

Markov Perfect Industry Dynamics with Many Firms*

Gabriel Y. Weintraub
Stanford University

C. Lanier Benkard
Stanford University
and NBER

Benjamin Van Roy
Stanford University

PRELIMINARY
October 16, 2005

Abstract

We propose an approximation method for analyzing Ericson and Pakes (1995)-style dynamic models of imperfect competition. We develop a simple algorithm for computing an “oblivious equilibrium,” in which each firm is assumed to make decisions based only on its own state and knowledge of the long run average industry state, but where firms ignore current information about competitors’ states. We prove that, as the market becomes large, if the equilibrium distribution of firm states obeys a certain “light-tail” condition, then oblivious equilibria closely approximate Markov perfect equilibria. We develop bounds that can be computed to assess the accuracy of the approximation for any given applied problem. Through computational experiments, we find that the method often generates useful approximations for industries with hundreds of firms and in some cases even tens of firms.

JEL classification numbers: C63, C73, L11, L13

*We have had very helpful conversations with Herman Bennett, José Blanchet, Uli Doraszelski, Liran Einav, Hugo Hopenhayn, Cristóbal Huneeus, Ken Judd, and Jon Levin, as well as with seminar participants at SITE, UCLA, and University of Chile. This research was supported by the National Science Foundation through CAREER Grant ECS-9985229, by the Office of Naval Research through Grant MURI N00014-00-1-0637, by General Motors through the GM-Stanford Collaborative Research Lab on Work Systems, and by the Lillie Fund. Correspondence: gweintra@stanford.edu; lanierb@stanford.edu; bvr@stanford.edu.

Contents

1	Introduction	1
2	A Dynamic Model of Imperfect Competition	4
2.1	Model and Notation	4
2.2	Model Primitives	6
2.3	Assumptions	6
2.4	Equilibrium	8
3	Oblivious Equilibrium	9
3.1	Oblivious Strategies and Entry Rate Functions	10
3.2	Oblivious Equilibrium	10
3.3	The Invariant Industry Distribution	12
4	Asymptotic Results	13
4.1	Assumptions about the Sequence of Profit Functions	14
4.2	Asymptotic Markov Equilibrium	14
4.3	Uniform Law of Large Numbers	16
4.4	A Light-Tail Condition Implies AME	16
5	Algorithm and Bounds	18
5.1	Computing Oblivious Equilibria	19
5.2	Error Bounds	20
6	Computational Experiments	22
6.1	The Computational Model	22
6.2	Numerical Results: Behavior of the Bound	23
6.3	Closeness to Markov Perfect Equilibrium	25
6.4	Example: The Evolution of Industry Structure as Market Size Increases	26
7	Conclusions and Future Research	27
A	Proofs and Mathematical Arguments	29
A.1	Proofs and Mathematical Arguments for Section 2	29
A.1.1	The Poisson Entry Model	29
A.1.2	Bellman's Equation	30
A.2	Proofs and Mathematical Arguments for Section 3	30
A.3	Proofs and Mathematical Arguments for Section 4	35
A.3.1	Preliminary Lemmas	35
A.4	Proof of Theorem 4.2	36
A.4.1	Proof of Theorem 4.4	37
A.4.2	Example: Logit Demand System with Price Competition	41
A.5	Proofs and Mathematical Arguments for Section 5	42
B	Tables and Figures	49

1 Introduction

Ericson and Pakes (1995) (hereafter EP) introduced an approach to modeling industry dynamics in which entry, exit, and investment, together with idiosyncratic shocks, result in heterogeneity among firms. The analysis of such models — which we will refer to as EP-type models — relies on computation of Markov perfect equilibria (MPE) using dynamic programming algorithms (e.g., Pakes and McGuire (1994)). A great advantage of the EP framework is that it is easily extended to cover many important dynamic phenomena.¹ A major shortcoming, however, is the computational complexity of solving for MPE. Methods that accelerate these computations have been proposed (Pakes and McGuire (2001) and Doraszelski and Judd (2003)). However, even with such improvements, in practice computational concerns have typically limited the analysis to industries with just two or three firms. Such limitations have made it difficult to construct realistic empirical models, and application of the EP framework to empirical problems has been rare (exceptions include Benkard (2004), Gowrisankaran and Town (1997), Jenkins, Liu, Matzkin, and McFadden (2004), and Ryan (2005)). More generally, model details are often dictated as much by computational concerns as economic ones.

In an EP-type model, at each time, each firm has a state variable that captures its competitive advantage. Though more general state spaces can be considered, we focus on the simple case where the firm state is an integer. The value of this integer can represent, for example, a measure of product quality, the firm’s current productivity level, or its capacity. Each firm’s state evolves over time based on investments and random shocks. The *industry state* is a vector representing the number of firms with each possible value of the firm state variable. Even if firms are restricted to symmetric strategies, the number of relevant industry states (and thus, the compute time and memory required for computing a MPE) becomes enormous very quickly. For example, most industries contain more than 20 firms, but it would require approximately 22 million gigabytes of computer memory to store the policy function for an industry with just 20 firms and 40 firm states. As a result, it seems unlikely that exact computation of equilibria will ever be possible in many applied problems of interest.

With this motivation, in this paper we instead propose an approximation method, one that dramatically reduces the computational complexity of EP-type models in industries with many firms. The intuition behind our approach is as follows. Consider an EP-type model in which firm shocks are idiosyncratic. In each period, some firms receive positive shocks and some receive negative shocks. Now suppose there are a large

¹See, for example, Benkard (2004), Berry and Pakes (1993), Besanko and Doraszelski (2004), Besanko, Doraszelski, Kryukov, and Satterthwaite (2005), Doraszelski and Markovich (2003), Fershtman and Pakes (2000), Goettler, Parlour, and Rajan (2004), Gowrisankaran (1999), Jenkins, Liu, Matzkin, and McFadden (2004), Judd, Schmedders, and Yeltekin (2002), Langohr (2003), Markovich (2003), de Roos (2002), and Song (2003), as well as Pakes (2000) for a survey.

number of firms. It is natural to think that changes in individual firms' states average out at the industry level, such that the industry state does not change much over time. In that case, each firm could make near-optimal decisions knowing only its own firm state and the long run average industry state. We call *oblivious* strategies, strategies for which a firm considers only its own state and the long run average industry state, and we will define a new solution concept, called *oblivious equilibrium*, in which firms use oblivious strategies. Computing an oblivious equilibrium is simple because dynamic programming algorithms that optimize over oblivious strategies require compute time and memory that scale only with the number of firm states, and not with the number of firms. Indeed, it is easy to compute oblivious equilibria for industries with thousands of firms and hundreds of firm states.

To formalize the intuition above, we prove an asymptotic result that provides sufficient conditions for oblivious equilibria to closely approximate MPE as the market size grows. It may seem that this would be true provided that the average number of firms in the industry grows to infinity as the market size grows. However, this is not sufficient. If the market is highly concentrated — for example, as is the case with Microsoft in the software industry — then the approximation is unlikely to be accurate. A strategy that does not keep track of the dominant firm's state will not perform well. Instead, we show that, alongside some technical requirements, a sufficient condition for oblivious equilibria to well approximate MPE asymptotically is that they generate a firm size distribution that is “light-tailed,” in a sense that we will make precise. For example, if the demand system is given by a logit model and the spot market equilibrium is Nash in prices, then the condition holds if the average firm size remains bounded by the same number for all market sizes.

We provide an algorithm based on dynamic programming that computes oblivious equilibria. The algorithm is computationally light, often terminating within a couple minutes of run time on a common laptop computer even for industries with thousands of firms. It is also easy to implement, requiring, typically, fewer than two hundred lines of Matlab code. This represents a considerable savings over existing algorithms. Another distinguishing feature of the algorithm is that it places no a priori restrictions on the number of firms or the number of firm states. Instead, these are determined endogenously and computed alongside the oblivious equilibrium.

Our asymptotic result provides a condition under which the approximation is accurate for large markets. We also derive bounds on the approximation error that can be efficiently computed for any given applied problem. Despite the practical importance of such error bounds, there are very few cases in the approximate dynamic programming literature where researchers have been able to provide useful bounding techniques for high-dimensional stochastic control problems. Possibly the only relevant examples involve optimal stopping

(Haugh, Kogan, and Wang 2004) and portfolio optimization (Haugh, Kogan, and Wang 2005). As such, our bounding technique represents a significant contribution. We show that these error bounds can be derived quite generally; they do not require many of our modeling assumptions. Furthermore, while it is important that there be no aggregate shocks for the asymptotic results to hold, we are able to derive error bounds even for models that incorporate aggregate shocks. Using this bounding algorithm, we find that oblivious equilibria often offer accurate approximations for industries involving hundreds of firms, and in some cases even tens of firms.

Though our emphasis is on the use of oblivious equilibrium as an approximation of MPE, oblivious equilibrium can also be motivated as a behavioral model in its own right. If observing the industry state and designing strategies that keep track of it are costly, and do not lead to significant increases in profit, firms may be better off using oblivious strategies.

The concept of oblivious equilibrium is closely related to Hopenhayn (1992). Hopenhayn models an industry that hosts an infinite number of firms, each of which garners an infinitesimal fraction of the market. His model is tractable because it assumes that the industry state is constant over time, implicitly assuming a law of large number holds. This assumption is based on the same intuition that motivates our consideration of oblivious equilibrium. However, our goal is to analyze models that closely reflect real world industries that have finite numbers of firms. Also, our EP-type model is more general because the transitional dynamics resulting from firms' investment strategies are generated by equilibrium behavior that is explicitly modelled. Hopenhayn abstracts from this aspect of the model and instead assumes that firms' state trajectories (their productivities) follow exogenous Markov processes.²

Our approach also has similarities to a number of other past literatures. The light-tail condition we consider is analogous to notions of diffuse industry structure in large markets associated with Sutton (1991)'s exogenous sunk cost model. Though our goals are different, our asymptotic results are close in spirit to the work of Novshek and Sonnenschein (1978) who, in a static setting, provide conditions under which Cournot-Nash equilibrium converges to Walras competitive equilibrium when there is free entry. The notion that, for asymptotically large markets, strategies can remain effective while ignoring information and strategic behavior also appears in Vives (2002), in the context of Cournot models. Finally, Jovanovic and Rosenthal (1988) consider sequential equilibria in a model where the actual equilibrium strategies can be represented as simple functions of summary statistics of the distribution of firm states. This differs from our context, in which simple functions provide only an approximation to equilibrium strategies.

There are a variety of relevant economic issues that can be studied using our methods. As an example,

²Note that Klette and Kortum (2003) and Melitz (2003) model investment explicitly in Hopenhayn-style models.

we briefly explore an important question in industrial organization: What features of an industry determine whether an industry becomes fragmented or remains concentrated as the market grows in size? Sutton (1991) sought to identify simple features of an industry that distinguish models that lead to these two outcomes. In our model, through computational experiments, we show that an arbitrarily small increase in a single model parameter that identifies the extent of vertical product differentiation, can turn an asymptotically fragmented market into an asymptotically concentrated one. These results show that the predictions of Sutton (1991) may sometimes be quite sensitive. For example, they imply that very different market structures might be observed in the same industry across markets that are the same size and that have indistinguishable characteristics. Note that our results remain broadly consistent with Sutton (1991) since the extent of vertical product differentiation impacts the returns to investment. Consistent with Sutton (1991), industries with higher returns to investment (even if the difference may be arbitrarily small) tend to be more concentrated.

The paper is organized as follows. In Section 2 we outline the dynamic industry model. In Section 3 we introduce the concept of oblivious strategies and oblivious equilibrium. In Section 4 we provide conditions under which oblivious strategies approximate MPE strategies asymptotically as the market size grows. In Section 5 we provide methods for computing oblivious equilibria and error bounds. In Section 6 we report results from computational experiments. Finally, Section 7 presents conclusions and a discussion of future research directions.

2 A Dynamic Model of Imperfect Competition

In this section we formulate a model of an industry in which firms compete in a single-good market. Our model is close in spirit to that of Ericson and Pakes (1995), but with some differences. Most notably, we modify the entry and exit processes in Ericson and Pakes (1995) so as to make them more realistic when there are a large number of firms. Additionally, our asymptotic results do not hold with aggregate industry shocks so our model includes only idiosyncratic shocks.

2.1 Model and Notation

The industry evolves over discrete time periods and an infinite horizon. We index time periods with non-negative integers $t \in \mathbb{N}$ ($\mathbb{N} = \{0, 1, 2, \dots\}$). All random variables are defined on a probability space $(\Omega, \mathcal{F}, \mathcal{P})$ equipped with a filtration $\{\mathcal{F}_t : t \geq 0\}$. We adopt a convention of indexing by t variables that are \mathcal{F}_t -measurable.

Each firm that enters the industry is assigned a unique positive integer-valued index. The set of indices of incumbent firms at time t is denoted by S_t . At each time $t \in \mathbb{N}$, we denote the number of incumbent firms as n_t .

Firm heterogeneity is reflected through firm states. To fix an interpretation, we will refer to a firm's state as its quality level. However, firm states might more generally reflect productivity, capacity, the size of its consumer network, or any other aspect of the firm that affects its profits. At time t , the quality level of firm $i \in S_t$ is denoted by $x_{it} \in \mathbb{N}$.

We define the *industry state* s_t to be a vector over quality levels that specifies, for each quality level $x \in \mathbb{N}$, the number of incumbent firms at quality level x in period t . Though there are a countable number of industry states, we will consider an extended state space $\mathcal{S} = \left\{ s \in \mathfrak{R}_+^\infty \mid \sum_{x=0}^\infty s(x) < \infty \right\}$. This will allow us, for example, to consider derivatives of functions with respect to the industry state. For each $i \in S_t$, we define $s_{-i,t} \in \mathcal{S}$ to be the state of the *competitors* of firm i ; that is, $s_{-i,t}(x) = s_t(x) - 1$ if $x_{it} = x$, and $s_{-i,t}(x) = s_t(x)$, otherwise. Similarly, $n_{-i,t}$ denotes to the number of competitors of firm i .

In each period, each incumbent firm earns profits on a spot market. A firm's single period expected profit $\pi(x_{it}, s_{-i,t})$ depends on its quality level x_{it} and its competitors' state $s_{-i,t}$.

The model also allows for entry and exit. In each period, each incumbent firm $i \in S_t$ observes a positive real-valued sell-off value ϕ_{it} that is private information to the firm. If the sell-off value exceeds the value of continuing in the industry then the firm may choose to exit, in which case it earns the sell-off value and then ceases operations permanently.

If the firm instead decides to remain in the industry, then it can invest to improve its quality level. If a firm invests $\iota_{it} \in \mathfrak{R}_+$, then the firm's state at time $t + 1$ is given by,

$$x_{i,t+1} = x_{it} + w(\iota_{it}, \zeta_{i,t+1}),$$

where the function w captures the impact of investment on quality and $\zeta_{i,t+1}$ reflects uncertainty in the outcome of investment. Uncertainty may arise, for example, due to the risk associated with a research and development endeavor or a marketing campaign. We denote the unit cost of investment by d .

In each period new firms can enter the industry by paying a setup cost κ . Entrants do not earn profits in the period that they enter. They appear in the following period at state $x^e \in \mathbb{N}$ and can earn profits thereafter.

Each firm aims to maximize expected net present value. The interest rate is assumed to be positive and constant over time, resulting in a constant discount factor of $\beta \in (0, 1)$ per time period.

In each period, events occur in the following order:

1. Each incumbent firms observes its sell-off value and then makes exit and investment decisions.
2. The number of entering firms is determined and each entrant pays an entry cost of κ .
3. Incumbent firms compete in the spot market and receive profits.
4. Exiting firms exit and receive their sell-off values.
5. Investment outcomes are determined, new entrants enter, and the industry takes on a new state s_{t+1} .

2.2 Model Primitives

The model as specified is general enough to encompass numerous applied problems in economics. Indeed, similar models have been applied to advertising, auctions, collusion, consumer learning, environmental policy, international trade policy, learning-by-doing, limit order markets, mergers, network externalities, and other applied problems. To study any particular applied problem it is necessary to further specify the primitives of the model, including:

profit function	π
sell-off value distribution	$\sim \phi_{it}$
investment impact function	w
investment uncertainty distribution	$\sim \zeta_{it}$
unit investment cost	d
entry cost	κ
discount factor	β

Note that in most applied problems the profit function would not be specified directly, but would instead result from a deeper set of primitives that specify a demand function, a cost function, and a static equilibrium concept. An important parameter of the demand function, that we will focus on below, is the size of the relevant market, which we will denote as m .

2.3 Assumptions

We make several assumptions about the model primitives, beginning with the profit function. An industry state $s \in \mathcal{S}$ is said to *dominate* $s' \in \mathcal{S}$ if for all $x \in \mathbb{N}$, $\sum_{x \leq z} s(z) \geq \sum_{x \leq z} s'(z)$. We will denote this relation by $s \succeq s'$. Intuitively, competition associated with s is no weaker than competition associated with s' .

Assumption 2.1.

1. For all $s \in \mathcal{S}$, $\pi(x, s)$ is increasing in x .
2. For all $x \in \mathbb{N}$ and $s, s' \in \mathcal{S}$, if $s \succeq s'$ then $\pi(x, s) \leq \pi(x, s')$.
3. For all $x \in \mathbb{N}$ and $s \in \mathcal{S}$, $\pi(x, s) > 0$, and $\sup_{x,s} \pi(x, s) < \infty$.
4. For all $x \in \mathbb{N}$, $y \in \mathbb{N}$, and $s \in \mathcal{S}$, $\pi(x, s)$ is differentiable with respect to $s(y)$. Further, for any $x \in \mathbb{N}$, $y \in \mathbb{N}$, $s \in \mathcal{S}$, and $h \in \mathcal{S}$ such that $s + \gamma h \in \mathcal{S}$ for $\gamma > 0$ sufficiently small, if

$$\sum_{y \in \mathbb{N}} h(y) \left| \frac{\partial \ln \pi(x, s)}{\partial s(y)} \right| < \infty,$$

then

$$\left. \frac{d \ln \pi(x, s + \gamma h)}{d\gamma} \right|_{\gamma=0} = \sum_{y \in \mathbb{N}} h(y) \frac{\partial \ln \pi(x, s)}{\partial s(y)}.$$

The assumptions are natural. Assumption 2.1.1 ensures that increases in quality lead to increases in profit. Assumption 2.1.2 states that strengthened competition cannot result in increased profit. Assumption 2.1.3 ensures that profits are positive and bounded. The first part of Assumption 2.1.4 requires partial differentiability of the profit function with respect to each $s(y)$. Profit functions that are smooth, such as the ones arising from random utility demand models like a logit model, will satisfy this assumption. The second part of Assumption 2.1.4 is technical and essentially requires that the profit function is Fréchet differentiable.

