# A Framework for Applied Dynamic Analysis in IO<sup>\*</sup>

Ulrich Doraszelski and Ariel Pakes

September 1, 2006

#### Abstract

This paper reviews a framework for numerically analyzing dynamic interactions in imperfectly competitive industries. The framework dates back to Ericson & Pakes (1995), but it is based on equilibrium notions that had been available for some time before, and it has been extended in many ways by different authors since. The framework requires as input a set of primitives which describe the institutional structure in the industry to be analyzed. The framework outputs profits and policies for every incumbent and potential entrant at each possible state of the industry. These policies can be used to simulate the distribution of sample paths for all firms from any initial industry structure. The sample paths generated by the model can be quite different depending on the primitives, and most of the extensions were designed to enable the framework to accommodate empirically relevant cases that required modification of the initial structure. The sample paths possess similar properties to those observed in (the recently available) panel data sets on industries. These sample paths can be used either for an analysis of the likely response to a policy or an environmental change, or as the model's implication in an estimation algorithm.

We begin with a review of an elementary version of the framework and a report on what is known about its analytic properties. Much of the rest of the paper deals with computational issues. We start with an introduction to iterative techniques for computing equilibrium that are analogous to the techniques used to compute the solution to single agent dynamic programming problems. This includes discussions of the determinants of the computational burden of these techniques, and the mechanism implicitly used to select an equilibrium when multiple equilibria are possible. We then outline a number of techniques that might be used to reduce the computational burden of the iterative algorithm. This section includes discussions of both the implications of differences in modeling assumptions used in the alternative techniques, and a discussion of the likely relevance of the different techniques for different institutional structures. A separate section reports on a technique for computing multiple equilibria from the same set of primitives. The paper concludes with a review of applications of the framework

<sup>&</sup>lt;sup>\*</sup>We would like to thank numerous colleagues for helpful discussions. Joe Harrington, Ken Judd, Rob Porter, and Michael Riordan gave us useful feedback on an earlier draft of this chapter. Pakes would like to thank his NSF grant for financial support and acknowledge the important role of Richard Ericson and Paul McGuire in developing the techniques reported on here. This paper supercedes his NBER discussion paper with the same title (Pakes 2000).

and a brief discussion of areas where further development of the framework would seem warranted.

# Contents

1	Introduction	4				
<b>2</b>	odel					
3	Equilibrium					
	3.1 Existence	. 16				
	3.2 Characterization	. 17				
	3.3 Multiplicity	. 19				
4	Introduction to Computation					
	4.1 Gaussian Methods	. 23				
	4.2 Computational Burden	. 29				
	4.3 Equilibrium Selection	. 31				
<b>5</b>	Alleviating the Computational Burden					
	5.1 Overview	. 32				
	5.2 Continuous-Time Models	. 34				
	5.3 Stochastic Approximation Algorithm	. 39				
	5.4 Function Approximation Methods	. 52				
	5.5 Oblivious Equilibrium	. 53				
6	Computing Multiple Equilibria	53				
7	Applications and Extensions	58				
	7.1 Empirics	. 59				
	7.2 Capacity and Advertising Dynamics	. 61				
	7.3 Mergers	. 64				
	7.4 Learning-by-Doing and Network Effects	. 66				
	7.5 Collusion	. 69				
8	Topics for Further Study	71				
9	Conclusions 7					

## 1 Introduction

The applied analysis of dynamic interactions in imperfectly competitive industries is just beginning. This paper reviews a framework which has been developed to facilitate this analysis. The framework dates back to Ericson & Pakes (1995) (hereafter EP), and has been improved and generalized by a series of authors since that article. The EP framework worked to integrate a set of baseline facts from the empirical literature with notions of equilibrium from the theoretical literature.

The facts, drawn in part from the newly available panel data sets on the evolution of firms and industries, made it clear that for a framework to be rich enough to be taken to data it had to allow for; heterogeneity among firms within a market (even for very narrowly defined markets), both firm and market (or sometimes industry) specific sources of uncertainty, and entry and exit (see, e.g., Dunne, Roberts & Samuelson (1988) and Davis & Haltiwanger (1992)). The firm specific uncertainty is needed to account for the fact that we often see both simultaneous entry and exit and rank reversals in the fortunes of firms within a market (no matter how we define "fortunes"). The market (and/or) industry specific uncertainty is needed to rationalize the fact that often the firms competing in a given market (or industry) are subject to changes in costs or demand conditions which cause their profits to be positively correlated.

The EP framework is designed to track an oligopolistic industry over time. In each period, incumbent firms decide whether to remain in the industry and how much to invest, and potential entrants decide whether to enter the industry. Once the investment, entry, and exit decisions are made, firms compete in the product market. Firm heterogeneity is accounted for by encoding all payoff-relevant characteristics of a firm in its "state." Typically a firm's state describes its production capacity, cost structure, and/or the quality of its product. A firm is able to change its state over time through investment, and there is both firm and industry specific variability in the outcomes of the investment. This can generate both diversity of fortunes of seemingly similar firms and positive correlation in their profits.

Formally the EP framework is a dynamic stochastic game with a discrete state space. Dynamic stochastic games have a long tradition in economics. Dating back to Shapley (1953), these games have become central to the analysis of strategic interactions among forward-looking players in dynamic environments. See Filar & Vrieze (1997) and Basar & Olsder (1999) for textbook treatments.<sup>1</sup>

The equilibrium notion used in the EP framework is that of Markov perfect equilibrium. For applied work there were at least two virtues of the Markov perfect notion. First it

<sup>&</sup>lt;sup>1</sup>This approach differs from continuous-time games with a continuum of states which date back to Isaacs (1954) (zero-sum games) and Starr & Ho (1969) (nonzero-sum games). See Basar & Olsder (1999) for a standard presentation of differential games and Dockner, Jorgensen, Van Long & Sorger (2000) for a survey of applications.

involved familiar notions, as Markov processes are used intensively in applied work in related fields. Perhaps more important, however, was the fact that Markov perfect equilibria deliver an empirically tractable way of analyzing outcomes. That is, they allow us to condition on a current state, hopefully a state that we might be able to read off the data, and generate a probability distribution of the subsequent state. That distribution can then be used for either estimation or for numerical analysis.<sup>2</sup>

The Markov perfect notion was used intensively in an influential set of theory papers examining dynamic issues in oligopolistic settings by Maskin & Tirole (1987, 1988*a*, 1988*b*). One lesson from these articles was just how rich a set of outcomes could be generated even from extremely stylized environments; far too stylized for anyone to think of using them to closely approximate behavior in any market. Partly this was the result of the possibility of multiple equilibria, a topic we examine in some detail below. However even more evident was the fact that very different types of behavior could be generated by changes in the primitives of the problem (the nature of the demand or cost functions, or the characteristics of investment processes allowed). Moreover the diverse patterns of outcomes illustrated by different industry studies reinforced the impression that, depending on the details of the institutional structure, different industries could have very different evolutionary patterns.

The EP framework copes with this result by providing a framework with an ability to "plug in" different primitives, and then numerically analyze their result. The framework delivers very little in the way of analytic results of applied interest; i.e. just about anything can happen. Indeed by adopting the framework the researcher essentially gives up on analytic elegance in favor of an ability to numerically analyze the more complex situations that might better approximate what we seem to observe in real data sets. Theoretical results from stylized environments are often used at a later stage as a guide to understanding the economics underlying the phenomena generated by the numerical results.

This paper provides an introduction to the EP framework and then considers issues that arise in using it to numerically analyze results from different specifications for its primitives. We begin by outlining the model in Section 2, and then, in Section 3, summarizing what is known about questions of existence of equilibrium, the characteristics of the equilibrium when it does exist, and the potential for multiple equilibria. Section 4 provides an introduction to techniques for computing an equilibria, and then considers their computational burden. Section 5 considers techniques for alleviating this computational burden. In Section 6 we consider what is known about computing multiple equilibria (from the same set of primitives). Section 7 reviews applications and extensions of the framework. Section 8 points out some (of the many) topics that require further study, and Section 9 concludes.

 $<sup>^{2}</sup>$ In fact there is some debate as to whether the concept of Markov perfect equilibrium restricts policies and outcomes in untestable ways, and hence weaker notions of equilibria, such as the notion of self-confirming equilibrium in Fudenberg & Levine (1993), are more appropriate for applied work. This is largely a topic beyond the scope of this paper, but we come back to a related issue in Section 5 below, where we provide a way of computing a weaker equilibrium notion.

# 2 Model

The EP framework is designed to capture the evolution of an industry with heterogeneous firms. The model is dynamic, time is discrete, and the horizon is infinite. There are two groups of firms, incumbent firms and potential entrants. An incumbent firm has to decide each period whether to remain in the industry and, if so, how much to invest. A potential entrant has to decide whether to enter the industry and, if so, how much to invest. We assume that entry, exit, and investment decisions are made simultaneously at the beginning of the period.

Once these decisions are made, product market competition takes place. For simplicity, we assume that the price or quantity that a firm sets in the product market has no effect on the dynamics of the industry. This reflects the traditional "static-dynamic" breakdown in teaching IO. Due to this "static-dynamic" breakdown, the profit function can be computed "off line" and fed into the algorithm for computing the equilibrium of the dynamic stochastic game. Hence, we essentially treat the per-period profit function as a primitive of the dynamic stochastic game.

For the reader who is familiar with this literature, we note that the model below differs from the model in EP in a few details. First, we treat setup costs and scrap values as privately known random variables in order to ensure the existence of an equilibrium.<sup>3</sup> Second, we assume that exit decisions are implemented after incumbent firms compete in the product market. That is, while entry, exit, and investment decisions are made at the beginning of the period, we assume that their realizations occur at the end of the period. Hence, a firm's current profit from product market competition is completely determined by the current state of the industry. Third, when we allow more than one potential entrant per period to come into the industry we assume that entry decisions, like exit decisions, are made simultaneously. Moreover, we allow a potential entrant to make an initial investment in order to improve the odds that it comes into industry in a more favorable state. Most of these changes make the model easier to compute.

**Incumbent firms.** Incumbent firm *i* is described by its state  $\omega_i \in \Omega$ . EP assume that the state takes on integer values and provide conditions which insure that there are finite upper and lower bounds to the states that can occur in equilibrium (see Section 3.2 for details). Thus, without loss of generality, we take  $\Omega = \{1, 2, \dots, \bar{\omega}\}$ . Typically the state of a firm encodes the characteristics of the products the firm sells or of the production process used to produce those products (e.g., the firm's capital stock or productivity), but it may also include variables that are not as directly "payoff relevant" (as, e.g., in models of collusion, see Section 7.5). A firm is able to change its state over time through its investment  $x_i \geq 0$ .

 $<sup>^{3}</sup>$ Pakes & McGuire (1994) suggest treating a potential entrant's setup cost as a random variable to overcome convergence problems in their algorithm. Gowrisankaran (1995) is the first to make the connection between existence of equilibrium and randomization of both entry and exit decisions and Doraszelski & Satterthwaite (2003) provide a formal proof.

