

ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

*An International Society for the Advancement of Economic
Theory in its Relation to Statistics and Mathematics*

<http://www.econometricsociety.org/>

Econometrica, Vol. 76, No. 6 (November, 2008), 1375–1411

MARKOV PERFECT INDUSTRY DYNAMICS WITH MANY FIRMS

GABRIEL Y. WEINTRAUB

Columbia Business School, Columbia University, New York, NY 10027, U.S.A.

C. LANIER BENKARD

Graduate School of Business, Stanford University, Stanford, CA 94305-5015, U.S.A.

BENJAMIN VAN ROY

Stanford University, Stanford, CA 94305-4026, U.S.A.

The copyright to this Article is held by the Econometric Society. It may be downloaded, printed and reproduced only for educational or research purposes, including use in course packs. No downloading or copying may be done for any commercial purpose without the explicit permission of the Econometric Society. For such commercial purposes contact the Office of the Econometric Society (contact information may be found at the website <http://www.econometricsociety.org> or in the back cover of *Econometrica*). This statement must be included on all copies of this Article that are made available electronically or in any other format.

MARKOV PERFECT INDUSTRY DYNAMICS WITH MANY FIRMS

BY GABRIEL Y. WEINTRAUB, C. LANIER BENKARD,
AND BENJAMIN VAN ROY¹

We propose an approximation method for analyzing Ericson and Pakes (1995)-style dynamic models of imperfect competition. We define a new equilibrium concept that we call 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. The great advantage of oblivious equilibria is that they are much easier to compute than are Markov perfect equilibria. Moreover, we show 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. This theorem justifies using oblivious equilibria to analyze Markov perfect industry dynamics in Ericson and Pakes (1995)-style models with many firms.

KEYWORDS: Dynamic games, oblivious equilibrium, approximation, industrial organization, imperfect competition.

1. INTRODUCTION

JUST OVER A DECADE AGO, [Ericson and Pakes \(1995\)](#) (hereafter EP) introduced an approach to modeling a dynamic industry with heterogeneous firms. The goal of their paper, stated in the title, was to promote empirical work to evaluate the effects of policy and environmental changes on things like job turnover, market structure, and welfare. The EP model and its extensions have proven to be both useful and broadly applicable. In the EP model, dynamics came in the form of firm investment, entry, and exit. However, subsequent work extended the model to many other contexts, including dynamics in the product market, dynamic demand, collusion, and asymmetric information.² With the advent of new ways to estimate such models (see [Pesendorfer and Schmidt-Dengler \(2003\)](#), [Bajari, Benkard, and Levin \(2007\)](#), [Pakes, Ostrovsky, and Berry \(2007\)](#), [Aguirregabiria and Mira \(2007\)](#)), this has become an active area for frontier level applied research.

¹We have had very helpful conversations with José Blanchet, Uli Doraszelski, Liran Einav, Hugo Hopenhayn, Ken Judd, Jon Levin, and Ariel Pakes, as well as seminar participants at Berkeley, Columbia, Duke, IIOC, Iowa, Informs, Kellogg, Minnesota, NYU, SITE, Stanford, Rochester, UCLA, UIUC, University of Chile, UT Austin, and Yale. We thank the editor and three anonymous referees for valuable suggestions. This research was supported by the Federal Reserve Bank of San Francisco, General Motors, the Lillie Fund, the National Science Foundation, and the Office of Naval Research.

²See [Berry and Pakes \(1993\)](#), [Gowrisankaran \(1999\)](#), [Fershtman and Pakes \(2000, 2005\)](#), [Judd, Schmedders, and Yeltekin \(2002\)](#), [Langohr \(2003\)](#), [Song \(2003\)](#), [Besanko and Doraszelski \(2004\)](#), [de Roos \(2004\)](#), [Besanko, Doraszelski, Kryukov, and Satterthwaite \(2005\)](#), [Goettler, Parlour, and Rajan \(2005\)](#), [Doraszelski and Markovich \(2007\)](#), [Markovich \(2008\)](#), [Noel \(2008\)](#), and [Schivardi and Schneider \(2008\)](#), as well as [Doraszelski and Pakes \(2007\)](#) for an excellent survey.

There remain, however, some substantial hurdles in the application of EP-type models. Because EP-type models are analytically intractable, analyzing market outcomes is typically done by solving for Markov perfect equilibria (MPE) numerically on a computer, using dynamic programming algorithms (e.g., [Pakes and McGuire \(1994\)](#)). This is a computational problem of the highest order. Note that even if it is possible to estimate the model parameters without computing an equilibrium, as in the papers listed above, equilibrium computation is still required to analyze the effects of a policy or other environmental change. Methods that accelerate these equilibrium computations have been proposed ([Judd \(1998\)](#), [Pakes and McGuire \(2001\)](#), [Doraszelski and Judd \(2006\)](#)). However, in practice computational concerns have typically limited the analysis to industries with just a few firms, much less than the real world industries at which the analysis is directed. Such limitations have made it difficult to construct realistic empirical models, and application of the EP framework to empirical problems is still quite difficult (see [Gowrisankaran and Town \(1997\)](#), [Benkard \(2004\)](#), [Jenkins, Liu, Matzkin, and McFadden \(2004\)](#), [Ryan \(2005\)](#), [Collard-Wexler \(2006\)](#)). Furthermore, even where applications have been deemed feasible, model details are often dictated as much by computational concerns as economic ones.

In this paper we consider a new notion of equilibrium that we call *oblivious equilibrium* (henceforth, OE). Our hope is that OE will serve as the basis for more refined equilibrium notions that can be used to closely approximate MPE market outcomes in realistically sized industries. OE is a natural starting point for approximating MPE outcomes because, as we show in this paper, under fairly general assumptions OE is close to MPE in markets with many firms. We make this statement more precise below. In addition, OE is much simpler to solve for than MPE, and it can be computed easily even for industries with hundreds or thousands of firms. Finally, we believe that for some industries OE and its extensions may provide an appealing behavioral model in its own right. By using OE in place of MPE, we hope it will become possible to analyze many empirical problems that would have been intractable before.

In an EP-type model, at each time, each firm has a state variable that captures its competitive advantage. Although 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 that represents 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 computer time and memory required for computing MPE) becomes enormous very quickly as the numbers of firms and firm states grows. For example, most industries contain more than 20 firms, but it would require more than 20 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 industries.

The intuition behind our alternative 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 is a large 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 by knowing only its own firm state and the long-run average industry state. Strategies for which a firm considers only its own state and the long-run average industry state we call *oblivious* strategies, and an equilibrium in which firms use oblivious strategies we call *oblivious equilibrium* (OE). Computing OE is simple because dynamic programming algorithms that optimize over oblivious strategies require computer time and memory that scale only with the number of firm states, not with the number of firms. Hence, computational considerations place very few constraints on model details.

A fundamental question becomes whether OE provides meaningful approximations of MPE behavior. To address this question, in the main theorem of the paper we prove an asymptotic result that provides sufficient conditions for OE 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, it turns out that 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 precise. In that case a strategy that does not keep track of the dominant firm's state would not perform well even if there were many firms in the industry. Instead, we show that, alongside some technical requirements, a sufficient condition for OE 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. To make this notion more concrete, 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 is uniformly bounded for all market sizes.

Since the main result of this paper is a limit theorem, it is natural to wonder whether there are any real world industries in which OE can be shown to approximate MPE well. We address this issue more directly in a companion paper ([Weintraub, Benkard, and Van Roy \(2007\)](#)), where we provide a computational algorithm that solves for OE and approximation bounds that can be computed to provide researchers with a numerical measure of how close OE is to MPE in their particular application. We find in that paper that OE approximate MPE better in some industries than others, depending on industry concentration and the nature of competitive interactions. In one example the approximation works well with a C4 of close to 100%. In another, the approximation fails to work well until there are over a thousand firms in the industry. These results suggest that, by using OE to approximate MPE, it is possible to

greatly increase the set of problems that can be analyzed using EP-type models. As further evidence of the usefulness of OE, Xu (2006) has already used OE in a novel study of research and development investment in the Korean electric motor industry. This application would not have been possible using exact computation of MPE.

While our computational results suggest that OE will be useful in many applications on its own, there are also many industries that are too large to compute MPE, but that are too concentrated for OE to provide a precise approximation. Thus, we believe that a major contribution of OE will be as a starting point with which to build even better approximations. We suggest some possible extensions to OE in the conclusions.

