_S(1, c)$, 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 $\bar{w}_E(n, c)$ from the probability of 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 $\bar{w}_E(1, c)$ in (14) and (15) can be used to identify the sunk cost of entry up to the scale parameter ω :

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

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.¹⁰ 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 $n_E \in \{2, \dots, \check{n}\}$,

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

By (15), this is equivalent to requiring that

$$v_S(1, c) = \dots = v_S(n_E - 1, c) > v_S(n_E, c)$$

¹⁰For example, normalizing a coefficient to equal one allows identification of the error's variance.

for some c and n_E . This is a very weak condition, in particular given Lemma 2's result that $v_S(n, c)$ is *always* weakly decreasing 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_E incumbents simultaneously exiting:

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

Because $\Pr[N' = 0 | N = n_E, C = c]$ and $\Pr[N' = 0 | N = 1, C = c]$ are known, this identifies the integral in the right hand side of (17).

We will now show that this integral can be written as a known monotone function of ω , so that, in turn, it identifies ω . To this end, first note that, using $v_S(1, c) = \dots = v_S(n_E - 1, c)$, we can explicitly solve (5) for the mixing probability $a_S(n_E, c, w)$:

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

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

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

with

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

Because k_1 and k_{n_E} are known,

$$\frac{\exp(\omega k_1) - \exp(\omega \varepsilon)}{\exp(\omega k_1) - \exp(\omega k_{n_E})} = \frac{1 - \exp(-\omega(k_1 - \varepsilon))}{1 - \exp(-\omega(k_1 - k_{n_E}))} \quad (19)$$

is a known function of ω . Moreover, it is straightforward to verify that it is strictly

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

With ω identified, we immediately recover φ and $\bar{w}_S(n, c) = \log v_S(n, c)$. With this value function, ω , the observed transition probabilities for C' and N' , and the assumed value for ρ , we can recover $\pi(n, c)$ as the unique solution to a linear functional equation constructed from the definition of a value function.

We summarize this result in a theorem.

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

$$\Pr[N' = 0|N = 1, C = c] = \dots = \Pr[N' = 0|N = n_E - 1, C = c] < \Pr[N' = 0|N = n_E, C = c].$$

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

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

We take three lessons away from this identification argument. First, it is possible to identify the model's parameters without examining the cross-sectional relationship between N and C that [Bresnahan and Reiss \(1990, 1991\)](#) use in their estimation. Second, estimation of our model need not follow the nested fixed point approach that we adopt. In the spirit of [Hotz et al. \(1994\)](#), we could instead estimate the equilibrium value functions directly from observed transition probabilities and from these deduce the underlying primitives. 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 $\pi(c, n)$ (for example by requiring linearity in c), we have found that these do not create straightforward restrictions on the equilibrium continuation values useful for extending static identification arguments to this dynamic setting. We do not

doubt the practical usefulness of restrictions on payoffs for estimation with a finite sample of data. We pursued the identification of ω through mixed strategies for its feasibility, not its novelty.

4.4 Estimation

We have created C++ and Matlab code for computing a full information maximum likelihood estimator of θ . As in Rust (1994), computation proceeds in three steps:

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

Note that the partial likelihood estimator $\tilde{\theta}$ computed in the first two steps is consistent, but not efficient. The third step's estimator $\hat{\theta}$ is asymptotically efficient. To compute estimated standard errors, we use the outer-product-of-the-gradient estimator of the (full) information matrix. In particular, we assume that R is large and T is small and use the average over markets of the outer products of the market-specific gradients, evaluated at $\hat{\theta}$.

The C++ code provides a full implementation of this three-step NFXP procedure for specifications with and without covariates. It uses Knitro for the optimization, with analytical gradients. We use this code for the Monte Carlo experiments in Section 5 and the empirical illustration in Section 6. The Matlab code provides a more user friendly implementation of the NFXP procedure that can be used as a sandbox for experimentation and teaching.

5 Monte Carlo Experiments

In this section we investigate the statistical properties and computational performance of our estimation procedure with Monte Carlo experiments. For these, we set the maximum number of firms entering any market to $\tilde{n} = 5$, let the cost shocks be log normally distributed with unit mean and scale parameter ω , normalize the mean per period fixed costs κ to one, and fix the discount factor

ρ at $\frac{1}{1.05}$. The statistical process governing the demand state has support on 200 grid points that are equally spaced on the logarithmic scale with distance d , $\{c_{[1]}, c_{[2]} = c_{[1]}e^d, \dots, c_{[200]} = c_{[1]}e^{200d}\}$. So that the growth of C_t is approximately normally distributed with mean μ and variance σ^2 , we follow [Tauchen \(1986\)](#) and specify the probability of transitioning to $c_{[i]}$ from $c_{[j]}$ for any $i = 2, \dots, 199$ and $j = 1, \dots, 200$ with

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

The probabilities of transitioning to the grid's end points equal $\Pr(C' = c_{[1]}|C = c_{[j]}) = \Phi\left(\frac{\log c_{[1]} - \log c_{[j]} + \frac{d}{2} - \mu}{\sigma}\right)$ and $\Pr(C' = c_{[200]}|C = c_{[j]}) = 1 - \Phi\left(\frac{\log c_{[200]} - \log c_{[j]} - \frac{d}{2} - \mu}{\sigma}\right)$ respectively. We set $\mu = 0$ and $\sigma = 0.02$.

Each Monte Carlo experiment consists of 1,000 synthetic samples. We use four different sample sizes, each of them with ten time periods and between 100 and 1,000 identical markets. We compute the equilibrium and simulate the evolution of (N, C) , beginning with a draw from the model's ergodic distribution. Since this specification includes no regressors in X , a single equilibrium calculation can support the likelihood function calculations for all of a sample's observations. We then use each sample to estimate the model's parameters with the three step procedure presented in [Section 4](#). The starting parameter vector used for the likelihood function's maximization equals a vector of ones multiplied by one random variable uniformly distributed on the interval $[1, 10]$. [Dubé et al. \(2012\)](#) caution that a nested fixed point algorithm can falsely converge when the tolerance criterion for the inner loop (which calculates the equilibrium) is set too loosely relative to that of the outer loop (which maximizes the likelihood function). We fix the convergence tolerance for the value function iteration at a value that is multiple orders of magnitude smaller than that for the likelihood maximization to avoid this potential pitfall.¹¹

We first simulate data from a model where the surplus function is parameterized as $\pi(c, n) = (c/n)k$, which means that per consumer surplus is constant in the number of active firms. We set the true values of k , φ , and ω , to 1.5, 10, and 1 respectively. [Table 1](#) reports the corresponding Monte Carlo experiments' results. Its first panel gives the averages of the 1,000 estimates for each parameter, and it shows that the NFXP estimator is essentially without bias, even for the sample with

¹¹We set the tolerance value to 10^{-10} for the inner loop and to 10^{-6} for the outer loop.

Table 1: Monte Carlo Results with Constant Profits

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

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

only 100 markets. The second panel reports the averages of the estimated standard errors. For the sample with 100 markets, the average estimated standard error for the estimate of the sunk cost is 2.924. Therefore, we would expect a 95 percent confidence interval to approximately correspond to (4, 16). This is possibly too wide for empirical usefulness, but the other estimates' standard errors are relatively small. As expected, increasing the sample size decreases the standard errors approximately at the rate $\sqrt{\check{r}}$. So for $\check{r} = 500$ the standard error on $\hat{\varphi}$ is only 1.254. The table's final panel reports the Monte Carlo estimates of 95% confidence intervals' coverage probabilities. With the exception of that for φ with $\check{r} = 100$, these are all within 1.5 probability points of their common nominal value. Apparently, the estimated standard errors provide accurate inference.

Table 2: Monte Carlo Results with Decreasing Profits

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

Note: Results of a Monte Carlo experiment using the three step NFXP estimator with 1,000 repetitions estimating the model with five profit parameters $k(1), k(2), \dots, k(5)$ and one entry cost parameter φ . 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 cost shock, ω , equals 1. Demand is discretized into 200 states. The demand process is governed by the drift parameter μ , which is set to zero, and the standard deviation σ , which equals 0.02. The bottom-most panel displays the fraction of samples for which the estimated 95% confidence interval contained the parameter's true value.