While a higher investment today is no guarantee of a more favorable state tomorrow, it does ensure a more favorable distribution over future states. Since a firm's transition from one state to another is subject to an idiosyncratic shock, there is variability in the fortunes of firms even if they carry out identical strategies. This variability in outcomes is necessary for the model to be able to rationalize the data on the evolution of firms.

Turning from investment to exit, we assume that at the beginning of each period each incumbent firm draws a random scrap value from a distribution  $F(\cdot)$ . Scrap values are independently and identically distributed across firms and periods.<sup>4</sup> Incumbent firm *i* learns its scrap value  $\phi_i$  prior to making its exit and investment decisions, but the scrap values of its rivals remain unknown to it. If the incumbent decides to exit the industry, it collects its scrap value  $\phi_i$  and perishes. Since this decision is conditioned on the privately known  $\phi_i$ , it is a random variable from the perspective of other firms, and we use  $r_i$  to denote the probability that incumbent firm *i* remains in the industry.

**Potential entrants.** In addition to incumbent firms, there are potential entrants. Entry has been treated differently in different papers. This proliferation of entry models is largely because there is no agreement, and very little in the way of empirical guidance, on the appropriate way to model entry; indeed this is one of our suggested directions for future research (see section 8). For concreteness we assume here that there is a finite number  $\mathcal{E}$  of potential entrants in each period and they make simultaneous entry decisions.

Potential entrants are short-lived and base their entry and investment decisions on the net present value of entering today; potential entrants do not take the option value of delaying entry into account. At the beginning of each period each potential entrant draws a random setup cost from a distribution  $F^e(\cdot)$ . Like scrap values, setup costs are privately known and independently and identically distributed across firms and periods. If potential entrant i enters the industry, it incurs its setup cost  $\phi_i^e$  and chooses its initial investment  $x_i^e \geq 0$ . It takes the entrant a period to set up so it does not earn profits until the next period. In that period the entrant becomes an incumbent with an initial state whose distribution depends on  $x_i^e$ . We use  $r_i^e$  to denote the probability that potential entrant i enters the industry.

**States.** At any point in time the industry is completely characterized by the list of states of the incumbent firms.<sup>5</sup> We refer to  $\omega \equiv (\omega_1, \ldots, \omega_n)$  as the state of the industry or as

<sup>&</sup>lt;sup>4</sup>In all the models that have been computed to date the random draws on scrap values are assumed to be independent over time. If this were not the case, then a firm's scrap value and its rivals' beliefs about it would become state variables of the model.

 $<sup>^{5}</sup>$ To be precise, *after* incumbent firms and potential entrants have learned the realization of their scrap value and setup cost for the period, respectively, a complete description of the industry requires a list of scrap value and setup cost in addition to a list of states of the incumbent firms. Since these additional state variables would add to the computational burden, we integrate out over them. In effect, this means that we write down the equations that characterize the equilibrium *before* firms learn their realizations.

the "industry structure". The set of possible industry structures is

$$S = \{ (\omega_1, \dots \omega_n) : \omega_i \in \Omega, \ n \le \bar{n} \},\$$

where *n* is the number of currently active firms and  $\bar{n}$  is the maximum number of firms that are ever active in the industry. We adopt the usual notation for partitioning vectors, e.g., if  $\omega = (\omega_1, \ldots, \omega_n)$ , then  $\omega_{-i} = (\omega_1, \ldots, \omega_{i-1}, \omega_{i+1}, \ldots, \omega_n)$ .

In order to simplify the modeling of entry and exit, it is convenient to use a special state, say state ø. This is the state from which an entrant enters or to which an exitor exits.

**Symmetry and anonymity.** As does virtually all the applied literature, we restrict attention to models with symmetric and anonymous primitives and focus on equilibria in which values (i.e., payoffs) and policies (i.e., strategies) are also symmetric and anonymous.

We say that a set of functions  $\{f_i(\cdot)\}_{i=1}^n$  is symmetric if

$$f_i(\omega_i, \omega_{-i}) = f_j(\omega_i, \omega_{-i})$$

for all indices i and j. Hence, there is no need to index the functions by i and j and we simply write  $f(\cdot)$  from hereon. The function  $f(\cdot)$  is anonymous (a.k.a., exchangeable) if

$$f(\omega_i, \omega_{-i}) = f(\omega_i, \omega_{\pi(-i)})$$

for all permutations  $\pi(-i)$  of the indices in -i.

The importance of the symmetry and anonymity assumptions for applied work is that they insure that all relevant differences among firms are encoded in the firms' states (including possible differences in "types" of firms). This allows us to connect differences in equilibrium responses among firms to differences in the characteristics of those firms.<sup>6</sup>

Symmetry and anonymity reduce the size of the state space. Without symmetry and anonymity S grows as an exponential in  $\bar{n}$ , i.e., its cardinality (#S) is greater than  $\bar{\omega}^{\bar{n}}$ . Once we restrict attention to symmetric and anonymous games the relevant state space grows as a polynomial in  $\bar{n}$ . That is, symmetry and exchangeability enables us to restrict the state space to

$$S^{\circ} \equiv \{(\omega_1, \omega_2, \dots, \omega_n) : \omega_1 \in \Omega, \omega_2 \le \omega_3 \le \dots \le \omega_n, n \le \bar{n}\} \subset S,$$

where  $\subset$  is notation for a proper subset.

<sup>&</sup>lt;sup>6</sup>In general one could be concerned with asymmetric and nonanonymous equilibria even in symmetric and anonymous games. The burden on applied work then becomes much greater. On the other hand once we define the state variables of the problem the symmetry and anonymity assumptions should, at least in principle, be testable. Of course if the test rejects we are left with the problem of distinguishing whether the rejection is due to a misspecification of the state variables or the inappropriateness of the equilibrium assumptions.

Given symmetry and anonymity we can characterize industry structures more compactly as a set of "counting measures", i.e., as vectors of positive integers in which each integer represents the number of firms in state  $\omega \in \Omega$ . Formally, if  $\mathcal{Z}^+$  is the set of non-negative integers then the state space can also be written as

$$S^{\circ} = \{(\omega_i, s) : \omega_i \in \Omega, s = (s_1, \dots s_{\bar{\omega}}), s_{\omega} \in \mathcal{Z}^+, \sum_{\omega \in \Omega} s_{\omega} \le \bar{n}\}.$$

Typically a program used to compute an equilibrium stores the detail it needs on the different industry structures in locations which are one-to-one with the elements of  $S^{\circ}$ . However the program has to compute continuation values by integrating over possible future states from given initial states. Since the same future industry structure can arise from two different sets of firm-specific outcomes, it is convenient to compute continuation values by summing the probabilities of all possible firm-specific outcomes from that state.<sup>7</sup> As a result when we compute the state-to-state transitions we use the less concise representation of the state space, and for that reason we will use it here also.

**Timing.** We assume that all decisions (pricing, investment, entry, and exit) are made as a function of the state of the industry at the beginning of the period. For simplicity we also assume that incumbents firms that decide to exit compete in the product market before they exit.<sup>8</sup> During the period the outcomes of the investment process are realized, potential entrants enter, and incumbent firms exit. At the end of the period the state of the industry is updated with the results of these processes.

Throughout we use  $\omega$  to denote the state of the industry at the beginning of the period and  $\omega'$  to denote its state at the end of the period (after the state-to-state transitions are realized). Firms observe the state at the beginning of the period as well as the outcomes of the entry, exit, and investment decisions during the period.

**Product market competition.** Let  $\pi(\omega_i, \omega_{-i})$  denote the current profit of incumbent firm *i* from product market competition when the industry is in state  $\omega$ . Until we get to our discussion of extensions to the basic framework (our Section 7) we treat this per-period profit function as a primitive of the dynamic stochastic game. Up to "regularity conditions" (which we come back to when they are needed, in our Section 3) the EP framework leaves

<sup>&</sup>lt;sup>7</sup>The alternative here would be to explicitly formulate the multinomial formula for the probabilities of vectors of outcomes for a firm's competitors. This would undoubtedly be the efficient way of computing continuation values if the number of firms was large enough, but for the number of firms typical for current applications the savings in the number of elements in the summand does not compensate for the computational burden of setting up the multinomial probabilities.

<sup>&</sup>lt;sup>8</sup>If exitors did not compete in the product market in the period they exit, then we would have to use expected profits as our profit measure where the expectation takes account of the exit decisions of a firm's competitors. Similarly, if we were to amend our earlier assumption and allow entrants to compete in the product market in the period they enter, then we would have to compute expected profits where the expectation takes account of entry decisions.

the specification the profit function to the analyst. This typically requires the researcher to specify a demand system, a cost function, and an equilibrium assumption.

The existing literature has used a variety of different specifications including: price competition with differentiated products where a firm's state indexes the quality of its product (Pakes & McGuire 1994); price competition with differentiated products where a firm's state represents the share of consumers who are aware of the firm's product (Doraszelski & Markovich 2006); and quantity or price competition with homogenous products where a firm's state determines its cost schedule through either its capacity or its capital stock (Berry & Pakes 1993, Gowrisankaran 1999, Besanko & Doraszelski 2004). It is important to keep in mind that the actual dynamics of the equilibria we compute (and no doubt of data on actual industries) depend on the properties of the per-period profit function in complex ways that we are only beginning to understand (see, e.g., Besanko & Doraszelski (2004) and Doraszelski & Markovich (2006)).

Below we use the Pakes & McGuire (1994) quality ladder model to illustrate the modeling of the product market. Incumbent firm *i* produces a product of quality  $\omega_i$ . The consumers behave according to a standard discrete choice model (and so either purchase one product during the period or use all their money on the outside alternative). Consumer *k*'s utility from choosing good *i* is given by  $g(\omega_i) - p_i + \epsilon_{ik}$ , where  $g(\cdot)$  is an increasing bounded function of  $\omega_i$ .<sup>9</sup> Here  $g(\omega_i)$  provides the mean (across consumers) of the utility from choice *i* and  $\epsilon_{i,k}$  represents taste differences among consumers. The no-purchase alternative or outside good is denoted as product 0 and has utility  $\epsilon_{0k}$ . The consumer specific terms  $\epsilon_{0k}, \epsilon_{1k}, \ldots, \epsilon_{nk}$  are assumed to be independently and identically extreme value distributed across products and consumers. As is well known since the work of McFadden (1974) this results in the demands

$$q_i(p_1,\ldots,p_n;\omega) = M \frac{\exp\left(g(\omega_i) - p_i\right)}{1 + \sum_{j=1}^n \exp\left(g(\omega_j) - p_j\right)},$$

where M > 0 is the size of the market (the measure of consumers).