Finally, we note that, while our emphasis is on the use of OE as an approximation to MPE, in many cases OE (and its extensions) may also provide an appealing behavioral model on its own. 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 rest of paper is organized as follows. In Section 2 we describe the relationship between our work and previous literature. In Section 3 we outline the dynamic industry model. In Section 4 we introduce the concept of oblivious strategies and oblivious equilibrium. In Section 5 we provide our main result, namely, we give conditions under which oblivious equilibrium approximates MPE asymptotically as the market size grows. Finally, Section 6 presents conclusions and a discussion of future research directions. All proofs and mathematical arguments are provided in the Appendix and the Supplemental material (Weintraub, Benkard, and Van Roy (2008)).

2. PREVIOUS LITERATURE

Our approach is related to some past work. The concept of OE is most closely related to Hopenhayn (1992).³ Hopenhayn (1992) modeled an industry with a continuum of firms, each of which garners an infinitesimal fraction of the market. His model is tractable because the industry state is constant over time, implicitly assuming a law of large numbers holds. This assumption is based on the same intuition that motivates our consideration of OE. However, our goal is to apply our model directly to data, matching such industry statistics as the number of firms, the market shares of leading firms, the level of markups, and the correlation between investment and firm size. Thus, we are forced to consider models that more closely reflect real world industries that have finite numbers of firms, with strictly positive market shares. Relative to Hopenhayn (1992), another difference is that in our EP-type model investment is explicitly modeled, though it is also possible to achieve this in a Hopenhayn-type model.⁴

³See also Jovanovic and Rosenthal (1988).

⁴See Klette and Kortum (2004) and Melitz (2003).

Our approach is also similar in spirit to [Krusell and Smith \(1998\)](#), whose method has become popular in the recent macroeconomics literature. [Krusell and Smith \(1998\)](#) were concerned with stochastic growth models with a continuum of heterogeneous agents, in which case the state space of the model is infinite dimensional and consists of the entire wealth distribution in the economy. Since solving for a true equilibrium to the model would be analytically and computationally intractable, they instead solved for an equilibrium in which strategies are a function of some simple summary statistics of the wealth distribution. However, rather than attempting to prove that their algorithm closely approximates the true equilibrium, as we do here, they viewed their model as a behavioral one involving agents with bounded rationality. Their work raises the question of why we did not consider a similar approach in EP-type models. In fact, we considered such an approach, but were unable to prove that equilibrium behavior of this sort would be close to MPE behavior in an EP-type model. The approximation theorem in this paper essentially shows that, asymptotically, the concept of OE can only eliminate equilibria relative to MPE, not create any. We were unable to establish a similar result for equilibria based on summary statistics.

Our results are also related to work by [Levine and Pesendorfer \(1995\)](#), [Fudenberg, Levine, and Pesendorfer \(1998\)](#), and [Al-Najjar and Smorodinsky \(2001\)](#) on the negligibility of players in repeated interactions. These papers were motivated by the fact that equilibria in a game with a large but finite number of agents can be very different from equilibria in the limit model with a continuum of agents. In a model with a finite number of players, deviations can be detected. Therefore, equilibria where agents act nonmyopically can be sustained using schemes of rewards and punishments based on past play. On the other hand, in the continuum model, all agents are negligible and an individual deviation cannot be detected and punished. These papers provide conditions under which this paradox is resolved (e.g., if players' actions are observed with noise). Our light-tail condition is related to the idea that any individual player has little effect on the subsequent play of others. However, our work differs from this literature in several ways. First, we restrict our attention to Markovian (payoff-relevant, history independent) strategies. This restriction already rules out many equilibria that the negligibility literature is concerned about.⁵ Additionally, unlike the papers listed above, in our model the number and size of firms (or agents) is endogenous. Finally, another difference between our work and the literature above is that the notion of a limit game in our model could, in general, be complex. Thus, we do not attempt to consider equilibria to the limit game itself. Instead we focus on the idea that OE becomes close to MPE as the market size increases.

⁵To see this more clearly, note that in an infinitely repeated game the only Markov perfect equilibria correspond to the infinite repetition of one of the Nash equilibria of the stage game. In this case, the paradox studied by the above-mentioned papers is immediately resolved: all agents act myopically.

An implication of our results is that, for asymptotically large markets, a simple strategy that ignores current market information can be close to optimal. In this sense, our results are also related to those of [Roberts and Postlewaite \(1976\)](#), [Novshek and Sonnenschein \(1978, 1983\)](#), [Mas-Colell \(1982, 1983\)](#), [Novshek \(1985\)](#), [Allen and Hellwig \(1986a, 1986b\)](#), and [Jones \(1987\)](#). Roughly speaking, these papers establish conditions in different static models under which the set of oligopolistic Nash equilibria approaches, in some sense, the set of (Walrasian) competitive equilibria as the size of individual firms (or agents) becomes small relative to the size of the market. There are some notable differences with our work, though. Our interests lie in approximating dynamic firm behavior in large markets, not in showing that the product market is perfectly competitive in the limit. In particular, while the above papers study *static* models, in which the main strategic decisions are usually prices or quantities, we study *dynamic* models, in which the main decisions are, for example, investment, entry, and exit. Thus, while we show that firm investment, entry, and exit strategies become simple in markets with many firms, we do not rule out that a small fraction or even all firms may still have some degree of market power in the product market even in the limit.

3. 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. The model is general enough to encompass numerous applied problems in economics. Indeed, a blossoming recent literature on EP-type models has applied similar models to advertising, auctions, collusion, consumer learning, environmental policy, international trade policy, learning-by-doing, limit order markets, mergers, network externalities, and other applied problems.

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, the asymptotic theorem in this paper does not hold when there are aggregate industry shocks, so our model includes only idiosyncratic shocks.⁶

3.1. Model and Notation

The industry evolves over discrete time periods and an infinite horizon. We index time periods with nonnegative 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

⁶In [Weintraub, Benkard, and Van Roy \(2007\)](#) we extended the model and analysis to include aggregate shocks.

filtration $\{\mathcal{F}_t : t \geq 0\}$. We adopt the 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 . We define the state space $\bar{\mathcal{S}} = \{s \in \mathbb{N}^\infty \mid \sum_{x=0}^\infty s(x) < \infty\}$. Although in principle there are a countable number of industry states, we will also consider an extended state space $\mathcal{S} = \{s \in \mathbb{R}_+^\infty \mid \sum_{x=0}^\infty s(x) < \infty\}$. 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 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 $\iota_{it} \in \mathbb{R}_+$, then the firm's state at time $t+1$ is given by

$$x_{i,t+1} = \max(0, x_{it} + w(\iota_{it}, \zeta_{i,t+1})),$$

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. Note that this specification is very general: w may take on either positive or negative values (e.g., allowing for positive depreciation). 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.⁷

⁷Note that it would not change any of our results to assume that the entry state was a random variable.

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 firm 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} .

3.2. Model Primitives

Our model above allows for a wide variety of applied problems. To study any particular problem it is necessary to further specify the primitives of the model, including the profit function π , the distribution of the sell-off value ϕ_{it} , the investment impact function w , the distribution of the investment uncertainty ζ_{it} , the unit investment cost d , the entry cost κ , and the discount factor β .

Note that in most applications 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 (and hence the profit function) that we will focus on below, is the size of the relevant market, which we will denote as m . Later on in the paper we subscript the profit function with the market size parameter, π_m , to explicitly recognize the dependence of profits on market size. For expositional clarity, the subscript is omitted in the assumptions listed below, implying that the market size is being held fixed.

3.3. 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 $x \in \mathbb{N}$, $\sum_{z \geq x} s(z) \geq \sum_{z \geq x} s'(z)$. We will denote this relation by $s \succeq s'$. Intuitively, competition associated with s is no weaker than competition associated with s' .

ASSUMPTION 3.1:

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 $x \in \mathbb{N}$, the function $\ln \pi(x, \cdot) : \mathcal{S} \rightarrow \mathfrak{R}_+$ is continuously Fréchet differentiable. Hence, for all $x \in \mathbb{N}$, $y \in \mathbb{N}$, and $s \in \mathcal{S}$, $\ln \pi(x, s)$ is continuously dif-

ferentiable with respect to $s(y)$. Further, for any $x \in \mathbb{N}$, $s \in \mathcal{S}$, 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 fairly weak. Assumption 3.1.1 ensures that increases in quality lead to increases in profit. Assumption 3.1.2 states that strengthened competition cannot result in increased profit. Assumption 3.1.3 ensures that profits are positive and bounded. Assumption 3.1.4 is technical and requires that log profits be Fréchet differentiable. Note that it requires partial differentiability of the profit function with respect to each $s(y)$. Profit functions that are “smooth,” such as ones arising from random utility demand models like the logit model, will satisfy this assumption.

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

ASSUMPTION 3.2:

1. The random variables $\{\phi_{it} | t \geq 0, i \geq 1\}$ are independent and identically distributed (i.i.d.) and have finite expectations and well-defined density functions with support $\mathfrak{R}_+$.

