

Industry Dynamics: Foundations For Models with an Infinite Number of Firms*

Gabriel Y. Weintraub
Columbia University

C. Lanier Benkard
Stanford University
and NBER

Benjamin Van Roy
Stanford University

August, 2010

Abstract

This paper explores the connection between three important threads of economic research offering different approaches to studying the dynamics of an industry with heterogeneous firms. *Finite models* of the form pioneered by Ericson and Pakes (1995) capture the dynamics of a finite number of heterogeneous firms as they compete in an industry, and are typically analyzed using the concept of Markov perfect equilibrium (MPE). *Infinite models* of the form pioneered by Hopenhayn (1992), on the other hand, consider an infinite number of infinitesimal firms, and are typically analyzed using the concept of stationary equilibrium (SE). A third approach uses oblivious equilibrium (OE), which maintains the simplifying benefits of an infinite model but within the more realistic setting of a finite model. The paper relates these three approaches. The main result of the paper provides conditions under which SE of infinite models approximate MPE of finite models arbitrarily well in asymptotically large markets. Our conditions require that the distribution of firm states in SE obeys a certain “light-tail” condition. In a second set of results, we show that the set of OE of a finite model approaches the set of SE of the infinite model in large markets under a similar light-tail condition.

JEL classification numbers: E22, L11, L13

*We have had very helpful conversations with Daron Acemoglu, Sachin Adlakha, José Blanchet, Eduardo Engel, Juan Escobar, Hugo Hopenhayn, Ramesh Johari, Chad Jones, Boyan Jovanovic, Andy Skrzypacz, and Assaf Zeevi. We thank the Associate Editor and two anonymous referees for suggestions that greatly improved the paper. Correspondence: gweintraub@columbia.edu; lanier.benkard@yale.edu; bvr@stanford.edu.

1 Introduction

This paper explores the connection between three threads of economic research, each of which offers a different approach to studying the dynamics of an industry with heterogeneous firms. One is based on Ericson and Pakes (1995), which captures interactions among a finite number of individual firms as they enter, compete in, and eventually exit an industry. We refer to such models as *finite*, as they explicitly account for the evolution of a finite set of firms. While this is an attractive feature, the analysis of finite models typically involves computation of Markov perfect equilibria (MPE) using dynamic programming algorithms (see Doraszelski and Pakes (2007) for an excellent survey). Such analyses restrict the number of firms to a small number because the computational requirements become unmanageable when there are more than that.

A second approach, pioneered by Hopenhayn (1992), assumes that there are an infinite number of firms, each of which garners an infinitesimal market share. Because of averaging effects across firms, these *infinite models* have stationary equilibria (SE) in which the aggregate industry state is constant over time. Such equilibria can typically be analyzed and computed efficiently; the simplification (relative to finite models) arising from the fact that individual firms need not keep track of the industry state since it is constant.

A third approach, due to Weintraub, Benkard, and Van Roy (2008), analyzes oblivious equilibria (OE) in finite models. In OE, firms optimize assuming that the industry state is constant over time and equal to its long-run expected value. As a result, OE shares the computational advantages of SE in infinite models. Additionally, as with MPE, the finite model setting of OE is useful because the model can more easily be related to industry data such as the number of firms and the market shares of leading firms. On the other hand, because the industry state is not truly constant in a finite model, OE represents only an approximation to optimal fully informed MPE behavior in the model. OE has been developed only recently, but has already seen wide use in applications.¹

This paper explores the connection between the three approaches. One of the main goals of the paper is to provide theoretical foundations for infinite models, and justification for their use. An important literature on infinite models has developed in macroeconomics, international trade, and industrial organization, studying diverse dynamic phenomena such as the size distribution of firms (Luttmer 2007), the intra-industry effects of international trade (Melitz 2003), R&D investments (Klette and Kortum 2004), firms' technological learning (Mitchell 2000), and job turnover (Hopenhayn and Rogerson 1993) to name a few. While

¹For example, Xu (2008), Mukherjee and Kadiyali (2008), Qi (2008), Iacovone, Javorcik, Keller, and Tybout (2009), Thürk (2009a) and Thürk (2009b) have used OE (or closely related equilibrium concepts) as a basis for their estimation methods and/or counterfactual analysis in different applications.

infinite models have become increasingly popular, they are an idealization of real-world industries because they assume an infinite number of firms. To our knowledge, there is as yet no rigorous justification in the literature for their use.

We address this issue by providing conditions under which infinite models give useful approximations of dynamic behavior in large but finite industries. More precisely, our main result (in Section 5) provides conditions under which stationary equilibria in infinite models approximate finite model MPE for large markets. Our conditions require that (1) the number of firms grows with the market size; and (2) that the distribution of firms across firm states exhibits a “light-tail” condition that we define precisely. Under these conditions one can justify use of an infinite model as a tractable proxy of a finite model. An infinite model SE can be a poor approximation of a finite model MPE if either of these conditions is violated.

Previous work (Weintraub, Benkard, and Van Roy 2008) has already shown that OE approximate MPE in large markets (under a similar light-tail condition as the one mentioned above). Thus, it remains to describe the link between OE in finite models and SE in infinite models. This link turns out to be very close. First, SE can be understood as the OE of an infinite model. Moreover, in Section 6 we show that, under the appropriate light-tail condition and as the market size grows, the set of finite model OE approaches the set of infinite model SE in the following sense: (1) for large markets, every OE is close to an SE of the infinite model; and (2) all sequences of strategies that approach an SE of the infinite model satisfy the OE conditions asymptotically. Point (1) corresponds to the upper-hemicontinuity and point (2) is related to the lower-hemicontinuity of the OE correspondence, respectively, at the point where the number of firms and market size become infinite. These results imply not only that OE and SE approximate MPE asymptotically (under the light-tailed condition), but also that they essentially approximate the *same subset of MPE*.

The light-tail condition is important for both of our main results. In Section 5.5 we provide an example of an infinite model SE that is not light-tailed and that does not approximate MPE asymptotically. The example illustrates that problems can arise when the market being studied is dominated by a small fraction of relatively large firms. In that case it is not appropriate to treat the industry state as constant asymptotically since dominant firms can exert significant influence on industry dynamics and therefore all firms should factor dominant firm states into their decisions. Our light-tail condition rules out such situations by ensuring that asymptotically large markets do not exhibit large concentration.

On the positive side, the light-tail condition is satisfied by many commonly used economic models. For example, it is satisfied by a model with a logit demand system and firms competing Nash in prices in the spot market if the average firm size is finite. Moreover, for many infinite models (e.g., those that assume monopolistic competition (Dixit and Stiglitz 1977)), the mere existence of an equilibrium with positive entry

rates immediately implies the light-tail condition. In these cases, the use of an infinite model as a tractable proxy of a finite model is well justified.

Note that unlike the rest of our assumptions the light-tail condition is a condition on the equilibrium distribution of firms, not a condition on the model primitives. We use it because it captures the core convergence issue and therefore provides very general results. Moreover, the light-tail condition is quite useful in practice because it is easy to check once an equilibrium to the model has been found.

A broad implication of our main result is that, for asymptotically large markets, a simple strategy that ignores current market information can be close to optimal. In this sense, our results contribute to the vast and classic literature on the convergence to competitive equilibria (Roberts and Postlewaite (1976), Novshek and Sonnenschein (1978), Mas-Colell (1982), Mas-Colell (1983), Novshek and Sonnenschein (1983), Novshek (1985), Allen and Hellwig (1986a), Allen and Hellwig (1986b), and Jones (1987)). Roughly speaking, these papers establish conditions in different static models under which the set of oligopolistic Nash equilibria approaches, in some sense, the set of (Walrasian) competitive equilibria as the size of individual firms (or agents) becomes small relative to the size of the market. There are some notable differences with our work, though. Our interests lie in approximating dynamic firm behavior in large markets, not in showing that the product market is perfectly competitive in the limit. In particular, while the above papers study *static* models, in which the main strategic decisions are usually prices or quantities, we study *dynamic* models, in which the main decisions are, for example, investment, entry, and exit. Thus, while we show that firm investment, entry, and exit strategies become simple in markets with many firms, we do not rule out that a small fraction or even all firms may still have some degree of market power in the product market even in the limit. Indeed, in the examples provided in the paper the limit product market is given by monopolistic competition. Additionally, differently from the papers above, in our model the size of firms is endogenously determined in equilibrium through investment. For that reason, we impose a light-tail condition that controls for the appearance of dominant firms to obtain our asymptotic result.

Our analysis demonstrates that if the light-tail condition is not satisfied, a model that averages out idiosyncratic shocks can provide a poor approximation to industry dynamics. In this sense, our work is related to Jovanovic (1987) and Gabaix (2008) that provide conditions for which idiosyncratic fluctuations can generate aggregate shocks even in an economy with a large number of firms. In particular, Gabaix (2008) argues that if the distribution of firm sizes is heavy-tailed, idiosyncratic shocks to large firms can lead to non-trivial aggregate shocks. He empirically shows that the movements of the 100 largest firms in the US appear to explain one third of variations in output.

Finally, there is a closely related paper by Adlakha, Johari, Weintraub, and Goldsmith (2010) that proves

that SE approximates MPE asymptotically in a model with multidimensional firm states but in a different setting. Their model does not include entry and exit and, instead of considering asymptotics in the market size, they increase the number of agents directly.

2 A Model of an Industry with a Finite Number of Firms

The model in this section is close to Weintraub, Benkard, and Van Roy (2008) which in turn is similar in spirit to Ericson and Pakes (1995). An important difference, however, is that we consider multidimensional firm states. Adlakha, Johari, Weintraub, and Goldsmith (2010) also consider this and other generalizations in a model without entry and exit. In addition, we note that an important difference with Ericson and Pakes (1995) is that our model includes only idiosyncratic shocks. This simplification is important to relate our finite model to an infinite model.

2.1 Model and Notation

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

Each firm that enters the industry is assigned a unique positive integer-valued index denoted by i . Firm heterogeneity is reflected through firm states. A firm's state is a multidimensional vector of characteristics that affect the firm's profits, such as the quality levels of its products, productivity, capacity, or the size of its consumer network. The state of firm i at time t is denoted by $x_{it} \in \mathcal{X}$, where $\mathcal{X} \subseteq \mathbb{N}^q$, $q \geq 1$.

The *industry state* s_t represents a complete list of all incumbent firms and the value of their state variables at time t . Formally, we define s_t to be a vector over individual states that specifies, for each individual state $x \in \mathcal{X}$, the number of incumbent firms at state x in period t . At each time $t \in \mathbb{N}$, we denote the number of incumbent firms as n_t .

We define the state space $\bar{\mathcal{S}} = \left\{ s : \mathcal{X} \rightarrow \mathbb{N} \mid \sum_{x \in \mathcal{X}} s(x) < \infty \right\}$. Though in principle there are a countable number of industry states, we will also consider an extended state space $\mathcal{S} = \left\{ s : \mathcal{X} \rightarrow \mathbb{R}_+ \mid \sum_{x \in \mathcal{X}} s(x) < \infty \right\}$ that allows the number of firms at any state to be non-integer valued. This will be useful to define the infinite model and will allow us, for example, to consider derivatives of functions with respect to the industry state. Note that an element of \mathcal{S} defines a measure over \mathcal{X} and an element of $\bar{\mathcal{S}}$ defines a counting measure.

For each incumbent firm i , we define $s_{-i,t} \in \mathcal{S}$ to be the state of the *competitors* of firm i ; that is, $s_{-i,t}(x) = s_t(x) - 1$ if $x_{it} = x$, and $s_{-i,t}(x) = s_t(x)$, otherwise. Similarly, $n_{-i,t}$ denotes the number of competitors of firm i .

In each period, each incumbent firm earns profits $\pi_m(x_{it}, s_{-i,t})$ that depend on its state x_{it} and its competitors' state $s_{-i,t}$. Note that in most applied problems the profit function would not be specified directly, but would instead result from a deeper set of primitives that specify a demand function, a cost function, and a static equilibrium concept. We make explicit the dependence of profits on an important parameter of the demand system that we focus on later; the size of the relevant market, denoted by m . Profits would typically increase with market size for a firm at a given state (x, s) .

The model also allows for entry and exit. In each period, each incumbent firm i 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 take an action to improve its individual state. Let $\mathcal{I} \subseteq \mathfrak{R}_+^p$ ($p \geq 1$) be the multidimensional action space; for concreteness, we refer to this action as an investment. If a firm invests $l_{it} \in \mathcal{I}$, then the firm's state at time $t + 1$ is given by,

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

where the function w captures the impact of investment and $\zeta_{i,t+1}$ is a finite dimensional random vector that 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. The function w takes values in \mathbb{Z}^q . This specification is very general as w may take on either positive or negative values in any of its components (e.g., allowing for positive depreciation). The investment cost is given by the function $d : \mathcal{I} \rightarrow \mathfrak{R}_+$.

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 $x^e \in \mathcal{X}$ and can earn profits thereafter. The entry model is described more precisely below.

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

In each period, events occur in the following order:

1. Each incumbent firms observes its sell-off value and then makes exit and investment decisions.

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

2.2 Assumptions

We begin with some assumptions about the model primitives. Note that, relative to Weintraub, Benkard, and Van Roy (2008), we weaken some assumptions in order to accommodate multidimensional firm states.