For our second set of simulations we parameterize the flow surplus function as $\pi(c, n) = (c/n)k(n)$, where $(k(1), k(2), k(3), k(4), k(5))$ is set to $(1.8, 1.4, 1.2, 1.0, 0.9)$. This specification has the average surplus per consumer decrease in the number of active firms.¹² Table 2 reports the results of the corresponding Monte Carlo experiments. Again, all parameter estimates are essentially without bias, the estimated standard errors are small enough to be empirically useful, and the 95% 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% confidence level in all of our Monte Carlo samples regardless of the sample size. Overall, we conclude that the NFXP procedure has the potential to be empirically useful. Verifying that potential with observations not created by the model is the subject of the next section.

Since our equilibrium computation algorithm finds fixed points to relatively low dimensional contraction mappings, one would expect the estimation procedure to be relatively fast. Table 3 shows that this in fact the case. Even in the largest of our Monte Carlo samples, the average computation of the maximum likelihood estimator takes about two minutes using the C++ code.

Su and Judd’s (2012) results suggest that we might be able to improve on the already rapid performance of our estimation procedure by using a mathematical programming with equilibrium constraints (MPEC) procedure in lieu of a nested fixed point algorithm. The MPEC estimator treats the value functions as a vector of nuisance parameters to be estimated subject to the equilibrium constraint implied by the sequence of Bellman equations and thereby omits the inner loop. We implemented the MPEC estimator of our model in C++ using analytical gradients of both the objective function and the constraints. The MPEC estimator always yielded the same estimates as our NFXP procedure, but we found it to be more than ten times slower than the NFXP implementation.¹³ MPEC’s relatively poor

¹²The parameter values for the Monte Carlo simulation are chosen such that the specification with decreasing average surplus per consumer generates a realistic distribution of firms per market. No firm is active in about 5% of the markets, a monopolist serves about 30% of the markets, and five firms serve about 5% of the markets. In contrast, the previous specification with constant per consumer surplus generates a distribution with an additional local mode at five firms, because competition does not get “tougher” when the number of firms increases.

¹³Su and Judd (2012) emphasize the usefulness of passing “sparsity patterns” to the optimizer,

Table 3: 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 (in seconds)	30.46	38.27	50.17	73.97
<i>one entry cost parameter, five profit parameters</i>				
time per run (in seconds)	65.79	74.90	90.50	123.73

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

performance arises from the computation of the objective function’s gradients with respect to the nuisance parameters, which requires repeatedly retrieving information from very large and relatively dense matrices. These computational challenges might not be insurmountable, but our NFXP estimator seems to balance the costs of programmer time and execution time well.

6 Empirical Illustration

The Monte Carlo results suggest that we can use observations from a few hundred markets over ten time periods to estimate the model’s parameters accurately enough for differentiating between economically distinct hypotheses about the magnitude of sunk costs and the effects of additional competition on producers’ surplus. In this section we take the model beyond “data” of its own making and estimate its parameters with observations from the Motion Picture Theaters industry. Although we intend this application to be illustrative, we have tried to make our results slightly more useful by including in X_r a measure of the geographic diversity of a market’s consumers. [Davis \(2006\)](#) finds that theater location substantially influences consumers’ decisions about whether and where to attend film screenings. Indeed, one’s probability of attending a given theater declines substantially when the travel

which indicate which derivatives of the constraints with respect to the nuisance parameters are identically zero. We followed their advise, and we also initialized the nuisance parameters at their true values. For the parameters of interest, we used the same starting values as in the NFXP estimation.

distance moves from between zero and five miles to between five and ten miles.¹⁴ Davis (2002) finds that the concomitant low cross-price elasticities from such spatial preferences impact firms’ pricing behavior. Using observations from a New Haven area theater that experimented with a temporary price cut, Davis (2002) found that rivals five to seven miles away responded with lower prices but those ten to twelve miles away did not. This spatial differentiation of Motion Picture Theaters makes market definitions based on readily available geographic data plausibly applicable to this industry. Since we place a measure of how far apart consumers are from each other into X_r , our estimates quantify how the spatial structure of demand impacts the level of profits and the toughness of competition.

6.1 The Data

This analysis equates a market with a Micropolitan Statistical Area (μ SA) as defined by the Office of Management and Budget. Each one is based around of an urban core of at least 10,000 but less than 50,000 inhabitants.¹⁵ For this paper, 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 circle, that all travel within a tract must pass through its center, and that straightline roads connect the tracts’ centers to each other. We then measured *geographic preference diversity* with the average distance between two randomly-chosen residents of the μ SA. With the same methodology, 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 4 displays the five standard quantiles for population,

¹⁴See the logit model estimates reported in Table 5 of Davis (2006).

¹⁵We use the release of the “Annual Estimates of the Population of Metropolitan and Micropolitan Statistical Areas from April 1, 2000 to July 1, 2009” from the US Census Bureau as baseline for our analysis, which includes information on 574 μ SAs.

Table 4: 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	32.77	37.10	41.29	45.83	51.10
Geographic Preference Diversity	9.24	11.17	13.37	16.81	21.23
Geographic Market Isolation	23.94	28.77	37.61	51.83	72.70

Note: All variables are measured as of 2000 for the 573 μ SAs in our sample. Population is expressed in thousands of people, Median Household Income is expressed in thousands of dollars, and the remaining variables are expressed in miles. Please see the text for further details.

median per capita income, geographic preference diversity, and geographic market isolation. Population varies by about a factor of four from the 10th to the 90th percentiles. For the United States as a whole, median household income equalled \$42,148 in 2000. This is very close to the median value across the μ SA's, \$41,288. About 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. The 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 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 in each μ SA from the County Business Patterns (CBP), beginning in 2000 and ending in 2009. The top panel of [Table 5](#) reports the frequencies of the number of theaters across all of the μ SA-year observations. No theaters serve the market

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

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

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

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

in about twenty percent of the observations, a single theater serves about half of them, and about thirty percent of our observations have more than one theater. The maximum number of theaters observed is nine, but only 4.9 percent of the observations have four or more. Each row of Table 5 reports the observed frequencies of the number of theaters conditional on its previous year's value. Regardless of the initial number of theaters, the most common outcome is for it to remain unchanged. Nevertheless, the number of theaters changes in about 15 percent of the observed annual transitions.

In addition to this panel of producer counts, our estimation requires repeated measurements of the demand indicator C and cross-sectional measurements of time-invariant regressors X . The time-invariant regressors we employ are the median income and geographic preference diversity from the 2000 Census described above as well as dummy variables indicating membership in the nine U.S. Census Divisions. For the time-varying demand indicator, we use annual population for each μSA as published by the Census Bureau. For our sample from 2000 to 2009, the mean and standard deviation for the annual population growth rate in this sample 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 overwhelmingly reflect differences in measurement methodology rather than true population adjustments, we end our estimation sample in 2009.

6.2 **Bresnahan and Reiss (1990, 1991) Ordered Probits**

Our model is an infinite-horizon extension of the two-period free-entry model of [Bresnahan and Reiss \(1990, 1991\)](#), so it behooves us to estimate that simpler model before proceeding with our own. That model shares the specification for one-period profits used by our Monte Carlo experiments, with the effects of market-specific regressors X_r captured by a log-linear term.

$$\pi_r(n, c) = \exp(\beta' X_r) \frac{c}{n} k(n) \tag{20}$$

Just as in our model, the cost of entry has a log-normal distribution. Each market begins each period with no incumbents, and entry occurs until its cost exceeds its benefit. With this, the probability of observing N_r firms serving a market with C_r customers and other regressors X_r equals

$$\Pr[N_r = n | C_r, X_r] = \Phi \left(\frac{\ln(C_r) + \beta' X_r + \ln(k(n)/n)}{\omega} \right) - \Phi \left(\frac{\ln(C_r) + \beta' X_r + \ln(k(n+1)/(n+1))}{\omega} \right),$$

where ω is the standard deviation for the cost of entry's logarithm. Estimation of this ordered-probit specification is straightforward.

