

Computational Methods for Oblivious Equilibrium

Gabriel Y. Weintraub

Business School, Columbia University, New York, New York 10027, gweintraub@columbia.edu

C. Lanier Benkard

Department of Economics, Yale University, New Haven, Connecticut 06511, lanier.benkard@yale.edu

Benjamin Van Roy

Stanford University, Stanford, California 94305, bvr@stanford.edu

Oblivious equilibrium is a new solution concept for approximating Markov-perfect equilibrium in dynamic models of imperfect competition among heterogeneous firms. In this paper, we present algorithms for computing oblivious equilibrium and for bounding approximation error. We report results from computational case studies that serve to assess both efficiency of the algorithms and accuracy of oblivious equilibrium as an approximation to Markov-perfect equilibrium. We also extend the definition of oblivious equilibrium, originally proposed for models with only firm-specific idiosyncratic random shocks, and our algorithms to accommodate models with industry-wide aggregate shocks. Our results suggest that, by using oblivious equilibrium to approximate Markov-perfect equilibrium, it is possible to greatly increase the set of dynamic models of imperfect competition that can be analyzed computationally.

Subject classifications: games/group decisions: stochastic; dynamic programming/optimal control: Markov.

Area of review: Special Issue on Computational Economics.

History: Received April 2007; revisions received July 2008, April 2009; accepted September 2009. Published online in *Articles in Advance*.

1. Introduction

Recently, Ericson and Pakes (1995) (hereafter EP) introduced a framework for modeling a dynamic industry with heterogeneous firms. The stated goal of this work was to facilitate empirical research analyzing the effects of policy and environmental changes on things like market structure and consumer welfare in different markets. Due to the importance of dynamics in determining policy outcomes, and also because the EP model has proved to be quite adaptable and broadly applicable, the model has received much attention in the literature. Indeed, recent work has applied the framework to studying problems as diverse as advertising, auctions, collusion, consumer learning, environmental policy, firm mergers, industry dynamics, limit order markets, network externalities, and R&D investment.¹

Despite this activity, there remain some substantial hurdles in the application of EP-type models. Because EP-type models are analytically intractable, their analysis relies on solving numerically for Markov-perfect equilibrium (MPE) on a computer, and this computation suffers from the curse of dimensionality. 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. The *industry state* is a vector encoding the number of firms with each possible value of the

firm state variable. At each time, each firm makes a decision based on its firm state and the industry state, and its subsequent firm state is determined by the current state, the current decision, and a random shock. Even if firms are restricted to symmetric strategies, the number of relevant industry states (and thus, the computer time and memory required for computing a MPE) becomes enormous very quickly as the numbers of firms and firm states grows. For a model with just 20 firms and 40 firm states, there are quadrillions of industry states. This renders commonly used dynamic programming algorithms infeasible in many problems of practical interest (see the related literature section for a discussion on some of these methods).

As a result, computational concerns have typically limited 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 (the original intent) has been rare (see Gowrisankaran and Town 1997, Benkard 2004, Jenkins et al. 2004, Ryan 2005, Collard-Wexler 2006).

In a recent paper (Weintraub et al. 2008b), we introduced a new notion of equilibrium for EP-type models called *oblivious equilibrium* (OE) that has the attractive feature that it is not subject to the curse of dimensionality. In an OE, each firm makes decisions based only on its own firm state and the long-run average industry state that will prevail in equilibrium, while ignoring the current industry state. Because firms' decisions are not a function

of the industry state, computing an OE only requires solving single-dimensional dynamic programming problems. The main result of Weintraub et al. (2008b) is an asymptotic theorem that establishes conditions under which OE well approximates MPE asymptotically as the market size grows. Intuitively, in a large market the random evolution of individual firms will average out, such that variations in the normalized industry state are small in equilibrium. Given this, firms can make near-optimal decisions based on the average equilibrium industry state rather than the current industry state.

However, although the results of Weintraub et al. (2008b) motivate consideration of OE as a solution concept, they are incomplete. First, to accommodate practice, we need methods for computing OE. In this paper we develop an algorithm for computing OE, and demonstrate its efficiency through computational experiments. A nice feature of the algorithm is that, unlike existing methods (see §2 for references), there is no need to place a priori restrictions on the number of firms in the industry or the number of allowable states per firm. These are determined by the algorithm as part of the equilibrium solution. We find that the algorithm is typically able to compute OE in less than a minute, even for industries with thousands of firms, a task that is far beyond what is computationally feasible for MPE computation with common dynamic programming algorithms.

Second, although the asymptotic theorem in our previous paper provides conditions for which OE approximates MPE well as the market size increases, it does not guarantee a good approximation in empirical applications whose market size is observed and fixed. Thus, again to accommodate practice, once an OE is obtained for a particular problem instance, we need methods to verify its accuracy as an approximation to MPE. To address this issue, in this paper we derive bounds on the approximation error that can be computed for each problem instance. Specifically, we measure error in terms of the expected incremental value that an individual firm in the industry can capture by unilaterally deviating from the OE strategy to an optimal Markov best response. The bounds can be computed using an efficient simulation-based algorithm that requires knowledge only of the OE strategies. Our algorithm for bounding approximation error is novel and represents a significant contribution. Not only is it important for the practical use of our model and solution concepts, but we believe the ideas should generalize to a broader class of games and large-scale stochastic control problems.

The bounds on the approximation error give us the ability to evaluate whether OE provides close approximations to MPE in problem instances of practical interest. The third contribution of the paper is to show through a computational study that OE does indeed offer useful approximations for many relevant industries. Specifically, we find that the approximation is often good for industries involving hundreds of firms, and in some cases even tens of firms.

These results support the conclusion that OE can be useful in empirical applications.

Finally, it is natural to think about extending the basic notion of OE in many directions, and we provide one important extension here. In the model of Weintraub et al. (2008b), all the random shocks were assumed to be idiosyncratic across firms. However, in many problems of practical interest, it is important to incorporate shocks that are common to all firms in a market. These “aggregate” shocks are important, for example, when analyzing the dynamic effects of industry-wide business cycles (see Dunne et al. 1988, Davis et al. 1998). As a fourth contribution, we extend the model and also the computational algorithms and error bounds discussed earlier to accommodate aggregate shocks. We show through a computational study that OE can also offer useful approximations for many relevant models that incorporate aggregate shocks.

Our results suggest that OE opens the door to a much broader range of applications for EP-type models. Indeed, our algorithms have already been applied in an empirical study of R&D investment in the Korean electric motor industry (Xu 2006) and in a study of the impact of advertising regulation in the cigarette industry (Qi 2008). We have also done a computational study to determine conditions under which an industry becomes fragmented or remains concentrated as the market size grows (Weintraub et al. 2005a). These studies would not have been possible using exact computation of MPE.

It is worth mentioning that, in practice, problem instances will arise in which the bounds suggest that OE may not offer an accurate approximation. Our hope is that in such cases it will be possible to extend the basic notion of OE further in order to improve the approximation at the expense of computational cost. It is encouraging that OE has already proven useful in many cases, and this motivates further work to design extensions that address additional cases. Some work along these lines is presented in Weintraub et al. (2007) and further discussed in the conclusions.

Finally, we note that although our emphasis is on the use of OE as an approximation to MPE, in many cases OE 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 paper is organized as follows. In §2 we discuss related literature. In §3 we describe the dynamic industry model. In §4 we introduce the concept of oblivious strategies and oblivious equilibrium. Note that §§3 and 4 primarily restate assumptions and definitions from our previous work, and are included here for completeness. In §5 we provide a new method for computing OE. In §6 we derive several novel error bounds. In §7 we report the results of the computational experiments. In §8 we extend OE to treat a model with aggregate shocks. Section 9 presents conclusions and a discussion of future research directions.

2. Related Literature

We discuss the relation between our approach and some past work (see also Weintraub et al. 2008b).

In the past literature on EP models, MPE are usually computed using iterative dynamic programming algorithms (e.g., Pakes and McGuire 1994). However, as discussed above, computational requirements grow with the number of industry states, making dynamic programming infeasible in many problems of practical interest. With this motivation, Judd (1998), Pakes and McGuire (2001), and Doraszelski and Judd (2006) have proposed methods that accelerate MPE computation that are capable of addressing models with several additional firms. In this paper, we take a different tack and consider algorithms that can efficiently deal with any number of firms, but aim to compute an approximation rather than an exact MPE and to bound its error. Our view is that for most models of practical interest, exact computation of MPE is unlikely to ever become feasible, and given that, approximations may offer the best available guidance for policy and strategy decisions.

OE-type approximations are based on state aggregation. Each firm predicts expected discounted profits based on partial information about the current state and uses a piecewise-constant approximation to the value function. The broader approximate dynamic programming literature makes use of other families of functions such as linear combinations of arbitrary basis functions and nonlinearly parameterized approximators. An alternative approach to the one offered in this paper is to make use of such approximations and various approximate dynamic programming algorithms to address EP-type models (see Farias et al. 2008, de Farias and Van Roy 2003).

In a theory literature related to our work, several papers have studied dynamic rational expectations equilibrium models of competition, but without explicitly modeling agents' heterogeneity. For example, Deaton and Laroque (1996) and Routledge et al. (2000) use this type of model to study the impact of speculative storage in commodity price dynamics. Because heterogeneity among agents is not explicitly modeled, the rational expectations equilibrium is described by a low-dimensional dynamic programming problem that does not suffer from the curse of dimensionality.

Most closely related to our work is Hopenhayn (1992) who develops a competitive equilibrium model of industry dynamics with heterogeneous firms. In every period firms receive productivity shocks, make entry and exit decisions, and compete in a perfectly competitive product market.² The industry holds a continuum of firms, each of which garners an infinitesimal fraction of the market, and productivity shocks generate heterogeneity among firms. The model is tractable because the industry state is constant over time, implicitly assuming that a law of large numbers holds. This assumption is based on the same intuition that motivates our consideration of OE. However, there are some notable differences between our approach and Hopenhayn (1992).

