Last-In First-Out Oligopoly Dynamics*

Jaap H. Abbring† Jeffrey R. Campbell‡

September 25, 2006

Abstract

This paper extends the static analysis of oligopoly structure into an infinite-horizon setting with sunk costs and demand uncertainty. The observation that exit rates decline with firm age motivates the assumption of last-in first-out dynamics: An entrant expects to exit *before* any incumbent firms. This selects a unique Markov-perfect equilibrium. With mild restrictions on the demand shocks, a sequence of demand thresholds describes firms' equilibrium entry and survival decisions. Bresnahan and Reiss's (1993) empirical analysis of oligopolists' entry and exit assumes that such thresholds govern the evolution of the number of competitors. Our analysis provides an infinite-horizon game-theoretic foundation for that structure.

*We are grateful to Meredith Crowley and Allan Collard-Wexler for their insightful comments on earlier drafts of this paper. The National Science Foundation supported this research through Grant 0137042 to the NBER.

†Vrije Universiteit Amsterdam and Tinbergen Institute. E-mail:jabbring@feweb.vu.nl

‡Federal Reserve Bank of Chicago and NBER. E-mail: jcampbell@frbchi.org

JEL Code L13

Keywords: Sunk costs, Demand uncertainty, Markov-Perfect equilibrium

1 Introduction

This paper develops and presents a simple and tractable model of oligopoly dynamics. The model's firms make entry and exit decisions in an infinite-horizon setting with stochastic demand. Calculating its equilibrium dynamics requires only a few *seconds* on a standard personal computer. With mild restrictions on the demand shocks, *threshold* rules govern firms' entry and exit decisions. That is, entry occurs whenever demand passes above one in a sequence of entry thresholds, and exit occurs if it subsequently passes below a corresponding exit threshold. Bresnahan and Reiss's (1993) empirical analysis of oligopolists' entry and exit assumes that such thresholds govern the evolution of the number of competitors. Our analysis provides an infinite-horizon game-theoretic foundation for that structure, so it can be used to extend their earlier structural estimation of static oligopoly models to a fully dynamic setting.

The model industry is a dynamic version of the static entry game used by Bresnahan and Reiss (1990). A random number of customers demands the industry's services, and this state evolves stochastically. Entry possibly requires paying a sunk cost, and continued operation incurs fixed costs. The wish to avoid these per-period fixed costs in markets that are no longer profitable motivates firms to exit.

Bresnahan and Reiss (1991a) noted that the static version of this game can have multiple equilibria, which obviously complicates prediction. To select a unique equilibrium, both Bresnahan and Reiss (1990) and Berry (1992) assume that firms' move sequentially. We take a similar approach by allowing older firms to commit to continuation before their younger counterparts. We also restrict attention to equilibria in which firms correctly believe that firms exit always in the reverse order of their entry. That is, the equilibria have a last-in first-out structure. Three considerations motivate this focus. First, it is consistent with the widespread observation that young firms exit more frequently than their older counterparts. Second, the equilibrium approximates the "natural" Markov-perfect equilibrium in an extension of the model in which firm's costs decrease with age and the most efficient firms survive. Third and perhaps most importantly, this restriction vastly simplifies the equilibrium analysis. We prove that there always exists such an equilibrium and that it is (essentially) unique.

The model's theoretical simplicity makes it well-suited for exploring how parameter changes impact equilibrium dynamics and long-run market structure. To illustrate this, we calculate the effects of increasing demand uncertainty on firms' equilibrium entry and exit thresholds. Non-strategic analysis of the firm life cycle suggests that additional uncertainty should raise the value of the option to exit and thereby substantially lower both entry and exit thresholds. The oligopolistic exit thresholds do indeed fall with uncertainty, but the entry thresholds do not. Their relative invariance to demand uncertainty reflects an offsetting effect which a monopolist does not face: Increasing demand uncertainty raises the probability of further entry and thereby reduces a new firm's value. We also calculate the population "estimates" of oligopoly profit margins using the ordered Probit procedure of Bresnahan and Reiss (1990). We find that the entry threshold estimates from this static estimation procedure are biased downwards, and this bias leads to a downward bias in the estimated rate that profits fall with additional competition. That is, the static procedure can find "evidence" that profit margins decline with entry when in fact they are constant.

The remainder of this paper proceeds as follows. The next section presents the model's primitives and demonstrates the uniqueness of a Markov-perfect equilibrium with a last-in-first-out structure. It closes with an examination of a particular specification for the model's demand shocks that yields a pencil-and-paper solution. Section 3 gives sufficient conditions for firms to use threshold rules for their equilibrium entry and exit decisions, and Section 4 illustrates its application. Section 5 considers extensions of our analysis to include a learning curve and firm-specific technology shocks. Section 6 relates the our work with previous analyses of dynamic games with timing restrictions and with the extensive literature on oligopoly with Markov-perfect equilibrium, and Section 7 contains some concluding remarks.

2 The Model

The model consists of a single oligopolistic market in discrete time $t \in \{0, 1, \ldots\}$. There is a large number of firms which are potentially active in the market. We index these firms by $\mathbb{N}$. At time 0, $N_0 = 0$ firms are active. Entry and subsequent exit of firms determines the number of active firms in each later period t, N_t . The number of consumers in the market, C_t, evolves exogenously according to a nonnegative first-order Markov process bounded between $\hat{C}$ and $\check{C}$. We denote the conditional distribution of C_t with $Q(c|C_{t-1}) \equiv \Pr[C_t \leq c|C_{t-1}]$.

Figure 1 illustrates the sequence of events and actions within a period. It begins with the inherited values of N_t and C_{t-1}. First, all participants observe the realization of C_t; and all active firms receive profits equal to $(C_t/N_t) \times \pi(N_t) - \kappa$. Here, each firm serves C_t/N_t customers, and $\pi(N_t)$ is the producer surplus earned from each one. The term $\kappa \geq 0$ represents fixed costs of production.

After serving the market, active firms decide whether they will remain so. These decisions are sequential and begin with the oldest firm. After this, any remaining firms make the same

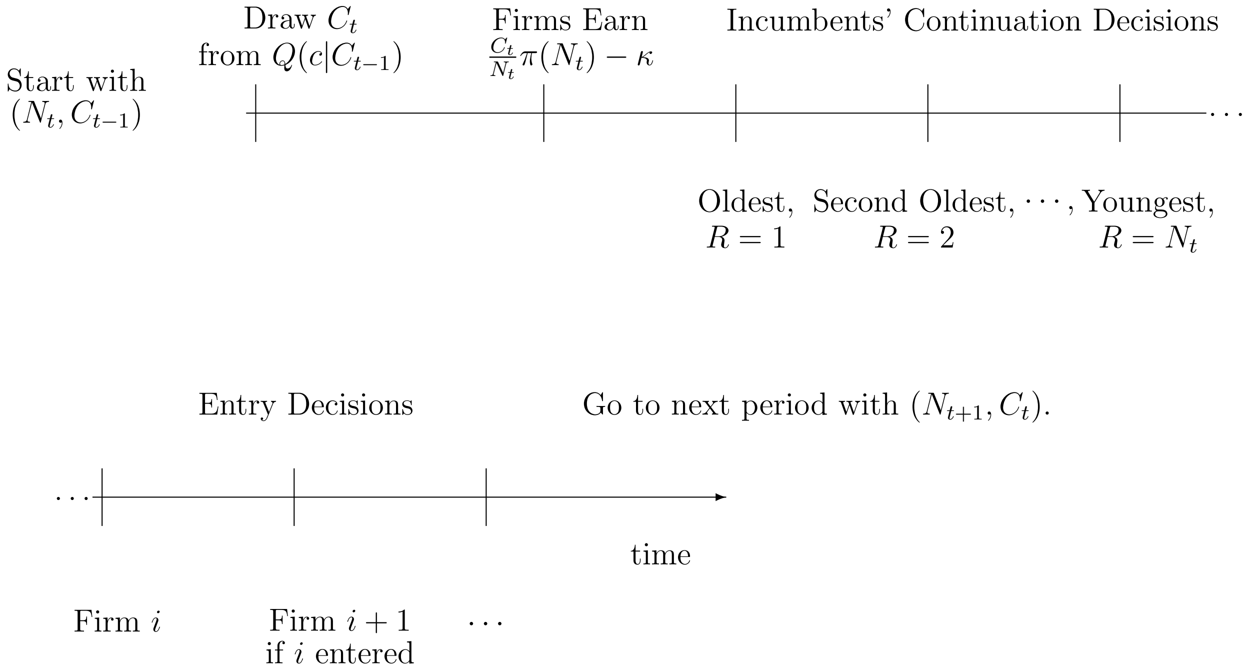

Figure 1: The Sequence of Actions within a Period

decision in the order of their entry. If firm i is active, then R_t^i denotes its rank in this sequence. Exit is costless but irreversible and allows the firm to avoid future periods' fixed production costs.

After active firms' continuation decisions, those firms that have not yet had an opportunity to enter make entry decisions in the order of their names. The cost of entry potentially depends on the number of firms already committed to serving the market in the next period. We denote this cost with $\varphi(N)$, and we assume that it is weakly increasing in N. This allows for, but does not require, later entrants to face a "barrier to entry" in the form of elevated sunk costs. The payoff to staying out of the industry is always zero.[1] The period ends when some potential entrant decides to stay out of the industry. Both active firms' and potential entrants' decisions maximize their expected stream of profits discounted with a factor β.

[1] Option value considerations will not impact the entry decision, because a firm with an entry opportunity cannot delay its choice.

2.1 Markov-perfect equilibrium

We choose as our equilibrium concept symmetric Markov-perfect equilibrium. When firm i decides whether to stay or exit, C_t, $N_t - R_t^i$ (the number of active firms following it in the sequence) and R_{t+1}^i (its rank in the next period's sequence of active firms) are available and payoff-relevant. Collect these into $H_{it} \equiv (N_t - R_t^i, C_t, R_{t+1}^i)$. Similarly, the payoff-relevant state to a potential entrant is $H_{it} \equiv (C_t, R_{t+1}^i)$. Note that H_{it} takes its values in $\mathcal{H}_S \equiv \mathbb{N} \times \left[\hat{C}, \check{C}\right] \times \mathbb{N}$ for firms i active in period t and in $\mathcal{H}_E \equiv \left[\hat{C}, \check{C}\right] \times \mathbb{N}$ for potential entrants. Here and below, we use S and E to denote survivors and entrants.

