Very Simple Markov-Perfect Industry Dynamics: Theory

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

March 2017

Abstract

This paper develops a simple model of firm entry, competition, and exit in oligopolistic markets. It features toughness of competition, sunk entry costs, and market-level demand and cost shocks, but assumes that firms' expected payoffs are identical when entry and survival decisions are made. We prove that this model has an essentially unique symmetric Markov-perfect equilibrium, and we provide an algorithm for its computation. Because this algorithm only requires finding the fixed points of a finite sequence of contraction mappings, it is guaranteed to converge quickly.

[*]CentER, Department of Econometrics & OR, Tilburg University, and CEPR. E-mail: jaap@abbring.org.

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

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

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

JEL Codes: L13, C25, C73.

Keywords: demand uncertainty, dynamic oligopoly, firm entry and exit, sunk costs, toughness of competition.

We thank Martin Pesendorfer for his insightful conference discussion. The research of Jaap Abbring is financially supported by the Netherlands Organisation for Scientific Research (NWO) through Vici grant 453-11-002. Nan Yang gratefully acknowledges financial support from National University of Singapore Start-up Grant R-313-000-100-133 and an HSS Faculty Research Fellowship. The views expressed herein are those of the authors, and they do not necessarily represent those of the Federal Reserve Bank of Chicago, the Federal Reserve System, or its Board of Governors.

1 Introduction

In this paper, we present a model of firm entry, competition, and exit in oligopolistic markets. It features toughness of competition, sunk entry costs, and market-level demand and cost shocks. We allow firms to use mixed strategies and close the model by focusing on symmetric Markov-perfect equilibria. The model's key simplifying assumption is that firms' expected payoffs are identical when entry and survival decisions are made. Using this and the equilibrium implications of mixed strategies for payoffs, we construct an algorithm for equilibrium computation that calculates the fixed points of a finite sequence of low-dimensional contraction mappings. Since it relies only on contraction mappings, the algorithm is *guaranteed* to calculate an equilibrium. We prove that adding a competitor cannot increase incumbents' equilibrium continuation values. This result in turn ensures that the symmetric equilibrium calculated by our algorithm is essentially unique. The algorithm converges sufficiently quickly to be embedded in a nested fixed point estimation procedure and used for large-scale computational experiments.

Our model can be viewed as a special case of Ericson and Pakes's (1995) Markov-perfect industry dynamics framework. Ericson and Pakes focused on symmetric equilibria in pure strategies, but Doraszelski and Satterthwaite (2010) showed that such equilibria might not exist in their original framework. To address this problem, Gowrisankaran (1999) added privately-observed firm-specific shocks to the costs of continuation; and Doraszelski and Satterthwaite provided sufficient conditions for such an augmented framework to have a symmetric equilibrium in pure strategies. Research following Ericson and Pakes (summarized by Doraszelski and Pakes, 2007) has generally adopted this augmented version of their framework.

We instead return to Ericson and Pakes's original complete-information approach. We show that the firm-specific shocks that guarantee existence of an equilibrium in pure strategies in the augmented framework obscure a useful consequence of firms employing mixed strategies. In equilibrium, firms earn the value of the outside option (zero) whenever they nontrivially randomize over exit and survival. Therefore, symmetric equilibrium payoffs to incumbents contemplating survival equal either zero or the value of all incumbents choosing certain continuation. This insight allows us to calculate continuation values from some nodes of the game tree without knowing everything about the game's subsequent play. Combining this insight with a demonstration that continuation values weakly decrease with the number of active firms yields the contraction mappings that we use both to calculate the equilibrium and to demonstrate its uniqueness. In contrast, there is no guarantee that the augmented Ericson and Pakes framework has a unique equilibrium; and computing even one of its equilibria can be onerous.

Earlier research similarly exploited the structure of specific games to enable their theoretical and computational analysis. Abbring and Campbell (2010) considered a dynamic oligopoly model like ours, but assumed that incumbent firms make continuation decisions sequentially in the order of their entry. Moreover, they restricted attention to Markov-perfect equilibria in which older firms always outlive their younger rivals, which they called "last-in first-out" dynamics. Our equilibrium characterization and computation rely neither on sequential timing assumptions nor on a restriction to last-in first-out dynamics.

Another strand of the literature applied backward induction to compute the equilibria of dynamic directional games (e.g. Cabral and Riordan, 1994; Judd, Schmedders, and Yeltekin, 2012). Iskhakov, Rust, and Schjerning (2016) systemized this familiar procedure into an algorithm for computing *all* these games' equilibria. In the games considered, the state space can be partially ordered using primitive restrictions on state transitions: State B comes after state A if B can be reached from A but not the other way around. Their algorithm iterates backwards through this partially ordered set of states. Transitions from states considered in a given iteration to states considered in later iterations are impossible, so the algorithm can calculate equilibrium outcomes and continuation values recursively. Our algorithm similarly iterates over an ordered partition of our game's state space. However, our game is *not* directional and in each iteration transitions to states not yet visited by our algorithm *are* possible. Instead of exploiting the directionality of state transitions hardwired into the primitives of Iskhakov, Rust, and Schjerning's framework, we rely on the fact that the expected symmetric equilibrium payoff in any survival subgame in which firms exit with positive probability must be zero. This allows us to order state D after state C if state D can be reached from state C but the opposite transition requires firms to choose exit with positive probability.

2 The Model

Consider a market in discrete time indexed by $t \in \mathbb{N} \equiv \{1, 2, \ldots\}$, in which firms make entry and exit decisions. In period t, firms that have entered in the past and not yet exited serve the market. Each firm has a name $f \in \mathcal{F} \equiv \mathcal{F}_0 \cup (\mathbb{N} \times \{1, 2, \ldots, \check{\jmath}\})$. Initial incumbents have distinct names in $\mathcal{F}_0$, while potential entrants' names are from $\mathbb{N} \times \{1, 2, \ldots, \check{\jmath}\}$. The first component of a potential entrant's name gives the period in which it has its only opportunity to enter the market, and the second component gives its position in that period's queue of $\check{\jmath} < \infty$ firms. 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 stages. Play in period t begins on the left with the entry stage. If $t = 1$, nature sets the number N_1 of firms serving the market in period 1 and the initial demand state Y_1. If $t > 1$, these are inherited from the previous period. We use $\mathcal{Y}$ to denote the support of Y_t. Although we consistently refer to Y_t as "demand," it can encompass any market characteristics that may affect, but are not affected by, firms' decisions. For instance, Y_t may be vector-valued and include cost shocks. It follows a first-order Markov process.

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

A1. $\exists \check{\pi} < \infty$ such that $\forall (n, y) \in \mathbb{N} \times \mathcal{Y}$, $-\infty < \mathbb{E}\left[\pi(n, Y') | Y = y\right] < \check{\pi}$;

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

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

Here and throughout; we denote the next period's value of a generic variable Z with Z', random variables with capital Roman letters, and their realizations with the corresponding small Roman letters. The first assumption ensures that expected discounted profits (values) in all entry or survival decision nodes are bounded from above. Because firms will, in equilibrium, limit losses by exiting, this will allow us to restrict our analysis of equilibrium values to the space of bounded functions. We will

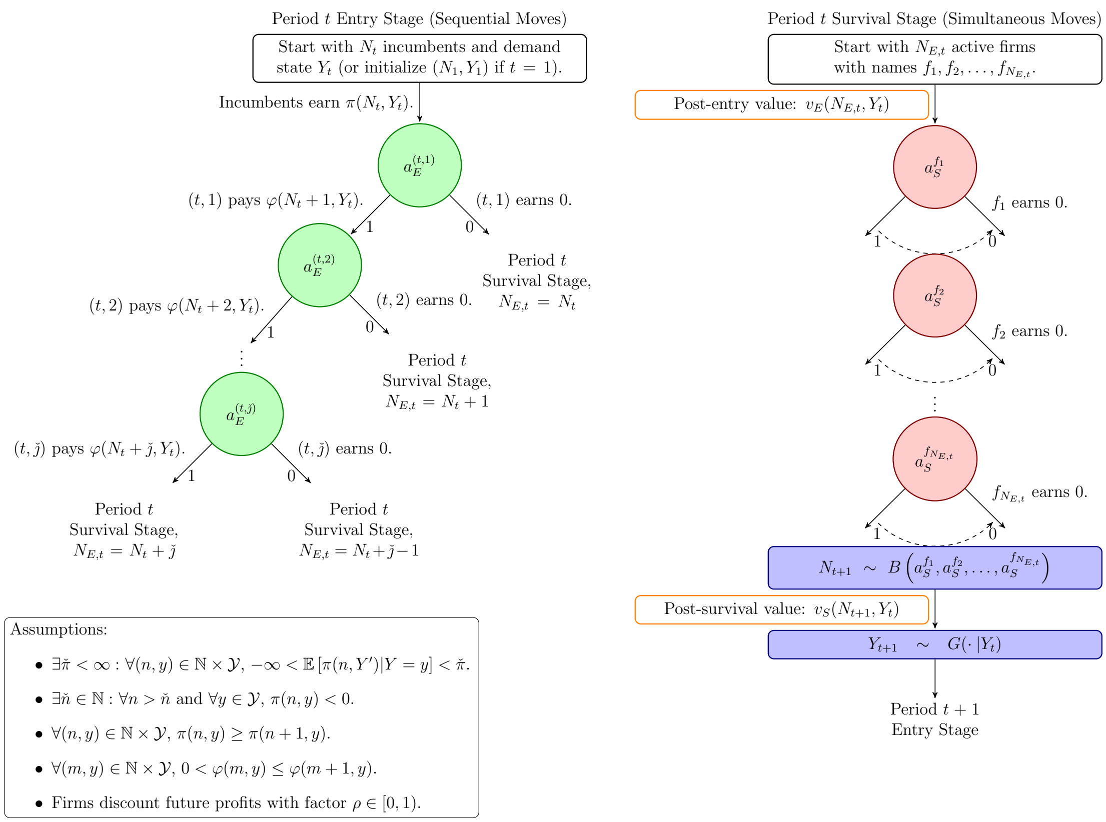

Figure 1: The Model's Recursive Extensive Form

use the second assumption to bound the number of firms that will participate in the market simultaneously. It is not restrictive in empirical applications to oligopolistic markets. The third assumption requires the addition of a competitor to reduce weakly each incumbent's profit. That is, what Sutton (1991) labelled the *toughness of competition* must dominate any complementarities between firms' activities.

After incumbents earn their profits, entry may occur. The period t entry cohort consists of firms with names in $\{t\} \times \{1, \ldots, \check{\jmath}\}$. These firms make their entry decisions sequentially in the order of their names' second components. We denote firm f's entry decision with $a_E^f \in \{0, 1\}$. A firm in the j'th position of the current period's entry queue that enters pays the sunk cost $\varphi(N_t + j, Y_t)$. This satisfies

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

If $\varphi(m, y) < \varphi(m + 1, y)$, then the $m + 1$'th firm faces an *economic barrier to entry* (McAfee, Mialon, and Williams, 2004). A firm choosing not to enter earns a payoff of zero and never has another entry opportunity. Such a refusal to enter also ends the entry stage, so firms remaining in this period's entry cohort that have not yet had an opportunity to enter *never* get to do so.

The total number of firms in the market after the entry stage equals $N_{E,t}$, which sums the incumbents with the actual entrants. Denote their names with $f_1, \ldots, f_{N_{E,t}}$. In the survival stage, these firms simultaneously choose probabilities of remaining active, $a_S^{f_1}, \ldots, a_S^{f_{N_{E,t}}} \in [0,1]$.[1] Subsequently, all survival outcomes are realized independently across firms according to the chosen Bernoulli distributions. Firms that exit earn a payoff of zero and never again participate in the market. The N_{t+1} surviving firms continue to the next period, $t+1$.[2] To end the period, nature draws a new demand state Y_{t+1} from the Markov transition distribution $G(\cdot\, |Y_t)$.

The timing of our game is similar to that in Ericson and Pakes (1995). Like them, we allow for sequential entry.[3] Moreover, like Ericson and Pakes, and unlike Abbring and Campbell (2010), we assume simultaneous survival decisions. Because we allow for mixed survival rules, this may lead to excessive exits. Since entry precedes exit, potential entrants cannot take immediate advantage of such "exit mistakes" and thereby outmaneuver incumbents. This is not so relevant to Ericson and Pakes, who restrict strategies to be pure (at the expense of losing equilibrium existence; see Doraszelski and Satterthwaite, 2010). To establish the robustness of our results to the game's timing assumptions, we considered a variant of our model in which at most one firm enters each period and entry and survival decisions are all taken simultaneously. In this paper's online supplement, we demonstrate that this alternative game has a unique equilibrium in which potential entrants never displace incumbents; and we provide an algorithm for its rapid calculation.

---

[1] We do not explicitly model the firms' randomization devices. A more complete development would assign each active firm an independent uniformly-distributed random variable and have each firm choose a set of realizations that direct it to survive. In this extension, a survival probability equal to one could indicate either that the firm chooses to exit *never* or that it chooses to exit whenever its random variable falls into a particular non-empty set of measure zero. Throughout this paper, we will assume the former and interpret $a_S = 0$ and $a_S = 1$ as dictating *certain* exit and survival.

[2] The assumption that entrants immediately contemplate exit might seem strange, but exit immediately following entry never occurs in equilibrium. Furthermore, this timing assumption removes an unrealistic possibility. If entrants did not make these continuation decisions, then they could effectively commit to continuation. This would allow an entrant to displace an incumbent only by virtue of this commitment power.

[3] See their page 60: "We assume that, in each period, ex ante identical firms decide to enter sequentially until the expected value of entry falls sufficiently to render further entry unprofitable."

3 Equilibrium

We assume that firms play a symmetric Markov-perfect equilibrium, a subgame-perfect equilibrium in which all firms use the same Markov strategy.

3.1 Symmetric Markov-Perfect Equilibrium

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

In a symmetric Markov-perfect equilibrium, a firm's expected continuation value at a particular node of the game can be written as a function of that node's payoff-relevant state variables. Two of these value functions are particularly useful for the model's equilibrium analysis: the *post-entry* value function, v_E, and the *post-survival* value function, v_S. The post-entry value $v_E(N_E, Y)$ equals the expected discounted profits of a firm in a market with demand state Y and N_E firms just after all entry decisions are made. The post-survival value $v_S(N', Y)$ equals the expected discounted profits from being active in the same market with N' firms just after the survival outcomes are realized. Figure 1 shows the points in the survival stage when these value functions apply.