Table 6 reports the results from estimating this model using our data. The regressors X_r include the logarithm of median income as well as Census Division dummies, but the table omits their corresponding coefficient estimates. Because there are few observations with four or more firms, we truncate the number of firms (our dependent variable) from above at four. The table's first column reports estimates based on data from using the 5730 observations from all of the sample's μ SAs. The next two columns report estimates from a model in which $k(n)$ was allowed to depend on whether geographic preference diversity was above or below its median value, 13.4 miles. The first line reports $1/(k(1) \times 10^3)$, which gives the

Table 6: Estimates of [Bresnahan and Reiss \(1990, 1991\)](#) Ordered-Probit Model

	All μ SAs	Sample Selection	
		Diversity > 13.4 miles	Diversity \leq 13.4 miles
$1/(k(1) \times 10^3)$	12.83 (1.22)	15.02 (1.42)	11.31 (1.09)
$k(2)/k(1)$	0.50 (0.02)	0.61 (0.03)	0.43 (0.02)
$k(3)/k(2)$	0.71 (0.02)	0.75 (0.02)	0.65 (0.03)
$k(4)/k(3)$	0.85 (0.02)	0.87 (0.02)	0.80 (0.04)
ω	0.85 (0.02)		0.83 (0.03)
Observations	5730	2860	2870

Note: Standard errors are reported in parentheses. In the ordered probit model, $1/(k(1) \times 10^3)$ gives the median value of the population (in units of 1,000) above which at least one firm enters the market. The standard deviation of the entry cost's logarithm is ω . The coefficient on each μ SA's log population is constrained to equal one, and both specifications included the logarithm of median income and census division dummies as regressors. Please see the text for further details.

population (in thousands of people) that is required to support one producer at the entry cost's median value, 1, in a market with characteristics $X_r = 0$. For the specification that ignores geographic product diversity, this is slightly under 13,000 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 about 15,000 people to support a monopolist in a μ SA with geographic preference diversity above the median and only about 11,300 to support a monopolist in a μ SA with preference diversity below the median. A Wald test indicates that this difference is statistically significant at every conventional level, but we take this with a grain of salt since the model assumes that the errors are *i.i.d.* over time for a given μ SA.

The estimates of $k(n+1)/k(n)$ indicate very tough competition. Those from the first specification indicate that duopolists receive half of a monopolist's producers'

surplus per customer. Adding even more competitors erodes this surplus further, but at a decreasing rate. 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 support this prior conjecture. For high-diversity markets, the duopolists' producers' surplus per customer equals about 60 percent of a monopolist's, while the analogous estimate for low-diversity markets is only 43 percent. Again, a Wald test indicates that this difference is statistically significant, but we caution against uncritically accepting its conclusion.

6.3 The Dynamic Model

These estimates of tough competition are difficult to reconcile with other evidence from this industry. Movie theaters potentially compete in the market for customers and the market for films. [Davis \(2005\)](#) provides evidence on competition for customers from regressions of theaters' admissions prices against indicators of the presence of other theaters at various distances using data from large (relative to μ SAs) U.S. cities in the 1990s. Based on both across-market and within-market-over-time variation, he concludes that

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

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

give owners indirect and extralegal control over box-office prices.¹⁶ Supporting the view that film distributors constrain theaters’ pricing choices, [Davis \(2006\)](#) finds that

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

Although the literature says little about theaters’ potential monopsony power in the market for films, it is possible that it is substantial. On the other hand, misspecification from omitting sunk costs and non-trivial dynamic interactions between theater owners might be responsible for the finding of substantial toughness of competition where in fact there is none.¹⁷ The estimation of our dynamically richer model can determine whether such a misspecification is responsible for this conflict.

This dynamic model’s estimation uses the demand-process specification from our Monte Carlo exercise. We restrict $C_{r,t}$ to a grid of 200 points equally spaced on a logarithmic scale. Its minimum value equals the minimum population observed in our data, 11,011, divided by 1.25. Analogously, its maximum value equals the maximum population, 197,912 multiplied by 1.25. For estimation, we replace each observation of 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 give the cost shocks’ logarithms a normal distribution with standard deviation ω and mean $-\omega^2/2$, normalize κ to one, and fix the discount factor ρ at $\frac{1}{1.05}$. The specification for the producers’ surplus function comes from [\(20\)](#). As with the ordered probit model, we include the logarithm of the μ SA’s median income in 2000 and Census Division dummies in X_r .¹⁸

[Table 7](#) reports the estimated parameters for two specifications that mimic those applied to the [Bresnahan and Reiss \(1990, 1991\)](#) ordered probit model. In the first specification (that ignores geographic preference diversity), the full-information maximum likelihood estimates of the demand process drift and innovation standard

¹⁶[Orbach and Einav \(2007\)](#) review the legal environment in which theater owners negotiate with film distributors and set admissions prices.

¹⁷[Abbring and Campbell \(2010\)](#) illustrate this possibility with a Monte Carlo study based on their Last-In First-Out model of oligopoly.

¹⁸Our estimation code constrains the flow surplus to weakly decrease with the number of firms, as required by the monotonicity assumption in [Section 2](#). This constraint does not bind at the maximum-likelihood estimates.

deviation are very close to the unconditional sample mean and standard deviation of population growth; 0.34 and 1.21 percent versus 0.34 and 1.11 percent. The coefficients in β are jointly and (with the exceptions of those multiplying three division dummies) individually significant. The average realization of the sunk cost of entry, φ is over fifty times the average realization of the 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 the cost shock is low. To calculate more informative measures of fixed and sunk costs, we simulated the estimated model with median income set equal to its (cross μ SA) mean and the dummy for the East North Central division (which includes Illinois, Michigan, Wisconsin, Indiana, and Ohio) set to one. In the simulation, the average fixed cost of continuation and sunk cost of entry *paid* were 0.72 and 1.59. The estimates of all these parameters from the specification that accounts for geographic preference diversity are similar, but not identical, to these baseline estimates.

Although Table 7 contains estimates of $k(n)$ for $n = 1, \dots, 4$; it is easier to interpret these if they are presented as in Table 6. Table 8 reports these transformations and their standard errors. In Table 6, $1/(k(1) \times 10^3)$ gave the population (in thousands) necessary to attract one firm at the cost shock's median realization in a market with characteristics $X_r = 0$. In this dynamic model, the same object has only the narrower interpretation of the population that sets the monopolist's current profit equal to zero at the cost shock's median realization in a market with characteristics $X_r = 0$. The baseline specification's estimate of this is 16,600 people. This greatly exceeds the analogous estimate from the ordered probit model, 12,830 people. [Abbring and Campbell's \(2010\)](#) simulation results suggest that the option value of delaying exit lies behind this difference. Since exit is irreversible, firms exit only after producer surplus earned is much less than the fixed cost of operating. The continued operation of such firms biases the static model's estimate of $1/(k(1) \times 10^3)$ downwards. The analogous estimates from the specification that accounts for geographic preference diversity equal 18,520 and 16,250 people for high-diversity and low-diversity μ SA's, respectively. A Wald test indicates that this difference is significant at the 10 percent level. Just as with the static ordered probit model, concentrating customers' locations makes a market more profitable for a monopolist.

Table 7: Parameter Estimates of the Dynamic Model

	All μ SAs	Sample Selection	
		Diversity > 13.4 miles	Diversity \leq 13.4 miles
$k(1) \times 10^5$	6.01 (0.80)	5.40 (0.73)	6.16 (0.86)
$k(2) \times 10^5$	3.23 (0.47)	3.24 (0.48)	2.93 (0.49)
$k(3) \times 10^5$	2.65 (0.41)	2.72 (0.42)	2.29 (0.43)
$k(4) \times 10^5$	2.04 (0.33)	2.16 (0.35)	1.53 (0.35)
φ	50.66 (9.68)		48.36 (9.56)
Median Income	0.87 (0.17)		0.84 (0.17)
Mid Atlantic	-0.63 (0.15)		-0.59 (0.15)
East North Central	-0.48 (0.14)		-0.44 (0.14)
West North Central	0.06 (0.16)		0.10 (0.16)
South Atlantic	-0.73 (0.15)		-0.71 (0.15)
East South Central	-0.51 (0.16)		-0.49 (0.16)
West South	-0.32 (0.15)		-0.28 (0.15)
Mountain	-0.12 (0.15)		-0.07 (0.15)
Pacific	-0.10 (0.15)		-0.08 (0.15)
ω	1.75 (0.07)		1.74 (0.08)
$\mu \times 10^2$	0.34 (0.02)		0.34 (0.02)
$\sigma \times 10^2$	1.21 (0.01)		1.21 (0.01)
$-\mathcal{L}$	9199.22		9192.20
Number of Markets	573	287	286

Note: Standard errors are reported in parentheses. The data include 573 μ SAs from 2000 to 2009. \tilde{n} equals nine, which is the maximum of the number of active firms observed in the data. The geographic entities are dummy variables that refer to Census Divisions. The baseline is New England.