We also make assumptions about the distributions of the private shocks:

Assumption 2.2.

1. The variables $\{\phi_{it} | t \geq 0, i \geq 1\}$ are i.i.d. and have finite expectations and well-defined density functions with support \mathbb{R}_+ .
2. The random variables $\{\zeta_{it} | t \geq 0, i \geq 1\}$ are i.i.d. and independent of $\{\phi_{it} | t \geq 0, i \geq 1\}$.
3. For all ζ , $w(\iota, \zeta)$ is nondecreasing in ι .
4. For all $\iota > 0$, $\mathcal{P}[w(\iota, \zeta_{i,t+1}) > 0] > 0$.
5. There exists a positive constant $\bar{w} \in \mathbb{N}$ such that $|w(\iota, \zeta)| \leq \bar{w}$, for all (ι, ζ) .
6. For all $k \in \{-\bar{w}, \dots, \bar{w}\}$, $\mathcal{P}[w(\iota, \zeta_{i,t+1}) = k]$ is continuous in ι .
7. The transitions generated by $w(\iota, \zeta)$ are unique investment choice admissible .

Assumptions 2.2.1 and 2.2.2 imply that investment and exit outcomes are idiosyncratic conditional on the state. Assumption 2.2.3 and 2.2.4 imply that investment is productive. Assumption 2.2.5 places a finite bound on how much progress can be made or lost in a single period through investment. Assumption 2.2.6 ensures that the impact of investment on transition probabilities is continuous. Assumption 2.2.7 was introduced by Doraszelski and Satterthwaite (2003) to guarantee that the solution of the firms' investment decision problem is unique. It is used to prove existence of pure strategy equilibrium and it is satisfied by many of the commonly used specifications in the literature.

We assume that there are a large number of potential entrants who play a symmetric mixed entry strategy. In that case the number of actual entrants is well approximated by the Poisson distribution (see appendix for a derivation of this result). This leads to the following assumptions:

Assumption 2.3.

1. *The number of firms entering during period t is a Poisson random variable that is conditionally independent of $\{\phi_{it}, \zeta_{it} | t \geq 0, i \geq 1\}$, conditioned on s_t .*
2. *$\kappa > \beta \cdot \bar{\phi}$, where $\bar{\phi}$ is the expected net present value of entering the market, investing zero each period, and then exiting at an optimal stopping time.*

We denote the expected number of firms entering at industry state s_t , by $\lambda(s_t)$. This state-dependent entry rate will be endogenously determined, and our solution concept will require that it satisfies a zero expected profit condition. Modeling the number of entrants as a Poisson random variable has the advantage that it leads to more elegant asymptotic results. Assumption 2.3.2 ensures that the sell-off value by itself is not sufficient reason to enter the industry.

2.4 Equilibrium

As a model of industry behavior we focus on pure strategy Markov perfect equilibrium (MPE), in the sense of Maskin and Tirole (1988). We further assume that equilibrium is symmetric, such that all firms use a common stationary investment/exit strategy. In particular, there is a function ι such that at each time t , each incumbent firm $i \in S_t$ invests an amount $\iota_{it} = \iota(x_{it}, s_{-i,t})$. Similarly, each firm follows an exit strategy that takes the form of a cutoff rule: there is a real-valued function ρ such that an incumbent firm $i \in S_t$ exits at time t if and only if $\phi_{it} \geq \rho(x_{it}, s_{-i,t})$. In the appendix we show that there always exists an optimal exit strategy of this form even among very general classes of exit strategies. Let \mathcal{M} denote the set of exit/investment strategies such that an element $\mu \in \mathcal{M}$ is a pair of functions $\mu = (\iota, \rho)$, where $\iota : \mathbb{N} \times \mathcal{S} \rightarrow \mathfrak{R}_+$ is an investment strategy and $\rho : \mathbb{N} \times \mathcal{S} \rightarrow \mathfrak{R}_+$ is an exit strategy. Similarly, we denote the set of entry rate functions by Λ , where an element of Λ is a function $\lambda : \mathcal{S} \rightarrow \mathfrak{R}_+$.

We define the value function $V(x, s | \mu', \mu, \lambda)$ to be the expected net present value for a firm at state x when its competitors' state is s , given that its competitors each follows a common strategy $\mu \in \mathcal{M}$, the entry rate function is $\lambda \in \Lambda$, and the firm itself follows strategy $\mu' \in \mathcal{M}$. In particular,

$$V(x, s | \mu', \mu, \lambda) = E_{\mu', \mu, \lambda} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} (\pi_m(x_{ik}, s_{-i,k}) - d_{ik}) + \beta^{\tau_i-t} \phi_{i, \tau_i} \Big| x_{it} = x, s_{-i,t} = s \right],$$

where i is taken to be the index of a firm at quality level x at time t , τ_i is a random variable representing the time at which firm i exits the industry, and the subscripts of the expectation indicate the strategy followed

by firm i , the strategy followed by its competitors, and the entry rate function. In an abuse of notation, we will use the shorthand, $V(x, s|\mu, \lambda) \equiv V(x, s|\mu, \mu, \lambda)$, to refer to the expected discounted value of profits when firm i follows the same strategy μ as its competitors.

An equilibrium to our model comprises an investment/exit strategy $\mu = (\iota, \rho) \in \mathcal{M}$, and an entry rate function $\lambda \in \Lambda$ that satisfy the following conditions:

1. Incumbent firm strategies represent a MPE:

$$(2.1) \quad \sup_{\mu' \in \mathcal{M}} V(x, s|\mu', \mu, \lambda) = V(x, s|\mu, \lambda) \quad \forall x \in \mathbb{N}, \forall s \in \mathcal{S}.$$

In the appendix, we show that this supremum can always be attained simultaneously for all x and s by a common strategy μ' .

2. At each state, either entrants have zero expected profits or the entry rate is zero (or both):

$$\begin{aligned} \sum_{s \in \mathcal{S}} \lambda(s) (\beta E_{\mu, \lambda} [V(x^e, s_{t+1}|\mu, \lambda)|s_t = s] - \kappa) &= 0 \\ \beta E_{\mu, \lambda} [V(x^e, s_{t+1}|\mu, \lambda)|s_t = s] - \kappa &\leq 0 \quad \forall s \in \mathcal{S} \\ \lambda(s) &\geq 0 \quad \forall s \in \mathcal{S}. \end{aligned}$$

Doraszelski and Satterthwaite (2003) establish existence of an equilibrium in pure strategies for a closely related model. We do not provide an existence proof here because it is long and cumbersome and would replicate this previous work. With respect to uniqueness, in general we presume that our model may have multiple equilibria.³

Dynamic programming algorithms can be used to optimize firm strategies, and equilibria to our model can be computed via their iterative application. However, these algorithms require compute time and memory that grow proportionately with the number of relevant industry states, which is often intractable in contexts of practical interest. This difficulty motivates our alternative approach.

3 Oblivious Equilibrium

We will propose a method for approximating MPE based on the idea that when there are a large number of firms, simultaneous changes in individual firm quality levels can average out such that the normalized

³Doraszelski and Satterthwaite (2003) also provide an example of multiple equilibria in their closely related model.

industry state remains roughly constant over time. In this setting, each firm can potentially make near-optimal decisions based only on its own quality level and the long run average industry state. With this motivation, we consider restricting firm strategies so that each firm's decisions depend only on the firm's quality level. We call such restricted strategies *oblivious* since they involve decisions made without full knowledge of the circumstances — in particular, the state of the industry.

3.1 Oblivious Strategies and Entry Rate Functions

Let $\tilde{\mathcal{M}} \subset \mathcal{M}$ and $\tilde{\Lambda} \subset \Lambda$ denote the set of oblivious strategies and the set of oblivious entry rate functions. Since each strategy $\mu = (\iota, \rho) \in \tilde{\mathcal{M}}$ generates decisions $\iota(x, s)$ and $\rho(x, s)$ that do not depend on s , with some abuse of notation, we will often drop the second argument and write $\iota(x)$ and $\rho(x)$. Similarly, for an entry rate function $\lambda \in \tilde{\Lambda}$, we will denote by λ the real-valued entry rate which persists for all industry states.

3.2 Oblivious Equilibrium

Note that if all firms use a common strategy $\mu \in \tilde{\mathcal{M}}$, the quality level of each evolves as an independent transient Markov chain. Let the k -period transition sub-probabilities of this transient Markov chain be denoted by $P_\mu^k(x, y)$. Then, the expected time that a firm spends at a quality level x is given by $\sum_{k=0}^{\infty} P_\mu^k(x^e, x)$, and the expected lifespan of a firm is $\sum_{k=0}^{\infty} \sum_{x \in \mathbb{N}} P_\mu^k(x^e, x)$. Denote the expected number of firms at quality level x at time t by $\tilde{s}_t(x) = E[s_t(x)]$. The following result offers an expression for the long-run expected industry state when dynamics are governed by oblivious strategies and entry rate functions.

Lemma 3.1. *If firms make decisions according to an oblivious strategy $\mu \in \tilde{\mathcal{M}}$ and enter according to an oblivious entry rate function $\lambda \in \tilde{\Lambda}$, and the expected time that a firm spends in the industry is finite, then*

$$(3.1) \quad \lim_{t \rightarrow \infty} \tilde{s}_t(x) = \lambda \sum_{k=0}^{\infty} P_\mu^k(x^e, x),$$

for all $x \in \mathbb{N}$.

We omit the proof, which is straightforward. To abbreviate notation, we let $\tilde{s}_{\mu, \lambda}(x) = \lim_{t \rightarrow \infty} \tilde{s}_t(x)$ for $\mu \in \tilde{\mathcal{M}}$, $\lambda \in \tilde{\Lambda}$, and $x \in \mathbb{N}$. For an oblivious strategy $\mu \in \tilde{\mathcal{M}}$ and an oblivious entry rate function $\lambda \in \tilde{\Lambda}$ we define an *oblivious value function*

$$\tilde{V}(x | \mu', \mu, \lambda) = E_{\mu'} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} (\pi_m(x_{ik}, \tilde{s}_{\mu, \lambda}) - d_{ik}) + \beta^{\tau_i-t} \phi_{i, \tau_i} \Big| x_{it} = x \right].$$

This value function should be interpreted as the expected net present value of a firm that is at quality level x and follows oblivious strategy μ' , under the assumption that its competitors' state will be $\tilde{s}_{\mu,\lambda}$ for all time. Note that only the firm's own strategy μ' influences the firm's state trajectory because neither the profit function nor the strategy μ' depends on the industry state. Hence, the subscript in the expectation only reflects this dependence. Importantly, however, the oblivious value function remains a function of the competitors' strategy μ and the entry rate λ through the expected industry state $\tilde{s}_{\mu,\lambda}$. Again, we abuse notation by using $\tilde{V}(x|\mu, \lambda) \equiv \tilde{V}(x|\mu, \mu, \lambda)$ to refer to the oblivious value function when firm i follows the same strategy μ as its competitors.

We now define a new solution concept: an *oblivious equilibrium* consists of a strategy $\mu \in \tilde{\mathcal{M}}$ and an entry rate function $\lambda \in \tilde{\Lambda}$ that satisfy the following conditions:

1. Firm strategies optimize an oblivious value function:

$$(3.2) \quad \sup_{\mu' \in \tilde{\mathcal{M}}} \tilde{V}(x|\mu', \mu, \lambda) = \tilde{V}(x|\mu, \lambda), \quad \forall x \in \mathbb{N}.$$

2. Either the oblivious expected value of entry is zero or the entry rate is zero (or both):

$$\begin{aligned} \lambda \left(\beta \tilde{V}(x^e|\mu, \lambda) - \kappa \right) &= 0 \\ \beta \tilde{V}(x^e|\mu, \lambda) - \kappa &\leq 0 \\ \lambda &\geq 0. \end{aligned}$$

It is straightforward to show that an oblivious equilibrium exists under mild technical conditions. Furthermore, if the entry cost is not prohibitively high then an oblivious equilibrium with a positive entry rate exists. We omit the proof of this for brevity. With respect to uniqueness, we have been unable to find multiple oblivious equilibria in any of the applied problems we have considered, but similarly with the case of MPE, we have no reason to believe that in general there is a unique oblivious equilibrium.⁴

Note that we assume that, even if firms are using oblivious strategies, the sequence of single-period profits received are the ones associated with Nash, and not oblivious, static behavior. For example, if single-period profits are derived from a game in which firms compete in prices, we assume firms price according to Nash equilibrium strategies, and are not oblivious with respect to pricing behavior. We make this assumption because we are interested in the dynamic behavior of the industry, and we believe that assuming oblivious pricing behavior, which would be cumbersome, would not change our conclusions in that respect.

⁴However, since oblivious strategies rule out strategies that are dependent on competitors' states, there are likely to be fewer oblivious equilibria than there are MPE.

We will later explore situations where the number of firms is large and oblivious equilibria approximate MPE, in a sense that we will define precisely. First, we analyze the long-run behavior of an industry where strategies and the entry rate function are oblivious.

3.3 The Invariant Industry Distribution

In Lemma 3.1, we characterized the long-run expected industry state. Our next result characterizes the long-run distribution. The symbol \Rightarrow denotes weak convergence as $t \rightarrow \infty$.

Lemma 3.2. *Let Assumptions 2.2 and 2.3 hold. Assume that firms follow a common oblivious strategy $\mu \in \tilde{\mathcal{M}}$, the expected entry rate is $\lambda \in \tilde{\Lambda}$, and the expected time that each firm spends in the industry is finite. Let $\{Z_x : x \in \mathbb{N}\}$ be a sequence of independent Poisson random variables with means $\{\tilde{s}_{\mu,\lambda}(x) : x \in \mathbb{N}\}$, and let Z be a Poisson random variable with mean $\sum_{x \in \mathbb{N}} \tilde{s}_{\mu,\lambda}(x)$. Then,*

- (a) $\{s_t : t \geq 0\}$ is an irreducible, aperiodic and positive recurrent Markov chain;
- (b) the invariant distribution of s_t is a product form of Poisson random variables;
- (c) for all x , $s_t(x) \Rightarrow Z_x$;
- (d) $n_t \Rightarrow Z$.

To conclude this section we state an important result for later use. First note that

$$s_t(x) = \sum_{i \in S_t} \mathbf{1}_{\{x_{it}=x\}} = \sum_{j=1}^{n_t} \mathbf{1}_{\{x_{(j)t}=x\}},$$

where $\mathbf{1}_A$ denotes the indicator of event A . Hence, for example, $\mathbf{1}_{\{x_{it}=x\}} = 1$ if $x_{it} = x$, and $\mathbf{1}_{\{x_{it}=x\}} = 0$, otherwise. $\{x_{(j)t} : j = 1, \dots, n_t\}$ is a random permutation of $\{x_{it} : i \in S_t\}$. That is, we randomly pick a firm from S_t and assign to it the index $j = 1$; from the remaining firms we randomly pick another firm and assign to it the index $j = 2$, and so on.

Lemma 3.3. *Let Assumptions 2.2 and 2.3 hold. Assume that firms follow a common oblivious strategy $\mu \in \tilde{\mathcal{M}}$, the expected entry rate is $\lambda \in \tilde{\Lambda}$, and the expected time that each firm spends in the industry is finite. Let $\{Y_n : n \in \mathbb{N}\}$ be a sequence of integer-valued i.i.d. random variables, each distributed according to $\tilde{s}_{\mu,\lambda}(\cdot) / \sum_{x \in \mathbb{N}} \tilde{s}_{\mu,\lambda}(x)$. Then, for all $n \in \mathbb{N}$,*

$$\left(x_{(1)t}, \dots, x_{(n_t)t} \middle| n_t = n \right) \Rightarrow (Y_1, \dots, Y_n).$$

Proofs of Lemmas 3.2 and 3.3 can be found in the appendix. Assuming a Poisson entry process is key to proving these results. Lemma 3.3 ensures that if we sample a firm randomly from those firms currently

in the industry and the industry state is distributed according to the invariant distribution, the firm's state will be distributed according to the normalized expected industry state. Further, each subsequent time we sample without replacement we get an independent sample from the same distribution. For brevity, when we consider sampling a random firm from among those currently in the industry and the industry state is distributed according to the invariant distribution, we will say that we are sampling a firm from the industry's invariant distribution.

It is straightforward to show that if per-period profit is bounded, say by some quantity $\bar{\pi}$, then the expected time a firm spends in the industry is finite for any oblivious strategy $\mu \in \tilde{\mathcal{M}}$ that comprises an oblivious equilibrium. This follows from the fact that the sell-off value has support in \mathfrak{R}_+ and the continuation value from every state is bounded above by $\frac{\bar{\pi}}{1-\beta} + \bar{\phi}$. Hence, the probability of exiting in each period is bounded below by a positive constant. This implies that the previous lemmas apply when firms use oblivious equilibrium strategies.

4 Asymptotic Results

In this section, we establish asymptotic results that provide conditions under which oblivious equilibria offer close approximations to MPE as the market size grows. As specified above, our model does not explicitly depend on market size. However, market size would typically enter the profit function, $\pi(x_{it}, s_{-i,t})$, through the underlying demand system; in particular, profit for a firm at a given state (x, s) would typically increase with market size. Therefore, in this section we consider a sequence of markets, indexed by market sizes $m \in \mathbb{N}$. All other model primitives are assumed to remain constant within this sequence except for the profit function, which depends on m . To convey this dependence, we denote profit functions by π_m .

We index functions and random variables associated with market size m with a superscript (m) . From this point onward we let $(\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)})$ denote an oblivious equilibrium for market size m . Let $V^{(m)}$ and $\tilde{V}^{(m)}$ represent the value function and oblivious value function, respectively, when the market size is m . To further abbreviate notation we denote the expected industry state associated with $(\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)})$ by $\tilde{s}^{(m)} \equiv \tilde{s}_{\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}}$.

The random variable $s_t^{(m)}$ denotes the industry state at time t when every firm uses strategy $\tilde{\mu}^{(m)}$ and the entry rate is $\tilde{\lambda}^{(m)}$. We denote the invariant distribution of $\{s_t^{(m)} : t \geq 0\}$ by $q^{(m)}$. In order to simplify our analysis, we assume that the initial industry state $s_0^{(m)}$ is sampled from $q^{(m)}$. Hence, $s_t^{(m)}$ is a stationary process; $s_t^{(m)}$ is distributed according to $q^{(m)}$ for all $t \geq 0$. Note that this assumption does not affect asymptotic results, since for any initial condition, the process approaches stationarity as time progresses.