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

$$v_S(n', y) = \rho \mathbb{E}_{a_E} \left[\pi(n', Y') + v_E(N'_E, Y') \middle| N' = n', Y = y \right]. \tag{1}$$

Here, $\mathbb{E}_{a_E}$ is an expectation over the next period's demand state Y' and post-entry number of firms N'_E. This expectation operator's subscript indicates its dependence on a_E. In particular, given $N' = n'$, N'_E is a deterministic function of $a_E(\cdot, Y')$. Note that Assumption A1 implies that v_E is bounded from above. This ensures that the expectation in the right-hand side of (1) exists.[4]

Because the payoff from leaving the market is zero, a firm's post-entry value in a state (n_E, y) equals the probability that it survives, $a_S(n_E, y)$, times the expected payoff from surviving:[5]

$$v_E(n_E, y) = a_S(n_E, y)\, \mathbb{E}_{a_S}\left[v_S(N', y)\middle|N_E = n_E, Y = y\right], \tag{2}$$

The expectation $\mathbb{E}_{a_S}$ over N' takes survival of the firm of interest as given. That is, it takes N' to equal one plus the outcome of $n_E - 1$ independent Bernoulli (survival) trials with success probability $a_S(n_E, y)$. Its subscript makes its dependence on a_S explicit. It conditions on the current value of Y because this influences the survival probability's value.

If a strategy (a_E, a_S) forms a symmetric Markov-perfect equilibrium with payoffs (v_E, v_S), then no firm can gain from a one-shot deviation from its prescriptions:[6]

$$a_E(m, y) \in \arg\max_{a \in \{0,1\}} a\big(\mathbb{E}_{a_E}\left[v_E(N_E, y)|M = m, Y = y\right] - \varphi(m, y)\big) \text{ and} \tag{3}$$

$$a_S(n_E, y) \in \arg\max_{a \in [0,1]} a\mathbb{E}_{a_S}\left[v_S(N', y)|N_E = n_E, Y = y\right]. \tag{4}$$

Together with Assumption A2, (3) and (4) bound the long-run number of firms in equilibrium.

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

The Appendix provides this Lemma's proof. Intuitively, firms cannot survive for sure with $n > \check{n}$ firms because this would give them negative payoffs. To see this, note that if all firms continue for sure, each would earn a negative profit one or more times (due to our assumption $\pi(n, y) < 0$ for all $n > \check{n}$). In the first future period in which firms leave with positive probability, (4) requires the post-entry value to equal zero. Therefore, continuing for sure with $n > \check{n}$ yields a negative

[4] At this point in our model's development, we cannot exclude the possibility that it equals $-\infty$.

[5] We define the right-hand side of (2) to equal zero if the firm collects a payoff of zero by exiting for sure ($a_S(n_E, y) = 0$), even if $\mathbb{E}_{a_S}\left[v_S(N', y)\middle|N_E = n_E, Y = y\right] = -\infty$.

[6] We define the maximands in the right-hand sides of (3) and (4) to equal zero if $a = 0$.

expected payoff. Any firm could avoid this by exiting instead, so $a_S(n, y) < 1$ and $v_E(n, y) = 0$. Because no firm would be willing to pay a positive sunk cost to enter a survival subgame with zero expected payoff, $a_E(n, y) = 0$.

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

Lemma 1 implies that setting the number of potential entrants per period ($\check{\jmath}$) to exceed $\check{n}$ guarantees that at least one potential entrant per period refuses an entry opportunity. In this sense, the model becomes one of free entry, as in Ericson and Pakes (1995). This is a standard and convenient assumption in applications without an identifiable and finite set of potentially active firms. The remaining development of our model imposes this free entry assumption ($\check{\jmath} > \check{n}$).

In the online supplement, we show that (3) and (4) are not only necessary, but also sufficient for (a_E, a_S) to be an equilibrium strategy. Proofs of this "one-shot deviation principle" (e.g. Fudenberg and Tirole, 1991, Theorem 4.2) typically make assumptions on payoffs that bound from both above and below the value gains from deviating in the distant future from *any* strategy, whether that strategy satisfies (3) and (4) or not. Our model does not satisfy these assumptions, because it imposes no lower bound on profits.[7] Conditions (3) and (4) do, however, imply a lower bound (corresponding to the outside option of zero) on the values in the survival and entry nodes in equilibrium (including v_E). Because the expected discounted profits in these decision nodes are bounded from above under any strategy profile, the gains from deviating from a strategy (a_E, a_S) that satisfies (3) and (4) are bounded from above. Using this, we adapt existing proofs of the one-shot deviation principle.

Before proceeding to characterize the equilibrium set, we wish to note and dispense with an uninteresting source of equilibrium multiplicity. If a potential

[7]The absence of a lower bound on profits is important when we bring the model to the data as we do in Abbring et al. (2017), where we provide a full econometric development of the model presented here. There, y is vector-valued and includes two elements, a demand state that is observed by the econometrician and a cost shock that is unobserved by the econometrician. These cost shocks serve as the model's econometric error. Permitting profits to be unbounded from below (and therefore permitting cost shocks to become arbitrarily large) is critical for ensuring that the model is statistically nondegenerate.

entrant is indifferent between entering and staying out, we may be able to construct one equilibrium from another by varying only that choice. Similarly, an incumbent monopolist can be indifferent between continuation and exit, and we can possibly construct one equilibrium from another by changing that choice alone. To avoid such uninteresting caveats to our results, we follow Abbring and Campbell (2010) by focusing on equilibria that *default to inactivity.* In such an equilibrium, a potential entrant that is indifferent between entering or not stays out,

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

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

$$v_S(1, y) = \cdots = v_S(n_E, y) = 0 \Rightarrow a_S(n_E, y) = 0.$$

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. When Y follows a continuous distribution, an exact indifference between activity and inactivity occurs with probability zero. For this reason, the restriction to equilibria that default to inactivity is very weak.

3.2 Existence, Uniqueness, and Computation

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

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

The Appendix contains Lemma 2's proof. It says that no endogenous complementarity between firms arises in equilibrium. To appreciate its implications, consider a one-shot simultaneous-moves survival game played by n_E active firms. In it, each of the n' survivors earns $v_S(n', y)$, where v_S is the post-survival value in a symmetric Markov-perfect equilibrium of our dynamic game, and each exiting firm earns zero. A survival probability $a_S(n_E, y)$ is a symmetric Nash equilibrium strategy of this one-shot game if and only if it satisfies (4). Thus, a survival rule a_S from a symmetric

Markov-perfect equilibrium gives a symmetric Nash equilibrium survival probability $a_S(n_E, y)$ for each one-shot game defined by $n_E \in \{1, \ldots, \check{n}\}$ and $y \in \mathcal{Y}$. The converse also holds good: A collection of Nash equilibrium survival probabilities from survival games can be assembled into a survival rule.

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

- $v_S(1, y) \leq 0$. Lemma 2 implies that $v_S(n', y) \leq 0$ for all $n' > 1$. Therefore, exiting for sure is a weakly dominant strategy. Since $v_S(n', y) \neq 0$ for at least one $n' \in \{1, \ldots, n_E\}$, we also know that $v_S(n_E, y) < 0$. This makes exiting for sure the unique best response to any positive symmetric continuation probability, so there is only one symmetric equilibrium. In it, $a_S(n_E, y) = 0$.

- $v_S(n_E, y) \geq 0$. Lemma 2 implies that $v_S(n', y) \geq 0$ for $n' < n_E$. Therefore, continuing for sure is a weakly dominant strategy. Since $v_S(n', y) \neq 0$ for at least one $n' \in \{1, \ldots, n_E\}$, we also know that $v_S(1, y) > 0$. This makes continuing for sure the unique best response to any continuation probability less than one, so there is only one symmetric equilibrium. In it, $a_S(n_E, y) = 1$.

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

$$\sum_{n'=1}^{n_E} \binom{n_E - 1}{n' - 1} a^{n'-1} (1 - a)^{n_E - n'} v_S(n', y) = 0,$$

  where the left-hand side gives the expected value from survival when all other $n_E - 1$ firms survive with probability a and the right-hand side gives the value from exit. This establishes existence of a mixed strategy equilibrium. Lemma 2 and this case's preconditions together guarantee that the left-hand side strictly decreases with a. Therefore, there is only one symmetric mixed strategy equilibrium.

For future reference, we state this equilibrium uniqueness result with

Corrollary 1 *Fix* $n_E \in \{1, \ldots, \check{n}\}$ *and* $y \in \mathcal{Y}$*, let* v_S *be the post-survival value function associated with a symmetric Markov-perfect equilibrium, and suppose that* $v_S(n', y) \neq 0$ *for at least one* $n' \in \{1, \ldots, n_E\}$*. In the one-shot survival game in which* n_E *firms simultaneously choose between survival and exit (as in the survival stage of Figure 1), each of the* n' *survivors earns* $v_S(n', y)$*, and each exiting firm earns zero; there is a unique symmetric Nash equilibrium, possibly in mixed strategies.*

When the individual payoff from joint continuation is positive, this unique Nash-equilibrium strategy from Corollary 1 guarantees that firms survive for sure and receive this payoff. In all other states, each firm is either indifferent between surviving and exiting or prefers to exit for sure; and following the strategy gives each of them an expected payoff of zero. The post-entry payoff is always zero in the trivial case with $v_S(1, y) = \cdots = v_S(n_E, y) = 0$ excluded by Corollary 1. Thus,

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

$$v_E(n_E, y) = \max\{0, v_S(n_E, y)\}.$$

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

We proceed to demonstrate equilibrium existence constructively using Algorithm 1. Equilibrium uniqueness follows from this as a byproduct. Denote the candidate equilibrium values that the algorithm calculates with ν_E and ν_S, and the corresponding candidate equilibrium strategy with (α_E, α_S).

First, consider states with $\check{n}$ firms. By Lemma 1, there will be no entry in a period starting with $\check{n}$ firms. With (1), this implies that any possible candidate equilibrium post-survival value must satisfy

$$\nu_S(\check{n}, y) = \rho \mathbb{E}\big[\pi(\check{n}, Y') + \nu_E(\check{n}, Y')\big| Y = y\big].$$

With Corollary 2, this constrains the candidate post-entry value to satisfy

$$\nu_E(\check{n}, y) = \max\left\{0, \rho \mathbb{E}\big[\pi(\check{n}, Y') + \nu_E(\check{n}, Y')\big| Y = y\big]\right\}. \tag{5}$$