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.

For asymptotically large markets we have shown that OE in EP-type models coincide with equilibria of Hopenhayn-type models (see Weintraub et al. 2008a), so in that sense the two concepts are similar. However, because OE is an equilibrium concept that is applied directly to a finite industry, whereas Hopenhayn-type models consider an infinite number of firms, the former offers approximations superior to the latter for finite markets. Our computational results show that OE can provide accurate approximations even in industries with tens of firms. In these cases, a model with a continuum of firms is likely to provide less accurate approximations. Also, note that our algorithms to compute OE, the error bounds, and the extension to aggregate shocks are all novel in the context of analyzing industry dynamics.

Finally, the notion that strategies can remain effective by considering the aggregate behavior of the competitors only is related to the literature in aggregative games (see Dindos and Mezzetti 2006, Novshek 1985). These are static games for which equilibrium strategies can be represented as simple functions of the aggregate competitors' behavior, in particular, the sum of everybody else's actions. This differs from our context, where equilibrium strategies are in reality very complex, but can sometimes be well approximated using simple functions.

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 identical to the one in Weintraub et al. (2008b), which in turn is close in spirit to Ericson and Pakes (1995). We restate it here for completeness. The model is general enough to encompass numerous applied problems in economics (see above for examples). The basic model includes only idiosyncratic shocks. In §8 we extend the model to incorporate aggregate shocks that are common to all firms.

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 filtration $\{\mathcal{F}_t; t \geq 0\}$. We adopt a convention of indexing by t variables that are \mathcal{F}_t -measurable.

Each firm that enters the industry is assigned a unique positive integer-valued index. The set of indices of incumbent firms at time t is denoted by S_t . 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 $\mathcal{S} = \{s \in \mathbb{R}_+^\infty \mid \sum_{x=0}^\infty s(x) < \infty\}$. 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.

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 $x_{it} \in \mathbb{N}$ and its competitors' state $s_{-i,t} \in \mathcal{S}$.

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

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

where the function h 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 because h 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.³

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

In each period, events occur in the following order:

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

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 h , 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.

3.3. Assumptions

We make several assumptions about the model primitives, beginning with the profit function.

ASSUMPTION 3.1. For all $s \in \mathcal{S}$, $\pi(x, s)$ is increasing in x . Furthermore, $\sup_{x,s} \pi(x, s) < \infty$.

The assumption is natural. It ensures that increases in quality lead to increases in profit and that profits are bounded.

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

ASSUMPTION 3.2.

1. The random variables $\{\phi_{it} \mid t \geq 0, i \geq 1\}$ are i.i.d. and have finite expectations and well-defined density functions with support \mathbb{R}_+ .
2. The random variables $\{\zeta_{it} \mid t \geq 0, i \geq 1\}$ are i.i.d. and independent of $\{\phi_{it} \mid t \geq 0, i \geq 1\}$.
3. For all ζ , $h(\iota, \zeta)$ is nondecreasing in ι .
4. For all $\iota > 0$, $\mathcal{P}[h(\iota, \zeta_{i,t+1}) > 0] > 0$.
5. There exists a positive constant $\bar{h} \in \mathbb{N}$ such that $|h(\iota, \zeta)| \leq \bar{h}$, for all (ι, ζ) . There exists a positive constant $\bar{\iota}$ such that $\iota_{it} \leq \bar{\iota}$, $\forall i, \forall t$.
6. For all $k \in \{-\bar{h}, \dots, \bar{h}\}$, $\mathcal{P}[h(\iota, \zeta_{i,t+1}) = k]$ is continuous in ι .
7. The transitions generated by $h(\iota, \zeta)$ 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 (2010) that ensures a unique solution to the firms' investment decision problem. In particular, it ensures the firms' investment decision problem is strictly concave or that the unique maximizer is a corner solution. The assumption is used to guarantee 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 an asymptotically large number of potential entrants who play a symmetric mixed entry strategy. This results in a Poisson-distributed number of entrants (see Weintraub et al. 2008b for a derivation of this result). Our associated modeling assumptions are as follows:

ASSUMPTION 3.3.

1. *The number of firms entering during period t is a Poisson random variable that is conditionally independent of $\{\phi_{it}, \zeta_{it} \mid t \geq 0, i \geq 1\}$, 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 simpler dynamics. However, our results can accommodate other entry processes as well. Assumption 3.3.2 ensures that the sell-off value by itself is not 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 Weintraub et al. (2008b), 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 $V(x, s \mid \mu', \mu, \lambda)$ 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,

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

where i is taken to be the index of a firm at quality level x at time t , τ_i is a random variable representing the time at which firm i exits the industry, and the subscripts of the expectation indicate the strategy followed by firm i , the strategy followed by its competitors, and the entry rate function. In an abuse of notation, we will use the shorthand, $V(x, s \mid \mu, \lambda) \equiv V(x, s \mid \mu, \mu, \lambda)$, 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 of 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 an MPE:

$$\sup_{\mu' \in \mathcal{M}} V(x, s \mid \mu', \mu, \lambda) = V(x, s \mid \mu, \lambda) \quad \forall x \in \mathbb{N}, \forall s \in \bar{\mathcal{S}}. \quad (1)$$

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

$$\sum_{s \in \bar{\mathcal{S}}} \lambda(s) (\beta \mathbb{E}_{\mu, \lambda} [V(x^e, s_{-i,t+1} \mid \mu, \lambda) \mid s_t = s] - \kappa) = 0$$

$$\beta \mathbb{E}_{\mu, \lambda} [V(x^e, s_{-i,t+1} \mid \mu, \lambda) \mid s_t = s] - \kappa \leq 0 \quad \forall s \in \bar{\mathcal{S}}$$

$$\lambda(s) \geq 0 \quad \forall s \in \bar{\mathcal{S}}.$$

In Weintraub et al. (2008b), 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 (2010) establish existence of an equilibrium in pure strategies for a closely related model. We do not provide an existence proof here because it is long and cumbersome and would replicate this previous work. With respect to uniqueness, in general we presume that our model may have multiple equilibria.⁴

Dynamic programming algorithms can be used to optimize firm strategies and equilibria to our model can be computed via their iterative application. Stationary points of such iterations are MPE, but there is no guarantee of convergence. Nevertheless, such methods have proven to be effective in practical contexts and consequently have seen broad use. However, they require computing time and memory that grow proportionately with the number of relevant industry states, which is intractable in many applications. This difficulty motivates our alternative approach.

4. Oblivious Equilibrium

We now formally define the concept of oblivious equilibrium as in Weintraub et al. (2008b), restated here for completeness. The motivation for considering OE is 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* because 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. Because 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*

$$\lim_{t \rightarrow \infty} \tilde{s}_t(x) = \lambda \sum_{k=0}^{\infty} P_\mu^k(x^e, x), \quad (2)$$

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

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

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

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

using $\tilde{V}(x | \mu, \lambda) \equiv \tilde{V}(x | \mu, \mu, \lambda)$ to refer to the oblivious value function when firm i follows the same strategy μ as its competitors.

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

1. Firm strategies optimize an oblivious value function:

$$\sup_{\mu' \in \tilde{\mathcal{M}}} \tilde{V}(x | \mu', \mu, \lambda) = \tilde{V}(x | \mu, \lambda), \quad \forall x \in \mathbb{N}. \quad (3)$$

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

$$\lambda(\beta \tilde{V}(x^e | \mu, \lambda) - \kappa) = 0$$

$$\beta \tilde{V}(x^e | \mu, \lambda) - \kappa \leq 0$$

$$\lambda \geq 0.$$

It is straightforward to show that OE exists under mild technical conditions. Furthermore, if the entry cost is not prohibitively high relative to single-period profits, 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 OE 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.⁵

Finally, in Weintraub et al. (2008b) we show that when strategies and entry rate functions are oblivious, the Markov process $\{s_t; t \geq 0\}$ admits a unique invariant distribution. Moreover, we show that when firms play OE strategies, the invariant distribution of the industry state is such that the number of firms in each state $x \in \mathbb{N}$ is given by a Poisson random variable with mean $\tilde{s}_{\mu, \lambda}(x)$, independent across states $x \in \mathbb{N}$.

5. Algorithm

In this section we propose an algorithm to solve for OE.

Algorithm 1 (below) is designed to compute an OE with a positive entry rate. It starts with two extreme entry rates: $\underline{\lambda} = 0$ and $\bar{\lambda} = (1/\kappa)(\sup_{x,s} \pi(x, s)/(1-\beta) + \bar{\phi})$. Under mild assumptions, any oblivious equilibrium entry rate must lie between these two extremes. The algorithm searches over entry rates between these two extremes for one that leads to an OE. For each candidate entry rate λ , an inner loop (steps 6–10) computes an OE firm strategy for that fixed entry rate. Strategies are updated “smoothly” (step 9).⁶ If the termination conditions of both the inner and outer loops are satisfied with $\epsilon_1 = \epsilon_2 = 0$, we have an OE. Small values of ϵ_1 and ϵ_2 allow for small errors associated with limitations of numerical precision.