It will be helpful to decompose $s_t^{(m)}$ according to $s_t^{(m)} = f_t^{(m)} n_t^{(m)}$, where $f_t^{(m)}$ is the random vector that represents the fraction of firms in each state and $n_t^{(m)}$ is the total number of firms, respectively. Similarly, let $\tilde{f}^{(m)} \equiv E[f_t^{(m)}]$ denote the expected fraction of firms in each state and $\tilde{n}^{(m)} \equiv E[n_t^{(m)}] = \sum_{x \in \mathbb{N}} \tilde{s}^{(m)}(x)$ denote the expected number of firms. Using Lemma 3.3, it is easy to check that $\tilde{f}^{(m)} = \frac{\tilde{s}^{(m)}}{\tilde{n}^{(m)}}$. With some abuse of notation, we define $\pi_m(x_{it}, f_{-i,t}, n_{-i,t}) \equiv \pi_m(x_{it}, n_{-i,t} \cdot f_{-i,t})$.

4.1 Assumptions about the Sequence of Profit Functions

In addition to Assumption 2.1, which applies to individual profit functions, we will make the following assumptions, which apply to sequences of profit functions. Let $\mathcal{S}_1 = \{f \in \mathcal{S} \mid \sum_{x \in \mathbb{N}} f(x) = 1\}$ and $\mathcal{S}_{1,z} = \{f \in \mathcal{S}_1 \mid \forall x > z, f(x) = 0\}$.

Assumption 4.1.

1. $\sup_{x \in \mathbb{N}, s \in \mathcal{S}} \pi_m(x, s) = O(m)$.
2. For all increasing sequences $\{m_k \in \mathbb{N} \mid k \in \mathbb{N}\}$, $n : \mathbb{N} \mapsto \mathbb{N}$ with $n(m_k) = o(m_k)$, $x, z \in \mathbb{N}$ with $x > z$, and $f \in \mathcal{S}_{1,z}$, $\lim_{k \rightarrow \infty} \pi_{m_k}(x, f, n(m_k)) = \infty$.

3.

$$\sup_{m \in \mathbb{N}, x \in \mathbb{N}, f \in \mathcal{S}_1, n > 0} \left| \frac{d \ln \pi_m(x, f, n)}{d \ln n} \right| < \infty.$$

The assumptions are again natural. Assumption 4.1.1 should hold for virtually all relevant classes of profit functions. It is satisfied, for example, if the total disposable income of the consumer population grows linearly in market size.⁵ Assumption 4.1.2 is also natural. It states that if the number of firms grows slower than the market size then the largest firm's profit becomes arbitrarily large as the market grows. Assumption 4.1.3 states that the relative rate of change of profit with respect to relative changes in the number of firms is uniformly bounded. To provide a concrete example, we show in the Appendix that these assumptions are satisfied by a single period profit function derived from a demand system given by a logit model and where the spot market equilibrium is Nash in prices. Later, we use this profit function for the computational experiments and describe it in detail in Subsection 6.1.

4.2 Asymptotic Markov Equilibrium

Our aim is to establish that, under certain conditions, oblivious equilibria well-approximate MPE as the market size grows. We define the following concept to formalize the sense in which this approximation becomes exact.

⁵For example, if each consumer has income that is less than some upper bound \bar{Y} then total disposable income of the consumer population (an upper bound to firm profits) is always less than $m \cdot \bar{Y}$.

Definition 4.1. A sequence $(\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) \in \mathcal{M} \times \Lambda$ possesses the asymptotic Markov equilibrium (AME) property if for all $x \in \mathbb{N}$,

$$\lim_{m \rightarrow \infty} E_{\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\sup_{\mu' \in \mathcal{M}} V^{(m)}(x, s_t^{(m)} | \mu', \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) - V^{(m)}(x, s_t^{(m)} | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) \right] = 0.$$

Recall that the process s_t is taken to be stationary, and therefore, this expectation does not depend on t . The definition of AME assesses approximation error at each firm state x in terms of the amount by which a firm at state x can increase its net present value by deviating from the oblivious equilibrium strategy $\tilde{\mu}^{(m)}$, and instead following an optimal (non-oblivious) best response that keeps track of the true industry state. The notion of AME relates to the more common notion that ϵ -equilibria approximate true equilibria in games as $\epsilon \rightarrow 0$ (Fudenberg and Levine 1986).

Note that standard MPE solution algorithms (e.g., Pakes and McGuire (1994)) aim to compute point-wise ϵ -equilibria; that is, where a firm cannot improve its net present value by more than ϵ starting from *any* state (x, s) . The AME property instead considers the benefit of deviating to an optimal strategy starting from each firm state x , averaged over the invariant distribution of industry states. It would not be possible to obtain our results point-wise. This is because in an oblivious equilibrium firms may be making poor decisions in states that are far from the expected state. Offsetting this effect is the fact that these states have very low probability of occurrence, so they have a small impact on expected discounted profits.⁶

If a sequence of oblivious equilibria has the AME property then, as m grows, $\sup_{\mu' \in \mathcal{M}} V^{(m)}(x, s | \mu', \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) \approx V^{(m)}(x, s | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)})$ for states s with probabilities that are not vanishing. This implies that, asymptotically, $\tilde{\mu}^{(m)}$ is a near optimal strategy when the industry starts in any state that has significant probability of occurrence. Further, since $\beta \tilde{V}^{(m)}(x^e) = \kappa$ for all m , asymptotically, $\beta V^{(m)}(x^e, s | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) \approx \kappa$ for states s with probabilities that are not vanishing. Hence, asymptotically, $\tilde{\lambda}^{(m)}$ satisfies the zero profit condition at such states. In summary, if $E_{\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\sup_{\mu' \in \mathcal{M}} V^{(m)}(x, s_t^{(m)} | \mu', \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) - V^{(m)}(x, s_t^{(m)} | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) \right]$ is small, MPE strategies and entry rates at relevant states should be well approximated by oblivious ones. We will later present computational results that support this point.⁷

⁶Note that Pakes and McGuire (2001) use a similar concept as a stopping rule in their stochastic algorithm. They compare the difference in the value function from two consecutive iterates, weighting states according to their empirical distribution of visits in the long run.

⁷One might attempt to formalize this heuristic argument by following a similar line of reasoning as Fudenberg and Levine (1986) or Altman, Pourtallier, Haurie, and Moresino (2000). They provide conditions under which, if a sequence of restricted games, G_m , converges to a game of interest, G in an appropriate way, then any convergent sequence of ϵ_m -equilibria of G_m with $\epsilon_m \rightarrow 0$ converges to an equilibrium of G . In our case, we do not have a well defined limit game. If $m = \infty$, profits are infinite and we would have infinite entry. This makes it hard to use their line of thought.

4.3 Uniform Law of Large Numbers

The following theorem establishes that when the number of firms is large, the industry state becomes approximately constant (i.e., $s_t^{(m)} \approx \tilde{s}^{(m)}$) with high probability. We use \rightarrow_p to denote convergence in probability as $m \rightarrow \infty$.

Theorem 4.1. *If $\lim_{m \rightarrow \infty} \tilde{n}^{(m)} = \infty$ then*

$$\sup_{x \in \mathbb{N}} \left| \frac{s_t^{(m)}(x)}{\tilde{n}^{(m)}} - \frac{\tilde{s}^{(m)}(x)}{\tilde{n}^{(m)}} \right| \rightarrow_p 0.$$

The theorem can be proved by invoking a uniform law of large numbers (Vapnik and Chervonenkis (1971)) and using Lemma 3.3. It suggests that when the expected number of firms is large, using an oblivious strategy might be close to optimal, and that a sequence of oblivious equilibria possesses the AME property. However, for this to be the case it turns out that additional conditions are required.

4.4 A Light-Tail Condition Implies AME

Even when there are a large number of firms, if the market tends to be concentrated — for example, if the market is usually dominated by a single firm — the AME property is unlikely to hold. To ensure the AME property, we need to impose a “light-tail” condition that rules out this kind of market concentration.

Note that $\frac{d \ln \pi_m(y, f, n)}{df(x)}$ is the semi-elasticity of one period profits with respect to the fraction of firms in state x . We define the *maximal absolute semi-elasticity function*:

$$g(x) = \max_{m \in \mathbb{N}, y \in \mathbb{N}, f \in \mathcal{S}_1, n \in \mathbb{N}} \left| \frac{d \ln \pi_m(y, f, n)}{df(x)} \right|.$$

For each x , $g(x)$ is the maximum rate of relative change of any firm’s expected single-period profit that could result from a small changes in the fraction of firms at quality level x . We can also interpret $g(x)$ as the maximal possible impact of the presence of a firm at state x on the profit of any other firm. Note that since larger competitors tend to have greater influence on firm profits, $g(x)$ typically increases with x and can be unbounded.

Finally, we introduce our light-tail condition. For each m , let $\tilde{x}^{(m)} \sim \tilde{f}^{(m)}$. According to Lemma 3.3, $\tilde{x}^{(m)}$ can be interpreted as the quality level of a firm that is randomly sampled from among all incumbents while the industry is in steady state.

Assumption 4.2. For all quality levels x , $g(x) < \infty$. For all $\epsilon > 0$, there exists a quality level z such that

$$E \left[g(\tilde{x}^{(m)}) \mathbf{1}_{\{\tilde{x}^{(m)} > z\}} \right] \leq \epsilon,$$

for all market sizes m .

Recall that $g(x)$ is the maximum rate of relative change of any firm's expected single-period profit that could result from a small change in the fraction of firms at quality level x . The first part of the assumption requires that for any x , this quantity is finite. If this condition is not satisfied, a small change in the number of firms at quality level x can have an arbitrarily large impact on other firms as the market size grows. It is unlikely that an oblivious equilibrium will provide a good approximation in this situation. Note that the assumption imposes that $g(x)$ is finite for each x , however, $g(x)$ could grow arbitrarily large as x grows.

To interpret the second part of the assumption it is helpful to first understand a weaker condition: $E[g(\tilde{x}^{(m)})] < \infty$. This weaker condition ensures that the expected impact of a randomly sampled incumbent is finite. Note that the impact of a sampled firm should grow with its quality level and can become arbitrarily high as the quality level does. This weaker condition can be viewed as a “light tail” condition, since it requires that the probability of sampling firms at large quality levels dies off sufficiently quickly so that the expected impact remains finite.

Note that, for any x and z , the product $g(x) \mathbf{1}_{\{x > z\}}$ is equal to 0 if $x \leq z$ but otherwise is equal to $g(x)$. Hence, $E[g(\tilde{x}^{(m)}) \mathbf{1}_{\{\tilde{x}^{(m)} > z\}}]$ is similar to $E[g(\tilde{x}^{(m)})]$ but ignores the impact of any sampled firm if its quality level is z or lower. Consequently, $E[g(\tilde{x}^{(m)}) \mathbf{1}_{\{\tilde{x}^{(m)} > z\}}]$ bounds the expected impact of the presence of a randomly sampled firm if the impact of any firm with quality level z or lower is ignored.

It is easy to see that the condition $E[g(\tilde{x}^{(m)})] < \infty$ is equivalent to a condition that, for any $\epsilon > 0$, there exists a quality level z such that $E[g(\tilde{x}^{(m)}) \mathbf{1}_{\{\tilde{x}^{(m)} > z\}}] \leq \epsilon$. This is because increasing z sufficiently will result in ignoring a larger and larger number of firms in computing the expected impact and the expected impact when none of the firms is ignored is finite. Assumption 4.2 poses a stronger condition in that it requires that a quality level z can be chosen such that $E \left[g(\tilde{x}^{(m)}) \mathbf{1}_{\{\tilde{x}^{(m)} > z\}} \right] \leq \epsilon$ for all market sizes m simultaneously. This is like the “light tail condition” $E[g(\tilde{x}^{(m)}) \mathbf{1}_{\{\tilde{x}^{(m)} > z\}}] \leq \epsilon$, or equivalently $E[g(\tilde{x}^{(m)})] < \infty$, which applies to a fixed market size, but it precludes the possibility that the tail becomes arbitrarily “fat” as the market size increases. In a sense, it requires that the tails of quality distributions $\tilde{f}^{(m)}$ are uniformly “light” over market sizes m .

As a different way to explain the light-tail condition, we note that if $g(x)$ is unbounded, then if there exists $\gamma > 0$, such that, $\sup_m E[g(\tilde{x}^{(m)})^{1+\gamma}] < \infty$, Assumption 4.2 is satisfied. The condition is slightly

stronger than requiring uniformly bounded first moments of $g(\tilde{x}^{(m)})$.

To give a concrete illustration about our light-tail condition, in the Appendix we show that for a demand system given by a logit model and where the spot market equilibrium is Nash in prices (see Subsection 6.1), Assumption 4.2 is satisfied if $\sup_m E[\tilde{x}^{(m)}] < \infty$. That is, if the average firm quality level remains uniformly bounded over all market sizes. For example, if $\tilde{x}^{(m)}$ converges in distribution to a lognormal distribution, the light-tail condition is satisfied. On the other hand, if $\tilde{x}^{(m)}$ converges to a Pareto distribution with parameter one (which does not have first moment), then the light-tail condition is not satisfied.

The following theorem establishes that, asymptotically, the average number of firms grows at least linearly in the market size.

Theorem 4.2. *Under Assumptions 2.1, 2.2, 2.3, 4.1.2, and 4.2, $\tilde{n}^{(m)} = \Omega(m)$ ⁸.*

The proof can be found in the appendix. The intuition behind the result is simple. If the number of firms were to grow slower than the market size, profits would blow up and the zero profit condition at the entry state would not be met.

The next result, which is also proved in the appendix, establishes a stronger form of convergence than Theorem 4.1.

Theorem 4.3. *Let Assumptions 2.1, 2.2, 2.3, 4.1.2, and 4.2, hold. Then, as m grows, $n_t^{(m)}/\tilde{n}^{(m)} \rightarrow_p 1$ and $\|f_t^{(m)} - \tilde{f}^{(m)}\|_{1,g} \rightarrow_p 0$.*

This new form of convergence allows us to ensure the AME property, which leads to the main result of this section.

Theorem 4.4. *Under Assumptions 2.1, 2.2, 2.3, 4.1, and 4.2, the sequence $\{\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}\}$ of oblivious equilibria possesses the AME property.*

This result is proved in the appendix.

5 Algorithm and Bounds

In this section we propose an algorithm that computes oblivious equilibria for a particular set of model primitives. We then derive expressions that bound the approximation error associated with a particular oblivious equilibrium.

⁸I.e., $\liminf_m \frac{\tilde{n}^{(m)}}{m} > 0$. With an additional technical regularity condition, it is straightforward to show that $\tilde{n}^{(m)} = O(m)$; i.e., $\limsup_m \frac{\tilde{n}^{(m)}}{m} < \infty$. Hence, $n^{(m)} = \Theta(m)$; i.e., asymptotically, $\tilde{n}^{(m)}$ grows linearly in m .

5.1 Computing Oblivious Equilibria

Algorithm 1 (below) is designed to compute an oblivious equilibrium. It starts with two extreme entry rates: $\underline{\lambda} = 0$ and $\bar{\lambda} = \sup_{x,s} \pi(x, s)/\kappa$. Under mild assumptions, any oblivious equilibrium entry rate must lie between these two extremes. The algorithm searches over entry rates between these two extremes for one that leads to an oblivious equilibrium.⁹ For each candidate entry rate λ , an inner loop (lines 5-8) computes an oblivious equilibrium firm strategy for that fixed entry rate. If the termination conditions of both the inner and outer loops are satisfied with $\epsilon_1 = \epsilon_2 = 0$, we have an oblivious equilibrium. Small values of ϵ_1 and ϵ_2 allow for small errors associated with limitations of numerical precision.

Algorithm 1 Oblivious Equilibrium Solver

```

1:  $\underline{\lambda} := 0; \bar{\lambda} := \sup_{x,s} \pi(x, s)/\kappa$ 
2:  $\mu(x) := 0$  for all  $x$ 
3: repeat
4:    $\lambda := (\underline{\lambda} + \bar{\lambda})/2$ 
5:   repeat
6:      $\mu' := \mu$ 
7:     Choose  $\mu \in \mathcal{M}$  to maximize  $\tilde{V}(x|\mu, \mu', \lambda)$  simultaneously for all  $x \in \mathbb{N}$ 
8:     until  $\|\mu - \mu'\|_\infty \leq \epsilon_1$ 
9:     if  $\beta\tilde{V}(x^e|\mu, \lambda) - \kappa \geq 0$  then
10:       $\underline{\lambda} := \lambda$ 
11:     else
12:       $\bar{\lambda} := \lambda$ 
13:     end if
14: until  $|\beta\tilde{V}(x^e|\mu, \lambda) - \kappa| \leq \epsilon_2$ 

```

The algorithm is easy to program and computationally efficient. In each iteration of the inner loop, the optimization problem to be solved is a one dimensional dynamic program. The state space in this dynamic program is the set of quality levels a firm can achieve. In principle, there could be an infinite number of them. However, beyond a certain quality level the optimal strategy for a firm is not to invest, so its quality cannot increase to beyond that level. In the numerical experiments we present in Section 6, the state space never had more than two hundred states per firm. The exact number of states is determined during execution of the algorithm.

Whether this algorithm is guaranteed to terminate in a finite number of iterations remains an open issue. However, in over 90% of the numerical experiments we present in the next section, it converged in less than five minutes. In the rest, it converged in less than fifteen minutes.¹⁰

⁹Note that there are potentially many alternative methods for searching over entry rates for an oblivious equilibrium. For example, one alternative would be to start at an arbitrary entry rate and then implement small increments and decrements to the entry rate until an entry rate is found that leads to an oblivious equilibrium.

¹⁰The algorithm was programmed in Matlab. The experiments were executed on two UNIX shared machines. A SunEnterprise

5.2 Error Bounds

To bound approximation error, we first need to define what is meant by *approximation error*. Consider an oblivious strategy and entry rate function $(\tilde{\mu}, \tilde{\lambda}) \in \tilde{\mathcal{M}}$. We will quantify approximation error at each firm state $x \in \mathbb{N}$ by

$$E \left[\sup_{\mu' \in \mathcal{M}} V(x, s_t | \mu', \tilde{\mu}, \tilde{\lambda}) - V(x, s_t | \tilde{\mu}, \tilde{\lambda}) \right].$$

Note that the expectation is over the invariant distribution of s_t . Hence, approximation error is the amount by which a firm at state $x \in \mathbb{N}$ can improve its expected net present value by unilaterally deviating from the oblivious equilibrium strategy $\tilde{\mu}$, and instead following an optimal (non-oblivious) best response, averaged over competitor states drawn from the invariant industry state distribution.

The next theorem provides two bounds on the approximation error. Recall that \tilde{s} is the expected state in equilibrium ($E[s]$). Let $a_x(y)$ be the expected discounted sum of an indicator of visits to state y for a firm starting at state x that uses strategy $\tilde{\mu}$.