We first make some natural assumptions about the profit function.

Assumption 2.1. *For each $m \in \mathbb{N}$, π_m satisfies the following:*

1. For all $x \in \mathcal{X}$ and $s \in \mathcal{S}$, $\pi_m(x, s) > 0$, and $\sup_{x \in \mathcal{X}, s \in \mathcal{S}} \pi_m(x, s) < \infty$.
2. For all $x \in \mathcal{X}$, the function $\ln \pi_m(x, \cdot) : \mathcal{S} \rightarrow \mathbb{R}_+$ is continuously Fréchet differentiable. Hence, for all $x \in \mathcal{X}$, $y \in \mathcal{X}$, and $s \in \mathcal{S}$, $\pi_m(x, s)$ is continuously differentiable with respect to $s(y)$. Further, for any $x \in \mathcal{X}$, $s \in \mathcal{S}$, and $h \in \mathcal{S}$ such that $s + \gamma h \in \mathcal{S}$ for $\gamma > 0$ sufficiently small, if

$$\sum_{y \in \mathcal{X}} h(y) \left| \frac{\partial \ln \pi_m(x, s)}{\partial s(y)} \right| < \infty,$$

then

$$\frac{d \ln \pi_m(x, s + \gamma h)}{d\gamma} \Big|_{\gamma=0} = \sum_{y \in \mathcal{X}} h(y) \frac{\partial \ln \pi_m(x, s)}{\partial s(y)}.$$

We assume that profits are positive and bounded (2.1.1). The last part of assumption 2.1 is technical and requires that log-profits are Fréchet differentiable. Note that this requires partial differentiability of the profit function with respect to each $s(y)$. Profit functions that are “smooth”, such as ones arising from random utility demand models like the logit model, will satisfy this.

Next we make assumptions about the investment function and the distributions of the private shocks.

Assumption 2.2.

1. The random variables $\{\phi_{it} | t \geq 0, i \geq 1\}$ are i.i.d. and have finite expectations and well-defined density functions with support \mathbb{R}_+ .
2. The random vectors $\{\zeta_{it} | t \geq 0, i \geq 1\}$ are i.i.d. and independent of $\{\phi_{it} | t \geq 0, i \geq 1\}$.
3. There exists a positive constant $\bar{w} \in \mathbb{N}$ such that $\|w(x, \iota, \zeta)\|_\infty \leq \bar{w}$, for all (x, ι, ζ) . The set \mathcal{I} is compact and convex.
4. For all $x \in \mathcal{X}$ and $y \in \{y \in \mathbb{Z}^q : \|y\|_\infty \leq 1\}$, there exists $\iota \in \mathcal{I}$, such that $\mathcal{P}[w(x, \iota, \zeta_{i,t+1}) = y] > 0$.

5. For all $x \in \mathcal{X}$ and $y \in \mathbb{Z}^q$, $\mathcal{P}[w(x, \iota, \zeta_{i,t+1}) = y]$ is continuous in ι . The function d is continuous and $d(0) = 0$.
6. For all competitors' decisions and all terminal values, a firm's one time-step ahead optimization problem to determine its optimal investment has a unique solution.

As stated already above, we need to assume that investment and exit outcomes are idiosyncratic conditional on the state (2.2.1 and 2.2.2), ruling out aggregate shocks. We require this assumption in order to compare the model to the infinite model.

We also place a finite (but possibly large) bound on how much a firm's state can change in one period (2.2.3), an assumption that seems weak. Assumption 2.2.4 ensures that the investment process is rich enough to allow firms to move across the state space. Assumption 2.2.5 ensures that the impact of investment on transition probabilities and the investment cost function are continuous. Finally, assumption 2.2.6 is a generalization of the unique investment choice admissibility assumption in Doraszelski and Satterthwaite (2010) that is used to guarantee existence of an equilibrium to the model in pure strategies. It is satisfied by many of the commonly used specifications in the literature.

We assume that there are a large number of short-lived potential entrants each period who play a symmetric mixed entry strategy. In that case one can show that the number of actual entrants is well approximated by the Poisson distribution (see Weintraub, Benkard, and Van Roy (2008) for a derivation of this result). Modeling the number of entrants as a Poisson random variable provides a realistic model of entry in markets with varying sizes, and also has the advantage that it leads to more elegant asymptotic results. We also require that the entry and exit cost parameters be such that firms would not want to enter simply to collect the exit value.

Assumption 2.3.

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

Assumptions 2.1, 2.2, and 2.3 are kept throughout the paper.

2.3 Equilibrium

We consider symmetric pure strategy Markov perfect equilibria (MPE). We denote firms' investment strategies as $\iota(x_{it}, s_{-i,t})$. Weintraub, Benkard, and Van Roy (2008) show that in this model there always exists an optimal exit strategy that takes the form of a cutoff rule $\rho(x_{it}, s_{-i,t})$ such that an incumbent firm i exits at time t if and only if $\phi_{it} \geq \rho(x_{it}, s_{-i,t})$.

To simplify some of the mathematical expressions, it helps to group the exit and investment strategies together notationally. Let \mathcal{M} denote the set of investment/exit strategies such that an element $\mu \in \mathcal{M}$ is a pair of functions $\mu = (\iota, \rho)$, where $\iota : \mathcal{X} \times \mathcal{S} \rightarrow \mathcal{I}$ is an investment strategy and $\rho : \mathcal{X} \times \mathcal{S} \rightarrow \mathfrak{R}_+$ is an exit strategy.

Similarly, we denote the expected number of firms entering at industry state s_t , by $\lambda(s_t)$. This state-dependent entry rate will be endogenously determined, and our solution concept will require that it satisfies a zero expected profit condition. We denote the set of entry rate functions by Λ , where an element of Λ is a function $\lambda : \mathcal{S} \rightarrow \mathfrak{R}_+$.

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

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

where i is taken to be the index of a firm at state x at time t , τ_i is a random variable representing the time at which firm i exits the industry, and the subscripts of the expectation indicate the strategy followed by firm i , the strategy followed by its competitors, and the entry rate function. In an abuse of notation, we will use the shorthand, $V(x, s|\mu, \lambda) \equiv V(x, s|\mu, \mu, \lambda)$, to refer to the expected discounted value of profits when firm i follows the same strategy μ as its competitors. Note that in the notation we are suppressing the dependence of V on m .

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

1. Incumbent firm strategies represent an MPE:

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

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

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

Weintraub, Benkard, and Van Roy (2008) 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' . They also discuss that existence of MPE can be established using similar arguments to previous work (like Doraszelski and Satterthwaite (2010) and Escobar (2008)).

Dynamic programming algorithms can be used to optimize firm strategies, and equilibria to our model can be computed via their iterative application. However, these algorithms require compute time and memory that grow proportionately with the number of relevant industry states, which is often intractable in contexts of practical interest. This computational complexity has led researchers to explore models that are simpler to analyze. One way to simplify the model is to assume that there are an infinite number of firms.

3 A Model of an Industry with an Infinite Number of Firms

In this section we formulate an infinite model that represents an asymptotic regime within the finite model where the number of firms and the market size become infinite. The model we present is motivated by and very close in spirit to that proposed by Hopenhayn (1992). However, there are some differences that result from making the infinite model comparable to the finite model presented above. We mention the most significant ones. First, in Hopenhayn's model firms' states are real-valued, but in our model they are multidimensional vectors of positive integers. Note, however, that the key characteristic of a continuum of firms, that there are an infinite number of infinitesimal firms, is maintained. Second, Hopenhayn restricts a-priori the set of feasible individual states to lie on a compact set (the interval $[0, 1]$), whereas we do not. Finally, in our model, firm investment decisions are endogenously determined, whereas Hopenhayn's model focusses on entry and exit and assumes firms' trajectories follow exogenous Markov processes.

Because of averaging effects across firms, in an infinite model the industry state evolves deterministically. Further, following Hopenhayn (1992), we propose an equilibrium concept in which the state of the industry is constant over time, corresponding to the steady-state behavior of the industry. Together these two things imply that each firm's strategy in equilibrium depends only on its own individual state, alleviating the curse of dimensionality. We now define the infinite model.

3.1 Model and Notation

The infinite model represents an asymptotic regime within the finite model above where the number of firms and the market size become infinite. In order to study this regime, we consider a sequence of markets indexed by market sizes $m \in \mathbb{N}$. All other model primitives except the market size are assumed to remain

constant within this sequence.

As in the finite model, firms' states take on values in \mathcal{X} , and the state s_t of the industry at each time t is an element of \mathcal{S} . Note that unlike the finite model, in the infinite model, there is no need for a distinction between the industry state and the state of the competitors of a specific firm. This is because there are an infinite number of firms so if one firm is removed from the industry, there is no change to the industry state.

Profits in the infinite model are defined by the single-period profit function $\pi_\infty : \mathcal{X} \times \mathcal{S} \rightarrow \mathfrak{R}_+$. Relative to the profit function π_m of the finite model, the infinite model's profit function is a limit: $\pi_\infty(x, s) = \lim_{m \rightarrow \infty} \pi_m(x, ms)$ (see Assumption 3.1 below for a rigorous definition). Note that in order to handle an infinite number of firms, we overload the notation so that s_t as an argument to π_∞ is interpreted not as a counting measure, but as the number of firms at individual state x normalized by the market size m . In the infinite model, if $s_t(x)$ is nonzero for some x , then there are an infinite number of firms at individual state x . For large markets, the product $ms_t(x)$ is an approximation to the number of firms that would be at individual state x . Similarly, in our infinite model, an entry rate λ is interpreted as the ratio between the number of firms entering and the market size. The idea is that for a large market size m , the product $m\lambda$ is an approximation to the number of entering firms. These are the only two variables that we need to normalize for the infinite model. This overloaded notation is useful because it allows us to use the same definition of equilibrium for both models (see Section 4).

Investment and exit decisions are generated by strategies ι and ρ and entry is controlled by a (normalized) entry rate function λ . The exit and investment processes are the same as in the finite model.

Because of averaging effects across firms, in the infinite model the industry state evolves deterministically. It is possible to define industry state transitions and a value function over the state space $\mathcal{X} \times \mathcal{S}$ that take this into account. However, to be consistent with past literature and because we are interested in the long-run behavior of the industry, we instead define an equilibrium concept that assumes the industry state reaches a steady state. First, we introduce assumptions over the sequence of profit functions.

3.2 Assumptions

It will be helpful to decompose s according to $s = n \cdot f$, where f is a vector representing the fraction of firms in each state and n the total number of firms. Let $\mathcal{S}_1 = \{f \in \mathcal{S} \mid \sum_{x \in \mathcal{X}} f(x) = 1\}$. With some abuse of notation, we define $\pi_m(x_{it}, f_{-i,t}, n_{-i,t}) \equiv \pi_m(x_{it}, n_{-i,t} \cdot f_{-i,t})$, where $f_{-i,t}$ is a vector representing the fraction of competitors of firm i in each state at time period t .

We require some assumptions about the sequence of profit functions (indexed by m) in addition to Assumption 2.1, which applies to individual profit functions. The assumptions are discussed below.

Assumption 3.1.

1. $\sup_{x \in \mathcal{X}, s \in \mathcal{S}} \pi_m(x, s) = O(m)$.²
2. *There exists a non-negative real-valued function π_∞ such that, for all $x \in \mathcal{X}$, $f \in \mathcal{S}_1$, $c > 0$, and all sequences $\{n(m) | m \in \mathbb{N}\}$,*

$$\lim_{m \rightarrow \infty} \pi_m(x, f, n(m)) = \begin{cases} \pi_\infty(x, cf) & \text{if } \lim_m n(m)/m = c \in (0, \infty), \\ \infty & \text{if } \lim_m n(m)/m = 0, \\ 0 & \text{if } \lim_m n(m)/m = \infty. \end{cases}$$

3. *For all sequences $n(m)$ satisfying $\liminf_{m \rightarrow \infty} n(m)/m > 0$, there exists constants $d, e, k > 0$, such that $\pi_m(x, f, n(m)) \leq d \|x\|_\infty^k + e$, for all $x \in \mathcal{X}$, $f \in \mathcal{S}_1$, and $m \in \mathbb{N}$.*

4.

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

We first assume (3.1.1) that profits increase at most linearly with market size. This assumption should hold for virtually any relevant class of profit functions. It is satisfied, for example, if the total disposable income of the consumer population grows linearly in market size.

Next we assume (3.1.2) that, if for a given normalized industry state, the number of firms grows proportionally with the market size, then profits converge to a non-negative number given by the limit profit function π_∞ as the market size grows to infinity. Recall that π_∞ is the single-period profit function of the infinite model. Hence, this assumption essentially defines the asymptotic regime for the infinite model. If the number of firms increases slower than the market size, profits grow to infinity; if the number of firms increases faster than the market size, profits converge to zero.

Assumption 3.1.3 imposes a condition over the growth rate of single-period profits. It is used to simplify several arguments involving dynamics. For example, it implies that expected discounted profits remain uniformly bounded over all market sizes when the number of firms and the market size grow to infinity at the same rate and that expected discounted profits are finite in the limit model. Assumption 3.1.4 requires that profits are “smooth” with respect to the number of firms and, in particular, that the respective elasticity is uniformly bounded.