Algorithm 1 (Oblivious equilibrium solver)

```

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

The algorithm is easy to program and computationally efficient. In each iteration of the inner loop, the optimization problem to be solved is a one-dimensional dynamic program. The state space in this dynamic program is the set of quality levels a firm can achieve. In principle, there could be an infinite number of them. However, beyond a certain quality level the optimal strategy for a firm is not to invest, so its quality cannot increase to beyond that level. In the numerical experiments we present in §7, the state space never had more than 100 states per firm. The exact number of states is determined during execution of the algorithm.

Although there are alternative algorithms to solve for OE, we chose the one previously described, motivated by a few observations:

1. For $\lambda \in \tilde{\Lambda}$, let $\mu^* \in \tilde{\mathcal{M}}$ be the OE strategy associated with that entry rate. In our numerical experiments, we observed that $\tilde{V}(x^e | \mu^*, \lambda)$ is decreasing in λ . Therefore, the line search method proposed provides an efficient way to find an entry rate that satisfies the zero oblivious expected value condition and yields an OE.⁷

2. For given $\lambda \in \tilde{\Lambda}$ we use a myopic best-response algorithm to find the OE strategy associated with that entry rate. Although in many cases, including ours, the theoretical convergence properties of these algorithms are not well understood (Fudenberg and Levine 1998), positive practice experience supports their use. Moreover, our computational experiments showed that a smooth update of strategies (step 9) speeds up convergence.

Whether this algorithm is guaranteed to terminate in a finite number of iterations remains an open issue. However, in the numerical experiments we present in the next section, it always terminated in fewer than 30 seconds.⁸

6. Error Bounds

In this section we derive expressions that can be computed via simulation and that bound approximation error

associated with a particular OE. While the asymptotic results in Weintraub et al. (2008b) provide conditions under which the approximation will work well as the market size grows, the error bound can be used to evaluate the OE as an approximation of MPE for a particular set of model primitives.

To bound approximation error, we first need to define what is meant by *approximation error*. Consider an oblivious strategy and entry rate function $(\tilde{\mu}, \tilde{\lambda}) \in \tilde{\mathcal{M}} \times \tilde{\Lambda}$. We assume that the initial industry state s_0 is sampled from the invariant distribution of $\{s_t: t \geq 0\}$. Hence, s_t is a stationary process; s_t is distributed according to its invariant distribution for all $t \geq 0$. We will quantify approximation error at each firm state $x \in \mathbb{N}$ by

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

The expectation is over the invariant distribution of s_t . Hence, approximation error is the amount by which a firm at state $x \in \mathbb{N}$ can improve its expected net present value by unilaterally deviating from the OE strategy $\tilde{\mu}$, and instead following an optimal (nonoblivious) best response. Recall that an MPE requires that the expression in square brackets equals zero for all states (x, s) . Approximation error 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 useful bounds pointwise. This is because in an 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 very low probability of occurrence, so they have a small impact on expected discounted profits. The idea is that when approximation error is small, MPE strategies and entry rates at relevant states should be well approximated by oblivious ones. In §7 we present computational results that support this point.⁹

The next theorem provides two bounds on the approximation error. Recall that \tilde{s} is the long-run expected state in OE ($E[s_t]$). Let $a_x(y)$ be the expected discounted sum of an indicator of visits to state y for a firm starting at state x that uses strategy $\tilde{\mu}$. Let $[x]^+ = \max(x, 0)$ and $\underline{x}(k, t) = [x - (k - t)\bar{h}]^+$. Finally, let $|\Delta|(s) = \sup_{y \in \mathbb{N}} |\pi(y, s) - \pi(y, \tilde{s})|$ and $\Delta_A(s) = \sup_{y \in A} (\pi(y, s) - \pi(y, \tilde{s}))$.

THEOREM 6.1. *Let Assumptions 3.1, 3.2, and 3.3 hold. Then, for any OE $(\tilde{\mu}, \tilde{\lambda})$ and firm state $x \in \mathbb{N}$,*

$$E \left[\sup_{\mu' \in \mathcal{M}} V(x, s_t | \mu', \tilde{\mu}, \tilde{\lambda}) - V(x, s_t | \tilde{\mu}, \tilde{\lambda}) \right] \leq \frac{2}{1-\beta} E[|\Delta|(s_t)], \quad (4)$$

and

$$E \left[\sup_{\mu' \in \mathcal{M}} V(x, s_t | \mu', \tilde{\mu}, \tilde{\lambda}) - V(x, s_t | \tilde{\mu}, \tilde{\lambda}) \right] \leq \sum_{k=t}^{\infty} \beta^{k-t} E \left[[\Delta_{\{\underline{x}(k,t), \dots, x+(k-t)\bar{h}\}}(s_k)]^+ \right] + \sum_{y \in \mathbb{N}} a_x(y) (\pi(y, \tilde{s}) - E[\pi(y, s_t)]). \quad (5)$$

The derivation of these bounds can be found in the electronic companion to this paper that is available as part of the online version at <http://or.journal.informs.org>. It is worth mentioning that the result can be generalized a great deal. In particular, many of the prior assumptions can be dropped; for instance, most alternative entry processes will not change the result.

The first bound is simpler, so we will use it to provide an explanation of the main steps of the derivation here. First, we compare the value functions in the definition of approximation error through the OE value function. Formally,

$$\begin{aligned} & \mathbb{E}[V(x, s_t | \mu^*, \tilde{\mu}, \tilde{\lambda}) - V(x, s_t | \tilde{\mu}, \tilde{\lambda})] \\ &= \mathbb{E}[V(x, s_t | \mu^*, \tilde{\mu}, \tilde{\lambda}) - \tilde{V}(x | \tilde{\mu}, \tilde{\lambda})] \\ &+ \mathbb{E}[\tilde{V}(x | \tilde{\mu}, \tilde{\lambda}) - V(x, s_t | \tilde{\mu}, \tilde{\lambda})], \end{aligned} \quad (6)$$

where $\mu^* \in \mathcal{M}$ is a Markovian (nonoblivious) best response to an OE $(\tilde{\mu}, \tilde{\lambda})$ for a firm that is keeping track of the industry state. We now explain how we bound the first expectation in the right-hand side above. A similar argument can be used for the second one. First, we observe that because $\tilde{\mu}$ and $\tilde{\lambda}$ attain an OE, for all x ,

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

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,

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

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 $V(x, s | \mu^*, \tilde{\mu}, \tilde{\lambda})$ with $\tilde{V}(x | \tilde{\mu}, \tilde{\lambda})$ by only taking into consideration the difference between single-period profits (actual versus oblivious). Because strategies are the same, the terms associated with expected investments and sell-off value cancel out. Hence, using the previous inequality, we obtain that the difference between value functions can be bounded by a discounted sum of expected differences between actual and oblivious single-period profits as required by the error bounds. To obtain a bound that does not depend on μ^* , we use the fact that under OE strategies firms' trajectories are independent.

By doing a more careful account on profits' differences and using the fact that a firm can change by at most \bar{h} quality units per time period, we obtain the second bound that is tighter. Note that the right-hand side of the second bound depends on the initial firm state x , whereas the right-hand side of the first bound does not.

Both bounds can be easily estimated via simulation algorithms. Computing the bounds involves computing

expectations over the industry state s_t under its invariant distribution. Once the OE has been computed, the industry state has a known distribution, namely, the product form of Poisson random variables with mean \tilde{s} (see Weintraub et al. 2008b). In particular, note that the bounds are not a function of the true MPE or even of the optimal nonoblivious best-response strategy. Computing either of these strategies could require solving a high-dimensional dynamic program.

If the dynamics of the model include depreciation—that is, there is a positive probability the quality level of the firm goes down even if investment is arbitrarily large—tighter bounds can be derived. Let $\Delta(y, s) = \pi(y, s) - \pi(y, \tilde{s})$. Let $\hat{\mu}$ be a strategy such that the firm never exits the industry and invests an infinite amount at every state, and let $\{\hat{x}_t: t \geq 0\}$ be a process that describes the state evolution of a firm that uses strategy $\hat{\mu}$. We have the following result which we prove in the electronic companion.

THEOREM 6.2. *Let Assumptions 3.1, 3.2, and 3.3 hold. Suppose that, for all $s \in \mathcal{F}$, the function $\Delta(y, s)^+$ is nondecreasing in y . Then, for all OE $(\tilde{\mu}, \tilde{\lambda})$, and firm state $x \in \mathbb{N}$,*

$$\begin{aligned} & \mathbb{E}\left[\sup_{\mu' \in \mathcal{M}} V(x, s_t | \mu', \tilde{\mu}, \tilde{\lambda}) - V(x, s_t | \tilde{\mu}, \tilde{\lambda})\right] \\ & \leq \sum_{k=t}^{\infty} \beta^{k-t} \mathbb{E}_{\tilde{\mu}, \tilde{\mu}, \tilde{\lambda}}[\Delta(\hat{x}_k, s_k)^+ | \hat{x}_t = x] \\ & \quad + \sum_{y \in \mathbb{N}} a_x(y) (\pi(y, \tilde{s}) - \mathbb{E}[\pi(y, s_t)]). \end{aligned} \quad (7)$$

As before, s_k is distributed according to the invariant distribution for all $k \geq 0$. The expectation over \hat{x}_k can be written in closed form for the model in §7.1, facilitating its computation (see the electronic companion for this result). If there is depreciation and $\Delta(y, s)^+$ is nondecreasing in y , as is the case in the model we introduce below, bound (7) is generally tighter than bound (5). The latter takes a maximum over achievable states in the first sum. The former takes an expectation with respect to $\hat{\mu}$, and because of depreciation, larger achievable states have smaller weights, reducing the magnitude of the bound.