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) . There exists a positive constant $\bar{\iota}$ such that $\iota_{it} < \bar{\iota}$, $\forall i, \forall t$.

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.

Again the assumptions are natural and fairly weak. Assumptions 3.2.1 and 3.2.2 imply that investment and exit outcomes are idiosyncratic conditional on the state. Assumptions 3.2.3 and 3.2.4 imply that investment is productive. Note that positive depreciation is neither required nor ruled out. Assumption 3.2.5 places a finite bound on how much progress can be made or lost in a single period through investment. Assumption 3.2.6 ensures that the impact of investment on transition probabilities is continuous. Assumption 3.2.7

is an assumption introduced by Doraszelski and Satterthwaite (2007) that ensures a unique solution to the firms' investment decision problem. It is used to guarantee the existence of an equilibrium in pure strategies and 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 the Appendix for a derivation of this result). This leads to the following assumptions:

ASSUMPTION 3.3:

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 and earning zero profits 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 discounted profits condition. Modeling the number of entrants as a Poisson random variable has the advantage that it leads to more elegant asymptotic results. Assumption 3.2.2 ensures that the sell-off value by itself is not a sufficient reason to enter the industry.

3.4. 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,

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

where i is taken to be the index of a firm at quality level x at time t , τ_i is a random variable that represents 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:

$$(3.1) \quad \sup_{\mu' \in \mathcal{M}} V(x, s|\mu', \mu, \lambda) = V(x, s|\mu, \lambda) \quad \forall x \in \mathbb{N}, \forall s \in \overline{\mathcal{S}}.$$

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

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

In the [Appendix](#), we show that the supremum in part 1 of the definition above can always be attained simultaneously for all x and s by a common strategy μ' .

[Doraszelski and Satterthwaite \(2007\)](#) established the 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.⁸

Dynamic programming algorithms can be used to optimize firm strategies, and equilibria to our model can be computed via their iterative application.

⁸[Doraszelski and Satterthwaite \(2007\)](#) also provided an example of multiple equilibria in their closely related model.

However, these algorithms require computer 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.

4. 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 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.

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 that persists for all industry states.

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 subprobabilities 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 4.1: *Let Assumption 3.2 hold. 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*

$$(4.1) \quad \lim_{t \rightarrow \infty} \tilde{s}_t(x) = \lambda \sum_{k=0}^{\infty} P_\mu^k(x^e, x)$$

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*

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

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:

$$(4.2) \quad \sup_{\mu' \in \tilde{\mathcal{M}}} \tilde{V}(x|\mu', \mu, \lambda) = \tilde{V}(x|\mu, \lambda) \quad \forall x \in \mathbb{N}.$$

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

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

It is straightforward to show that OE exists under mild technical conditions. Furthermore, if the entry cost is not prohibitively high, then an OE 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 OEs 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 OE.⁹

⁹However, since oblivious strategies rule out strategies that are dependent on competitors' states, there are likely to be fewer OEs than there are MPE.

Finally, in the [Appendix](#) we show that when strategies and entry rate functions are oblivious, the Markov chain $\{s_t : t \geq 0\}$ admits an invariant distribution. Moreover, we show that when firms play OE strategies, the long-run behavior of the industry is such that the number of firms in each state x is given by a Poisson random variable with mean $\tilde{s}_{\mu,\lambda}(x)$, independent across states x .

5. ASYMPTOTIC RESULTS

As stated in the [Introduction](#) (see also [Weintraub, Benkard, and Van Roy \(2007\)](#)), OEs are simple to compute because they only involve one dimensional dynamic programs. However, how well do OEs approximate MPEs? In this section, we study this question and prove the main result of the paper. In particular, we establish an asymptotic result that provides conditions under which OEs offer close approximations to MPEs in large markets. The main condition is that the sequence of OEs generates firm size (or quality level) distributions that are “light-tailed” in a sense that we will make precise.

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 $m \in \mathbb{R}_+$. 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 OE 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 $q^{(m)}$. 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 long-run 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 let $\tilde{n}^{(m)} \equiv E[n_t^{(m)}] = \sum_{x \in \mathbb{N}} \tilde{s}^{(m)}(x)$ denote the expected number of firms. It is easy to check that $\tilde{f}^{(m)} = \tilde{s}^{(m)} / \tilde{n}^{(m)}$. With some abuse of notation, we define $\pi_m(x_{it}, f_{-i,t}, n_{-i,t}) \equiv \pi_m(x_{it}, n_{-i,t} \cdot f_{-i,t})$.

5.1. Assumptions About the Sequence of Profit Functions

In addition to Assumption 3.1, 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 5.1:

1. $\sup_{x \in \mathbb{N}, s \in \mathcal{S}} \pi_m(x, s) = O(m)$.¹⁰
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. The following holds

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

The assumptions are again fairly weak. Assumption 5.1.1, which states that profits increase at most linearly with market size, 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.¹¹ Assumption 5.1.2 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 5.1.3 requires that profits be smooth with respect to the number of firms and, in particular, 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 Section 5.5 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.

5.2. Asymptotic Markov Equilibrium

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

¹⁰In this notation, $f(m) = O(m)$ denotes $\limsup_m \frac{f(m)}{m} < \infty$ and $f(m) = o(m)$ denotes $\limsup_m \frac{f(m)}{m} = 0$.

¹¹For 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 5.1: 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 the approximation error at each firm state x in terms of the amount by which a firm at state x can increase its expected net present value by deviating from the OE strategy $\tilde{\mu}^{(m)}$ and instead following an optimal (nonoblivious) best response that keeps track of the true industry state. Recall that MPE require that the expression in square brackets equals zero for all states (x, s) . In that sense, the notion of AME relates to the more common notion that ε -equilibria approximate true equilibria in games as $\varepsilon \rightarrow 0$ (Fudenberg and Levine (1986)).¹²

Note that standard MPE solution algorithms (e.g., Pakes and McGuire (1994)) aim to compute pointwise 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 pointwise. This is because in OE, 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 a very low probability of occurrence, so they have a small impact on expected discounted profits.

If a sequence of OEs 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 that have significant probability of occurrence. This implies that, asymptotically, $\tilde{\mu}^{(m)}$ is a near-optimal strategy (so it approximately satisfies the MPE equation) when the industry starts in any state that has a significant probability of occurrence. Additionally, one can show that when the sequence of OEs has the AME property, then it is also the case that $\lim_{m \rightarrow \infty} E_{\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} [|V^{(m)}(x^e, s_t^{(m)} | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) - \tilde{V}(x^e | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)})|] = 0$. Since $\beta \tilde{V}(x^e | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) = \kappa$ for all m , asymptotically, $\beta V^{(m)}(x^e, s | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) \approx \kappa$ for states s that have a significant probability of occurrence. Hence, asymptotically, $\tilde{\lambda}^{(m)}$ satisfies the zero profit condition at such states. In summary, if $E_{\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} [\sup_{\mu' \in \mathcal{M}} V^{(m)}(x, s_t^{(m)} | \mu', \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) -$

¹²Similar notions of closeness based on ε -equilibria are used by Roberts and Postlewaite (1976), Fudenberg, Levine, and Pesendorfer (1998), and Al-Najjar and Smorodinsky (2001). Pakes and McGuire (2001) also used a concept similar to the AME property as a stopping rule in their stochastic algorithm.

$V^{(m)}(x, s_t^{(m)} | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)})]$ is small, MPE strategies and entry rates at relevant states should be well approximated by oblivious ones. Weintraub, Benkard, and Van Roy (2007) presented computational results that support this point.¹³

Since OE will only have the ability to approximate MPE strategies in states that have a significant probability of occurrence, OE will not be able to approximate equilibria in which behavior depends on off-the-equilibrium-path values, as is typically the case in tacit collusive schemes. In that sense, OE can be also understood (at least heuristically) as a mechanism that selects MPE that are “noncollusive.”¹⁴

5.3. 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 5.1: *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 A.4 in the Appendix. 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 OEs possesses the AME property. However, for this to be the case it turns out that additional conditions are required.

5.4. 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 few firms—the AME property is unlikely to hold. A strategy that does not keep track of the dominant firms will perform poorly. To ensure the AME property, we need to impose a light-tail condition that rules out this kind of market concentration.

¹³One might attempt to formalize this argument by following a line of reasoning similar to Fudenberg and Levine (1986) or Altman, Pourtaillier, Haurie and Moresino (2000), who provided 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 $\varepsilon_m \rightarrow 0$ converges to an equilibrium of G .

¹⁴Note that even under the Markov perfection assumption, tacit collusive agreements could be supported as equilibria. For example, firms may remain small by the threat of escalation in investment outlays (see Nocke (2007)).