Theorem 5.1. *Let Assumptions 2.1, 2.2, and 2.3 hold. Then, for any oblivious equilibrium $(\tilde{\mu}, \tilde{\lambda})$ and firm state $x \in \mathbb{N}$,*

$$(5.1) \quad E \left[\sup_{\mu' \in \mathcal{M}} V(x, s_t | \mu', \tilde{\mu}, \tilde{\lambda}) - V(x, s_t | \tilde{\mu}, \tilde{\lambda}) \right] \leq \frac{2}{1-\beta} E[\Delta |\pi_m|(s)],$$

where $\Delta |\pi|(s) = \max_{y \in \mathbb{N}} |\pi(y, s) - \pi(y, \tilde{s})|$, and

$$(5.2) \quad E \left[\sup_{\mu' \in \mathcal{M}} V(x, s_t | \mu', \tilde{\mu}, \tilde{\lambda}) - V(x, s_t | \tilde{\mu}, \tilde{\lambda}) \right] \leq \frac{1}{1-\beta} E[\Delta \pi(s)] + \sum_{y \in \mathbb{N}} a_x(y) (\pi(y, \tilde{s}) - E[\pi(y, s)]),$$

where $\Delta \pi(s) = \max_{y \in \mathbb{N}} (\pi(y, s) - \pi(y, \tilde{s}))$.

The derivation of these bounds can be found in the appendix. It is worth mentioning that the result can be generalized a great deal. In particular, most of the prior assumptions can be dropped; for instance, most alternative entry processes will not change the result. In addition, in the appendix we provide a version of the bounds for a model that incorporates aggregate shocks.

The first bound is simpler so we will use it to build the intuition behind the bounds. According to the first bound, if one period profits when the state is sampled from the invariant distribution are close to one period profits evaluated at the expected state, then the oblivious equilibrium offers a close approximation.

6500, Solaris 2.8, with 16 GB RAM, 12 GB swap, and 18 GB tmp; and a SunEnterprise 5500, Solaris 8, with 4 GB RAM, 4 GB swap, and 8.7 GB tmp.

This is because, in that case, firms cannot improve their expected discounted profits much by unilaterally deviating from the oblivious strategy and keeping track of the industry state.

The second bound is more involved. We provide it because it is much tighter. Note that the right-hand-side of the second bound depends on the initial firm state x , whereas the right-hand-side of the first bound does not.

Both bounds can be easily estimated via simulation algorithms. Computing the bounds involves computing an expectation over the industry state s_t . Once the oblivious equilibrium has been computed, the industry state has a known distribution, namely, the product form of Poisson random variables with mean \tilde{s} (by Lemma 3.2). In particular, note that the bounds are not a function of the true MPE or even of the optimal non-oblivious best response strategy. Computing either of these strategies could require solving a high-dimensional dynamic program.

In our computational experiments, we also make use of a stronger bound that applies when the quality level of a firm can change by at most one unit per time period ($\bar{w} = 1$), a common feature of EP-type models.

Theorem 5.2. *Let Assumptions 2.1, 2.2, and 2.3 hold. Further, assume that the quality level of any firm can change by at most one unit per time period. Then, for any oblivious equilibrium $(\tilde{\mu}, \tilde{\lambda})$, firm state $x \in \mathbb{N}$, and $\bar{a} \in \mathbb{N}$,*

$$(5.3) \quad E \left[\sup_{\mu' \in \mathcal{M}} V(x, s | \mu', \tilde{\mu}, \tilde{\lambda}) - V(x, s | \tilde{\mu}, \tilde{\lambda}) \right] \leq \frac{1}{1 - \beta} [E[\Delta_A \pi_m(s)] + \beta^{\bar{a}-x} E[\Delta_{A^c} \pi_m(s)]] + \sum_{y \in \mathbb{N}} a_x(y) (\pi_m(y, \tilde{s}) - E[\pi_m(y, s)]) .$$

where $\Delta_A \pi_m(s) = \max_{y \in A} (\pi_m(y, s) - \pi_m(y, \tilde{s}))$, $A = \{0, \dots, \bar{a} - 1\}$ and $x \in A$.

The proof of this result is similar to that of Theorem 5.1 and is omitted. In general, $E[\pi_m(y, s)] - \pi_m(y, \tilde{s}) \geq 0$ and is increasing in y . Therefore, $E[\max_{y \in A} (\pi_m(y, s) - \pi_m(y, \tilde{s}))]$ increases as the set A increases, making the first term our previous bound (5.2) large. In this new bound (5.3), the term that involves $\max_{y \in A^c}$ is the larger one because it involves the maximum over the larger states. However, this is offset by the factor of $\beta^{\bar{a}-x}$. The term $\beta^{\bar{a}-x}$ appears because it takes at least $\bar{a} - x$ periods to transition from state x to any state in A^c . When using this bound in the computational experiments, we choose the set A to minimize the right-hand-side.

6 Computational Experiments

In this section we conduct some computational experiments to evaluate how oblivious equilibrium performs in practice, and also to demonstrate the range of applications that are possible using our methods. We begin with the model to be analyzed. The model is similar to Pakes and McGuire (1994). However, it differs in the entry and exit processes, in the demand system, and in that we do not consider an aggregate shock.

6.1 The Computational Model

SINGLE-PERIOD PROFIT FUNCTION. We consider an industry with differentiated products, where each firm's state variable represents the quality of its product. There are m consumers in the market. In period t , consumer j receives utility u_{ijt} from consuming the good produced by firm i given by:

$$u_{ijt} = \theta_1 \ln\left(\frac{x_{it}}{\psi} + 1\right) + \theta_2 \ln(Y - p_{it}) + \nu_{ijt}, \quad i \in S_t, \quad j = 1, \dots, m,$$

where Y is the consumer's income, p_{it} is the price of the good produced by firm i , and ψ is a scaling factor. ν_{ijt} are i.i.d. random variables distributed Gumbel that represent unobserved characteristics for each consumer-good pair. There is also an outside good that provides consumers zero utility. We assume consumers buy at most one product each period and that they choose the product that maximizes utility. Under these assumptions our demand system is a classical logit model.

Let $N(x_{it}, p_i) = \exp(\theta_1 \ln(x_{it} + 1) + \theta_2 \ln(Y - p_i))$. Then, the expected market share of each firm is given by:

$$\sigma(x_{it}, s_{-i,t}, p) = \frac{N(x_{it}, p_i)}{1 + \sum_{j \in S_t} N(x_{jt}, p_j)}, \quad \forall i \in S_t.$$

We assume that firms set prices in the spot market. If there is a constant marginal cost c , the Nash equilibrium of the pricing game satisfies the first-order conditions,

$$(6.1) \quad Y - p_i + \theta_2(p_i - c)(\sigma(x_{it}, s_{-i,t}, p) - 1) = 0, \quad \forall i \in S_t.$$

There is a unique Nash equilibrium in pure strategies, denoted p_i^* (Caplin and Nalebuff (1991)). Expected profits are given by:

$$\pi_m(x_{it}, s_{-i,t}) = m\sigma(x_{it}, s_{-i,t}, p^*)(p_i^* - c), \quad \forall i \in S_t.$$

SELL-OFF PRICE. ϕ_{it} are i.i.d. exponential random variables with mean K .

TRANSITION DYNAMICS. A firm's investment is successful with probability $\frac{au}{1+au}$, in which case the quality

of its product increases by one level. The firm’s product depreciates one quality level with probability δ , independently each period. Note that our model differs from Pakes and McGuire (1994) here because the depreciation shocks in our model are idiosyncratic. Combining the investment and depreciation processes, it follows that the transition probabilities for a firm in state x that invests ι are given by:

$$\mathcal{P} \left[x_{t+1}^i = y \mid x_t^i = x, \iota \right] = \begin{cases} \frac{(1-\delta)a\iota}{1+a\iota} & \text{if } y = x + 1 \\ \frac{(1-\delta)+\delta a\iota}{1+a\iota} & \text{if } y = x \\ \frac{\delta}{1+a\iota} & \text{if } y = x - 1 . \end{cases}$$

6.2 Numerical Results: Behavior of the Bound

Our first set of results investigate the behavior of the approximation error bound under several different model specifications. A wide range of parameters for our model could reasonably represent different real world industries of interest. In practice the parameters would either be estimated using data from a particular industry or chosen to reflect an industry under study. We begin by investigating a particular set of representative parameter values. Following Pakes and McGuire (1994) we fix $a = 3$ and $\delta = 0.7$. Additionally, we fix marginal cost at $c = 0.5$, income at $Y = 1$, $\theta_2 = 0.5$, and $\psi = 1$. The discount factor is $\beta = 0.95$. The entry cost is $\kappa = 35$ and the entry state is $x^e = 10$. The average sell-off value is $K = 10$. In this case, $\beta \cdot \bar{\phi} < \kappa$, so sell-off value by itself is not sufficient reason to enter the industry (assumption 2.3.2). Additionally, both sell-off values and entry costs are substantially larger than marginal costs, consistent with empirical evidence.

In our computational experiments we found that the most important parameter affecting the approximation error bounds was θ_1 , which determines the importance that consumers place on product quality. If θ_1 is small, the degree of vertical differentiation between products is small. This reduces the impact of changes in the industry state on profits, making the MPE strategies less sensitive to the industry state. Additionally, when θ_1 is small it turns out that the invariant distribution \tilde{s} is very “light-tailed”. Oblivious strategies work well in this case, and the approximation error bound is small. If θ_1 is large, we get the reverse implications and the approximation error bound is larger.

Based on these experiments, here we consider two different values of θ_1 and the investment cost d : (θ_1, d) : (0.1, 0.1) and (0.5, 0.5). The former (“Low”) is a situation where the level of vertical differentiation is low and it is inexpensive to invest to improve quality. The latter (“High”) is the opposite. As a point of comparison, if a firm increases its state from $x = 10$ to $x = 20$, its single-period profits increase by 7% and 40% respectively in the two cases (holding competitors constant).

For each set of parameters, we use the approximation error bound in Theorem 5.2 to compute an upper bound on the percentage error in the value function, $\frac{E[\sup_{\mu' \in \mathcal{M}} V(x^e, s | \mu', \tilde{\mu}, \tilde{\lambda}) - V(x^e, s | \tilde{\mu}, \tilde{\lambda})]}{E[V(x^e, s | \tilde{\mu}, \tilde{\lambda})]}$, where $(\tilde{\mu}, \tilde{\lambda})$ are the OE strategy and entry rate, respectively, and the expectations are taken with respect to s . We estimate the expectations using simulation.¹¹ We compute the previously mentioned percentage approximation error bound for different market sizes. As the market size increases, the expected number of firms increases and the approximation error bound decreases.

In Figure 1 we present the percentage approximation error bound as a function of the expected number of firms for the two levels of vertical differentiation (the two curves are obtained by varying the market size). For the low vertical differentiation case it takes around 60 firms to bring the bound down to 2%, and 250 firms to bring it below 1%. For the high case it takes around 250 firms to bring the bound to 5% and 1000 firms to bring it near 2%.

When the level of vertical differentiation is high, the number of firms required to have a good approximation is large, requiring hundreds and even thousands of firms. The approximation would be better if the industry state s were always close to its mean, \tilde{s} . One aspect of the model that interferes with this is the Poisson entry process, which leads to a large amount of variability in the number of firms inside the industry. Recall that we chose to model the entry process this way because it facilitated the asymptotic arguments in Section 4. However, the expressions for the approximation error bounds remain correct for a wide range of entry models. To investigate this issue further, as an alternative, we tried using an entry process where the number of entrants each period is deterministic, but still satisfies a zero profits condition.¹² This entry process implies a smaller variability in the number of firms.

Figure 2 presents the results with the new entry process. In the case of low vertical differentiation, the approximation error bound is less than 2% with just 30 firms and it is around 1% with 100 firms. When the level of vertical differentiation is high the approximation error bound is around 5% when there are 120 firms and around 2% for 700 firms.

Going one step further in reducing the variability of the industry dynamics, we tried shutting down entry and exit altogether and considered an industry with a fixed number of firms. This situation can be viewed as a particular case of the model presented in Section 2, with a constant sell-off value equal to zero and a very high entry cost. See Figure 3 for the results. For the low case the approximation error bound is less than

¹¹The expected value function is estimated with a relative precision of 1% and a confidence level of 99%. The bound is estimated with a relative precision of at most 10% and a confidence level of 99% (in cases where the bound is very small it is difficult to achieve better precision than this). Note that the percentage approximation error bound depends on the state x so for the purposes of this section we consider the percentage bound evaluated at the entry state.

¹²Note that the zero profits condition typically requires a fractional number of entrants to be satisfied exactly, so to accommodate this we instead randomized the number of entrants between the two neighboring integers. For example, if the equilibrium entry rate is 2.5, then the number of entrants is 2 or 3 with probability 0.5. Allowing for fractional numbers ensures existence of equilibrium.

0.5% with just 5 firms, while for the high case it is 5% for 5 firms, less than 3% with 40 firms, and less than 1% with 400 firms.

Most economic applications would involve from less than ten to several hundred firms. These results show that the approximation error bound may sometimes be small (<2%) in these cases, though this would depend on the model and parameter values for the industry under study.

6.3 Closeness to Markov Perfect Equilibrium

Having gained some insight into what features of the model lead to low values of the approximation error bound, the question arises as to what value of the error bounds is required to obtain a good approximation. To shed light on this issue we compare long-run statistics for the same industry primitives under oblivious equilibrium and MPE strategies. A major constraint on this exercise is that it requires the ability to actually compute the MPE. With the current algorithms we are able to compute MPE for industries with a maximum of five to ten firms — to keep computation manageable we use four firms here. We therefore limit our analysis to the case of a fixed number of firms (no entry and exit), because only for that case were the approximation error bounds small under oblivious strategies (with only four firms). We use the same parameter specifications as in the previous subsection. Because of computational constraints in computing the MPE, we also impose a maximum state that a firm can reach of $x_{max} = 15$, at which point investment is assumed to have no further effect. The market size is fixed, $m = 30$.¹³

Recall that under oblivious equilibrium strategies, the industry state is described by an ergodic Markov chain (Lemma 3.2). This is also true under MPE strategies (Ericson and Pakes (1995)). Therefore, both systems have a well defined invariant distribution that describes their long-run behavior. We compare the average values of several economic statistics of interest under the oblivious equilibrium and the MPE invariant distributions. The quantities compared are: average investment, average producer surplus, average consumer surplus, average share of the largest firm (C1), and average share of the largest two firms (C2). Table 1 reports these statistics for a wide range of parameters. We also report the maximum value (across all states) and weighted average value of the approximation error bound, as well as the maximum and weighted average of the actual benefit from deviating and keeping track of the industry state (the actual difference $\frac{E[\sup_{\mu' \in \mathcal{M}} V(x^e, s | \mu', \tilde{\mu}, \tilde{\lambda}) - V(x^e, s | \tilde{\mu}, \tilde{\lambda})]}{E[V(x^e, s | \tilde{\mu}, \tilde{\lambda})]}$). Note that the latter quantity should always be smaller than the approximation error bound. In the results below we concentrate on the maximum values of these quantities. From the computational experiments we conclude the following (see Table 1):

¹³The code used to compute MPE was generously provided by Uli Doraszelski.

1. When the error bound is less than 1% the long-run quantities estimated under oblivious equilibrium and MPE strategies are very close.
2. Performance of the approximation depends on the shape of the average industry state. When the bound is between 1-20% and the average industry state is symmetric (see Figure 4), the long-run quantities estimated under oblivious equilibrium and MPE strategies remain close. When the bound is above 1% and the average industry state is very skewed (see Figure 5), the long-run quantities can be quite different on a percentage basis (5% to 20%). However, as shown in Table 2, this large percentage difference partly reflects the fact that there is very little investment in these industries, so a small amount of error results in a large percentage error.
3. The approximation error bound is not tight. For a wide range of parameters the approximation error bound is as much as 10 to 20 times larger than the actual benefit from deviating.

The previous results suggest that MPE dynamics are well-approximated by oblivious equilibrium strategies when the approximation error bound is small (less than 1-2% and in some cases even up to 20 %). These results, together with those from Subsection 6.2, demonstrate that the oblivious equilibrium approximation significantly expands the range of applied problems that can be analyzed computationally.

6.4 Example: The Evolution of Industry Structure as Market Size Increases

In this subsection we analyze the structure of the industry as market size grows under different parameter specifications. We note that the intention of this subsection is to show the potential of our methods and to provoke interest in future research in this area, and not to make robust and conclusive statements.

An important question in industrial organization is to understand what features of an industry are most important to determining the industry's structure. Of particular interest (Sutton 1991) is whether a market becomes fragmented, or remains concentrated, as the market grows in size. Sutton (1991) and Shaked and Sutton (1987) suggest that the presence of "endogenous" sunk costs imposes a lower bound to concentration. Our model is a dynamic model of endogenous sunk costs because, while entrants pay the same cost of entry regardless of market size (an "exogenous" sunk cost), in larger markets they may need to invest large amounts to become as large as the incumbent firms (an "endogenous" sunk cost). In what follows we show that under slightly different parameterizations of our model we get two extremely different outcomes: a fragmented and a very concentrated industry. an approach is not

We use the same parameter specifications as in Subsection 6.2 with the only difference that now we consider $\psi = 7$ in the utility of the consumers. This specification makes it easier to obtain the two different

situations. We fix $d = 0.5$ and we compute the oblivious equilibrium for $\theta_1 = 0.9$ and $\theta_1 = 1.2$. In Figures 6 and 7 we present the average industry states obtained for different market sizes in the two cases. For $\theta_1 = 0.9$, as the market size grows the industry becomes fragmented. The number of firms grows but firms do not grow in size, so the market share for every firm converges to zero. When $\theta_1 = 1.2$ the entry rate decreases and firms grow larger as market size grows. In this second case, the industry state is heavy-tailed.¹⁴

These results have a clear intuition: when $\theta_1 > 1$ consumers care more about quality and this increases the returns to investment. What results is a quality race between firms, so there are a small number of firms in the industry that become ever larger with the market size. When $\theta_1 < 1$ returns to investment are lower and investment is not worthwhile above a certain point regardless of the market size. (Note that with further experimentation it is easy to show that for all $\theta_1 < 1$ the industry fragments, whereas for $\theta_1 > 1$ the industry remains concentrated as market size increases.)

These results contradict Sutton (1991) in one sense because Sutton (1991) sought predictions about market structure that are robust across a wide class of models. Here, different market structures result from the same model through an arbitrarily small change in a single parameter. An implication of this result is that the same industry might be observed with very different market structures across markets that are the same size and that have indistinguishable characteristics.

However, our results remain broadly consistent with Sutton's predictions because the parameter that determines market structure in our model pertains directly to the returns to investment. Consistent with the predictions of Sutton (1991), industries with higher returns to investment tend to be concentrated. Our results show, however, that this relationship can be very fragile.