Even tighter bounds can be derived for industries where there is no exit of incumbent firms and no entry of new firms. These bounds are described in the electronic companion.

7. Computational Experiments

In this section we conduct some computational experiments to evaluate how OE performs in practice. We begin with the model to be analyzed. The model is similar to Pakes and McGuire (1994). However, it differs in the entry and exit processes, in the demand system, and in that we do not consider an aggregate shock.¹⁰ The model satisfies Assumptions 3.1, 3.2, and 3.3.

7.1. The Computational Model

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

$$u_{ijt} = \theta_1 \ln\left(\frac{x_{it}}{\psi} + 1\right) + \theta_2 \ln(Y - p_{it}) + v_{ijt},$$

$$i \in S_t, j = 1, \dots, m,$$

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

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

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

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

$$Y - p_{it} + \theta_2(p_{it} - c)(\sigma(x_{it}, s_{-i,t}, p_t) - 1) = 0, \quad \forall i \in S_t. \quad (8)$$

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

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

Sell-Off Price. ϕ_{it} are i.i.d. exponential random variables with mean K .

Transition Dynamics. A firm's investment is successful with probability $au/(1+au)$, in which case the quality of its product increases by one level. The firm's product depreciates one quality level with probability δ , independently each period. Note that our model differs from Pakes and McGuire (1994) here because the depreciation shocks in our model are idiosyncratic. Combining the investment and depreciation processes, it follows that the transition probabilities for a firm in state x that does not exit and invests ι are given by:

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

7.2. Numerical Results: Behavior of the Bound

Our first set of results investigates the behavior of the approximation error bound under several different model specifications. A wide range of parameters for our model could reasonably represent different real-world industries of interest. In practice, the parameters would either be estimated using data from a particular industry or chosen to reflect an industry under study. We begin by investigating a particular set of representative parameter values. Following Pakes and McGuire (1994), we fix $a = 3$ and $\delta = 0.7$. Additionally, we fix marginal cost at $c = 0.5$, income at $Y = 1$, $\theta_2 = 0.5$, and $\psi = 1$. The discount factor is $\beta = 0.95$. The entry cost is $\kappa = 35$ and the entry state is $x^e = 10$. The average sell-off value is $K = 10$. In this case, $\beta \cdot \bar{\phi} < \kappa$, so the sell-off value by itself is not sufficient reason to enter the industry (Assumption 3.3.2). Additionally, both sell-off values and entry costs are substantially larger than marginal costs, consistent with empirical evidence.

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

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

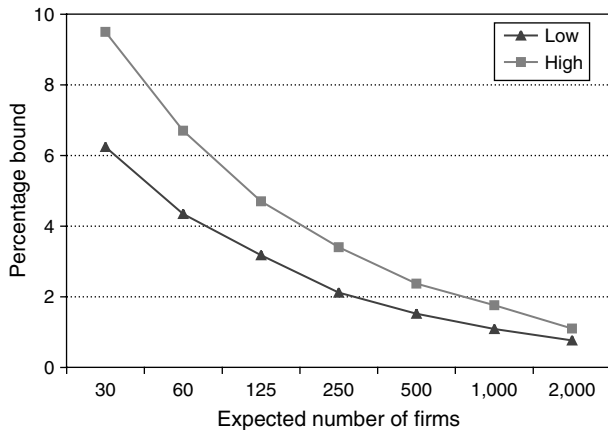
For each set of parameters, we use the approximation error bound in Theorem 6.2 to compute an upper bound on the percentage error in the value function,

$$\frac{E[\sup_{\mu' \in \mathcal{M}} V(x, s \mid \mu', \tilde{\mu}, \tilde{\lambda}) - V(x, s \mid \tilde{\mu}, \tilde{\lambda})]}{E[V(x, s \mid \tilde{\mu}, \tilde{\lambda})]},$$

where $(\tilde{\mu}, \tilde{\lambda})$ are the OE strategy and entry rate, respectively, and the expectations are taken with respect to s .¹¹ We estimate the expectations using simulation.¹² We compute the previously mentioned percentage approximation error bound for different market sizes. As the market size increases, the expected number of firms increases and the approximation error bound decreases.

In Figure 1 we present the percentage approximation error bound as a function of the expected number of firms for the two levels of vertical differentiation (the two curves

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



are obtained by varying the market size). For the low vertical differentiation case, it takes around 150 firms to bring the bound down to 3%, and 250 firms to bring it to 2%. For the high case, it takes around 200 firms to bring the bound to 4% and 700 firms to bring it to 2%.


When the level of vertical differentiation is high, the number of firms required to have a good approximation is large, requiring hundreds and even thousands of firms. The approximation would be better if the industry state s was always close to its mean, \bar{s} . One aspect of the model that interferes with this is the Poisson entry process, which leads to a large amount of variability in the number of firms inside the industry. Recall that we chose to model the entry process this way because it simplified the dynamics. However, the expressions for the approximation error bounds remain correct for a wide range of entry models. To investigate this issue further, as an alternative, we tried using an entry process where the number of entrants each period is “almost deterministic,” but still satisfies a zero-profits condition.¹³ This entry process implies a smaller variability in the number of firms.

Figure 2 presents the results with the new entry process. In the case of low vertical differentiation, the approximation error bound is around 3% with just 60 firms, around 2% with 125 firms, and around 1% with 500 firms. When the level of vertical differentiation is high, the approximation error bound is around 3% when there are 125 firms and around 2% for 350 firms.

Going one step further in reducing the variability of the industry dynamics, we tried shutting down entry and exit altogether and considered an industry with a fixed number of firms. We used the error bound in Corollary B.2 in the electronic companion. See Figure 3 for the results.¹⁴ For the low case, the approximation error bound is less than 0.5% with just 5 firms, whereas for the high case it is 2% for 20 firms, and around 1% with 100 firms.

Most economic applications would involve from fewer than ten to several hundred firms. These results show that

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




the approximation error bound may sometimes be small (<2%) in these cases, although this would depend on the model and parameter values for the industry under study.

7.3. Closeness to MPE Economic Indicators of Interest

Having gained some insight into what features of the model lead to low values of the approximation error bound, the question arises as to what value of the error bounds is required to obtain a good approximation of economic indicators like the ones in which researchers are usually interested. To shed light on this issue, we compare long-run statistics for the same industry primitives under OE and MPE strategies. A major constraint on this exercise is that it requires the ability to actually compute the MPE. With the current methods we are able to compute MPE for industries with a maximum of 5 to 10 firms. Because we require the ability to compute equilibria for many different parameter values, to keep computation manageable we use

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




4 firms here. We therefore limit our analysis to the case of a fixed number of firms (no entry and exit), because only for that case were the approximation error bounds small under oblivious strategies (with only 4 firms). We use the same parameter specifications as in the previous subsection. Because of computational constraints in computing the MPE, we also impose a maximum state that a firm can reach of $x_{\max} = 15$, at which point investment is assumed to have no further effect. The market size is fixed, $m = 30$.¹⁵

Recall that under OE strategies, the industry state is described by an ergodic Markov process (see Weintraub et al. 2008b). Under our assumptions, this is also true under MPE strategies (see also Ericson and Pakes 1995). Therefore, both systems have a well-defined invariant distribution that describe their long-run behavior. We compare the expected values of several economic statistics of interest with respect to the OE and the MPE invariant distributions. The quantities compared are: average investment, average producer surplus, average consumer surplus, average share of the largest firm (C1), and average share of the largest two firms (C2). Table 1 reports these statistics for a wide range of parameters. The table also reports the maximum value (across all states) and weighted average value (according to the invariant distribution) of the approximation percentage error bound, as well as the maximum and weighted average of the actual benefit from deviating and keeping track of the industry state (the actual difference $(E[\sup_{\mu' \in \mathcal{M}} V(x, s | \mu', \tilde{\mu}, \tilde{\lambda}) - V(x, s | \tilde{\mu}, \tilde{\lambda})]) / (E[V(x, s | \tilde{\mu}, \tilde{\lambda})])$). Note that the latter quantity should always be smaller than the approximation error bound. In the results below we concentrate on the maximum value of the error bound across all states.

The table is separated into two groups. The first five rows correspond to industries with a relatively low cost of investment (low value of d relative to θ_1). In these industries the industry state tends to have a symmetric distribution (see Figure 4) reflecting a rich investment process. The last five rows of the table correspond to industries with a relatively high cost of investment. In these industries the

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



industry state tends to be skewed (see Figure 5), reflecting low levels of investment.

From the computational experiments we conclude the following:

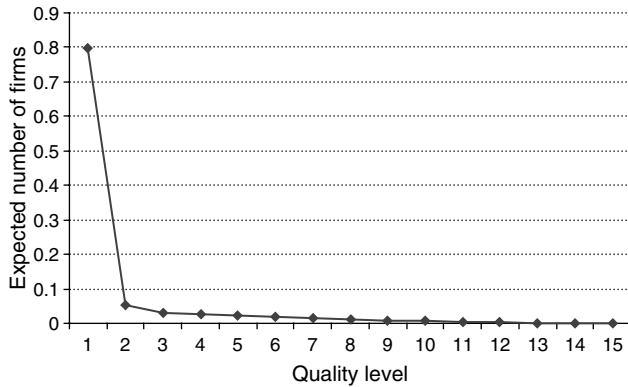
1. When the bound is less than 1% the long-run quantities estimated under OE and MPE strategies are very close. The relation between long-run economic indicators and the error bound can be seen more clearly in Figure 6, where we plot the results for the industries with a relatively high cost of investment. We observe that if the error bound is small, the differences in long-run average producer and consumer surplus under OE and MPE strategies are small. As the level of vertical differentiation increases, the error bounds increase, and the differences increase as well. The results suggest that error bounds provide a useful indicator of whether economic quantities of interest are being well approximated.