We note that the limit profit function π_∞ does not necessarily represent a static competitive equilibrium. For example, it could represent a monopolistically competitive market. In fact, an example of a standard economic model that satisfies these assumptions is a logit demand model where the spot market equilibrium

²In this notation, $n(m) = O(h(m))$ denotes $\limsup_m \frac{n(m)}{h(m)} < \infty$.

is Nash in prices. In this case, the limit profit function π_∞ corresponds to a logit model of monopolistic competition (Besanko, Perry, and Spady 1990).³ In Section 5.6 we discuss this example in more detail.

Equilibrium in the infinite model is defined in the following section.

4 Oblivious Equilibrium

In this section we define a notion of *oblivious equilibrium* that applies to both the finite model and the infinite model. In the infinite model, our definition of equilibrium will correspond with the standard definition of *stationary equilibrium* (SE) used in the literature. In a finite model, our equilibrium definition corresponds to the definition of *oblivious equilibrium* (OE) in Weintraub, Benkard, and Van Roy (2008).

4.1 Notation

In the infinite model, if the industry is in steady state then a firm's optimal strategy will depend only on its own individual state (and a single industry state). Similarly, in the finite model we can think about restricting firm strategies to be functions only of the the firm's individual state, even if the optimal Markov strategy may not necessarily satisfy this restriction. Following Weintraub, Benkard, and Van Roy (2008) we call such restricted strategies *oblivious*. Let $\tilde{\mathcal{M}} \subset \mathcal{M}$ and $\tilde{\Lambda} \subset \Lambda$ denote the set of oblivious strategies and the set of oblivious entry rate functions. Since each strategy $\mu = (\iota, \rho) \in \tilde{\mathcal{M}}$ generates decisions $\iota(x, s)$ and $\rho(x, s)$ that do not depend on s , with some abuse of notation, we will often drop the second argument and write $\iota(x)$ and $\rho(x)$. Similarly, for an entry rate function $\lambda \in \tilde{\Lambda}$, we will denote by λ the real-valued entry rate that persists for all industry states.

Suppose all firms use a common strategy $\mu \in \tilde{\mathcal{M}}$ and the entry rate is $\lambda \in \tilde{\Lambda}$. In the infinite model, though each firm evolves stochastically, the percentage of firms that transition from any given individual state to another is deterministic. Similarly, the percentage of firms that exit is deterministic. Let $\mathcal{P}_\mu[x_{i,t+1} = x | x_{i,t} = y]$ be the probability that a firm in state y evolves to state x next period when using strategy μ . The constant industry state $\tilde{s}_{\mu,\lambda}$ that represents the steady state satisfies the following balance equations:

$$(4.1) \quad \tilde{s}_{\mu,\lambda}(x) = \begin{cases} \sum_{y \in \mathcal{X}} \mathcal{P}_\mu[x_{i,t+1} = x | x_{i,t} = y] \tilde{s}_{\mu,\lambda}(y) + \lambda & \text{if } x = x^e, \\ \sum_{y \in \mathcal{X}} \mathcal{P}_\mu[x_{i,t+1} = x | x_{i,t} = y] \tilde{s}_{\mu,\lambda}(y) & \text{otherwise.} \end{cases}$$

In the finite model, this equation can be interpreted as defining the long-run expected industry state under

³More precisely, the standard logit model requires a modest and natural extension to allow for fractional states. Also, this single-period profit function satisfies the assumptions in the subset of industry states where it is strictly positive.

strategies (μ, λ) .

For an oblivious strategy $\mu \in \tilde{\mathcal{M}}$, an oblivious entry rate function $\lambda \in \tilde{\Lambda}$, and market size $m \in \mathbb{N} \cup \{\infty\}$ we define an *oblivious value function*

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

This value function should be interpreted as the expected net present value of a firm that is at individual state x and follows oblivious strategy μ' , under the assumption that its competitors' state will be $\tilde{s}_{\mu, \lambda}$ for all time. In the infinite model this assumption is correct. In the finite model, it is approximately correct when there are a large number of firms. Again, we abuse notation by using $\tilde{V}^{(m)}(x|\mu, \lambda) \equiv \tilde{V}^{(m)}(x|\mu, \mu, \lambda)$ to refer to the oblivious value function when firm i follows the same strategy μ as its competitors.

4.2 Definition of Equilibrium

We define an oblivious equilibrium (OE) for market size $m \in \mathbb{N} \cup \{\infty\}$. A *stationary equilibrium (SE)* of our infinite model corresponds to an OE for $m = \infty$. An *oblivious equilibrium* for market size $m \in \mathbb{N} \cup \{\infty\}$ comprises an investment/exit strategy $\mu = (\iota, \rho) \in \tilde{\mathcal{M}}$ and an entry rate function $\lambda \in \tilde{\Lambda}$ that satisfy the following conditions:

1. The industry state is constant and given by equations (4.1). The expected time each firm spends inside the industry is finite: $\sum_{x \in \mathcal{X}} \tilde{s}_{\mu, \lambda}(x) < \infty$.⁴
2. Firm strategies optimize an oblivious value function:

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

3. The value of entry is zero or the entry rate is zero (or both):

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

⁴Because there could be a countably infinite number of individual states, the system of equations below may not have a unique solution. If that is the case, we choose the minimum non-negative solution. The industry state $\tilde{s}_{\mu, \lambda}(x) = \lambda \tilde{T}_x$, where \tilde{T}_x is the expected number of visits to state x under strategy μ (Kemeny, Snell, and Knapp 1976). Also, note that under our assumptions, the latter part of the condition is immediately satisfied for all $m \in \mathbb{N}$.

Using the fact that expected discounted profits are bounded (by Assumption 3.1) together with Assumption 2.2, it is possible to show that the supremum in part 2 of the definition can always be attained simultaneously for all x by a common strategy μ' , for all $m \in \mathbb{N} \cup \{\infty\}$ (see Puterman (1994)).

In some infinite models there will not exist any SE with positive entry, i.e., such that $\lambda > 0$. For example, consider an industry where single-period profits exhibit increasing returns to investment, so that every incumbent firm has incentives to grow arbitrarily large (for a fixed industry state). In this case, new entrants might not be able to recover the entry cost (even if it is arbitrarily small), because they will need to invest an arbitrarily large amount of resources to catch up with incumbents. Because we are interested in situations where competition between incumbent firms is actually observed, we assume throughout the paper that there always exists an SE with positive entry and focus on this case from here on.

5 Stationary Equilibrium Approximates MPE Asymptotically

Infinite industry models are intended to approximate finite models with large numbers of firms. In this section we formalize this notion. In particular, we establish that if a light-tail condition is satisfied, then infinite model SE offer close approximations to finite model MPE when markets are large. Note that we do asymptotics in the market size and not the number of firms because, in our model, the number of firms is endogenous.

We begin with some notation and a definition of the *asymptotic Markov equilibrium property*, which formalizes the sense in which the approximation becomes exact. Next, we introduce the light-tail condition and we prove the main result of the paper. Finally, we provide examples to illustrate some of the main ideas behind the results.

5.1 Notation

We index functions and random quantities associated with market size $m \in \mathbb{N}$ with a superscript (m) . Let $V^{(m)}$ represent the value function when the market size is m . The random vector $s_t^{(m)}$ denotes the industry state at time t when every firm uses strategy $\mu^{(m)}$ and the entry rate is $\lambda^{(m)}$. We assume that the process $\{s_t^{(m)} : t \geq 0\}$ is ergodic and admits a unique invariant distribution. In order to simplify our analysis, we assume that the initial industry state $s_0^{(m)}$ is sampled from the invariant distribution. Hence, $s_t^{(m)}$ is a stationary process; $s_t^{(m)}$ is distributed according to its invariant distribution for all $t \geq 0$. Note that this assumption does not affect long-run asymptotic results since for any initial condition the process approaches

stationarity as time progresses. Let $xs^{(m)} \equiv \tilde{s}_{\mu^{(m)}, \lambda^{(m)}} = E[s_t^{(m)}]$. It will be helpful to decompose $s_t^{(m)}$ according to $s_t^{(m)} = f_t^{(m)} n_t^{(m)}$, where $f_t^{(m)}$ is the random vector that represents the fraction of firms in each state and $n_t^{(m)}$ is the total number of firms, respectively. Similarly, let $\tilde{f}^{(m)} \equiv E[f_t^{(m)}]$ denote the expected fraction of firms in each state and $\tilde{n}^{(m)} \equiv E[n_t^{(m)}] = \sum_{x \in \mathcal{X}} \tilde{s}^{(m)}(x)$ denote the expected number of firms. It is easy to check that $\tilde{f}^{(m)} = \frac{\tilde{s}^{(m)}}{\tilde{n}^{(m)}}$.

5.2 Asymptotic Markov Equilibrium Property

Similarly to Weintraub, Benkard, and Van Roy (2008), we define the following concept to formalize the sense in which the approximation becomes exact.

Definition 5.1. A sequence $\{(\mu^{(m)}, \lambda^{(m)}) | m \in \mathbb{N}\} \in \mathcal{M} \times \Lambda$ possesses the asymptotic Markov equilibrium (AME) property if for all $x \in \mathcal{X}$,

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

Recall that the process s_t is taken to be stationary, and therefore, this expectation does not depend on t . The definition of AME assesses approximation error at each firm state x in terms of the amount by which a firm at state x can increase its actual expected net present value by deviating from the strategy $\mu^{(m)}$, and instead following a best response strategy. Recall that an MPE requires that the expression in square brackets equals zero for all states (x, s) . The AME property instead considers the benefit of deviating to an optimal strategy starting from each firm state x , averaged over the invariant distribution of industry states. If a sequence possesses the AME property, then, asymptotically firms make near-optimal decisions in states that have significant probability of occurrence. Hence, MPE strategies should be well-approximated in the set of relevant states.

5.3 Light-Tail Condition

It turns out that even if the number of firms becomes infinite along a sequence of oblivious strategies, that is not enough to guarantee that the AME property will hold. If the market also tends to be concentrated along the sequence, for example if the market is usually dominated by a small fraction of very large firms, then the AME property may not hold. To ensure the AME property, we need to impose a ‘‘light-tail’’ condition that rules out this kind of market concentration.

Note that $\frac{d \ln \pi_m(y, f, n)}{df(x)}$ is the semi-elasticity of one period profits with respect to the fraction of firms in

state x . We define the *maximal absolute semi-elasticity function*:

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

For each x , $g(x)$ is the maximum rate of relative change of any firm's single-period profit that could result from a small change in the fraction of firms at individual state x . If \mathcal{X} is single-dimensional and a firm's state represents size, then since larger competitors tend to have greater influence on firm profits, $g(x)$ would typically increase with x (and could be unbounded).

We introduce a light-tail condition that extends the condition of Weintraub, Benkard, and Van Roy (2008) to multidimensional firm states. Consider a sequence of oblivious strategies and entry rates $\{(\mu^{(m)}, \lambda^{(m)}) | m \in \mathbb{N}\}$. For each m , let $\tilde{x}^{(m)} \sim \tilde{f}^{(m)}$, that is, $\tilde{x}^{(m)}$ is a random vector with probability mass function $\tilde{f}^{(m)}$. The random vector $\tilde{x}^{(m)}$ can be interpreted as the individual state of a firm that is randomly sampled from among all incumbents when the industry state $s_t^{(m)}$ is distributed according to its invariant distribution.

Assumption 5.1. *For all individual states $x \in \mathcal{X}$, $g(x) < \infty$. For all $\epsilon > 0$, there exists $z \in \mathbb{N}$ such that*

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

for all market sizes $m \in \mathbb{N}$.

The assumption first requires that for all x the maximum rate of relative change of any firm's single-period profit that could result from a small change in the fraction of firms at individual state x is finite. The second part of the assumption controls for the appearance of "dominant firms". To better understand this part, consider the special case where there exists a random vector \tilde{x} , such that, $\tilde{x}^{(m)} = \tilde{x}$, for all $m \in \mathbb{N}$ (which will be the case in Theorem 5.1 for SE). In that case, the second part is equivalent to $E[g(\tilde{x})] < \infty$. The assumption essentially requires that states where a small change in the fraction of firms has a large impact on the profits of other firms must have a small probability under the invariant distribution, so that the expected impact of a randomly sampled incumbent is finite. In practice this typically means that very large firms (and hence high concentration) rarely occur under the invariant distribution. We provide an example in Section 5.6. Note also that if the set of individual states \mathcal{X} is finite, then the light-tail condition is immediately satisfied.