In this section, we establish the main result of the paper; under an appropriate light-tail condition, the sequence of OEs possesses the AME property.

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

$$g(x) = \sup_{\substack{m \in \mathbb{N}_+, y \in \mathbb{N}, \\ f \in S_1, n > 0}} \left| \frac{d \ln \pi_m(y, f, n)}{df(x)} \right|.$$

For each x , $g(x)$ is the maximum rate of relative change of any firm's single-period profit that could result from a small change in the fraction of firms at quality level x . 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 $\tilde{x}^{(m)} \sim \tilde{f}^{(m)}$, that is, $\tilde{x}^{(m)}$ is a random variable with probability mass function $\tilde{f}^{(m)}$. The random variable $\tilde{x}^{(m)}$ can be interpreted as the quality level of a firm that is randomly sampled from among all incumbents while the industry state is distributed according to its invariant distribution.

ASSUMPTION 5.2: *For all quality levels x , $g(x) < \infty$. For all $\varepsilon > 0$, there exists a quality level z such that*

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

for all market sizes m .

Put simply, the light-tail condition requires that states where a small change in the fraction of firms has a large impact on the profits of other firms must have a small probability under the invariant distribution. In practice, this typically means that very large firms (and hence high concentration) rarely occur under the invariant distribution.

Recall that $g(x)$ is the maximum rate of relative change of any firm's 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' profits as the market size grows. It is unlikely that OE will provide a good approximation in this situation. Note that the assumption imposes that $g(x)$ is finite for each x , but allows $g(x)$ to 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. It 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. 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 $\varepsilon > 0$, there exists a quality level z such that $E[g(\tilde{x}^{(m)})\mathbf{1}_{\{\tilde{x}^{(m)}>z\}}] \leq \varepsilon$. This is because increasing z sufficiently will result in ignoring a larger and larger fraction of firms in computing the expected impact and the expected impact when none of the firms is ignored is finite. Assumption 5.2 poses a stronger condition in that it requires that a quality level z can be chosen such that $E[g(\tilde{x}^{(m)})\mathbf{1}_{\{\tilde{x}^{(m)}>z\}}] \leq \varepsilon$ for all market sizes m simultaneously. This is like the light-tail condition $E[g(\tilde{x}^{(m)})\mathbf{1}_{\{\tilde{x}^{(m)}>z\}}] \leq \varepsilon$ 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 .

We note that if $g(x)$ is strictly increasing and unbounded, the light-tail condition is equivalent to the condition that the class of random variables $\{g(\tilde{x}^{(m)}) | m \in \mathfrak{H}_+\}$ is uniformly integrable. In this case, if there exists $\gamma > 0$, such that $\sup_m E[g(\tilde{x}^{(m)})^{1+\gamma}] < \infty$, Assumption 5.2 is satisfied. The condition is slightly stronger than requiring uniformly bounded first moments of $g(\tilde{x}^{(m)})$.

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

THEOREM 5.2: *Under Assumptions 3.1, 3.2, 3.3, 5.1.2, and 5.2, $\tilde{n}^{(m)} = \Omega(m)$.*¹⁵

The proof can be found in the [Appendix](#). The result implies that if the light-tail condition is satisfied, then the expected number of firms under OE strategies grows to infinity as the market size grows. 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 5.1. First, we define $\|f\|_{1,g} = \sum_x |f(x)|g(x)$.

¹⁵In this notation, $f(m) = \Omega(m)$ implies that $\liminf_m \frac{f(m)}{m} > 0$. With an additional technical regularity condition, it is straightforward to show that $\tilde{n}^{(m)} = O(m)$. Hence, in this case, $\tilde{n}^{(m)} = \Theta(m)$, that is, $\tilde{n}^{(m)}$ grows linearly in m asymptotically.

THEOREM 5.3: *Let Assumptions 3.1, 3.2, 3.3, 5.1.2, and 5.2 hold. Then, as m grows, $n_t^{(m)}/\tilde{n}^{(m)} \rightarrow_p 1$ and $\|f_t^{(m)} - \tilde{f}^{(m)}\|_{1,g} \rightarrow_p 0$.*

The light-tail condition is key to proving the second part of the result, namely, convergence of the normalized industry states in the $\|\cdot\|_{1,g}$ weighted norm. This new form of convergence allows us to ensure the AME property, which leads to the main result of the paper.

THEOREM 5.4: *Under Assumptions 3.1, 3.2, 3.3, 5.1, and 5.2, the sequence $\{\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}\}$ of OEs possesses the AME property.*

This result is proved in the [Appendix](#). We provide an explanation of the main steps of the proof here. First, we compare the value functions in the definition of the AME property through the OE value function. Formally,

$$\begin{aligned}
 (5.1) \quad & E_{\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\sup_{\mu' \in \mathcal{M}} V^{(m)}(x, s_t^{(m)} | \mu', \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) - V^{(m)}(x, s_t^{(m)} | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) \right] \\
 &= E_{\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\sup_{\mu' \in \mathcal{M}} V^{(m)}(x, s_t^{(m)} | \mu', \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) - \tilde{V}^{(m)}(x | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) \right] \\
 &\quad + E_{\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} [\tilde{V}^{(m)}(x | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) - V^{(m)}(x, s_t^{(m)} | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)})].
 \end{aligned}$$

We now explain how we prove that the first expectation in the right-hand side above converges to zero; the argument for the second one is analogous. First, we observe that because $\tilde{\mu}^{(m)}$ and $\tilde{\lambda}^{(m)}$ attain an OE, for all x ,

$$\begin{aligned}
 \tilde{V}^{(m)}(x | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) &= \sup_{\mu' \in \mathcal{M}} \tilde{V}^{(m)}(x | \mu', \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) \\
 &= \sup_{\mu' \in \mathcal{M}} \tilde{V}^{(m)}(x | \mu', \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}),
 \end{aligned}$$

where the last equation follows because there will always be an optimal oblivious strategy when optimizing an oblivious value function even if we consider more general strategies (a key feature of oblivious strategies). Hence,

$$\begin{aligned}
 & V^{(m)}(x, s | \mu^{*(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) - \tilde{V}^{(m)}(x | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) \\
 &\leq V^{(m)}(x, s | \mu^{*(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) - \tilde{V}^{(m)}(x | \mu^{*(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}),
 \end{aligned}$$

where $\mu^{*(m)} \in \mathcal{M}$ achieves the supremum in equation (5.1). Note that in the right-hand side of the above inequality both value functions are evaluated at the same set of strategies. This allows us to compare $V^{(m)}(x, s | \mu^{*(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)})$ with $\tilde{V}^{(m)}(x | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)})$ by only taking into consideration the difference between single-period profits (actual versus oblivious). The quantities associated with the difference between strategies ($\mu^{*(m)}$ versus $\tilde{\mu}^{(m)}$) can be

neglected, greatly simplifying the analysis. Hence, using the previous inequality we obtain one of the key insights of the proof; the difference between value functions can be bounded by a discounted sum of expected differences between actual and oblivious single-period profits. Formally, the first expectation in the right-hand side of equation (5.1) can be bounded by

$$(5.2) \quad \sum_{k=t}^{\infty} \beta^{k-t} E_{\mu^{*(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} [|\pi_m(x_{ik}, s_{-i,k}^{(m)}) - \pi_m(x_{ik}, \tilde{s}^{(m)})|] \\ x_{it} = x, s_{-i,t}^{(m)} \sim q^{(m)}].$$

Hence, if expected actual and oblivious single-period profits are close, then the respective expected value functions are also close.

To finish the proof, we show that the discounted sum (5.2) converges to zero. For this, first we argue that $\pi_m(x, s)$ and $\pi_m(x, s')$ are close when $|n/n' - 1|$ and $\|f - f'\|_{1,g}$ are small. The previous statement embodies another key insight. Single-period profits evaluated at different industry states are close if the states are close *under the appropriate norm*. For example, even if $\|f - f'\|_{\infty}$ is small (where $\|\cdot\|_{\infty}$ denotes the sup norm), the difference between $\pi_m(x, f, n) - \pi_m(x, f', n)$ could still be huge. Differences in profits are generally determined by the cumulative difference in the fraction of firms at each quality level. Moreover, it is likely that differences at larger states have a bigger impact on profits. The weighted $\|\cdot\|_{1,g}$ norm captures these effects. Finally, the previous continuity statement together with Theorem 5.3 is used to show that the discounted sum (5.2) converges to zero.

5.5. Example: Logit Demand System With Price Competition

To provide a concrete example of the conditions required for the asymptotic results above, we consider a model similar to that of [Pakes and McGuire \(1994\)](#). 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

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

where Y is the consumer's income and p_{it} is the price of the good produced by firm i . v_{ijt} are i.i.d. random variables that are Gumbel distributed and 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.