2. Performance of the approximation depends on the richness of the equilibrium investment process. When the bound is between 1%–20% and there is a rich investment process, the long-run quantities estimated under OE and MPE strategies are still quite close. When the bound is

Table 1. Comparison of MPE and OE strategies (four firms, no entry and exit).

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

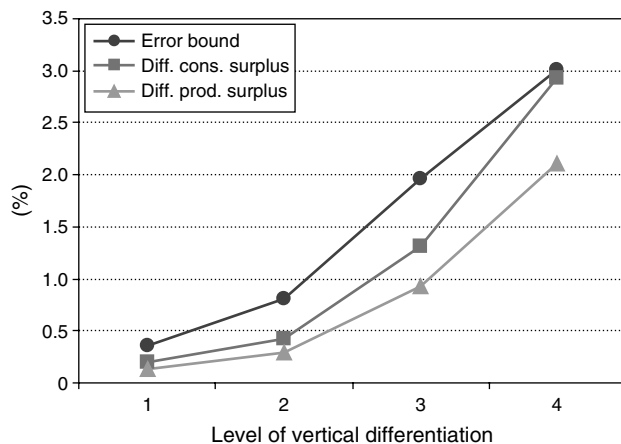
Notes. Long-run statistics and value functions simulated with a relative precision of 1.0% and a confidence level of 99%. Error bound simulated with a relative precision of at most 10% and a confidence level of 99%.

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

above 1% and there is little investment, the long-run quantities can be quite different on a percentage basis (5% to 20%), but still remain fairly close in absolute terms (see Table 2).

3. The performance bound is not tight. For a wide range of parameters the performance bound is as much as 10 to 20 times larger than the actual benefit from deviating.

The previous results suggest that economic indicators of interest under MPE strategies are well approximated with OE strategies when the approximation error bound is small (less than 1%–2%, and in some cases even up to 20%). These results, together with those from §7.2, demonstrate that the OE approximation significantly expands the range of applied problems that can be analyzed computationally.

Figure 6. Percentage differences between OE and MPE long-run average consumer and producer surplus and percentage error bound.

Notes. The four levels of vertical differentiation correspond to the first four rows of the bottom part of Table 1, that is, to industries with a relatively high cost of investment. As the level of vertical differentiation increases from 1 to 4, θ_1 and d increase.

Table 2. Comparison of MPE and OE investment (four firms, no entry and exit).

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

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

8. Aggregate Shocks

Our base model does not allow for shocks to firm profitability that are common across firms. In some contexts, for example, when studying how industry dynamics evolve over a business cycle, it is important to account for aggregate shocks. In this section, we extend OE and our computational methods to allow for aggregate shocks at the expense of greater modeling and computational complexity.

In §8.1 we introduce a dynamic industry model with aggregate shocks. In §8.2 we extend the notion of OE to accommodate aggregate shocks. Error bounds for this model are introduced in §8.3. We provide computational experiments in §8.4.

8.1. Model with Aggregate Profit Shocks

In this section we extend the model in §3 to incorporate a profit shock, z_t , that is common to all firms in the industry. z_t might represent a common demand shock, a common shock to input prices, or a common technology shock. These common shocks will serve to generate periods over which profits are high (or low) for all firms in the industry simultaneously.

The following assumption defines the aggregate shock process and replaces our earlier assumptions on the profit function.

ASSUMPTION 8.1. Let $Z = \{z_t \in A: t \geq 0\}$ be a finite and ergodic Markov chain. Single-period profits for firm i at time t are given by $\pi(x_{it}, s_{-i,t}, z_t)$. For all z , $\pi(x, s, z)$ satisfies Assumption 3.1. Additionally, $\sup_{x,s,z} \pi(x, s, z) < \infty$.

In this model, a strategy is a function $\mu(x, s, z)$ that depends on the firm's own state, the competitors' state, and the level of the aggregate shock. An entry rate is a function, $\lambda(s, z)$, that depends on the industry state and the level of the aggregate shock. To formalize these notions, let \mathcal{M}_z denote the set of exit/investment strategies such that an element $\mu \in \mathcal{M}_z$ is a pair of functions $\mu = (\iota, \rho)$, where $\iota: \mathbb{N} \times \mathcal{S} \times A \rightarrow \mathfrak{R}_+$ is an investment strategy and

$\rho: \mathbb{N} \times \mathcal{S} \times A \rightarrow \mathfrak{R}_+$ is an exit strategy. We denote the set of entry rate functions by Λ_z , where an element of Λ_z is a function $\lambda: \mathcal{S} \times A \rightarrow \mathfrak{R}_+$.

Define the value function, $V(x, s, z | \mu', \mu, \lambda)$, to be the expected net present value for a firm at state x when its competitors' state is s , and the value of the aggregate shock is z , given that its competitors each follow a common strategy $\mu \in \mathcal{M}_z$, the entry rate function is $\lambda \in \Lambda_z$, and the firm itself follows strategy $\mu' \in \mathcal{M}_z$. Because this definition is analogous to the one in §3.4, we omit the details here for brevity. An equilibrium in this model comprises an investment/exit strategy $\mu = (\iota, \rho) \in \mathcal{M}_z$, and an entry rate function $\lambda \in \Lambda_z$ such that:

1. Incumbent firm strategies represent a MPE:

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

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

$$\sum_{s \in \bar{\mathcal{S}}, z \in A} \lambda(s, z) (\beta E_{\mu, \lambda} [V(x^e, s_{-i, t+1}, z_{t+1} | \mu, \lambda) | s_t = s, z_t = z] - \kappa) = 0$$

$$\beta E_{\mu, \lambda} [V(x^e, s_{-i, t+1}, z_{t+1} | \mu, \lambda) | s_t = s, z_t = z] - \kappa \leq 0, \quad \forall s \in \bar{\mathcal{S}}, \forall z \in A$$

$$\lambda(s, z) \geq 0, \quad \forall s \in \bar{\mathcal{S}}, \forall z \in A.$$

8.2. Oblivious Equilibrium with Aggregate Shocks

In this section we extend the notion of OE to incorporate aggregate shocks. Recall that an OE was based on the idea that when there are a large number of firms (and no aggregate shocks), simultaneous changes in individual firm quality levels can average out such that in the long run the industry state remains roughly constant over time. Because the aggregate shocks are likely to be of first-order importance to strategies, in extending the notion of oblivious strategies to this model it makes sense to make strategies a function of the current value of the shock. In this case, even if there are a large number of firms, the industry state will not necessarily be close to a constant state; it will vary with the aggregate shock. However, we can still take advantage of averaging effects among many firms to restrict firm's strategies so that they do not depend on the industry state. Actually, if there are many firms and the time period is large, because of averaging effects, firms should be able to accurately predict the industry state based on the entire history of realizations of the aggregate shock. This is computationally impractical; instead, we will allow firms to predict the industry state based on a finite set of statistics that depend on the history of realizations of the shock.

Based on this motivation, we will restrict firm strategies so that each firm's decisions depend only on the firm's quality level, the current value of the aggregate shock, and a finite set of statistics that depend on the history of realizations of the aggregate shock. We call such restricted strategies *extended oblivious strategies*. To convey this dependence we introduce a N -dimensional Markov chain $\{w_t \in \mathcal{W} = \mathcal{W}_1 \times \dots \times \mathcal{W}_N: t \geq 0\}$. We make the following assumption.

ASSUMPTION 8.2. *We assume that $\{w_t: t \geq 0\}$ is a finite Markov chain adapted to the filtration generated by $\{z_t: t \geq 0\}$. $\{w_t: t \geq 0\}$ has a single recurrent class $\bar{\mathcal{W}} \subseteq \mathcal{W}$ and admits a unique invariant distribution. For all $t \geq 0$, $w_t(1) = z_t$.*

Note that Assumption 8.2, together with Assumptions 3.2, 3.3, and 8.1, imply that when firms use Markov strategies, $\{(s_t, w_t): t \geq 0\}$ is also a Markov chain that admits a unique invariant distribution.

Let $\bar{\mathcal{M}}_z$ and $\bar{\Lambda}_z$ denote the set of extended oblivious strategies and the set of extended oblivious entry rate functions. If firm i uses strategy $\mu \in \bar{\mathcal{M}}_z$, then at time period t , firm i takes action $\mu(x_{it}, w_t)$, where x_{it} is the state of firm i at time t . Similarly, if $\lambda \in \bar{\Lambda}_z$, then at time t , the entry rate is equal to $\lambda(w_t)$. Because by the second part of Assumption 8.2, $w_t(1) = z_t$, firms keep track of the current level of the aggregate shock when making decisions with extended oblivious strategies. The state variables $w_t(2), \dots, w_t(N)$ allow firms to incorporate additional information about the history of realizations of the aggregate shock into the strategies. As we show in §8.4, accounting for past shocks will generally improve firms' decisions. As a consequence, strategies that depend on past shocks may provide a more appealing behavioral model. It is worth mentioning here, though, that past aggregate shocks are not payoff relevant (hence, they do not influence MPE strategies), so allowing extended oblivious strategies to depend on them may give rise to extended oblivious equilibria that are poor approximations to MPE.

Different extended oblivious strategies can be defined depending on the specification of w_t . We provide a few examples below. Both examples satisfy Assumption 8.2.

EXAMPLE 8.1. Suppose that for $j \in \{1, \dots, N\}$, $w_t(j) = z_{t-j+1}$. Hence, $w_t = \{z_t, z_{t-1}, \dots, z_{t-N+1}\}$; the aggregate shocks statistics correspond to the last realizations of the shock.

One disadvantage of the previous scheme is that realizations of the shock that appear in a certain window of time influence the strategy, but if a realization occurs even slightly outside this window, it has no influence. With this motivation, we introduce an alternative scheme based on exponentially weighted averages of past shocks.