Table 8: Estimates of the Toughness of Competition in the Dynamic Model

	All μ SAs	Sample Selection	
		Diversity > 13.4 miles	Diversity \leq 13.4 miles
$1/(k(1) \times 10^3)$	16.65 (2.21)	18.52 (2.51)	16.25 (2.27)
$k(2)/k(1)$	0.54 (0.14)	0.60 (0.14)	0.48 (0.21)
$k(3)/k(2)$	0.82 (0.06)	0.84 (0.06)	0.78 (0.10)
$k(4)/k(3)$	0.77 (0.08)	0.79 (0.08)	0.67 (0.21)
Number of Markets	573	287	286

Note: This Table is based on the model's estimates as reported in Table 7. Standard errors are reported in parentheses. The ratio $1/(k(1) \times 10^3)$ can be interpreted as the population (in units of 1,000) that sets the monopolist's current profit equal to zero at the cost shock's median realization. The ratio $k(n+1)/k(n)$ is an indicator of the toughness of competition. Please see the text for further details.

Leaving aside this bias in the level of profits, the dynamic model's estimates of how profits *change* with entry closely resemble the analogous estimates from the static ordered probit. In the baseline specification, duopolists' producers' surplus per consumer equals 54 percent of a monopolist's (as opposed to 50 percent), and the analogous estimates from the alternative specification equal 60 and 48 percent respectively (as opposed to 61 and 43 percent). However, the difference between these two estimates is not statistically significant at any conventional level. The remaining estimates in Table 8 also differ little from their analogues in Table 6. We conclude that the dynamic model finds evidence that lowering preference diversity substantially increases profits, but there is little evidence that it intensifies competition between theaters.

Apparently, adding competitors strongly reduces producers' surplus per customer. Since [Davis \(2002, 2006\)](#) found little effect of competition on ticket prices, our results suggest that adding theaters increases competition in the market for screening rights. Such a finding, if supported by more direct observations of vertical contracts, would open an interesting policy question. How should a local social planner that licenses entry trade off the benefits of additional product diversity against the cost of transferring surplus to film producers. Although our model might contribute to answering this and similar policy questions, its full consideration lies well beyond the scope of this illustrative estimation exercise.

7 Conclusion

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

We anticipate three applications of our model and its maximum-likelihood estimator. First, they can be used to estimate the impact of observable cross-market heterogeneity on the primitive determinants of industry dynamics. Our examination of geographic preference diversity above exemplifies such an application. Second, our model is simple enough for inclusion as a moving part in general equilibrium models with entry, exit, and endogenous markups, such as [Jaimovich’s \(2007\)](#). Third, the estimated model can serve as a point of departure for an analysis with a more computationally and theoretically demanding model. By estimating our model first, one can gain familiarity with the industry’s dynamics and obtain starting values for homotopy-based estimation and equilibrium calculation.

Appendix

The model in the main text embodies structure on the stochastic processes for profitability shocks and how they influence profits that contributes nothing to its theoretical analysis but nevertheless makes maximum likelihood estimation more tractable. This appendix presents a *general model* without this structure that encompasses the *special* model of the main text, and it proves the appropriate generalizations of Lemmas 1 and 2, Corollaries 1 and 2, and Theorem 1.

A Primitives

In the general model, firms' profits depend on $Y_t \in \mathcal{Y}$, which could be vector valued. Figure 2 gives the model's recursive extensive form. Period t starts in state (N_t, Y_{t-1}) , and then Y_t is drawn from the Markov transition distribution $\tilde{G}(\cdot|Y_{t-1})$. (Here and throughout this appendix, we place a tilde over any primitive object with a similar name in the special model.) Next, each of the N_t incumbent firms 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 the entry queue. In the entry subgame of period t , firm (t, j) pays the sunk cost $\tilde{\varphi}(N_t + j, Y_t)$ upon entry. Otherwise, the potential entrant earns 0. Progressing to the period t survival subgame, an active firm choosing survival incurs no cost *during period t* . The expected profits from operating in period $t + 1$ subsume the special model's costs of continuation.

The restrictions we place on $\tilde{\pi}(n, y)$ are

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

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

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

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

Assumption A1 differs slightly from the uniform bound of $\mathbb{E}[\pi(n, C')|C = c]$ in the text because $\tilde{\pi}(n, y)$ encompasses the costs of continuation, but it serves the same function of restricting equilibrium post-entry expected continuation values to the space of bounded non-negative functions. The remaining assumptions have direct analogues in the text.




Figure 2: The General Model's Recursive Extensive Form

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

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

B Equilibrium

For both a potential entrant and a firm contemplating survival, the payoff-relevant variables are the number of firms committed to play that period's survival subgame

and the current state of demand. A Markovian strategy is a pair of functions, $\tilde{a}_E : N^+ \times \mathcal{Y} \rightarrow \{0, 1\}$ and $\tilde{a}_S : N^+ \times \mathcal{Y} \rightarrow [0, 1]$. An equilibrium strategy and its associated continuation values $\tilde{v}_E(n, y)$ and $\tilde{v}_S(n', y)$ satisfy conditions analogous to the special model's equations (1), (2), (3) and (4):

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

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

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

and

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

The expectation operators condition on the deciding firm choosing continuation and all other firms using the strategy in the operator's subscript. As in (3), the M in (23) denotes the number of active firms if the potential entrant chooses entry.

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

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

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

Corollary 1 *Let \tilde{v}_S be the post-survival value function associated with a symmetric Markov-perfect equilibrium. Consider the one-shot survival game in which n_E firms simultaneously choose between survival and exit (as in the survival subgame of Figure 2), each of the n' survivors earns $\tilde{v}_S(n', y)$, and each exiting firm earns zero. This game has a unique symmetric Nash equilibrium, possibly in mixed strategies.*

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

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

To demonstrate equilibrium existence and uniqueness in the general model using these results, begin with a number of incumbent firms n , hypothesized values for the equilibrium post-survival continuation value $\tilde{w}_S(n', y)$ and entry strategy $\tilde{\alpha}_E(n', y)$ for all $n' \in \{n+1, \dots, \tilde{n}\}$ and $y \in \mathcal{Y}$. With the entry strategy, construct the transition rule

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

with

$$\tilde{\Lambda}_E(n', n, y) \equiv \prod_{j=n+1}^{n'} \tilde{\alpha}_E(j, y).$$

Lemma 1 tells us that for any strategy consistent with an equilibrium, $\mu(n, y) \leq \tilde{n}$ if $n \leq \tilde{n}$.

From Corollary 2, we know that all n incumbent firms choose survival with certainty whenever the individual value to joint continuation is positive. Suppose that whenever $\mu(n, y) = n$, incumbents expect to receive a (trial) post-survival continuation value $f(y)$. When entry brings the number of firms to n' , they expect to receive $\tilde{w}_S(n', y)$. Under these conditions, the Bellman operator used to update $f(y)$ is

$$\begin{aligned} T_n(f)(y) &= \rho \mathbb{E}[\tilde{\pi}(n, Y')] \\ &\quad + \max\{0, f(Y') + \mathbb{1}\{\mu(n, Y') > n\}(\tilde{w}_S(\mu(n, Y'), Y') - f(Y'))\} | Y = y \end{aligned}$$

If $\tilde{w}_S(n', y) \in [0, \frac{\hat{\pi}}{1-\rho}]$ for all $n' > n$ and $y \in \mathcal{Y}$, then T_n maps the set $\{f : \mathcal{Y} \rightarrow [0, \frac{\hat{\pi}}{1-\rho}]\}$ into itself. Furthermore, $T_n(f)$ satisfies Blackwell's sufficient conditions for a contraction mapping, so it has a unique fixed point $f^*(y)$. Of course, this fixed point depends on the hypothesized values of $\tilde{w}_S(n', \cdot)$ and $\tilde{\alpha}_E(n', \cdot)$.