A Markov strategy for firm i is a pair $(A_S^i(H_S), A_E^i(H_E))$ for each $H_S \in \mathcal{H}_S$ and $H_E \in \mathcal{H}_E$. These represent the probability of being active in the next period given that the firm is currently active $(A_S^i(\cdot))$ and given that the firm has an entry opportunity $(A_E^i(\cdot))$. A symmetric Markov-perfect equilibrium is a subgame-perfect equilibrium in which all firms follow the same Markov strategy.

When firms use Markov strategies, the payoff-relevant state variables determine their expected discounted profits. We denote these values with $v_S(H_S)$ and $v_E(H_E)$. In Markov-perfect equilibrium, these satisfy the Bellman equations

$$v_S(H_S) = \max_{a \in [0,1]} a\beta\mathbb{E}\left[\frac{C'}{N'}\pi(N') - \kappa + v_S(H_S') \,\middle|\, H_S\right], \tag{1}$$

and

$$v_E(H_E) = \max_{a \in [0,1]} a\beta\mathbb{E}\left[\frac{C'}{N'}\pi(N') - \kappa + v_E(H_E') \,\middle|\, H_E\right] - a\varphi(N(H_E) + 1). \tag{2}$$

Here and throughout, we adopt conventional notation and denote the variable corresponding to X in the next period with X'. In Equations (1) and (2), the expectation of N' is calculated using all firms' strategies conditional on the particular firm of interest choosing to be active. In equation (2), $N(H_E)$ is the number of firms already committed to serving the market in the next period when the state equals H_E.

Examining the effects of changing the model primitives requires an equilibrium concept specific enough to determine a unique outcome, and it is well known that multiple Markov-perfect equilibria might exist. To overcome this standard difficulty, we restrict attention to equilibria in which firms' entry and exit policies obey a Last In/First Out (LIFO) rule.

Definition 1. *A LIFO strategy is a strategy (A_S, A_E) in which $A_S(H_S) \in \{0, 1\}$, $A_E(H_E) \in \{0, 1\}$, and $A_S(N - R, C, R')$ is weakly decreasing in R'.*

If all firms adopt a common LIFO strategy, then they exit in the reverse order of their entry. As we mentioned in the paper's introduction, this embodies in an extreme form the empirical regularity that young firms exit more frequently than their older counterparts.

With this definition, we can demonstrate existence of a Markov-perfect equilibrium in a LIFO strategy.

Proposition 1. *There exists a symmetric Markov-perfect equilibrium in a LIFO strategy.*

This paper's appendix contains the proposition's constructive proof, which has two critical steps. First, we note that the upper bound on C implies that the number of firms that ever produce in a subgame-perfect equilibrium cannot exceed some bound, which we call $\check{N}$. Because this firm expects to outlive none of its older competitors, this firm's optimal exit rule corresponds to that from a simple dynamic programming problem. Second, we solve exit decision problems for firms with ranks $\check{N}, \check{N}-1, \ldots, 1$, which embody the assumption that other firms follow a LIFO strategy. In these, a firm with rank R forms its expectations about the behavior of firms with higher ranks using the solutions of those firms' decision problems. With the solutions to these standard dynamic programming problems in hand, we construct a LIFO strategy and then verify that it indeed forms a Markov-perfect equilibrium.

The constructive existence proof strongly suggests that the Markov-perfect equilibrium in a LIFO strategy is unique, because the decision problems used in its construction have unique solutions to their Bellman equations. However, we can construct multiple LIFO equilibria by varying a firm's actions in states of indifference between activity and inactivity. We sidestep this difficulty by concentrating on equilibria in which a firm defaults to inactivity.

Definition 2. *A symmetric Markov-perfect equilibrium strategy* (A_S, A_E) defaults to inactivity *if* $A_S(H_S) = 0$ *whenever* $v_S(H_S) = 0$ *and* $A_E(H_E) = 0$ *whenever* $v_E(H_E) = \phi(N(H_E)+1)$.

Proposition 2. *There exists a unique symmetric Markov-perfect equilibrium in a LIFO strategy that defaults to inactivity.*

Other symmetric Markov-perfect equilibria which default to inactivity might exist, but in these the apparent advantage of early entrants to commit to continuation does not translate into longevity. Henceforth, we constrain our attention to the unique symmetric Markov-perfect equilibrium in a LIFO strategy that defaults to inactivity.

2.2 A Pencil-and-Paper Example

If we assume that $C_t = C_{t-1}$ with probability $1-\lambda$ and that it equals a draw from a uniform distribution on $[\hat{C}, \check{C}]$ with the complementary probability, then we can calculate the model's equilibrium value functions and decision rules with pencil and paper. Before proceeding, we examine this special case to illustrate the model's moving parts. For further simplification, suppose that $\pi(N) = 0$ for $N \geq 3$, so at most two firms serve the industry. To ensure that the equilibrium dynamics are not trivial, we also assume that no firm will serve the industry if demand is low enough and that two firms will serve the industry if it is sufficiently high.[2]

To begin, consider an incumbent firm with rank 2. In an equilibrium in a LIFO strategy, its profit equals $(C/2)\pi(2) - \kappa$. It will earn this until the next time that C_t changes, at which point the new demand value will be statistically independent of its current value. It is straightforward to use these facts to show that this firm's value function is the following piecewise linear function of C.

$$v_S(0, C, 2) = \begin{cases} \beta \frac{(1-\lambda)\left(\frac{C}{2}\pi(2)-\kappa\right)+\lambda\widetilde{v}(0,2)}{1-\beta(1-\lambda)} & \text{if } C > \underline{C}_2 \\ 0 & \text{otherwise,} \end{cases}$$

where

$$\widetilde{v}(0, 2) = \frac{1}{2}\left(\frac{\hat{C} + \check{C}}{2}\right)\pi(2) - \kappa + \int_{\hat{C}}^{\check{C}} \frac{v_S(0, C', 2)}{(\check{C} - \hat{C})} dC'.$$

Here, $\widetilde{v}(0, 2)$ is the firm's average continuation value given a new draw of C_t and $\underline{C}_2$ is the largest value of C that satisfies $v_S(0, C, 2) = 0$. Optimality requires the firm to exit if $C < \underline{C}_2$. This value function is monotone in C, so there is a unique entry threshold $\overline{C}_2$ which equates the continuation value with the entry cost. Thus, a second duopolist enters whenever C_t exceeds $\overline{C}_2$ and exits if it subsequently falls at or below $\underline{C}_2$.

Next, consider the problem of an incumbent with rank 1. If this firm is currently a monopolist, it expects to remain so until $C_t > \overline{C}_2$; and if it is currently a duopolist, it expects to become a monopolist when C_t falls below $\underline{C}_2$. This firm's value function is also piecewise linear. If the firm begins the period as the sole incumbent, it is

$$v_S(0, C, 1) = \begin{cases} \beta \frac{(1-\lambda)(C\pi(1)-\kappa)+\lambda\widetilde{v}(0,1)}{1-\beta(1-\lambda)} & \text{if } \underline{C}_1 < C \leq \overline{C}_2, \\ \beta \frac{(1-\lambda)\left(\frac{C}{2}\pi(2)-\kappa\right)+\lambda\widetilde{v}(1,1)}{1-\beta(1-\lambda)} & \text{if } C > \overline{C}_2, \\ 0 & \text{otherwise;} \end{cases}$$

[2] Sufficient conditions for these two properties are, respectively, $(1-\lambda)\left(\hat{C}\pi(1) - \kappa\right) + \lambda \frac{\frac{\hat{C}+\check{C}}{2}\pi(1)-\kappa}{1-\beta} < 0$ and $\beta \frac{\frac{\check{C}}{2}\pi(2)-\kappa}{1-\beta(1-\lambda)} > \varphi(2)$.

and if it begins as one of two incumbents it is

$$v_S(1,C,1) = \begin{cases} \beta\frac{(1-\lambda)(C\pi(1)-\kappa)+\lambda\widetilde{v}(0,1)}{1-\beta(1-\lambda)} & \text{if } \underline{C}_1 < C \leq \underline{C}_2, \\ \beta\frac{(1-\lambda)\left(\frac{C}{2}\pi(2)-\kappa\right)+\lambda\widetilde{v}(1,1)}{1-\beta(1-\lambda)} & \text{if } C > \underline{C}_2, \\ 0 & \text{otherwise.} \end{cases}$$

The exit threshold $\underline{C}_1$ is the greatest value of C such that $v_S(0,C,1) = 0$, and the average continuation values following a change in C_t for a monopolist and a duopolist are

$$\begin{aligned} \widetilde{v}(0,1) &= \left(\frac{\hat{C}+\check{C}}{2}\right)\pi(1) - \kappa + \int_{\hat{C}}^{\check{C}} \frac{v_S(0,C',1)}{(\check{C}-\hat{C})} dC', \\ \widetilde{v}(1,1) &= \frac{1}{2}\left(\frac{\hat{C}+\check{C}}{2}\right)\pi(2) - \kappa + \int_{\hat{C}}^{\check{C}} \frac{v_S(1,C',1)}{(\check{C}-\hat{C})} dC'. \end{aligned}$$

This value function does not always increase with C, because slightly raising C from $\overline{C}_2$ induces entry by the second firm and causes both current profits and the continuation value to discretely drop. Nevertheless, we know that they drop to a value above $\varphi(1)$, because at this point the second firm chooses to enter. Hence, it is still possible to find a unique entry threshold $\overline{C}_1$ which equates the value of entering with rank 1 to the cost of doing so.

Figure 2 visually represents the example industry's equilibrium. In each panel, the demand state C runs along the horizontal axis, while the vertical axis gives the value of a firm at the time that entry and exit decisions are made. The top panel plots the value of a firm with rank 1, while the bottom plots the value for a competitor with rank 2. For visual clarity, the two panels have different vertical scales.

Consider first the bottom panel. The value of a duopolist with rank 2 equals zero for $C < \underline{C}_2$, and thereafter increases linearly with C. The entry threshold $\overline{C}_2$ equates this value with $\varphi(2)$. The value of an older firm with rank 1 has two branches. The upper monopoly branch gives the value of a monopolist expecting no further entry. If C increases above $\overline{C}_2$ and thus induces a entry, the firm's value drops to the lower duopoly branch. This has the same slope as the value function in the lower panel. Its intercept is higher, because the incumbent expects to eventually become a monopolist the first time that C passes below $\underline{C}_2$.[3] When this occurs, the firm's value returns to the monopoly branch. The entry and exit thresholds for this firm occur where the monopoly branch intersects $\varphi(1)$ and 0.

[3]The two panels' different vertical scales mask these results.