EXAMPLE 8.2. Suppose that $w_t(1) = z_t$ and that for $j \in \{2, \dots, N\}$, $w_{t+1}(j) = \alpha_j g_j(z_t) + (1 - \alpha_j) w_t(j)$ and $w_0(j) = g_j(z_0)$, where $\alpha_j \in [0, 1]$ and $g_j: A \rightarrow \mathfrak{R}$.¹⁶

Suppose that all firms use a common strategy $\mu \in \tilde{\mathcal{M}}_z$ and that entry occurs according to the entry rate function $\lambda \in \tilde{\Lambda}_z$. We assume that (s_0, w_0) is distributed according to the invariant distribution of $\{(s_t, w_t): t \geq 0\}$. Hence, (s_t, w_t) is a stationary process.

Firms predict the industry state based on the current realization of w_t . Accordingly, for all $w \in \tilde{\mathcal{W}}$, we define $\tilde{s}_{\mu, \lambda}(w) = E[s_t | w_t = w]$. In words, $\tilde{s}_{\mu, \lambda}(w)$ is the long-run expected industry state when dynamics are governed by extended oblivious strategy μ and extended oblivious entry rate function λ , conditional on the current realization of w_t being w .

With some abuse of notation we define an extended oblivious value function as,

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

This value function should be interpreted as the expected net present value of a firm that is at quality level x , when the aggregate shocks statistics have value w , and the firm follows extended oblivious strategy μ' . The firm assumes that its competitors' state will be $\tilde{s}_{\mu, \lambda}(w_k)$ for all time periods k .

An *extended oblivious equilibrium* consists of a strategy $\mu \in \tilde{\mathcal{M}}_z$ and an entry rate function $\lambda \in \tilde{\Lambda}_z$ that satisfy the following conditions:

1. Firm strategies optimize an extended oblivious value function:

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

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

$$\sum_{w \in \tilde{\mathcal{W}}} \lambda(w) (\beta E[\tilde{V}(x^e, w_{t+1} | \mu, \lambda) | w_t = w] - \kappa) = 0,$$

$$\beta E[\tilde{V}(x^e, w_{t+1} | \mu, \lambda) | w_t = w] - \kappa \leq 0, \quad \forall w \in \tilde{\mathcal{W}},$$

$$\lambda(w) \geq 0, \quad \forall w \in \tilde{\mathcal{W}}.$$

In the electronic companion, we also provide an algorithm for computing an extended OE. Note that the state space of the firm's dynamic programming problem scales with the number of firm states and with the size of $\tilde{\mathcal{W}}$, the recurrent class of the aggregate shock statistics process. As the set $\tilde{\mathcal{W}}$ becomes richer, more computation time and memory is needed.

8.3. Error Bounds

We derive error bounds for this model. As before, approximation error is the amount by which a firm at state $x \in \mathbb{N}$ can improve its expected net present value by unilaterally

deviating from the extended OE strategy, and instead following an optimal (nonoblivious) best response.

Because an optimal strategy for a firm that unilaterally deviates from an extended OE strategy depends on the aggregate shock statistics, (because its competitors are using extended OE strategies), we introduce extended Markov strategies. We define \mathcal{M}_{ze} and Λ_{ze} as the set of extended Markov strategies and extended entry rate functions, respectively. An extended Markov strategy is a function of the firm's own state, the industry state, the aggregate shock, and the aggregate shock statistics. If firm i uses strategy $\mu \in \mathcal{M}_{ze}$, then at time period t , firm i takes action $\mu(x_{it}, s_{-i,t}, w_t)$. Similarly, if $\lambda \in \Lambda_{ze}$, then the entry rate at time t is $\lambda(s_t, w_t)$.¹⁷

For extended Markov strategy $\mu', \mu \in \mathcal{M}_{ze}$ and extended entry rate function $\lambda \in \Lambda_{ze}$, with some abuse of notation, we define the extended value function,

$$V(x, s, w | \mu', \mu, \lambda) = E_{\mu', \mu, \lambda} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} (\pi(x_{ik}, s_{-i,k}, z_k) - dt_{ik}) + \beta^{\tau_i-t} \phi_{i, \tau_i} \Big| x_{it} = x, s_{-i,t} = s, w_t = w \right], \quad (9)$$

where i is taken to be the index of a firm at quality level x at time t . The extended value function generalizes the value function defined in §3.4, allowing for dependence on extended strategies. We use this value function to evaluate the *actual* expected discounted profits garnered by a firm that uses an extended Markov strategy.

Consider an extended oblivious strategy and entry rate $(\tilde{\mu}, \tilde{\lambda}) \in \tilde{\mathcal{M}}_z \times \tilde{\Lambda}_z$. We assume, the initial state (s_0, w_0) is sampled from the invariant distribution of $\{(s_t, w_t): t \geq 0\}$. Hence, (s_t, w_t) is a stationary process, it is distributed according to its invariant distribution for all $t \geq 0$. To abbreviate, let $\tilde{s} = \tilde{s}_{\tilde{\mu}, \tilde{\lambda}}$. With some abuse of notation, let $\Delta_A(s, w) = \sup_{y \in A} (\pi(y, s, w(1)) - \pi(y, \tilde{s}(w), w(1)))$ and let $\Delta(y, s, w) = \pi(y, s, w(1)) - \pi(y, \tilde{s}(w), w(1))$. We have the following result that we prove in the electronic companion.

THEOREM 8.1. *Let Assumptions 8.1, 8.2, 3.2, and 3.3 hold. Then, for any extended OE $(\tilde{\mu}, \tilde{\lambda})$, and firm state $x \in \mathbb{N}$,*

$$E \left[\sup_{\mu' \in \mathcal{M}_{ze}} V(x, s_t, w_t | \mu', \tilde{\mu}, \tilde{\lambda}) - V(x, s_t, w_t | \tilde{\mu}, \tilde{\lambda}) \right] \leq \sum_{k=t}^{\infty} \beta^{k-t} E \left[[\Delta_{\{x(k,t), \dots, x+(k-t)\tilde{h}\}}(s_k, w_k)]^+ \right] + E \left[E_{\tilde{\mu}, \tilde{\lambda}} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} (\pi(x_{ik}, \tilde{s}(w_k), z_k) - \pi(x_{ik}, s_{-i,k}, z_k)) \Big| x_{it} = x, s_{-i,t} = s_t, w_t \right] \right]. \quad (10)$$

Suppose that for all $s \in \tilde{\mathcal{F}}$ and $w \in \tilde{\mathcal{W}}$, the function $\Delta(y, s, w)^+$ is nondecreasing in y . Then, for any extended

OE $(\tilde{\mu}, \tilde{\lambda})$, and firm state $x \in \mathbb{N}$,

$$\begin{aligned}
 & E \left[\sup_{\mu' \in \mathcal{M}_{z_e}} V(x, s_t, w_t | \mu', \tilde{\mu}, \tilde{\lambda}) - V(x, s_t, w_t | \tilde{\mu}, \tilde{\lambda}) \right] \\
 & \leq \sum_{k=t}^{\infty} \beta^{k-t} E_{\tilde{\mu}, \tilde{\mu}, \tilde{\lambda}} [\Delta(\hat{x}_k, s_k, w_k)^+ | \hat{x}_t = x] \\
 & \quad + E \left[E_{\tilde{\mu}, \tilde{\lambda}} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} (\pi(x_{ik}, \tilde{s}(w_k), z_k) - \pi(x_{ik}, s_{-i,k}, z_k)) \right. \right. \\
 & \quad \quad \left. \left. \mid x_{it} = x, s_{-i,t} = s_t, w_t \right] \right]. \quad (11)
 \end{aligned}$$

Note that \hat{x}_k in the second bound is controlled by strategy $\hat{\mu}$, that is, a strategy in which the firm never exits the industry and invests an infinite amount at every state. Recall that (s_k, w_k) is distributed according to the invariant distribution for all $k \geq 0$. The bound can be computed using simulation. As before, these bounds are quite general and do not rely on many of the detailed modeling assumptions. Finally, bounds that hold for fixed values of the aggregate shock (as opposed to averaging over them) can also be obtained.

8.4. Computational Experiments

In this section we present computational experiments for the model with aggregate shocks. First, we define the aggregate shocks model. Then, we characterize extended OE for different specifications of w_t . Finally, we study the behavior of the error bounds.

We consider the model presented in §7 with the parameters used for the case of high level of vertical differentiation and the “almost deterministic” entry process. We assume there is an aggregate shock to demand, z_t , that affects the market size in the following way:

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

We assume that z_t takes three values: $L = 0.8$, $M = 1$, and $H = 1.2$. Hence, L is an unfavorable state of the economy for all firms in which the total market size is reduced by 20%. On the contrary, H is a favorable state of the economy for all firms in which the total market size is increased by 20%.

We consider two different transition matrices for the Markov process that describes the evolution of z_t :

$$\begin{bmatrix} 0.6 & 0.4 & 0.0 \\ 0.2 & 0.6 & 0.2 \\ 0.0 & 0.4 & 0.6 \end{bmatrix} \quad \begin{bmatrix} 0.4 & 0.4 & 0.2 \\ 0.3 & 0.4 & 0.3 \\ 0.2 & 0.4 & 0.4 \end{bmatrix}$$


The first row and column of each matrix correspond to state L , the second to state M , and the third to state H . The matrix on the left represents a relatively more persistent aggregate shock in which the most likely value of the shock next period is its value in the current period. In contrast, the matrix on the right represents a relatively less persistent shock. We denote these cases as mp and lp , respectively.

We compute extended OE for the specification of w_t described in Example 8.1 for the cases $N = 1, 2, 3$. If $N = 1$, firms keep track of the current value of the shock only. If $N = 2$ and $N = 3$, firms keep track of the current and past realization of the shock, and of the current and two previous realizations, respectively. Similarly to OE, the algorithm presented in the electronic companion to compute extended OE is not guaranteed to terminate. However, for the examples below it terminated in less than one minute for the case $N = 1$ and less than 30 minutes for the case $N = 3$.

In Figure 7 we present the results obtained for the case $N = 1$ for both transition matrices mp and lp . For all values of w we show: (i) $\tilde{n}(w) = \sum_x \tilde{s}_{\mu, \lambda}(x, w)$, the conditional expected total number of firms when dynamics are governed by extended OE strategies and the current value of the shock is w ; and (ii) $\lambda(w)$, the extended OE entry rate. In both cases, the entry rate is large at state H and zero in the two other states. In addition, the average number of firms is larger at state H . Results are intuitive because H is the most profitable state for firms, and hence, attracts more of them. The extended OE outcome is consistent with the empirical fact that entry is procyclical with the business cycle (e.g., see Campbell 1998).