As noted in the text, this dynamic program is the centerpiece of our algorithm for equilibrium calculation. Algorithm 1 presents it in flowchart form. It begins with initializing n at \tilde{n} and assigning the initial values of zero to the equilibrium strategies, continuation values, and a dummy function $f^* : \mathcal{Y} \rightarrow [0, \frac{\hat{\pi}}{1-\rho}]$. The next step initializes the transition rule $\mu(n, y)$ to the constant \tilde{n} and the following step applies Bellman equation iteration to calculate the fixed point to $T_{\tilde{n}}$. The result

gets assigned to $\tilde{w}_S(\tilde{n}, \cdot)$ and is used to construct the corresponding the hypothesized post-entry continuation value $\tilde{w}_E(\tilde{n}, \cdot)$. Then, the entry strategy $\tilde{\alpha}_E(\tilde{n}, y)$ gets set to an indicator for positive post-entry payoffs. A loop repeats these calculations for all valid lower values of n , using the results of earlier calculations to set $\mu(n, y)$. After completing this first loop, a second loop sets the survival strategy $\tilde{\alpha}_S(n, y)$ to the highest probability consistent with equilibrium in the survival subgame given the candidate continuation values in memory. (Because these calculations do not build on each other, the second loop can be parallelized.) The appropriate generalization of Theorem 1 to this framework states that the candidate equilibrium strategies and payoffs arising from Algorithm 1 correspond to the unique Markov-perfect Nash equilibrium that defaults to inactivity.

Theorem 1 (Equilibrium existence and uniqueness in the general model)

There exists a unique symmetric Markov-perfect equilibrium that defaults to inactivity. Algorithm 1 computes it and its corresponding post-entry and post-survival continuation values.


C Proofs

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

Proof of Lemma 1. Let N_t denote the random sequence of the number of active firms at the beginning of period t and let M_t denote the random sequence of the number of active firms just before period t 's survival subgame. These arise from equilibrium play conditional on the initial state (m_0, y_0) . Define the random time τ (with $0 \leq \tau \leq \infty$) as the first period in which the firms playing the survival game choose exit with a positive probability. That is,

$$\tau = \min\{\{t : \tilde{\alpha}_S(M_t, Y_t) < 1\} \cup \{\infty\}\}.$$

If indeed $\tilde{\alpha}_S(m_0, y_0) < 1$ as asserted, then $\tau = 0$. Suppose to the contrary that $\tilde{\alpha}_S(m_0, y_0) = 1$, so $\Pr[\tau > 0] = 1$. By definition, exit can occur only in or after period τ , so we know that $N_{t+1} = M_t \geq m_0$ for $t < \tau$. Since $m_0 > \tilde{n}$, Assumption A2 implies $\tilde{\pi}(N_t, Y_t) < 0$ for $t \leq \tau$. If the realization of τ is infinite, then the incumbent firms



Algorithm 1: Equilibrium Calculation for the General Model

receive an infinite sequence of strictly negative payoffs. If instead the realization of τ is finite, then the incumbent firms receive a finite sequence of strictly negative payoffs followed by a zero expected continuation value from playing the period- τ survival subgame. Therefore, the expectation of the discounted sum of payoffs, $\tilde{v}_S(m_0, y_0)$ must be strictly negative. Since any incumbent firm can raise its payoff to zero by choosing certain exit, the supposition that $\tilde{a}_S(m_0, y_0) = 1$ must be incorrect. If $\tilde{\varphi}(m_0 + 1, y) > 0$, then the first potential entrant (firm $(0, 1)$) maximizes its payoff by staying out of the industry and earning zero. If instead $\tilde{\varphi}(m_0 + 1, y) = 0$, then the assumption that the equilibrium strategy defaults to inactivity dictates the same action. In either case, $\tilde{a}_E(m_0, y_0) = 0$ as asserted. ■

Proof of Lemma 2. We need to show that for all $y_0 \in \mathcal{Y}$ and all $n_1 \in \mathbb{N}$, $\tilde{v}_S(n_1, y_0) \geq \tilde{v}_S(n_1 + 1, y_0)$.

Let N_t and N_t^+ denote the number of firms at the beginning of period t from the games initialized with n_1 and $n_1 + 1$ firms respectively. Below, we refer to N_t and N_t^+ as the outcomes of the *original* and *perturbed* games. The corresponding numbers of active firms just before period t 's survival subgame are M_t and M_t^+ . The original and perturbed games share a common realization of demand shocks, y_0, Y_1, \dots

Define the random times τ and τ^+ as in the proof of Lemma 1.

$$\begin{aligned}\tau &= \min\{t : \tilde{a}_S(M_t, Y_t) < 1\} \cup \{\infty\} \\ \tau^+ &= \min\{t : \tilde{a}_S(M_t^+, Y_t) < 1\} \cup \{\infty\}\end{aligned}$$

By these definitions, N_t and N_t^+ are weakly increasing sequences for all $t < \tau$ and $t < \tau^+$, respectively. We know that $N_t \leq N_t^+$ for all $t \leq \tau^+$, because otherwise the two games would have potential entrants in the same payoff-relevant states making different entry decisions. That would violate the assumption that the equilibrium strategy is Markovian.

We first wish to show that $\tau^+ \leq \tau$ always. Suppose to the contrary that there exists realizations of $\{Y_t\}_{t=0}^\infty$, $\{M_t\}_{t=0}^\infty$, $\{N_t\}_{t=0}^\infty$, $\{M_t^+\}_{t=0}^\infty$, and $\{N_t^+\}_{t=0}^\infty$ such that $\tau < \tau^+$. In this case, $N_t < N_t^+$ for $t \leq \tau$. Otherwise the original and perturbed games would have the same payoff relevant state at some date before τ and so would evolve identically thereafter. This would imply that $\tau = \tau^+$.

By definition, $\tilde{a}_S(M_\tau, Y_\tau) < 1$. Suppose that one of the M_τ incumbents deviates from this strategy and chooses certain survival, and let \mathcal{N}_t denote the sequence of

firm numbers that arises in that game after that deviation conditional on (i) the deviating incumbent surviving until period t and (ii) all other players following the equilibrium strategy. Since $\tilde{a}_S(M_\tau, Y_\tau) < 1$, we know that no entry occurs in period τ of the original game. (Otherwise, either an entrant pays a positive sunk cost for a zero continuation value or the proposed equilibrium strategy does not default to inactivity.) Therefore, $N_\tau = M_\tau$. Since some of these N_τ active firms might not survive, we have $\mathcal{N}_{\tau+1} \leq N_\tau$. Combining this with $N_\tau < N_\tau^+$ and $N_\tau^+ \leq N_{\tau+1}^+$ yields $\mathcal{N}_{\tau+1} < N_{\tau+1}^+$. Since the equilibrium strategy is Markovian, $\mathcal{N}_t \leq N_t^+$ for $t = \tau + 1, \dots, \tau^+$. Assumption A3 then implies that for all possible realizations of $\{Y_t\}_{t=\tau+1}^\infty$, $\tilde{\pi}(\mathcal{N}_t, Y_t) \geq \tilde{\pi}(N_t^+, Y_t)$ for $t = \tau + 1, \dots, \tau^+$. Furthermore, this inequality is strict for $t = \tau + 1$.

By choosing survival until the realization of τ^+ and exiting at that date, the deviating incumbent in the original game earns a strictly greater payoff than does an incumbent in the perturbed game. Since $\tau^+ > \tau$ we know that the expected discounted sum of these payoffs from the perturbed game, $\tilde{v}_S(N_\tau^+, Y_\tau)$, is strictly positive. Therefore, we conclude that the deviating incumbent in the original game receives a payoff strictly greater than that earned from following the equilibrium strategy, zero. This is inconsistent with the supposition of equilibrium, so we conclude that $\tau^+ \leq \tau$.

Since $\tau^+ \leq \tau$, period 0 incumbents in both the original and perturbed games survive until τ^+ ; the original game's incumbent earns a weakly higher payoff until τ^+ . The original game's incumbent must have a weakly higher period- τ^+ post-survival continuation value because the analogous continuation value from the perturbed game is non-positive. These two conclusions together imply that $v_S(n_1, y_0) \geq v_S(n_1 + 1, y_0)$. ■

Proof of Corollaries 1 and 2. Suppose that the survival subgame starts with m active firms. One of three mutually-exclusive cases holds good.

- $\tilde{v}_S(m, y) > 0$. In this case, Lemma 2 implies that $\tilde{v}_S(n', y) > 0$ for all $n' < m$, so $\tilde{a}_S(m, y) = 1$ is a dominant strategy. Therefore, this is the unique symmetric equilibrium strategy. The realized payoff to each of the active firms equals $\tilde{v}_S(m, y)$.
- $\tilde{v}_S(1, y) < 0$. In this case, Lemma 2 implies that $\tilde{v}_S(n', y) < 0$ for all $n' > 1$, so $\tilde{a}_S(m, y) = 0$ is a dominant strategy. Therefore, this is the unique symmetric

equilibrium strategy. The realized payoff to each of the active firms equals zero.