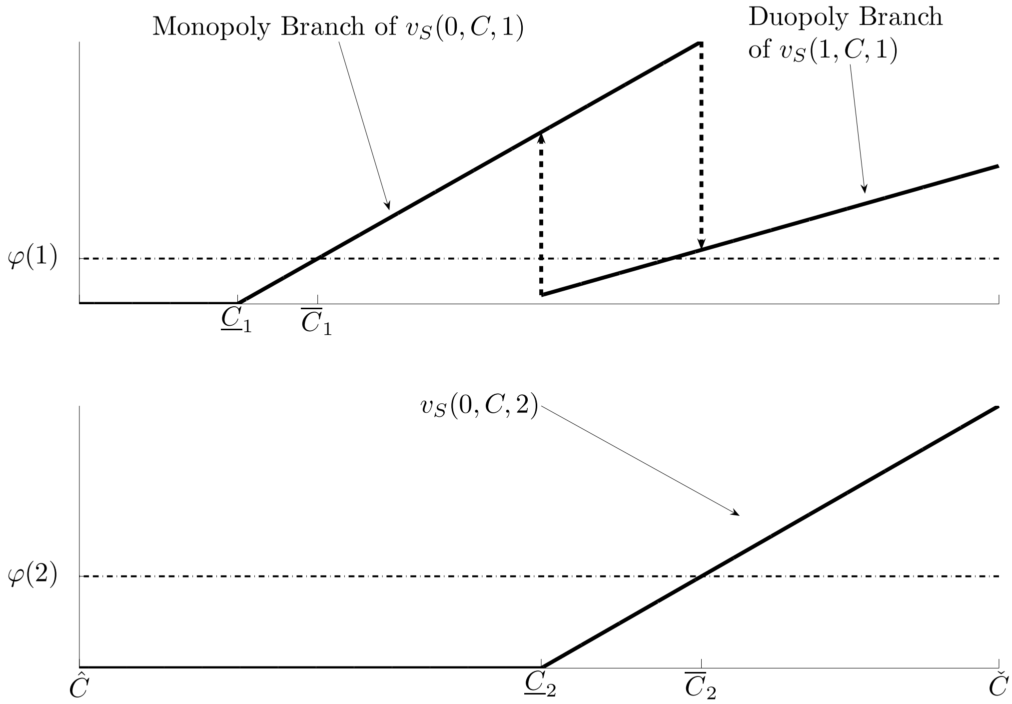

Figure 2: Equilibrium in the Pencil-and-Paper Example

3 Threshold Entry and Exit Rules

In the paper-and-pencil example, all firms follow *threshold rules* for their entry and continuation decisions.

Definition 3. *A firm with rank R' follows a* threshold rule *if there exist real numbers $\underline{C}_{R'}$ and $\overline{C}_{R'} \geq \underline{C}_{R'}$ such that $A_S(N - R, C, R') = I\{C > \underline{C}_{R'}\}$ and $A_E(C, R') = I\{C > \overline{C}_{R'}\}$.*

With such a rule, a potential entrant into a market with $R' - 1$ incumbents actually enters if and only if $C > \overline{C}_{R'}$, and this firm exits the first time that $C \leq \underline{C}_{R'}$.

There are three reasons to care about whether or not firms follow threshold rules. First, they pervade theoretical and empirical industrial organization. Second, they simplify the model's analysis, as the pencil-and-paper example illustrated. Third, as the next proposition shows, higher realizations of demand always result in more active firms if and only if all firms use threshold rules.

Proposition 3. *Consider a sequence of possible demand realizations, $C_1, C_2, \ldots$ and the* *corresponding number of operating firms from the equilibrium of Proposition 2,* $N_1, N_2, \ldots$. *Then increasing* C_t *weakly increases* N_{t+j} *for positive* t, *non-negative* j, *and* any *possible sequence of demand realizations if and only if firms of all ranks follow threshold rules.*

This proposition's proof is obvious. A monotonic influence of C_t on N_{t+j} appeals to us as "natural". It is straightforward to show that a firm with the highest possible rank always follows a threshold rule, but the following example illustrates that firms with lower ranks might not without further restrictions.

3.1 A Non-Monotonic Exit Rule

Consider the problem of an incumbent monopolist in a market with demand state C_1. With probability 1/2, this increases by Δ; and with the same probability it decreases the same amount. The demand state remains at its period 2 value thereafter.

The market can support two firms at most. A second firm enters at the end of period 2 if demand exceeds $\overline{C}_2$, which equates the value of perpetual duopoly profits with entry's cost.

$$\frac{\overline{C}_2}{2}\pi(2) - \kappa = \frac{1-\beta}{\beta}\varphi(2)$$

Similarly, an incumbent monopolist chooses to remain in production after period 2 if demand exceeds $\underline{C}_1$, given by

$$\underline{C}_1\pi(1) - \kappa = 0$$

Next, consider the incumbent's continuation decision at the end of period 1 if there is no threat of further entry. Assume that $C^\star - \Delta < \underline{C}_1$, so that the incumbent operates after period 1 only if the positive demand shock occurs. Both available options return the same expected profit if $C = C^\star$, where

$$\beta\left(C^\star\pi(1) - \kappa + \frac{1}{2}\frac{(C^\star + \Delta)\pi(1) - \kappa}{1-\beta}\right) = 0$$

The first two terms are the monopolist's expected profit in the second period, while the second term is the expected profit from all other periods.[4]

To construct an example of a non-monotonic decision rule, adjust $\pi(2)$ so that $C^\star + \Delta = \overline{C}_2 - \varepsilon$, where ε slightly exceeds zero. Clearly, it is optimal for the incumbent monopolist to exit immediately if $C \leq C^\star$, and continuation is optimal if C exceeds $C^\star$ and is less than

[4] By construction, $C^\star > \underline{C}_1$.

$C^\star + \varepsilon$. If C exceeds $C^\star + \varepsilon$, the incumbent receives *duopoly* profits following a positive demand shock. The incumbent's payoff to continuation in this case is

$$\beta\left(C\pi(1) - \kappa + \frac{1}{2}\beta\frac{\frac{C+\Delta}{2}\pi(2) - \kappa}{1-\beta}\right).$$

If C exceeds $C^\star + \varepsilon$ only slightly, then the fall in profits due to entry greatly exceeds the gain from the extra demand. This payoff is *less* than the continuation payoff when $C = C^\star$, so the incumbent's optimal decision is to exit if C slightly exceeds $C^\star + \varepsilon$.

In summary, increasing C from below $C^\star$ to between $C^\star$ and $C^\star + \varepsilon$ induces the incumbent to switch from exit to continuation, and further increasing C past $C^\star + \varepsilon$ induces the incumbent to exit again. In this sense, the incumbent's exit rule is non-monotonic, and there exists no equivalent threshold rule.

3.2 Demand Processes that Generate Threshold Rules

The above example illustrates that firms do not generically use threshold rules in equilibrium. In it, increasing the current value of C increases the likelihood of crossing $\overline{C}_2$, which would decrease the incumbent's profit. In contrast, increasing C in the pencil-and-paper example leaves the probability of future entry unchanged. Together, these examples suggest that firms will use threshold rules if the stochastic process limits the negative "potential entry" effect of increasing C on expected future profits. Here we present sufficient conditions for this to be so. We rely on the following class of stochastic processes for C_t.

Definition 4. *The transition function $Q(\cdot|C)$ is a* mixture of uniform autoregressions with bounded growth *if (i) there exists a sequence of transition functions*

$$Q_k(c|C) = \begin{cases} 1 & \text{if } c > \mu_k(C) + \sigma_k/2 \\ (c - \mu_k(C) + \sigma_k/2)/\sigma_k & \text{if } \mu_k(C) - \sigma_k/2 \leq c \leq \mu_k(C) + \sigma_k/2 \\ 0 & \text{otherwise,} \end{cases}$$

with both $\mu_k(C) \leq C + \sigma_k/2$ and $\mu_k(C)$ weakly increasing in C; and (ii) there exists a sequence of positive real numbers p_k such that $\lim_{K\to\infty} \sum_{k=1}^{K} p_k = 1$ and

$$\lim_{K\to\infty} \sup_{c,C} \left| Q(c|C) - \sum_{k=1}^{K} p_k Q_k(c|C) \right| = 0.$$

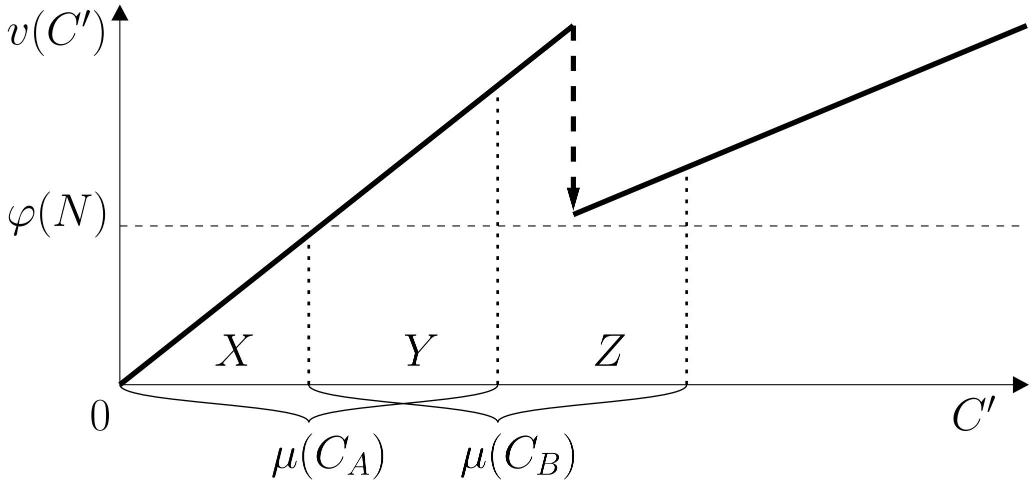

$$C' \sim U\left[\mu(C) - \sigma/2, \mu(C) + \sigma/2\right]$$
$$\mathbb{E}\left[v(C')|C = C_A\right] \geq \varphi(N) \Rightarrow \mathbb{E}\left[v(C')|C = C_B\right] \geq \varphi(N)$$

Figure 3: Increasing C from C_A to C_B Does Not Lower the Continuation Value

In this definition, each of the mixing distributions is a (possibly nonlinear) autoregression with conditional mean $\mu_k(C)$ and uniform innovations with variance $\sigma_k^2/12$. The coefficients p_k give the mixing probabilities. The condition that $\mu_k(C) \leq C + \sigma_k/2$ ensures that the current state is always in or above the support of each mixing distribution. This is the sense in which Definition 4 bounds the growth of C. With this definition in place, we can state this section's central result.