Assuming a constant marginal cost of  $c \ge 0$ , the price of firm *i* in state  $\omega$  solves

$$\max_{p_i\geq 0} q_i(p_1,\ldots,p_n;\omega) \left(p_i-c\right).$$

The first-order conditions are given by

$$0 = \frac{\partial}{\partial p_i} q_i(p_1, \dots, p_n; \omega) \left( p_i - c \right) + q_i(p_1, \dots, p_n; \omega),$$

$$g(\omega_i) = \frac{3\omega_i - 4}{12 + \ln\left(2 - \exp\left(16 - 3\omega_i\right)\right)} \quad \text{if} \quad \omega_i \ge 5.$$

<sup>&</sup>lt;sup>9</sup>Introducing the  $g(\cdot)$  function is an easy way to insure profits are bounded. Specifically, Pakes & McGuire (1994) set

and have a unique solution (Caplin & Nalebuff 1991). The Nash equilibrium prices are found by numerically solving the system of first-order conditions. Profits are then computed as

$$\pi(\omega_i, \omega_{-i}) = q_i(p_1(\omega), \dots, p_n(\omega); \omega)(p_i(\omega) - c)$$

State-to-state transitions. The probability that the industry transits from today's state  $\omega$  to tomorrow's state  $\omega'$  is determined jointly by the investment decisions of the incumbent firms that remain in the industry, the decisions of potential entrants that enter the industry, and industry-wide shocks that represent movements in the demand or cost conditions facing the industry. Again, subject to mild regularity conditions, the primitives governing these transitions can vary with what seems appropriate for the study at hand, though as EP emphasize, to get their upper bound to  $\Omega$  they use the assumption that a firm's state does not increase when it does not invest.

We come back to the relationship between the specification for the transition probabilities and computational and modeling issues below, but for now we suffice with the simple case where the transition probability for firm *i*'s state depends only on its own investment and current state (and not on the investments or states of its competitors). Then what the analyst has to specify is a family of probability distributions for  $\omega'_i$ ; one for each possible  $\omega_i$ , level of investment  $x_i$ , and industry-wide shock  $\eta$  (these represent factor prices or demand conditions that cause correlation between the outcomes of different firms in the same industry). Our notation for this family of distributions is

$$\mathcal{P}_{\omega'} \equiv \{ p(\cdot \mid \omega_i, x_i, \eta) : \omega_i \in \Omega, x_i \in \mathcal{R}^+, \eta \in \Upsilon \},\$$

where  $\Upsilon$  is the set of values for  $\eta$ . We assume that  $\mathcal{P}_{\omega'}$  is stochastically increasing (in the first-order stochastic dominance sense) in the natural order of both  $\omega_i$  and  $x_i$ .

Industry-wide demand shocks. Pakes & McGuire (1994), hereafter PM, among others assume that incumbent firm i's state conditional on remaining in the industry evolves according to

$$\omega_i' = \omega_i + \nu_i - \eta.$$

where  $\nu_i$  represents the outcome of the firm's investment process, and hence has a distribution which is stochastically increasing in  $x_i$  and independent across firms, and  $\eta$  represents the improvements in the outside alternative. These improvements in the mean utility from not purchasing any of the goods marketed occur with exogenously fixed probabilities and affect the desirability of all goods produced in the industry. We note that for any given application one might need to also allow for intertemporal changes in costs, or for a second industry wide random variable which moves all firms' marginal costs. Then the level of costs would also become a state variable in the problem. In Pakes & McGuire (1994)  $\mathcal{P}_{\omega'}$  is constructed from the family

$$\mathcal{P}_{\nu} \equiv \{ p(\cdot | x_i) : x_i \in \mathcal{R}^+ \}$$

and the probabilities  $\{p(\eta)\}_{\eta\in\Upsilon}$ . In the simplest case  $\nu\in\{0,1\}$  and  $\mathcal{P}_{\nu}$  is

$$\Pr(\nu|x_i) = \begin{cases} \frac{\alpha x_i}{1+\alpha x_i} & \text{if } \nu = 1\\ \frac{1}{1+\alpha x_i} & \text{if } \nu = 0, \end{cases}$$
(1)

where  $\alpha > 0$  parameterizes the effectiveness of investment. This family is stochastically increasing in  $x_i$  as required. To complete the specification we need to specify the distribution of  $\eta$  and the simplest case here is to assume  $\eta \in \{0, 1\}$ , and  $\eta = 1$  with probability  $\delta$ . Note that this simplistic specification can be made somewhat more realistic by noting that the length of the decision making interval can be made small relative to, say, the interval at which data becomes available. For example if the data is annual and we assume the firm makes decisions monthly, the distribution for  $\omega'_i$  given  $\omega_i$  is a twelve fold convolution of the increments modeled above.

To complete the specification we note that a firm which exits does not invest and transits to the special state  $\phi$  with probability one. A potential entrant who enters has the law of motion  $\omega'_i = \omega^e + \nu_i - \eta$ , where  $\omega^e \in \Omega$  is an (exogenously given) initial state and  $\nu_i$  is distributed with the same probabilities as are given above.

Firm-specific depreciation shocks. Note that in the model above the firm specific outcome (the  $\nu$ ) can only increase the firm's state, as may be appropriate for a research and development or perhaps an advertising game. Besanko & Doraszelski (2004) among others assume that firm *i*'s state evolves according to

$$\omega_i' = \omega_i + \nu_i - \eta_i,$$

where the realization of  $\eta_i$  is firm specific (in contrast to the industry-wide  $\eta$  in the previous specification). The probabilities for  $\nu_i$  and  $\eta_i$  are derived just as in the model with industry shocks given above. That is the distribution of  $\nu_i$  is given by a family of distributions which are increasing in  $x_i$ , and the distribution of  $\eta_i$  is determined by an exogenous "breakdown" probability.

Depending on what is encoded in a firm's state, more elaborate specifications of firmspecific depreciation shocks may be called for. In the context of physical capital, in particular, we often think of depreciation as being proportional to the stock. That is, absent investment, the capital stock tomorrow is with certainty a particular fraction of the capital stock today. Since fractional values of the capital stock are not allowed, we can not reproduce deterministic proportional decay exactly, however there are a number of ways of modeling decay which are similar in certain respects. One is to assume that firm i owns  $\omega_i$  machines and that each machine has a probability of  $\delta$  per period of breaking down independent of other machines. Then firm *i* will own anywhere from 0 to  $\omega_i$  machines next period, so that  $\omega'_i$  is binomially distributed with support  $\{0, 1, \ldots, \omega_i\}$  (before investment is taken into account). An alternative is to choose integer values that are on either side of the desired fraction and assign probabilities of transition to them which generate the right expected decay (see, e.g., Benkard 2004).

Note that since a positive outcome from the research of one firm will take profits away from its competitors, if there were no source of correlation between the transitions of firms within an industry, as in the model with just firm-specific shocks, the model would predict a negative correlation between the profitability of different firms. In fact, the profits of firms within an industry are often (though not always) positively correlated, i.e., industrywide shocks to demand or cost conditions often have larger impacts than the outcomes of rivalrous investment decisions. By allowing for both firm-specific and industry-wide shocks we let the specific application determine the profit correlations. For notational simplicity we focus on the industry-wide shock in what follows.

An incumbent's problem. Suppose that the industry is in state  $\omega$ . Incumbent firm *i* solves an intertemporal maximization problem to reach its exit and investment decisions. Let  $V(\omega_i, \omega_{-i}, \phi)$  denote the expected net present value of all future cash flows to incumbent firm *i*, when the industry structure is given by  $\omega$  and the firm has drawn a scrap value  $\phi$ .  $V(\omega_i, \omega_{-i}, \phi)$  is defined recursively by the solution to the following Bellman equation

$$V(\omega_i, \omega_{-i}, \phi) = \pi(\omega_i, \omega_{-i}) + \max\left\{\phi, \max_{x_i} - x_i + \beta E\left[V(\omega'_i, \omega'_{-i}, \phi') \mid \omega_i, \omega_{-i}, x_i\right]\right\}, \quad (2)$$

where  $\beta$  is the common discount factor, and where it is understood that E is the expectation operator which integrates out over the probability distribution of possible next period values for the firm's own state, its competitors' states, and its scrap value conditional on the current state and the firm's choice of investment. The firm's value consists of the current profit from product market competition  $(\pi(\omega_i, \omega_{-i}))$  plus the larger of the return to exiting the industry ( $\phi$ ) and the continuation value for remaining in the industry. The continuation value consists of the discounted (by  $\beta$ ) expectation of next period's value minus the cost of the investment incurred in the interim.

To take the expectation required to determine its continuation value the firm must have a perception of the likely future states of its competitors conditional on the different possible outcomes of the industry-wide shock,  $\eta$ . We let the firm's perceived probability of the next period value of its competitors' state  $(\omega'_{-i})$  conditional on  $\eta$  be  $q(\omega'_{-i}|\omega_i, \omega_{-i}, \eta)$ . In equilibrium these perceptions will have to satisfy certain conditions, and we will come back to these conditions below, but for now assume only that  $q(\cdot)$  is the distribution used by the firm. Then we can write

$$E\left[V(\omega_i', \omega_{-i}', \phi') \mid \omega_i, \omega_{-i}, x\right] = \sum_{\nu} W(\nu \mid \omega_i, \omega_{-i}) p\left(\nu \mid x_i\right)$$
(3)

where

$$W(\nu \mid \omega_i, \omega_{-i}) \equiv \sum_{\omega'_{-i}, \eta} \int_{\phi'} V(\omega_i + \nu - \eta, \omega'_{-i}, \phi') dF(\phi') q(\omega'_{-i} \mid \omega_i, \omega_{-i}, \eta) p(\eta).$$

 $W(\nu \mid \omega_i, \omega_{-i})$  is the expected discounted value of the firm conditional on the outcome of its investment being  $\nu$ . To obtain it we had to integrate out over the distribution of possible outcomes for; the firms' competitors, the firm's own future scrap value, and the outside alternative.

Note that neither  $W(\cdot)$  nor  $q(\cdot)$  are primitives of the problem. As a result different computational algorithms construct them in different ways, a point we will be much more explicit about in the computational section below.

The important point to note now is that if we (or the agent) were to know  $\{W(\cdot)\}$ , optimal behavior could be determined from a simple single agent optimizing problem. To do so substitute equation (3) into (2) and obtain a Kuhn-Tucker condition

$$x_i \left(\beta \sum_{\nu} W(\nu \mid \omega_i, \omega_{-i}) \frac{\partial p (\nu \mid x_i)}{\partial x_i} - 1\right) = 0 \quad \land \quad x_i \ge 0 \tag{4}$$

for the optimal  $x_i$ , say  $x(\omega_i, \omega_{-i})$ . Below we provide conditions where the solution to this equation is unique, as it will be, for example, if we use the  $\mathcal{P}_{\nu}$  in equation (1). In this case

$$x(\omega_i, \omega_{-i}) = \max\left\{0, \frac{-1 + \sqrt{\beta \alpha \left(W(1|\omega_i, \omega_{-i}) - W(0|\omega_i, \omega_{-i})\right)}}{\alpha}\right\}$$
(5)

if  $W(1|\omega_i, \omega_{-i}) \ge W(0|\omega_i, \omega_{-i})$ , and  $x(\omega_i, \omega_{-i}) = 0$  otherwise.