- $\tilde{v}_S(m, y) \leq 0$ and $\tilde{v}_S(1, y) \geq 0$. If $m = 1$, then $\tilde{a}_S(1, y) = 0$ by the default to inactivity assumption. Consider the case with $m > 1$. The expected payoff to survival given that all other firms survive with probability p is $v(p) \equiv \sum_{n'=1}^m \binom{m-1}{n'-1} p^{n'-1} (1-p)^{m-n'} \tilde{v}_S(n', y)$. A firm's best response to p is certain survival if $v(p) > 0$ and certain exit if $v(p) < 0$. If $v(p) = 0$ then any survival probability is a best response to p , including p itself. If $\tilde{v}_S(m, y) = \tilde{v}_S(1, y) = 0$, then $v(p) = 0$ for all p . In this case, the requirement that the equilibrium strategy defaults to inactivity requires that $p = 0$. If instead $\tilde{v}_S(1, y) > \tilde{v}_S(m, y)$, then $v(p)$ is strictly decreasing in p , and the conditions of the case can be expressed as $v(0) \geq 0$ and $v(1) \leq 0$. Therefore, there exists exactly one $p \in [0, 1]$ such that $v(p) = 0$. This is the unique equilibrium survival probability.

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

Proof of Theorem 1.

The proof is divided into three parts. First, we show that the candidate continuation values satisfy the monotonicity requirements of Lemma 2. Second, we demonstrate that these continuation values are indeed those associated with candidate equilibrium strategy. By construction, the candidate strategy admits no profitable one-shot deviation *if* the continuation values are indeed those given by the candidate equilibrium. Therefore, this second step establishes that the candidate strategy forms an equilibrium. Third, we apply Lemma 2 to demonstrate that an equilibrium's existence implies its uniqueness.

Fix $n \in \{1, 2, \dots, \tilde{n}\}$ and suppose that for all $n' \in \{1, 2, \dots, \tilde{n}\} \cap \{n+1, n+2, \dots\}$, we know that $\tilde{w}_S(n', y) \geq \tilde{w}_S(n'+1, y)$. (This condition is trivially true for $n = \tilde{n}$.) Consider evaluating T_n at the value of $f^*(y)$ in memory *after* the completion of the $n+1$ -indexed dynamic programming problem. Let $\mu(\cdot, \cdot)$ be the value of the transition rule at the conclusion of the algorithm. (Because $\mu(n', \cdot)$ gets set at the beginning of loop iteration with $n = n'$ and never is altered again, it has the same

value that was used to construct the Bellman operator $T_{n'}$.) With these definitions, we have

$$\begin{aligned}
T_n(f^*)(y) &= \rho \mathbb{E}[\tilde{\pi}(n, Y') + \max\{0, f^*(Y') \\
&\quad + \mathbb{1}\{\mu(n, Y') > n\}(\tilde{w}_S(\mu(n, Y'), Y') - f^*(Y'))\} | Y = y] \\
&> \rho \mathbb{E}[\tilde{\pi}(n+1, Y') + \max\{0, f^*(Y') \\
&\quad + \mathbb{1}\{\mu(n, Y') > n\}(\tilde{w}_S(\mu(n, Y'), Y') - f^*(Y'))\} | Y = y] \\
&= \rho \mathbb{E}[\tilde{\pi}(n+1, Y') + \max\{0, f^*(Y') \\
&\quad + \mathbb{1}\{\mu(n, Y') > n+1\}(\tilde{w}_S(\mu(n, Y'), Y') - f^*(Y'))\} | Y = y] \\
&= \rho \mathbb{E}[\tilde{\pi}(n+1, Y') + \max\{0, f^*(Y') \\
&\quad + \mathbb{1}\{\mu(n+1, Y') > n+1\}(\tilde{w}_S(\mu(n+1, Y'), Y') - f^*(Y'))\} | Y = y] \\
&= \tilde{w}_S(n+1, y)
\end{aligned}$$

The inequality follows from Assumption A3, and the first equality comes from the equivalence of $\tilde{w}_S(n+1, Y')$ with $f^*(Y')$. Since $\tilde{w}_S(n', Y)$ is weakly decreasing in n' for $n' > n$, so is $\tilde{w}_E(n', Y)$. Thus, we know that $\tilde{\alpha}_E(n', Y')$ weakly decreases with n' . Therefore, $\mu(n, Y') = \mu(n+1, Y')$ whenever $\mu(n, Y') > n+1$. This gives us the second equality. The operator T_n is a monotone contraction mapping, so $T_n(f^*)(n, y) > \tilde{w}_S(n+1, y)$ implies that its fixed point, $\tilde{w}_S(n, y)$, strictly exceeds $\tilde{w}_S(n+1, y)$. This is the desired monotonicity result for the proof's first part.

For the second part, we note that monotonicity of the candidate post-survival continuation value immediately implies that $\tilde{\alpha}_E(n+1, y) \leq \tilde{\alpha}_S(n, y)$ for all $n \in \mathbb{N}^+$. That is, the candidate strategy never prescribes entry immediately before an incumbent chooses a positive probability of exit. This result gives us a simple expression for the probability that the number of firms active after the current period's survival subgame is n' given the period starts with demand state y and n firms, one incumbent survives for sure, and all other firms follow the candidate

equilibrium strategy.

$$\lambda(n', n, y) = \begin{cases} 1 & \text{if } n' > n \text{ \& } \prod_{j=n+1}^{n'} \tilde{\alpha}_E(j, y) = 1, \\ & \text{\& } \tilde{\alpha}_E(n' + 1, y) = 0, \\ \binom{n-1}{n'-1} (\tilde{\alpha}_S(n, y))^{n'-1} & \text{if } 1 \leq n' \leq n \text{ \& } \tilde{\alpha}_E(n+1, y) = 0, \\ \times (1 - \tilde{\alpha}_S(n, y))^{n-n'} & \\ 0 & \text{otherwise.} \end{cases}$$

With this, we create a functional operator with its unique fixed point equal to the post-survival continuation value from playing the game when all players follow the candidate equilibrium strategy.

$$T^*(w)(n, y) = \rho \mathbb{E} \left[\tilde{\pi}(n, Y') + \tilde{\alpha}_S(n, Y') \sum_{n'=1}^{\infty} \lambda(n', n, Y') w(n', Y') \mid Y = y \right]$$

We wish to show that $T^*(\tilde{w}_S)(n, y) = \tilde{w}_S(n, y)$. For this, note whenever $\alpha_S(n, Y') = 0$, then the expected value of continuing when all other firms choose to exit for sure, $\tilde{w}(1, Y')$ is non-positive. If instead $\tilde{\alpha}(n, Y') \in (0, 1)$, then Algorithm 1 was used to compute this survival probability by equating the payoff to surviving for sure to that from exiting for sure, zero. Therefore, the second term inside the expectations operator of $T^*(\tilde{w}_S)(n, y)$ is non-zero only if $\tilde{\alpha}_S(n, Y') = 1$. In this case, $\lambda(n', n, Y) = 1$ for $n' = n + \sum_{j=0}^{\infty} \tilde{\alpha}_E(n + j, Y')$ and equals zero otherwise. Since the post-survival continuation value at this value of n' must weakly exceed zero, we have

$$\begin{aligned} T^*(\tilde{w}_S)(n, y) &= \rho \mathbb{E} \left[\tilde{\pi}(n, Y') + \max\{0, \tilde{w}_S(n + \sum_{j=0}^{\infty} \tilde{\alpha}_E(n + j, Y'), Y')\} \mid Y = y \right] \\ &= \rho \mathbb{E} [\tilde{\pi}(n, Y') + \max\{0, \tilde{w}_S(\mu^n(n, Y'), Y')\} \mid Y = y] \\ &= T_n(\tilde{w}_S(n, \cdot))(y) = \tilde{w}_S(n, y). \end{aligned}$$

The second equality follows from the definition of $\mu(n, y)$ in Algorithm 1 and the weak monotonicity of $\tilde{\alpha}_E(n', Y')$ in n' , and the third equality uses definition of $T_n(f)$.