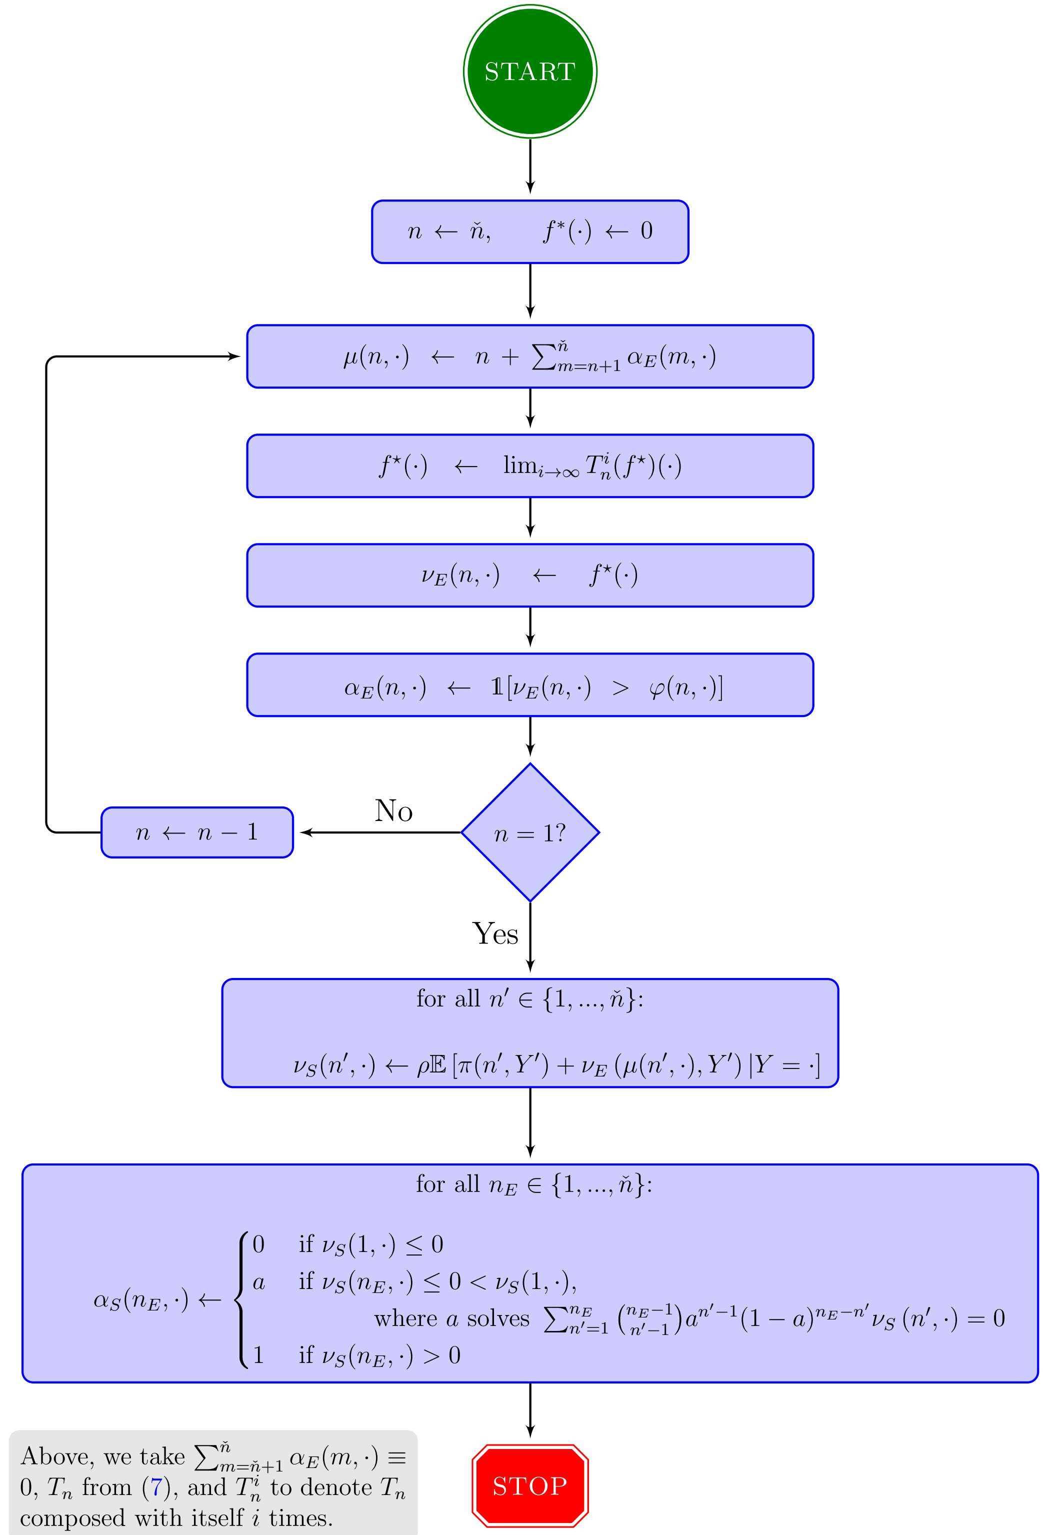

Algorithm 1: Equilibrium Calculation

The right-hand side of (5) defines a contraction mapping on the space of bounded functions on $\mathcal{Y}$, with a unique fixed point $\nu_E(\check{n}, \cdot)$. This $\nu_E(\check{n}, \cdot)$ is the only possible equilibrium post-entry value in a state with $\check{n}$ firms. Moreover, any entry rule that is (i) consistent with it, (ii) one-shot deviation proof as in (3), and (iii) defaults to inactivity must dictate entry into a market with $\check{n}-1$ incumbents if and only if the payoff from doing so is positive. Thus, the algorithm sets

$$\alpha_E\left(\check{n}, y\right) = \mathbb{1}\left[\nu_E(\check{n}, y) > \varphi(\check{n}, y)\right].$$

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

With $\nu_E(\check{n}, \cdot)$ and $\alpha_E(\check{n}, \cdot)$ calculated, the algorithm proceeds with the recursive construction of $\nu_E(n, \cdot)$ and $\alpha_E(n, \cdot)$ for n decreasing from $\check{n}-1$ to 1. For a given n, the algorithm has already calculated $\nu_E(n^\star, \cdot)$ and $\alpha_E(n^\star, \cdot)$ for $n^\star = n+1, n+2, ..., \check{n}$. Suppose that $\nu_E(n^\star, \cdot)$ and $\alpha_E(n^\star)$ weakly decrease with $n^\star$ (which, by Lemma 2 and Corollary 2, they will if ν_E is indeed an equilibrium post-entry value). Then,

$$\mu(n, y) \equiv n + \sum_{m=n+1}^{\check{n}} \alpha_E(m, y)$$

equals the number of firms that will be active in a period that starts with n firms after all that period's potential entrants have followed the candidate entry rule. Together, (1) and Corollary 2 require the candidate post-entry values to satisfy

$$\nu_E(n, y) = \max\left\{0, \rho\mathbb{E}\left[\pi(n, Y') + \nu_E(\mu(n, Y'), Y')\middle|Y = y\right]\right\}. \tag{6}$$

Given $\nu_E(n^\star, \cdot)$ for $n^\star = n+1, \ldots, \check{n}$, the right-hand side of (6) defines a contraction

$$\begin{aligned} T_n(f)(y) = \max\{0, \rho\mathbb{E}[\pi(n, Y') &+ \mathbb{1}[\mu(n, Y') = n]f(Y') \\ &+ \mathbb{1}[\mu(n, Y') > n]\nu_E(\mu(n, Y'), Y') \,|\, Y = y]\} \end{aligned} \tag{7}$$

with a unique fixed point $\nu_E(n, \cdot)$. This is the only possible post-entry value. Finally, a firm in state (n, y) enters if and only if $\nu_E(\mu(n, y), y) > \varphi(n, y)$. Again supposing that $\nu_E(n^\star, y)$ and $\alpha_E(n^\star, y)$ weakly decrease with n^*, and using that $\varphi(n^\star, y)$ weakly increases with n^*, this entry rule can be simplified to

$$\alpha_E(n, y) = \mathbb{1}\left[\nu_E(n, y) > \varphi(n, y)\right].$$

Once the algorithm's recursive part is complete, it has constructed a candidate post-entry value and entry rule. With (1), these imply a unique candidate post-survival value ν_S. After computing ν_S, the algorithm ends by setting the candidate survival rule α_S to a value consistent with ν_S and the analysis leading up to Corollary

1. Specifically, it sets $\alpha_S(n_E, y) = 0$ for all (n_E, y) such that $\nu_S(n_E, y) = \cdots = \nu_S(1, y) = 0$ (the algorithm subsumes this in the case that $\nu_S(1, y) \leq 0$) and finds an equilibrium to Corollary 1's one-shot survival game for all other (n_E, y). *If* the candidate is actually an equilibrium, then Corollary 1 guarantees that this candidate survival rule exists and is unique. This is indeed so.

Theorem 1 (Equilibrium existence and uniqueness) *There exists a unique symmetric Markov-perfect equilibrium. The equilibrium strategy and corresponding equilibrium payoffs are those computed by Algorithm 1.*

4 Conclusion

This paper's theoretical and computational results enable our model's empirical application. Since its key simplifying assumption imposes homogeneity of expected profits when firms make their entry and continuation choices, it is best suited for investigations that can be usefully undertaken while abstracting from persistent heterogeneity among competing firms. Examples of such studies include Bresnahan and Reiss's (1994) and Dunne, Klimek, Roberts, and Xu's (2013) estimations of oligopolists' sunk costs with panel data on firm counts and demand from cross sections of markets. In Abbring, Campbell, Tilly, and Yang (2017), we propose a simple procedure for empirically determining whether or not our model can be usefully applied to such data from a given industry. This decomposes the industry's Herfindahl-Hirschman Index (*HHI*) into its value with equally sized firms and a residual that we label the *contribution of heterogeneity.* Our procedure tests whether this heterogeneity measure contributes to forecasts of the number of active firms. If not, then our model can accommodate observed heterogeneity with *transitory* firm-specific disturbances. We applied this procedure to data from Motion Picture Theaters in 573 Micropolitan Statistical Areas in the United States. We found that heterogeneity's contribution to the *HHI* makes economically trivial contributions to Poisson regressions' forecasts of the number of firms serving that industry.

Our companion paper also demonstrates the practicality of applying our model to such data by estimating Motion Picture Theaters' sunk costs and the toughness of competition between them. The model's maximum likelihood estimation requires calculating a separate equilibrium for each market in the data for each trial value of its parameters, but this required only about thirty minutes using two Intel Xeon

E5-2699 v3 CPUs (released by Intel in 2014) on a single machine with C++ code. We were also able to conduct many policy experiments, which calculated the effects of large demand shocks and counterfactual competition policies. This experience leads us to conclude that structural investigations of oligopoly dynamics based on this paper's model can be done with few computational resources.

References

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

Abbring, J. H., J. R. Campbell, J. Tilly, and N. Yang (2017): "Very Simple Markov-Perfect Industry Dynamics: Empirics," Discussion Paper 2017-021, CentER, Tilburg.

Bresnahan, T. F. and P. C. Reiss (1994): "Measuring the Importance of Sunk Costs," *Annales d'Économie et de Statistique*, 34, 181–217.

Cabral, L. M. B. and M. H. Riordan (1994): "The Learning Curve, Market Dominance, and Predatory Pricing," *Econometrica*, 62, 1115–1140.

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.

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

Dunne, T., S. Klimek, M. J. Roberts, and D. Y. Xu (2013): "Entry, Exit, and the Determinants of Market Structure," *RAND Journal of Economics*, 44, 462–487.

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

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

Gowrisankaran, G. (1999): "A Dynamic Model of Endogenous Horizontal Mergers," *RAND Journal of Economics*, 30, 56–83.

Iskhakov, F., J. Rust, and B. Schjerning (2016): "Recursive Lexicographical Search: Finding All Markov Perfect Equilibria of Finite State Directional Dynamic Games," *Review of Economic Studies*, 83, 658–703.

Judd, K. L., K. Schmedders, and Ş. Yeltekin (2012): "Optimal Rules for Patent Races," *International Economic Review*, 53, 23–52.

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

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

Appendix: Proofs

Proof of Lemma 1. First, we will prove that $a_S(n, y) < 1$ for all $y \in \mathcal{Y}$ and $n > \check{n}$. Consider a period $t^\star$ survival subgame with $N_{E,t^\star} = n > \check{n}$ firms and demand state $Y_{t^\star} = y$. Define the random time τ as the first period weakly after $t^\star$ in which firms choose exit with positive probability, with $\tau \equiv \infty$ if they never do:

$$\tau \equiv \min\left(\{t \geq t^\star : a_S(N_{E,t}, Y_t) < 1\} \cup \{\infty\}\right).$$

Suppose that $a_S(n, y) = 1$, so $\tau > t^\star$. By definition, exit occurs only in or after period τ, so we know that $N_t = N_{E,t-1} \geq n$ for $t \in \{t^\star + 1, \ldots, \tau\}$. (Recall from Footnote 1 that we take $a_S(\cdot) = 1$ to dictate sure survival, not merely almost-sure survival.) Since $n > \check{n}$, this together with Assumption A2 implies that $\pi(N_t, Y_t) < 0$ for $t \in \{t^\star + 1, \ldots, \tau\}$. If $\tau = \infty$, then the incumbent firms receive an infinite sequence of strictly negative payoffs. If instead $\tau < \infty$, then the incumbent firms receive a finite sequence of strictly negative payoffs followed by the post-entry value from playing the period τ survival subgame $v_E(N_{E,\tau}, Y_\tau)$, which equals zero by (2), (4), and the definition of τ. Therefore, the period $t^\star$ post-survival value satisfies $v_S(n, y) < 0$. Since a period $t^\star$ incumbent firm can raise its payoff to zero by choosing certain exit, the supposition that $a_S(n, y) = 1$ must be incorrect.

Next, we will prove that $a_E(n, y) = 0$ for all $y \in \mathcal{Y}$ and $n > \check{n}$. Consider the decision of the first potential entrant, firm $(t^\star, 1)$, in a period $t^\star$ entry subgame that starts with $N_{t^\star} = n - 1 > \check{n} - 1$ incumbents and demand state $Y_{t^\star} = y$. Note that this firm pays $\varphi(n, y) > 0$ upon entry. In return, it earns a post-entry value of zero (as proven above). Therefore, it maximizes its payoff by staying out of the market and earning zero: $a_E(n, y) = 0$. ■

Proof of Lemma 2. It suffices to prove that $v_S(n', y) \geq v_S(n'+1, y)$ for all $n' \geq 1$ and $y \in \mathcal{Y}$. To this end, consider a subgame beginning immediately after the period $t^\star$'s simultaneous continuation and entry choices with $N_{t^\star+1} = n'$ and $Y_{t^\star} = y$. We call this the *original* subgame. Now consider a second period $t^\star$ subgame starting at the same point but with one additional firm. We refer to this as the *perturbed* subgame and use N_t^+ and $N_{E,t}^+$ to denote the initial and post-entry numbers of firms in this perturbed subgame in period t. Finally, define the random time τ^+ as the first period weakly after $t^\star + 1$, in which the firms in the perturbed subgame choose exit with positive probability, with $\tau^+ \equiv \infty$ if they never do:

$$\tau^+ \equiv \min\left(\{t \geq t^\star + 1 : a_S(N_{E,t}^+, Y_t) < 1\} \cup \{\infty\}\right).$$

There is no exit before period τ^+ in the perturbed subgame. Furthermore, we know that the period τ^+ post-entry value in that subgame equals zero. Therefore, we can write

$$v_S(n'+1, y) = \lim_{T\to\infty} \mathbb{E}\left[\sum_{t=t^\star+1}^{T} \rho^{t-t^\star} \mathbb{1}\left[t \leq \tau^+\right] \pi(N_t^+, Y_t) \middle| Y_{t^\star} = y\right].$$

Since τ^+ is a consequence of equilibrium choices, we know that $v_S(n'+1, y) > -\infty$.

Now consider an incumbent firm in the original subgame which (possibly) deviates after the period $t^\star$ survival stage by choosing to survive for sure as long as $t < \tau^+$ and to exit for sure if $t = \tau^+$. Let $\bar{N}_t$ denote the number firms serving the market during period t in the original subgame *with this deviation*. Since the original strategy was part of a subgame-perfect equilibrium, $v_S(n', y)$ exceeds the expected payoff from following this deviating strategy. That is

$$v_S(n', y) \geq \lim_{T\to\infty} \mathbb{E}\left[\sum_{t=t^\star+1}^{T} \rho^{t-t^\star} \mathbb{1}\left[t \leq \tau^+\right] \pi(\bar{N}_t, Y_t) \middle| Y_{t^\star} = y\right].$$

To show that the limit on the right-hand side is well defined, note that $\bar{N}_t \leq N_t^+$ for all $t \leq \tau^+$. Otherwise, the two subgames would have potential entrants in the same states making different entry choices. This would violate either the presumption that the equilibrium strategy is Markov or that it defaults to inactivity. This and Assumption A3 imply that $\pi(\bar{N}_t, Y_t) \geq \pi(N_t^+, Y_t)$ for all $t = t^\star + 1, \ldots, \tau^+$. Combining this with $v_S(n'+1, y) > -\infty$ gives the desired result.