Next we substitute  $x(\omega_i, \omega_{-i})$  into equation (2) and determine whether the firm continues. Letting  $\chi(\omega_i, \omega_{-i}, \phi)$  be the indicator function which takes the value of one if the firm continues and zero otherwise we have

$$\chi(\omega_i, \omega_{-i}, \phi) = \arg\max_{\chi \in \{0,1\}} (1-\chi)\phi + \chi \left(\beta \sum_{\nu} W(\nu \mid \omega_i, \omega_{-i}) \ p \ (\nu \mid x(\omega_i, \omega_{-i})) - x(\omega_i, \omega_{-i})\right).$$
(6)

The probability of drawing a  $\phi$  such that  $\chi(\omega_i, \omega_{-i}, \phi) = 1$  determines the probability of the firm remaining active or

$$r(\omega_i, \omega_{-i}) = F\left(\beta \sum_{\nu} W(\nu \mid \omega_i, \omega_{-i}) \ p \ (\nu \mid x(\omega_i, \omega_{-i})) - x(\omega_i, \omega_{-i})\right).$$
(7)

An entrant's problem. A potential entrant who chooses to enter must pay an entry fee and then becomes an incumbent firm in the next period. The entrant's decision on whether to enter is analogous to the incumbent's decision on whether to exit; i.e. it compares its continuation value to the cost of entry. Since the competitors of the potential entrant are given by the entire vector of active firms' states, or  $\omega$ , the value of potential entrants *i* is given by

$$V^{e}(\omega, \phi^{e}) = \max\left\{0, \ \max_{x_{i}^{e}} -\phi^{e} - x_{i}^{e} + \beta \sum_{\nu} W^{e}(\nu \mid \omega) \ p \ (\nu \mid x_{i}^{e})\right\},\tag{8}$$

where  $W^{e}(\cdot)$  is defined analogously to  $W(\cdot)$  for an incumbent firm.

The potential entrant solves for its optimal investment  $x^e(\omega)$  in a manner analogous to an incumbent firm, and then enters if and only if it is profitable to do so. Letting  $\chi^e(\omega, \phi^e)$  be the indicator which takes the value of one if the potential entrant enters and zero otherwise

$$\chi^{e}(\omega,\phi^{e}) = \arg\max_{\chi\in\{0,1\}} \chi\left(-\phi^{e} - x^{e}(\omega) + \beta\sum_{\nu} W^{e}(\nu\mid\omega) \ p\left(\nu\mid x^{e}(\omega)\right)\right). \tag{9}$$

So the probability that a potential entrant enters is

$$r^{e}(\omega) = F^{e}\left(-x^{e}(\omega) + \beta \sum_{\nu} W^{e}(\nu \mid \omega) \ p \ (\nu \mid x^{e}(\omega))\right).$$
(10)

## 3 Equilibrium

This section begins by providing a definition of equilibrium. We then provide sufficient conditions for the existence of an equilibrium that satisfy these definitions (section 3.1). We do not know conditions which insure uniqueness of the equilibrium but we do have results which characterize any of the possible equilibria (section 3.2), and we use those characteristics in our computational algorithms. To illustrate the issues underlying the multiplicity of equilibria, we consider a number of examples of multiple equilibria (section 3.3). Our discussion is mostly based on EP, Doraszelski & Satterthwaite (2003), and Besanko, Doraszelski, Kryukov & Satterthwaite (2004).

We consider Markov perfect equilibria (MPE). A Markov prefect equilibrium insures that at each  $\omega \in S^{\circ}$  each incumbent firm and each potential entrant

- chooses optimal policies given its perceptions on likely future industry structures, and
- those perceptions are consistent with the behavior of each agent's competitors.

One way of checking that these conditions are satisfied is to show that the equilibrium generates a set of value functions and policies, one for each potential entrant and each incumbent at each  $\omega \in S^{\circ}$ , such that

- given the policies, the value functions satisfies the Bellman equations in (2) and (8) for incumbents and potential entrants respectively, and
- given the value functions, the policies for investment and exit satisfy the optimality conditions in equations (4) and (6) for incumbents (augmented to include a check of whether the extreme point for investment is a global maximum in cases where functional forms do not guarantee that), and the analogous equations for potential entrants.

Note that since the horizon is infinite and the influence of past play is captured in the current state, there is a one-to-one correspondence between subgames and states. Hence, any Markov perfect equilibrium is subgame perfect. Further since a best reply to Markovian strategies is a Markovian strategy, a Markov perfect equilibrium remains a subgame perfect equilibrium even if more general strategies are considered.

#### 3.1 Existence

The extent to which the applied literature has faced the existence question is by testing whether the values and policies they computed in fact satisfy the equilibrium conditions up to some error. The fact that we allow for an error, makes this condition imperfect.<sup>10</sup> However since computers can only compute fixed points to machine precision, there is a sense in which we cannot determine whether numerical results satisfy any stronger notion of equilibrium than this. Indeed the reader that is only interested in computational issues should be able to skip this section and have no trouble with the rest of the paper.

This section asks a different question then that asked by the applied literature. In particular we ask for conditions that insure that there is an equilibrium in which the conditions given above hold exactly. The importance of an existence proof to applied work is that it provides conditions that ensure we are searching for an approximation to something that actually exists. The weakness is that the conditions we give for existence are sufficient but not necessary, so there may well be equilibria in situations which do not abide by the conditions in the proof.

We look for an equilibrium in pure strategies (computing mixed strategy equilibria increases the computational complexity of the algorithm markedly). The proof of existence requires a continuous mapping from policies into themselves. One reason for adding random scrap values/setup costs to the EP framework is that they allow us to treat the continuous exit and entry probabilities as the policies (in contrast to the discrete entry and exit deci-

<sup>&</sup>lt;sup>10</sup>The theoretical literature on game theory does use the notion of  $\epsilon$ -equilibrium, but it is not automatically satisfied by ensuring the equilibrium conditions up to a sufficiently small error. I.e. in an  $\epsilon$ -equilibrium a player's strategy brings the player within  $\epsilon$  of his best possible payoff, assuming that payoffs can be computed exactly. In the notion used in computation both payoffs and strategies are computed with error, and the connection between the computed strategies and an  $\epsilon$ -equilibria has not been shown. In addition it is well known that an  $\epsilon$ -equilibrium need not be close to an exact equilibrium.

sions; see Doraszelski & Satterthwaite (2003)). Continuity of this best-reply mapping can then be established under standard continuity assumptions on the transition functions.

Given continuity, one way to obtain a pure strategy equilibrium is to ensure that a firm's best reply is always unique.<sup>11</sup> As long as the densities of the scrap values and setup costs are continuous the entry and exit best responses are unique (up to a set of measure zero which is sufficient for the argument in Doraszelski & Satterthwaite (2003)). Turning to investment decisions, any family of distribution functions (our  $\mathcal{P}_{\nu}$ ) which guarantee that, for any given distribution of the outcomes from the actions of the firm's competitors, the continuation value for the firm is concave in its investment choice will guarantee a unique best reply. This is true of the simple family used above (which is taken from EP). Doraszelski & Satterthwaite (2003) go further. They define a class of transition functions that are called unique investment choice (UIC) admissible and prove that if the transition function is UIC admissible, then a firm's investment decision is indeed uniquely determined. The UIC class generalizes the EP example by allowing transitions to more than immediately adjacent states, and also allows for the construction of median preserving spreads.

#### 3.2 Characterization

EP provide conditions which insure that in equilibrium there is an  $\bar{n}$  and a  $\bar{\omega}$  such that if we start at an initial industry structure in which

- there are no more than  $\bar{n}$  firms active, and
- each active firm has an  $\omega_i \in \Omega$

then, with probability one,

- there will never be more than  $\bar{n}$  firms active, and
- we will never observe an active firm with an  $\omega_i \notin \Omega$ .

As a result  $\#S^{\circ}$  is finite, and equilibrium values and policies are computable.<sup>12</sup>

<sup>&</sup>lt;sup>11</sup>Escobar (2006) proposes a similar condition to ensure the existence of an equilibrium in pure strategies in the context of more general dynamic stochastic games with a countable state space and a continuum of actions. While Escobar's (2006) condition applies to games with continuous actions other than the investment decisions in the Ericson & Pakes (1995) model, there is no systematic treatment of incomplete information as a means to purify the discrete entry/exit decisions.

<sup>&</sup>lt;sup>12</sup>The proof restricts the profit function so that profits will be driven arbitrarily close to zero for n large enough regardless of the location of the incumbents, and that profits are bounded from above. It also assumes that the number of potential entrants in any given period is finite, that entry fees are bounded away from zero, and that if there is no investment, the firm's state cannot advance (i.e.,  $Pr(\nu = 1|x_i = 0) = 0$ ). The argument for the upper bound to  $\omega$  is roughly as follows. The value function is bounded if profits are. The incentive to invest is the increment to the value function from increasing  $\omega$ . Since the value function is bounded, for any given value of  $\omega_{-i}$  that increment can be made arbitrarily close to zero. So eventually the increment has a value less than than the cost of investment and investment shuts down. Once investment shuts down the firm's state can not increase. The argument is completed by showing that their is only a finite number of possible  $\omega_{-i}$ . The argument for the upper bound to n follows from the fact that profits can be driven arbitrarily small by increasing n, that there are finite number of potential entrants in each period, and that entry costs are bounded away from zero.

They also show that any equilibrium defines a time homogenous Markov process, for  $\omega'$  given  $\omega$ , defined by the family of distributions (or by the Markov transition kernel)

$$\mathcal{Q}_{\omega'} \equiv \{ Q(\cdot \mid \omega) : \omega \in S^{\circ} \}$$

Moreover every process generated by the equilibrium conditions satisfies a fixed point to an operator which takes the set of time homogenous Markov processes on  $S^{\circ}$  into itself, and one can view the procedures for computing the equilibria that we discuss below as algorithms for finding such a fixed point.<sup>13</sup>

Since  $\#S^{\circ}$  is finite, the process generating industry structures in the EP model is a finite state Markov chain. Consequently standard results in Markov chain theory (see, e.g., Freedman (1983)) insure that there is at least one recurrent class of states, say  $R \subset S^{\circ}$ , and with probability one each sample path (each sequence of  $\{\omega_t\}_t$ ) will enter one of these recurrent classes in finite time. Once in the recurrent class the sample path will stay within it forever. Consequently the points in the recurrent class are visited infinitely often. EP go further and provide conditions which insure that there is only one such class, so the process is ergodic, though ergodicity is not necessary for the discussion that follows and it does require substantive conditions (for an example with more than one recurrent class see Besanko & Doraszelski (2004)).