We assume that firms set prices in the spot market. If there is a constant marginal cost c , there is a unique Nash equilibrium in pure strategies, denoted p_t^* (Caplin and Nalebuff (1991)). Expected profits are given by

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

where σ represents the market share function from the logit model. In the Appendix we show that this single-period profit function satisfies Assumptions 3.1 and 5.1. In addition, we show that, in this model, the function $g(\tilde{x}^{(m)})$ takes a very simple form,

$$g(\tilde{x}^{(m)}) \propto (\tilde{x}^{(m)})^{\theta_1},$$

where θ_1 is the parameter on quality in the logit demand model. Therefore, the light-tail condition amounts to a simple condition on the equilibrium distribution of firm states. Under our assumptions, such a condition is equivalent to a condition on the equilibrium size distribution of firms.

The light-tail condition requires that for all $\varepsilon > 0$, there exists a quality level z such that $E[(\tilde{x}^{(m)})^{\theta_1} \mathbf{1}_{\{\tilde{x}^{(m)} > z\}}] \leq \varepsilon$ for all market sizes m . If $\theta_1 < 1$, then the light-tail condition 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. This condition allows for relatively fat-tailed distributions. For example, if $\tilde{x}^{(m)}$ is a sequence of log-normal distributions with uniformly bounded parameters, then the condition is satisfied. On the other hand, if $\tilde{x}^{(m)}$ converges to a Pareto distribution with parameter one (which does not have a finite first moment), then the condition would not be satisfied.

6. CONCLUSIONS AND FUTURE RESEARCH

The goal of this paper has been to provide foundations for an approximation method for analyzing Ericson and Pakes (1995)-style dynamic models of imperfect competition. Existing dynamic programming methods suffer from a curse of dimensionality that has limited application of these models to industries with a small number of firms and a small number of states per firm. We propose instead to approximate MPE behavior using OE, where firms make decisions only based on their own state and the long-run average industry state. Our main result lays a theoretical foundation for our method; we show that the approximation works well asymptotically, where asymptotics are taken in the market size. A key condition for an OE to well approximate MPE asymptotically is that the sequence of industry states generated by the OEs is light-tailed (as described by Assumption 5.2).

In a complementary paper (Weintraub, Benkard, and Van Roy (2007)), we have developed computational tools to facilitate the use of OE in practice.

That paper also provides computational evidence to support the conclusion that OE often yields good approximations of MPE behavior for industries like those that empirical researchers would like to study.

However, while we believe that the concept of OE will be useful in empirical applications on its own (see also Xu (2006)), we think that the greatest value of OE may come through some extensions that we now briefly describe. In a spirit similar to Krusell and Smith (1998), we envision algorithms that will progressively incorporate more information into firms' strategies, improving the accuracy of the approximation until an acceptable performance is reached. First, in Weintraub, Benkard, and Van Roy (2007) we extended the notion of OE to accommodate aggregate industry shocks. Additionally, to capture short-run transitional dynamics that may result, for example, from shocks or policy changes, we have also developed a nonstationary notion of OE in which every firm knows the industry state in the initial period, but does not update this knowledge after that point. Finally, in ongoing research, we are working on an extended notion of OE that allows for there to be a set of "dominant firms," whose firm states are always monitored by every other firm. Our hope is that the dominant firm OE will provide better approximations for more concentrated industries. Each of these extensions trades off increased computation time for a better behavioral model and a better approximation to MPE behavior. Note also that it is our belief that the theorem in this paper could be extended to cover each of these cases.

APPENDIX A: PROOFS AND MATHEMATICAL ARGUMENTS

A.1. *Proofs and Mathematical Arguments for Section 3*

A.1.1. *The Poisson Entry Model*

Suppose there are N potential entrants. Each potential entrant enters if the setup cost κ is less than the expected present value of future cash flows upon entry. Let $v_N(i)$ be the expected present value of future cash flows for each entering firm if i firms enter simultaneously; $v_N(i)$ 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_N . If we assume that every potential entrant employs the same strategy, the condition for a mixed strategy Nash equilibrium when $\kappa \in [v_N(N), v_N(1)]$ can be written as

$$(A.1) \quad \sum_{i=0}^{N-1} \binom{N-1}{i} p_N^i (1-p_N)^{N-1-i} v_N(i+1) - \kappa = 0,$$

which is solved by a unique $p_N \in [0, 1]$. If $\kappa < v_N(N)$, the equilibrium is a pure strategy with $p_N = 1$, whereas if $\kappa > v_N(1)$, the equilibrium is given by $p_N = 0$. 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 N grows large.

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

1. $v_N(i)$ is nonincreasing in $i \forall N$ and nonincreasing in $N \forall i$.
2. There exists a positive number M such that $|v_N(i)| < M \forall N, \forall i$.
3. There exists $\bar{N}$, such that $\forall N > \bar{N}$, $v_N(i) - \kappa$ changes sign in $\{1, N - 1\}$.
4. There exists a function $v(i)$ such that $\lim_{N \rightarrow \infty} \sup_{i \in \mathbb{N}} |v_N(i) - v(i)| = 0$.

Then the subsequent conditions hold:

1. For each $N > \bar{N}$, equation (A.1) has a unique solution $p_N^* \in (0, 1)$.
2. The limit $\lambda = \lim_{N \rightarrow \infty} N p_N^*$ exists, and if Y_N is a binomial random variable with parameters (N, p_N^*) and Z is a Poisson random variable with mean λ , then $Y_N \Rightarrow Z$.

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 playing a mixed strategy and entering the industry with a very small probability.

A.1.2. Bellman's Equation and Exit Cutoff Strategy

We define a dynamic programming operator:

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

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

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 3.1 and 3.2 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: Investment is bounded. Let $\bar{\pi} = \sup_{x, s} \pi(x, s) < \infty$. Then $\beta \bar{\phi} \leq \rho(x, s) \leq \frac{\bar{\pi}}{1-\beta} + \bar{\phi} \forall x \in \mathbb{N}, \forall s \in \bar{\mathcal{S}}$. 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

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

Note that $\pi(x, s) < \bar{\pi}$, investment is bounded, 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). *Q.E.D.*

By the lemma and the definition of the dynamic programming operator we observe that there exists an optimal exit strategy that has the form of a cutoff value.

A.2. Proofs and Mathematical Arguments for Section 4: Long-Run Behavior and the Invariant Industry Distribution

In this section we analyze the long-run behavior of an industry where strategies and the entry rate function are oblivious. The results will be useful to prove our main result. The proofs are provided in a separate technical appendix (Weintraub, Benkard, and Van Roy (2008)).

In Lemma 4.1, 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 A.3: *Let Assumptions 3.2 and 3.3 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 process;
- (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) $n_t \Rightarrow Z$.

We state another 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 A.4: *Let Assumptions 3.2 and 3.3 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 $n \in \mathbb{N}$,*

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

The Poisson entry process is key to proving these results. Lemma A.4 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.

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 OE. This follows from the fact that the sell-off value has support in $\Re_+$ 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 OE strategies.

A.3. Proofs and Mathematical Arguments for Section 5

A.3.1. Preliminary Lemmas

LEMMA A.5: *Under Assumptions 3.1, 3.2, and 3.3, for all x ,*

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

PROOF: We will assume that $x \geq x^e$; the case of $x < x^e$ is trivial. Assume for contradiction that $\sup_m \tilde{V}^{(m)}(x | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) = \infty$. 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 $\sup_m \tilde{V}^{(m)}(x^e | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) = \infty$. This contradicts the zero profit condition. *Q.E.D.*

Let $\ell_{1,g} = \{f \in \mathfrak{R}_+^\infty | \|f\|_{1,g} < \infty\}$. With some abuse of notation, let $\mathcal{S}_{1,g} = \mathcal{S}_1 \cap \ell_{1,g}$.

LEMMA A.6: *Let Assumptions 3.1, 3.2, 3.3, and 5.2 hold. Then, for all x ,*

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

PROOF: Let all expectations in this proof be conditioned on $x_{it} = x$. First note that

$$(A.2) \quad \sup_m \sup_{\mu \in \mathcal{M}} E_\mu \left[\sum_{k=t}^{\infty} \beta^{k-t} \pi_m(x_{ik}, \tilde{s}^{(m)}) \right] < \infty.$$

If not, $\sup_m \tilde{V}^{(m)}(x | \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}) = \infty$, because total investment cost is bounded by $\frac{\bar{c}}{1-\beta}$, violating Lemma A.5.