Because the difference of two convergent sequences' limits equals the limit of the sequences' difference, we can write

$$v_S(n',y) - v_S(n'+1,y) \geq \lim_{T\to\infty} \mathbb{E}\left[\sum_{t=t^\star+1}^{T} \rho^{t-t^\star} \mathbb{1}\left[t \leq \tau^+\right] \left(\pi(\bar{N}_t, Y_t) - \pi(N_t^+, Y_t)\right) \middle| Y_{t^\star} = y \right].$$

Each term in the partial sum on the right-hand side is non-negative, so we conclude that $v_S(n',y) - v_S(n'+1,y) \geq 0$. ■

Proof of Theorem 1. The proof is divided into three parts. First, we show that the candidate post-survival value from Algorithm 1 satisfies Lemma 2's monotonicity requirements. Second, we use this to demonstrate that the candidate strategy indeed forms an equilibrium. Third, we demonstrate equilibrium uniqueness.

Fix $n \in \{1, 2, \ldots, \check{n} - 1\}$ and suppose that we know that $\nu_E(n+1,\cdot) \geq \cdots \geq \nu_E(\check{n},\cdot)$. Evaluating T_n at $f^\star(\cdot) = \nu_E(n+1,\cdot)$ gives

$$\begin{aligned}
T_n(f^\star)(\cdot) =& \max\{0, \rho\mathbb{E}[\pi(n,Y') + f^\star(Y') \\
&+ \mathbb{1}[\mu(n,Y') > n](\nu_E(\mu(n,Y'),Y') - f^\star(Y'))|Y = \cdot]\} \\
\geq& \max\{0, \rho\mathbb{E}[\pi(n+1,Y') + f^\star(Y') \qquad (8)\\
&+ \mathbb{1}[\mu(n,Y') > n+1](\nu_E(\mu(n,Y'),Y') - f^\star(Y'))|Y = \cdot]\} \\
=& \max\{0, \rho\mathbb{E}[\pi(n+1,Y') + f^\star(Y') \qquad (9)\\
&+ \mathbb{1}[\mu(n+1,Y') > n+1](\nu_E(\mu(n+1,Y'),Y') - f^\star(Y'))|Y = \cdot]\} \\
=& \nu_E(n+1,\cdot).
\end{aligned}$$

The inequality in (8) follows from Assumption A3 and the assumed $f^\star(Y') = \nu_E(n+1,Y')$. Since $\nu_E(n^\star,Y')$ weakly decreases with $n^\star$ for $n^\star > n$, so does $\alpha_E(n^\star,Y')$. Therefore, $\mu(n,Y') = \mu(n+1,Y')$ whenever $\mu(n,Y') > n+1$. This gives us (9). The final equality follows again from $f^\star(Y') = \nu_E(n+1,Y')$. The operator T_n is a monotone contraction mapping, so $T_n(f^\star)(\cdot) \geq \nu_E(n+1,\cdot)$ implies that its fixed point, $\nu_E(n,\cdot)$, weakly exceeds $\nu_E(n+1,\cdot)$. Recursively applying this argument for n decreasing from $\check{n}-1$ to 1 proves that $\nu_E(1,\cdot) \geq \nu_E(2,\cdot) \cdots \geq \nu_E(\check{n},\cdot)$. With Assumption A3 and the now established fact that $\mu(n',\cdot)$ weakly increases with n', this monotonicity implies that

$$\nu_S(n',\cdot) = \rho\mathbb{E}\left[\pi(n',Y') + \nu_E\left(\mu(n',\cdot),Y'\right)|Y = \cdot\right]$$

weakly decreases with n'. This is the desired monotonicity result.

For the second part, we first verify that α_E, ν_S, and ν_E satisfy (3) and (1). Since $\nu_E(n_E, y)$ weakly decreases with n_E, $\alpha_E(m, y)$ weakly decreases with m, so that

$$\mathbb{E}_{\alpha_E}\left[\nu_E(N_E, y)|M = m, Y = y\right] = \nu_E(\mu(m, y), y). \tag{10}$$

Thus, to verify (3), it suffices to show that $\nu_E(\mu(m, y), y) > \varphi(m, y)$ if and only if $\nu_E(m, y) > \varphi(m, y)$. If $\nu_E(\mu(m, y), y) > \varphi(m, y)$ then, because $\mu(m, y) \geq m$ and $\nu_E(n_E, y)$ weakly decreases with n_E, $\nu_E(m, y) \geq \nu_E(\mu(m, y), y) > \varphi(m, y)$. Conversely, if $\nu_E(m, y) > \varphi(m, y)$, then $\nu_E(\mu(m, y), y) > \varphi(\mu(m, y), y)$ so that, by Assumption A4, $\nu_E(\mu(m, y), y) > \varphi(m, y)$. This establishes that α_E and ν_E satisfy (3). Using (10), it is easy to verify that α_E, ν_S, and ν_E satisfy (1).

Next, consider (4) and (2). For states (n_E, y) such that $\nu_S(1, y) = \cdots = \nu_S(n_E, y) = 0$, (4) imposes only the trivial requirement that $\alpha_S(n_E, y) \in [0, 1]$. Algorithm 1's selection of $\alpha_S(n_E, y) = 0$ (subsumed in the case $\nu_S(1, y) \leq 0$) satisfies this. For all other states, Algorithm 1 sets $\alpha_S(n_E, y)$ to the symmetric Nash equilibrium of Corollary 1's n_E-player one-shot survival game with payoffs $\nu_S(n', y)$ from survival with $n' = 1, \ldots, n_E$ firms, which satisfies (4). (If $\nu_S(n_E, y) = 0 < \nu_S(1, y)$, it sets $\alpha_S(n_E, y)$ to the unique mixing probability that makes firms indifferent, which indeed equals one as in Corollary 1's equilibrium.) Equation (2) requires $\nu_E(n, y)$ to equal the expected payoff to this game, $\max\{0, \nu_S(n, y)\}$, which is true by construction. We conclude that (α_E, α_S) indeed forms an equilibrium.

We end by demonstrating equilibrium uniqueness. First, Section 3.2's argument implies that any $v_E(\check{n}, \cdot)$ equals the unique fixed point $\nu_E(\check{n}, \cdot)$ of $T_{\check{n}}$. With (3), this gives a unique $a_E(n, \cdot)$ that defaults to inactivity, $\alpha(n, \cdot)$. Next, repeat the following argument for n decreasing from $\check{n}-1$ to 1. For given n, suppose that we have uniquely determined $v_E(n^\star, \cdot) = \nu_E(n^\star, \cdot)$ and $a_E(n^\star, \cdot) = \alpha_E(n^\star, \cdot)$ for $n^\star = n + 1, \ldots, \check{n}$. Then, Section 3.2's argument (which uses (10)) implies that any $v_E(n, \cdot)$ equals the unique fixed point $\nu_E(n, \cdot)$ of T_n. With (3), this gives a unique $a_E(n, \cdot)$ that defaults to inactivity. By the argument following (10), $a_E(n, \cdot) = \alpha_E(n, \cdot)$. This establishes that $v_E = \nu_E$ and $a_E = \alpha_E$. With (1), these imply a unique value of v_S, ν_S. Finally, Corollary 1 and the requirement that the strategy defaults to inactivity together imply that there is a unique a_S corresponding to this post-entry value, α_S. ■

Supplement to Very Simple Markov-Perfect Industry Dynamics: Theory

Jaap H. Abbring[*]

Jeffrey R. Campbell[†]

Jan Tilly[‡]

Nan Yang[§]

March 2017

This supplement to Abbring, Campbell, Tilly, and Yang (2017b) (hereafter referred to as the "main text") (i) proves a theorem we rely upon for the characterization of equilibria and (ii) develops results for an alternative specification of the model.

Section 1's Theorem S1 establishes that a strategy profile is subgame perfect if no player can benefit from deviating from it in one stage of the game and following it faithfully thereafter. Our proof very closely follows that of the analogous Theorem 4.2 in Fudenberg and Tirole (1991). That theorem only applies to games that are "continuous at infinity" (Fudenberg and Tirole, p. 110), which our game is not. In particular, we only bound payoffs from above (Assumption A1 in the main text) and not from below, because we want the model to encompass econometric specifications

[*]CentER, Department of Econometrics & OR, Tilburg University, and CEPR. E-mail: jaap@abbring.org.

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

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

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

JEL Codes: L13, C25, C73.

Keywords: demand uncertainty, dynamic oligopoly, firm entry and exit, sunk costs, toughness of competition.

The research of Jaap Abbring is financially supported by the Netherlands Organisation for Scientific Research (NWO) through Vici grant 453-11-002. Nan Yang gratefully acknowledges financial support from National University of Singapore Start-up Grant R-313-000-100-133 and an HSS Faculty Research Fellowship. The views expressed herein are those of the authors, and they do not necessarily represent those of the Federal Reserve Bank of Chicago, the Federal Reserve System, or its Board of Governors.

like Abbring, Campbell, Tilly, and Yang's (2017a) that feature arbitrarily large cost shocks. Instead, our proof leverages the presence of a repeatedly-available outside option with a fixed and bounded payoff, exit.

Section 2 presents the primitive assumptions and analysis for an alternative model of Markov-perfect industry dynamics in which one potential entrant makes its entry decision at the same time as incumbents choose between continuation and exit. We show that the alternative model always has a unique symmetric Markov-perfect equilibrium that satisfies an intuitive refinement criterion. This equilibrium can be computed rapidly by a simple algorithm based on contraction mappings similar to that in the main text.

1 One-Shot Deviations and Subgame Perfection

To prove that a strategy forms a subgame-perfect equilibrium if and only if one-shot deviations from it increase no player's payoff, we analyze a *general game* that encompasses that described in the recursive extensive form in Figure 1 of the main text. Like the main text's game, the general game

- is specified in discrete time $t \in \mathbb{N}$,
- is played by firms with names $f \in \mathcal{F}$,
- places an upper bound on each player's flow payoff, and
- regulary offers firms the option to collect a continuation value of zero.

The general game allows for arbitrary differences across players' payoffs, because symmetry is *not* required for this section's result.

Since a subgame of the general game may initialize from multiple points within a given period t, it is useful to focus on *stages* instead of periods. We index these stages with $k \in \mathbb{N}$ and define a function $\theta : \mathbb{N} \to \mathbb{N}$ such that $\theta(k) = t$ when stage k is located in period t. Each period contains the same, finite number $\check{k}$ of stages. In the main text's game, there are $\check{k} = \check{\jmath} + 2$ nontrivial stages each period: the $\check{\jmath} \in \mathbb{N}$ stages at which potential entrants $(t, 1), \ldots, (t, \check{\jmath})$ make their entry decisions, the stage at which all incumbents and new entrants make continuation decisions, and the stage in which nature draws Y'.

Denote the random variable for the entire *history* at stage k with H^k and its realization with h^k. At history h^k, firm f chooses an *action* from the finite (and possibly empty) set $\mathbb{A}^f(h^k)$. At each history h^k, only finitely many firms $f \in \mathcal{F}$ have $\mathbb{A}^f(h^k) \neq \emptyset$. For example, in the main text's game, $\mathbb{A}^f(h^k) = \{0, 1\}$ if firm f gets to decide on entry or survival at history h^k, and $\mathbb{A}^f(h^k) = \emptyset$ otherwise.

A *strategy* σ^f for firm f assigns a probability distribution $\sigma^f(h^k)$ over actions in $\mathbb{A}^f(h^k)$ to each history h^k. A *strategy profile* σ is a collection of strategies, one for each firm. We denote the random action taken by firm f in stage k with A_k^f, with possible realizations $a_k^f \in \mathbb{A}^f(h^k)$. We collect all firms' actions in stage k in the *action profile* A_k, with realizations a_k.

Firms can receive flow payoffs in one or more stages within each period. Let $g^f(a_k, h^k)$ denote the *flow payoff* that firm f receives in stage k with history h^k if firms take actions a_k. The analogue to Assumption A1 for this general game is

Assumption S1 (Flow payoff bounded from above) *There is a $\check{g} < \infty$ such that for any firm f, action profile a_k, and history h^k, we have $g^f(a_k, h^k) \leq \check{g}$.*

In the main text's game, $g^f(a_k, h^k) \equiv \rho \mathbb{E}\left[\pi(n', Y') | Y = y\right] \leq \rho \check{\pi} < \infty$ if stage k is the survival stage, the implied (by a_k and h^k) number of firms surviving that stage is n', the current demand equals y, and firm f is active. If instead stage k contains f's entry decision, the number of active firms prior to f's entry is n, and the current demand is y, then $g^f(a_k, h^k) \equiv -\varphi(n+1, y) \leq 0$ if firm f enters ($a_k^f = 1$). In all other cases, $g^f(a_k, h^k) \equiv 0$.

To complete the general model's specification, let $u^f(\sigma, h^k)$ denote firm f's expected payoff at history h^k, discounted to stage k, when all firms use the strategies in σ. This *continuation value* is defined as

$$u^f(\sigma, h^k) = \lim_{Q \to \infty} \mathbb{E}_\sigma \left[\sum_{q=k}^{Q} \rho^{\theta(q) - \theta(k)} g^f(A_q, H^q) \middle| H^k = h^k \right].$$

Here, the expectation operator's subscript indicates its dependence on all firms following their strategies in σ. Since Assumption S1 ensures that the flow payoff is bounded from above, the limit in the right-hand side is always well defined, either as a real number, which cannot exceed $\check{u} \equiv \frac{\check{k}\check{g}}{1-\rho}$, or as $-\infty$.

A formal statement of the assumption that each firm f in each stage k can collect a continuation payoff of zero within a finite number of stages, irrespective of the strategies followed by the other players, requires the following definition.

Definition S1 (*l*-shot deviation) *Given a firm f and a strategy σ^f, we say that an alternative strategy $\hat{\sigma}^f$ prescribes an l-*shot deviation *from σ^f starting in stage k if*

$$\hat{\sigma}^f(h^{k'}) = \sigma^f(h^{k'})$$

for all possible histories $h^{k'}$; $k' = k + l, k + l + 1, \ldots$.

If $l = 1$, $\hat{\sigma}^f$ prescribes a one-shot deviation. Note that Definition S1 only excludes deviations beyond stage $k+l$ and allows the deviation in earlier stages to be trivial.