In IO problems recurrent classes of points are typically much smaller then the state space or  $S^{\circ}$ . Consider our quality competition example. As we increase market size (our M) we will typically generate equilibria with a higher  $\bar{n}$ , and  $\#S^{\circ}$  will increase polynomially in  $\bar{n}$ . However as M increases the economics of the primitives will typically insure that, provided we start with a relatively large n at high enough states, we will never observe an industry structure with a small number of firms at low states, as entry will occur whenever the number and states of incumbents fall below certain thresholds. Thus #R can grow at a much slower rate that  $\#S^{\circ}$ . Indeed, at least in principle it need not grow in M at all (in our quality model it seems to grow less than linearly in M, see the example in Pakes & McGuire (2001)). Similarly if we allow two characteristics of products, say mpg and engine size of vehicles, then in equilibrium we may never observe vehicles with either both mpg and engine size very large (as it would be too expensive to produce) or both engines and mpg very small (as no one would buy them).

Recall that all subgames initiated from the recurrent class will stay within that class forever. So if we were trying to analyze likely future outcomes in an industry and are willing to assume that (or could check whether) the current industry structure were in the

<sup>&</sup>lt;sup>13</sup>To see this assume  $\mathcal{Q}_{\omega'}$  is generated by our equilibrium conditions. Take any  $\omega$  and use  $Q(\cdot | \omega) \in \mathcal{Q}_{\omega'}$  to form the marginal distribution of the likely locations of each firm's competitors in the next period. Substitute this for  $\{q(\cdot)\}$  in constructing the expected discounted value of future net cash flow conditional on investment outcomes, or the  $\{W(\cdot)\}$ , in equation (3). Then form optimal policies for all agents just as described above. In equilibrium these policies will generate an objective distribution of outcomes which coincide with the  $Q(\cdot|\omega) \in \mathcal{Q}_{\omega'}$  we started with.

recurrent class, then all we would require is information on policies on the recurrent class. Therefore a way to reduce the computational burden in applied problems is to design an algorithm which computes values and policies only on the recurrent class (see Pakes & McGuire (2001)).

Before leaving this section we want to emphasize that the nature of states in R, the transitions among those states, and the transitions into R all depend on the primitives of the problem and can vary greatly with different parameterizations for those primitives. As a result there is little in the way of characterizations of behavior that are directly relevant for policy or descriptive work that comes out of the general framework. Rather the goal is to have a framework which is flexible enough to generate a broad range of results and then let knowledge of the relevant institutions pick out the implications that seem relevant for the problems at hand.

### 3.3 Multiplicity

Doraszelski & Satterthwaite (2003) provide three examples that show that there need not be a unique equilibrium that is symmetric and anonymous. Multiplicity may arise from three sources:

- investment decisions;
- entry/exit decisions; and
- product market competition.

We now discuss each of these in turn, providing simple examples of the first two. Note that both the examples only consider equilibria that are symmetric and anonymous in the sense we defined earlier. Note also that the examples are all in the context of simple models where we can break out the static profit function and analyze it without considering investment, entry and exit decisions. Once this simplification is left behind, much richer examples of multiple equilibria can be generated, see Besanko et al. (2004).

**Investment decisions.** Consider a model with  $\bar{n} = 2$  firms and neither entry nor exit. Firm *i*'s capacity in state  $\omega_i \in \{1, 2, ..., 10\}$  is  $5(\omega_i - 1)$ . The state-to-state transitions are as in our example with a firm-specific depreciation shock with parameters  $\alpha = 2.375$  and  $\delta = 0.03$ . The discount factor is  $\beta = \frac{20}{21}$ . Products are undifferentiated and firms compete in prices subject to capacity constraints. There are M = 40 identical consumers with unit demand and reservation price v = 1. The equilibrium of this Bertrand-Edgeworth product market game is unique and symmetric (see Chapter 2 of Ghemawat 1997).

Figure 1 depicts two symmetric equilibria of the dynamic stochastic game. In both equilibria investment activity is greatest in states on or near the diagonal of the state space. That is, firms with equal or similar capacities are engaged in a "race" to become the



Figure 1: Two equilibria. Source: Doraszelski & Satterthwaite (2003).

industry leader. The difference in investment activity is greatest in state (5,5) where both firms invest 1.90 in the first equilibrium compared to 1.03 in the second one. Investment activity also differs considerably in states (1,6) and (6,1): in the first (second) equilibrium the smaller firm invests 2.24 (3.92) and the larger firm invests 1.57 (1.46). That is, in the second equilibrium, a firm that has fallen behind to state 1 strives to catch up and the industry leader (i.e., the firm in state 6) to some extent accommodates the laggard by reducing its investment.

Note that multiplicity rests on the dynamic nature of the game. Because product market competition takes place before investment decisions are carried out, a firm has no incentive to invest if  $\beta = 0$ . Hence, multiple equilibria cannot possibly arise in the static version of the game.

Entry/exit decisions. Consider a model with at most  $\bar{n} = 2$  firms,  $\bar{\omega} = 2$  active states plus one inactive state per firm, and the profits from product market competition given by

$$\pi(\omega_1) = 1, \quad \pi(\omega_1, \omega_2) = 0.$$

That is, monopoly profits are one and duopoly profits are zero. While entry is prohibited, exit is permissible with a scrap value that is uniformly distributed on [14, 16]. The discount factor is again  $\beta = \frac{20}{21}$ . Finally, a firm cannot transit between its active states (in our example with a firm-specific depreciation shock this corresponds to  $\alpha = 0$  and  $\delta = 0$ ). Notice that the parameters are chosen such that each firm would want to exit if its rival stays, i.e., exiting is more attractive than being a duopolist, but each firm would want to stay if its rival exits, i.e., being a monopolist is more attractive than exiting. This gives rise to a "war of attrition".

One equilibrium has both firms play the same strategies in states (1, 1), (1, 2), (2, 1), and (2, 2). The left panel of Table 1 summarize this equilibrium. Another equilibrium is illustrated in the right panels of Table 1. In states (1, 1) and (2, 2) both firms make the same exit decisions as before. In state (1, 2), however, firm 1 exits with high probability

$r(\omega_1,\omega_2)$	$\omega_2 = 1$	$\omega_2 = 2$	_	$r(\omega_1,\omega_2)$	$\omega_2 = 1$	$\omega_2 = 2$
$\omega_1 = 1$	0.8549	0.8549		$\omega_1 = 1$	0.8549	0.1542
$\omega_1 = 2$	0.8549	0.8549		$\omega_1 = 2$	1	0.8549

Table 1: Two more equilibria. Source: Doraszelski & Satterthwaite (2003).

and firm 2 stays for sure whereas in state (2, 1) firm 1 stays for sure and firm 2 exits with high probability. Note that the two equilibria differ starkly from each other: In the first equilibrium, a duopolistic industry may over time turn into either a monopolistic or an empty industry. In the second equilibrium, if the industry starts in states (1, 2) or (2, 1), then it always ends up as a monopoly.

**Product market competition.** To determine the profit function we need to model competition in the product market. Depending on the nature of demand and cost primitives, and the assumptions on strategies, it is well known that we can generate multiple solutions to this static game. The different profit functions will generate different investment, entry, and exit incentives and hence different equilibria with different distributions of sample paths for industry structures. This is yet another source of multiplicity.

We conclude with some further details on how multiplicity arises. The value of continued play to firm *i* is determined by its value function in states  $\omega'$  that may be reached from state  $\omega$ . Holding the value of continued play fixed, the strategic situation in state  $\omega$  is akin to a static game. If the reaction functions in this game intersect more than once, then multiple equilibria arise.

We say that the model satisfies stagewise uniqueness if the reactions functions intersect once irrespective of the value of continued play (Besanko et al. 2004), i.e., there is only one set of Nash equilibrium policies conditional on the value of continued play. Note that stagewise uniqueness requires more than just UIC admissibility. In our investment example above, the transition function satisfies UIC admissibility, so that given a firm's beliefs about the value of continued play, its investment decision is uniquely determined. The mere fact that each firm has a well-defined reaction function (as opposed to a correspondence) does not guarantee that these functions intersect only once.

If the model satisfies stagewise uniqueness and there are multiple equilibria, the multiplicity must arise from firms' expectations regarding the value of continued play. That is rational expectations are consistent with more than one value of continued play at a least one node of the game tree.

As Besanko et al. (2004) point out, a sufficient condition for uniqueness of equilibrium in a dynamic stochastic game is that the model satisfies stagewise uniqueness and that movements through the state space are unidirectional, i.e., a firm's state can only move in one direction (either up or down). Stagewise uniqueness precludes players' decisions from giving rise to multiple equilibria, and unidirectional movements preclude their expectations from doing so. Though applied problems with this property are rare, it is instructive to consider the reasoning behind this argument.

To do so let there be  $\bar{n} = 2$  firms and neither entry nor exit. Suppose that a firm can never move backward to a lower state (say because  $\delta = 0$  rules out depreciation shocks). Hence, once the industry reaches state  $(\bar{\omega}, \bar{\omega})$ , it remains there forever, so that the value of future play in state  $(\bar{\omega}, \bar{\omega})$  coincides with the value of being in this state. In conjunction with stagewise uniqueness, this uniquely determines the value of being in state  $(\bar{\omega}, \bar{\omega})$ . Next consider states  $(\bar{\omega} - 1, \bar{\omega})$  and  $(\bar{\omega}, \bar{\omega} - 1)$ . The value of future play in states  $(\bar{\omega} - 1, \bar{\omega})$  and  $(\bar{\omega}, \bar{\omega} - 1)$  depends on the value of being in state  $(\bar{\omega}, \bar{\omega})$ . Stagewise uniqueness ensures that firms' decisions in states  $(\bar{\omega} - 1, \bar{\omega})$  and  $(\bar{\omega}, \bar{\omega} - 1)$  as well as the value of being in these states are uniquely determined. Further, since the value of future play in state  $(\bar{\omega} - 1, \bar{\omega} - 1)$ depends on the value of being in states  $(\bar{\omega} - 1, \bar{\omega})$  and  $(\bar{\omega}, \bar{\omega} - 1)$ , firms' decisions and the value of being in states  $(\bar{\omega} - 1, \bar{\omega} - 1)$  are uniquely determined. Continuing to work backwards in this fashion establishes that the equilibrium is unique. Conversely, suppose a firm can never move forward to a higher state (say because  $\alpha = 0$  renders investment completely ineffective). A similar argument establishes uniqueness of equilibrium except that the argument is anchored in state (1, 1) rather than state  $(\bar{\omega}, \bar{\omega})^{14}$ 

# 4 Introduction to Computation

There have been a number of different algorithms proposed for computing equilibria to these games. One possibility is to look at the problem of computing equilibria as the problem of solving the system of nonlinear equations defined by our equilibrium conditions (see Chapter 5 of Judd (1998) for methods for solving nonlinear equations systems). However the applied problems we are interested in are typically too large to make direct application of Newton's method practical. Most applications therefore use some type of Gaussian method. The idea behind Gaussian methods is that it is harder to solve a large system of equations once than to solve smaller systems many times.