Proposition 4. *Let (A_S, A_E) be the unique symmetric Markov-perfect equilibrium in a LIFO strategy that defaults to inactivity. Assume that $Q(\cdot|C)$ is a mixture of uniform autoregressions with bounded growth. Then, firms with all ranks follow threshold policies.*

The key step in the proposition's proof is the demonstration that an increase in C that makes further entry more likely does not reduce the expected continuation value below the firm's cost of entry. To appreciate the contribution of the restriction on $Q(\cdot|C)$ to this, consider Figure 3. This plots a possible value function for a firm which would have rank

N with C' on the horizontal axis. As C' crosses a value that induces yet further entry, the function drops to a value *above* the cost of entry for this firm. We know that this will be so in equilibrium, because this firm's value will exceed that of any entrant; the entrant's continuation value must be above its entry cost; and $\varphi(N+1) \geq \varphi(N)$. The figure also marks the means and supports of two uniform distributions with identical variances but different means, which we have labelled $\mu(C_A) < \mu(C_B)$.

If $C = C_A$, then no realization of C' could induce further entry. Because the distribution of C' is uniform, the expected continuation value in this case equals the integral of the value function from $\mu(C_A) - \sigma/2$ to $\mu(C_A) + \sigma/2$ divided by σ. This is proportional to sum of the areas marked X and Y. Suppose that this continuation value exceeds $\varphi(N)$, so that this firm would choose to enter if $C = C_A$. Increasing C to C_B opens up the possibility of entry. The change in the continuation value from this increase is proportional to the integral of the function in the region marked Z, where it always exceeds $\varphi(N)$, minus the integral over the region X. The figure indicates that the function always lies below $\varphi(N)$ in region X. These two regions have equal width, so the increase in C from C_A to C_B does not change the firm's optimal entry decision.

A wide variety of demand processes are consistent with the requirements of Proposition 4. The stochastic process from the pencil-and-paper example satisfies the conditions with α and $1-\alpha$ serving as the mixing probabilities. In this case, one of the uniform distributions is degenerate at $\mu_1(C_t) = C_t$. To construct another example, consider a random walk reflected at $\hat{C}$ and $\check{C}$. That is, set

$$\mu(c) = \begin{cases} \hat{C} + \frac{\sigma}{2} & c < \hat{C} + \frac{\sigma}{2}, \\ c & \text{if} \quad \hat{C} + \frac{\sigma}{2} \leq c \leq \check{C} - \frac{\sigma}{2}, \text{ and} \\ \check{C} - \frac{\sigma}{2} & c > \check{C} - \frac{\sigma}{2}, \end{cases}$$

for some $0 < \sigma < \check{C} - \hat{C}$. By mixing over such reflected random walks, we can approximate any symmetric and continuous distribution for the growth rate of demand in the region away from the boundaries of $[\hat{C}, \check{C}]$.

4 Entry and Exit with Uncertainty

This section illustrates the application of our analysis by addressing two related questions: How does adding uncertainty impact oligopolist's entry and exit thresholds? How do estimates of oligopolists' profits per consumer based on static models without both uncertainty and sunk costs differ from their actual values?

A large literature characterizes competitive firms' entry and exit decisions with sunk costs and uncertain profits. Such a firm's value equals its *fundamental value*, the expected discounted profits from perpetual operation, plus the value of an option to sell this stream of (potentially negative) profits at a strike price of zero. The key insight of this literature is that uncertainty about future profits raises the value of this put option and thereby decreases the frequency of exit. Abbring and Campbell (2006b) estimated that this option value accounted for the majority of firm value in a particular competitive service-industry. Our model allows us to investigate how the insights from this well-studied decision theoretic problem apply to oligopolistic dynamics.

Our analysis of the second question follows a large literature based on static free-entry models of oligopoly structure, exemplified by Bresnahan and Reiss (1990, 1991b) and Berry (1992). They determined empirically how changing market size influenced the number of competitors using observations from cross-sections of local retail (Bresnahan and Reiss) and airline (Berry) markets. The models they used to structure their analysis can be viewed as versions of ours in which either demand remains unchanged over time or firms incur no sunk costs. These papers point to current demand as the key determinant of the number of firms: A market will attain N firms if N entrants can recover their fixed costs but $N+1$ entrants cannot. These authors emphasize that the observed relationship between C and N depends on the rate at which $\pi(N)$ decreases (which Sutton (1991) labelled the "toughness of competition") and the rate at which $\varphi(N)$ increases (which could be interpreted as a "barrier to entry"). If both of these functions are constant, then the number of active firms is roughly proportional to demand, $\overline{C}_j = j \times \overline{C}_1$. However, if either $\pi(N)$ decreases or $\varphi(N)$ increases, then N/C declines with C. In this sense, increasing the toughness of competition or imposing a sunk barrier to entry increases concentration.

Our approach to answering these questions is computational. Accordingly, we begin this section with an explicit presentation of the algorithm for equilibrium computation. We then document the impacts of demand uncertainty on equilibrium entry and exit thresholds for a particular model parameterization. The section concludes with the calculation of the entry thresholds and profits per consumer calculated from feeding data generated by our model's ergodic distribution through a static Probit model of long-run equilibrium like that of Bresnahan and Reiss (1990, 1991b).

4.1 Equilibrium Computation

The proof of Proposition 1 outlines a simple algorithm for computing the Markov-perfect equilibrium of interest. Given values for the model's primitives, we choose an evenly spaced grid of values for C in the interval $[\hat{C}, \check{C}]$ and a Markov chain over this grid to approximate $Q(c|C)$. Following the straightforward calculation of $\check{N}$, consider a firm entering with rank $\check{N}$. This firm rationally expects no further entry, so we can solve this firm's dynamic programming problem by setting $N'(\check{N}, C) = \check{N}$, beginning with a trial value for the firm's value function, and iterating on the Bellman equation (1). In practice, this takes very, very little computer time.

The rest of the computation proceeds recursively. This solution to the dynamic programming problem for a firm with rank R produces continuation and entry sets

$$\mathcal{S}_R = \{C|v_R(C) > 0\} \qquad \text{and} \qquad \mathcal{E}_R = \{C|v_R(C) > \varphi(R)\},$$

which we can use to calculate the expected evolution of N for a firm with rank $R-1$. Define $N'_R(N, C)$ to be the expected value of N_{t+1} given (i) $N_t = N$, (ii) $C_t = C$ and (iii) the firm with rank R decides to continue. Then we can calculate $N'_R(N, C)$ from $N'_{R+1}(N, C)$, $\mathcal{S}_R$, and $\mathcal{E}_R$ with

$$N'_R(N, C) = \begin{cases} N'_{R+1}(N, C) - 1 + I\{C \in \mathcal{S}_{R+1}\} & \text{if } N \geq R + 1, \\ R + I\{C \in \mathcal{E}_{R+1}\}\left(N'_{R+1}(N, C) - R\right) & \text{if } N = R. \end{cases}$$

A firm with rank $R - 1$ solves the dynamic programming problem in equation (1) given this particular rule for the evolution of N_t. Bellman equation iteration immediately yields its solution. Continuing yields the equilibrium entry and exit rules for all $\check{N}$ possible ranks. With these in hand, calculating observable aspects of industry dynamics, such as the ergodic distribution of N_t, is straightforward.

4.2 Equilibrium Entry and Exit Thresholds

With this algorithm, we have calculated the equilibrium entry and exit thresholds for a particular specification of the model which satisfies the sufficient conditions for firms to use threshold-based entry and exit policies. We set one model period to equal one year and chose β to replicate a 5% annual rate of interest. We set $\kappa = 1.75$ and $\varphi = 0.25(1 - \beta)/\beta$, so the fixed costs of a continuing establishment equal seven times sunk costs' rental equivalent value. We also set $\pi(N) = 4$ for all N. With these parameter values and no demand uncertainty, the entry thresholds are twice the corresponding exit thresholds and the entry threshold for a second firm equals one. We set $\hat{C} = e^{-1.5}$, $\check{C} = e^{1.5}$, and $Q(c|C)$ to equal a mixture over 51 reflected random walks in the *logarithm* of C with uniformly distributed innovations. The mixture approximates a normally distributed innovation. We denote the standard deviation of the normal distribution we seek to approximate with σ. Proposition 4 can be easily extended to the case where Definition 4 applies to a monotone transformation of C_t, so the logarithmic specification for demand has no direct theoretical consequences. We choose it because population and income measures typically require a logarithmic transformation to display homoskedasticity across time.

The first two panels of Table 1 report the equilibrium entry and exit thresholds for this specification for four values of σ — 0, 0.10, 0.20, and 0.30. Given the support of C_t, up to eight firms could populate the industry when $\sigma = 0$. Because C_t is reflected off of $\check{C}$, demand displays mean reversion. Thus, such high states of demand are somewhat temporary when $\sigma > 0$, so the maximum number of firms observed in the ergodic distribution decreases with σ. The two panels' cells for those missing firms' entry and exit thresholds are blank.

Consider first the impact of increasing σ on the entry thresholds. At least one firm enters an empty market with no demand uncertainty if $C_t > 0.50$. This threshold hardly changes as σ increases. Likewise, the entry threshold for a second firm remains very close to 1.00 as σ rises. The thresholds for higher-ranked entrants all rise with σ with one exception (to be discussed further below). Apparently, increasing demand uncertainty makes entry into an oligopoly less likely for a given value of C_t. Demand uncertainty has exactly the opposite impact on the entry of a potential monopolist. For such a firm, increasing uncertainty increases the value of the put option associated with exit, thereby raising profitability and *lowering* the firm's entry threshold.

This difference between oligopolists' and monopolists' entry decisions arises from the threat of potential entry. A monopolist captures all of the increased profit from a favorable demand shock. For an oligopolist, further entry chops this right tail off of the profit distribution and thereby reduces the firm's option value. This explanation squares with the single exception to the rule that increasing σ increases the entry threshold. Increasing σ from 0.20 to 0.30 simultaneously eliminates the possibility that a sixth firm enters and reduces $\overline{C}_5$ from 2.72 to 2.56 . The third panel of Table 1 further illustrates this effect. It reports the equilibrium entry thresholds for the case where $\pi(N) = 4 \times I\{N < 5\}$, so no more than four firms will populate the market. The entry thresholds for the first, second, and third firms are nearly identical to their values in the first panel. However, the entry thresholds for the fourth firm (facing no further entry) decline with σ.[5]

Next examine the exit thresholds in the table's second panel. Without demand uncertainty, these form a line out of the origin with a slope approximately equal to 0.44 . As expected, raising σ decreases all of the exit thresholds. This mimics the well-known effect of increased uncertainty on monopolists' exit decisions: Uncertainty raises the value of the firm's put option, and exit requires this option's exercise. The exit thresholds fall much more for later entrants. For example, raising σ from 0 to 0.3 lowers $\underline{C}_2$ from 0.87 to 0.63 and lowers $\underline{C}_4$ from 1.73 to 1.30 . Recall that C_t approximately follows a random walk in logarithms, so its level displays heteroskedasticity. The greater variance conditional on higher values of C_t make later entrant's option values larger than earlier entrants for any given value of σ, so increasing σ has a greater effect on these later entrants' exit thresholds. For completeness, Table 1 reports the equilibrium exit thresholds when $\pi(N) = 4 \times I\{N < 5\}$. As expected, this change has almost no impact on the exit thresholds for firms with ranks less than four. For the fourth firm, eliminating the possibility of further entry makes survival substantially more attractive and thereby lowers the exit threshold even further.