5.4 Main Result

Let $\mu \in \tilde{\mathcal{M}}$ and $\lambda \in \tilde{\Lambda}$ be an SE. We want to show that the SE approximates MPE as the market size grows. For these purposes, we define the following sequence of strategies and entry rate functions: $\mu^{(m)} = \mu$ and $\lambda^{(m)} = m\lambda$, for all $m \in \mathbb{N}$. The strategies $\mu^{(m)}$ are the same for all m and correspond to the SE strategy. The entry rate $\lambda^{(m)}$ takes into account the interpretation of the entry rate in an infinite industry model as the ratio of entering firms to the market size. It is simple to show that with these strategies and entry rates, each process $\{s_t^{(m)} : t \geq 0\}$ is ergodic and admits a unique invariant distribution for all $m \in \mathbb{N}$. Also note that, because $\mu^{(m)} = \mu$, $\forall m \in \mathbb{N}$, in this case $\tilde{f}^{(m)} = \tilde{s}_{\mu, \lambda} / \sum_{x \in \mathcal{X}} \tilde{s}_{\mu, \lambda}(x) \equiv \tilde{f}^{(\infty)}$, $\forall m \in \mathbb{N}$.

We now provide the main result of the paper. The proof can be found in the appendix.

Theorem 5.1. *Under Assumptions 3.1 and 5.1, the sequence $\{(\mu^{(m)}, \lambda^{(m)}) | m \in \mathbb{N}\}$ possesses the AME property.*

Theorem 5.1 states that under the light-tail condition, an SE of the infinite industry model approximates MPE of the finite model asymptotically as the market size grows. In the SE of the infinite industry model it is assumed that the industry state is constant because there are an infinite number of firms. However, in an industry with a large but finite expected number of firms, the industry state is not constant. The light-tail condition guarantees that in large markets, (1) the actual industry state is always close (in an appropriate sense) to its average; and (2) that, therefore, movements around the average industry state have a small impact on expected discounted profits. The AME property can be established using these facts.

A significant part of the proof follows a similar argument to the result in Weintraub, Benkard, and Van Roy (2008) that establishes that, under the light-tail condition, a sequence of oblivious equilibria over market sizes $m \in \mathbb{N}$ possesses the AME property. However, that proof must be extended in two ways. First, we extend it to deal with multidimensional firms' states (see Theorem 6.1 and Appendix C). In addition, because we are dealing with an SE (an OE for $m = \infty$), an additional argument is required to complete the proof. In particular, we need to show that expected discounted profits evaluated at the long-run expected industry state in market size m converge to expected discounted profits for $m = \infty$ as m grows.

Finally, note that the AME property is a statement about value functions. In addition, one can show that if the sequence $(\mu^{(m)}, \lambda^{(m)})$ possesses the AME property, then it is also true that $\lim_{m \rightarrow \infty} E_{\mu^{(m)}, \lambda^{(m)}} \left[\beta V^{(m)}(x^e, s_t^{(m)} | \mu^{(m)}, \lambda^{(m)}) \right] = \kappa$. Hence, entry rates at relevant states should also be well approximated.

5.5 Example: Importance of Light-Tail Condition

To illustrate the importance of the light-tail assumption, in this section we outline an example for which the SE is not light-tailed and which does not possess the AME property. The details of the example can be found in the Appendix B.

We consider an industry model with a profit function that can be written as $\pi_m(x, f, n) = \pi_m(x, \sum_x f(z)z)$. That is, the state is single-dimensional and profits depend only on the first moment of the normalized industry state. For the model considered, $g(x) \propto x$. In the appendix, we show that the SE (μ, λ) of the model is heavy-tailed, that is, $\sum_x \tilde{f}^{(\infty)}(x)x = \infty$.

A key step in arguing that the AME property holds in the example would be to show using a law of large numbers that the first moment of the actual normalized industry state sampled according to the invariant distribution converges to the first moment of the expected normalized industry state as the market size grows. If this is true then for large markets actual profits are close to profits evaluated at the average state. Recall that the SE strategy is based on the latter.

We show that this step fails in the example. For market size m the first moment of the actual normalized industry state is well approximated by $1/m \sum_{i=1}^m x_i$, where x_i 's are i.i.d. random variables sampled from the long-run distribution of firm states in equilibrium ($\tilde{f}^{(\infty)}$). The first moment of the average normalized industry state is $E[x_i]$. Using standard results in probability we show that $1/m \sum_{i=1}^m x_i - E[x_i \mathbf{1}[x_i \leq m]] \Rightarrow Y$ as $m \rightarrow \infty$ where Y is a random variable with a non-degenerate distribution (see Durrett (1996)). Hence, the first moment of the actual normalized industry state is not necessarily close to its average for large markets; if the SE is heavy-tailed the uncertainty does not vanish in the limit. From this, we argue that the AME property does not hold.

5.6 Example: Logit Demand System with Price Competition

In this section we illustrate the light-tail condition with an example. We consider an industry with differentiated products, where each firm's state variable is single-dimensional and represents the quality of its product. There are m consumers in the market. In period t , consumer j receives utility u_{ijt} from consuming the good produced by firm i given by:

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

where Y is the consumer's income and p_{it} is the price of the good produced by firm i . ν_{ijt} are i.i.d. random variables distributed Gumbel that represent unobserved characteristics for each consumer-good pair. There

is also an outside good that provides consumers zero utility. We assume consumers buy at most one product each period and that they choose the product that maximizes utility. Under these assumptions our demand system is a classical logit model.

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

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

where σ represents the market share function from the logit model.

One can show that, if $\lim_{m \rightarrow \infty} n(m)/m = \tilde{c}$, then, for all x ,

$$\lim_{m \rightarrow \infty} \pi_m(x, f, n(m)) = \pi_\infty(x, \tilde{c}f) = \frac{N(x, \tilde{p})}{\tilde{c} \sum_{y \in \mathbb{N}} f(y)N(y, \tilde{p})} (\tilde{p} - b),$$

where $\tilde{p} = (Y + b\theta_2)/(1 + \theta_2)$ and $N(y, p) = (y + 1)^{\theta_1}(Y - p)^{\theta_2}$. The limit profit function π_∞ corresponds to a logit model of monopolistic competition (Besanko, Perry, and Spady 1990).

Weintraub, Benkard, and Van Roy (2008) show that, in this model, the function $g(x)$ takes a very simple form, $g(x) \propto x^{\theta_1}$. Therefore, the light-tail condition amounts to a simple condition on the equilibrium distribution of firm states. Under our assumptions, such a condition is equivalent to a condition on the equilibrium size distribution of firms. If $\theta_1 \leq 1$ then the light-tail condition is satisfied if $E[\tilde{x}] < \infty$, i.e., if the average firm individual state is finite. This condition allows for relatively “fat-tailed” distributions. For example, if the tail of \tilde{x} decays like a log-normal distribution, then the condition is satisfied. On the other hand, if the tail of \tilde{x} decays like a Pareto distribution with parameter one, which does not have a finite first moment, then the condition would not be satisfied (like in the previous example).

It is interesting to note that in this example, for an SE with positive entry rate, (μ, λ) , it must be that $\pi_\infty(x, \tilde{s}_{\mu, \lambda}) > 0, \forall x \in \mathbb{N}$. If not, firms would not be able to recover the entry cost. Hence,

$$(5.1) \quad \sum_{y \in \mathbb{N}} \tilde{f}^{(\infty)}(y)N(y, \tilde{p}) = (Y - \tilde{p})^{\theta_2} \sum_{y \in \mathbb{N}} \tilde{f}^{(\infty)}(y)(y + 1)^{\theta_1} < \infty.$$

For this model, the light-tail assumption is satisfied if $E[\tilde{x}^{\theta_1}] = \sum_{y \in \mathbb{N}} \tilde{f}^{(\infty)}(y)y^{\theta_1} < \infty$, which is implied by expression (5.1). Hence, for this model, the light-tail condition is immediately satisfied for SE with positive entry rates. The same observation is also obtained in others models of monopolistic competition à la Dixit and Stiglitz (1977).

6 Relationship Between SE and OE

In this section we explore more closely the connection between OE of finite models and SE of infinite models. As discussed in Section 3, an SE of the infinite model is an OE of that model. In Section 5 we showed that under a light-tail condition, infinite model SE approximate MPE of finite models asymptotically as the market size grows. Similarly, Weintraub, Benkard, and Van Roy (2008) show that OE of finite models approximate MPE asymptotically as the market size grows. We review and extend this result in Section 6.1.

In Sections 6.2 and 6.3 we further show that under the light-tail condition the set of OE of a finite model approaches the set of SE of the infinite model as the market size grows. Our results are related to the upper-hemicontinuity and lower-hemicontinuity of the OE correspondence at the point where number of firms and market size become infinite. The results of this section imply that not only light-tailed OE and SE approximate MPE asymptotically, but also that they essentially approximate the *same subset of MPE*.

6.1 Asymptotic Results for Oblivious Equilibria

Let $(\mu^{(m)}, \lambda^{(m)})$ denote an oblivious equilibrium for market size m . We consider the same notation as in Section 5.1. Weintraub, Benkard, and Van Roy (2008) show the following result for a specialized version of our finite model in which individual states are single-dimensional; a similar proof with some changes that we describe in Appendix C is valid for our more general model.

Theorem 6.1. *Under Assumptions 3.1, and 5.1, the sequence $\{(\mu^{(m)}, \lambda^{(m)}) | m \in \mathbb{N}\}$ of oblivious equilibria possesses the AME property.*

The result is the analog to Theorem 5.1 for a sequence of OE in finite models.

Additionally, we prove the following result that we will use below. All proofs of this section can be found in the appendix. First, we define $\|f\|_{1,g} = \sum_x |f(x)|g(x)$. Note that if $f \in \mathcal{S}_1$ and X is a random vector with distribution f , then $E[g(X)] = \|f\|_{1,g}$. Let $A = \{f | f \in \mathcal{S}_1, \|f\|_{1,g} < \infty\}$ be a normed space endowed with the norm $\|\cdot\|_{1,g}$. If $f \in A$, we say f is light-tailed. Let $B = \{\tilde{f}^{(m)} | m \in \mathbb{N}\}$. Before stating the result we strengthen the light-tail condition. The additional assumption implies that $\|\cdot\|_{1,g}$ is indeed a norm.

Assumption 6.1. *There exists $\delta > 0$, such that for all individual states $x \in \mathcal{X}$, $g(x) > \delta$. Assumption 5.1 holds.*

Proposition 6.1. *Suppose Assumptions 3.1, and 6.1 hold. Then,*

1. The closure of $B \subseteq A$ is compact. Hence, the sequence of expected normalized industry states, $\{\tilde{f}^{(m)} | m \in \mathbb{N}\}$, has a convergent subsequence to an element in A , that is, to a light-tailed distribution.
2. Asymptotically, the expected number of firms $\tilde{n}^{(m)}$ and the OE entry rate $\tilde{\lambda}^{(m)}$ grow proportionally with the market size m .
3. The expected time inside the industry $\tilde{T}^{(m)}$ remains uniformly bounded over all market sizes.

Proposition 6.1.1 establishes that in a light-tailed sequence of OE, $\{\tilde{f}^{(m)} | m \in \mathbb{N}\}$ has a subsequence that converges to a light-tailed distribution. If the sequence of expected normalized industry states $\{\tilde{f}^{(m)} | m \in \mathbb{N}\}$ has a unique accumulation point, then it converges. Additionally, by the second part of the theorem, the expected number of firms grows proportionally to the market size asymptotically. In the limit the resulting market structure shares an important characteristic with the market structure assumed in an infinite industry model; in each individual state there will be an infinite number of firms. These observations underscore the close connection between finite model OE and infinite model SE.

6.2 Upper-Hemicontinuity

In this section we show that if a sequence of OE satisfies the light-tail condition, then it converges to a light-tailed infinite model SE. The previous statement corresponds to the upper-hemicontinuity of the OE correspondence at the point where number of firms and market size become infinite, when the sequence of OE satisfies the light-tail condition.

Recall that Proposition 6.1 states that if the light-tail condition is satisfied, then the sequence $\tilde{f}^{(m)}$ is contained in a compact set and that $\tilde{n}^{(m)}$ grows proportionally to the market size. To simplify our analysis, we will further assume that the sequences $\tilde{f}^{(m)}$ and $\tilde{n}^{(m)}/m$ for the OE that are being considered are “well-behaved” in the sense that they have one accumulation point each.

Assumption 6.2. *The sequences $\{\tilde{f}^{(m)} | m \in \mathbb{N}\}$ and $\{\tilde{n}^{(m)}/m | m \in \mathbb{N}\}$ have one accumulation point each.*

Assumption 6.2 together with Proposition 6.1 imply that both sequences converge. Let $\tilde{f} \equiv \lim_m \tilde{f}^{(m)}$ (with convergence in the $\|\cdot\|_{1,g}$ norm), and $\tilde{c} \equiv \lim_m \tilde{n}^{(m)}/m$ (note that $\tilde{c} > 0$). Note that by Proposition 6.1, $\|\tilde{f}\|_{1,g} < \infty$.

To state the next result we need one more assumption. Let $T^{(m)}$ be the random variable that represents the time a firm spends inside industry m when using strategy $\mu^{(m)}$. By definition, $\tilde{T}^{(m)} = E[T^{(m)}]$. In Proposition 6.1 we established that, under the light-tail condition, $\sup_m E[T^{(m)}] < \infty$. For technical reasons, we introduce a slightly stronger assumption.