7 Conclusions and Future Research

The goal of this paper has been to increase the set of applied economic problems that can be addressed using Ericson and Pakes (1995)-style dynamic models of imperfect competition. Due to the curse of dimensionality, existing dynamic programming methods have limited application of these models to industries with a small number of firms and a small number of states per firm. Even with accelerated methods (e.g., Pakes and McGuire (2001) and Doraszelski and Judd (2003)), it seems likely that it will never be possible to solve for

¹⁴To show the results more clearly, we forced the highest quality level to be 80. This leads the industry state to have a mass of firms near 80 for the larger market sizes. Without this arbitrary cutoff firms would grow even larger and it becomes difficult to solve for the oblivious equilibrium (because the expected lifespan of a firm tends to infinity and the industry is heavy-tailed) as well as to show the results graphically.

an MPE exactly in many problems of interest. As an alternative, we proposed a method for approximating MPE behavior using an oblivious equilibrium, where firms make decisions only based on their own state and the long run average industry state.

We began by showing that the approximation works well asymptotically, where asymptotics were taken in the market size. A sufficient condition for an oblivious equilibrium to well approximate a MPE asymptotically is that the sequence of industry states generated by the oblivious equilibria is “light-tailed” (as described by Assumption 4.2). This condition is also sufficient to establish that a model of the type introduced in Hopenhayn (1992) yields a good approximation of a finite industry. We also introduced a simple algorithm to compute an oblivious equilibrium. A nice feature of the algorithm is that there is no need to place a priori restrictions on the number of firms in the industry or the set of states that a firm can reach. As a result, computational considerations place very few constraints on model details.

To facilitate using oblivious equilibrium in practice, we derived approximation error bounds that indicate how good the approximation is in any particular problem under study. These approximation error bounds are quite general and thus can be used in a wide class of models. We believe them to be the first bounds of this type in this literature. Through computational experiments, we showed that oblivious equilibrium often yields a good approximation of MPE behavior for industries with a couple hundred firms, and sometimes even with just tens of firms.

Note that in our approximation we have considered very simple strategies that are functions only of a firm’s own state and the long run average industry state. The question naturally arises whether it may be possible to obtain better approximations using additional information, such as the total number of firms in the industry and/or the average state across firms. Solving for equilibria of this type would be more difficult than solving for oblivious equilibria, but is still likely to be computationally feasible. However, showing that such an approach would provide a good approximation is not a simple extension of our results. The approximation error bounds would be more difficult to obtain for this case because, even if the shocks to individual firms remained idiosyncratic, firms’ outcomes would be correlated through their strategies, which would be functions of commonly known variables. The asymptotic results would be similarly difficult to prove. This, and similar extensions, will be a subject of future research.

A Proofs and Mathematical Arguments

A.1 Proofs and Mathematical Arguments for Section 2

A.1.1 The Poisson Entry Model

Suppose there are K potential entrants at a given state. Each potential entrant enters if the setup cost κ is less than the present value of future cash flows upon entry. Let $v_K(i)$ be the present value of future cash flows for each entering firm if i of the K firms enter simultaneously. $v_K(i)$ is nonincreasing in i . One can then pose the problem faced by potential entrants as a game in which each entrant employs a mixed strategy and enters with some probability p_K . If we assume that every potential entrant employs the same strategy, the condition for a mixed strategy Nash equilibrium when $\kappa \in [v_K(K), v_K(1)]$ can be written as

$$(A.1) \quad \sum_{i=0}^{K-1} \binom{K-1}{i} p_K^i (1-p_K)^{K-1-i} v_K(i+1) - \kappa = 0,$$

which is solved by a unique $p_K \in [0, 1]$. If $\kappa < v_K(K)$, the equilibrium is a pure strategy with $p_K = 1$, whereas if $\eta > v_K(1)$, the equilibrium is given by $p_K = 0$. The following result, which we state without proof, establishes that our Poisson entry model can be viewed as a limiting case as the number of potential entrants K grows large.

Lemma A.1. *Let the following conditions hold:*

1. $v_K(i)$ is non-increasing in i , $\forall K$ and non-increasing in K , $\forall i$;
2. there exists a positive constant M such that $|v_K(i)| < M$, $\forall K$, $\forall i$;
3. there exists \bar{K} , such that $\forall K > \bar{K}$, $v_K(i)$ changes sign in $(0, K)$;
4. there exists a function $v(i)$ such that $\lim_{K \rightarrow \infty} \max_{i \leq \bar{i}} |v_K(i) - v(i)| = 0$, for all $\bar{i} \in \mathbb{N}$.

Then,

1. for each $K > \bar{K}$, Equation (A.1) has a unique solution $p_K^* \in (0, 1)$;
2. the limit $\lambda = \lim_{K \rightarrow \infty} K p_K^*$ exists, and if Y_K is a binomial random variable with parameters (K, p_K^*) and Z is a Poisson random variable with mean λ , then $Y_K \Rightarrow Z$.

The result states that if the number of potential entrants grows to infinity then the entry process converges to a Poisson random variable. Hence, Poisson entry can be understood as the result of a large population of potential entrants, each one entering the industry with a very small probability.

A.1.2 Bellman's Equation

We define a dynamic programming operator:

$$(T_{\mu,\lambda}V)(x, s) = \pi(x, s) + E \left[\max \left\{ \beta \phi_{it}, \sup_{\iota \geq 0} (-d\iota + \beta E_{\mu,\lambda} [V(x_{i,t+1}, s_{-i,t+1}) | x_{it} = x, s_{-i,t} = s, \iota_{it} = \iota]) \right\} \right],$$

for all $x \in \mathbb{N}$ and $s \in \mathcal{S}$.

To simplify the notation in this section we will let $V_{\mu,\lambda}^{\mu'} \equiv V(\cdot | \mu', \mu, \lambda)$.

Lemma A.2. *Let Assumptions ... hold. Then, for all $\mu \in \mathcal{M}$ and $\lambda \in \Lambda$, there exists $\mu^* \in \mathcal{M}$, such that:*

$$V_{\mu,\lambda}^{\mu^*} = \sup_{\mu' \in \mathcal{M}} V_{\mu,\lambda}^{\mu'} = T_{\mu,\lambda} V_{\mu,\lambda}^{\mu^*}.$$

Further, $V_{\mu,\lambda}^{\mu^*}$ is the unique fixed point of $T_{\mu,\lambda}$ within the class of bounded functions.

Proof. First note that, for some constant c , the value function $V_{\mu,\lambda}^{\mu^*}$ is bounded by $0 \leq V_{\mu,\lambda}^{\mu^*}(x, s) \leq \frac{c}{1-\beta} + \bar{\phi}, \forall x \in \mathbb{N}, s \in \mathcal{S}$. Therefore, it will never be optimal to invest more than a finite quantity, $\bar{\iota}$. Hence we can assume without loss of generality that $0 \leq \iota(x, s) \leq \bar{\iota}, \forall x \in \mathbb{N}, \forall s \in \mathcal{S}$. Additionally, $\beta \bar{\phi} \leq \rho(x, s) \leq \frac{c}{1-\beta} + \bar{\phi}, \forall x \in \mathbb{N}, \forall s \in \mathcal{S}$. Therefore the action space for each state is compact.

For a given state (x, s) , expected one period profits including investment and sell-off value can be written as:

$$\pi(x, s) - d\iota(x, s) \mathcal{P}[\phi < \rho(x, s)] + E[\phi | \phi \geq \rho(x, s)] \mathcal{P}[\phi \geq \rho(x, s)].$$

Note that $\pi(x, s) < c$, investment is bounded by the previous argument and ϕ has finite expectation. Therefore, expected one period profits including investment and the sell-off value are uniformly bounded for all states (x, s) . The result follows by Propositions 1.2.2 and 3.1.7 in Bertsekas (2001). \square

A.2 Proofs and Mathematical Arguments for Section 3

Lemma 3.2. *Let Assumptions 2.2 and 2.3 hold. Assume that firms follow a common oblivious strategy $\mu \in \tilde{\mathcal{M}}$, the expected entry rate is $\lambda \in \tilde{\Lambda}$, and the expected time that each firm spends in the industry is finite. Let $\{Z_x : x \in \mathbb{N}\}$ be a sequence of independent Poisson random variables with means $\{\tilde{s}_{\mu,\lambda}(x) : x \in \mathbb{N}\}$, and let Z be a Poisson random variable with mean $\sum_{x \in \mathbb{N}} \tilde{s}_{\mu,\lambda}(x)$. Then,*

- (a) $\{s_t : t \geq 0\}$ is an irreducible, aperiodic and positive recurrent Markov chain;
- (b) the invariant distribution of s_t is a product form of Poisson random variables;

(c) for all x , $s_t(x) \Rightarrow Z_x$;

(d) $n_t \Rightarrow Z$.

Proof. If every firm uses a strategy $\mu \in \tilde{\mathcal{M}}$ and entry is according to an entry rate function $\lambda \in \tilde{\Lambda}$, then $X \equiv \{s_t : t \geq 0\}$ is clearly an irreducible Markov chain. All states reach the state $\emptyset = \{0, 0, \dots\}$ with positive probability and all states can be reached from \emptyset as well. Now, take any state $s \in \mathcal{S}$. There is a positive probability that there is no entry, no exit and no firms' state transitions. Therefore, for all $s \in \mathcal{S}$, $\mathcal{P}[s_{t+1} = s | s_t = s] > 0$, so X is aperiodic. Moreover X is positive recurrent because the expected time to come back from state \emptyset to itself is finite (see Kleinrock (1975)).

Now, let us write:

$$(A.2) \quad s_t(x) = \sum_{\tau=0}^t \sum_{i=1}^{W_\tau} \mathbf{1}_{\{X_{i,t-\tau}=x\}},$$

where W_τ are i.i.d. Poisson random variables with mean λ , the first sum is taken over all periods previous to (and including) t , the second sum is taken over the firms that entered the industry in each period, and for each τ , $X_{i,t-\tau}$ are random variables that represent the state of firm i after $t - \tau$ periods inside the industry when using strategy μ . Since firms use oblivious strategy $\mu \in \tilde{\mathcal{M}}$ their state evolutions are independent, so $\mathbf{1}_{\{X_{i,t-\tau}=x\}}$ are i.i.d. It follows that $\sum_{i=1}^{W_\tau} \mathbf{1}_{\{X_{i,t-\tau}=x\}}$ is a filtered Poisson random variable, so it is a Poisson random variable. Thus $s_t(x)$, as a sum of independent Poisson random variables, is also Poisson. Given that the expected time a firm spends inside the industry is finite, using characteristic functions it is straightforward to show that $s_t(x) \Rightarrow Z_x, \forall x \in \mathbb{N}$. To show that $\{Z_x : x \in \mathbb{N}\}$ is a sequence of independent random variables note that by the filtering property of Poisson random variables (Durrett (1996)), for all t , $\{s_t(x) : x \in \mathbb{N}\}$ is a sequence of independent random variables.

By summing over $x \in \mathbb{N}$, we can show that $n_t \Rightarrow Z$. □

Lemma 3.3. *Let Assumptions 2.2 and 2.3 hold. Assume that firms follow a common oblivious strategy $\mu \in \tilde{\mathcal{M}}$, the expected entry rate is $\lambda \in \tilde{\Lambda}$, and the expected time that each firm spends in the industry is finite. Let $\{Y_n : n \in \mathbb{N}\}$ be a sequence of independent random variables, each distributed according to $\tilde{s}_{\mu,\lambda}(\cdot) / \sum_{x \in \mathbb{N}} \tilde{s}_{\mu,\lambda}(x)$. Then, for all $n \in \mathbb{N}$,*

$$\left(x_{(1)t}, \dots, x_{(n)t} \middle| n_t = n \right) \Rightarrow (Y_1, \dots, Y_n).$$

Proof. The proof is rather involved, but relies on a simple idea. For a Poisson process, conditional on n arrivals on an interval $[0, T]$, the unordered arrival times have the same distribution as n i.i.d. uniform

random variables in $[0, T]$.

Let us condition on $n_t = n$. $\{x_{(j)t} : j = 1, \dots, n\}$ are the random variables that represent the state of each of the n firms in the industry when they are sampled randomly. The expected time a firm spends inside the industry is finite, so the time a firm spends inside the industry is finite with probability one. Recall that a firm can increase its quality level by at most \bar{w} states each period. Therefore, $\forall \epsilon > 0$, there exists a state z , such that for all $j \in 1, \dots, n$, $\forall t$, $\mathcal{P}[x_{(j)t} > z] < \epsilon$. That is, the probability of observing a firm in a very large state is negligible, because firms “do not have time” to get that far. Hence,

$$\mathcal{P} \left[\bigcup_{j=1}^n \{x_{(j)t} > z\} \mid n_t = n \right] < \epsilon, \forall t.$$

Therefore, the sequence of random vectors $(x_{(1)t}, \dots, x_{(n)t} \mid n_t = n)$ is tight. Theorem 9.1 in Durrett (1996) states that $Z_n \Rightarrow Z$ if and only if for all sets A with $\mathcal{P}[Z \in \partial A] = 0$, $\lim_{n \rightarrow \infty} P[Z_n \in A] = P[Z \in A]$, where ∂A denotes the boundary of the set A . Using this result and the tightness previously established, it is enough to show that for all n , for all (z_1, \dots, z_n) :

$$(A.3) \quad \lim_{t \rightarrow \infty} \mathcal{P} \left[x_{(j)t} = z_j, j = 1, \dots, n \mid n_t = n \right] = \prod_{j=1}^n p(z_j),$$

where $p(\cdot)$ is the PMF $\tilde{s}_{\mu, \lambda}(\cdot) / \sum_{x \in \mathbb{N}} \tilde{s}_{\mu, \lambda}(x)$. Let \tilde{T}_j be the entry date of firm (j) and $T_j = t - \tilde{T}_j$ be its age. Then we can write:

$$(A.4) \quad \begin{aligned} \mathcal{P} \left[x_{(j)t} = z_j, j = 1, \dots, n \mid n_t = n \right] &= \sum_{\substack{0 \leq t_1 < \infty, \dots, \\ 0 \leq t_n < \infty}} \mathcal{P} [x_{(j)t} = z_j, j = 1, \dots, n \mid T_1 = t_1, \dots, T_n = t_n, n_t = n] \cdot \\ &\quad \mathcal{P} [T_1 = t_1, \dots, T_n = t_n \mid n_t = n] \\ &= \sum_{\substack{0 \leq t_1 < \infty, \dots, \\ 0 \leq t_n < \infty}} \prod_{j=1}^n \mathcal{P} [x_{(j)t} = z_j \mid T_j = t_j] \cdot \\ &\quad \mathcal{P} [T_1 = t_1, \dots, T_n = t_n \mid n_t = n], \end{aligned}$$

The last equation follows because the evolution of firms is independent across firms. Note that if any \tilde{t}_j has

a value greater than t , then $\mathcal{P} [T_1 = t_1, \dots, T_n = t_n | n_t = n] = 0$. We can write

$$\begin{aligned}
\mathcal{P} [x_{(j)t} = z_j | T_j = t_j] &= \frac{\mathcal{P} [x_{(j)t} = z_j, T_j = t_j]}{\mathcal{P} [T_j = t_j]} \\
&= \frac{\mathcal{P} [T_j = t_j, X_{j,t_j} = z_j]}{\mathcal{P} [T_j = t_j]} \\
&= \frac{\mathcal{P} [T_j = t_j] \mathcal{P} [X_{j,t_j} = z_j]}{\mathcal{P} [T_j = t_j]} \\
&= \mathcal{P} [X_{j,t_j} = z_j],
\end{aligned}
\tag{A.5}$$

where, similarly to above, X_{j,t_j} is a random variable that represents a firm's state after t_j periods conditional on having stayed in the industry. Note that for all k , $\{X_{j,k} : j \geq 1\}$ is i.i.d. The second to last equation follows because the evolution of a firm is independent of its entry time.

Now we show that

$$\lim_{t \rightarrow \infty} \mathcal{P} [T_1 = t_1, \dots, T_n = t_n | n_t = n] = \prod_{j=1}^n u[t_j],$$

for some probability mass function u . This equation can be derived directly. However, it is easier to invoke the relationship between n_t and a Poisson process and show that the equation holds using known results for Poisson processes.

Similarly to equation (A.2), we can write:

$$n_t = \sum_{\tau=0}^t \sum_{i=1}^{W_\tau} \mathbf{1}_{\{A_{i,t-\tau}=1\}},$$

where, for each τ , $A_{i,t-\tau}$ are i.i.d. Bernoulli random variables that equal one if the firm is still in the industry after $t - \tau$ periods when using strategy μ . Since $A_{i,t-\tau}$ are i.i.d., $n_{t,\tau} \equiv \sum_{i=1}^{Z_\tau} \mathbf{1}_{\{A_{i,t-\tau}=1\}}$ is a filtered Poisson random variable, and is therefore Poisson. Call its mean $\alpha_{t,\tau}$. It follows that n_t is a sum of independent Poisson random variables, so it is Poisson with mean $\sum_{\tau=0}^t \alpha_{t,\tau}$.

Consider $\{N(t) : t \geq 0\}$, a homogeneous Poisson process on the real line with rate 1. Note that $N(t)$ and n_t are equivalent in the sense that we can construct n_t using the process $\{N(s) : 0 \leq s \leq \sum_{\tau=0}^t \alpha_{t,\tau}\}$ in the following way: for each $0 \leq \tau \leq t$, consider the events in the interval $[\alpha_{t,\tau-1}, \alpha_{t,\tau-1} + \alpha_{t,\tau}]$, where $\alpha_{t,-1} = 0$. Assign those events to firms that are still in the industry at time t and entered at period τ , $n_{t,\tau}$. Note that: i) $N(\alpha_{t,\tau})$ and $n_{t,\tau}$ are both Poisson with mean $\alpha_{t,\tau}$; ii) $N(t)$ has independent increments and $n_{t,\tau}$ are independent, for all τ ; iii) $N(\sum_{\tau=0}^t \alpha_{t,\tau})$ and n_t are both sum of independent Poisson random

variables with mean $\sum_{\tau=0}^t \alpha_{t,\tau}$.