To summarize, adding uncertainty either leaves the equilibrium entry thresholds unchanged or raises them somewhat. This result embodies two effects: Increasing uncertainty alone would make entry more attractive, but the accompanying increase in the probability of future entry reduces expected future profits. On the other hand, adding uncertainty substantially reduces equilibrium exit thresholds. When demand is homoskedastic in logarithms, the resulting heteroskedasticity in levels implies that the exit threshold falls much more for higher-ranked entrants. Thus, the number of *surviving* firms becomes more sensitive to a given change in market size as uncertainty rises.

4.3 Static Analysis of Market Size and Entry

We now characterize how a static analysis of market size and entry interprets data generated by our dynamic model. For this, it is helpful to briefly review a stylized version of the entry model examined by Bresnahan and Reiss (Bresnahan and Reiss (1990, 1991b)). As in our dynamic model, the profits per firm are $(C/N) \times \pi(N)$ and at most $\check{N}$ firms serve the industry. The fixed costs to a firm serving the market are $e^{\varepsilon}\kappa$, where ε is a mean-zero

[5] Entry by an eighth firm does not occur when there is demand uncertainty, so this discussion begs the question of why increasing σ from 0 to 0.10 raises $\overline{C}_7$ from 3.49 to 3.60 . This change reflects the mean reversion noted above. The same principle explains the rise of $\overline{C}_6$ from 3.13 to 3.19 when σ goes from 0.10 to 0.20.

Table 1: Equilibrium Entry and Exit Thresholds(i,ii)

$\pi(N) = 4$

Entry Thresholds

σ	1	2	3	4	5	6	7	8
0.00	0.50	1.00	1.49	1.99	2.48	2.97	3.49	3.97
0.10	0.48	0.98	1.51	2.03	2.59	3.13	3.60	
0.20	0.46	0.99	1.57	2.12	2.72	3.19		
0.30	0.46	1.02	1.62	2.18	2.56			

Exit Thresholds

σ	1	2	3	4	5	6	7	8
0.00	0.44	0.87	1.31	1.73	2.18	2.61	3.03	3.49
0.10	0.38	0.76	1.15	1.54	1.92	2.32	2.69	
0.20	0.33	0.68	1.05	1.40	1.77	2.10		
0.30	0.29	0.63	0.97	1.30	1.58			

$\pi(N) = 4 \times I\{N < 5\}$

Entry Thresholds

σ	1	2	3	4	5	6	7	8
0.00	0.50	1.00	1.49	1.99				
0.10	0.48	0.98	1.52	1.90				
0.20	0.46	0.99	1.57	1.77				
0.30	0.45	1.03	1.57	1.75				

Exit Thresholds

σ	1	2	3	4	5	6	7	8
0.00	0.44	0.87	1.31	1.73				
0.10	0.38	0.76	1.15	1.51				
0.20	0.33	0.68	1.05	1.30				
0.30	0.29	0.64	0.96	1.19				

(i) The parameter values used were $\kappa = 1.75$, $\beta = 1.05^{-1}$, $\varphi = 0.25 \times (1 - \beta)/\beta$, and $Q(c|C)$ a mixture over uniformly-distributed random walks reflected off of $\hat{C}$ and $\check{C}$ with approximate innovation variance σ^2. (ii) An empty cell indicates that the ergodic distribution of N_t puts zero probability on the given value of N. Please see the text for details.

normally-distributed shock. There are no sunk costs. Free entry requires that all active firms earn a non-negative profit and that an additional firm would earn a non-positive profit. That is

$$\frac{C}{N} \times \pi(N) \geq e^{\varepsilon}\kappa \qquad \text{and} \qquad \frac{C}{N+1}\pi(N) \leq e^{\varepsilon}\kappa.$$

For each $N = 1, \ldots, \check{N}$ define the deterministic entry threshold $C^{\star}_N$ to be the unique solution to $(C/N)\pi(N) - \kappa = 0$. Exactly N firms will serve the industry if If $\ln C > \ln C^{\star}_N + \varepsilon$, and $\ln C \leq \ln C^{\star}_{N+1} + \varepsilon$. The probability that this occurs is $\Phi\left(\frac{\ln(C/C^{\star}_N)}{\nu}\right) - \Phi\left(\frac{\ln(C/C^{\star}_{N+1})}{\nu}\right)$. In this expression, we set $C^{\star}_0 = 0$ and $C^{\star}_{\check{N}+1} = \infty$. Given observations of C and N from a cross section of markets, ordered Probit estimation immediately yields estimates for $C^{\star}_1, \ldots, C^{\star}_{\check{N}}$ and ν. With these we can estimate how profits per customer fall with additional competitors. Specifically, the definition of $C^{\star}_N$ implies that $\pi(N)/\pi(1) = C^{\star}_1 \times N/C^{\star}_N$. If the level of demand required to support N firms equals N times the level required for a monopolist, then we infer that profits per customer do not fall with additional entry. On the other hand, if demand must exceed $N \times C^{\star}_1$ to induce N firms to enter, then profits per customer must decline with N. In this way, the Probit analysis infers the toughness of competition from the relationship between market size and the number of competitors.

For a given joint distribution of C and N, we can define the population counterparts to the estimated thresholds by minimizing the population analogue of the ordered Probit's log-likelihood function.

$$\mathcal{L}(C^{\star}_1, \ldots, C^{\star}_{\check{N}}, \nu) \equiv \mathbb{E}\left[\sum_{R=0}^{\check{N}} I\{N = R\} \ln\left(\Phi\left(\frac{\ln(C/C^{\star}_R)}{\nu}\right) - \Phi\left(\frac{\ln(C/C^{\star}_{R+1})}{\nu}\right)\right)\right]$$

Because the ordered Probit likelihood function is always concave (even if it does not represent the true data generating process), this function has a unique minimizer. Population "estimates" of $\pi(N)/\pi(1)$ correspond to these thresholds. Calculating these for our model's data generating process and comparing them to their true values indicates whether abstraction from dynamic considerations substantially biases the static estimates of the toughness of competition.

For the model specification examined above the first panel of Table 2 reports the ordered probit estimates of $C^{\star}_1, \ldots, C^{\star}_{\check{N}}$ for the positive values of σ considered above, and its bottom panel gives the implied estimates of $\pi(N)/\pi(1)$. For all three of values of σ used, the static entry thresholds almost exactly equal the average of the dynamic model's corresponding entry and exit thresholds. That is, the static analysis "splits the difference" between them.

Table 2: Static Probit Analysis of Market Structure[(i)]

Implied Static Entry Thresholds

σ	1	2	3	4	5	6	7
0.10	0.43	0.87	1.32	1.78	2.24	2.70	3.13
0.20	0.39	0.83	1.29	1.74	2.21	2.60	
0.30	0.35	0.81	1.26	1.69	2.03		

Implied $\pi(N)/\pi(1)$

σ	1	2	3	4	5	6	7
0.10	1.00	0.98	0.97	0.96	0.95	0.94	0.95
0.20	1.00	0.95	0.91	0.91	0.89	0.91	
0.30	1.00	0.86	0.83	0.82	0.85		

(i) The table's reports population values of probit-based entry thresholds from the static model of Bresnahan and Reiss calculated using the ergodic distribution of the dynamic model. The bottom panel reports the implied values of $\pi(N)$ normalized by $\pi(1)$. An empty cell indicates that the ergodic distribution of N_t puts zero probability on the given value of N. Please see the text for further details.

Consequentially, the implied thresholds fall with σ. Because the exit thresholds for higher ranked firms fall faster with σ, the static entry thresholds also grow somewhat less rapidly with N.

Recall that the true values of $\pi(N)/\pi(1)$ all equal one. That is, an additional competitor steals business from incumbents but does not lower the profit per customer. For the case with $\sigma = 0.10$, the implied values deviate little from the truth. However, raising σ further substantially lowers these "estimates". When $\sigma = 0.3$, the implied value of $\pi(2)/\pi(1)$ equals 0.85. Further increases in N change this little.

Apparently, the static probit analysis can find evidence that $\pi(N)$ falls from a dynamic model in which $\pi(N)$ is constant. How so? If we consider the ordered probit's thresholds as estimates of $\overline{C}_1, \overline{C}_2, \ldots, \overline{C}_{\check{N}}$, the dynamics impart a downward bias to them. Consider the impact of a very simple downward bias which subtracts off a constant Δ from each static threshold. Denote the new thresholds with $\hat{C}_1, \hat{C}_2, \ldots, \hat{C}_{\check{N}}$. Using these to infer the toughness of competition yields

$$\frac{\hat{\pi}(N)}{\hat{\pi}(1)} = \frac{\hat{C}_1 N}{\hat{C}_N} = \frac{(\overline{C}_1 - \Delta)N}{\overline{C}_N - \Delta} = \frac{\overline{C}_N - \Delta N}{\overline{C}_N - \Delta}$$

Clearly, this decreases with N, so a simple downward bias can lead to the inference that $\pi(N)$ falls with N. In our analysis, the delay in exit arising from option-value considerations imparts a substantial downward bias to each estimated threshold. In this way, omitting dynamic considerations from a long-run analysis of industry structure can lead to a finding of falling profits when in fact profits per customer are constant.

This bias is large in the specification under consideration. Determining the impact of this bias on empirical work must proceed on a case-by-case basis. For a particular application, it could be that substantial cross-market variation in permanent market size makes the dynamic considerations discussed here less important. However, we expect option-value considerations to pervade oligopolists' exit decisions. To the extent that this is so, inference based on the relationship between market size and the number of competitors requires a full dynamic analysis. The Last-In First-Out oligopoly model of this paper provides one framework for doing so.[6]

5 Technology Dynamics

In the model, $\pi(N)$ is stationary and does not depend on the identity of the firm or its history. The previous literature on industry dynamics suggests relaxing this in two ways. A firm's productivity could improve with experience, or it could be stochastic and require Bayesian learning on the part of owners. In this section, show that the basic approach we follow can accommodate these two extensions.