Assumption 6.3. *The sequence of random variables $\{T^{(m)}|m \in \mathbb{N}\}$ is uniformly integrable.*

Note that if there exists $\gamma > 0$, such that, $\sup_m E \left[(T^{(m)})^{1+\gamma} \right] < \infty$, then Assumption 6.3 holds. The condition is slightly stronger than requiring uniformly bounded first moments of $\tilde{T}^{(m)}$.

We have the following result.

Theorem 6.2. *Suppose Assumptions 3.1, 6.1, 6.2, and 6.3 hold. Then, the sequence of OE $\{(\mu^{(m)}, \lambda^{(m)})|m \in \mathbb{N}\}$ converges to a light-tailed infinite model SE. Formally, there exists an infinite model SE (μ, λ) , such that, for all $x \in \mathcal{X}$, $\lim_m \mu^{(m)}(x) = \mu(x)$ and $\lim_m \lambda^{(m)}/m = \lambda$. Moreover, $\tilde{s}_{\mu, \lambda} = \tilde{c}\tilde{f}$.*

The light-tail assumption is useful to show the AME property because it guarantees that in large markets, movements around the average industry state have a small impact on expected discounted profits. Similarly, the light-tail condition is key to obtain the continuity results in this section, because it guarantees that a sequence of OE long-run expected industry states converges in a way such that small deviations from the limit industry state have a small impact on profits.

6.3 Lower-Hemicontinuity

In this section we show that all sequences of oblivious strategies that approach a light-tailed SE satisfy the OE conditions asymptotically. This result is related to the lower-hemicontinuity of the OE correspondence at the point where number of firms and market size become infinite, when SE in the infinite model satisfies the light-tail condition. We begin with some definitions.

Definition 6.1. *A sequence $\{(\mu^{(m)}, \lambda^{(m)})|m \in \mathbb{N}\} \in \tilde{\mathcal{M}} \times \tilde{\Lambda}$ possesses the asymptotic oblivious equilibrium (AOE) property, if for all $x \in \mathcal{X}$,*

$$\lim_{m \rightarrow \infty} \sup_{\mu' \in \tilde{\mathcal{M}}} \tilde{V}^{(m)}(x|\mu', \mu^{(m)}, \lambda^{(m)}) - \tilde{V}^{(m)}(x|\mu^{(m)}, \lambda^{(m)}) = 0 \quad \text{and}$$

$$\lim_{m \rightarrow \infty} \beta \tilde{V}^{(m)}(x^e|\mu^{(m)}, \lambda^{(m)}) = \kappa.$$

The AOE property requires that the OE conditions are satisfied asymptotically.

Definition 6.2. *We say that the sequence $\{(\mu^{(m)}, \lambda^{(m)})|m \in \mathbb{N}\} \in \tilde{\mathcal{M}} \times \tilde{\Lambda}$ converges to an SE (μ, λ) if for all $x \in \mathcal{X}$, $\lim_{m \rightarrow \infty} \mu^{(m)}(x) = \mu(x)$, $\lim_{m \rightarrow \infty} \lambda^{(m)}/m = \lambda$, $\lim_{m \rightarrow \infty} \left\| \tilde{s}_{\mu^{(m)}, \lambda^{(m)}} / \sum_{x \in \mathcal{X}} \tilde{s}_{\mu^{(m)}, \lambda^{(m)}}(x) - \tilde{s}_{\mu, \lambda} / \sum_{x \in \mathcal{X}} \tilde{s}_{\mu, \lambda}(x) \right\|_{1, g} = 0$, and $\lim_{m \rightarrow \infty} \sum_{x \in \mathcal{X}} \tilde{s}_{\mu^{(m)}, \lambda^{(m)}}(x)/m = \sum_{x \in \mathcal{X}} \tilde{s}_{\mu, \lambda}(x)$.*

The definition establishes a norm under which a sequence of oblivious strategies and entry rate functions converges: strategies, entry rates, associated vector of expected fraction of firms, and expected number of firms should converge in an appropriate sense. We have the following result.

Theorem 6.3. *Suppose Assumption 3.1 holds. Suppose (μ, λ) is an infinite model SE that satisfies Assumption 6.1. Let $\{(\mu^{(m)}, \lambda^{(m)}) | m \in \mathbb{N}\} \in \tilde{\mathcal{M}} \times \tilde{\Lambda}$ be a sequence of oblivious strategies and entry rate functions that converges to (μ, λ) . Then, $(\mu^{(m)}, \lambda^{(m)})$ possesses the AOE property.*

The result establishes a weaker property than lower-hemicontinuity of the OE correspondence at the point where number of firms and market size become infinite, because we only require that the sequences of strategies possess the AOE property (hence, that the OE conditions are satisfied asymptotically); we do not require that they are sequences of OE. On the other hand, what is shown is stronger than lower-hemicontinuity which would only require that there *exists* a sequence of OE that converges to the SE.

7 Conclusions

Infinite models came into wide use because of their tractability, and because it was believed that they would provide a good approximation to real world industries with finite numbers of firms. In this paper we provided foundations for this approach. In a fairly general setting we have shown that there is in fact a very close relationship between SE in infinite models, and MPE and OE in finite models. In large markets it therefore should not matter which approach is taken.

There is, however, one important caveat to these results. If the equilibrium being analyzed is “light-tailed”, then the infinite model SE will be approximately the same as both MPE and OE for large markets. However, as our example in section 5.5 shows, if the equilibrium being considered is instead heavy-tailed, then the approximation may fail.

As mentioned above, the light-tail condition is a condition on equilibrium outcomes and not model primitives. While it is useful in practice because it can be checked directly after an equilibrium has been found, an interesting open question for future research would be to find conditions over the model primitives that guarantee that all SE are light-tailed.

Finally, note that the paper focuses on SE in infinite models. We focus on this case because it is of practical interest and has been extensively used in the literature. It seems likely, however, that the logic of our results extend more generally, so that infinite model equilibria approximate finite model MPE more generally, not just for the stationary case. We leave this extension to future research.

References

- Adlakha, S., R. Johari, G. Y. Weintraub, and A. Goldsmith (2010). Mean field analysis for large population stochastic games. In *IEEE Conference on Decision and Control (forthcoming)*.
- Allen, B. and M. Hellwig (1986a). Bertrand-Edgeworth oligopoly in large markets. *The Review of Economic Studies* 53(2), 175 – 204.
- Allen, B. and M. Hellwig (1986b). Price-setting firms and the oligopolistic foundations of perfect competition. *The American Economic Review* 76(2), 387 – 392.
- Bertsekas, D. P. (2001). *Dynamic Programming and Optimal Control, Vol. 2* (Second ed.). Athena Scientific.
- Besanko, D., M. K. Perry, and R. H. Spady (1990). The logit model of monopolistic competition: Brand diversity. *The Journal of Industrial Economics* 38(4), 397 – 415.
- Caplin, A. and B. Nalebuff (1991). Aggregation and imperfect competition - on the existence of equilibrium. *Econometrica* 59(1), 25 – 59.
- Dixit, A. K. and J. E. Stiglitz (1977). Monopolistic competition and optimum product diversity. *American Economic Review* 67(3), 297 – 308.
- Doraszelski, U. and A. Pakes (2007). A framework for applied dynamic analysis in IO. In *Handbook of Industrial Organization, Volume 3*. North-Holland, Amsterdam.
- Doraszelski, U. and M. Satterthwaite (2010). Computable markov-perfect industry dynamics. *RAND Journal of Economics* 41(2), 215 – 243.
- Durrett, R. (1996). *Probability: Theory and Examples* (Second ed.). Duxbury Press.
- Ericson, R. and A. Pakes (1995). Markov-perfect industry dynamics: A framework for empirical work. *Review of Economic Studies* 62(1), 53 – 82.
- Escobar, J. (2008). Existence of pure and behavior strategy equilibrium in dynamic stochastic games. Working Paper, Stanford University.
- Gabaix, X. (2008). The granular origins of aggregate fluctuations. Working Paper, New York University.
- Hopenhayn, H. and R. Rogerson (1993). Job turnover and policy evaluation - a general equilibrium-analysis. *Journal of Political Economy* 101(5), 915 – 938.
- Hopenhayn, H. A. (1992). Entry, exit and firm dynamics in long run equilibrium. *Econometrica* 60(5), 1127 – 1150.

- Iacovone, L., B. Javorcik, W. Keller, and J. Tybout (2009). Walmart in Mexico: The impact of FDI on innovation and industry productivity. working paper, Penn State University.
- Jones, L. E. (1987). The efficiency of monopolistically competitive equilibria in large economies: commodity differentiation with gross substitutes. *Journal of Economic Theory* 41, 356 – 391.
- Jovanovic, B. (1987). Micro shocks and aggregate risk. *The Quarterly Journal of Economics* 102(2), 395 – 410.
- Kemeny, J. G., J. L. Snell, and A. W. Knapp (1976). *Denumerable Markov Chains*. Springer.
- Klette, T. J. and S. Kortum (2004). Innovating firms and aggregate innovation. *Journal of Political Economy* 112(5), 986 – 1018.
- Luttmer, E. G. J. (2007). Selection, growth, and the size distribution of firms. *Quarterly Journal of Economics* 122(3), 1103 – 1144.
- Marsden, J. E. and M. J. Hoffman (1993). *Elementary Classical Analysis* (Second ed.). W. H. Freeman and Company.
- Mas-Colell, A. (1982). The Cournotian foundation of Walrasian equilibrium theory: an exposition of recent theory. In *Hildenbrand, W. (ed.), Advances in Economic Theory: Invited Papers for the Fourth World Congress of the Econometric Society*, pp. 183–224.
- Mas-Colell, A. (1983). Walrasian equilibria as limits of noncooperative equilibria. Part I: mixed strategies. *Journal of Economic Theory* 30, 153 – 170.
- Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica* 71(6), 1695 – 1725.
- Mitchell, M. F. (2000). The scope and organization of production: firm dynamics over the learning curve. *Rand Journal of Economics* 31(1), 180 – 205.
- Mukherjee, A. and V. Kadiyali (2008). The competitive dynamics of DVD release timing and pricing. working paper, Cornell University.
- Novshek, W. (1985). Perfectly competitive markets as the limits of Cournot markets. *Journal of Economic Theory* 35, 72 – 82.
- Novshek, W. and H. Sonnenschein (1978). Cournot and Walras equilibrium. *Journal of Economic Theory* 19, 223 – 266.

- Novshek, W. and H. Sonnenschein (1983). Walrasian equilibria as limits of noncooperative equilibria. Part II: pure strategies. *Journal of Economic Theory* 30, 171 – 187.
- Puterman, M. L. (1994). *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (First ed.). John Wiley & Sons, Inc.
- Qi, S. (2008). The impact of advertising regulation on industry: The cigarette advertising ban of 1971. Working paper, University of Minnesota.
- Roberts, D. J. and A. Postlewaite (1976). The incentives for price-taking behavior in large exchange economies. *Econometrica* 4(1), 115 – 127.
- Royden, H. (1988). *Real Analysis* (Third ed.). Pearson Education.
- Thurk, J. (2009a). International protection of intellectual property: A quantitative assessment. working paper, University of Texas at Austin.
- Thurk, J. (2009b). Market effects of patent reform in the U.S. semiconductor industry. working paper, University of Texas at Austin.
- Weintraub, G. Y., C. L. Benkard, and B. Van Roy (2008). Markov perfect industry dynamics with many firms. *Econometrica* 76(6), 1375–1411.
- Xu, Y. (2008). A structural empirical model of R&D, firm heterogeneity, and industry evolution. working paper, NYU University.

A Proof of Theorem 5.1

Proof of Theorem 5.1. Under Assumptions 3.1 and 5.1, the sequence $\{(\mu^{(m)}, \lambda^{(m)}) | m \in \mathbb{N}\}$ possesses the AME property.

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

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

Let

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

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

Let us write

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

To complete the proof, we will establish that $E_{\mu^{(m)}, \lambda^{(m)}}[A^{(m)}(x, s_t^{(m)})]$ converges to zero. An analogous argument that we omit for brevity establishes that $E_{\mu^{(m)}, \lambda^{(m)}}[B^{(m)}(x, s_t^{(m)})]$ converges to zero as well.