By Assumption 3.1.4, the fundamental theorem of calculus, and the chain rule for Fréchet differentiable functions (Luenberger (1969)), it follows that for any $y \in \mathbb{N}$, $f, f' \in \mathcal{S}_{1,g}$, and $m, n \in \mathfrak{R}_+$,

$$\begin{aligned} (A.3) \quad & |\ln \pi_m(y, f, n) - \ln \pi_m(y, f', n)| \\ &= \left| \int_{\gamma=0}^1 \sum_{x \in \mathbb{N}} (f(x) - f'(x)) \left(\frac{\partial \ln \pi_m(y, f' + \gamma(f - f'), n)}{\partial f(x)} \right) d\gamma \right| \\ &\leq \int_{\gamma=0}^1 \sum_{x \in \mathbb{N}} (|f(x) - f'(x)|) g(x) d\gamma \\ &= \|f - f'\|_{1,g} \\ &< \infty. \end{aligned}$$

Letting $\underline{f} = (1, 0, 0, \dots)$, using Assumption 5.2, it follows that $\sup_{x \in \mathbb{N}, m \in \mathfrak{R}_+} |\pi_m(x, \underline{f}, \tilde{n}^{(m)}) - \pi_m(x, \tilde{s}^{(m)})| \equiv C < \infty$. By Assumption 3.1.2, 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.2). *Q.E.D.*

A.3.2. Proof of Theorem 5.2

THEOREM 5.2: *Under Assumptions 3.1, 3.2, 3.3, 5.1.2, and 5.2, $\tilde{n}^{(m)} = \Omega(m)$.*

PROOF: Assume for contradiction that $\liminf_m (\tilde{n}^{(m)}/m) = 0$. Then there exists an increasing sequence m_k such that $\lim_k (\tilde{n}^{(m_k)}/m_k) = 0$, and by Assumption 5.1.2, for all $x, z \in \mathbb{N}$ with $x > z$, and $f \in \mathcal{S}_{1,z}$, $\lim_{k \rightarrow \infty} \pi_{m_k}(x, f, \tilde{n}^{(m_k)}) = \infty$.

Using equation (A.3), it follows from Assumptions 3.1.4 and 5.2 that

$$\limsup_{z \rightarrow \infty} \inf_m \sup_{f \in \mathcal{S}_{1,z}} \sup_x |\ln \pi_m(x, \tilde{f}^{(m)}, \tilde{n}^{(m)}) - \ln \pi_m(x, \hat{f}, \tilde{n}^{(m)})| = 0.$$

It follows that

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

and, therefore,

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

for a sequence $\hat{f}^{(m)} \in \mathcal{S}_{1,z}$ with sufficiently large z , and 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.5. It follows that $\tilde{n}^{(m)} = \Omega(m)$. Q.E.D.

A.3.3. Proof of Theorem 5.3

LEMMA A.7: Let Assumptions 3.1, 3.2, 3.3, 5.1.2, and 5.2 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 A.3, $n_t^{(m)}$ is a Poisson random variable with mean $\tilde{n}^{(m)}$. By Theorem 5.2, $\tilde{n}^{(m)} = \Omega(m)$. The result follows. Q.E.D.

THEOREM 5.3: Let Assumptions 3.1, 3.2, 3.3, 5.1.2, and 5.2 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.7. To complete the proof, we will establish convergence of $\|f_t^{(m)} - \tilde{f}^{(m)}\|_{1,g}$. Note that

for any $z \in \mathbb{N}$,

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

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 $\varepsilon > 0$, for sufficiently large z , $\limsup_{m \rightarrow \infty} \mathcal{P}[B_z^{(m)} \geq \delta] \leq \varepsilon/\delta$. The assertion that $\|f_t^{(m)} - \tilde{f}^{(m)}\|_{1,g} \rightarrow_p 0$ follows from these facts.

By Lemma A.4, for any x , $(f_t^{(m)}(x)|n_t^{(m)} = n)$ is distributed as the empirical mean of n i.i.d. Bernoulli random variables with expectation $\tilde{f}^{(m)}(x)$. It follows that for any x , $(|f_t^{(m)}(x) - \tilde{f}^{(m)}(x)||n_t = n)$ converges in probability to zero uniformly over m as n grows. By Theorem 5.2 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 Assumption 5.2, for any $\delta > 0$, for sufficiently large z , $\limsup_{m \rightarrow \infty} C_z^{(m)} < \delta$ 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 $\varepsilon > 0$, for sufficiently large z , $\limsup_{m \rightarrow \infty} \mathcal{P}[B_z^{(m)} \geq \delta] \leq \varepsilon/\delta$. The result follows. *Q.E.D.*

A.3.4. Proof of Theorem 5.4

The following technical lemma follows immediately from Assumption 5.1.3. The proof is similar to Lemma A.10 and is omitted.

LEMMA A.8: *Let Assumptions 3.1.3 and 5.1.3 hold. Then, for all $\varepsilon > 0$, there exists $\delta > 0$ such that for all $n, \hat{n} \in \mathbb{R}_+$ satisfying $|n/\hat{n} - 1| < \delta$,*

$$\sup_{m \in \mathbb{N}_+, x \in \mathbb{N}, f \in \mathcal{S}_1} \left| \frac{\pi_m(x, f, n) - \pi_m(x, f, \hat{n})}{\pi_m(x, f, \hat{n})} \right| \leq \varepsilon.$$

LEMMA A.9: *Let Assumptions 3.1, 3.2, 3.3, 5.1, and 5.2 hold. Then, for all sequences $\{\mu^{(m)} \in \mathcal{M}\}$,*

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

PROOF: For the purpose of this proof, we will assume that all expectations are conditioned on $x_{it} = x$ and $s_{-i,t}^{(m)} \sim q^{(m)}$. Note that if $s_{-i,t}^{(m)}$ is distributed according to $q^{(m)}$, then the marginal distribution of $s_{-i,k}^{(m)}$ is also $q^{(m)}$ for all $k > t$. Let $\Delta_{it}^{(m)} = |\pi_m(x_{it}, s_{-i,t}^{(m)}) - \pi_m(x_{it}, f_{-i,t}^{(m)}, \tilde{n}^{(m)})|$. Fix $\varepsilon > 0$ and let $\delta > 0$ satisfy the assertion of Lemma A.8. Let $Z^{(m)}$ denote the event $|n_t^{(m)}/\tilde{n}^{(m)} - 1| \geq \delta$. Applying Tonelli's theorem, we obtain

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

where the second to last inequality follows from Lemma A.8 and Assumption 5.1.1. Since ε is arbitrary, the expected sum is finite (by Lemma A.6), and $\mathcal{P}[Z^{(m)}] \leq e^{-\Omega(m)}$ (by Lemma A.7), the result follows. *Q.E.D.*

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

LEMMA A.10: *Let Assumptions 3.1.3 and 3.1.4 hold. Then, for all $\varepsilon > 0$, there exists $\delta > 0$ such that for $f, \hat{f} \in \mathcal{S}_{1,g}$ satisfying $\|f - \hat{f}\|_{1,g} < \delta$,*

$$\sup_{m \in \mathfrak{N}_+, x \in \mathbb{N}, n \in \mathfrak{N}_+} \left| \frac{\pi_m(x, f, n) - \pi_m(x, \hat{f}, n)}{\pi_m(x, \hat{f}, n)} \right| \leq \varepsilon.$$

PROOF: By equation (A.3), for any $x \in \mathbb{N}$, $f, \hat{f} \in \mathcal{S}_{1,g}$, and $m, n \in \mathfrak{N}_+$,

$$(A.4) \quad \left| \ln(\pi_m(x, f, n)/\pi_m(x, \hat{f}, n)) \right| \leq \|f - \hat{f}\|_{1,g}.$$

Let $B = \pi_m(x, f, n)/\pi_m(x, \hat{f}, n)$, where we have ignored the dependence on $m, x, f, \hat{f}, n$ to simplify notation. Take $\varepsilon > 0$ and let δ in the statement of

the lemma be $\delta = \ln(1 + \varepsilon) > 0$. Suppose $B > 1$. By equation (A.4), $\ln(B) = |\ln(B)| < \ln(1 + \varepsilon)$; hence, $|B - 1| < \varepsilon$ for all m, x, n . Suppose $B < 1$. By equation (A.4), $-\ln(B) = |\ln(B)| < \ln(1 + \varepsilon)$; hence, $1/B < 1 + \varepsilon$. Then $|B - 1| = 1 - B < \varepsilon/(1 + \varepsilon) < \varepsilon$, for all m, x, n , where the last inequality follows because $\varepsilon > 0$. The result follows. *Q.E.D.*

LEMMA A.11: *Let Assumptions 3.1, 3.2, 3.3, 5.1, and 5.2 hold. Then, for all sequences $\{\mu^{(m)} \in \mathcal{M}\}$,*

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

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 $\varepsilon > 0$ and let δ satisfy the assertion of Lemma A.10. Let $Z^{(m)}$ denote the event $\|f_{-i,k}^{(m)} - \tilde{f}^{(m)}\|_{1,g} \geq \delta$. Applying Tonelli's theorem, we obtain

$$E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} \Delta_{ik}^{(m)} \right] \\ \leq \sum_{k=t}^{\infty} \beta^{k-t} E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} [\Delta_{ik}^{(m)}] \\ = \sum_{k=t}^{\infty} \beta^{k-t} (E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} [\Delta_{ik}^{(m)} \mathbf{1}_{-Z^{(m)}}] + E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} [\Delta_{ik}^{(m)} \mathbf{1}_{Z^{(m)}}]) \\ \leq \varepsilon C + \sum_{k=t}^{\infty} \beta^{k-t} E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} [\Delta_{ik}^{(m)} \mathbf{1}_{Z^{(m)}}]$$

for some constant $C > 0$. The last inequality follows from Lemmas A.6 and A.10.