Assumption S2 (Exit option) *For each stage k and firm f, there exists a finite $k' \geq k$ such that, for all possible histories $h^{k'}$ at stage k' and strategy profiles σ, $u^f(\hat{\sigma}, h^{k'}) = 0$ for a strategy profile $\hat{\sigma}$ obtained from σ by replacing σ^f with some strategy $\hat{\sigma}^f$ that prescribes a (possibly trivial) one-shot deviation from it in stage k'.*

The main text's game satisfies this assumption. In particular, in that game, a firm f active in stage k will have an option to exit and collect a zero continuation value within $\check{k}$ stages. A firm f that will have an entry opportunity in or after stage k will be able to forgo that entry opportunity and collect zero within a finite number of periods (and therefore stages). A firm f that has exited before stage k will trivially collect $u^f(\sigma, h^k) = 0$ for all possible histories at stage k (in this trivial case, k' can simply be taken equal to k, and no one-shot deviation is needed).

With the general game's specification complete, we begin its analysis with two further definitions.

Definition S2 (*l*-shot-deviation proof) *A strategy profile σ is l-*shot-deviation proof *if for any stage k, history h^k, and firm f; there is no strategy $\hat{\sigma}^f$ that prescribes an l-shot deviation from σ^f starting in stage k such that*

$$u^f(\hat{\sigma}, h^k) > u^f(\sigma, h^k),$$

where $\hat{\sigma}$ is the strategy profile obtained by replacing σ^f with $\hat{\sigma}^f$ in σ.

Definition S3 (Subgame perfection) *A strategy profile* σ *is* subgame perfect *if for any stage* k*, history* h^k *and firm* f*; there is no strategy* $\hat{\sigma}^f$ *such that*

$$u^f(\hat{\sigma}, h^k) > u^f(\sigma, h^k)$$

where $\hat{\sigma}$ *is the strategy profile obtained by replacing* σ^f *with* $\hat{\sigma}^f$ *in* σ*.*

We begin our demonstration that a strategy profile is one-shot-deviation proof if and only if it is subgame perfect with the following lemma. Its proof mimics Fudenberg and Tirole (1991)'s inductive proof of their Theorem 4.1.

Lemma S1 *Any strategy profile* σ *that is one-shot-deviation proof is also* l*-shot deviation proof for any* $l \in \mathbb{N}$*.*

Proof. The lemma is true by assumption for $l = 1$. Suppose that it is true for some $l \in \mathbb{N}$. We wish to demonstrate that it is also true for $l + 1$. Consider a strategy for some firm f, $\hat{\sigma}^f$, that prescribes a single $(l+1)$-shot deviation starting in some stage k. Use $\hat{\sigma}$ to denote the strategy profile obtained from σ by replacing σ^f with $\hat{\sigma}^f$. Next, construct a second strategy $\tilde{\sigma}^f$ that agrees with $\hat{\sigma}^f$ in all stages $k' \leq k + l - 1$ and agrees with σ^f otherwise. The strategy profile obtained from σ by replacing σ^f with $\tilde{\sigma}^f$ is $\tilde{\sigma}$.

Note that $\tilde{\sigma}^f$ prescribes an l-shot deviation from σ^f starting in stage k. Fix an arbitrary history h^k. The presumption that the proposition is true for l implies that

$$\begin{aligned} u^f(\hat{\sigma}, h^k) - u^f(\sigma, h^k) &= u^f(\hat{\sigma}, h^k) - u^f(\tilde{\sigma}, h^k) + u^f(\tilde{\sigma}, h^k) - u^f(\sigma, h^k) \\ &\leq u^f(\hat{\sigma}, h^k) - u^f(\tilde{\sigma}, h^k). \end{aligned} \tag{1}$$

The definition of $u^f(\cdot, h^k)$ gives us

$$\begin{aligned} u^f(\hat{\sigma}, h^k) = &\sum_{q=k}^{k+l-1} \rho^{\theta(q)-\theta(k)} \mathbb{E}_{\hat{\sigma}} \left[g^f(A_q, H^q) \,\middle|\, H^k = h^k \right] \\ &+ \rho^{\theta(k+l)-\theta(k)} \mathbb{E}_{\hat{\sigma}} \left[u^f(\hat{\sigma}, H^{k+l}) \,\middle|\, H^k = h^k \right]. \end{aligned}$$

The same definition and the construction of $\tilde{\sigma}$ yields

$$\begin{aligned}u^f(\tilde{\sigma}, h^k) = &\sum_{q=k}^{k+l-1} \rho^{\theta(q)-\theta(k)} \mathbb{E}_{\hat{\sigma}} \left[g^f(A_q, H^q) \,\middle|\, H^k = h^k \right] \\ &+ \rho^{\theta(k+l)-\theta(k)} \mathbb{E}_{\hat{\sigma}} \left[u^f(\sigma, H^{k+l}) \,\middle|\, H^k = h^k \right] \\ = &u^f(\hat{\sigma}, h^k) + \rho^{\theta(k+l)-\theta(k)} \mathbb{E}_{\hat{\sigma}} \left[u^f(\sigma, H^{k+l}) \,\middle|\, H^k = h^k \right] \\ &- \rho^{\theta(k+l)-\theta(k)} \mathbb{E}_{\hat{\sigma}} \left[u^f(\hat{\sigma}, H^{k+l}) \,\middle|\, H^k = h^k \right].\end{aligned}$$

Since $\hat{\sigma}^f$ prescribes a one-shot deviation from σ^f in stage $k+l$, we know that $u^f(\hat{\sigma}, h^{k+l}) \leq u^f(\sigma, h^{k+l})$ for all possible histories h^{k+l} in stage $k+l$. Therefore,

$$u^f(\hat{\sigma}, h^k) - u^f(\tilde{\sigma}, h^k) = \rho^{\theta(k+l)-\theta(k)} \mathbb{E}_{\hat{\sigma}} \left[u^f(\hat{\sigma}, H^{k+l}) - u^f(\sigma, H^{k+l}) \,\middle|\, H^k = h^k \right] \leq 0.$$

The above inequality and (1) jointly imply that $u^f(\hat{\sigma}, h^k) - u^f(\sigma, h^k) \leq 0$. That is, the proposed $(l+1)$-shot deviation fails to increase firm f's payoff. Continuing inductively establishes the desired conclusion. ■

With this lemma in hand, we state and prove this section's central theorem, which establishes the necessity and sufficiency of one-shot-deviation proofness for subgame perfection.

Theorem S1 *A strategy profile σ is subgame perfect if and only if it is one-shot-deviation proof.*

Proof. Necessity directly follows from the definition of subgame perfection. To demonstrate sufficiency, assume that firm f uses a strategy $\hat{\sigma}^f$ that prescribes a deviation from some stage k and use $\hat{\sigma}$ to denote the strategy profile obtained from σ by replacing σ^f with $\hat{\sigma}^f$. If $\hat{\sigma}^f$ is an l-shot deviation from σ^f, Lemma S1 ensures that it cannot improve on σ^f.

Suppose that $\hat{\sigma}^f$ instead prescribes a deviation in infinitely many stages. Fix an arbitrary history h^k and some finite $k' > k$. Assumption S2 guarantees that there is a finite stage $k'' \geq k'$, such that firm f can guarantee a payoff $u^f(\cdot, h^{k''}) = 0$ for all possible histories $h^{k''}$ and all strategy profiles by a (possibly trivial) one-shot deviation in stage k''. Now construct an alternative deviating strategy $\tilde{\sigma}^f$ that agrees with $\hat{\sigma}^f$ until (but not including) k'' and agrees with σ^f afterwards. The strategy profile obtained from σ by replacing σ^f with $\tilde{\sigma}^f$ is $\tilde{\sigma}$. The construction of $\tilde{\sigma}$ gives us

$$u^f(\hat{\sigma}, h^k) - u^f(\tilde{\sigma}, h^k) = \rho^{\theta(k'') - \theta(k)} \mathbb{E}_{\hat{\sigma}} \left[u^f(\hat{\sigma}, H^{k''}) - u^f(\sigma, H^{k''}) \,\middle|\, H^k = h^k \right].$$

Since σ is one-shot-deviation proof and Assumption S2 ensures that firm f can earn a continuation value of zero by a one-shot deviation from σ in stage k'', we have

$$\mathbb{E}_{\hat{\sigma}} \left[u^f(\sigma, H^{k''}) \,\middle|\, H^k = h^k \right] \geq 0.$$

So

$$u^f(\hat{\sigma}, h^k) - u^f(\tilde{\sigma}, h^k) \leq \rho^{\theta(k'') - \theta(k)} \mathbb{E}_{\hat{\sigma}} \left[u^f(\hat{\sigma}, H^{k''}) \,\middle|\, H^k = h^k \right] \leq \rho^{\theta(k'') - \theta(k)} \breve{u}.$$

With this inequality in hand, we can write

$$\begin{aligned} u^f(\hat{\sigma}, h^k) - u^f(\sigma, h^k) =& u^f(\hat{\sigma}, h^k) - u^f(\tilde{\sigma}, h^k) + u^f(\tilde{\sigma}, h^k) - u^f(\sigma, h^k) \\ \leq& u^f(\hat{\sigma}, h^k) - u^f(\tilde{\sigma}, h^k) \\ \leq& \rho^{\theta(k'') - \theta(k)} \breve{u} \leq \rho^{\theta(k') - \theta(k)} \breve{u}. \end{aligned}$$

Here, the first inequality follows from Lemma S1. Since k' was arbitrary and $\lim_{k' \to \infty} \theta(k') = \infty$, we conclude that $u^f(\hat{\sigma}, h^k) - u^f(\sigma, h^k) \leq 0$. Therefore, the proposed deviation fails to improve firm f's payoff and the strategy profile is subgame perfect. ■

2 An Alternative Model

The remainder of this supplement presents an alternative model in which one potential entrant per period makes an entry decision at the same time incumbent firms choose between continuation and exit.

2.1 Primitives

The alternative model considered here differs from the one presented in the main text in two primitive assumptions. First, each period has exactly one potential entrant. Second, the single potential entrant makes its entry decision at the same time incumbents make their continuation decisions. The recursive extensive form in

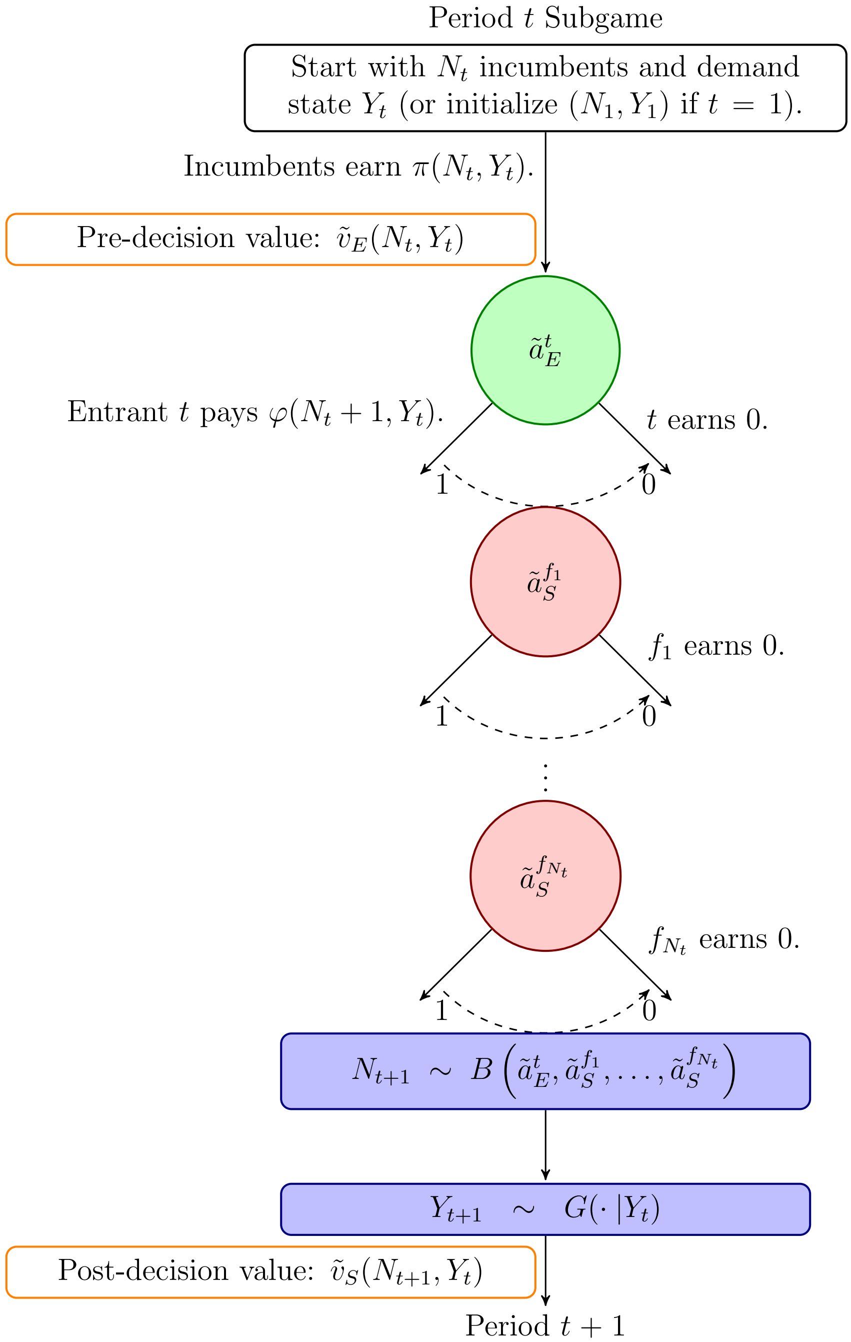

Figure 1: The Alternative Model's Recursive Extensive Form

Figure 1 presents the timing. All primitive assumptions from the main text on per period profit $\pi(n, y)$ and the cost of entry $\varphi(n, y)$ remain in place.

In this model, there is no distinction between the entry stage and the survival stage. Each period of the alternative model starts with incumbent firms earning the per period profit $\pi(N_t, Y_t)$. Then, one potential entrant with name t contemplates becoming active *at the same time* incumbents make their continuation decisions. (Incumbent firms in period 0 have arbitrary names in $\mathcal{F}_0$.) Entry requires paying the sunk cost $\varphi(N_t + 1, Y_t)$. As in the model presented in the main text, simultaneously moving firms, including the potential entrant, can use mixed strategies. All entry and survival outcomes are realized independently across firms according to the chosen Bernoulli distributions. Firms that exit or fail to enter earn zero and never again participate in the market. Active firms continue to the next period.