Suppose $N(\sum_{\tau=0}^t \alpha_{t,\tau}) = n$. Conditional on that event, the unordered arrival times of $N(t)$ have the same distribution as n i.i.d. uniform random variables in $[0, \sum_{\tau=0}^t \alpha_{t,\tau}]$ (Durrett (1996)). By the equivalence argument described above, conditional on $n_t = n$, the unordered arrival times of the n firms are i.i.d. discrete random variables with PMF:

$$v_t(\tau) = \frac{\alpha_{t,\tau}}{\sum_{j=0}^t \alpha_{t,j}}, 0 \leq \tau \leq t.$$

Recall that $\alpha_{t,\tau}$ is the expected number of firms that entered at time τ and are still inside the industry at time t . Since the entry rate is the same every period and every firm uses the same strategy which is independent of every other firm, $\alpha_{t,\tau} = \tilde{\alpha}_k$, where $\tilde{\alpha}_k$ is the expected number of firms that entered the industry at time s , for any s , and are still inside the industry at time $s + k$. This suggests making a change of variable and defining:

$$u_t(k) = \frac{\tilde{\alpha}_k}{\sum_{j=0}^t \tilde{\alpha}_j}, 0 \leq k \leq t.$$

$u_t(k)$ is the probability a random sampled firm from the industry at time t entered k periods ago, conditional on $n_t = n$. Taking the limit as t tends to infinity, we get that:

$$\lim_{t \rightarrow \infty} u_t(k) \equiv u(k) = \frac{\tilde{\alpha}_k}{\sum_{j=0}^{\infty} \tilde{\alpha}_j}, 0 \leq k < \infty,$$

provided that $\lim_{t \rightarrow \infty} E[n_t] = \sum_{j=0}^{\infty} \tilde{\alpha}_j < \infty$, which is true by assumption. $u(k)$ is the probability a random sampled firm from the industry at time $t \gg 0$ entered k periods before the sampling period. Therefore

$$\lim_{t \rightarrow \infty} \mathcal{P}[T_1 = t_1, \dots, T_n = t_n \mid n_t = n] = \prod_{j=1}^n u[t_j].$$

Replacing the previous equation and equation (A.5) in equation (A.4) we obtain:

$$\lim_{t \rightarrow \infty} \mathcal{P}[x_{(j)t} = z_j, j = 1, \dots, n \mid n_t = n] = \prod_{j=1}^n \sum_{0 \leq t < \infty} \mathcal{P}[X_{j,t} = z_j] u(t)$$

where the interchange between the infinite sum and the limit follows by dominated convergence. The sum yields the PMF $p(\cdot)$. The previous equation proves that Y_1, \dots, Y_n are i.i.d. with PMF $p(\cdot)$ which is independent of n .

Consider $t \gg 0$. More formally, suppose that s_0 is sampled from the invariant distribution, which is well defined by Lemma 3.2. In this case, s_t is a stationary process; s_t is distributed according to the invariant

distribution for all $t \geq 0$.

$$\tilde{s}_{\mu,\lambda}(x) = E[s_t(x)] = E \left[\sum_{j=1}^{n_t} \mathbf{1}_{\{x_{(j)t}=x\}} \right].$$

Conditioning on n_t , and considering that we already proved that $\{x_{(j)t} : j = 1, \dots, n\}$ are i.i.d. with PMF $p(\cdot)$ which is independent of n , we get that $p(\cdot) = \tilde{s}_{\mu,\lambda}(\cdot) / \sum_{x \in \mathbb{N}} \tilde{s}_{\mu,\lambda}(x)$. \square

A.3 Proofs and Mathematical Arguments for Section 4

A.3.1 Preliminary Lemmas

We will use the following lemma, which follows immediately from Assumption 4.2. We omit the proof.

First, we define $\|f\|_{1,g} = \sum_x |f(x)|g(x)$.

Lemma A.3. *Let $\{\tilde{f}^{(m)} | m \in \mathbb{N}\}$ be a sequence of distributions satisfying Assumption 4.2. Then,*

$$\lim_{z \rightarrow \infty} \sup_m \inf_{\hat{f} \in \mathcal{S}_{1,z}} \|\tilde{f}^{(m)} - \hat{f}\|_{1,g} = 0.$$

Lemma A.4. *Under Assumptions 2.1, 2.2, and 2.3,*

$$\sup_m \tilde{V}(x | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) < \infty.$$

Proof. We will assume that $x \geq x^e$; the case of $x < x^e$ is trivial. Assume for contradiction that $\sup_m \tilde{V}(x | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) = \infty$. We will argue that this contradicts the zero profit condition for entering firms. If a firm invests $\iota > 0$, there is a probability $p(\iota) > 0$ that the firm will increase its quality level by at least one unit. Let τ be the time a firm takes to transition from state x^e to state x . If a firm invests $\iota > 0$ in each period, by a geometric trials argument, $E[\tau] < \infty$. Therefore, there exists an investment strategy for which the expected time and cost to transition from x^e to x are uniformly bounded above over m . It follows that $\sup_m \tilde{V}(x^e | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) = \infty$. This contradicts the zero profit condition. \square

Let $\ell_{1,g} = \{h : \mathbb{N} \mapsto \mathfrak{R} \mid \|h\|_{1,g} < \infty\}$. With some abuse of notation, let $\mathcal{S}_{1,g} = \mathcal{S}_1 \cap \ell_{1,g}$.

Lemma A.5. *Let Assumptions 2.1, 2.2, and 2.3, hold. Then, for all x ,*

$$\sup_m \sup_{\mu \in \mathcal{M}} E_\mu \left[\sum_{k=t}^{\infty} \beta^{k-t} \sup_{f \in \mathcal{S}_1} \pi_m(x_{ik}, f, \tilde{n}^{(m)}) \mid x_{it} = x \right] < \infty.$$

Proof. Let all expectations in this proof be conditioned on $x_{it} = x$. The industry state, and therefore $\pi_m(\cdot, \tilde{s}^{(m)})$, is not influenced by firm i 's investment decisions. Hence,

$$(A.6) \quad \sup_m \sup_{\mu \in \mathcal{M}} E_\mu \left[\sum_{k=t}^{\infty} \beta^{k-t} \pi_m(x_{ik}, \tilde{s}^{(m)}) \right] \leq \sup_m \tilde{V}(x | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) < \infty,$$

where the final relation follows from Lemma A.4.

By Assumption 2.1.4, for any $y \in \mathbb{N}$, $f, f' \in \mathcal{S}_{1,g}$, $n \in \mathfrak{R}_+$, and $m \in \mathbb{N}$,

$$\begin{aligned} |\pi_m(y, f, n) - \pi_m(y, f', n)| &= \left| \int_{\gamma=0}^1 \sum_{x \in \mathbb{N}} \left(\frac{f(x) - f'(x)}{\|f - f'\|_2} \right) \left(\frac{\partial \ln \pi_m(y, f' + \gamma(f - f'), n)}{\partial f(x)} \right) d\gamma \right| \\ &\leq \int_{\gamma=0}^1 \sum_{x \in \mathbb{N}} \left(\frac{|f(x) - f'(x)|}{\|f - f'\|_2} \right) g(x) d\gamma \\ &= \frac{\|f - f'\|_{1,g}}{\|f - f'\|_2} \\ &< \infty. \end{aligned}$$

Letting $\underline{f} = (1, 0, 0, \dots)$, it follows that $\sup_{x \in \mathbb{N}, m \in \mathbb{N}} |\pi_m(x, \underline{f}, \tilde{n}^{(m)}) - \pi_m(x, \tilde{s}^{(m)})| \equiv C < \infty$. By Assumption 2.1.2, for all $m \in \mathfrak{R}_+$, $x \in \mathbb{N}$, and $f \in \mathcal{S}_1$, $\pi_m(x, f, \tilde{n}^{(m)}) \leq \pi_m(x, \underline{f}, \tilde{n}^{(m)}) \leq \pi_m(x, \tilde{s}^{(m)}) + C$. The result then follows from (A.6). \square

A.4 Proof of Theorem 4.2

Theorem 4.2. *Under Assumptions 2.1, 2.2, 2.3, 4.1.2, and 4.2, $\tilde{n}^{(m)} = \Omega(m)$.*

Proof. Assume for contradiction that $\liminf_m \frac{\tilde{n}^{(m)}}{m} = 0$. Then there exists an increasing sequence m_k such that $\lim_k \frac{\tilde{n}^{(m_k)}}{m_k} = 0$, and by Assumption 4.1.2, for all $x, z \in \mathbb{N}$ with $x > z$, and $f \in \mathcal{S}_{1,z}$, $\lim_{k \rightarrow \infty} \pi_{m_k}(x, f, \tilde{n}^{(m_k)}) = \infty$.

Note that for any $h : \mathcal{S} \mapsto \mathfrak{R}_+$, $y \in \mathbb{N}$, $f \in \mathcal{S}_{1,g}$, and $n \in \mathbb{N}$,

$$\sum_{x \in \mathbb{N}} h(x) \left| \frac{\partial \pi_m(y, f, n)}{\partial f(x)} \right| \leq \|h\|_{1,g}.$$

It follows from Assumptions 2.1.4 and 4.2 that

$$\lim_{z \rightarrow \infty} \sup_m \inf_{\hat{f} \in \mathcal{S}_{1,z}} \sup_x \left| \ln \pi_m(x, \tilde{f}^{(m)}, \tilde{n}^{(m)}) - \ln \pi_m(x, \hat{f}, \tilde{n}^{(m)}) \right| = 0.$$

It follows that

$$\lim_{m \rightarrow \infty} \sup_m \inf_{\hat{f} \in \mathcal{S}_{1,z}} \sup_x \left| \frac{\pi_m(x, \tilde{f}^{(m)}, \tilde{n}^{(m)})}{\pi_m(x, \hat{f}^{(m)}, \tilde{n}^{(m)})} - 1 \right| = 0,$$

and therefore,

$$\pi_m(x, \tilde{f}^{(m)}, \tilde{n}^{(m)}) \geq (1 - \epsilon) \pi_m(x, \hat{f}^{(m)}, \tilde{n}^{(m)}),$$

for all x and m . This implies that for all $x > z$, $\lim_{k \rightarrow \infty} \pi_{m_k}(x, \tilde{f}^{(m_k)}, \tilde{n}^{(m_k)}) = \infty$, which contradicts Lemma A.4. It follows that $\tilde{n}^{(m)} = \Omega(m)$. \square

A.4.1 Proof of Theorem 4.4

Lemma A.6. *Let Assumptions 2.1, 2.2, 2.3, 4.1.2, and 4.2, hold. Then, for any $\delta > 0$,*

$$\mathcal{P} \left[\left| \frac{n_t^{(m)}}{\tilde{n}^{(m)}} - 1 \right| \geq \delta \right] \leq e^{-\Omega(m)}.$$

Proof. By a simple analysis of the Poisson distribution, it is easy to show that if n is a Poisson random variable with mean \tilde{n} ,

$$\mathcal{P} \left[\left| \frac{n}{\tilde{n}} - 1 \right| \geq \delta \right] \leq e^{-\Theta(\tilde{n})}.$$

By Lemma 3.2, $n_t^{(m)}$ is a Poisson random variable with mean $\tilde{n}^{(m)}$. By Theorem 4.2, $\tilde{n}^{(m)} = \Omega(m)$. The result follows. \square

Theorem 4.3. *Let Assumptions 2.1, 2.2, 2.3, 4.1.2, and 4.2, hold. Then, as m grows, $n_t^{(m)}/\tilde{n}^{(m)} \rightarrow_p 1$ and $\|f_t^{(m)} - \tilde{f}^{(m)}\|_{1,g} \rightarrow_p 0$.*

Proof. Convergence of $n_t^{(m)}/\tilde{n}^{(m)}$ follows from Lemma A.6. To complete the proof, we will establish convergence of $\|f_t^{(m)} - \tilde{f}^{(m)}\|_{1,g}$. Note that for any $z \in \mathbb{N}$,

$$\begin{aligned} \|f_t^{(m)} - \tilde{f}^{(m)}\|_{1,g} &\leq z \max_{x \leq z} g(x) |f_t^{(m)}(x) - \tilde{f}^{(m)}(x)| + \sum_{x > z} g(x) f_t^{(m)}(x) + \sum_{x > z} g(x) \tilde{f}^{(m)}(x) \\ &\equiv A_z^{(m)} + B_z^{(m)} + C_z^{(m)}. \end{aligned}$$

We will show that for any z , $A_z^{(m)}$ converges in probability to zero, that for any $\delta > 0$, for sufficiently large z , $\lim_{m \rightarrow \infty} \mathcal{P}[C_z^{(m)} \geq \delta] = 0$, and that for any $\delta > 0$ and $\epsilon > 0$, for sufficiently large z , $\limsup_{m \rightarrow \infty} \mathcal{P}[B_z^{(m)} \geq \delta] \leq \epsilon/\delta$. The assertion that $\|f_t^{(m)} - \tilde{f}^{(m)}\|_{1,g} \rightarrow_p 0$ follows from these facts.

By Lemma 3.3, for any x , $(f_t^{(m)}(x) \mid n_t^{(m)} = n)$ is distributed as the empirical mean of n i.i.d. Bernoulli samples with expectation $\tilde{f}^{(m)}(x)$. It follows that for any x , $(|f_t^{(m)}(x) - \tilde{f}^{(m)}(x)| \mid n_t = n)$ converges in

probability to zero uniformly over m as n grows. By Lemma 4.2 and the fact that $n_t^{(m)}/\tilde{n}^{(m)}$ converges in probability to 1, for any n , $\lim_{m \rightarrow \infty} \mathcal{P}[n_t^{(m)} \leq n] = 0$. It follows that for any z , $A_z^{(m)}$ converges in probability to zero.

By Lemma A.3, for any $\delta > 0$, for sufficiently large z , $\limsup_{m \rightarrow \infty} C_z^{(m)} \leq \delta$, and therefore, $\lim_{m \rightarrow \infty} \mathcal{P}[C_z^{(m)} \geq \delta] = 0$. By Tonelli's Theorem, $E[B_z^{(m)}] = C_z^{(m)}$. Invoking Markov's inequality, for any $\delta > 0$ and $\epsilon > 0$, for sufficiently large z , $\limsup_{m \rightarrow \infty} \mathcal{P}[B_z^{(m)} \geq \delta] \leq \epsilon/\delta$. The result follows. \square

The following technical lemma follows immediately from Assumption 4.1.3. We omit the proof.

Lemma A.7. *Let Assumptions 2.1.3 and 4.1.3 hold. Then, for all $\epsilon > 0$ there exists $\delta > 0$ such that for all $n, \hat{n} \in \mathfrak{R}_+$ satisfying $|n/\hat{n} - 1| < \delta$,*

$$\sup_{x, f, m} \left| \frac{\pi_m(x, f, n) - \pi_m(x, f, \hat{n})}{\pi_m(x, f, \hat{n})} \right| < \epsilon.$$

Lemma A.8. *Let Assumptions 2.1, 2.2, 2.3, 4.1, and 4.2 hold. Then, for all sequences $\{\mu^{(m)} \in \mathcal{M}\}$,*

$$\lim_{m \rightarrow \infty} E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} |\pi_m(x_{ik}, s_{-i,k}^{(m)}) - \pi_m(x_{ik}, f_{-i,k}^{(m)}, \tilde{n}^{(m)})| \mid x_{it} = x, s_{-i,t}^{(m)} \sim q^{(m)} \right] = 0.$$

Proof. For the purpose of this proof, we will assume that all expectations are conditioned on $x_{it} = x$ and $s_{-i,t}^{(m)} \sim q^{(m)}$. Let $\Delta_{it}^{(m)} = |\pi_m(x_{it}, s_{-i,t}^{(m)}) - \pi_m(x_{it}, f_{-i,t}^{(m)}, \tilde{n}^{(m)})|$. Fix $\epsilon > 0$ and let $\delta > 0$ satisfy the assertion of Lemma A.7. Let Z denote the event $|n_k^{(m)}/\tilde{n}^{(m)} - 1| > \delta$. Applying Tonelli's Theorem, we obtain

$$\begin{aligned} E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} \Delta_{ik}^{(m)} \right] &= \sum_{k=t}^{\tau_i} \beta^{k-t} E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\Delta_{ik}^{(m)} \right] \\ &= \sum_{k=t}^{\tau_i} \beta^{k-t} \left(E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\Delta_{ik}^{(m)} \mathbf{1}_{\neg Z} \right] + E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\Delta_{ik}^{(m)} \mathbf{1}_Z \right] \right) \\ &\leq \sum_{k=t}^{\tau_i} \beta^{k-t} \left(\epsilon E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\pi_m(x_{ik}, f_{-i,k}^{(m)}, \tilde{n}^{(m)}) \right] + O(m) \mathcal{P}[Z] \right) \\ &\leq \epsilon E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} \sup_{f \in \mathcal{S}_1} \pi_m(x_{ik}, f, \tilde{n}^{(m)}) \right] + \frac{O(m) \mathcal{P}[Z]}{1 - \beta}, \end{aligned}$$

where the first inequality follows from Lemma A.7 and Assumption 4.1.1. Since ϵ is arbitrary, the expected sum is finite (by Lemma A.5), and $\mathcal{P}[Z] \leq e^{-\Omega(m)}$ (by Lemma A.6), the result follows. \square

The following technical lemma follows immediately from assumptions on the profit function. We omit

the proof.

Lemma A.9. *Let Assumptions 2.1.3 and 2.1.4 hold. Then, for all $\epsilon > 0$ there exists $\delta > 0$ such that for $f, \hat{f} \in \mathcal{S}_{1,g}$ satisfying $\|f - \hat{f}\|_{1,g} < \delta$,*

$$\sup_{x,f,m} \left| \frac{\pi_m(x, f, n) - \pi_m(x, \hat{f}, n)}{\pi_m(x, \hat{f}, n)} \right| < \epsilon.$$

Lemma A.10. *Let Assumptions 2.1, 2.2, 2.3, 4.1, and 4.2 hold. Then, for all sequences $\{\mu^{(m)} \in \mathcal{M}\}$,*

$$\lim_{m \rightarrow \infty} E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} |\pi_m(x_{ik}, f_{-i,k}^{(m)}, \tilde{n}^{(m)}) - \pi_m(x_{ik}, \tilde{s}^{(m)})| \mid x_{it} = x, s_{-i,t}^{(m)} \sim q^{(m)} \right] = 0.$$

Proof. For the purpose of this proof, we will assume that all expectations are conditioned on $x_{it} = x$ and $s_{-i,t}^{(m)} \sim q^{(m)}$. Let $\Delta_{it}^{(m)} = |\pi_m(x_{it}, f_{-i,t}^{(m)}, \tilde{n}^{(m)}) - \pi_m(x_{it}, \tilde{s}^{(m)})|$. Fix $\epsilon > 0$ and let δ satisfy the assertion of Lemma A.9. Let Z denote the event $\|f_{-i,k}^{(m)} - \tilde{f}^{(m)}\|_{1,g} > \delta$. Applying Tonelli's Theorem, we obtain

$$\begin{aligned} E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} \Delta_{ik}^{(m)} \right] &= \sum_{k=t}^{\tau_i} \beta^{k-t} E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\Delta_{ik}^{(m)} \right] \\ &= \sum_{k=t}^{\tau_i} \beta^{k-t} E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\Delta_{ik}^{(m)} \mathbf{1}_{\neg Z} \right] + E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\Delta_{ik}^{(m)} \mathbf{1}_Z \right] \\ &\leq \frac{\epsilon}{1-\beta} + \sum_{k=t}^{\tau_i} \beta^{k-t} E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\Delta_{ik}^{(m)} \mathbf{1}_Z \right], \end{aligned}$$

where the inequality follows from Lemma A.9.