5.1 Learning by Ageing

We begin with the learning curve. The most popular specification for technology which displays such intertemporal economies of scope is "learning-by-doing." That is, the production frontier shifts out with the level of previous cumulative output. Benkard (2000) estimated

[6] It is worth noting that Bresnahan and Reiss (1991b) report in their abstract that "Our empirical results suggest that competitive conduct changes quickly as the number of incuments increases. In markets with five or fewer incumbents, almost all variation in competitive conduct occurs with the entry of the second or third firm." This is exactly the pattern displayed in Table 2.

such a specification which also includes "forgetting-by-not-doing" using data from the production of a wide-bodied aircraft, and he investigated the consequences of such a technology for oligopolistic dynamics in Benkard (2004). Following this approach requires explicitly modelling firms' production decisions and incorporating them into the dynamic game. This would be an interesting extension, but it is overly ambitious for the present paper. Instead, we adopt here a specification of *learning by ageing*. That is, a firm's technology frontier expands deterministically with the passage of time. Bahk and Gort (1993) estimated such a specification for the learning curve using a panel of U.S. manufacturing plants, and Cabral (1993) examines how such learning impacts oligopoly dynamics in a model similar to ours but with constant demand and simultaneous entry and continuation decisions.

To incorporate a learning curve into our model, we alter two of its assumptions. First, firms have heterogeneous fixed costs. Each period $\check{N}$ potential entrants have an opportunity to enter. The first has type $T = 1$, the second, $T = 2$, etc. A firm's fixed costs in its first period of operation are $\xi = \kappa + \nu(T)$, where $\nu(\cdot)$ is positive and strictly increasing. Thereafter, the firm's fixed costs evolve deterministically according to $\xi' - \kappa = \lambda(\xi - \kappa)$, where $0 < \lambda < 1$. So that older firms always have lower costs than their younger rivals, we assume that $\lambda\nu(T) < \nu(1)$. Second, incumbent producers and potential entrants make their continuation decisions *simultaneously* instead of sequentially. The model without learning by ageing relies on sequential continuation decisions to make a firm's rank payoff relevant. Incorporating learning by ageing allows us to dispense with this assumption.[7]

In this environment, the payoff relevant state for entry and continuation decisions is C, and the vector of incumbents fixed costs. A Markov-perfect equilibrium is a pair of functions of this state giving the probabilities of survival and entry. We follow Cabral (1993) and focus on a "natural" equilibrium in which low-cost firms never exit while leaving behind a high-cost competitor and no high-cost potential entrant actually enters at the same time that a low-cost potential entrant remains inactive. Because a firm's cost decreases strictly with its age, any such equilibrium has a last-in first-out structure. Hence, it is straightforward to demonstrate analogous results to Propositions 1 and 2 using simple extensions of their proofs.

To see the relationship between the natural equilibrium in this setting and the last-in first-out equilibrium in our model, consider a sequence of specifications for firms' fixed costs in which $\nu(\cdot)$ converges to zero. Because the value functions are continuous in firms' fixed costs,

[7]The analysis of LIFO equilibrium assigns firms entering in the same period different ranks. In the model with the learning curve, the technology types distinguish simultaneous entrants.

the limits of the equilibrium value functions equal their counterparts from the equilibrium in LIFO strategies. In this specific sense, the assumption of sequential continuation decisions and a restriction to LIFO strategies "stand in" for learning by ageing.

5.2 Bayesian Learning and Technology Shocks

Entry entails risk. Jovanovic (1982) modelled this risk as imperfect information about a time-invariant productivity parameter. A firm's owner optimally infers its value given noisy observations and makes continuation decisions based on this inference. Hopenhayn (1992) uses a similar specification with observable but continuously evolving productivity to generate a declining hazard rate for exit. Here, we demonstrate that such technology shocks can be added to our model without destroying its simplicity.

To do so, we again focus on the case where a firm's fixed cost varies over time. Denote firm i's fixed cost at time t with κ_{it}. This can take on one of two values, $\hat{\kappa} \leq \check{\kappa}$. At the time of entry, this fixed cost is drawn from a distribution with probability p on $\kappa_{it} = \hat{\kappa}$. Thereafter, it evolves according to a Markov chain. This fixed cost is observable to all market participants after production takes place. In this sense, this specification is closer to Hopenhayn's (1992) than Jovanovic's (1982), but the firm's owner does learn a substantial amount about productivity after its first period of production.

If the difference between $\check{\kappa}$ and $\hat{\kappa}$ is large enough, then it is possible that there does not exist a symmetric Markov-perfect equilibrium in a LIFO strategy. To see this, suppose that p is close to one – so that the realization $\check{\kappa}$ is very unlikely – the probability of transiting from $\hat{\kappa}$ to $\check{\kappa}$ is small, and $\check{\kappa}$ is an absorbing state. If $\check{\kappa}$ greatly exceeds $\hat{\kappa}$ than an old firm with such a poor draw of κ_{it} might find continuation unprofitable, even given LIFO expectations. If such a firm exits, then its younger competitors' ranks decrease. In this case, calculating a firm's continuation value requires understanding the policies of firms' with lower ranks.

This difficulty disappears in two cases. In the first, the Markov transition matrix equals the identity matrix and survival as a perpetual monopolist is only profitable if $\kappa_{it} = \hat{\kappa}$. Allowing for such a possibility would add the realistic feature that adding a long-lived competitor potentially requires many entrepreneurs to try and fail first. In the second case, the difference between $\check{\kappa}$ and $\hat{\kappa}$ is small. Because incumbents' value function is *strictly* decreasing in the firm's rank except in degenerate and uninteresting cases, we know that such small shocks will never induce a low ranked incumbent to exit before a high-ranked rival. Thus, the expectation that LIFO always holds is rational. Adding such a small technology shock would alter the timing of firms' exits but not their order.

6 Related Literature

This paper's analysis implicitly relies upon a great deal of previous work on the theory of dynamic games and the empirics of industry structure and dynamics. This section serves to acknowledge this dependence explicitly. There are two areas of previous research that are particularly important for us.

6.1 Timing and Expectational Assumptions

The sequential nature of firms' entry and exit decisions allows Markovian strategies to themselves depend on a firm's rank. This and the assumption that firms rationally expect last-in-first-out dynamics substantially structures our analysis. In some previous work, the assumption that firms move sequentially gives early movers a form of commitment to their actions. Examples are Dixit's (1980) two-period Stackelberg investment game and Maskin and Tirole's (1988) infinite-horizon alternating-moves quantity game. In other work with finite-horizon games, ordering players moves selects a unique Nash equilibrium for empirical analysis. As Berry (1992) notes, this approach is particularly useful when firm-specific observable variables are of substantial interest. Sequencing firms' actions need not select a single Markov-perfect equilbrium in an infinite-horizon setting like ours. In this case, researchers sometimes structure expectations with assumptions – such as *LIFO* – to select a "natural" equilibrium. Cabral's (1993) restriction that high-cost firms exit before their low-cost counterparts provides one example of such an expectational assumption. [8]

The most common defense of timing assumptions, that incumbents can take actions earlier simply by virtue of their incumbency, applies to our work as well. In the equilibrium we consider, these assumptions on timing and expectations make older firms more valuable than their otherwise identical younger counterparts. For this reason, we expect incumbent firms' to use the tools available to them (such as cheap talk) to move potential entrants' expectations towards those we consider. A formal consideration of equilibrium selection is, however, well beyond the scope of this paper.

[8] Assumptions on agents' expectations can also select a "natural" equilibrium in finite-stage games of dynamic oligopoly with incomplete or private information. For example, Bagwell et al. (1997) assume that imperfectly informed customers rationally expect the firms that charged the lowest price previously will do so again. This selects an equilibrium in which otherwise static price decisions have dynamic consequences.

6.2 Dynamic Empirical Analysis of Oligopolistic Markets

Ericson and Pakes (1995) proposed a framework for the empirical analysis of Markov-perfect dynamics which is particularly well-suited for modelling oligopolists' investment choices. Benkard (2004) provides one example of its application. It allows for a wide variety of such dynamic investment decisions, but there is no characterization of its equilibrium set beyond the existence proof due to Doraszelski and Satterthwaite (2005). Accordingly, the estimation of this framework's unknown parameters either occurs "offline" as in Benkard (2004) or by considering each firm's decision problem in isolation and letting the data reveal which equilibrium was played in sample, as in Bajari, Benkard, and Levin (Forthcoming).

Strictly speaking, the Ericson and Pakes framework encompasses our model, but we abstract from its most compelling feature, technological change arising from investment decisions. Nevertheless, we expect that the possibility of business-stealing future entry in those models will also severely reduce a new firm's option value and thereby reduce entry. Bresnahan and Reiss (1993) take a different empirical approach to dynamic oligopoly to which our framework can also contribute. They consider panel observations of the numbers of consumers and producers from concentrated markets for dental services. Their goal is to estimate oligopolists' fixed and sunk costs (κ and $\varphi(N)$ in our model). They acknowledge the numerous theoretical difficulties associated with an infinite horizon model of oligopolistic entry and exit (such as ours) and they then proceed to estimate a much more tractable two-period model in which entry and exit thresholds determine the number of operating firms given its previous value and the current number of consumers. To us, the loss from using a two-period model seems small in their case, because eight years separate their observations. However, such an approach becomes unacceptable with more frequent observations, such as the monthly entry, sales, and exit data used by Abbring and Campbell (2006b). We believe that a simple extension of this paper's LIFO equilibrium model which includes econometric error could be appropriate for the estimation of oligopolists' sunk costs with such a data set.

7 Conclusion

Because there is a unique Markov-perfect equilibrium in LIFO strategies, we can conduct comparative dynamics experiments such as that above exploring the effects of uncertainty on entry and exit rules. Raising late entrant's sunk costs is another experiment of interest for industrial organization which our framework naturally accommodates. A companion paper to this (Abbring and Campbell, 2006a) applies this framework to another experiment of interest for industrial organization: raising late entrant's sunk costs. For the case of an industry with at most two firms, we prove that raising barriers to a second producer's entry increases the probability that some firm will serve the industry and decreases its long-run entry and exit rates. In numerical examples with more than two firms, imposing a barrier to entry stabilizes industry structure. Another natural application of this framework is the estimation of oligopoly entry and exit thresholds, as in Bresnahan and Reiss (1993). This awaits future research.