Note that $\Delta_{ik}^{(m)} \leq 2 \sup_{f \in S_1} \pi_m(x_{ik}, f, \tilde{n}^{(m)})$. Hence,

$$\sum_{k=t}^{\infty} \beta^{k-t} E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} [\Delta_{ik}^{(m)} \mathbf{1}_{Z^{(m)}}] \\ \leq \sum_{k=t}^{\infty} \beta^{k-t} E_{\mu^{(m)}, \tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[2 \sup_{f \in S_1} \pi_m(x_{ik}, f, \tilde{n}^{(m)}) \mathbf{1}_{Z^{(m)}} \right]$$

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

because $\sup_{\mu \in \mathcal{M}}$ is attained by an oblivious strategy, so $f_{-i,k}^{(m)}$ evolves independently from $x_{ik}^{(m)}$. Since ε is arbitrary, $\mathcal{P}[Z^{(m)}] \rightarrow 0$ (by Theorem 5.3), and the expected sum is uniformly bounded over all m (by Lemma A.6), the result follows. *Q.E.D.*

THEOREM 5.4: *Under Assumptions 3.1, 3.2, 3.3, 5.1, and 5.2, the sequence $\{\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}\}$ of OE possesses the AME property.*

PROOF: Let $\mu^{*(m)}$ be an optimal (nonoblivious) 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 OE, for all x ,

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

where the last equation follows because there will always be an optimal oblivious strategy when optimizing an oblivious value function, even if we consider more general strategies. It follows that

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

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 exits 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

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

$$B^{(m)}(x, s) \leq E_{\tilde{\mu}^{(m)}, \tilde{\lambda}^{(m)}} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} \Delta_{ik}^{(m)} \middle| x_{it} = x, s_{-i,t}^{(m)} = s \right],$$

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

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

By the triangle inequality,

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

The result therefore follows from Lemmas A.9 and A.11. Q.E.D.

A.3.5. Derivations for the Logit Demand Model Outlined in Section 5.5

For brevity, we omit the proofs that show Assumptions 3.1, 5.1.1, and 5.1.2 are satisfied. Now, we show that Assumption 5.1.3 is satisfied. Let $\psi = 1$. Note that

$$\begin{aligned} \text{(A.5)} \quad \frac{d \ln \pi_m(x, f, n)}{d \ln n} &= \frac{\partial \ln \pi_m(x, f, n)}{\partial \ln n} + \frac{\partial \ln \pi_m(x, f, n)}{\partial p_x} \frac{\partial p_x}{\partial \ln n} \\ & \quad + \sum_{i \in S} \frac{\partial \ln \pi_m(x, f, n)}{\partial p_i} \frac{\partial p_i}{\partial \ln n}, \end{aligned}$$

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 in the number of firms keeping prices fixed. The second and third terms consider the change in profits implied by the change of prices. Now

$$\begin{aligned} & \frac{\partial \ln \pi_m(x, f, n)}{\partial \ln n} \\ &= - \frac{n \sum_{z \in \mathbb{N}} f(z)(1+z)^{\theta_1}(Y-p_z)^{\theta_2}}{1 + n \sum_{z \in \mathbb{N}} f(z)(1+z)^{\theta_1}(Y-p_z)^{\theta_2} + (1+x)^{\theta_1}(Y-p_x)^{\theta_2}}. \end{aligned}$$

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 $\theta_2 \leq \frac{1}{2}$, then

$$\begin{aligned} & \sup_{m \in \mathbb{N}_+, x \in \mathbb{N}, f \in \mathcal{S}_1, n > 0} \left| \frac{\partial \ln \pi_m(x, f, n)}{\partial p_x} \frac{\partial p_x}{\partial \ln n} + \sum_{i \in S} \frac{\partial \ln \pi_m(x, f, n)}{\partial p_i} \frac{\partial p_i}{\partial \ln n} \right| \\ & < \infty. \end{aligned}$$

The complete derivation is long and algebraically cumbersome, so it is omitted. However, we note a couple of important points. To compute $\partial p_i / \partial \ln n$ we use the first-order condition for profit maximization together with the implicit function theorem. Each term in the sum is $\Theta(\frac{1}{n})$; hence, the sum remains bounded, even if it includes an infinite number of terms.

Now we derive the maximal absolute semielasticity function, $g(x)$. Similarly to equation (A.5), we have

$$\begin{aligned} \text{(A.6)} \quad \frac{d \ln \pi_m(y, f, n)}{df(x)} &= \frac{\partial \ln \pi_m(y, f, n)}{\partial f(x)} + \frac{\partial \ln \pi_m(y, f, n)}{\partial p_y} \frac{\partial p_y}{\partial f(x)} \\ &+ \sum_{i \in S} \frac{\partial \ln \pi_m(y, f, n)}{\partial p_i} \frac{\partial p_i}{\partial f(x)}. \end{aligned}$$

Now,

$$\begin{aligned} & \frac{\partial \ln \pi_m(y, f, n)}{\partial f(x)} \\ &= - \frac{n(1+x)^{\theta_1}(Y-p_x)^{\theta_2}}{1 + n \sum_{z \in \mathbb{N}} f(z)(1+z)^{\theta_1}(Y-p_z)^{\theta_2} + (1+y)^{\theta_1}(Y-p_y)^{\theta_2}}. \end{aligned}$$

Therefore,

$$\sup_{m \in \mathbb{N}_+, y \in \mathbb{N}, f \in \mathcal{S}_1, n > 0} \left| \frac{\partial \ln \pi_m(y, f, n)}{\partial f(x)} \right| \propto x^{\theta_1}.$$

The second and third terms in equation (A.6) can be bound in a similar way to (A.5). The result follows.

REFERENCES

- AGUIRREGABIRIA, V., AND P. MIRA (2007): "Sequential Estimation of Dynamic Discrete Games," *Econometrica*, 75, 1–54. [1375]
- ALLEN, B., AND M. HELLWIG (1986a): "Bertrand–Edgeworth Oligopoly in Large Markets," *The Review of Economic Studies*, 53, 175–204. [1380]
- (1986b): "Price-Setting Firms and the Oligopolistic Foundations of Perfect Competition," *The American Economic Review*, 76, 387–392. [1380]
- AL-NAJJAR, N. I., AND R. SMORODINSKY (2001): "Large Nonanonymous Repeated Games," *Games and Economic Behavior*, 37, 26–39. [1379,1390]
- ALTMAN, E., O. POURTALLIER, A. HAURIE, AND F. MORESINO (2000): "Approximating Nash Equilibria in Nonzero-Sum Games," *International Game Theory Review*, 2, 155–172. [1391]
- BAJARI, P., C. L. BENKARD, AND J. LEVIN (2007): "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, 75, 1331–1370. [1375]
- BENKARD, C. L. (2004): "A Dynamic Analysis of the Market for Wide-Bodied Commercial Aircraft," *Review of Economic Studies*, 71, 581–611. [1376]
- BERRY, S., AND A. PAKES (1993): "Some Applications and Limitations of Recent Advances in Empirical Industrial Organization: Merger Analysis," *American Economic Review*, 8, 247–252. [1375]
- BERTSEKAS, D. P. (2001): *Dynamic Programming and Optimal Control*, Vol. 2 (Second Ed.). Nashua, NH: Athena Scientific. [1399]
- BESANKO, D., AND U. DORASZELSKI (2004): "Capacity Dynamics and Endogenous Asymmetries in Firm Size," *RAND Journal of Economics*, 35, 23–49. [1375]
- BESANKO, D., U. DORASZELSKI, Y. KRYUKOV, AND M. SATTERTHWAIT (2005): "Learning-by-Doing, Organizational Forgetting, and Industry Dynamics," Working Paper, Northwestern University. [1375]
- CAPLIN, A., AND B. NALEBUFF (1991): "Aggregation and Imperfect Competition—On the Existence of Equilibrium," *Econometrica*, 59, 25–59. [1396]
- COLLARD-WEXLER, A. (2006): "Productivity Dispersion and Plant Selection in the Ready-Mix Concrete Industry," Working Paper, Northwestern University. [1376]
- DE ROOS, N. (2004): "A Model of Collusion Timing," *International Journal of Industrial Organization*, 22, 351–387. [1375]
- DORASZELSKI, U., AND K. JUDD (2006): "Avoiding the Curse of Dimensionality in Dynamic Stochastic Games," Working Paper, Hoover Institution. [1376]
- DORASZELSKI, U., AND S. MARKOVICH (2007): "Advertising Dynamics and Competitive Advantage," *RAND Journal of Economics*, 38, 557–592. [1375]
- DORASZELSKI, U., AND A. PAKES (2007): "A Framework for Applied Dynamic Analysis in IO," in *Handbook of Industrial Organization*, Vol. 3. Amsterdam: North-Holland. [1375]
- DORASZELSKI, U., AND M. SATTERTHWAIT (2007): "Computable Markov-Perfect Industry Dynamics: Existence, Purification, and Multiplicity," Working Paper, Harvard University. [1384, 1385]
- ERICSON, R., AND A. PAKES (1995): "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *Review of Economic Studies*, 62, 53–82. [1375,1380,1396]
- FERSHTMAN, C., AND A. PAKES (2000): "A Dynamic Oligopoly With Collusion and Price Wars," *RAND Journal of Economics*, 31, 207–236. [1375]
- (2005): "Finite State Dynamic Games With Asymmetric Information: A Framework for Applied Work," Working Paper, Harvard University. [1375]
- FUDENBERG, D., AND D. K. LEVINE (1986): "Limit Games and Limit Equilibria," *Journal of Economic Theory*, 38, 261–279. [1390,1391]