Note that $\Delta_{ik}^{(m)} \leq 2 \sup_{f \in \mathcal{S}_1} \pi_m(x_{ik}, f, \tilde{n}^{(m)})$. Hence,

$$\begin{aligned} \sum_{k=t}^{\tau_i} \beta^{k-t} E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\Delta_{ik}^{(m)} \mathbf{1}_Z \right] &\leq \sum_{k=t}^{\tau_i} \beta^{k-t} E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[2 \sup_{f \in \mathcal{S}_1} \pi_m(x_{ik}, f, \tilde{n}^{(m)}) \mathbf{1}_{Z_\delta} \right] \\ &= 2\mathcal{P}[Z] E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} \sup_{f \in \mathcal{S}_1} \pi_m(x_{ik}, f, \tilde{n}^{(m)}) \right], \end{aligned}$$

where the final expression follows from the fact that $f_{ik}^{(m)}$ evolves independently from $x_{ik}^{(m)}$. Since ϵ is arbitrary, $\mathcal{P}[Z] \rightarrow 0$ (by Theorem 4.3), and the expected sum is finite (by Lemma A.5), the result follows. \square

Theorem 4.4. *Under Assumptions 2.1, 2.2, 2.3, 4.1, and 4.2, the sequence $\{\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}\}$ of oblivious equilibria possesses the AME property.*

Proof. Let $\mu^{*(m)}$ be an optimal (non-oblivious) best response to $(\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)})$; in particular,

$$V^{(m)}(x, s | \mu^{*(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) = \sup_{\mu \in \mathcal{M}} V^{(m)}(x, s | \mu, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}).$$

Let

$$\hat{V}^{(m)}(x, s) = V^{(m)}(x, s | \mu^{*(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) - V^{(m)}(x, s | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) \geq 0.$$

The AME property, which we set out to establish, asserts that for all $x \in \mathbb{N}$, $\lim_{m \rightarrow \infty} E_{\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}}[\hat{V}^{(m)}(x, s_t^{(m)})] = 0$.

For any m , because $\tilde{\mu}^{(m)}$ and $\tilde{\lambda}^{(m)}$ attain an oblivious equilibrium,

$$\tilde{V}^{(m)}(x | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) = \sup_{\tilde{\mu} \in \tilde{\mathcal{M}}} \tilde{V}^{(m)}(x | \tilde{\mu}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}).$$

It follows that

$$\begin{aligned} \hat{V}^{(m)}(x, s) &= \left(V^{(m)}(x, s | \mu^{*(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) - \tilde{V}^{(m)}(x | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) \right) \\ &\quad + \left(\tilde{V}^{(m)}(x | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) - V^{(m)}(x, s | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) \right) \\ &\leq \left(V^{(m)}(x, s | \mu^{*(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) - \tilde{V}^{(m)}(x | \mu^{*(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) \right) \\ &\quad + \left(\tilde{V}^{(m)}(x | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) - V^{(m)}(x, s | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) \right) \\ &\equiv A^{(m)}(x, s) + B^{(m)}(x, s). \end{aligned}$$

To complete the proof, we will establish that $E_{\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}}[A^{(m)}(x, s_t^{(m)})]$ and $E_{\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}}[B^{(m)}(x, s_t^{(m)})]$ converge to zero.

Let τ_i be the time at which firm i exists, and let $\Delta_{it}^{(m)} = |\pi_m(x_{it}, s_{-i,t}^{(m)}) - \pi_m(x_{it}, \tilde{s}^{(m)})|$. It is easy to see that

$$\begin{aligned} A^{(m)}(x, s) &\leq E_{\mu^{*(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} \Delta_{ik}^{(m)} \mid x_{it} = x, s_{-i,t}^{(m)} = s \right] \\ B^{(m)}(x, s) &\leq E_{\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} \Delta_{ik}^{(m)} \mid x_{it} = x, s_{-i,t}^{(m)} = s \right], \end{aligned}$$

and letting $q^{(m)}$ be the invariant distribution of $s_t^{(m)}$ with the oblivious strategy $\tilde{\mu}^{(m)}$ and the oblivious entry

rate $\tilde{\lambda}^{(m)}$,

$$\begin{aligned} E_{\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[A^{(m)}(x, s_t^{(m)}) \right] &\leq E_{\mu^{*(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} \Delta_{ik}^{(m)} \mid x_{it} = x, s_{-i,t}^{(m)} \sim q^{(m)} \right] \\ E_{\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[B^{(m)}(x, s_t^{(m)}) \right] &\leq E_{\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} \Delta_{ik}^{(m)} \mid x_{it} = x, s_{-i,t}^{(m)} \sim q^{(m)} \right]. \end{aligned}$$

By the triangle inequality,

$$\Delta_{ik}^{(m)} \leq |\pi_m(x_{ik}, s_{-i,k}) - \pi_m(x_{ik}, f_{-i,k}^{(m)}, \tilde{n}^{(m)})| + |\pi_m(x_{ik}, f_{ik}^{(m)}, \tilde{n}^{(m)}) - \pi_m(x_{ik}, \tilde{s}^{(m)})|.$$

The result therefore follows from Lemmas A.8 and A.10. \square

A.4.2 Example: Logit Demand System with Price Competition

We consider a single period profit function like the one introduced in Subsection 6.1. It is direct to check that Assumptions 2.1, 4.1.1, and 4.1.2 are satisfied. Now, we show that Assumption 4.1.3 is satisfied. To simplify notation we assume $\psi = 1$. Now,

$$(A.7) \quad \frac{d \ln \pi_m(x, f, n)}{d \ln n} = \frac{\partial \ln \pi_m(x, f, n)}{\partial \ln n} + \frac{\partial \ln \pi_m(x, f, n)}{\partial p_x} \frac{\partial p_x}{\partial \ln n} + \sum_{i \in S} \frac{\partial \ln \pi_m(x, f, n)}{\partial p_i} \frac{\partial p_i}{\partial \ln n},$$

where S is the set of firms in state $s = fn$, and p_x is the price charged by the firm in state x . The first term takes into account the direct change in profits due to the change on the number of firms keeping prices fixed. The second and third terms consider the change in profits implied by the change of prices. Now,

$$\frac{\partial \ln \pi_m(x, f, n)}{\partial \ln n} = - \frac{n \sum_{z \in \mathbb{N}} f(z) (1+z)^{\delta_1} (Y - p_z)^{\delta_2}}{1 + n \sum_{z \in \mathbb{N}} f(z) (1+z)^{\delta_1} (Y - p_z)^{\delta_2} + (1+x)^{\delta_1} (Y - p_x)^{\delta_2}}.$$

Therefore, $\sup_{m \in \mathbb{N}, x \in \mathbb{N}, f \in \mathcal{S}_1, n > 0} \left| \frac{\partial \ln \pi_m(x, f, n)}{\partial \ln n} \right| = 1$. Similarly, it is possible to show that if $\delta_2 \leq \frac{1}{2}$, $\sup_{m \in \mathbb{N}, x \in \mathbb{N}, f \in \mathcal{S}_1, n > 0} \left| \frac{\partial \ln \pi_m(x, f, n)}{\partial p_x} \frac{\partial p_x}{\partial \ln n} + \sum_{i \in S} \frac{\partial \ln \pi_m(x, f, n)}{\partial p_i} \frac{\partial p_i}{\partial \ln n} \right| < \infty$. The complete derivation is long and algebraically cumbersome so it is omitted. However, we note a couple of important points. To compute $\frac{\partial p_i}{\partial \ln n}$ we use equation (6.1) together with the implicit function theorem. Each term in the sum is $\Theta\left(\frac{1}{n}\right)$, hence the sum, even if it includes an infinite number of terms, remains bounded.

Now we derive the maximal absolute semi-elasticity function, $g(x)$. Similarly to equation (A.7) we have

$$(A.8) \quad \frac{d \ln \pi_m(y, f, n)}{df(x)} = \frac{\partial \ln \pi_m(y, f, n)}{\partial f(x)} + \frac{\partial \ln \pi_m(y, f, n)}{\partial p_y} \frac{\partial p_y}{\partial f(x)} + \sum_{i \in S} \frac{\partial \ln \pi_m(y, f, n)}{\partial p_i} \frac{\partial p_i}{\partial f(x)}.$$

Now,

$$\frac{\partial \ln \pi_m(y, f, n)}{\partial f(x)} = -\frac{n(1+x)^{\delta_1}(Y-p_x)^{\delta_2}}{1+n \sum_{z \in \mathbb{N}} f(z)(1+z)^{\delta_1}(Y-p_z)^{\delta_2} + (1+y)^{\delta_1}(Y-p_y)^{\delta_2}}$$

Therefore,

$$(A.9) \quad \max_{m \in \mathbb{N}, y \in \mathbb{N}, f \in \mathcal{S}_1, n \in \mathbb{N}} \left| \frac{\partial \ln \pi_m(y, f, n)}{\partial f(x)} \right| = O\left(x^{\delta_1}\right).$$

The second and third terms in equation (A.8) can be bound in a similar way to (A.7). Assumption 4.2 is satisfied if $\sup_m E[g(\tilde{x}^{(m)})^{1+\gamma}] < \infty$, for some $\gamma > 0$. Take $\delta_1 = \frac{1}{1+\gamma}$ to conclude that $\sup_m E[\tilde{x}^{(m)}] < \infty$ implies the light-tail condition.¹⁵

A.5 Proofs and Mathematical Arguments for Section 5

Theorem 5.1. *Let Assumptions 2.1, 2.2, and 2.3 hold. Then, for any oblivious equilibrium $(\tilde{\mu}, \tilde{\lambda})$ and firm state $x \in \mathbb{N}$,*

$$(A.10) \quad E \left[\sup_{\mu' \in \mathcal{M}} V(x, s_t | \mu', \tilde{\mu}, \tilde{\lambda}) - V(x, s_t | \tilde{\mu}, \tilde{\lambda}) \right] \leq \frac{2}{1-\beta} E[\Delta |\pi_m|(s)],$$

where $\Delta |\pi|(s) = \max_{y \in \mathbb{N}} |\pi(y, s) - \pi(y, \tilde{s})|$, and

$$(A.11) \quad E \left[\sup_{\mu' \in \mathcal{M}} V(x, s_t | \mu', \tilde{\mu}, \tilde{\lambda}) - V(x, s_t | \tilde{\mu}, \tilde{\lambda}) \right] \leq \frac{1}{1-\beta} E[\Delta \pi(s)] + \sum_{y \in \mathbb{N}} a_x(y) (\pi(y, \tilde{s}) - E[\pi(y, s)]),$$

where $\Delta \pi(s) = \max_{y \in \mathbb{N}} (\pi(y, s) - \pi(y, \tilde{s}))$.

Proof. We derive the second bound, beginning with the following proposition. Let μ^* be an optimal (non-oblivious) best response to an oblivious equilibrium $(\tilde{\mu}, \tilde{\lambda})$ for a firm that is keeping track of the industry state.

Proposition A.1.

$$E[V(x, s | \mu^*, \tilde{\mu}, \tilde{\lambda}) - \tilde{V}(x, s | \tilde{\mu}, \tilde{\lambda})] \leq \frac{1}{1-\beta} E[\Delta \pi_m(s)], \quad \forall x \in \mathbb{N}.$$

¹⁵Note that for large n , $\pi_m(x, f, n)$ grows roughly like x^{δ_1} as x grows. Hence, taking $\delta_1 = \frac{1}{1+\gamma}$ for small $\gamma > 0$ means that profits increase almost linearly with the quality level.

Proof. By a similar argument to the one at the beginning of the proof of Theorem 4.4, we have that

$$(A.12) \quad V(x, s | \mu^*, \tilde{\mu}, \tilde{\lambda}) - \tilde{V}(x, s | \tilde{\mu}, \tilde{\lambda}) \leq E_{\mu^*, \tilde{\mu}, \tilde{\lambda}} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} (\pi_m(x_{ik}, s_{-i,k}) - \pi_m(x_{ik}, \tilde{s})) \mid x_{it} = x, s_{-i,t} = s \right].$$

The equation can be rewritten as:

$$(A.13) \quad V(x, s | \mu^*, \tilde{\mu}, \tilde{\lambda}) - \tilde{V}(x, s | \tilde{\mu}, \tilde{\lambda}) \leq \sum_{k=t}^{\infty} \beta^{k-t} \sum_{\substack{y \in \mathbb{N} \\ s' \in \mathcal{S}}} P_{\mu^*, \tilde{\mu}, \tilde{\lambda}}[x_{ik} = y, s_{-i,k} = s' \mid x_{it} = x, s_{-i,t} = s] \cdot (\pi_m(y, s') - \pi_m(y, \tilde{s})),$$

where $P_{\mu^*, \tilde{\mu}, \tilde{\lambda}}[x_{ik} = y, s_{-i,k} = s' \mid x_{it} = x, s_{-i,t} = s]$ is the probability firm i , currently in state x with competitors in state s , will be in state y and s' , respectively, $k - t$ periods from now.

We can write:

$$\begin{aligned} P_{\mu^*, \tilde{\mu}, \tilde{\lambda}}[x_{ik} = y, s_{-i,k} = s' \mid x_{it} = x, s_{-i,t} = s] &= P_{\mu^*, \tilde{\mu}, \tilde{\lambda}}[x_{ik} = y \mid s_{-i,k} = s', x_{it} = x, s_{-i,t} = s] \cdot \\ &P_{\mu^*, \tilde{\mu}, \tilde{\lambda}}[s_{-i,k} = s' \mid x_{it} = x, s_{-i,t} = s] \\ &= P_{\mu^*, \tilde{\mu}, \tilde{\lambda}}[x_{ik} = y \mid s_{-i,k} = s', x_{it} = x, s_{-i,t} = s] \cdot \\ &P_{\tilde{\mu}, \tilde{\lambda}}[s_{-i,k} = s' \mid s_{-i,t} = s]. \end{aligned}$$

The last equation follows because rival firms use strategy $\tilde{\mu}$, which only depends on their own state, and the entry rate is $\tilde{\lambda}$ independent of the industry state. Substituting into equation (A.13) gives:

$$(A.14) \quad \begin{aligned} V(x, s | \mu^*, \tilde{\mu}, \tilde{\lambda}) - \tilde{V}(x, s | \tilde{\mu}, \tilde{\lambda}) &\leq \sum_{k=t}^{\infty} \beta^{k-t} \sum_{s' \in \mathcal{S}} P_{\tilde{\mu}, \tilde{\lambda}}[s_{-i,k} = s' \mid s_{-i,t} = s] \\ &\sum_{y \in \mathbb{N}} P_{\mu^*, \tilde{\mu}, \tilde{\lambda}}[x_{ik} = y \mid s_{-i,k} = s', x_{it} = x, s_{-i,t} = s] (\pi_m(y, s') - \pi_m(y, \tilde{s})) \\ &\leq \sum_{k=t}^{\infty} \beta^{k-t} \sum_{s' \in \mathcal{S}} P_{\tilde{\mu}, \tilde{\lambda}}[s_{-i,k} = s' \mid s_{-i,t} = s] \max_{y \in \mathbb{N}} (\pi_m(y, s') - \pi_m(y, \tilde{s})), \end{aligned}$$

Recall that $q(s)$ is the invariant distribution of $\{s_t : t \geq 0\}$, where s_t is the industry state at time t when every firms uses strategy $\tilde{\mu}$ and the entry rate is $\tilde{\lambda}$. Note that $\{s_{-i,t} : t \geq 0\}$ is the same process as $\{s_t : t \geq 0\}$. Hence for any $k \geq t$:

$$(A.15) \quad q(s') = \sum_{s \in \mathcal{S}} q(s) P_{\tilde{\mu}, \tilde{\lambda}}[s_{-i,k} = s' \mid s_{-i,t} = s]$$

Multiplying equations (A.14) by $q(s)$ and summing over all $s \in \mathcal{S}$ we obtain:

$$\begin{aligned}
\sum_{s \in \mathcal{S}} q(s) \left(V(x, s | \mu^*, \tilde{\mu}, \tilde{\lambda}) - \tilde{V}(x, s | \tilde{\mu}, \tilde{\lambda}) \right) &\leq \sum_{s \in \mathcal{S}} q(s) \sum_{k=t}^{\infty} \beta^{k-t} \sum_{s' \in \mathcal{S}} P_{\tilde{\mu}, \tilde{\lambda}}[s_{-i, k} = s' \mid s_{-i, t} = s] \cdot \\
&\quad \max_{y \in \mathbb{N}} (\pi_m(y, s') - \pi_m(y, \tilde{s})) \\
&= \sum_{k=t}^{\infty} \beta^{k-t} \sum_{s' \in \mathcal{S}} \sum_{s \in \mathcal{S}} q(s) P_{\tilde{\mu}, \tilde{\lambda}}[s_{-i, k} = s' \mid s_{-i, t} = s] \cdot \\
&\quad \max_{y \in \mathbb{N}} (\pi_m(y, s') - \pi_m(y, \tilde{s})) \\
\text{(A.16)} \quad &= \sum_{k=t}^{\infty} \beta^{k-t} \sum_{s' \in \mathcal{S}} q(s') \max_{y \in \mathbb{N}} (\pi_m(y, s') - \pi_m(y, \tilde{s}))
\end{aligned}$$

The second equation follows by Fubini and the last one by equation (A.15). The previous argument is valid for any $x \in \mathbb{N}$, therefore:

$$E[V(x, s | \mu^*, \tilde{\mu}, \tilde{\lambda}) - \tilde{V}(x, s | \tilde{\mu}, \tilde{\lambda})] \leq \frac{1}{1 - \beta} E[\Delta \pi_m(s)], \quad \forall x \in \mathbb{N},$$

where s is a random vector distributed according to q . □

Returning to the derivation of the bound, we have that:

$$\begin{aligned}
\text{(A.17)} \quad E[V(x, s | \mu^*, \tilde{\mu}, \tilde{\lambda}) - V(x, s | \tilde{\mu}, \tilde{\lambda})] &= E[V(x, s | \mu^*, \tilde{\mu}, \tilde{\lambda}) - \tilde{V}(x, s | \tilde{\mu}, \tilde{\lambda})] + E[\tilde{V}(x, s | \tilde{\mu}, \tilde{\lambda}) - V(x, s | \tilde{\mu}, \tilde{\lambda})]
\end{aligned}$$