References

Abbring, J. and J. R. Campbell (2006a). Oligopoly dynamics with barriers to entry. Technical report, Federal Reserve Bank of Chicago.

Abbring, J. H. and J. R. Campbell (2006b). A firm's first year. Technical report, Federal Reserve Bank of Chicago.

Bagwell, K., G. Ramey, and D. F. Spulber (1997). Dynamic retail price and investment competition. *RAND Journal of Economics 28*(2), 207–227.

Bahk, B. and M. Gort (1993). Decomposing learning by doing in new plants. *Journal of Political Economy 101*, 561–583.

Bajari, P. C., C. L. Benkard, and J. Levin. Estimating dynamic models of imperfect competition. *Econometrica.*

Benkard, C. L. (2000, September). Learning and forgetting: The dynamics of aircraft production. *American Economic Review 90*(4), 1034–1054.

Benkard, C. L. (2004). A dynamic analysis of the market for wide-bodied commercial aircraft. *Review of Economic Studies 71*(3), 581–611.

Berry, S. (1992). Estimation of a model of entry in the airline industry. *Econometrica 60*, 889–917.

Bresnahan, T. F. and P. C. Reiss (1990). Entry in monopoly markets. *Review of Economic Studies 57*(4), 531–553.

Bresnahan, T. F. and P. C. Reiss (1991a). Empirical models of discrete games. *Journal of Econometrics 48*(1-2), 57–81.

Bresnahan, T. F. and P. C. Reiss (1991b, October). Entry and competition in concentrated markets. *Journal of Political Economy 99*(5), 977–1009.

Bresnahan, T. F. and P. C. Reiss (1993). Measuring the importance of sunk costs. *Annales d'Economie et de Statistique 31*, 181–217.

Cabral, L. M. (1993). Experience advantages and entry dynamics. *Journal of Economic Theory 59*, 403–416.

Dixit, A. (1980, March). The role of investment in entry-deterrence. *Economic Journal 90*(357), 95–106.

Doraszelski, U. and M. Satterthwaite (2005). Foundations of markov-perfect industry dynamics: Existence, purification, and multiplicity. Technical report, Harvard University and Northwestern University.

Ericson, R. and A. Pakes (1995). Markov-perfect industry dynamics: A framework for empirical work. *Review of Economic Studies 62*, 53–82.

Hopenhayn, H. A. (1992). Entry, exit, and firm dynamics in long run equilibrium. *Econometrica 60*, 1127–1150.

Jovanovic, B. (1982). Selection and the evolution of industry. *Econometrica 50*, 649–670.

Maskin, E. and J. Tirole (1988). A theory of dynamic oligopoly, i: Overview and quantity competition with large fixed costs. *Econometrica 56*(3), 549–569.

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

Appendix

A Proofs of Results in Section 2

Proof of Proposition 1. The proof proceeds by first constructing a candidate equilibrium strategy and then verifying that it does indeed form an equilibrium. To construct the candidate strategy, we consider first the exit decision problem of a firm that entered with rank $\check{N}$ and expects all incumbent firms at the time of its entry to remain active indefinitely. For such a firm, the only relevant state is C, so we can simplify the Bellman equation (1) in this case to

$$v_{\check{N}}(C) = \max\left\{\beta\mathbb{E}\left[\pi(\check{N})\frac{C'}{\check{N}} - \kappa + v_{\check{N}}(C')|C\right], 0\right\}. \tag{3}$$

It is straightforward to verify that the Bellman operator in (3) satisfies Blackwell's sufficient conditions for a contraction mapping. Furthermore, we can restrict the Bellman equation to the set of functions from $[\hat{C}, \check{C}]$ into $\left[0, (\beta/(1-\beta))\pi(\check{N})(\check{C}/\check{N})\right]$. This is a complete metric space, so the contraction mapping theorem implies that there exists a unique value function satisfying (3). Denote the set $\{C|v_{\check{N}}(C) > 0\}$ with $\mathcal{S}_{\check{N}}$ and the set $\{C|v_{\check{N}}(C) > \varphi\}$ with $\mathcal{E}_{\check{N}}$. Under the maintained hypotheses of this maximization problem, this firm chooses to remain active if and only if $C \in \mathcal{S}_{\check{N}}$ and it chooses to enter the industry if and only if $C \in \mathcal{E}_{\check{N}}$.[9]

The construction of the candidate equilibrium continues by considering the related exit decision problem of a firm that entered with rank $R < \check{N}$, expects all incumbent firms at the time of its entry to remain active indefinitely, and expects the number of firms to evolve according to the deterministic $\{R, R+1, \ldots, \check{N}\}$-valued transition function $N'_R(C, N)$. The relevant state for this firm's exit decision problem reduces to (C, N) and the Bellman equation can be simplified to

$$v_R(C, N) = \max\left\{\beta\mathbb{E}\left[\pi(N'_R(C, N))\frac{C'}{N'_R(C, N)} - \kappa + v_R(C', N'_R(C, N))\right], 0\right\}. \tag{4}$$

For the same reasons that firm $\check{N}$'s associated functional equation has a unique solution, so does this one. We denote the set $\{C|v_R(C, R) > 0\}$ with $\mathcal{S}_R$ and the set $\{C|v_R(C, R) > \varphi\}$ with $\mathcal{E}_R$.

[9] Note that the specification of $\mathcal{S}_{\check{N}}$ and $\mathcal{E}_{\check{N}}$ ensures that firms default to inactivity in the case of indifference.

We finish the construction of the candidate equilibrium by solving these $\check{N}$ decision problems using the following recursive definition of N':

$$N'_R(N,C) = \begin{cases} \check{N} & \text{if } R = \check{N} \\ N'_{R+1}(N,C) - I(N \leq R, C \notin \mathcal{E}_{R+1}) - I(N > R, C \notin \mathcal{S}_{R+1}) & \text{if } R < \check{N}. \end{cases}$$

With the resulting sets $\mathcal{S}_R$ and $\mathcal{E}_R$, the candidate equilibrium strategy (A_S, A_E) is given by the continuation and entry rules

$$A_S(N-R, C, R') = \begin{cases} 1 & \text{if } C \in \mathcal{S}_{R'} \\ 0 & \text{otherwise} \end{cases}$$

and

$$A_E(C, R') = \begin{cases} 1 & \text{if } C \in \mathcal{E}_{R'} \\ 0 & \text{otherwise.} \end{cases}$$

We now proceed to verify that this strategy forms a symmetric Markov-perfect equilibrium. We begin by demonstrating that the candidate strategy is a LIFO strategy. For this, it is sufficient and necessary that $\mathcal{S}_R \subseteq \mathcal{S}_{R-1}$ and $\mathcal{E}_R \subseteq \mathcal{E}_{R-1}$. To do so, consider iterating on the Bellman equation for firm $R-1$ beginning with

$$v^1_{R-1}(C,N) = \begin{cases} v_R(C,R) & \text{if } N = R-1 \\ v_R(C,N) & \text{if } N > R-1. \end{cases}$$

This is the value function for firm R's exit decision problem extended to $\left[\hat{C}, \check{C}\right] \times \{R-1, R, \ldots, \check{N}\}$. Applying the Bellman operator in Equation (4) to this yields

$$v^2_{R-1}(C,N) = \begin{cases} \beta \max\{0, \frac{\pi(R-1)\mu(C)}{R-1} - \kappa + \mathbb{E}\left[v_R(C',R) \mid C\right]\} & \text{if } N = R-1, \\ \beta \max\{0, \frac{\pi(N'_{R-1}(C,N))\mu(C)}{N'_{R-1}(C,N)} - \kappa + \mathbb{E}\left[v_R(C', N'_{R-1}(C,N)) \,\middle|\, C\right]\} & \text{if } N > R-1. \end{cases}$$

Note that because $N'_{R-1}(C,N) = N'_R(C,N)$ if $N > R-1$, $v^2_{R-1}(C,N) = v_R(C,N)$ in this case. Furthermore, $v^2_{R-1}(C,R-1) > v^1_{R-1}(C,R-1)$. Denote the value function arising from the j'th iteration of the Bellman equation with v^j_{R-1}. The Bellman operator in (4) is monotonic, so the fact that $v^2_{R-1} \geq v^1_{R-1}$ implies that $v^j_{R-1} \geq v^{j-1}_{R-1}$. The limit of this sequence equals v_{R-1}, so we conclude that $v_{R-1}(C,N) \geq v_R(C,N)$. From this and the definitions of $\mathcal{S}_R$ and $\mathcal{E}_R$, the conjecture above follows immediately.

The next step is to show that deviations of potential entrants from the strategy do not increase their payoffs. Consider a potential entrant with state $H_E = (C, R')$. The strategy directs this firm to enter if $C \in \mathcal{E}_{R'}$ and to remain out of the industry otherwise. If the strategy calls for entry, then the payoff to entry is $v_{R'}(C, N'_{R'}(C, R'))$, because the strategy dictates that the number of firms evolves according to $N'_{R'}(C, N)$ after this firm's entry. The strategy calls for entry if and only if this payoff exceeds the entry cost, so deviation from the strategy cannot increase the firm's payoff.

The final step is to show that deviations of incumbent firms from the strategy do not increase their payoffs. Consider an incumbent with state $H_S = (N - R, C, R')$. The strategy directs this firm to remain active if $C \in \mathcal{S}_{R'}$ and to leave the industry otherwise. If the strategy calls for survival, then following it yields a payoff of $v_{R'}(C, N'_{R'}(C, N - R + R'))$, which is strictly greater than zero if and only if the strategy calls for survival. Hence, deviation from the strategy cannot increase the firm's payoff. $\square$

Proof of Proposition 2. The LIFO strategy constructed in the proof of Proposition 1 defaults to inactivity. Thus, a symmetric Markov-perfect equilibrium in a LIFO strategy that defaults to inactivity exists.

To obtain uniqueness, first note that $v_S(N - R, C, R')$ equals $v_{R'}(C, N'_{R'}(C, N - R + R'))$ for all $N - R$, C, and R' in any symmetric Markov-perfect equilibrium in a LIFO strategy, and that $v_{R'}$ is uniquely determined by Equation (3) for $R' = \check{N}$ and by Equation (4) for $R' < \check{N}$. So, the the strategy constructed in the proof of Proposition 1 is the unique LIFO strategy that is both optimal given these value functions and defaults to inactivity. $\square$

B Proofs of Results in Section 3

We first develop some auxiliary results.