Since the candidate continuation values are indeed those induced by the candidate equilibrium strategy, the candidate strategy's construction guarantees that there is no profitable one-shot deviation from the strategy given that all other firms follow it. We conclude that Algorithm 1's candidate equilibrium is

indeed an equilibrium. The remainder of this proof demonstrates its uniqueness. Corollary 1 says that there is a unique equilibrium survival strategy corresponding to every equilibrium continuation value function, and Lemma 2 implies that an equilibrium entry strategy must prescribe entry if and only if the continuation value from entering and immediately proceeding to the survival subgame is positive. Therefore, each pair of equilibrium continuation value functions \tilde{v}_S and \tilde{v}_E has exactly one corresponding equilibrium strategy. Corollary 2 requires any equilibrium continuation values to be fixed-points to the Bellman operators used in Algorithm 1. These fixed points are unique because the operators are contractions. Therefore, there is a unique pair of equilibrium continuation value functions (those constructed by Algorithm 1) and a unique corresponding equilibrium strategy. ■

References

- ABBRING, J. H. AND J. R. CAMPBELL (2010): “Last-In First-Out Oligopoly Dynamics,” *Econometrica*, 78, 1491–1527. [7](#), [31](#), [32](#)
- ABBRING, J. H., J. R. CAMPBELL, AND N. YANG (2010): “Simple Markov-Perfect Industry Dynamics,” Working Paper 2010–21, Federal Reserve Bank of Chicago, Chicago. [1](#), [2](#), [9](#)
- BAJARI, P. L., C. L. BENKARD, AND J. LEVIN (2007): “Estimating Dynamic Models of Imperfect Competition,” *Econometrica*, 75, 1331–1370. [2](#)
- BRESNAHAN, T. F. AND P. C. REISS (1990): “Entry in Monopoly Markets,” *Review of Economic Studies*, 57, 531–553. [1](#), [18](#), [28](#), [29](#), [31](#)
- (1991): “Entry and Competition in Concentrated Markets,” *Journal of Political Economy*, 99, 977–1009. [1](#), [18](#), [28](#), [29](#), [31](#)
- DAVIS, P. (2002): “Estimating multi-way error components models with unbalanced data structures,” *Journal of Econometrics*, 106, 67–95. [25](#), [26](#), [34](#)
- (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. [2](#), [30](#)
- (2006): “Spatial Competition in Retail Markets: Movie Theaters,” *RAND Journal of Economics*, 37, 964–982. [2](#), [24](#), [25](#), [26](#), [31](#), [34](#)
- DORASZELSKI, U. AND A. PAKES (2007): “A Framework for Applied Dynamic Analysis in IO,” in *Handbook of Industrial Organization*, ed. by M. Armstrong and R. H. Porter, Amsterdam: Elsevier Science, vol. 3, chap. 4. [2](#)
- 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. [20](#)
- ERICSON, R. AND A. PAKES (1995): “Markov-Perfect Industry Dynamics: A Framework for Empirical Work,” *Review of Economic Studies*, 62, 53–82. [2](#)

- FUDENBERG, D. AND J. TIROLE (1991): *Game Theory*, Cambridge, MA: MIT Press. [7](#)
- HOPENHAYN, H. A. (1992): “Entry, Exit, and Firm Dynamics in Long Run Equilibrium,” *Econometrica*, 60, 1127–1150. [1](#)
- HOTZ, V. J. AND R. A. MILLER (1993): “Conditional Choice Probabilities and the Estimation of Dynamic Models,” *Review of Economic Studies*, 60, 497–529. [2](#)
- HOTZ, V. J., R. A. MILLER, S. G. SANDERS, AND J. A. SMITH (1994): “A Simulation Estimator for Dynamic Models of Discrete Choice,” *Review of Economic Studies*, 61, 265–289. [15](#), [18](#)
- JAIMOVICH, N. (2007): “Firm dynamics and markup variations: Implications for sunspot equilibria and endogenous economic fluctuations,” *Journal of Economic Theory*, 137, 300–325. [35](#)
- MASKIN, E. AND J. TIROLE (1988): “A Theory of Dynamic Oligopoly, I: Overview and Quantity Competition with Large Fixed Costs,” *Econometrica*, 56, 549–569. [5](#)
- MELITZ, M. J. (2003): “The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity,” *Econometrica*, 71, 1695–1725. [1](#)
- 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. [31](#)
- RUST, J. (1987): “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55, 999–1033. [0](#), [1](#), [4](#)
- (1994): “Structural Estimation of Markov Decision Processes,” in *Handbook of Econometrics*, ed. by R. Engle and D. McFadden, Amsterdam: North-Holland, vol. 4, 3081–3143. [15](#), [19](#)
- SU, C.-L. AND K. L. JUDD (2012): “Constrained Optimization Approaches to Estimation of Structural Models,” *Econometrica*, 80, 2213–2230. [2](#), [23](#)
- SUTTON, J. (1991): *Sunk Costs and Market Structure*, MIT Press. [3](#)

TAUCHEN, G. (1986): “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions,” *Economics letters*, 20, 177–181. [20](#)

WEINTRAUB, G. Y., C. L. BENKARD, AND B. VAN ROY (2008): “Markov Perfect Industry Dynamics with Many Firms,” *Econometrica*, 76, 1375–1411. [2](#)

XU, Y. (2008): “A Structural Empirical Model of R&D, Firm Heterogeneity and Industry Evolution,” Mimeo, Department of Economics, New York University, New York, NY. [2](#)

Working Paper Series

A series of research studies on regional economic issues relating to the Seventh Federal Reserve District, and on financial and economic topics.

Comment on “Letting Different Views about Business Cycles Compete” <i>Jonas D.M. Fisher</i>	WP-10-01
Macroeconomic Implications of Agglomeration <i>Morris A. Davis, Jonas D.M. Fisher and Toni M. Whited</i>	WP-10-02
Accounting for non-annuitization <i>Svetlana Pashchenko</i>	WP-10-03
Robustness and Macroeconomic Policy <i>Gadi Barlevy</i>	WP-10-04
Benefits of Relationship Banking: Evidence from Consumer Credit Markets <i>Sumit Agarwal, Souphala Chomsisengphet, Chunlin Liu, and Nicholas S. Souleles</i>	WP-10-05
The Effect of Sales Tax Holidays on Household Consumption Patterns <i>Nathan Marwell and Leslie McGranahan</i>	WP-10-06
Gathering Insights on the Forest from the Trees: A New Metric for Financial Conditions <i>Scott Brave and R. Andrew Butters</i>	WP-10-07
Identification of Models of the Labor Market <i>Eric French and Christopher Taber</i>	WP-10-08
Public Pensions and Labor Supply Over the Life Cycle <i>Eric French and John Jones</i>	WP-10-09
Explaining Asset Pricing Puzzles Associated with the 1987 Market Crash <i>Luca Benzoni, Pierre Collin-Dufresne, and Robert S. Goldstein</i>	WP-10-10
Prenatal Sex Selection and Girls’ Well-Being: Evidence from India <i>Luoqia Hu and Analía Schlosser</i>	WP-10-11
Mortgage Choices and Housing Speculation <i>Gadi Barlevy and Jonas D.M. Fisher</i>	WP-10-12
Did Adhering to the Gold Standard Reduce the Cost of Capital? <i>Ron Alquist and Benjamin Chabot</i>	WP-10-13
Introduction to the <i>Macroeconomic Dynamics</i> : Special issues on money, credit, and liquidity <i>Ed Nosal, Christopher Waller, and Randall Wright</i>	WP-10-14
Summer Workshop on Money, Banking, Payments and Finance: An Overview <i>Ed Nosal and Randall Wright</i>	WP-10-15
Cognitive Abilities and Household Financial Decision Making <i>Sumit Agarwal and Bhashkar Mazumder</i>	WP-10-16