Because μ and λ attain an SE, we have

$$\sup_{\mu' \in \tilde{\mathcal{M}}} \tilde{V}^{(\infty)}(x | \mu', \mu, \lambda) = \sup_{\mu' \in \tilde{\mathcal{M}}} \tilde{V}^{(\infty)}(x | \mu', \mu, \lambda) = \tilde{V}^{(\infty)}(x | \mu, \lambda), \quad \forall x \in \mathcal{X},$$

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

Therefore,

$$\begin{aligned}\tilde{V}^{(\infty)}(x | \mu, \lambda) &\geq E_{\mu^{*(m)}, \mu^{(m)}, \lambda^{(m)}} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} (\pi_{\infty}(x_{ik}, \tilde{s}_{\mu, \lambda}) - d(\iota_{ik})) + \beta^{\tau_i-t} \phi_{i, \tau_i} \mid x_{it} = x, s_{-i, t} = s \right] \\ \text{(A.1)} \quad &\equiv \tilde{V}^{*(m)}(x, s | \mu, \lambda).\end{aligned}$$

Hence,

$$A^{(m)}(x, s) \leq V^{(m)}(x, s | \mu^{*(m)}, \mu^{(m)}, \lambda^{(m)}) - \tilde{V}^{*(m)}(x, s | \mu, \lambda).$$

Let $\Delta_{it}^{(m)} = |\pi_m(x_{it}, s_{-i, t}^{(m)}) - \pi_{\infty}(x_{it}, \tilde{s}_{\mu, \lambda})|$. It follows that

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

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

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

By the triangle inequality,

$$\Delta_{ik}^{(m)} \leq |\pi_m(x_{ik}, s_{-i,k}^{(m)}) - \pi_m(x_{ik}, \tilde{s}^{(m)})| + |\pi_m(x_{ik}, \tilde{s}^{(m)}) - \pi_\infty(x_{ik}, \tilde{s}_{\mu,\lambda})|.$$

Using Assumption 5.1 and following a similar argument to the extension of the result in Weintraub, Benkard, and Van Roy (2008) to multidimensional firm states that establishes that, under the light-tail condition, a sequence of oblivious equilibria possesses the AME property (see Proposition 6.1 and Appendix C), one can show that:

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

Note that $\tilde{s}^{(m)} = m\tilde{s}_{\mu,\lambda}$. Hence, by Assumptions 3.1.2, for all $x \in \mathcal{X}$, $\lim_{m \rightarrow \infty} |\pi_m(x, \tilde{s}^{(m)}) - \pi_\infty(x, \tilde{s}_{\mu,\lambda})| = 0$. Assumption 3.1.3 together with the dominated convergence theorem imply that,

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

The result follows. □

B Importance of Light-Tail Condition

In this appendix we discuss in detail the example mentioned in Section 5.5. To simplify the exposition of the example we assume that investment decisions and entry rates are determined exogenously; only exit decisions are determined endogenously in equilibrium. It is straightforward to construct similar examples for which investment decision and entry rates are also derived endogenously in equilibrium.

Specifically, we consider the following version of our dynamic industry model:

- $\mathcal{X} = \{\underline{x}, \underline{x} + 1, \underline{x} + 2, \dots\}$, where \underline{x} is a positive integer.
- Investment decisions are exogenous. In particular, a firm in state x at time t that does not exit the industry will transition to state $x + 1$ at time $t + 1$ for sure.
- Entry decisions are exogenous. In particular, each period there is an entry rate equal to λ to individual state \underline{x} .

- Exit decisions are determined endogenously in equilibrium. With out loss of generality, we assume the distribution of the sell-off value ϕ_{it} is uniform[0, 1].
- The single period profit function is given by $\pi_m(x, f, n) = \pi^1(x) + \pi_m^2(x, f, n)$, where

$$\begin{aligned}\pi^1(x) &= \frac{1}{\beta} \left(\frac{x-1}{x} \right)^2 - E \left[\max \left(\phi, \left(\frac{x}{x+1} \right)^2 \right) \right], \\ \pi_m^2(x, f, n) &= \frac{x}{\left(\sum_z f(z)z - \sum_{z < m} f^*(z)z \right)^2 + x},\end{aligned}$$

where f^* is a pre-determined pmf such that $\sum_z f^*(z)z = \infty$ (but note that $\sum_{z < m} f^*(z)z < \infty, \forall m$). Note that π^1 is only a function of the firms' own individual state. We interpret π^1 as the profits each incumbent firm garners from serving a different local market as a monopoly. π_m^2 is similar to a typical profit function derived from a Dixit and Stiglitz (1977) model of monopolistic competition, but where profits depend on the difference between $\sum_z f(z)z$ and $\sum_{z < m} f^*(z)z$ (as oppose to just $\sum_z f(z)z$). The specific functional forms for π^1 and π_m^2 are chosen for convenience in the example.

It is simple to show that π_m satisfy Assumptions 2.1 and 3.1, except for Assumption 3.1.2. However, it is simple to modify π_m so that the latter assumption is also satisfied. For clarity of presentation, we do not make that modification. Also, note that for all $c > 0$, and for all $f \in \mathcal{S}_1$, $\lim_{m \rightarrow \infty} \pi_m^2(x, f, cm) \equiv \pi_\infty^2(x, cf) = 0$. Therefore, $\pi_\infty(x, cf) = \pi^1(x)$.

We now show that the SE of this model does not satisfy the light-tail condition; we refer to the SE as heavy-tailed. For the sequence of profit functions defined above, we have that $g(x) \propto x$. Therefore, if an SE (μ, λ) is heavy-tailed, then $\sum_x \tilde{f}^{(\infty)}(x)x = \infty$. Recall that in our notation $\tilde{f}^{(\infty)} = \tilde{s}_{\mu, \lambda} / \sum_x \tilde{s}_{\mu, \lambda}(x)$.

Given single period profits π_∞ , it is simple to show using standard dynamic programming arguments that under the optimal strategy, a firm in state x stays in the industry with probability $(x/(x+1))^2$. Given these transitions, $\tilde{s}_{\mu, \lambda}$ can be derived recursively and is given by (see equation (4.1)):

$$\begin{aligned}\tilde{s}_{\mu, \lambda}(\underline{x}) &= \lambda \\ \tilde{s}_{\mu, \lambda}(\underline{x} + 1) &= \tilde{s}_{\mu, \lambda}(\underline{x})(\underline{x}/(\underline{x} + 1))^2 = \lambda(\underline{x}/(\underline{x} + 1))^2 \\ \tilde{s}_{\mu, \lambda}(\underline{x} + 2) &= \tilde{s}_{\mu, \lambda}(\underline{x} + 1)((\underline{x} + 1)/(\underline{x} + 2))^2 = \lambda(\underline{x}/(\underline{x} + 1))^2((\underline{x} + 1)/(\underline{x} + 2))^2 = \lambda(\underline{x}/(\underline{x} + 2))^2 \\ &\dots \\ \tilde{s}_{\mu, \lambda}(x) &= \lambda(\underline{x}/x)^2.\end{aligned}$$

Note that $\sum_x \tilde{s}_{\mu,\lambda}(x) < \infty$, as required by the definition of SE. Also, $\sum_x \tilde{f}^{(\infty)}(x)x = \infty$, so the SE is heavy-tailed.

We now show that the AME property does not hold for a particular specification of the model. Recall that f^* in π_m^2 is a pre-determined pmf such that $\sum_z f^*(z)z = \infty$. In particular, we let $f^* = \tilde{f}^{(\infty)}$. In this case,

$$(B.1) \quad \begin{aligned} E[\pi_m^2(x, s_{-i,k}^{(m)})] &= E \left[\frac{x}{\left(\sum_z f_{-i,k}^{(m)}(z)z - \sum_{z < m} \tilde{f}^{(\infty)}(z)z \right)^2 + x} \right] \\ &\approx E \left[\frac{x}{\left((1/m) \sum_{i=1}^m x_i - E[x_i \mathbf{1}[x_i < m]] \right)^2 + x} \right], \end{aligned}$$

where $s_{-i,k}^{(m)}$ is distributed according to the invariant distribution of $\{s_t^{(m)} : t \geq 0\}$, and x_i are i.i.d. random variables with pmf $\tilde{f}^{(\infty)}$. The last approximation is valid for large m , because $n_k^{(m)}$ is a Poisson random variable with a mean that grows proportionally with m (see Weintraub, Benkard, and Van Roy (2008)). Taking the limit as $m \rightarrow \infty$, using the continuous mapping theorem together with the bounded convergence theorem, and Theorem 7.7. in Durrett (1996), we obtain that the previous expectation converges to $E \left[\frac{x}{Y^2 + x} \right]$, where Y is a non-degenerate random variable. This is strictly positive. Hence, for all market sizes m , there exists a set of states visited with a probability that is uniformly bounded away from zero under the invariant distribution, for which π_m^2 is uniformly bounded away from zero over all market sizes. On the other hand, $\pi_\infty^2(x, \tilde{s}_{\mu,\lambda}) = 0$. From these two arguments, we observe that the AME property does not hold. An exit strategy that keeps track of the industry state and considers actual profits will yield longer times inside the industry (to garner profits given by π_m^2) compared to the SE exit strategy (that considers π_∞^2 that is always zero).

C Extension of AME Property for Sequence of OE

In this appendix we extend the result in Weintraub, Benkard, and Van Roy (2008) for our model with multidimensional firm states that establishes the AME property for a sequence of OE. The proof follows similar steps to Weintraub, Benkard, and Van Roy (2008). We only present the steps and results for which the analysis differs significantly from that paper.

Theorem 6.1. *Under Assumptions 3.1, and 5.1, the sequence $\{(\mu^{(m)}, \lambda^{(m)}) | m \in \mathbb{N}\}$ of oblivious equilibria possesses the AME property.*

Proof. We prove some preliminary results.

Lemma C.1. (Lemma A.5 in Weintraub, Benkard, and Van Roy (2008)) For all $x \in \mathcal{X}$,

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

Proof. Assume for contradiction that $\sup_m \tilde{V}^{(m)}(x|\mu^{(m)}, \lambda^{(m)}) = \infty$. By Assumption 2.2.4, there is a sequence of investment decisions that induces a strictly positive probability of reaching x from x^e in a finite number of periods, for all m . These decisions involve an investment cost that is uniformly bounded above over all m . Therefore, $\sup_m \tilde{V}^{(m)}(x^e|\mu^{(m)}, \lambda^{(m)}) = \infty$, contradicting the zero profit condition. \square

Lemma C.2. (Theorem 5.2 in Weintraub, Benkard, and Van Roy (2008)) Suppose Assumptions 3.1 and 5.1 hold. Then, the sequence of expected number of firms $\tilde{n}^{(m)}$ satisfies $\liminf_m \tilde{n}^{(m)}/m > 0$.

Proof. Assume for contradiction that $\liminf_m \tilde{n}^{(m)}/m = 0$. Hence, there is an increasing subsequence m_k such that $\lim_k \tilde{n}^{(m_k)}/m_k = 0$. By Assumption 3.1.2, $\lim_k \pi_{m_k}(x, f, \tilde{n}^{(m_k)}) = \infty$, for all $x \in \mathcal{X}$ and $f \in \mathcal{S}_1$. Now, for all m ,

$$\begin{aligned} |\ln \pi_m(x, \tilde{f}^{(m)}, \tilde{n}^{(m)}) - \ln \pi_m(x, f, \tilde{n}^{(m)})| &\leq \|\tilde{f}^{(m)} - f\|_{1,g} \\ &\leq \sum_{\{x \in \mathcal{X}: \|x\|_\infty \leq z\}} |\tilde{f}^{(m)}(x) - f(x)|g(x) + \epsilon \\ &\leq 2 \max_{\{x \in \mathcal{X}: \|x\|_\infty \leq z\}} g(x) + \epsilon < \infty. \end{aligned}$$

The first inequality follows by equation (A.3) in Lemma A.6 in Weintraub, Benkard, and Van Roy (2008). The second follows by choosing z and ϵ according to Assumption 5.1, and choosing $f(x) = 0, \forall x \in \mathcal{X}$ such that $\|x\|_\infty > z$.

Hence, $\lim_m \pi_m(x, \tilde{f}^{(m)}, \tilde{n}^{(m)}) = \infty$, contradicting Lemma C.1. It follows that $\liminf_m \tilde{n}^{(m)}/m > 0$. \square

Lemma C.3. (Lemma A.6 in Weintraub, Benkard, and Van Roy (2008)) Suppose Assumptions 3.1 and 5.1 hold. Then, for all $x \in \mathcal{X}$,

$$\sup_m \sup_{\mu \in \mathcal{M}} E_\mu \left[\sum_{t=0}^{\infty} \beta^t \sup_{f \in \mathcal{S}_1} \pi_m(x_{it}, f, \tilde{n}^{(m)}) \Big| x_{i0} = x \right] < \infty.$$

Proof. By Lemma C.2, $\liminf_m \tilde{n}^{(m)}/m > 0$. By Assumption 3.1.3, there exists $d, e, k > 0$, such that $\pi_m(x, f, \tilde{n}^{(m)}) \leq d\|x\|_\infty^k + e$, for all $x \in \mathcal{X}$, $f \in \mathcal{S}_1$ and $m \in \mathbb{N}$. This together with assumption 2.2.3 implies,

$$\sup_m \sup_{\mu \in \mathcal{M}} E_\mu \left[\sum_{t=0}^{\infty} \beta^t \sup_{f \in \mathcal{S}_1} \pi_m(x_{it}, f, \tilde{n}^{(m)}) \Big| x_{i0} = x \right] \leq \sum_{t=0}^{\infty} \beta^t \left(d(\|x\|_\infty + t\bar{w})^k + e \right) < \infty.$$