The Gaussian method we begin with is the backward solution method used in Pakes & McGuire (1994). This section begins by describing their algorithm, and concludes with a general discussion of two problems that can make it necessary to use more sophisticated algorithms in applied work. The first is the burden of computing equilibria, the second is the possibility of multiple equilibria. The later sections provide tools that ameliorate these two problems. First we discuss a number of computational methods that alleviate the computational burden. Among others we explain the continuous-time stochastic games

<sup>&</sup>lt;sup>14</sup>As an aside we note unidirectional movements through the state space not only make a dynamic stochastic game much more tractable (see, e.g., Cabral & Riordan 1994) but also greatly simplify computing an equilibrium (see, e.g., Judd, Schmedders & Yeltekin 2002).

due to Doraszelski & Judd (2004) and the stochastic approximation algorithm due to Pakes & McGuire (2001). Then we explain the homotopy methods used in Besanko et al. (2004) to address the multiplicity problem. We note however that neither problem is currently "solved", i.e., both computational burden and multiplicity can effectively limit the extent to which we can analyze particular applied problems.

### 4.1 Gaussian Methods

Like all backward solution techniques the Pakes & McGuire (1994) algorithm is an iterative procedure which starts each iteration with a set of numbers in memory, provides a rule which updates those numbers in the course of the iteration, and checks to see if the updated numbers satisfy a convergence criterion at the end of the iteration. If not the algorithm begins a new iteration.

The algorithm holds in memory an expected value function and policy for each incumbent firm and each potential entrant in each state  $\omega \in S^{\circ}$ . The expected value function is defined as

$$V(\omega_i, \omega_{-i}) \equiv \int_{\phi} V(\omega_i, \omega_{-i}, \phi) dF(\phi).$$
(11)

By keeping these expected values, instead of the values themselves in memory, we can ignore the continuous state variable  $\phi$  in storing and recalling objects from memory.

From the definition of the Bellman equation in (2) and of the  $W(\cdot)$  in (3) we have

$$V(\omega_i, \omega_{-i}) = \pi(\omega_i, \omega_{-i}) +$$
(12)

$$(1 - r(\omega_i, \omega_{-i}))\phi(\omega_i, \omega_{-i}) + r(\omega_i, \omega_{-i}) \left\{ -x(\omega_i, \omega_{-i}) + \beta \sum_{\nu} W(\nu \mid \omega_i, \omega_{-i}) p(\nu \mid x_i) \right\},\$$

where

$$W(\nu \mid \omega_i, \omega_{-i}) \equiv \sum_{\omega'_{-i}, \eta} V(\omega_i + \nu - \eta, \omega'_{-i}) q(\omega'_{-i} \mid \omega_i, \omega_{-i}, \eta) p(\eta),$$
(13)

and  $\phi(\omega_i, \omega_{-i})$  is the expectation of  $\phi$  conditional on exiting in state  $(\omega_i, \omega_{-i})$  (i.e on  $\chi(\omega_i, \omega_{-i}, \phi) = 0$ ).

An iteration circles through the states in some fixed order, updating *all* policies and expected values associated with each state every time it visits the state (any algorithm which updates all points at each iteration is a "synchronous" algorithm, we come back to this below).

**Iterations.** Iterative techniques begin with an initial set of values and policies for all states in  $S^{\circ}$ . There are a number of alternative ways to obtain the initial guesses. If no other information is available, reasonable initial guesses are often obtained by assuming the

agent gets the profits from its state forever, and discounting that sum. That is, if we let the superscript 0 to denote guesses, then

$$V^{0}(\omega_{i}, \omega_{-i}) = \frac{\pi(\omega_{i}, \omega_{-i})}{1 - \beta}, \quad r^{0}(\omega_{i}, \omega_{-i}) = 1, \quad x^{0}(\omega_{i}, \omega_{-i}) = 0,$$
$$V^{e,0}(\omega) = r^{e,0}(\omega) = x^{e,0}(\omega) = 0.$$

Of course if there is some information available that allows one to form initial conditions which are likely to be closer to the true values than this, e.g., if one has available the equilibrium values from a set of parameters that are close to the set that underlies the current calculation, one would use them.

Iteration l determines  $V^{l}(\cdot), x^{l}(\cdot), r^{l}(\cdot), v^{e,l}(\cdot), r^{e,l}(\cdot)$  from  $V^{l-1}(\cdot), x^{l-1}(\cdot), r^{l-1}(\cdot), V^{e,l-1}(\cdot), x^{e,l-1}(\cdot), r^{e,l-1}(\cdot), r^{e,l-1}(\cdot)$ . We now provide detail on how the values and policies are updated from one iteration to the next.

Updating policies and values. Fix  $\omega$ . For each incumbent and potential entrant we update the  $W(\cdot)$ 's by treating the continuation values in memory as the true expected values of being in a given state in the next period, and the policies in memory as the true policies of the firm's competitors. We then mimic the single agent problem discussed above for computing updated values and policies conditional on these  $\{W(\cdot)\}$  (or  $\{W^e(\cdot)\}$ ).

If we let l index iterations, the first step is constructing the  $\{W^l(\cdot)\}\$  from the policies and expected values in memory at iteration l-1. The calculation is simplified if we begin by constructing the probabilities of a firm's future state, counting exit as a possible outcome, conditional on the current states, the policies in memory and the outside shock. Recall that a firm either exits (to state  $\emptyset$ ) or continues with state ( $\omega + \nu - \eta$ ), so if we use the probabilities for  $\nu$  given x in equation (1)

$$p^{l-1}(\omega'_{i}|\omega_{i},\omega_{-i},\eta) = \begin{cases} 1 - r^{l-1}(\omega_{i},\omega_{-i}) & \text{if} & \omega'_{i} = \phi, \\ r^{l-1}(\omega_{i},\omega_{-i})p\left(\nu_{i}|x^{l-1}(\omega_{i},\omega_{-i})\right) & \text{if} & \omega'_{i} = \omega_{i} + 1 - \eta, \\ r^{l-1}(\omega_{i},\omega_{-i})\left(1 - p\left(\nu_{i}|x^{l-1}(\omega_{i},\omega_{-i})\right)\right) & \text{if} & \omega'_{i} = \omega_{i} - \eta. \end{cases}$$

Assuming that there is a fixed number, say  $\mathcal{E}$ , of potential entrants each period, we will also need the binomial probability that precisely  $\mathcal{E}'$  of the  $\mathcal{E}$  potential entrants enter conditional on iteration l - 1's policies which is constructed as

$$r^{e,l-1}(\mathcal{E}',\omega) \equiv \begin{pmatrix} \mathcal{E}\\ \mathcal{E}' \end{pmatrix} r^{e,l-1}(\omega)^{\mathcal{E}'} [1-r^{e,l-1}(\omega)]^{\mathcal{E}-\mathcal{E}'}.$$

Next take the  $p^{l-1}(\cdot)$ 's and the  $r^{e,l-1}(\cdot)$ 's and compute the firm's perceived probabilities of next periods values of its competitors state as

$$q^{l}(\omega_{-i}'|\omega_{i},\omega_{-i},\eta) = \prod_{j\neq i} p^{l-1}(\omega_{j}'|\omega_{j},\omega_{-j},\eta) \prod_{j=1}^{\mathcal{E}'} p(\nu_{j}|x^{e,l-1}(\omega)) \ r^{e,l-1}(\mathcal{E}',\omega),$$

where  $p(\nu_j|\cdot)$  is as defined in equation (1), and  $x^{e,l-1}(\cdot)$  is the entry policy at iteration l-1.

Substituting  $q^{l}(\cdot)$  into equation (13) and calling up expected values from memory we compute

$$W^{l}(\nu \mid \omega_{i}, \omega_{-i}) \equiv \sum_{\omega'_{-i}, \eta} V^{l-1}(\omega_{i} + \nu - \eta, \omega'_{-i}) q^{l}(\omega'_{-i} \mid \omega_{i}, \omega_{-i}, \eta) p(\eta).$$

We now update continuation values and policies conditional on the  $\{W^l(\cdot)\}\$  for a given  $\omega$ . For each incumbent *i* we solve the resultant single agent optimization problem for the optimal investment and exit decisions as follows

• Substitute  $W^{l}(\cdot)$  for the generic  $W(\cdot)$  in the Kuhn-Tucker equation (4) to obtain the optimal investment of the firm should it continue. In the simple case where  $\nu$  only takes on the values  $\{0, 1\}$ , the optimal x satisfies

$$x\left(-1+\beta[W^{l}(1;\omega_{i},\omega_{-i})-W^{l}(0;\omega_{i},\omega_{-i})]\frac{\partial p(1|x)}{\partial x}\right)=0, \text{ and } x\geq 0.$$

Depending on the chosen functional forms for the transition probabilities, the solution to this problem, which we label  $x^{l}(\omega_{i}, \omega_{-i})$ , may or may not have a closed-form solution. Equation (5) exhibits the solution in case the success probability  $p(1|x) = \frac{\alpha x}{1+\alpha x}$ . The second-order condition is satisfied automatically here, but is has to be explicitly checked if more general functional forms are used.

Substitute x<sup>l</sup>(ω<sub>i</sub>, ω<sub>-i</sub>), W<sup>l-1</sup>(ν; ω<sub>i</sub>, ω<sub>-i</sub>) into equation (6) to determine the exit policy, i.e. χ<sup>l</sup>(ω<sub>i</sub>, ω<sub>-i</sub>, φ), and use that policy to determine the probability that incumbent firm i will draw a φ that induces it to remain in the industry as

$$r^{l}(\omega_{i},\omega_{-i}) = F\left(-x^{l}(\omega_{i},\omega_{-i}) + \beta \sum_{\nu} W^{l}(\nu;\omega_{i},\omega_{-i})p(\nu|x^{l}(\omega_{i},\omega_{-i}))\right).$$

• Substitute  $W^{l}(\nu; \omega_{i}, \omega_{-i})$ ,  $x^{l}(\omega_{i}, \omega_{-i})$ , and  $r^{l}(\omega_{i}, \omega_{-i})$  into (12) and take the value of the resultant expression as the continuation value  $V^{l}(\omega_{i}, \omega_{-i})$ . Place  $V^{l}(\cdot), x^{l}(\cdot), r^{l}(\cdot)$  in memory. This completes the update for the incumbent firms at  $\omega$ .