- FUDENBERG, D., D. K. LEVINE, AND W. PESENDORFER (1998): "When Are Nonanonymous Players Negligible?" *Journal of Economic Theory*, 79, 46–71. [1379,1390]
- GOETTLER, R. L., C. A. PARLOUR, AND U. RAJAN (2005): "Equilibrium in a Dynamic Limit Order Market," *Journal of Finance*, 60, 351–387. [1375]
- GOWRISANKARAN, G. (1999): "A Dynamic Model of Endogenous Horizontal Mergers," *RAND Journal of Economics*, 30, 56–83. [1375]
- GOWRISANKARAN, G., AND R. TOWN (1997): "Dynamic Equilibrium in the Hospital Industry," *Journal of Economics and Management Strategy*, 6, 45–74. [1376]
- HOPENHAYN, H. A. (1992): "Entry, Exit and Firm Dynamics in Long Run Equilibrium," *Econometrica*, 60, 1127–1150. [1378]
- 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. [1376]
- JONES, L. E. (1987): "The Efficiency of Monopolistically Competitive Equilibria in Large Economies: Commodity Differentiation With Gross Substitutes," *Journal of Economic Theory*, 41, 356–391. [1380]
- JOVANOVIĆ, B., AND R. W. ROSENTHAL (1988): "Anonymous Sequential Games," *Journal of Mathematical Economics*, 17, 77–87. [1378]
- JUDD, K. (1998): *Numerical Methods in Economics*. Cambridge, MA: MIT Press. [1376]
- JUDD, K. L., K. SCHMEDDERS, AND S. YELTEKIN (2002): "Optimal Rules for Patent Races," Working Paper, Hoover Institution. [1375]
- KLETTE, T. J., AND S. KORTUM (2004): "Innovating Firms and Aggregate Innovation," *Journal of Political Economy*, 112, 986–1018. [1378]
- KRUSELL, P., AND A. A. SMITH JR. (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106, 867–896. [1379,1397]
- LANGOHR, P. (2003): "Competitive Convergence and Divergence: Capability and Position Dynamics," Working Paper, Northwestern University. [1375]
- LEVINE, D. K., AND W. PESENDORFER (1995): "When Are Agents Negligible?" *The American Economic Review*, 85, 1160–1170. [1379]
- LUNENBERGER, D. G. (1969): *Optimization by Vector Space Methods* (First Ed.). New York: Wiley. [1401]
- MARKOVICH, S. (2008): "Snowball: A Dynamic Oligopoly Model With Indirect Network Effects," *Journal of Economic Dynamics and Control*, 32, 909–938. [1375]
- MAS-COLELL, A. (1982): "The Cournotian Foundation of Walrasian Equilibrium Theory: An Exposition of Recent Theory," in *Advances in Economic Theory: Invited Papers for the Fourth World Congress of the Econometric Society*, 183–224, ed. by W. Hildenbrand. Cambridge, U.K.: Cambridge University Press. [1380]
- (1983): "Walrasian Equilibria as Limits of Noncooperative Equilibria. Part I: Mixed Strategies," *Journal of Economic Theory*, 30, 153–170. [1380]
- MASKIN, E., AND J. TIROLE (1988): "A Theory of Dynamic Oligopoly, I and II," *Econometrica*, 56, 549–570. [1384]
- MELITZ, M. J. (2003): "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71, 1695–1725. [1378]
- NOCKE, V. (2007): "Collusion and Dynamic (Under-)Investment in Quality," *RAND Journal of Economics*, 38, 227–249. [1391]
- NOEL, M. D. (2008): "Edgeworth Cycles and Focal Prices: Computational Dynamic Markov Equilibria," *Journal of Economics and Management Strategy*, 17, 345–377. [1375]
- NOVSHEK, W. (1985): "Perfectly Competitive Markets as the Limits of Cournot Markets," *Journal of Economic Theory*, 35, 72–82. [1380]
- NOVSHEK, W., AND H. SONNENSCHN (1978): "Cournot and Walras Equilibrium," *Journal of Economic Theory*, 19, 223–266. [1380]
- (1983): "Walrasian Equilibria as Limits of Noncooperative Equilibria. Part II: Pure Strategies," *Journal of Economic Theory*, 30, 171–187. [1380]

- PAKES, A., AND P. MCGUIRE (1994): "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model," *RAND Journal of Economics*, 25, 555–589. [1376,1390,1395]
- (2001): "Stochastic Algorithms, Symmetric Markov Perfect Equilibrium, and the 'Curse' of Dimensionality," *Econometrica*, 69, 1261–1281. [1376,1390]
- PAKES, A., M. OSTROVSKY, AND S. BERRY (2007): "Simple Estimators for the Parameters of Discrete Dynamic Games, With Entry/Exit Examples," *RAND Journal of Economics*, 38, 373–399. [1375]
- PESENDORFER, M., AND P. SCHMIDT-DENGLER (2003): "Identification and Estimation of Dynamic Games," Mimeo, London School of Economics. [1375]
- ROBERTS, D. J., AND A. POSTLEWAITE (1976): "The Incentives for Price-Taking Behavior in Large Exchange Economies," *Econometrica*, 4, 115–127. [1380,1390]
- RYAN, S. (2005): "The Costs of Environmental Regulation in a Concentrated Industry," Mimeo, MIT. [1376]
- SCHIVARDI, F., AND M. SCHNEIDER (2008): "Strategic Experimentation and Disruptive Technological Change," *Review of Economic Dynamics*, 11, 386–412. [1375]
- SONG, M. (2003): "A Dynamic Model of Cooperative Research in the Semiconductor Industry," Working Paper, Harvard University. [1375]
- 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, 264–280. [1391]
- WEINTRAUB, G. Y., C. L. BENKARD, AND B. VAN ROY (2007): "Computational Methods for Oblivious Equilibrium," Working Paper, Stanford University. [1377,1380,1388,1391,1396,1397]
- (2008): "Supplement to 'Markov Perfect Industry Dynamics With Many Firms,'" *Econometrica Supplemental Material*, 76, http://www.econometricsociety.org/ecta/Supmat/6158_proofs.pdf. [1378,1399]
- XU, Y. (2006): "A Structural Empirical Model of R&D, Firm Heterogeneity, and Industry Evolution," Working Paper, Penn State University. [1378,1397]

Decisions, Risk, and Operations Division, Columbia Business School, Columbia University, Uris Hall 402, 3022 Broadway, New York, NY 10027, U.S.A.; gweintraub@columbia.edu,

Graduate School of Business, Stanford University, 518 Memorial Way, Stanford, CA 94305-5015, U.S.A.; lanierb@stanford.edu,

and

Dept. of Management Science and Engineering, Stanford University, 380 Panama Way, Stanford, CA 94305-4026, U.S.A.; bvr@stanford.edu.

Manuscript received November, 2005; final revision received May, 2008.