Definition 5. *A function $f : (\hat{C}, \check{C}] \mapsto \mathbb{R}$ is* $\tilde{c}$-separable*, $\tilde{c} \in (\hat{C}, \check{C}]$, if (i) $f(c) \geq f(\tilde{c})$ for all $c > \tilde{c}$ and (ii) $f(c) \leq f(\tilde{c})$ for all $c < \tilde{c}$.*

Lemma 1. *Let $f : (\hat{C}, \check{C}] \mapsto \mathbb{R}$ be integrable with respect to $Q(\cdot|c)$, $\tilde{c}$-separable, and non-decreasing on $(\hat{C}, \tilde{c}]$, for some $\tilde{c} \in (\hat{C}, \check{C}]$. Suppose that μ is non-decreasing and that either*

(i). $Q(\cdot|c)$ is degenerate at $\mu(c) \leq c$, $c \in (\hat{C}, \check{C}]$, or

(ii). $Q(\cdot|c)$ is uniform on $(\mu(c) - \frac{\sigma}{2}, \mu(c) + \frac{\sigma}{2}] \subseteq (\hat{C}, \check{C}]$ with $\mu(c) - \frac{\sigma}{2} \leq c$ and $\sigma > 0$, $c \in (\hat{C}, \check{C}]$, and f is $\tilde{c}$-separable.

Then, $g(c) := \int_{\hat{C}}^{\check{C}} f(c')dQ(c'|c)$ is non-decreasing on $(\hat{C}, \tilde{c}]$.

Proof. In Case (i), the result follows immediately from $g(c) = f(\mu(c))$. Now consider Case (ii). First, note that $g(c) = \int_{\mu(c)-\sigma/2}^{\mu(c)+\sigma/2} f(u)du$. Because f is non-decreasing on $(\hat{C}, \tilde{c}]$, it immediately follows that g is non-decreasing on $\{c \in (\hat{C}, \tilde{c}] | \mu(c) + \sigma/2 \leq \tilde{c}\}$. Next, for $c' \leq c'' \leq \tilde{c}$ such that $\mu(c') + \sigma/2 \geq \tilde{c}$, we have that

$$\begin{aligned}\sigma\left(g(c'') - g(c')\right) &= \int_{\mu(c')+\sigma/2}^{\mu(c'')+\sigma/2} f(u)du - \int_{\mu(c')-\sigma/2}^{\mu(c'')-\sigma/2} f(u)du \\ &\geq \int_{\mu(c')+\sigma/2}^{\mu(c'')+\sigma/2} f(\tilde{c})du - \int_{\mu(c')-\sigma/2}^{\mu(c'')-\sigma/2} f(\tilde{c})du \\ &= 0.\end{aligned}$$

Taken together, this implies that g is non-decreasing on $(\hat{C}, \tilde{c}]$. $\square$

Lemma 2. *Let $f : (\hat{C}, \check{C}] \mapsto \mathbb{R}$ be integrable with respect to $Q(\cdot|c)$, $\tilde{c}$-separable, and non-decreasing on $(\hat{C}, \tilde{c}]$, for some $\tilde{c} \in (\hat{C}, \check{C}]$. Let $Q^K = \sum_{k=1}^{K} p_k Q_k$ for some $p_1, \ldots, p_k \geq 0$ such that $\sum_{k=1}^{K} p_k = 1$ and Q_k, and corresponding μ_K, satisfying the conditions of Lemma 1. Then, $g^K(c) := \int_{\hat{C}}^{\check{C}} f(c')dQ^K(c'|c)$ is non-decreasing on $(\hat{C}, \tilde{c}]$.*

Proof. Lemma 1 implies that $g_k(c) = \int_{\hat{C}}^{\check{C}} f(c')dQ_k(c'|c)$ is non-decreasing on $(\hat{C}, \tilde{c}]$, $k = 1, \ldots, K$. In turn, because $g^K = \sum_{k=1}^{K} p_k g_k$, this implies that g^K is non-decreasing on $(\hat{C}, \tilde{c}]$. $\square$

Lemma 3. *Let $f : (\hat{C}, \check{C}] \mapsto \mathbb{R}$ be bounded, $\tilde{c}$-separable, and non-decreasing on $(\hat{C}, \tilde{c}]$, for some $\tilde{c} \in (\hat{C}, \check{C}]$. Let $Q^1, Q^2, \ldots$ be a sequence of mixture Markov transition functions satisfying the conditions of Lemma 2 such that $\sup|Q^K - Q| \to 0$ for some Markov transition distribution function Q as $K \to \infty$. Then, $g(c) := \int_{\hat{C}}^{\check{C}} f(c')dQ(c'|c)$ is non-decreasing on $(\hat{C}, \tilde{c}]$.*

Proof. Lemma 2 implies that the function g^K corresponding to each Q^K, $K = 1, 2, \ldots$, is non-decreasing on $(\hat{C}, \tilde{c}]$. Because f is bounded, $g^K \to g$ as $K \to \infty$ and g is non-decreasing on $(\hat{C}, \tilde{c}]$. $\square$

Proof of Proposition 4. The proof begins with a characterization of $\mathcal{S}_{\check{N}} = \{C | v_{\check{N}}(C) > 0\}$ and $\mathcal{E}_{\check{N}} = \{C | v_{\check{N}}(C) > \varphi\}$. The Bellman operator in Equation (3) maps the space of value functions that are non-decreasing in C into itself, so the value function $v_{\check{N}}(C)$ is non-decreasing in C. It immediately follows that there exist thresholds $\underline{C}_{\check{N}}$ and $\overline{C}_{\check{N}}$ such that $\mathcal{S}_{\check{N}} = \{C \mid C > \underline{C}_{\check{N}}\}$ and $\mathcal{E}_{\check{N}} = \{C \mid C > \overline{C}_{\check{N}}\}$. Note that either of these thresholds might equal $\hat{C}$ or $\check{C}$.

Next, consider the characterization of $\mathcal{S}_{R-1} = \{C|v_{R-1}(C, R-1) > 0\}$ and $\mathcal{E}_{R-1} = \{C|v_{R-1}(R-1) > \varphi\}$ given that, for all $\tilde{R} \geq R$, there exist thresholds $\overline{C}_{\tilde{R}}$ and $\underline{C}_{\tilde{R}}$ such that $\mathcal{S}_{\tilde{R}} = \{C \,|\, C > \underline{C}_{\tilde{R}}\}$ and $\mathcal{E}_{\tilde{R}} = \{C \,|\, C > \overline{C}_{\tilde{R}}\}$ and $v_{\tilde{R}}(C, \tilde{R})$ is non-decreasing for all $C < \overline{C}_{\tilde{R}}$. There are two cases to consider.

In the first, $\overline{C}_{\tilde{R}} = \check{C}$ for all $\tilde{R} \geq R$, so that a firm entering with rank $R-1$ expects no further entry to occur during its lifetime. In that case, $v_{R-1}(C, R-1)$ is non-decreasing in C, so there exist thresholds $\underline{C}_{R-1}$ and $\overline{C}_{R-1}$ such that $\mathcal{S}_{R-1} = \{C \mid C > \underline{C}_{R-1}\}$ and $\mathcal{E}_{R-1} = \{C \,|\, C > \overline{C}_{R-1}\}$.

In the second case, $\overline{C}_R < \check{C}$. If $v_{R-1}(C, R-1) \geq \varphi$ for all C, then we can set $\underline{C}_{R-1} = \overline{C}_{R-1} = \hat{C}$. Now consider the case that $v_{R-1}(C, R-1) < \varphi$ for some C. The remainder of the proof utilizes the increasing sequence v^j_{R-1} of value functions constructed in the proof of Proposition 1. First, note that $v^1_{R-1}(c, R-1)$, as a function of c, $\overline{C}^1_{R-1} := \overline{C}_R$, and φ satisfy the conditions of Lemma 3. Next, suppose that $v^{j-1}_{R-1}(c, R-1)$, as a function of c, $\overline{C}^{j-1}_{R-1}$, and φ satisfy the conditions of Lemma 3, for some $\overline{C}^{j-1}_{R-1}$. Then, $\mathbb{E}[v^{j-1}_{R-1}(C', R)|C = c]$ is non-decreasing in c on $(\hat{C}, \overline{C}^{j-1}_{R-1}]$. With this, inspection of Equation (4) determines that $v^j_{R-1}(c, R-1)$ is non-decreasing in c on the same interval. Furthermore, the proof of Proposition 1 shows that $v^j_{R-1}(C, R-1) \geq v^{j-1}_{R-1}(C, R-1)$. Together, this implies that $\overline{C}^j_{R-1} := \inf\{C|v^j_{R-1}(C, R-1) > \varphi\} \leq \overline{C}^{j-1}_{R-1}$ and that $v^j_{R-1}(c, R-1)$, as a function of c, $\overline{C}^j_{R-1}$, and φ satisfy the conditions of Lemma 3.

Define $\overline{C}_{R-1} = \lim_{j\to\infty} \overline{C}^j_{R-1}$. First, note that $v^j_{R-1}(c, R-1)$ is non-decreasing and weakly smaller than φ on $(\hat{C}, \overline{C}^j_{R-1})$ for all j, so that $v_{R-1}(C, R-1)$ is non-decreasing and weakly smaller than φ on $(\hat{C}, \overline{C}_{R-1})$. Next, note that $v^1_{R-1}(c, R-1) > \varphi$ for $c \in (\overline{C}^1_{R-1}, \check{C}]$. Suppose that $v^{j-1}_{R-1}(c, R-1) > \varphi$ for all $c \in (\overline{C}^{j-1}_{R-1}, \check{C}]$. Because the Bellman operator in Equation (4) is monotone, $v^j_{R-1}(c, R-1) > \varphi$ for all $c \in (\overline{C}^{j-1}_{R-1}, \check{C}]$ as well. Furthermore, because $v^j_{R-1}(c, R-1)$ is non-decreasing on $(\hat{C}, \overline{C}^j_{R-1})$, $v^j_{R-1}(c, R-1) > \varphi$ for all $c \in (\overline{C}^j_{R-1}, \overline{C}^{j-1}_{R-1}]$ by the definition of $\overline{C}^j_{R-1}$. Because the sequence $\{v^j_{R-1}\}$ is non-decreasing, $v_{R-1}(c, R-1) > \varphi$ for all $c \in (\overline{C}_{R-1}, \check{C}]$. Thus, $\mathcal{S}_{R-1} = \{C|C > \overline{C}_{R-1}\}$. Finally, because $v_{R-1}(c, R-1)$ is non-decreasing in c for $c < \overline{C}_{R-1}$, $\mathcal{E}_{R-1} = \{C|C > \underline{C}_{R-1}\}$ for $\underline{C}_{R-1} := \inf\{C|v_{R-1}(C, R-1) > 0\}$. $\square$