The first term is bounded by the previous proposition. Let us analyze the second term:

$$\begin{aligned}
\tilde{V}(x, s | \tilde{\mu}, \tilde{\lambda}) - V(x, s | \tilde{\mu}, \tilde{\lambda}) &= E_{\tilde{\mu}, \tilde{\lambda}} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} (\pi_m(x_{ik}, \tilde{s}) - dt_k^i) + \beta^{\tau_i-t} \phi_{\tau_i}^i \mid x_{it} = x, s_{-i, t} = s \right] \\
&- E_{\tilde{\mu}, \tilde{\lambda}} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} (\pi_m(x_{ik}, s_{-i, k}) - dt_k^i) + \beta^{\tau_i-t} \phi_{\tau_i}^i \mid x_{it} = x, s_{-i, t} = s \right] \\
&= E_{\tilde{\mu}, \tilde{\lambda}} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} (\pi_m(x_{ik}, \tilde{s}) - \pi_m(x_{ik}, s_{-i, k})) \mid x_{it} = x, s_{-i, t} = s \right] \\
&= \sum_{k=t}^{\infty} \beta^{k-t} \sum_{s' \in \mathcal{S}} P_{\tilde{\mu}, \tilde{\lambda}}[s_{-i, t} = s' \mid s_{-i, t} = s] \cdot \\
\text{(A.18)} \quad &\sum_{y \in \mathbb{N}} P_{\tilde{\mu}}[x_{ik} = y \mid x_{it} = x] (\pi_m(y, \tilde{s}) - \pi_m(y, s')).
\end{aligned}$$

The last equation follows because under oblivious strategies firms' trajectories are independent. Multiplying

each term by $q(s)$, summing over all $s \in \mathcal{S}$ and interchanging sums in the right hand side using Fubini we obtain:

$$(A.19) \quad E[\tilde{V}(x, s|\tilde{\mu}, \tilde{\lambda}) - V(x, s|\tilde{\mu}, \tilde{\lambda})] = \sum_{k=t}^{\infty} \beta^{k-t} \sum_{y \in \mathbb{N}} P^{\tilde{\mu}}[x_{ik} = y \mid x_{it} = x] (\pi_m(y, \tilde{s}) - E[\pi_m(y, s)]) ,$$

where s is a random vector distributed according to q . Finally, interchanging the sums

$$(A.20) \quad E[\tilde{V}(x, s|\tilde{\mu}, \tilde{\lambda}) - V(x, s|\tilde{\mu}, \tilde{\lambda})] = \sum_{y \in \mathbb{N}} a_x(y) (\pi_m(y, \tilde{s}) - E[\pi_m(y, s)]) .$$

The second bound follows by equations (A.17), (A.20) and the proposition. The first one follows by a similar argument, but with the difference that we take $\max_{y \in \mathbb{N}}$ in equation (A.18) and we take absolute value of the difference of one period profits in equations (A.14) and (A.18). \square

References

- Altman, E., O. Pourtallier, A. Haurie, and F. Moresino (2000). Approximating Nash equilibria in nonzero-sum games. *International Game Theory Review* 2(2), 155 – 172.
- 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. and A. Pakes (1993). Some applications and limitations of recent advances in empirical industrial organization: Merger analysis. *American Economic Review* 83(2), 247 – 252.
- Bertsekas, D. P. (2001). *Dynamic Programming and Optimal Control, Vol. 2* (Second ed.). Athena Scientific.
- Besanko, D. and U. Doraszelski (2004). Capacity dynamics and endogenous asymmetries in firm size. *RAND Journal of Economics* 35(1), 23 – 49.
- Besanko, D., U. Doraszelski, Y. Kryukov, and M. Satterthwaite (2005). Learning-by-doing, organizational forgetting, and industry dynamics. Working Paper Kellogg University.
- Caplin, A. and B. Nalebuff (1991). Aggregation and imperfect competition - on the existence of equilibrium. *Econometrica* 59(1), 25 – 59.
- de Roos, N. (2002). A model of collusion timing. Working Paper Yale University.
- Doraszelski, U. and K. Judd (2003). Continuous-time stochastic games with discrete states. Working Paper Hoover Institution.
- Doraszelski, U. and S. Markovich (2003). Advertising dynamics and competitive advantage. Working Paper Hoover Institution.
- Doraszelski, U. and M. Satterthwaite (2003). Foundations of Markov-perfect industry dynamics: Existence, purification, and multiplicity. Working Paper, Hoover Institution.
- Durrett, R. (1996). *Probability: Theory and Examples* (Second ed.). Duxbury Press.
- Ericson, R. and A. Pakes (1995). Markov-perfect industry dynamics: A framework for empirical work. *Review Of Economic Studies* 62(1), 53 – 82.
- Fershtman, C. and A. Pakes (2000). A dynamic oligopoly with collusion and price wars. *RAND Journal of Economics* 31(2), 207 – 236.
- Fudenberg, D. and D. Levine (1986). Limit games and limit equilibria. *Journal of Economic Theory* 38(2), 261 – 279.

- Goettler, R. L., C. A. Parlour, and U. Rajan (2004). Equilibrium in a dynamic limit order market. Forthcoming, *Journal of Finance*.
- Gowrisankaran, G. (1999). A dynamic model of endogenous horizontal mergers. *RAND Journal of Economics* 30(1), 56 – 83.
- Gowrisankaran, G. and R. Town (1997). Dynamic equilibrium in the hospital industry. *Journal of Economics and Management Strategy* 6(1), 45 – 74.
- Haugh, M. B., L. Kogan, and J. Wang (2004). Pricing american options: A duality approach. *Operations Research* 52(2), 258–270.
- Haugh, M. B., L. Kogan, and J. Wang (2005). Portfolio evaluation: A duality approach. to appear in *Operations Research*.
- Hopenhayn, H. A. (1992). Entry, exit and firm dynamics in long run equilibrium. *Econometrica* 60(5), 1127 – 1150.
- Jenkins, M., P. Liu, R. L. Matzkin, and D. L. McFadden (2004). The browser war - econometric analysis of Markov perfect equilibrium in markets with network effects. Working Paper.
- Jovanovic, B. and R. W. Rosenthal (1988). Anonymous sequential games. *Journal of Mathematical Economics* 17(1), 77 – 87.
- Judd, K. L., K. Schmedders, and S. Yeltekin (2002). Optimal rules for patent races. Working Paper, Hoover Institution.
- Kleinrock, L. (1975). *Queueing Systems, Volume 1: Theory* (First ed.). Wiley-Interscience.
- Klette, T. J. and S. Kortum (2003). Innovating firms and aggregate innovation. accepted in *Journal of Political Economy*.
- Langohr, P. (2003). Competitive convergence and divergence: Capability and position dynamics. Working Paper Northwestern University.
- Markovich, S. (2003). Rolling vs. melting: The snowball-effect in a dynamic oligopoly model with network externalities. Working Paper Tel-Aviv University.
- Maskin, E. and J. Tirole (1988). A theory of dynamic oligopoly, i and ii. *Econometrica* 56(3), 549 – 570.
- Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica* 71(6), 1695 – 1725.

- Novshek, W. and H. Sonnenschein (1978). Cournot and Walras equilibrium. *Journal of Economic Theory* 19, 223 – 266.
- Pakes, A. (2000). A framework for applied dynamic analysis in I.O. NBER Working Paper 8024.
- Pakes, A. and P. McGuire (1994). Computing Markov-perfect Nash equilibria: Numerical implications of a dynamic differentiated product model. *RAND Journal of Economics* 25(4), 555 – 589.
- Pakes, A. and P. McGuire (2001). Stochastic algorithms, symmetric Markov perfect equilibrium, and the 'curse' of dimensionality. *Econometrica* 69(5), 1261 – 1281.
- Ryan, S. (2005). The costs of environmental regulation in a concentrated industry. MIT, Mimeo.
- Shaked, A. and J. Sutton (1987). Product differentiation and industrial structure. *Journal of Industrial Economics* 36(2), 131 – 146.
- Song, M. (2003). A dynamic model of cooperative research in the semiconductor industry. Working Paper Harvard University.
- Sutton, J. (1991). *Sunk Costs and Market Structure* (First ed.). MIT Press.
- Vapnik, V. N. and A. Y. Chervonenkis (1971). On the Uniform Convergence of Relative Frequencies of Events to their Probabilities. *Theory of Probability and its Applications* 16(2), 264–280.
- Vives, X. (2002). Private information, strategic behavior, and efficiency in cournot markets. *RAND Journal of Economics* 33(3), 361 – 376.

B Tables and Figures

Table 1: Comparison of MPE and OE strategies (4 firms, no entry and exit)

Parameters		Long Run Statistics (% Diff)					Perf Bound (% Diff)		Actual (% Diff)	
θ_1	d	Inv.	Prod Surp	Cons Surp	C1	C2	Max Diff	Weighted Avg	Max Diff	Weighted Avg
0.10	0.10	-0.26	-0.01	-0.02	0.03	0.03	0.14	0.13	0.08	0.07
0.15	0.27	3.54	0.14	0.2	1.22	0.46	0.36	0.35	0.1	0.1
0.20	0.35	4.18	0.29	0.42	1.93	1.03	0.81	0.77	-0.09	-0.05
0.30	0.30	-0.13	0.06	0.08	0.08	0.16	1.67	1.22	0.04	0.01
0.30	0.55	9.28	0.93	1.31	5.10	2.45	1.96	1.85	0.26	0.25
0.40	0.80	21.02	2.10	2.93	11.58	4.12	3.01	2.92	0.30	0.29
0.50	1.00	18.62	3.30	4.33	15.69	5.94	6.29	5.86	0.32	0.30
0.50	0.50	-0.11	0.20	0.28	0.18	0.50	6.64	3.61	0.21	0.06
0.70	0.70	-2.21	0.40	0.15	1.08	2.09	18.85	8.35	1.60	0.67
0.85	0.70	-2.19	0.23	-0.28	1.37	2.10	30.80	9.64	1.80	0.20
1.20	1.00	-8.41	1.59	0.36	2.73	8.41	131.30	48.50	23.70	7.90

Long run statistics and value functions simulated with a relative precision of 1.0% and a confidence level of 99%. Performance bound simulated with a relative precision of at most 10% and a confidence level of 99%.

Table 2: Comparison of MPE and OE Investment (4 firms, no entry and exit)

Parameters		Investment		% Diff
θ_1	d	MPE	OE	
0.10	0.10	0.752	0.754	-0.26
0.15	0.27	0.192	0.185	3.54
0.20	0.35	0.261	0.250	4.18
0.30	0.30	0.754	0.755	-0.13
0.30	0.55	0.238	0.216	9.28
0.40	0.80	0.168	0.133	21.02
0.50	1.00	0.195	0.158	18.62
0.50	0.50	0.741	0.742	-0.11
0.70	0.70	0.694	0.709	-2.21
0.85	0.70	0.748	0.765	-2.19
1.20	1.00	0.553	0.599	-8.41

Investment simulated with a relative precision of 1.0% and a confidence level of 99%.

Figure 1: Percentage approximation error bound for Poisson entry process for different market sizes.

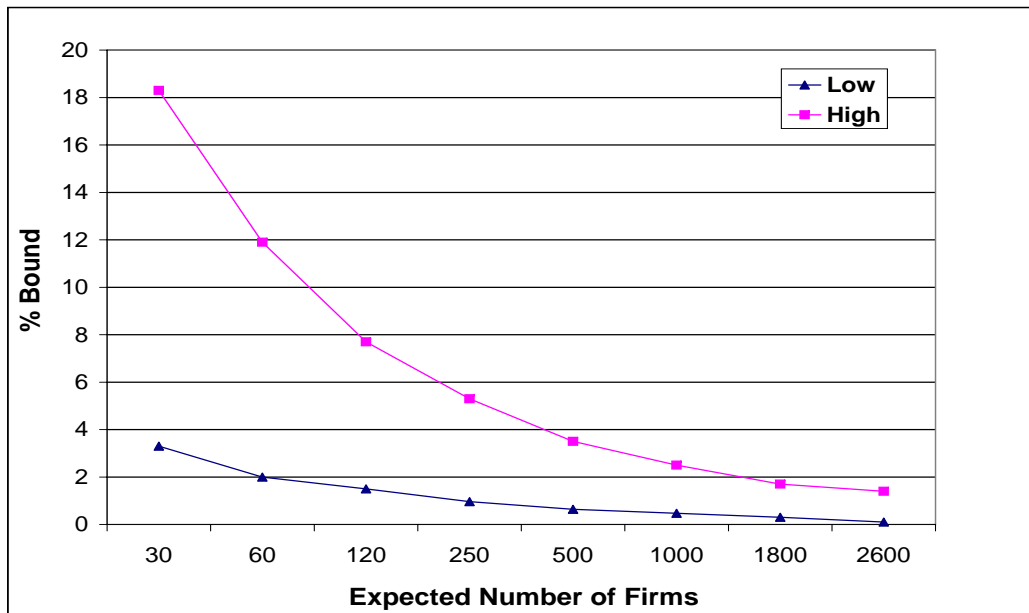


Figure 2: Percentage approximation error bound for deterministic entry process for different market sizes.

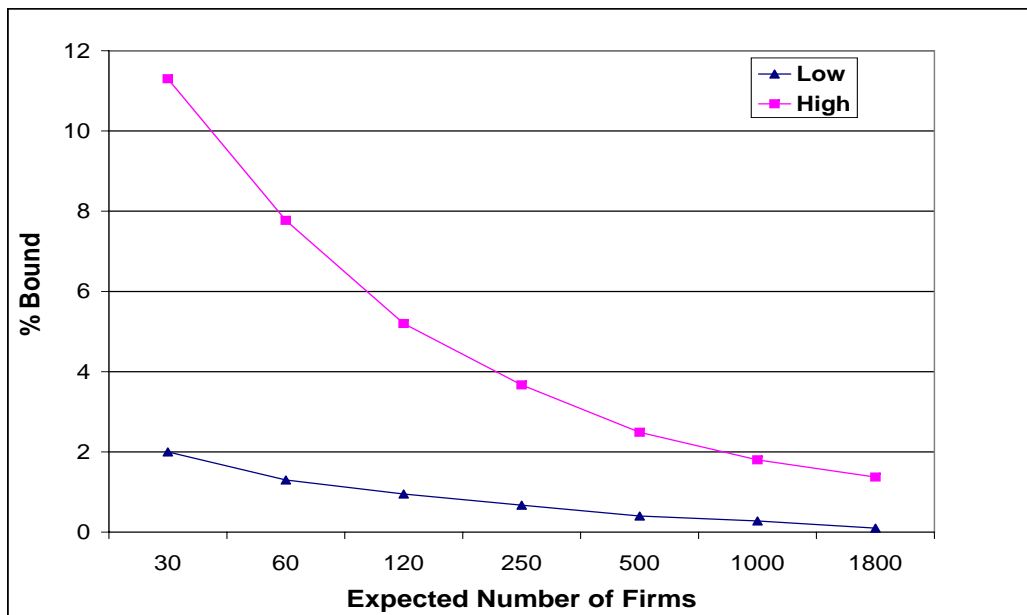


Figure 3: Percentage approximation error bound for fixed number of firms.




Figure 4: Average industry state for $\theta_1 = 0.5$ and $d = 0.5$.




Figure 5: Average industry state for $\theta_1 = 0.4$ and $d = 0.8$.

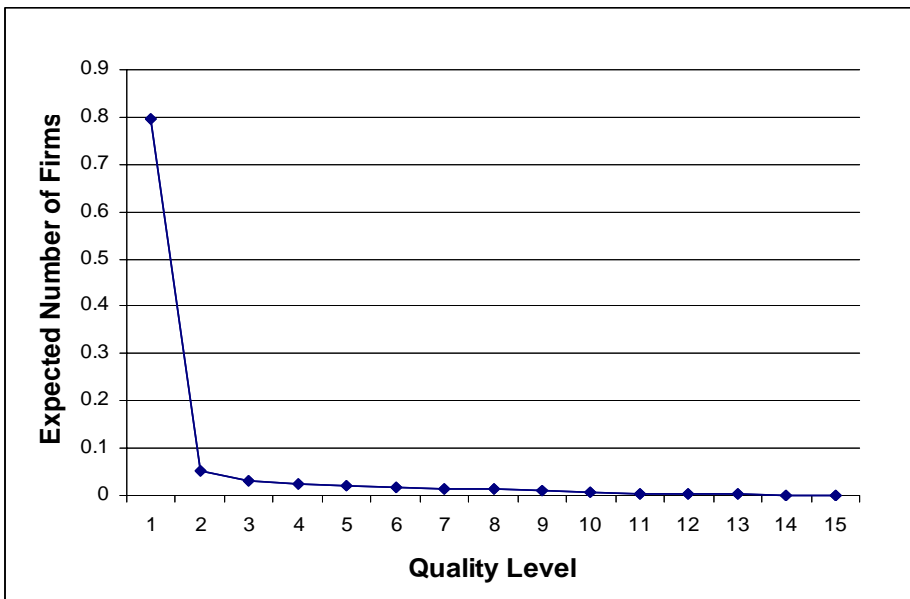


Figure 6: Average industry state for different market sizes for $\theta_1 = 0.9$.

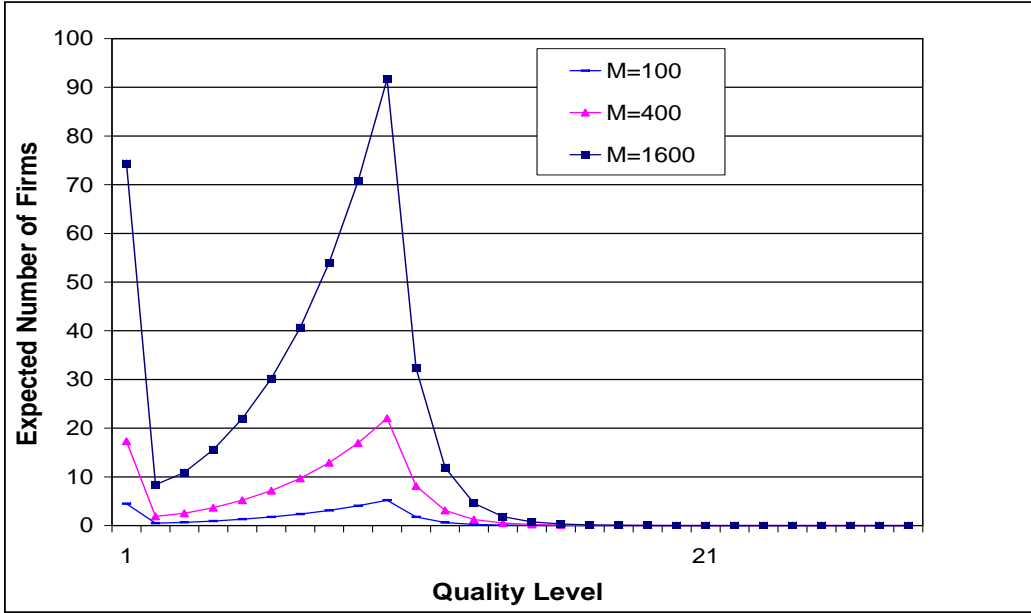


Figure 7: Average industry state for different market sizes for $\theta_1 = 1.2$.