2.2 Equilibrium

For incumbent firms contemplating survival, the payoff-relevant variables are the number of active firms at the beginning of the period and the current demand state. For a potential entrant, they are this same number of firms plus one (that is, the number of firms that would serve the market next period if this firm would enter and all incumbents would continue) and the current demand state. Note that this implies that incumbents that decide on survival in a state (n, y) move jointly with a potential entrant in state $(n + 1, y)$. We have made this, perhaps counterintuitive, choice, because it is consistent with the specification of the state variables in the main text and will simplify notation later.

As in the main text, we focus our analysis on symmetric Markov-perfect equilibria: subgame perfect equilibria in which all players use a common Markov strategy. In this context, a Markov strategy is a pair of functions, $\tilde{a}_E : \mathbb{N}\times\mathcal{Y} \to [0, 1]$ and $\tilde{a}_S : \mathbb{N} \times \mathcal{Y} \to [0, 1]$. Firms' values at two nodes of the game tree—just before firms' entry/continuation decisions and just after the decisions' realizations—are important for the equilibrium analysis. Although there are no separate entry and survival stages, we use notation analogous to the main text and denote the *pre-decision* and *post-decision* values as $\tilde{v}_E$ and $\tilde{v}_S$. Figure 1 shows the points where these value functions apply.

An equilibrium strategy and its associated values satisfy

$$\tilde{v}_E(n,y) = \max_{a\in[0,1]} a\mathbb{E}_{\tilde{a}_S,\tilde{a}_E}\left[\tilde{v}_S(N',y)|N=n,Y=y\right], \tag{2}$$

$$\tilde{v}_S(n',y) = \rho\mathbb{E}\left[\pi(n',Y')+\tilde{v}_E(n',Y')|Y=y\right], \tag{3}$$

$$\tilde{a}_E(n+1,y) \in \arg\max_{a\in[0,1]} a\big(\mathbb{E}_{\tilde{a}_S}\left[\tilde{v}_S(N',y)|N=n,Y=y\right]-\varphi(n+1,y)\big), \tag{4}$$

$$\tilde{a}_S(n,y) \in \arg\max_{a\in[0,1]} a\mathbb{E}_{\tilde{a}_S,\tilde{a}_E}\left[\tilde{v}_S(N',y)|N=n,Y=y\right]. \tag{5}$$

The expectation operators condition on the deciding firm choosing to be active in the next period and on all other firms using the entry or exit rule in the operator's subscript. Theorem S1 ensures that (2)–(5) are not only necessary but also sufficient conditions for a symmetric Markov-perfect equilibrium.

2.3 Equilibrium Multiplicity and Refinement

There can be more than one symmetric Markov-perfect equilibrium in this model. Unlike the model in the main text, it involves simultaneous moves by firms with heterogeneous payoffs (incumbents on the one hand and a potential entrant on the other). We demonstrate how this might lead to equilibrium multiplicity with a numerical example.

Example S1 *Suppose that the aggregate state Y follows a deterministic first-order Markov process such that $Y_1=3$, $Y_2=2$, and $Y_t=0$ for all $t\geq 3$. Specify the per period profit as*

$$\pi(n,y)=y/n-1.5,$$

and assume that $\rho>0$ and $\varphi(n,y)=\rho/4$. In this game, any firm serving the market in period 3 or beyond will earn a negative payoff. Therefore, in equilibrium, no firms will enter and all incumbents will exit in period 2 (demand state 2) and beyond (demand state 0).

In period 1 (demand state 3), equilibrium play is not so trivial. Suppose that one incumbent is active at the beginning of period 1. The static game (the analogue of Corollary 1's static game in the main text) between the potential entrant and the incumbent has two symmetric equilibria in pure strategies—either the potential entrant enters and the incumbent exits, or the incumbent stays active and the *potential entrant stays out. The implied equilibrium entry and survival rules satisfy (1)* $(\tilde{a}_E(2,3) = 1, \tilde{a}_S(1,3) = 0)$ *and (2)* $(\tilde{a}_E(2,3) = 0, \tilde{a}_S(1,3) = 1)$. *There is also a mixed-strategy equilibrium, which implies* $(\tilde{a}_E(2,3) = 0.5, \tilde{a}_S(1,3) = 0.25)$. *These are all symmetric equilibria because all players use the same strategy, even though that strategy's actions depend nontrivially on whether the player is a potential entrant or an incumbent at a particular node of the game tree.*

Since an entrant always needs to pay a sunk cost to become active, its expected payoff from becoming active is strictly lower than that of an incumbent from remaining active. If an equilibrium strategy dictated that the potential entrant enters while an incumbent chooses to exit, that incumbent could make a side payment to the potential entrant in return for focusing on an alternative equilibrium strategy in which their roles were reversed. With this in mind, we focus on equilibria in which entry only occurs when all incumbents choose sure continuation. We label these "natural."

Definition S4 (Natural Markov-perfect equilibrium) *A* natural *Markov-perfect equilibrium is a symmetric Markov-perfect equilibrium* $(\tilde{a}_E, \tilde{a}_S)$ *such that for all* $(n, y) \in \mathbb{N} \times \mathcal{Y}$, $\tilde{a}_E(n+1, y) > 0$ *implies* $\tilde{a}_S(n, y) = 1$.

A analogue of the main text's Lemma 1 establishes that we can again restrict the equilibrium analysis to states (n, y) in $\{1, \ldots, \check{n}\} \times \mathcal{Y}$.

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

Proof of Lemma 1*. First, we prove that $\tilde{a}_S(n, y) < 1$ for all $y \in \mathcal{Y}$ and $n > \check{n}$. Consider a period $t^\star$ subgame with $N_{t^\star} = n > \check{n}$ incumbent firms and demand state $Y_{t^\star} = y$. Define the random time τ as the first period weakly after $t^\star$ in which the firms choose exit with positive probability, with $\tau \equiv \infty$ if they never do:

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

Suppose that $\tilde{a}_S(n, y) = 1$, so $\tau > t^\star$. By definition, exit occurs only in or after period τ, so we know that $N_t \geq n$ for $t \in \{t^\star + 1, \ldots, \tau\}$. (As in the main text, we take $\tilde{a}_S(\cdot) = 1$ to dictate sure survival, not merely almost-sure survival.) Since $n > \check{n}$, this together with Assumption A2 implies that $\pi(N_t, Y_t) < 0$ for

$t \in \{t^\star+1,\ldots,\tau\}$. If $\tau=\infty$, then the incumbent firms receive an infinite sequence of strictly negative payoffs. If instead $\tau<\infty$, then the incumbent firms receive a finite sequence of strictly negative payoffs followed by the pre-decision value from playing the period τ subgame $\tilde{v}_E(N_\tau, Y_\tau)$, which equals zero by (2), (5), and the definition of τ. Therefore, the period $t^\star$ post-decision value satisfies $\tilde{v}_S(n,y)<0$. Since a period $t^\star$ incumbent firm can raise its payoff to zero by choosing certain exit, the supposition that $\tilde{a}_S(n,y)=1$ must be incorrect.

Next, we will prove that $\tilde{a}_E(n,y)=0$ for all $y \in \mathcal{Y}$ and $n>\check{n}$. Suppose that $\tilde{a}_E(n,y)>0$ for some $n>\check{n}$. Since we are considering a natural equilibrium, we know that $\tilde{a}_S(n,y)=1$. Therefore, the entrant pays $\varphi(n+1,y)>0$ to enter and receives $\pi(n+1,Y')<0$ from production in the next period. We have already proven that this firm's pre-decision continuation value in the next period equals zero, so the firm earns a strictly negative payoff from entry. Since it could improve its payoff by choosing not to enter for sure ($\tilde{a}_E(n,y)=0$), the supposition that $\tilde{a}_E(n,y)>0$ must be incorrect. ■

As in the main text, we hereafter restrict attention to equilibria in strategies that *default to inactivity*. A strategy that defaults to inactivity requires a potential entrant that is indifferent between continuing with all n incumbents (as it would if it continued in any natural equilibrium) or not entering to stay out:

$$\tilde{v}_S(n+1,y)=\varphi(n+1,y) \Rightarrow \tilde{a}_E(n+1,y)=0.$$

Similarly, such a strategy requires an incumbent firm that is indifferent between exit and continuing with any combination of the current incumbents, and that weakly prefers exit over continuing with all potentially active firms, to exit:

$$\tilde{v}_S(1,y)=\cdots=\tilde{v}_S(n,y)=0 \text{ and } \tilde{v}_S(n+1,y)\leq 0 \Rightarrow \tilde{a}_S(n,y)=0.$$

2.4 Equilibrium Analysis

The natural equilibrium refinement together with the default-to-inactivity restriction on equilibrium strategies (which, as in the main text, we will keep implicit in what follows) give us the following lemma.

Lemma S2 *The entry strategy in a natural Markov-perfect equilibrium is pure:* $\tilde{a}_E(n,y)\in\{0,1\}$ *for all* $n\in\{1,\ldots,\check{n}\}$ *and all* $y\in\mathcal{Y}$.

Proof. Since a natural equilibrium requires $\tilde{a}_S(n,y) = 1$ when $\tilde{a}_E(n,y) > 0$, equilibrium condition (4) can be rewritten as

$$\tilde{a}_E(n,y) \in \arg\max_{a\in[0,1]} a\big(\tilde{v}_S(n+1,y) - \varphi(n,y)\big).$$

Therefore, $\tilde{a}_E(n,y) = 1$ if $\tilde{v}_S(n+1,y) > \varphi(n,y)$ and $\tilde{a}_E(n,y) = 0$ if $\tilde{v}_S(n+1,y) < \varphi(n,y)$. When $\tilde{v}_S(n+1,y) = \varphi(n,y)$, the default-to-inactivity restriction requires $\tilde{a}_E(n,y) = 0$. ■

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

The proof of Lemma 2⋆ is identical to that of Lemma 2 in the main text, except for minor changes in terminology and state variables related to the change in timing. We repeat it here for completeness only.

Proof of Lemma 2⋆. It suffices to prove that $\tilde{v}_S(n',y) \geq \tilde{v}_S(n'+1,y)$ for all $n' \geq 1$ and demand states $y \in \mathcal{Y}$. To this end, consider a subgame beginning immediately after the period $t^\star$'s simultaneous continuation and entry choices with $N_{t^\star+1} = n'$ and $Y_{t^\star} = y$. We call this the *original* subgame. Now consider a second period $t^\star$ subgame starting at the same point but with one additional firm. We refer to this as the *perturbed* subgame and use N_t^+ to denote the number of firms serving the market during period t within it. Finally, define the random time τ^+ as the first period weakly after $t^\star+1$, in which the firms in the perturbed subgame choose exit with positive probability, with $\tau^+ \equiv \infty$ if they never do:

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

There is no exit before period τ^+ in the perturbed subgame and the period τ^+ pre-decision continuation value equals zero. Therefore, we can write

$$\tilde{v}_S(n'+1,y) = \lim_{T\to\infty} \mathbb{E}\left[\sum_{t=t^\star+1}^{T} \rho^{t-t^\star} \mathbb{1}\left[t \leq \tau^+\right] \pi(N_t^+, Y_t) \middle| Y_{t^\star} = y\right].$$

Since τ^+ is a consequence of equilibrium choices, we know that $v_S(n'+1,y) > -\infty$.

Now consider an incumbent firm in the original subgame which (possibly) deviates after the period $t^\star$ survival stage by choosing to survive for sure as long as $t < \tau^+$ and to exit for sure if $t = \tau^+$. Let $\bar{N}_t$ denote the number firms serving the market during period t in the original subgame *with this deviation.* Since the original strategy was part of a subgame perfect equilibrium, $\tilde{v}_S(n', y)$ exceeds the expected payoff from following this deviating strategy. That is

$$\tilde{v}_S(n', y) \geq \lim_{T \to \infty} \mathbb{E}\left[\sum_{t=t^\star+1}^{T} \rho^{t-t^\star} \mathbb{1}\left[t \leq \tau^+\right] \pi(\bar{N}_t, Y_t) \middle| Y_{t^\star} = y\right].$$

To show that the limit on the right-hand side is well defined, note that $\bar{N}_t \leq N_t^+$ for all $t \leq \tau^+$. Otherwise, the two subgames would have potential entrants in the same states making different entry choices. This would violate either the presumption that the equilibrium strategy is Markov or that it defaults to inactivity. This and Assumption A3 imply that $\pi(\bar{N}_t, Y_t) \geq \pi(N_t^+, Y_t)$ for all $t = t^\star + 1, \ldots, \tau^+$. Combining this with $v_S(n' + 1, y) > -\infty$ gives the desired result.

Because the difference of two convergent sequences' limits equals the limit of the sequences' difference, we can write

$$\begin{aligned} v_S(n', y) - & v_S(n' + 1, y) \\ & \geq \lim_{T \to \infty} \mathbb{E}\left[\sum_{t=t^\star+1}^{T} \rho^{t-t^\star} \mathbb{1}\left[t \leq \tau^+\right] \left(\pi(\bar{N}_t, Y_t) - \pi(N_t^+, Y_t)\right) \middle| Y_{t^\star} = y\right]. \end{aligned}$$

Each term in the partial sum on the right-hand side is non-negative, so we conclude that $v_S(n', y) - v_S(n' + 1, y) \geq 0$. ■

Corollary 1* *Fix $n \in \{1, \ldots, \check{n}\}$ and $y \in \mathcal{Y}$, let $\tilde{v}_S$ be the post-decision value function associated with a natural Markov-perfect equilibrium that defaults to inactivity, and suppose that $\tilde{v}_S(n', y) \neq 0$ for at least one $n' \in \{1, \ldots, n\}$ and that $\tilde{v}_S(n + 1, y) \neq \varphi(n + 1, y)$. Consider the one-shot game in which n incumbent firms and one potential entrant simultaneously choose between activity and inactivity. This game has a symmetric Nash equilibrium—possibly in mixed strategies—in which the potential entrant's chosen probability of entry does not exceed the incumbents' chosen probability of survival. Furthermore, there is only one symmetric Nash equilibrium with this property. In it, the entry strategy is pure.*

Proof of Corollary 1***.** Corollary 1*'s one-shot game falls into one of four mutually-exclusive cases.