It is interesting to observe some properties of extended OE if $N > 1$. First, because z_t is a Markov process, it follows that under extended oblivious strategies, s_t is conditionally independent of z_t , conditioned on z_{t-1} . Therefore, $\tilde{s}_{\mu, \lambda}(w_t)$ does not vary with the values of z_t , it only varies with the values of $z_{t-1}, \dots, z_{t-N+1}$. Second, again because z_t is Markov, $w_{t+1} = \{z_{t+1}, z_t, \dots, z_{t-N+2}\}$ is conditionally independent of z_{t-N+1} , conditioned on z_t, \dots, z_{t-N+2} . Likewise, $\tilde{s}_{\mu, \lambda}(w_{t+1})$ is conditionally independent of z_{t-N+1} , conditioned on z_t, \dots, z_{t-N+2} . Recall that investment and entry rate decisions at time period t in OE are made in a

Figure 7. Expected number of firms and entry rate for extended OE.



way that optimizes discounted future profits based on the future trajectory $\{(x_\tau, w_\tau) \mid \tau > t\}$, assuming industry states $\tilde{s}_{\mu, \lambda}(w_\tau)$ at each time τ . As such, investment and entry decisions made at time t in an extended OE do not depend on z_{t-N+1} . For example, for $N = 2$ an extended OE strategy and entry rate are functions of z_t only.¹⁸

In Table 3 we compare the extended OE entry rate functions for the different values of N in the case mp . As expected, entry rates change when going from $N = 1$ to $N = 2$. In the latter, there is entry at state M , whereas in the former there is not. Also, note that when $N = 3$ there could be entry for all values of z_t . However, interestingly, the larger entry rates are obtained in states $\{z_t = H, z_{t-1} = M\}$ and $\{z_t = M, z_{t-1} = L\}$.¹⁹ Again consistent with the empirical evidence mentioned above, these are situations in which the industry has evolved from a less favorable state to a more favorable one attracting more entrants. Even though strategies change significantly as N increases, implied long-run averages change less. The long-run average entry rate decreases from 4.3 in the case $N = 1$ to 4.0 in the case $N = 3$. The long-run average number of firms decreases from 113.5 in the case $N = 1$ to 112.7 in the case $N = 3$.

Having analyzed extended OE, we now study the behavior of the error bound in the two cases mp and lp and for the different values of N . We consider error bound (11) evaluated at the entry state and in percentage terms relative to the expected value function.²⁰ Results can be found in Figure 8. We observe that both the persistence of the shock and the number of realizations of the shock firms keep track of have a first-order impact in the magnitude of the error bound. If the shock is less persistent, history becomes irrelevant more rapidly, and firms can make relatively better decisions by just keeping track of one or few values of the shock. This is reflected in the results; error bounds for the case lp are smaller than for the case mp . Also, error bounds decrease significantly as N increases and firms keep track of more shock values; the additional information allows firms to better predict the conditional industry state and make better decisions.

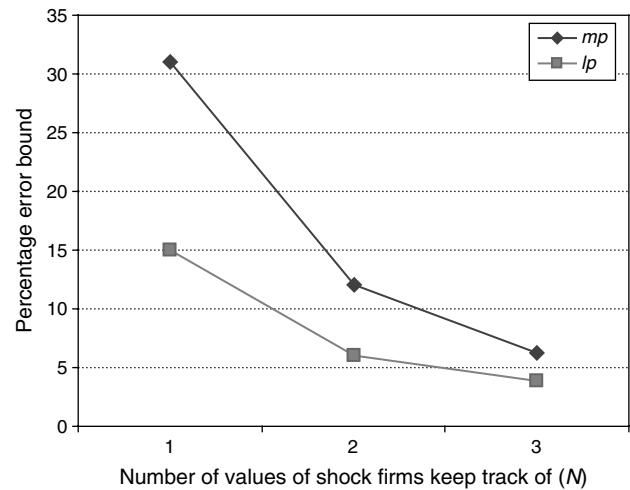
Note that if $N = 3$, error bounds are close to 5%. In the example, the industry accommodates 113 firms on average. Therefore, the magnitude of extended OE error bounds are not too different from those from OE with the same number of firms (see Figure 2). We conclude that by keeping

Table 3. Extended OE entry rate functions for case mp and different values of N .

N	Entry rate function
1	$\lambda(H) = 17.3$
2	$\lambda(M) = 3.5, \lambda(H) = 9.9$
3	$\lambda(L, L) = 0.6, \lambda(M, M) = 2.8, \lambda(H, H) = 4.6,$ $\lambda(M, L) = 11.7, \lambda(H, M) = 14.4$

Notes. For $N = 1, 2$, entry rate is a function of $z_t, \lambda(z_t)$. For $N = 3$, entry rate is a function of $(z_t, z_{t-1}), \lambda(z_t, z_{t-1})$. Only nonzero values are reported.

Figure 8. Percentage approximation error bound for extended OE.



track of few shock values the accuracy of the approximation based on extended OE in an industry with aggregate shocks can be comparable to that of OE in an industry without aggregate shocks.

The results show that extended OE further extends the set of applied problems that can be analyzed computationally. First, extended OE provide a computationally feasible approach to approximating dynamic behavior in industries with aggregate shocks. Second, extended OE provide sensible results that qualitatively coincide with patterns observed in industry data.

9. Conclusions and Future Research

The goal of this paper has been to increase the set of applied economic problems that can be addressed using dynamic models of imperfect competition in the style of Ericson and Pakes (1995). Due to the curse of dimensionality, existing dynamic programming methods have limited application of these models to industries with a small number of firms and a small number of states per firm. As an alternative, we proposed a method for approximating MPE behavior using an OE, where firms make decisions based only on their own state and the long-run average industry state.

We introduced a simple algorithm to compute an OE. A nice feature of the method is that there is no need to place a priori restrictions on the number of firms in the industry or the set of states that a firm can reach. As a result, computational considerations place very few constraints on model details. To facilitate using OE in practice, we derived approximation error bounds that indicate how good the approximation is in any particular problem instance. These approximation error bounds are quite general, and thus can be used in a wide class of models. Through computational experiments, we showed that OE often yields a good approximation of MPE behavior for

industries like those that empirical researchers would like to study. We also extended the notion of OE, our algorithms and error bounds for a model with an aggregate shock common to all firms. Through computational experiments, we showed that the extended notion of OE further expands the set of dynamic industries that can be analyzed. Note that even though the emphasis in this paper is on using OE to approximate industry dynamics, OE can also be used to approximate equilibria in more general stochastic games (see Weintraub et al. 2005b, Abhishek et al. 2007).

Together, these methods provide a toolkit that facilitates the application of dynamic oligopoly models to a wide range of empirical problems. Computation of OE is light enough that it is even feasible to embed OE computation into an estimation algorithm (see Xu 2006 for an example). The estimation algorithm searches over model parameters and computes OE for many candidate sets of parameters. It stops when it finds the set of parameters that most closely match a set of statistics generated by the OE of the model and those observed in the data. Examples of statistics that we might want to match in such an algorithm include the number of firms in the industry; sales, profits, prices, and/or market shares; entry and exit rates; rates of investment, the correlation between investment and market share, etc. The estimated model can then be used for many purposes including, for example, to evaluate the effects of a proposed policy change on industry structure, prices, and welfare. For example, one might want to use such a technique to evaluate the effects of a carbon tax or cap-and-trade system on various polluting industries (Ryan 2005). Our hope is that through the use of OE, empirical researchers will be able to utilize richer industry models than was previously possible, and to tackle important problems that were previously intractable.

Although we believe that the concept of OE will be useful in applications on its own, there are also some important extensions (see Weintraub et al. 2007). In order to capture short-run transitional dynamics that may result, for example, from shocks or policy changes, we have 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. Additionally, 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. This extension trades off increased computation time and memory for a better behavioral model and a better approximation to MPE behavior. Our hope is that the dominant firm OE will provide better approximations for more concentrated industries.

10. Electronic Companion

An electronic companion to this paper is available as part of the online version that can be found at <http://or.journal.informs.org/>.

Endnotes

1. See Berry and Pakes (1993), Gowrisankaran (1999), Fershtman and Pakes (2000), Judd et al. (2002), Langohr (2003), Song (2003), Besanko and Doraszelski (2004), de Roos (2004), Besanko et al. (2010), Fershtman and Pakes (2005), Goettler et al. (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.
2. Luttmer (2007) and Melitz (2003) extended Hopenhayn’s model to a setting with monopolistic competition in the product market.
3. Note that it would not change any of our results to assume that the entry state was a random variable.
4. Doraszelski and Satterthwaite (2010) also provide an example of multiple equilibria in their closely related model.
5. However, because oblivious strategies rule out strategies that are dependent on competitors’ states, there are likely to be fewer OE than there are MPE.
6. The parameters γ and N were set after some experimentation to speed up convergence to the values $2/3$ and 0 , respectively.
7. Note that there are potentially many alternative methods for searching over entry rates for an OE. For example, one alternative would be to start at an arbitrary entry rate and then implement small increments and decrements to the entry rate until an entry rate is found that leads to an OE.
8. The algorithm was programmed in C++ and the experiments were executed on a UNIX shared machine with a CPU Intel 2.66 GHz and 32 GB of RAM.
9. The AME property defined in Weintraub et al. (2008b) for their asymptotic result required that the approximation error converged to zero as market size grows for all states $x \in \mathbb{N}$.
10. In §8 we extend the model to incorporate aggregate shocks.