Working Paper Series *(continued)*

Complex Mortgages <i>Gene Amromin, Jennifer Huang, Clemens Sialm, and Edward Zhong</i>	WP-10-17
The Role of Housing in Labor Reallocation <i>Morris Davis, Jonas Fisher, and Marcelo Veracierto</i>	WP-10-18
Why Do Banks Reward their Customers to Use their Credit Cards? <i>Sumit Agarwal, Sujit Chakravorti, and Anna Lunn</i>	WP-10-19
The impact of the originate-to-distribute model on banks before and during the financial crisis <i>Richard J. Rosen</i>	WP-10-20
Simple Markov-Perfect Industry Dynamics <i>Jaap H. Abbring, Jeffrey R. Campbell, and Nan Yang</i>	WP-10-21
Commodity Money with Frequent Search <i>Ezra Oberfield and Nicholas Trachter</i>	WP-10-22
Corporate Average Fuel Economy Standards and the Market for New Vehicles <i>Thomas Klier and Joshua Linn</i>	WP-11-01
The Role of Securitization in Mortgage Renegotiation <i>Sumit Agarwal, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, and Douglas D. Evanoff</i>	WP-11-02
Market-Based Loss Mitigation Practices for Troubled Mortgages Following the Financial Crisis <i>Sumit Agarwal, Gene Amromin, Itzhak Ben-David, Souphala Chomsisengphet, and Douglas D. Evanoff</i>	WP-11-03
Federal Reserve Policies and Financial Market Conditions During the Crisis <i>Scott A. Brave and Hesna Genay</i>	WP-11-04
The Financial Labor Supply Accelerator <i>Jeffrey R. Campbell and Zvi Hercowitz</i>	WP-11-05
Survival and long-run dynamics with heterogeneous beliefs under recursive preferences <i>Jaroslav Borovička</i>	WP-11-06
A Leverage-based Model of Speculative Bubbles (Revised) <i>Gadi Barlevy</i>	WP-11-07
Estimation of Panel Data Regression Models with Two-Sided Censoring or Truncation <i>Sule Alan, Bo E. Honoré, LuoJia Hu, and Søren Leth-Petersen</i>	WP-11-08
Fertility Transitions Along the Extensive and Intensive Margins <i>Daniel Aaronson, Fabian Lange, and Bhashkar Mazumder</i>	WP-11-09
Black-White Differences in Intergenerational Economic Mobility in the US <i>Bhashkar Mazumder</i>	WP-11-10

Working Paper Series *(continued)*

Can Standard Preferences Explain the Prices of Out-of-the-Money S&P 500 Put Options? <i>Luca Benzoni, Pierre Collin-Dufresne, and Robert S. Goldstein</i>	WP-11-11
Business Networks, Production Chains, and Productivity: A Theory of Input-Output Architecture <i>Ezra Oberfield</i>	WP-11-12
Equilibrium Bank Runs Revisited <i>Ed Nosal</i>	WP-11-13
Are Covered Bonds a Substitute for Mortgage-Backed Securities? <i>Santiago Carbó-Valverde, Richard J. Rosen, and Francisco Rodríguez-Fernández</i>	WP-11-14
The Cost of Banking Panics in an Age before “Too Big to Fail” <i>Benjamin Chabot</i>	WP-11-15
Import Protection, Business Cycles, and Exchange Rates: Evidence from the Great Recession <i>Chad P. Bown and Meredith A. Crowley</i>	WP-11-16
Examining Macroeconomic Models through the Lens of Asset Pricing <i>Jaroslav Borovička and Lars Peter Hansen</i>	WP-12-01
The Chicago Fed DSGE Model <i>Scott A. Brave, Jeffrey R. Campbell, Jonas D.M. Fisher, and Alejandro Justiniano</i>	WP-12-02
Macroeconomic Effects of Federal Reserve Forward Guidance <i>Jeffrey R. Campbell, Charles L. Evans, Jonas D.M. Fisher, and Alejandro Justiniano</i>	WP-12-03
Modeling Credit Contagion via the Updating of Fragile Beliefs <i>Luca Benzoni, Pierre Collin-Dufresne, Robert S. Goldstein, and Jean Helwege</i>	WP-12-04
Signaling Effects of Monetary Policy <i>Leonardo Melosi</i>	WP-12-05
Empirical Research on Sovereign Debt and Default <i>Michael Tomz and Mark L. J. Wright</i>	WP-12-06
Credit Risk and Disaster Risk <i>François Gourio</i>	WP-12-07
From the Horse’s Mouth: How do Investor Expectations of Risk and Return Vary with Economic Conditions? <i>Gene Amromin and Steven A. Sharpe</i>	WP-12-08
Using Vehicle Taxes To Reduce Carbon Dioxide Emissions Rates of New Passenger Vehicles: Evidence from France, Germany, and Sweden <i>Thomas Klier and Joshua Linn</i>	WP-12-09
Spending Responses to State Sales Tax Holidays <i>Sumit Agarwal and Leslie McGranahan</i>	WP-12-10

Working Paper Series *(continued)*

Micro Data and Macro Technology <i>Ezra Oberfield and Devesh Raval</i>	WP-12-11
The Effect of Disability Insurance Receipt on Labor Supply: A Dynamic Analysis <i>Eric French and Jae Song</i>	WP-12-12
Medicaid Insurance in Old Age <i>Mariacristina De Nardi, Eric French, and John Bailey Jones</i>	WP-12-13
Fetal Origins and Parental Responses <i>Douglas Almond and Bhashkar Mazumder</i>	WP-12-14
Repos, Fire Sales, and Bankruptcy Policy <i>Gaetano Antinolfi, Francesca Carapella, Charles Kahn, Antoine Martin, David Mills, and Ed Nosal</i>	WP-12-15
Speculative Runs on Interest Rate Pegs The Frictionless Case <i>Marco Bassetto and Christopher Phelan</i>	WP-12-16
Institutions, the Cost of Capital, and Long-Run Economic Growth: Evidence from the 19th Century Capital Market <i>Ron Alquist and Ben Chabot</i>	WP-12-17
Emerging Economies, Trade Policy, and Macroeconomic Shocks <i>Chad P. Bown and Meredith A. Crowley</i>	WP-12-18
The Urban Density Premium across Establishments <i>R. Jason Faberman and Matthew Freedman</i>	WP-13-01
Why Do Borrowers Make Mortgage Refinancing Mistakes? <i>Sumit Agarwal, Richard J. Rosen, and Vincent Yao</i>	WP-13-02
Bank Panics, Government Guarantees, and the Long-Run Size of the Financial Sector: Evidence from Free-Banking America <i>Benjamin Chabot and Charles C. Moul</i>	WP-13-03
Fiscal Consequences of Paying Interest on Reserves <i>Marco Bassetto and Todd Messer</i>	WP-13-04
Properties of the Vacancy Statistic in the Discrete Circle Covering Problem <i>Gadi Barlevy and H. N. Nagaraja</i>	WP-13-05
Credit Crunches and Credit Allocation in a Model of Entrepreneurship <i>Marco Bassetto, Marco Cagetti, and Mariacristina De Nardi</i>	WP-13-06

Working Paper Series *(continued)*

Financial Incentives and Educational Investment: The Impact of Performance-Based Scholarships on Student Time Use <i>Lisa Barrow and Cecilia Elena Rouse</i>	WP-13-07
The Global Welfare Impact of China: Trade Integration and Technological Change <i>Julian di Giovanni, Andrei A. Levchenko, and Jing Zhang</i>	WP-13-08
Structural Change in an Open Economy <i>Timothy Uy, Kei-Mu Yi, and Jing Zhang</i>	WP-13-09
The Global Labor Market Impact of Emerging Giants: a Quantitative Assessment <i>Andrei A. Levchenko and Jing Zhang</i>	WP-13-10
Size-Dependent Regulations, Firm Size Distribution, and Reallocation <i>François Gourio and Nicolas Roys</i>	WP-13-11
Modeling the Evolution of Expectations and Uncertainty in General Equilibrium <i>Francesco Bianchi and Leonardo Melosi</i>	WP-13-12
Rushing into American Dream? House Prices, Timing of Homeownership, and Adjustment of Consumer Credit <i>Sumit Agarwal, Luojia Hu, and Xing Huang</i>	WP-13-13
The Earned Income Tax Credit and Food Consumption Patterns <i>Leslie McGranahan and Diane W. Schanzenbach</i>	WP-13-14
Agglomeration in the European automobile supplier industry <i>Thomas Klier and Dan McMillen</i>	WP-13-15
Human Capital and Long-Run Labor Income Risk <i>Luca Benzoni and Olena Chyruk</i>	WP-13-16
The Effects of the Saving and Banking Glut on the U.S. Economy <i>Alejandro Justiniano, Giorgio E. Primiceri, and Andrea Tambalotti</i>	WP-13-17
A Portfolio-Balance Approach to the Nominal Term Structure <i>Thomas B. King</i>	WP-13-18
Gross Migration, Housing and Urban Population Dynamics <i>Morris A. Davis, Jonas D.M. Fisher, and Marcelo Veracierto</i>	WP-13-19
Very Simple Markov-Perfect Industry Dynamics <i>Jaap H. Abbring, Jeffrey R. Campbell, Jan Tilly, and Nan Yang</i>	WP-13-20