- $\tilde{v}_S(1, y) \leq 0$. Lemma 2* implies that $\tilde{v}_S(n', y) \leq 0$ for all $n' > 1$. Therefore, exiting for sure (setting $\tilde{a}_S(n, y) = 0$) is a weakly dominant strategy for an incumbent. Furthermore, since $\tilde{v}_S(n', y) \neq 0$ for at least one $n' \in \{1, \ldots, n\}$, we know that $\tilde{v}_S(n + 1, y) \leq \tilde{v}_S(n, y) < 0$. Therefore, exiting for sure is also an incumbent's unique best response to any positive symmetric continuation probability and any entry probability. Since $\tilde{v}_S(n + 1, y) < 0 < \varphi(n + 1, y)$, not entering for sure (setting $\tilde{a}_E(n + 1, y) = 0$) is a strictly dominant strategy for a potential entrant. Therefore, there is only one symmetric equilibrium, in which all incumbents exit for sure and the potential entrant stays out for sure.

- $\tilde{v}_S(1, y) > 0$ and $\tilde{v}_S(n, y) < 0$. To construct the equilibrium of interest, set $\tilde{a}_E(n + 1, y) = 0$. No symmetric equilibrium exists with this entry choice and a pure continuation strategy, because an incumbent's best response to all other incumbent firms continuing for sure is to exit for sure, and vice versa. In a mixed strategy equilibrium, incumbent firms must be indifferent between continuation and exit. By the intermediate value theorem, there is some $a \in (0, 1)$ that solves

$$\sum_{n'=1}^{n} \binom{n-1}{n'-1} a^{n'-1} (1-a)^{n-n'} \tilde{v}_S(n', y) = 0. \tag{6}$$

  Lemma 2* guarantees that the left-hand side of (6) weakly decreases with a, and the subcase's conditions strengthen that conclusion so that the left-hand side of (6) strictly decreases with a. So, there is only one value of a that solves (6). The pair $\tilde{a}_S(n, y) = a$ and $\tilde{a}_E(n + 1, y) = 0$ form one Nash equilibrium to this game.

  To see that this is the only equilibrium in which $\tilde{a}_E(n + 1, y) \leq \tilde{a}_S(n, y)$, suppose to the contrary that there exists another such equilibrium in which $\tilde{a}_E(n + 1, y) > 0$. Lemma 2* and the conditions of this case guarantee that $\tilde{v}_S(n, y) < 0$ and $\tilde{v}_S(n + 1, y) < 0$, so $\tilde{a}_S(n, y) < 1$. Since $\tilde{a}_S(n, y) \geq \tilde{a}_E(n + 1, y) > 0$ by assumption, we therefore know that both the potential entrant and the incumbents are indifferent between continuation and exit. The payoff to a potential entrant is strictly below the payoff to an incumbent in any outcomes of the game in which both are active, because only the potential entrant pays a positive entry cost. With Lemma 2*, this implies that both can only be indifferent if the potential entrant is more likely to end up with fewer competitors, $\tilde{a}_S(n,y) < \tilde{a}_E(n+1,y)$. This violates the restriction that $\tilde{a}_S(n,y) \geq \tilde{a}_E(n+1,y)$.

- $\tilde{v}_S(n,y) \geq 0$ and $\tilde{v}_S(n+1,y) < \varphi(n+1,y)$. To construct the equilibrium of interest, set $\tilde{a}_E(n+1,y) = 0$. Since $\tilde{v}_S(n,y) \geq 0$, pairing this with $\tilde{a}_S(n,y) = 1$ forms one Nash equilibrium. To show that this is the only equilibrium in which $\tilde{a}_E(n+1,y) \leq \tilde{a}_S(n,y)$, we look at two cases.

  - First, suppose that there exists another such equilibrium in which $\tilde{a}_E(n+1,y) = 0$ and $\tilde{a}_S(n,y) \in [0,1)$. By assumption, there exists an $n' \leq n$ such that $\tilde{v}_S(n',y) \neq 0$. Lemma 2* and the supposition that $\tilde{v}_S(n,y) \geq 0$ together imply that $\tilde{v}_S(n^\star,y) > 0$ for all $n^\star \leq n'$. Therefore, the payoff to continuing with probability a when all other incumbents continue with probability $\tilde{a}_S(n,y) \in [0,1)$ is strictly increasing in a, so continuing with any probability greater than $\tilde{a}_S(n,y)$ increases a firm's profit. Therefore, the original value of $\tilde{a}_S(n,y) \in [0,1)$ paired with $\tilde{a}_E(n+1,y) = 0$ cannot have formed an equilibrium.

  - Second, suppose that there exists another such equilibrium in which $\tilde{a}_E(n+1,y) > 0$. If $\tilde{a}_S(n,y) = 1$, then the payoff to the potential entrant would equal $\tilde{a}_E(n+1,y)\left(\tilde{v}_S(n+1,y) - \varphi(n+1,y)\right) < 0$ in this equilibrium. Hence, the potential entrant could profitably deviate to $\tilde{a}_E(n+1,y) = 0$, so there cannot be an equilibrium as supposed with $\tilde{a}_S(n,y) = 1$. If instead $\tilde{a}_S(n,y) < 1$, then both the potential entrant and all incumbents are indifferent between activity and inactivity. With Lemma 2*, this implies that both can only be indifferent if the potential entrant is more likely to end up with fewer competitors, so $\tilde{a}_S(n,y) < \tilde{a}_E(n+1,y)$. This violates the restriction that $\tilde{a}_S(n,y) \geq \tilde{a}_E(n+1,y)$ Therefore, there cannot be an equilibrium as supposed with $\tilde{a}_S(n,y) < 1$ either.

  Therefore, $\tilde{a}_E(n+1,y) = 0$ and $\tilde{a}_S(n,y) = 1$ form the only symmetric equilibrium in which $\tilde{a}_E(n,y) \leq \tilde{a}_S(n,y)$.

- $\tilde{v}_S(n+1,y) > \varphi(n+1,y)$. Lemma 2* and the case's precondition together give us that $v(n', y) > 0$ for all $n' \in \{1, \ldots, n+1\}$, so sure continuation (setting $\tilde{a}_S(n, y) = 1$) is a dominant strategy for each incumbent. For the potential entrant, since $\tilde{v}_S(n+1, y) > \varphi(n+1, y)$, entering for sure (setting $\tilde{a}_E(n, y) = 1$) is a strictly dominant strategy. Therefore, there is a unique Nash equilibrium. In it, all incumbents choose sure continuation and the potential entrant chooses sure entry.

This establishes the equilibrium existence and uniqueness asserted by Corollary $1^\star$.

■

Corollary $2^\star$ *If $\tilde{v}_E$ and $\tilde{v}_S$ are the pre-decision and post-decision value functions associated with a natural Markov-perfect equilibrium and $\tilde{a}_E$ is that equilibrium's entry rule, then*

$$\tilde{v}_E(n, y) = \max\{0, \tilde{a}_E(n+1, y)\tilde{v}_S(n+1, y) + (1 - \tilde{a}_E(n+1, y))\tilde{v}_S(n, y)\}.$$

Proof of Corollary $2^\star$. There are two cases to consider:

- $\tilde{v}_S(1, y) = \tilde{v}_S(2, y) = \cdots \tilde{v}_S(n, y) = 0$. Lemma $2^\star$ implies that $\tilde{v}_S(n+1, y) \leq 0$, so setting $\tilde{a}_E(n+1, y) = 0$ is a dominant strategy for the potential entrant. Therefore, the value of certain continuation for any incumbent equals 0 and the value functions and entry rule satisfy the stated equality.

- There exists an $n' \in \{1, \ldots, n\}$ such that $\tilde{v}_S(n', y) \neq 0$. This case contains two subcases:

  - $\tilde{v}_S(n+1, y) \neq \varphi(n+1, y)$. Corrolary $1^\star$ applies. Its proof demonstrates that

    $$\begin{aligned} &\mathbb{E}_{\tilde{a}_E, \tilde{a}_S}\left[\tilde{v}_S(N', y) | N = n, Y = y\right] \\ &\qquad = \tilde{a}_E(n+1, y)\tilde{v}_S(n+1, y) + (1 - \tilde{a}_E(n+1, y))\tilde{v}_S(n, y) \end{aligned}$$

    if $\tilde{v}_S(n, y) \geq 0$ (in which case $\tilde{a}_S(n, y) = 1$) and

    $$\mathbb{E}_{\tilde{a}_E, \tilde{a}_S}\left[\tilde{v}_S(N', y) | N = n, Y = y\right] = 0$$

    if $\tilde{v}_S(n, y) < 0$ (in which case $\tilde{a}_S(n, y) < 1$). Thus, the value functions and entry rule satisfy the stated equality.

– $\tilde{v}_S(n+1,y) = \varphi(n+1,y)$. Lemma 2⋆ guarantees that $\tilde{v}_S(n',y) > 0$ for all $n' \in \{1,\ldots,n+1\}$, so sure continuation ($\tilde{a}_S(n,y) = 1$) is a dominant strategy for each incumbent. Since the equilibrium defaults to inactivity, $\tilde{a}_E(n+1,y) = 0$. Thus, the value functions and entry rule satisfy the stated equality.

■

Just as with the model of the main text, we constructively demonstrate equilibrium existence and uniqueness. We present the algorithm for equilibrium calculation in Supplementary Algorithm 1. It begins by initializing the number of firms under consideration (n) to $\check{n}$ and both the candidate equilibrium entry rule $\tilde{\alpha}_E : \mathbb{N} \times \mathcal{Y} \to \{0,1\}$ and a dummy function $f^\star : \mathcal{Y} \to [0, \frac{\rho\check{\pi}}{1-\rho}]$ to zero. We will denote the candidate equilibrium pre-decision and post-decision value functions by $\tilde{\nu}_E$ and $\tilde{\nu}_S$, respectively. The algorithm then enters its main loop, which begins by using Bellman equation iteration (on the dummy function $f^\star$) to solve a dynamic programming problem. The relevant Bellman operator is

$$\begin{aligned} T_n(f^\star)(y) = &\max\{0, (1-\tilde{\alpha}_E(n+1,y))\rho\mathbb{E}[\pi(n,Y') + f^\star(Y') \,|\, Y=y] \\ &+ \tilde{\alpha}_E(n+1,y)\rho\mathbb{E}[\pi(n+1,Y') + \tilde{\nu}_E((n+1,Y'),Y') \,|\, Y=y]\}. \end{aligned} \tag{7}$$

The next two steps assign the fixed point of T_n stored in $f^\star(\cdot)$ to $\tilde{\nu}_E(n,\cdot)$ and use (3) to construct $\tilde{\nu}_S(n,\cdot)$. In the loop's final step, $\tilde{\alpha}_E(n,y)$ is set to $\mathbb{1}[\tilde{\nu}_S(n,y) > \varphi(n,y)]$. If the current value of n exceeds one, it is decremented and the algorithm returns to the top of the main loop. If instead n equals one, then the algorithm proceeds to its final task of setting the candidate equilibrium survival rule, which closely follows its calculation by Algorithm 1 in the main text. Supplementary Algorithm 1 only computes candidate post-entry and post-survival values and a candidate survival rule on $\{1,\ldots,\check{n}\} \times \mathcal{Y}$. As noted in the main text, it is straightforward to extend them to the full state space $\mathbb{N}\times\mathcal{Y}$. Because this extension is of little applied interest, we keep it implicit and simply refer to Supplementary Algorithm 1 as computing a candidate equilibrium.

The appropriately modified version of Theorem 1 states that the candidate equilibrium strategies and payoffs arising from Supplementary Algorithm 1 correspond to the unique natural Markov-perfect equilibrium.

Theorem 1* **(Equilibrium existence and uniqueness)** *There exists a unique natural Markov-perfect equilibrium. The equilibrium strategy and corresponding equilibrium payoffs are those computed by Supplementary Algorithm 1.*

Proof of Theorem 1*. The proof is divided into three parts. First, we show that the candidate continuation values from Supplementary Algorithm 1 satisfy the monotonicity requirements of Lemma 2*. Second, we use this to demonstrate that the candidate strategy indeed forms an equilibrium. Third, we establish equilibrium uniqueness.