11. Although we are not able to show that $\Delta(y, s)^+$ is non-decreasing in y , we check it computationally for all sampled states s in the simulation.
12. The expected value function is estimated with a relative precision of 1% and a confidence level of 98%. The bound is estimated with a relative precision of at most 10% and a confidence level of 98% (in cases where the bound is very small it is difficult to achieve better precision than this). Note that the percentage approximation error bound depends on the state x , so for the purposes of this section we consider the percentage bound evaluated at the entry state. For the computations we took the maximum achievable state, \tilde{x}_{\max} , to be a state such that the expected number of visits of a firm using $\tilde{\mu}$ was at most 10^{-5} . In computing the bounds, we assumed that the maximum achievable state under the best-response (nonoblivious) strategy was also \tilde{x}_{\max} .
13. Note that the zero-profit condition typically requires a fractional number of entrants to be satisfied exactly, so to

accommodate this we instead randomized the number of entrants between the two neighboring integers. For example, if the equilibrium entry rate is 2.5, then the number of entrants is 2 or 3, with probability 0.5. Allowing for fractional numbers ensures existence of equilibrium. Note that with this entry process $\{s_t: t \geq 0\}$ also admits a unique invariant distribution.

14. Because now there is no entry state, we report the percentage error bound evaluated at the most visited state.

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

16. In principle, \mathcal{W}_j is an uncountable set that takes values between $\underline{a}_j = \min_{a \in A} g_j(a)$ and $\bar{a}_j = \max_{a \in A} g_j(a)$. However, for computational purposes we could assume that \mathcal{W}_j is a finite grid contained in $[\underline{a}_j, \bar{a}_j]$ and we could approximate the values of $w_t(j)$ with its closest element in the grid.

17. Recall that $w_t(1) = z_t$, hence, strategies are a function of z_t .

18. Note, however, that extended OE with $N = 2$ are different to extended OE with $N = 1$.

19. Note that state $\{z_t = 3, z_{t-1} = 1\}$ is not possible given the transition matrix mp .

20. Although we are not able to show that $\Delta(y, s, w)^+$ is nondecreasing in y , we check it computationally for all sampled states (s, w) in the simulation.

Acknowledgments

The authors have had very helpful conversations with Uli Doraszelski, Liran Einav, Feryal Erhun, Paul Glasserman, Hugo Hopenhayn, Ken Judd, Jon Levin, and Ariel Pakes, as well as seminar participants at the University of California–Berkeley, Columbia University, Duke University, the University of Iowa, INFORMS, the Kellogg School of Management, the University of Minnesota, New York University, SITE, Stanford University, the University of Rochester, the University of California at Los Angeles, the University of Illinois at Urbana–Champaign, the University of Chile, the University of Texas at Austin, and Yale University. They also thank the associate editor and two anonymous referees for valuable suggestions. Przemyslaw Jezierski provided exemplary research assistance. 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.

References

Abhishek, V., S. Adlakha, R. Johari, G. Y. Weintraub. 2007. Oblivious equilibrium for general stochastic games with many players. *Allerton Conf. Comm., Control Comput. University of Illinois at Urbana–Champaign*.

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

Berry, S., A. Pakes. 1993. Some applications and limitations of recent advances in empirical industrial organization: Merger analysis. *Amer. Econom. Rev.* **83**(2) 247–252.

Besanko, D., U. Doraszelski. 2004. Capacity dynamics and endogenous asymmetries in firm size. *RAND J. Econom.* **35**(1) 23–49.

Besanko, D., U. Doraszelski, Y. Kryukov, M. Satterthwaite. 2010. Learning-by-doing, organizational forgetting, and industry dynamics. *Econometrica* **78**(2) 453–508.

Campbell, J. R. 1998. Entry, exit, embodied technology, and business cycles. *Rev. Econom. Dynam.* **1** 371–408.

Caplin, A., B. Nalebuff. 1991. Aggregation and imperfect competition—On the existence of equilibrium. *Econometrica* **59**(1) 25–59.

Collard-Wexler, A. 2006. Productivity dispersion and plant selection in the ready-mix concrete industry. Working paper, New York University.

Davis, S. J., J. C. Haltiwanger, S. Schuh. 1998. *Job Creation and Destruction*, 1st ed. The MIT Press, Cambridge, MA.

Deaton, A., G. Laroque. 1996. Competitive storage and commodity price dynamics. *J. Political Econom.* **104**(5) 896–923.

de Farias, D. P., B. Van Roy. 2003. The linear programming approach to approximate dynamic programming. *Oper. Res.* **51**(6) 850–865.

de Roos, N. 2004. A model of collusion timing. *Internat. J. Indust. Organ.* **22** 351–387.

Dindos, M., C. Mezzetti. 2006. Better-reply dynamics and global convergence to Nash equilibrium in aggregative games. *Games Econom. Behav.* **54** 261–292.

Doraszelski, U., K. Judd. 2006. Avoiding the curse of dimensionality in dynamic stochastic games. Working paper, Hoover Institution, Stanford, CA.

Doraszelski, U., S. Markovich. 2007. Advertising dynamics and competitive advantage. *RAND J. Econom.* **38**(3) 557–592.

Doraszelski, U., A. Pakes. 2007. A framework for applied dynamic analysis in IO. *Handbook of Industrial Organization*, Vol. 3. North-Holland, Amsterdam.

Doraszelski, U., M. Satterthwaite. 2010. Computable Markov-perfect industry dynamics: Existence, purification, and multiplicity. *RAND J. Econom.* **41**(2) 215–243.

Dunne, T., M. J. Roberts, L. Samuelson. 1988. Patterns of firm entry and exit in US manufacturing industries. *RAND J. Econom.* **19**(4) 495–515.

Ericson, R., A. Pakes. 1995. Markov-perfect industry dynamics: A framework for empirical work. *Rev. Econom. Stud.* **62**(1) 53–82.

Farias, V., D. Saure, G. Y. Weintraub. 2008. An approximate dynamic programming approach to solving large scale dynamic oligopoly models. Working paper, Columbia University, New York.

Fershtman, C., A. Pakes. 2000. A dynamic oligopoly with collusion and price wars. *RAND J. Econom.* **31**(2) 207–236.

Fershtman, C., A. Pakes. 2005. Finite state dynamic games with asymmetric information: A framework for applied work. Working paper, Harvard University, Cambridge, MA.

Fudenberg, D., D. K. Levine. 1998. *The Theory of Learning in Games*. MIT Press, Cambridge, MA.

Goettler, R. L., C. A. Parlour, U. Rajan. 2005. Equilibrium in a dynamic limit order market. *J. Finance* **60**(5) 351–387.

Gowrisankaran, G. 1999. A dynamic model of endogenous horizontal mergers. *RAND J. Econom.* **30**(1) 56–83.

Gowrisankaran, G., R. Town. 1997. Dynamic equilibrium in the hospital industry. *J. Econom. Management Strategy* **6**(1) 45–74.

Hopenhayn, H. A. 1992. Entry, exit and firm dynamics in long run equilibrium. *Econometrica* **60**(5) 1127–1150.

Jenkins, M., P. Liu, R. L. Matzkin, D. L. McFadden. 2004. The browser war—Econometric analysis of Markov perfect equilibrium in markets with network effects. Working paper, University of California, Berkeley.

Judd, K. 1998. *Numerical Methods in Economics*. MIT Press, Cambridge, MA.

- Judd, K. L., K. Schmedders, S. Yeltekin. 2002. Optimal rules for patent races. Working paper, Hoover Institution, Stanford, CA.
- Langohr, P. 2003. Competitive convergence and divergence: Capability and position dynamics. Working paper, Northwestern University, Evanston, IL.
- Luttmer, E. G. J. 2007. Selection, growth, and the size distribution of firms. *Quart. J. Econom.* **122**(3) 1103–1144.
- Markovich, S. 2008. Snowball: A dynamic oligopoly model with indirect network effects. *J. Econom. Dynam. Control* **32** 909–938.
- Maskin, E., J. Tirole. 1988. A theory of dynamic oligopoly, I and II. *Econometrica* **56**(3) 549–570.
- Melitz, M. J. 2003. The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica* **71**(6) 1695–1725.
- Noel, M. D. 2008. Edgeworth cycles and focal prices: Computational dynamic Markov equilibria. *J. Econom. Management Strategy* **17**(2) 345–377.
- Novshek, W. 1985. Perfectly competitive markets as the limits of Cournot markets. *J. Econom. Theory* **35** 72–82.
- Pakes, A., P. McGuire. 1994. Computing Markov-perfect Nash equilibria: Numerical implications of a dynamic differentiated product model. *RAND J. Econom.* **25**(4) 555–589.
- Pakes, A., P. McGuire. 2001. Stochastic algorithms, symmetric Markov perfect equilibrium, and the “curse” of dimensionality. *Econometrica* **69**(5) 1261–1281.
- Qi, S. 2008. The impact of advertising regulation on industry: The cigarette advertising ban of 1971. Working paper, University of Minnesota, Minneapolis.
- Routledge, B. R., D. J. Seppi, C. S. Spatt. 2000. Equilibrium forward curves for commodities. *J. Finance* **55**(3) 1297–1338.
- Ryan, S. 2005. The costs of environmental regulation in a concentrated industry. Mimeo, MIT, Cambridge, MA.
- Schivardi, F., M. Schneider. 2008. Strategic experimentation and disruptive technological change. *Rev. Econom. Dynam.* **11**(2) 386–412.
- Song, M. 2003. A dynamic model of cooperative research in the semiconductor industry. Working paper, Harvard University, Cambridge, MA.
- Weintraub, G. Y., C. L. Benkard, B. Van Roy. 2005a. Markov perfect industry dynamics with many firms. Working paper, Stanford University, Stanford, CA.
- Weintraub, G. Y., C. L. Benkard, B. Van Roy. 2005b. Oblivious equilibrium: A mean field approximation for large-scale dynamic games. *Advances in Neural Information Processing Systems*. MIT Press, Cambridge, MA.
- Weintraub, G. Y., C. L. Benkard, B. Van Roy. 2007. Extensions to oblivious equilibrium. Working paper, Stanford University, Stanford, CA.
- Weintraub, G. Y., C. L. Benkard, B. Van Roy. 2008a. Industry dynamics: Foundations for models with an infinite number of firms. Working paper, Stanford University, Stanford, CA.
- Weintraub, G. Y., C. L. Benkard, B. Van Roy. 2008b. Markov perfect industry dynamics with many firms. *Econometrica* **76**(6) 1375–1411.
- Xu, Y. 2006. A structural empirical model of R&D, firm heterogeneity, and industry evolution. Working paper, Penn State University, University Park, PA.