Updating values and policies for potential entrants, i.e.,  $V^{e,l}(\omega)$ ,  $x^{e,l}(\omega)$ ,, and  $r^{e,l}(\omega)$ , is analogous except that we use  $W^{e,l-1}(\nu|\omega)$  to evaluate the expected discounted value of future net cash flows for the potential entrants, and the entrant bases its entry decision, and its investment should it enter, on the Bellman equation (8) instead of the incumbent's Bellman equation in (2). That completes our discussion of the updates that occur during an iteration. Stopping rule. The algorithm cycles through the state space until the changes in the values and policies from one iteration to the next are "small."<sup>15</sup> To quantify this, we require a measure of the distance between two sets of values and policies. Different distance measures or norms have been used in the literature, and as we shall see some of them make more sense for one algorithm than for another. For the Gaussian scheme described above, there is an argument for using the  $L_{\infty}$  or sup norm, as it is the strongest norm, and to use a norm which is close to "unit free", so we divide by  $||V^{l}(\omega)||$ . Therefore we take the relative difference in values at iteration l to be

$$\left\|\frac{V^{l} - V^{l-1}}{1 + |V^{l}|}\right\| = \max_{\omega \in S^{\circ}} \left|\frac{V^{l}(\omega) - V^{l-1}(\omega)}{1 + |V^{l}(\omega)|}\right|.$$

where the plus one is to avoid division by zero, and similarly for policies.<sup>16</sup>

There remains the issue of when to terminate the iterations. We would like a criterion that is based on the distance between the current iteration's policies and values and the true fixed point, and not just on the distance between subsequent iterates. Note that if the algorithm does converge (and there is no guarantee that it will), then convergence is linear as in all Gaussian schemes (Ortega & Rheinboldt 1970, p. 301). A sequence  $\{z^l\}_{l=0}^{\infty}$  is said to converge linearly to the limit  $z^{\infty}$  if and only if

$$\lim_{l \to \infty} \frac{\|z^{l+1} - z^{\infty}\|}{\|z^l - z^{\infty}\|} \le \theta < 1.$$

Assume temporarily that the first inequality can be strengthened to hold along the entire sequence of iterates, i.e.,

$$\|z^{l+1} - z^{\infty}\| \le \theta \|z^l - z^{\infty}\| \tag{14}$$

for all l. This is similar to the contraction property of a single-agent dynamic programming problem. Then the distance to the limit can be shown to be related to the distance between

<sup>&</sup>lt;sup>15</sup>If one does not fix  $\Omega$  and  $\bar{n}$  at the outset, but rather allows the algorithm to solve for them (as did Pakes & McGuire (1994)), the sequence of iterations must typically be done more than once. For a fixed  $\Omega$ we compute the fixed point with the constraint that there are never more that  $\tilde{n}$  firms active. Those values and policies are then used as initial conditions for a sequence of iterations that allows for a larger number of active firms. We repete this procedure until  $\tilde{n}$  is so large that there are no states at which a potential entrant wants to enter.  $\Omega$  is set as follows. First we compute values and policies for the monopoly problem. The value of  $\omega$  at which a monopolist would exit and the value of  $\omega$  at which the monopolist would stop investing become the lowest and highest values in  $\Omega$ . Though a value of  $\omega$  at which the monopolist would exit would induce exit no matter the market structure, it is possible for an oligopolist to still invest at an  $\omega$  at which a monopolist would not. Thus once one computes the equilibrium values one must check to see that no firm wants to invest at  $\bar{\omega}$ . If this is not the case  $\bar{\omega}$  is increased, and the computation is done again. For more detail see Pakes & McGuire (1994).

<sup>&</sup>lt;sup>16</sup>Note that in a single agent dynamic programming problem we would define a norm on absolute differences (instead of relative differences) and then derive error bounds from the contraction mapping theorem. Seeing that our problem is not a contraction we can not derive error bounds in this way, and the argument for using relative differences is that at least it makes the distance measure, and the tolerance we will eventually use, scale invariant.

subsequent iterates by

$$||z^{l} - z^{\infty}|| \le \frac{||z^{l+1} - z^{l}||}{1 - \theta}.$$

Consequently to ensure that the current iterate is within a prespecified tolerance  $\epsilon$  of the limit, we can stop once

$$||z^{l+1} - z^l|| \le \epsilon(1 - \theta). \tag{15}$$

Of course the algorithm described above is *not* a contraction mapping and the condition in equation (14) does not, in general, hold along the entire sequence of iterates. On the other hand, if the algorithm converges, our experience is that it does hold after a "burn in" phase of several iterations. Figure 2 illustrates this point; it plots the relative difference in values and policies between iterations (on a log scale) versus the number of iterations. After iteration 50 or so, the distance decreases as a straight line until it hits machine precision around iteration 325. Unlike a single-agent dynamic programming problem, however, we do not have an obvious guess for the convergence factor. One way to resolve this problem is to estimate  $\theta$  from the iteration history. There are many ways of doing this. For example, the slope of the straight line in Figure 2 is a direct estimate of  $\log_{10} \theta$  because if the current and previous errors are related by  $||z^{l+1} - z^l|| = \theta ||z^l - z^{l-1}||$ , then the current and initial errors are related by  $||z^{l+1} - z^l|| = \theta^l ||z^1 - z^0||$  and hence  $\log_{10} ||z^{l+1} - z^l|| = l \log_{10} \theta + \log_{10} ||z^1 - z^0||$ . Alternatively, suppose we want to continue iterating until we are with  $\epsilon$  of the limit. Then we could let k be the first iteration such that  $||z_k - z_{k-1}|| < 10\epsilon$  and l the first iteration such that  $||z_l - z_{l-1}|| < \epsilon$  to produce the estimate

$$\hat{\theta} = \left(\frac{||z_l - z_{l-1}||}{||z_k - z_{k-1}||}\right)^{\frac{1}{l-k}} \approx 10^{\frac{1}{k-l}}.$$

For further discussion see pp. 42–44 of Judd (1998) and Doraszelski & Judd (2004).

**Gauss-Jacobi vs. Gauss-Seidel.** Pakes & McGuire (1994) describe a Gauss-Jacobi scheme. That is when we compute the *l*th iteration values and policies we only use the information in memory from the (l-1)st iteration. An alternative is to use a so-called Gauss-Seidel scheme (used, e.g., in Benkard 2004, Doraszelski & Judd 2004). Consider the *l*th iteration. When we are at state  $\omega$  we know the new guesses  $V^l(\omega_i^{\dagger}, \omega_{-i}^{\dagger})$ ,  $r^l(\omega_i^{\dagger}, \omega_{-i}^{\dagger})$ ,  $x^l(\omega_i^{\dagger}, \omega_{-i}^{\dagger})$ ,  $r^{e,l}(\omega_i^{\dagger}, \omega_{-i}^{\dagger})$ , and  $x^{e,l}(\omega_i^{\dagger}, \omega_{-i}^{\dagger})$ ) in states  $\omega^{\dagger}$  that are updated prior to state  $\omega$  in each iteration. This information can be used when updating the values and policies at state  $\omega$  requires values and policies from state  $\omega^{\dagger}$ .

In systems of linear equations it can be shown that if both Gauss-Jacobi and Gauss-Seidel schemes converge, then Gauss-Seidel schemes are faster (see, e.g., Section 2.6.2 in Bertsekas & Tsitsiklis 1997). There is some experience suggesting that this is also true in the systems of nonlinear equations that characterize the equilibria of dynamic stochastic games.



Figure 2: Relative difference in values and policies between iterations. Source: Doraszelski & Judd (2004) and own calculations.

**Convergence.** We note that there is no proof that either of these algorithms converge. Moreover when a Gauss-Jacobi scheme converges, a Gauss-Seidel scheme may not, and *vice versa*. Again there is some experience that suggests that Gauss-Seidel schemes are more prone to convergence failures. Pakes & McGuire (1994) discuss a number of problems and fixes. In particular, it is frequently helpful to introduce dampening into the algorithm. In the *l*th iteration after computing  $V^l(\cdot)$ ,  $x^l(\cdot)$ ,  $r^l(\cdot) V^{e,l}(\cdot)$ ,  $x^{e,l}(\cdot)$ ,  $r^{e,l}(\cdot)$  dampening forms a convex combination of the new and the old guesses and uses these as input for the (l+1)th iteration. This may facilitate convergence by "smoothing out" the path that the algorithm takes; see footnote 18 of Pakes & McGuire (1994) and Chapter 3 of Judd (1998) for details.

**Software.** Pakes, Gowrisankaran & McGuire (1993) make Gauss and Matlab code publicly accessible. The code allows for three types of product market competition. One is a differentiated product model. Here equilibrium in the "spot" market is Nash in prices, and firms invest to improve the qualities of their products. The other two models are homogeneous product models; one has differences in marginal costs across firms and the other has differences in capacities. In these models equilibrium in the spot market is Nash in quantities and investment is directed at decreasing marginal cost (in the model with differences in marginal costs) or at increasing capacity (in the model with capacity constraints). For each model the user is allowed to specify values for parameters that set demand, production costs, sunk entry and exit fees, the efficacy of investments, and the discount rate. In addition the user specifies an initial market structure for the analysis. Finally there are additional modules which compute the solutions to a social planner's problem and to a "perfect cartel" problem based on the same parameters, and provide statistics which allow the user to compare these outcomes to those from equilibrium of the dynamic stochastic game. The publicly accessible code has been used extensively as a teaching tool, but it is not designed for computational efficiency and flexibility. For this reason researchers using variants of the EP framework have largely developed their own programs.

Doraszelski & Judd (2004) make a set of C and Matlab codes available upon request.

## 4.2 Computational Burden

The computational burden of computing equilibria is large enough to often limit the type of applied problems that can be analyzed. There are two aspects of the computation that can limit the complexity of the models we analyze; the computer memory required to store the value and policies, and the CPU time required to compute the equilibrium. Here we assume that CPU time is the binding constraint, and analyze its properties.

The compute time is determined by three factors:

- the number of states per iteration;
- the time per state required to update values and policies; and
- the number of iterations.

Together the first two factors determine the time per iteration. The time to convergence is then given by the product of the time per iteration and the number of iterations. We examine the three factors in turn.

Number of states. It is convenient to consider the number of states as a function of the maximum number of firms ever active, our  $\bar{n}$ , and of the number of possible values the firm specific state variable can take on, say  $\#\Omega$ .<sup>17</sup> Recall that provided we restrict attention to symmetric and anonymous equilibria the number of states grows polynomially in both  $\bar{n}$  and in  $\#\Omega$ . Even though there is no curse of dimensionality in the formal sense, the polynomial growth is often a challenge. Moreover, without special structure,  $\#\Omega$  grows exponentially in the number of state variables that take on firm specific values. Suppose, for example, we want to model firms that differ in terms of both product quality and production capacity. If quality and capacity each take on one of  $\bar{\omega}$  values, then  $\#\Omega = \bar{\omega}^2$ . Thus allowing for multiple sources of heterogeneity among firms can be especially costly. A similar problem arises if firms market more than one product, see Langohr (2003) and Song (2002).

<sup>&</sup>lt;sup>17</sup>We note that in fact  $\bar{n}$  and  $\#\Omega$  are endogenous, but the map from primitives is complex to be useful, and for many applied problems there is a *priori* information available on likely values for  $\bar{n}$  and  $\#\Omega$ .

**Time per state.** This consists of computing the summation underlying the expected discounted value of future net cash flow given an investment outcome, or the relevant elements of  $\{W(\cdot)\}$ , and then solving for optimal policies and the probabilities of entry and exit given  $\{W(\cdot)\}$ . As noted for simple enough functional forms the optimal investment can be written down as an analytic function of the  $\{W(\cdot)\}$ , and provided the distributions of entry and exit fees can be evaluated analytically, so can the entry and exit probabilities.