Fix $n \in \{1, 2, \ldots, \check{n} - 1\}$ and suppose that we know that $\tilde{\nu}_E(n+1, \cdot) \geq \cdots \geq \nu_E(\check{n}, \cdot)$. This immediately implies that $\tilde{\alpha}_E(n+1, \cdot) \geq \cdots \geq \tilde{\alpha}_E(\check{n}, \cdot)$. Evaluating T_n at $f^\star(\cdot) = \tilde{\nu}_E(n+1, \cdot)$ gives

$$\begin{aligned}
T_n(f^\star)(y) = &\max\{0, (1 - \tilde{\alpha}_E(n+1, y))\rho\mathbb{E}[\pi(n, Y') + f^\star(Y') \,|\, Y = y] \\
&\quad + \tilde{\alpha}_E(n+1, y)\rho\mathbb{E}[\pi(n+1, Y') + \tilde{\nu}_E(n+1, Y') \,|\, Y = y]\} \\
\geq &\max\{0, (1 - \tilde{\alpha}_E(n+1, y))\rho\mathbb{E}[\pi(n+1, Y') + f^\star(Y') \,|\, Y = y] \qquad (8) \\
&\quad + \tilde{\alpha}_E(n+1, y)\rho\mathbb{E}[\pi(n+2, Y') + \tilde{\nu}_E(n+2, Y') \,|\, Y = y]\} \\
\geq &\max\{0, (1 - \tilde{\alpha}_E(n+2, y))\rho\mathbb{E}[\pi(n+1, Y') + f^\star(Y') \,|\, Y = y] \qquad (9) \\
&\quad + \tilde{\alpha}_E(n+2, y)\rho\mathbb{E}[\pi(n+2, Y') + \tilde{\nu}_E(n+2, Y') \,|\, Y = y]\} \\
= &\,\tilde{\nu}_E(n+1, y).
\end{aligned}$$

for any $y \in \mathcal{Y}$. (For the case with $n = \check{n}-1$, we define $\tilde{\nu}_E(\check{n}+1, \cdot) = \tilde{\alpha}_E(\check{n}+1, \cdot) = 0$.) The inequality in (8) follows from Assumption A3 and the presumption that $\tilde{\nu}_E(n+1, \cdot) \geq \tilde{\nu}_E(n+2, \cdot)$. The inequality in (9) follows from $\tilde{\alpha}_E(n+2, y) \geq \tilde{\alpha}_E(n+1, y)$ and

$$\mathbb{E}[\pi(n+1, Y') + f^\star(Y') \,|\, Y = y] \geq \mathbb{E}[\pi(n+2, Y') + \tilde{\nu}_E(n+2, Y') \,|\, Y = y].$$

The final equality follows from the equivalence of $\tilde{\nu}_E(n+1, Y')$ with $f^\star(Y')$. The operator T_n is a monotone contraction mapping, so $T_n(f^\star)(n, \cdot) \geq \tilde{\nu}_E(n+1, \cdot)$ implies that $\tilde{\nu}_E(n, \cdot) \geq \tilde{\nu}_E(n+1, \cdot)$. Recursively applying this argument for n decreasing from $\check{n}-1$ to 1 proves that $\tilde{\nu}_E(1, \cdot) \geq \tilde{\nu}_E(2, \cdot) \cdots \geq \tilde{\nu}_E(\check{n}, \cdot)$ and $\tilde{\alpha}_E(1, \cdot) \geq \tilde{\alpha}_E(2, \cdot) \geq$

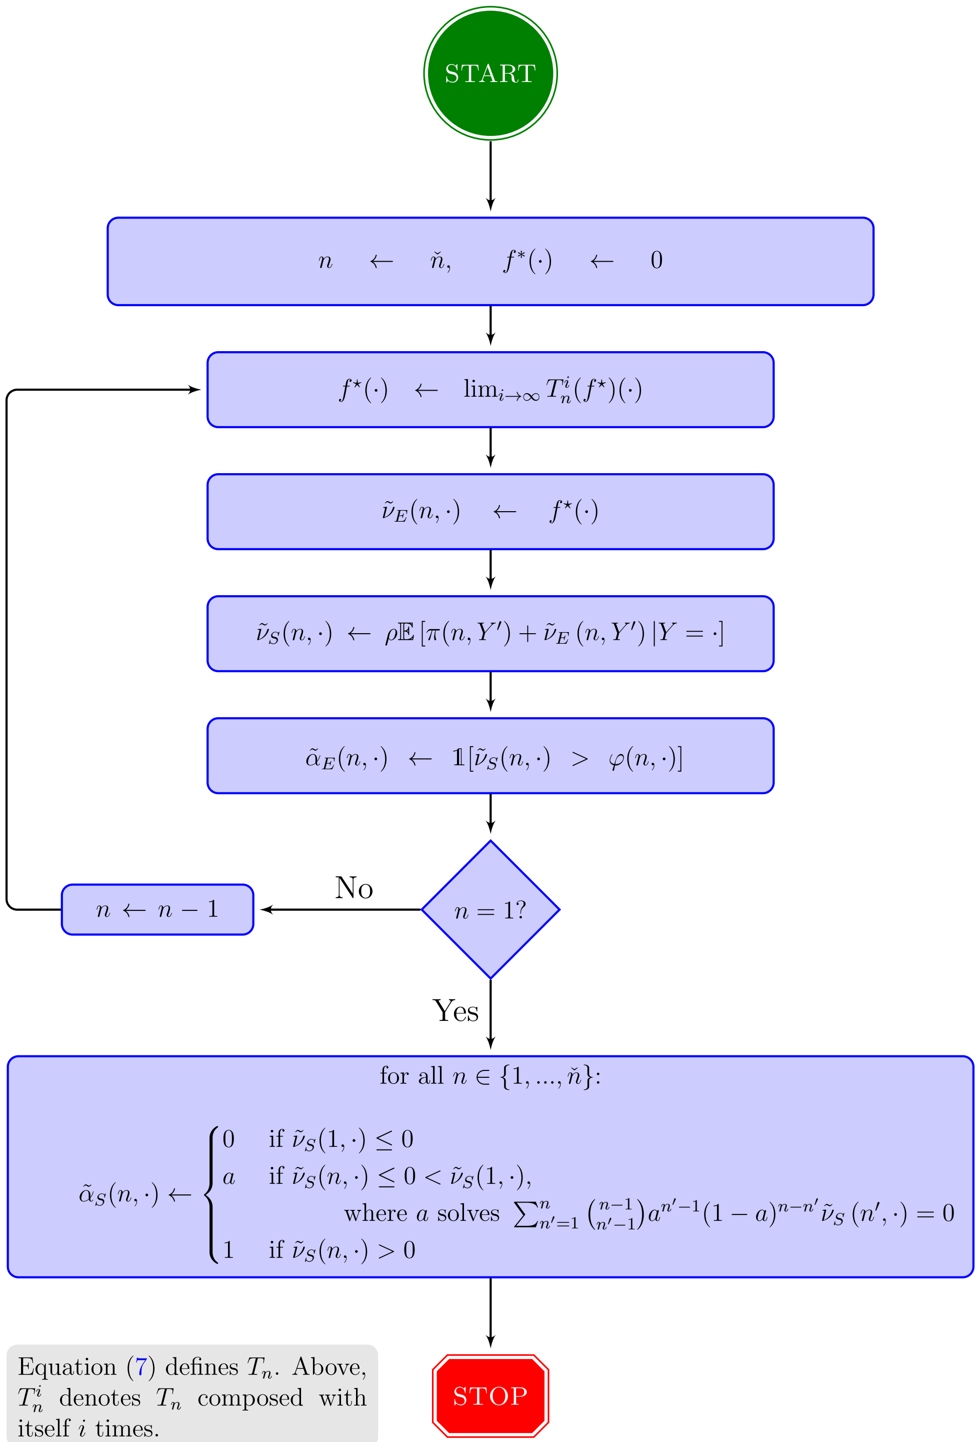

Supplementary Algorithm 1: Equilibrium Calculation for the Alternative Model

$\cdots \geq \tilde{\alpha}_E(\check{n}, \cdot)$. With Assumption A3 this monotonicity implies that

$$\tilde{\nu}_S(n', \cdot) = \rho\mathbb{E}\left[\pi(n', Y') + \tilde{\nu}_E\left(n', Y'\right)|Y = \cdot\right]$$

weakly decreases with n'. This is the desired monotonicity result.

For the second part, first note that the algorithm always sets $\tilde{\alpha}_E(n+1, y) = \mathbb{1}\{\tilde{\nu}_S(n+1, y) > \varphi(n+1, y)\}$. Next, consider the requirements of (4) and (5). For states (n, y) such that $\tilde{\nu}_S(n, y) = \cdots = \tilde{\nu}_S(1, y) = 0$, the monotonicity of $\tilde{\nu}_S(n^\star, y)$ in $n^\star$ ensures that $\tilde{\nu}_S(n+1, y) \leq 0$. Therefore, (4) requires $\tilde{\alpha}_E(n+1, y) = 0$, which is indeed the case. Given this, (5) imposes only the trivial requirement that $\tilde{\alpha}_S(n, y) \in [0, 1]$. Supplementary Algorithm 1's selection of $\tilde{\alpha}_S(n, y) = 0$ satisfies this. If instead $\tilde{\nu}_S(n+1, y) = \varphi(n+1, y)$, then the monotonicity of $\nu_S(n^\star, y)$ in $n^\star$ implies that $\nu_S(n^\star, y) > 0$ for $n^\star = 1, \ldots, n$. Supplementary Algorithm 1 sets $\alpha_E(n+1, y) = 0$ and $\alpha_S(n, y) = 1$ for these states, which satisfies both (4) and (5). For all other states (n, y), Supplementary Algorithm 1 sets $\tilde{\alpha}_E(n+1, y)$ and $\tilde{\alpha}_S(n, y)$ to the Nash equilibrium strategies from Corollary $1^\star$ game with n incumbents and payoffs $\tilde{\nu}_S(n', y)$ from continuation with $n' = 1, \ldots, n+1$ firms, which satisfies both (4) and (5). Equation (2) requires $\tilde{\nu}_E(n, y)$ to equal the expected payoff from this game to the potential entrant, which is true by construction. Similarly, Supplementary Algorithm 1 sets $\tilde{\nu}_S(n, y)$ so that it and $\tilde{\nu}_E(n, y)$ satisfy Equation (3) automatically. We conclude that Supplementary Algorithm 1's candidate strategy indeed forms an equilibrium.

The remainder of this proof demonstrates equilibrium uniqueness. Corollary $2^\star$ implies that any equilibrium $\tilde{v}_E(\check{n}, \cdot)$ equals the unique fixed point of $T_{\check{n}}$, $\tilde{\nu}_E(\check{n}, \cdot)$. Given this result, Equation (3) implies that $\tilde{v}_S(\check{n}, \cdot) = \tilde{\nu}_S(\check{n}, \cdot)$ in any equilibrium. In turn, this result and Lemma $2^\star$ together imply that $\tilde{v}_S(1, \cdot) \geq \tilde{v}_S(2, \cdot) \geq \cdots \geq \tilde{v}_S(\check{n}-1, \cdot) \geq \tilde{\nu}_S(\check{n}, \cdot)$. So when $\tilde{\nu}_S(\check{n}, y) > \varphi(\check{n}, y)$, a potential entrant's unique payoff-maximizing choice is sure entry. If instead $\tilde{\nu}_S(\check{n}, y) = \varphi(\check{n}, y)$, then the restriction that the equilibrium strategy defaults to inactivity requires that the potential entrant not enter for sure. Finally, if $\tilde{\nu}_S(\check{n}, y) < \varphi(\check{n}, y)$, then the restriction that the equilibrium is natural also requires the potential entrant not to enter for sure. Supplementary Algorithm 1's setting of $\tilde{\alpha}_E(\check{n}, y) = \mathbb{1}\left[\tilde{\nu}_S(\check{n}, y) > \varphi(\check{n}, y)\right]$ is the only choice for this entry rule that satisfies these restrictions.

Next, repeat the following argument for n decreasing from $\check{n}-1$ to 1. For given n, suppose that we have determined that $\tilde{v}_E(n^\star,\cdot) = \tilde{\nu}_E(n^\star,\cdot)$ and $\tilde{a}_E(n^\star,\cdot) = \tilde{\alpha}_E(n^\star,\cdot)$ for $n^\star = n+1,\ldots,\check{n}$. Corollary $2^\star$ implies that any equilibrium $\tilde{v}_E(n,\cdot)$ equals the unique fixed point of T_n, $\tilde{\nu}_E(n,\cdot)$. Given this result, Equation (3) implies that $\tilde{v}_S(n,\cdot) = \tilde{\nu}_S(n,\cdot)$ in any equilibrium. In turn, this result and Lemma $2^\star$ together imply that if $n > 1$, then $\tilde{v}_S(1,\cdot) \geq \cdots \geq \tilde{\nu}_S(n,\cdot)$. So when $\tilde{\nu}_S(n,y) > \varphi(n,y)$, a potential entrant's unique payoff-maximizing choice is sure entry. If instead $\tilde{\nu}_S(n,y) = \varphi(n,y)$, then the restriction that the equilibrium strategy defaults to inactivity requires that the potential entrant not enter for sure. Finally, if $\tilde{\nu}_S(n,y) < \varphi(n,y)$, then the restriction that the equilibrium is natural also requires the potential entrant not to enter for sure. Supplementary Algorithm 1's setting of $\tilde{\alpha}_E(n,y) = \mathbb{1}\left[\tilde{\nu}_S(n,y) > \varphi(n,y)\right]$ is the only choice for this entry rule that satisfies these restrictions.

The completion of this recursion establishes that $\tilde{v}_E = \tilde{\nu}_E$, $\tilde{v}_S = \tilde{\nu}_S$, and $\tilde{a}_E = \tilde{\alpha}_E$ in any equilibrium. Applying Corollary $1^\star$ then determines that $\tilde{a}_S(n,y) = \tilde{\alpha}_S(n,y)$ for any (n,y) such that there exists an $n' \in \{1,\ldots,n\}$ with $\tilde{\nu}_S(n',y) \neq 0$ and $\tilde{v}_S(n+1,y) \neq \varphi(n+1,y)$. States that do not satisfy its preconditions fall into two cases. If $\tilde{\nu}_S(1,y) = \tilde{\nu}_S(2,y)\cdots = \tilde{\nu}_S(n,y) = 0$, then the monotonicity of $\tilde{\nu}_S(n^\star,y)$ in $n^\star$ guarantees that $\tilde{\nu}_S(n+1,y) \leq 0$. Supplementary Algorithm 1 sets $\tilde{\alpha}_S(n,y) = 0$ for these states, which the assumption that the equilibrium strategy defaults to inactivity requires. If instead $\tilde{\nu}_S(n+1,y) = \varphi(n+1,y)$, then monotonicity of $\tilde{\nu}_S(n^\star,y)$ in $n^\star$ guarantees that $\tilde{\nu}_S(1,y) \geq \tilde{\nu}_S(2,y) \geq \cdots \geq \tilde{\nu}_S(n+1,y) > 0$. Given these values from survival, sure continuation is a dominant strategy for any incumbent. Supplementary Algorithm 1 indeed sets $\tilde{\alpha}_S(n,y)$ to one, as required. We conclude that the equilibrium strategy computed by Supplementary Algorithm 1 is the only natural Markov-perfect equilibrium strategy that defaults to inactivity.
■

References

ABBRING, J. H., J. R. CAMPBELL, J. TILLY, AND N. YANG (2017a): "Very Simple Markov-Perfect Industry Dynamics: Empirics," Discussion Paper 2017-021, CentER, Tilburg.

——— (2017b): "Very Simple Markov-Perfect Industry Dynamics: Theory," *Econometrica*, forthcoming.

FUDENBERG, D. AND J. TIROLE (1991): *Game Theory*, Cambridge: MIT Press.