The result follows. \square

Proof of Theorem 6.1. The rest of the proof is similar to Weintraub, Benkard, and Van Roy (2008) with minor modifications. We note one such modification in the last set of inequalities in Lemma A.11. For the purpose of this proof, we will assume that all expectations are conditioned on $x_{i0} = x$ and $s_{-i,0}^{(m)} \sim q^{(m)}$, where $q^{(m)}$ is the invariant distribution of $\{s_t^{(m)} : t \geq 0\}$. Let $\Delta_{it}^{(m)} = |\pi_m(x_{it}, f_{-i,t}^{(m)}, \tilde{n}^{(m)}) - \pi_m(x_{it}, \tilde{s}_{\mu^{(m)}, \lambda^{(m)}})|$. Let $Z_t^{(m)}$ denote the event $\|f_{-i,t}^{(m)} - \tilde{f}^{(m)}\|_{1,g} \geq \delta$, for $\delta > 0$. Recall that $(\mu^{(m)}, \lambda^{(m)})$ is an OE for market m . Consider any sequence $\{\mu'^{(m)} \in \mathcal{M}\}$. Then,

$$\begin{aligned} \sum_{t=0}^{\infty} \beta^t E_{\mu'^{(m)}, \mu^{(m)}, \lambda^{(m)}} \left[\Delta_{it}^{(m)} \mathbf{1}_{Z_t^{(m)}} \right] &\leq \sum_{t=0}^{\infty} \beta^t E_{\mu'^{(m)}, \mu^{(m)}, \lambda^{(m)}} \left[2 \sup_{f \in \mathcal{S}_1} \pi_m(x_{it}, f, \tilde{n}^{(m)}) \mathbf{1}_{Z_t^{(m)}} \right] \\ &\leq 2\mathcal{P}[Z_t^{(m)}] \sum_{t=0}^{\infty} \beta^t \left(d(\|x\|_\infty + t\bar{w})^k + e \right), \end{aligned}$$

where the last inequality follows by a similar argument to Lemma C.3. $\mathcal{P}[Z_t^{(m)}]$ is the same for all t and converges to zero as $m \rightarrow \infty$. The sum is finite, so the result follows. \square

D Proofs Section 6

Proof of Proposition 6.1. *Suppose Assumptions 3.1, and 6.1 hold. Then,*

1. *The closure of $B \subseteq A$ is compact. Hence, the sequence of expected normalized industry states, $\{\tilde{f}^{(m)} | m \in \mathbb{N}\}$, has a convergent subsequence to an element in A , that is, to a light-tailed distribution.*
2. *Asymptotically, the expected number of firms $\tilde{n}^{(m)}$ and the OE entry rate $\tilde{\lambda}^{(m)}$ grow proportionally with the market size m .*
3. *The expected time inside the industry $\tilde{T}^{(m)}$ remains uniformly bounded over all market sizes.*

Proof. Part (1). It is simple to show that the space A is complete. Using Assumption 6.1, it is straightforward to prove that B is a totally bounded subset of A . Therefore, the closure of B is compact (Marsden

and Hoffman 1993). The sequence $\{\tilde{f}^{(m)} | m \in \mathbb{N}\}$ has a convergent subsequence to an element in A by the Bolzano-Weierstrass Theorem.

Part (2). Lemma C.2 implies $\liminf_m \tilde{n}^{(m)}/m > 0$. Assumption 3.1.2 implies that if $\lim_m \tilde{n}^{(m)}/m = \infty$, then $\lim_m \pi_m(x, f, \tilde{n}^{(m)}) = 0$, for all $x \in \mathcal{X}$ and $f \in \mathcal{S}_1$. In this case, it is possible to show that the OE value function converges to zero as $m \rightarrow \infty$ and firms cannot recover the entry cost. Therefore, it must be that $\limsup_m \tilde{n}^{(m)}/m < \infty$ and indeed $\tilde{n}^{(m)}$ grows proportionally with the market size m . Additionally, we know that $\tilde{n}^{(m)} = \tilde{\lambda}^{(m)} \tilde{T}^{(m)}$. Hence, by part (3) of the theorem it must be that the OE entry rate also grows proportionally with the market size m .

Part (3). Suppose, for the sake of a contradiction, that the expected time inside the industry $\tilde{T}^{(m)}$ does not remain uniformly bounded over all market sizes. Because under our assumptions, for all m , $\tilde{T}^{(m)} < \infty$, this implies that $\limsup_{m \rightarrow \infty} \tilde{T}^{(m)} = \infty$. We will prove that in this case, for all $z \in \mathbb{N}$, $\limsup_{m \rightarrow \infty} \sum_{\{x \in \mathcal{X} : \|x\|_\infty > z\}} \tilde{f}^{(m)}(x) = 1$. Since by the first part of Assumption 6.1, there exists $\delta > 0$, such that for all $x \in \mathcal{X}$, $g(x) > \delta$, this contradicts the second part of Assumption 6.1 (light-tail assumption).

We define, for all $x \in \mathcal{X}$, $\tilde{T}_x^{(m)}$ as the expected number of visits a firm makes to state x when using strategy $\mu^{(m)}$. Note that $\tilde{T}^{(m)} = \sum_{x \in \mathcal{X}} \tilde{T}_x^{(m)}$.

Lemma C.1 shows that for all $x \in \mathcal{X}$, $\sup_m \tilde{V}^{(m)}(x | \mu^{(m)}, \lambda^{(m)}) < \infty$. Recall that by Assumption 2.2.1, the sell-off value has support in \mathfrak{R}_+ . Hence, for all $x \in \mathcal{X}$, each time a firm visits state x , there is a probability uniformly bounded away from zero over all market sizes, that the firm will exit the industry. The exit process from state x can be represented as a geometric random variable. It follows that, for all $x \in \mathcal{X}$, $\sup_m \tilde{T}_x^{(m)} < \infty$.

We can write $\tilde{f}^{(m)}(x) = \tilde{T}_x^{(m)} / \tilde{T}^{(m)}$. Therefore, for all $x \in \mathcal{X}$,

$$\sum_{\{x \in \mathcal{X} : \|x\|_\infty > z\}} \tilde{f}^{(m)}(x) = \frac{\sum_{\{x \in \mathcal{X} : \|x\|_\infty > z\}} \tilde{T}_x^{(m)}}{\sum_{x \in \mathcal{X}} \tilde{T}_x^{(m)}}.$$

Because for all $x \in \mathcal{X}$, $\sup_m \tilde{T}_x^{(m)} < \infty$ and, by assumption, $\limsup_{m \rightarrow \infty} \tilde{T}^{(m)} = \infty$, it must be that, for all $z \in \mathbb{N}$, $\limsup_m \sum_{\{x \in \mathcal{X} : \|x\|_\infty > z\}} \tilde{T}_x^{(m)} = \infty$. Hence, for all $z \in \mathbb{N}$, $\limsup_m \sum_{\{x \in \mathcal{X} : \|x\|_\infty > z\}} \tilde{f}^{(m)}(x) = 1$, and the second part of Assumption 6.1 (light-tail assumption) is violated. Therefore, it must be that the expected time inside the industry $\tilde{T}^{(m)}$ remains uniformly bounded over all market sizes. \square

Proof of Theorem 6.2. Suppose Assumptions 3.1, 6.1, 6.2, and 6.3 hold. Then, the sequence of OE $\{(\mu^{(m)}, \lambda^{(m)}) | m \in \mathbb{N}\}$ converges to a light-tailed infinite model SE. Formally, there exists an infinite model SE (μ, λ) , such that, for all $x \in \mathcal{X}$, $\lim_m \mu^{(m)}(x) = \mu(x)$ and $\lim_m \lambda^{(m)}/m = \lambda$. Moreover, $\tilde{s}_{\mu, \lambda} = \tilde{c}\tilde{f}$.

Proof. We start by proving some preliminary lemmas.

Lemma D.1. *Suppose Assumptions 3.1, 6.1, and 6.2 hold. Then, for all $x \in \mathcal{X}$, $\lim_{m \rightarrow \infty} \pi_m(x, \tilde{s}^{(m)}) = \pi_\infty(x, \tilde{c}\tilde{f})$. Moreover, there exists $d, e > 0$ and $k > 0$, such that, $\pi_\infty(x, \tilde{c}\tilde{f}) \leq d\|x\|_\infty^k + e$, for all $x \in \mathcal{X}$.*

Proof.

$$(D.1) \quad \begin{aligned} \pi_m(x, \tilde{s}^{(m)}) - \pi_\infty(x, \tilde{c}\tilde{f}) &= \left(\pi_m(x, \tilde{f}^{(m)}, \tilde{n}^{(m)}) - \pi_m(x, \tilde{f}, \tilde{n}^{(m)}) \right) \\ &+ \left(\pi_m(x, \tilde{f}, \tilde{n}^{(m)}) - \pi_\infty(x, \tilde{c}\tilde{f}) \right) \end{aligned}$$

The second term trivially converges to zero by Assumptions 3.1.2 and 6.2, and Proposition 6.1. For the first term, consider that by Assumption 6.2 and Proposition 6.1, $\lim_{m \rightarrow \infty} \|\tilde{f}^{(m)} - \tilde{f}\|_{1,g} = 0$. Then, by Assumptions 2.1.1 and 2.1.2, and Lemma A.10 in Weintraub, Benkard, and Van Roy (2008), it follows that $\lim_{m \rightarrow \infty} \pi_m(x, \tilde{f}^{(m)}, \tilde{n}^{(m)}) / \pi_m(x, \tilde{f}, \tilde{n}^{(m)}) = 1$. By Assumption 3.1.3, $\sup_m \pi_m(x, \tilde{f}, \tilde{n}^{(m)}) < \infty$. Therefore, $\lim_{m \rightarrow \infty} \pi_m(x, \tilde{f}^{(m)}, \tilde{n}^{(m)}) - \pi_m(x, \tilde{f}, \tilde{n}^{(m)}) = 0$. That there exists $d, e > 0$ and $k > 0$, such that, $\pi_\infty(x, \tilde{c}\tilde{f}) \leq d\|x\|_\infty^k + e$, for all $x \in \mathcal{X}$, follows directly from Assumption 3.1.3. \square

Now, we state that the oblivious equilibrium value function $\tilde{V}^{(m)}$ and the oblivious equilibrium strategy $\mu^{(m)}$ converge to the optimal value function and optimal strategy, respectively, of a firm's dynamic programming problem with profits given by $\pi_\infty(x, \tilde{c}\tilde{f})$. To abbreviate, with some abuse of notation, we let $\tilde{V}^{(m)}(x) \equiv \tilde{V}^{(m)}(x | \mu^{(m)}, \lambda^{(m)})$. and, for all $x \in \mathcal{X}$, we let

$$\tilde{V}^{(\infty)}(x) = \sup_{\mu' \in \tilde{\mathcal{M}}} E_{\mu'} \left[\sum_{k=t}^{\tau_i} \beta^{k-t} \left(\pi_\infty(x_{ik}, \tilde{c}\tilde{f}) - d(\iota_{ik}) \right) + \beta^{\tau_i-t} \phi_{i,\tau_i} \Big| x_{it} = x \right].$$

Let $\tilde{\mu}^{(\infty)} \in \tilde{\mathcal{M}}$ be the strategy that achieves the maximum above (the value function and optimal strategy are well defined by Assumption 2.2 and 3.1.3, and the results in Puterman (1994)). Hence, the value function $\tilde{V}^{(\infty)}$ and the strategy $\tilde{\mu}^{(\infty)}$ are the optimal value function and optimal strategy, respectively, of a firm's dynamic programming problem with profits given by $\pi_\infty(x, \tilde{c}\tilde{f})$.

Lemma D.2. *Suppose Assumptions 3.1, 6.1, and 6.2 hold. Then, for all $x \in \mathcal{X}$, $\lim_{m \rightarrow \infty} \tilde{V}^{(m)}(x) = \tilde{V}^{(\infty)}(x)$, and $\lim_{m \rightarrow \infty} \mu^{(m)}(x) = \tilde{\mu}^{(\infty)}(x)$.*

Proof. We prove convergence of the value functions. The proof of convergence of the strategy functions is analogous. The proof follows two main steps. First, we show that $\tilde{V}^{(m)}$ lies on a compact set. Then, we

prove that the limit of any convergent subsequence of $\tilde{V}^{(m)}$ must be $\tilde{V}^{(\infty)}$. This implies that, for all $x \in \mathcal{X}$, $\lim_{m \rightarrow \infty} \tilde{V}^{(m)}(x) = \tilde{V}^{(\infty)}(x)$.

Lemma C.1 establishes that $\bar{v}(x) \equiv \sup_m \tilde{V}^{(m)}(x) < \infty$, for all x . Therefore, $V^{(m)}(x) \in [0, \bar{v}(x)]$, for all m, x . By Tychonoff Theorem (Royden 1988), the product set $\times_{x \in \mathcal{X}} [0, \bar{v}(x)]$ is compact in the topology of pointwise convergence (or the product topology).

Suppose $\tilde{V}^{(m_n)}$ is a converging subsequence of $\tilde{V}^{(m)}$. Let, for all $x \in \mathcal{X}$, $\tilde{v}^{(\infty)}(x)$ be the (pointwise) limit of $\tilde{V}^{(m_n)}(x)$. We prove that, for all $x \in \mathcal{X}$, $\tilde{v}^{(\infty)}(x) = \tilde{V}^{(\infty)}(x)$. That is, the limit of any convergent subsequence of $\tilde{V}^{(m)}$ is $\tilde{V}^{(\infty)}$. Let us define the following sequence of dynamic programming operators for value functions $V \leq \bar{v}$:

$$(D.2) \quad \begin{aligned} F^{(m)}V(x) &= \pi_m(x, \tilde{s}^{(m)}) + E \left[\max_{\iota \in \mathcal{I}} \left\{ \phi_{it}, \sup_{\iota \in \mathcal{I}} (-d(\iota) + \beta E_{\mu, \lambda} [V(x_{i,t+1}) | x_{it} = x, \iota_{it} = \iota]) \right\} \right] \\ &\equiv \pi_m(x, \tilde{s}^{(m)}) + QV(x), \end{aligned}$$

for all $x \in \mathcal{X}$. The operator $F^{(\infty)}$ is defined as above, but with the profit function $\pi_\infty(x, \tilde{c}\tilde{f})$.

Using Assumption 2.2, the operator Q can be written as:

$$\begin{aligned} QV(x) &= \sup_{\iota \in \mathcal{I}, \rho \in [0, \bar{\rho}]} \left(-d(\iota) + \beta \sum_{\{y \in \mathcal{X}: y=x+z, \|z\|_\infty \leq \bar{w}\}} \mathcal{P} [x_{i,t+1} = y | x_{i,t} = x, \iota_{it} = \iota] V(y) \right) \\ &\times \mathcal{P}[\phi_{it} < \rho] + E[\phi_{it} | \phi_{it} \geq \rho] \mathcal{P}[\phi_{it} \geq \rho] \\ &\equiv \sup_{\iota \in \mathcal{I}, \rho \in [0, \bar{\rho}]} f_x(\iota, \rho, V), \end{aligned}$$

for $\bar{\rho} < \infty$. It is simple to check that, by Assumption 2.2, the operator f_x is continuous in the topology of pointwise convergence. Hence, by Berge's maximum theorem, the operator $QV(x)$ is continuous in the topology of pointwise convergence. Additionally, by Lemma D.1, for all $x \in \mathcal{X}$, $\lim_m \pi_m(x, \tilde{s}^{(m)}) = \pi_\infty(x, \tilde{c}\tilde{f})$. Therefore, for all $x \in \mathcal{X}$, $\lim_{n \rightarrow \infty} F^{(m_n)} \tilde{V}^{(m_n)}(x) = F^{(\infty)} \tilde{v}^{(\infty)}(x)$. Additionally, $\tilde{V}^{(m_n)}$ is the oblivious equilibrium value function, therefore, for all $n \in \mathbb{N}$, and for all $x \in \mathcal{X}$, it must solve Bellman's equation: $F^{(m_n)} \tilde{V}^{(m_n)}(x) = \tilde{V}^{(m_n)}(x)$ (because of Assumptions 2.1.1 and 2.2, and the results in Bertsekas (2001)). We conclude that, for all $x \in \mathcal{X}$, $F^{(\infty)} \tilde{v}^{(\infty)}(x) = \tilde{v}^{(\infty)}(x)$. Moreover, using the fact that $\pi_\infty(x, \tilde{c}\tilde{f})$ does not grow faster than a polynomial in $\|x\|_\infty$ as x grows (by Lemma D.1) and Assumption 2.2.3, Theorem 6.10.4 in Puterman (1994) implies that $\tilde{V}^{(\infty)}$ is the unique solution of $F^{(\infty)}v(x) = v(x)$, for all $x \in \mathcal{X}$, among the set of functions with finite $\|\cdot\|_{\infty, w}$ norm with $w(x) = 1/(d\|x\|_\infty^k + e)$.⁵ Using a

⁵The sup weighted norm is defined as $\|v\|_{\infty, w} = \sup_{x \in \mathcal{X}} w(x)|v(x)|$.

similar argument to Lemma C.3, it is simple to show that $v^{(\infty)}$ has finite $\|\cdot\|_{\infty, w}$ norm. Therefore, for all $x \in \mathcal{X}$, $\tilde{v}^{(\infty)}(x) = \tilde{V}^{(\infty)}(x)$. That is, the limit of any convergent subsequence of $\tilde{V}^{(m)}$ must be $\tilde{V}^{(\infty)}$. The result follows. \square

We prove one final lemma, about the expected time firms spend inside the industry. Like in the proof of Proposition 6.1, we define, for all $x \in \mathcal{X}$, $\tilde{T}_x^{(m)}$ as the expected number of visits a firm makes to state x when using strategy $\mu^{(m)}$. Note that the expected time inside the industry $\tilde{T}^{(m)} = \sum_{x \in \mathcal{X}} \tilde{T}_x^{(m)}$. Similarly, we define $\tilde{T}_x^{(\infty)}$ and $\tilde{T}^{(\infty)}$ as the expected number of visits a firm makes to state x and the expected time the firm spends inside the industry, respectively, when using strategy $\tilde{\mu}^{(\infty)}$. In the next lemma, we show that the expected number of visits to a state under oblivious equilibrium strategies $\mu^{(m)}$ converges to the expected number of visits under strategy $\tilde{\mu}^{(\infty)}$.

Lemma D.3. *Suppose Assumptions 3.1, 6.1, 6.2, and 6.3 hold. Then, for all $x \in \mathcal{X}$, $\lim_{m \rightarrow \infty} \tilde{T}_x^{(m)} = \tilde{T}_x^{(\infty)}$. Moreover, $\lim_{m \rightarrow \infty} \tilde{T}^{(m)} = \tilde{T}^{(\infty)} < \infty$.*

Proof. First we prove that $\lim_{m \rightarrow \infty} \tilde{T}_x^{(m)} = \tilde{T}_x^{(\infty)}$. Let $P_\mu^t(x, y)$ be the probability that a firm in state x will be in state y , t time periods from now when using strategy μ . The expected number of visits to state x can be written as $\tilde{T}_x^{(m)} = \sum_{t=0}^{\infty} P_{\mu^{(m)}}^t(x^e, x) = \sum_{t=0}^T P_{\mu^{(m)}}^t(x^e, x) + \sum_{t>T} P_{\mu^{(m)}}^t(x^e, x)$. Using Assumption 2.2 and Lemma D.2, it is simple to show that, for all t , $x \in \mathcal{X}$, $\lim_{m \rightarrow \infty} P_{\mu^{(m)}}^t(x^e, x) = P_{\mu^{(\infty)}}^t(x^e, x)$. Clearly, for all $t, m \in \mathbb{N}$, $x \in \mathcal{X}$, $P_{\mu^{(m)}}^t(x^e, x) \leq \mathcal{P}[T^{(m)} \geq t]$. If a firm is in state x after t time periods, it must be inside the industry after t time periods. It is simple to show that Assumption 6.3 implies that $\lim_{T \rightarrow \infty} \sup_m \sum_{t>T} \mathcal{P}[T^{(m)} \geq t] = 0$. Therefore, $\lim_{T \rightarrow \infty} \sup_m \sum_{t>T} P_{\mu^{(m)}}^t(x^e, x) = 0$. It follows that $\lim_{m \rightarrow \infty} \tilde{T}_x^{(m)} = \sum_{t=0}^{\infty} P_{\mu^{(\infty)}}^t(x^e, x) = \tilde{T}_x^{(\infty)}$.

Now we prove that $\lim_{m \rightarrow \infty} \tilde{T}^{(m)} = \tilde{T}^{(\infty)} < \infty$. Provided the limit exists, $\tilde{T}^{(\infty)}$ is finite, because $\tilde{T}^{(m)}$ remains uniformly bounded over all market sizes. Note that $\tilde{T}^{(m)} = \sum_{x \in \mathcal{X}} \tilde{T}_x^{(m)}$ and $\tilde{f}^{(m)}(x) = \tilde{T}_x^{(m)} / \tilde{T}^{(m)}$. Assumption 6.1 implies $\lim_{z \rightarrow \infty} \sup_m \sum_{\{x \in \mathcal{X} : \|x\|_{\infty} > z\}} \tilde{f}^{(m)}(x) = 0$. This together with $\sup_m \tilde{T}^{(m)} < \infty$ implies $\lim_{z \rightarrow \infty} \sup_m \sum_{\{x \in \mathcal{X} : \|x\|_{\infty} > z\}} \tilde{T}_x^{(m)} = 0$. Therefore, $\lim_{m \rightarrow \infty} \tilde{T}^{(m)} = \sum_{x \in \mathcal{X}} \tilde{T}_x^{(\infty)} = \tilde{T}^{(\infty)} < \infty$. \square

We use the previous lemmas to prove the theorem.

Proof of Theorem 6.2.

Lemmas D.1 and D.2 establish that, for all $x \in \mathcal{X}$, $\lim_{m \rightarrow \infty} \mu^{(m)}(x) = \tilde{\mu}^{(\infty)}(x)$ and $\lim_{m \rightarrow \infty} \tilde{V}^{(m)}(x) = \tilde{V}^{(\infty)}(x)$. It is simple to observe that $\tilde{n}^{(m)} = \lambda^{(m)} \tilde{T}^{(m)}$. By Assumption 6.2, Proposition 6.1, and Lemma

D.3, $\lim_m \lambda^{(m)}/m = \tilde{c}/\tilde{T}^{(\infty)} \equiv \tilde{\lambda}^{(\infty)}$. We prove that if $\mu(x) \equiv \tilde{\mu}^{(\infty)}(x)$, for all $x \in \mathcal{X}$, and $\lambda \equiv \tilde{\lambda}^{(\infty)}$, then (μ, λ) constitute an infinite model SE.

First, note that because $\lim_m \tilde{V}^{(m)}(x^e) = \tilde{V}^{(\infty)}(x^e)$ and $\tilde{V}^{(m)} = \kappa/\beta$, for all m , it must be that $\tilde{V}^{(\infty)}(x^e) = \kappa/\beta$. Second, recall that the value function $\tilde{V}^{(\infty)}$ and the strategy $\tilde{\mu}^{(\infty)}$ are the optimal value function and optimal strategy, respectively, of a firm's dynamic programming problem with profits given by $\pi_\infty(x, \tilde{c}\tilde{f})$.

Therefore, to establish that (μ, λ) defined above constitute an infinite model SE it is enough to show that $\tilde{s}_{\mu, \lambda} = \tilde{c}\tilde{f}$. It is simple to observe that the vector $\tilde{s}_{\mu, \lambda} = \lambda(\tilde{T}_x^{(\infty)})_{x \in \mathcal{X}} = \tilde{c}/\tilde{T}^{(\infty)} \cdot (\tilde{T}_x^{(\infty)})_{x \in \mathcal{X}}$. By Lemma D.3, $\tilde{T}_x^{(\infty)} = \lim_m \tilde{T}_x^{(m)}$ and $\tilde{T}^{(\infty)} = \lim_m \tilde{T}^{(m)}$. Because $\tilde{f}^{(m)} = (\tilde{T}_x^{(m)})_{x \in \mathcal{X}}/\tilde{T}^{(m)}$, we have that $\tilde{f} = (\tilde{T}_x^{(\infty)})_{x \in \mathcal{X}}/\tilde{T}^{(\infty)}$. Hence, $\tilde{s}_{\mu, \lambda} = \tilde{c}\tilde{f}$ as needed. \square

Proof of Theorem 6.3. *Suppose Assumption 3.1 holds. Suppose (μ, λ) is an infinite model SE that satisfies Assumption 6.1. Let $\{(\mu^{(m)}, \lambda^{(m)}) | m \in \mathbb{N}\} \in \tilde{\mathcal{M}} \times \tilde{\Lambda}$ be a sequence of oblivious strategies and entry rate functions that converges to (μ, λ) . Then, $(\mu^{(m)}, \lambda^{(m)})$ possesses the AOE property.*

Proof. The argument in Lemma D.1 establishes that $\lim_m \pi_m(x, \tilde{s}_{\mu^{(m)}, \lambda^{(m)}}) = \pi_\infty(x, \tilde{s}_{\mu, \lambda})$. Using a similar argument to Lemma D.2 one can show that, for all $x \in \mathcal{X}$, $\lim_{m \rightarrow \infty} \sup_{\mu' \in \tilde{\mathcal{M}}} \tilde{V}^{(m)}(x | \mu', \mu^{(m)}, \lambda^{(m)}) = \tilde{V}^{(\infty)}(x | \mu, \lambda)$. Similarly, since, for all $x \in \mathcal{X}$, $\lim_{m \rightarrow \infty} \mu^{(m)}(x) = \mu(x)$, it follows that $\lim_{m \rightarrow \infty} \tilde{V}^{(m)}(x | \mu^{(m)}, \lambda^{(m)}) = \tilde{V}^{(\infty)}(x, | \mu, \lambda)$. Hence, for all $x \in \mathcal{X}$, $\lim_{m \rightarrow \infty} \sup_{\mu' \in \tilde{\mathcal{M}}} \tilde{V}^{(m)}(x | \mu', \mu^{(m)}, \lambda^{(m)}) - \tilde{V}^{(m)}(x | \mu^{(m)}, \lambda^{(m)}) = 0$ and $\lim_{m \rightarrow \infty} \beta \tilde{V}^{(m)}(x^e | \mu^{(m)}, \lambda^{(m)}) = \kappa$. Therefore, $(\mu^{(m)}, \lambda^{(m)})$ possesses the AOE property. \square