However, as Pakes & McGuire (2001) point out, if we compute transition probabilities as we usually do using unordered states then the number of states that we need to sum over to compute continuation values grows exponentially in *both* the number of firms and the number of firm-specific state variables.<sup>18</sup> To see this, suppose that each firm can move to one of K states from one period to the next. For example if there were two firm specific state variables and each could move up by one, down by one, or stay the same then  $K = 3^2 = 9$ . Then if there are n firms continuing from that state and  $\mathcal{E}$  potential entrants, the total number of states that need to be summed over to obtain the continuation values for a given firm is  $K^{n+\mathcal{E}-1}$ .

Number of iterations. There is little known about what determines the number of iterations and how it scales in the size of the problem. Somewhat surprisingly our experience has been that the number of iterations grows slowly if at all in  $\#S^{\circ}$  but can vary greatly with the nature of the problem. Also the number of iterations does tend to increase with the discount factor  $\beta$ , and by quite a bit when  $\beta$  approaches one (Doraszelski & Judd 2004).

Memory requirements. Some of the more advanced algorithms to be discussed below have, at least in some applications, been able to decrease CPU time to such an extent that memory requirements typically become the binding constraint. This is a relatively recent phenomena and except for Pakes & McGuire (2001) very little attempt has been made to reduce memory requirements. This is an area which requires further attention.

**Summary.** It should be clear from the discussion above that the burden of computing equilibria can often be a significant consideration in choosing the model we analyze. Indeed we typically can add enough sources of heterogeneity to any applied problem to make it impossible to compute the required equilibrium.<sup>19</sup> Moreover this same caveat will apply to the more sophisticated algorithms introduced below. On the other hand all applied work can ever hope to accomplish is to form an approximation to the workings of real industries

<sup>&</sup>lt;sup>18</sup>The alternative, that is distinguishing among the elements in the smaller set  $S^{\circ}$  instead of S, which as we noted is not typically efficient for the kind of problems analyzed to date, would imply that that the sum grows exponentially in the number of state variables per firm and geometrically in the number of firms.

<sup>&</sup>lt;sup>19</sup>This also raises the question of what firms actually do and how close an approximation to actual behavior the Markov perfect framework provides. This is a topic which is largely beyond the scope of this paper, though it is touched on in Pakes & McGuire (2001). Our informal analysis indicates that perhaps the most frequent surprise in empirical work is just how good an approximation the Markov perfect framework gives (see, e.g., Benkard 2004).

that is adequate for the problem being analyzed. What the more sophisticated algorithms allow us to do is weaken the dominance of computational considerations in considering the relevant tradeoffs.

### 4.3 Equilibrium Selection

As we pointed out in Section 3.3, we generally have little reason to believe that there is a unique symmetric and anonymous equilibrium. This raises a number of issues for applied work. Some of them are not directly related to computation. Examples include the estimation of parameters in the presence of multiple equilibria, the ability of data to identify the equilibria (or the equilibrium selection rule) that is being played, and the issue of how to evaluate the possible outcomes of alternative policies in the presence of multiple equilibria (for more detail see, e.g., Pakes (2006) and the literature cited there).

Here we focus more narrowly on questions that involve the interaction between computation and multiple equilibria, starting with the fact that any algorithm which computes an equilibrium implicitly selects an equilibrium. Most applications have sufficed with the selected equilibrium, perhaps after performing some rudimentary search for alternatives.<sup>20</sup>

There is a sense in which any algorithm which converges is locally stable at least with respect to the path along which it converged. On the other hand, Besanko et al. (2004) prove, in the context of a dynamic model with learning-by-doing and organizational forgetting, that there is a set of equilibria that are unstable in the sense that they can not be computed by the Pakes & McGuire (1994) algorithm. An interesting and unexplored question is whether there is something special about the equilibria an algorithm selects. For example, would a learning algorithm either converge to it or at least remain near it if started in an appropriate neighborhood? Perhaps even more important is whether decision making processes used by actual firms would converge to it. If we could characterize stability in these terms, then there would be grounds for excluding the equilibria which are unstable as unlikely to occur.

There remains the question of the rationale for the selected equilibrium. There is a sense in which the building blocks of the Pakes & McGuire (1994) algorithm are intuitively appealing. The algorithm combines value function iteration as familiar from dynamic programming with best reply dynamics (akin to Cournot adjustment) as familiar from static games. Since the ordering of the states does not affect iterates in a Gauss-Jacobi algorithm, the Pakes & McGuire (1994) algorithm is independent of the order in which states are visited. This means that different computers which use the same parameter values and the same (admittedly arbitrary) initial conditions to the Pakes & McGuire (1994) algorithm

<sup>&</sup>lt;sup>20</sup>Neither the Pakes & McGuire (1994) algorithm, nor the modifications that have been made to it to enable it to be used for a broader range of applied problems, offers a systematic approach to computing multiple equilibria. So to the extent that authors have searched for multiplicity, the search has been rather *ad hoc*, usually consisting of computing equilibria many times from different initial conditions, or different orders for cycling through the states in  $S^{\circ}$ .

ought to obtain the same answer. Despite these features of the Pakes & McGuire (1994) algorithm, however, it is not clear that the equilibrium computed by it is the one we should focus on.

Undoubtedly the reasonableness of any selection depends on what the subsequent analysis is to be used for. If we were trying to demonstrate the feasibility of a theoretical proposition, or the feasibility of a particular change in outcomes as we vary some parameter, then probably any of the equilibria would be fine. On the other hand if we were trying to analyze how a particular industry is likely to respond to a proposed policy change, we need to be more careful. In this case we would want to try to figure out which of the possible equilibria is more likely, or at least to bound the outcomes of interest that could arise. As noted above, the question of how to figure out which equilibria are more likely to arise is beyond the scope of this paper. However the possibility of bounding outcomes raises distinct computational questions. In particular can we compute all possible equilibria, or at least trace out a large subset of them? This is the multiplicity issue we come back to below.

## 5 Alleviating the Computational Burden

The computational burden of computing equilibria is oftentimes substantial, thus limiting the applications of the EP framework. In particular, there is a curse of dimensionality in both; (i) computing expectations over successor states, and in (ii) the size of the state space. This section outlines approaches to alleviating the computational burden. We begin with an overview of the ideas behind the techniques and the differences between them. We then proceed to a more detailed description of each of them, and we conclude with some thoughts for future research. The subsections of this section correspond to the different algorithms, and each should be readable independent of whether the others are read.

#### 5.1 Overview

The first approach we describe, due to Doraszelski & Judd (2004), is designed to ease the burden of computing the expectation over successor states, and hence decreases the computation time needed to update values and policies at a particular state (it makes no attempt to alleviate the burden imposed by the number of states). Doraszelski & Judd (2004) set up a continuous-time model in which at any particular instant only one firm experiences a change in its state. As a result if each firm's transition can go to one of Kstates, and there are n firms, we only need to sum over  $(K - 1) \times n$  states to compute continuation values. This is in contrast to the  $K^n$  possible future states that we need to sum over in the computational algorithm described above.<sup>21</sup> This implies that the discrete-

 $<sup>^{21}</sup>$ Note that if we are willing to explicitly restrict players to move one at a time we could obtain similar gains from a discrete-time model in which decisions are made sequentially as a result of a random selec-

and continuous-time models discussed in this paper have different implications. As a result it may (but need not) be the case that one of the models provides a better approximation to behavior in a particular setting than the other. We come back to this point in our discussion of topics for further research as one of those topics is the role of timing assumptions in the models that have been used to date.

The second approach we describe for reducing the burden of computing equilibria is the stochastic algorithm of Pakes & McGuire (2001). This algorithm is designed to alleviate *both* the curse of dimensionality resulting from the need to calculate expectations over successor states at each point in the state space, and that arising from the growth in the number of points in that state space. It use simulations draws both to; (i) approximate expectations, and to (ii) determine a recurrent class of points which eventually become the only points at which it computes equilibrium policies.

Using simulation instead of explicit integration to compute expectations makes the time per state linear in the number of firms (just as in the continuous-time model). The stochastic algorithm, however, also deals with the large size of the state space. It updates only one state at each iteration (the state visited at that iteration), and the same simulation draws used to compute the expectation are used to update the state visited. The algorithm eventually wanders into a recurrent class of points, and once in that class stays within it. That is, after an initial run in period the algorithm only computes expectations on this recurrent class of points, and, as noted, the cardinality of the recurrent class can be much smaller than that of the state space itself. The relationship of the cardinality of the recurrent class to that of the state space depends on the details of the model. However as discussed below for I.O. models with even a reasonably large number of state variables the cardinality of the recurrent class is often only a tiny fraction of that of the state space.

From a conceptual point of view the Pakes & McGuire (2001) algorithm is an adaptation of artificial intelligence algorithms which have been used extensively for both single agent problems and zero sum games (see Bertsekas & Tsitsiklis (1996)) to the problem of computing the equilibrium of dynamic games. It is therefore related to the recent theoretical literature on learning, both in macroeconomics (see, e.g., Sargent (1993) and Lettau & Uhlig (1999) and in games (see, e.g., Fudenberg & Levine (1998)). An additional rationale for the algorithm, then, is that the algorithm's path has an interpretation as an outcome of a learning process that might be a good approximation to behavior in certain situations.

We note that in both these subsections we provide numerical examples with computational time. The examples refer to models in which we increase the number of firms holding the states per firm constant. As we increase the realism of our models we will want to also

tion mechanism and outcomes are realized before the next decision is made. Discrete-time models with a deterministic order of moves have been investigated in classic articles by Maskin & Tirole (1988 a, 1988b); see also Noel (2004) for a computable version of Maskin & Tirole's (1988b) Edgeworth cycle model. From a computational point of view the deterministic order prevents us from using anonymity to reduce the size of the state space. It is the idea of using a random order of moves which preserves anonymity, see Doraszelski (2005).

consider how the computational complexity algorithm varies as we increase the number of state variables per firm. Unfortunately there has been very little experimentation with models with more than one state variable per firm.

Finally we also briefly consider parametric approximation methods similar to those discussed in Judd (1998), and approximations to the decision making process which might be relevant when there are a large number of firms as in Weintraub, Benkard & Van Roy (2005). The first set of methods are designed to reduce the number of points which we have to update at each iteration of the iterative process, the second changes the equilibrium concept in a way that simplifies computation.

Many of the ideas behind the algorithms we describe are complementary. This suggests building hybrids. For example one might want to combine the continuous-time approach of Doraszelski & Judd (2004) with the idea of focusing on the recurrent class due to Pakes & McGuire (2001). Alternatively the idea in Weintraub et al. (2005) of changing the equilibrium conditions may be further exploited by allowing for aggregate shocks and more detailed responses to a few "large" firms, still keeping the idea of a "fringe" of small firms whose every move need not be considered in computing future expectations.

#### 5.2 Continuous-Time Models

Doraszelski & Judd (2004) develop continuous-time stochastic games with a finite number of states. We consider a variant of their game which mimics the assumptions of our the discrete-time model as closely as possible.

The major difference is that in the discrete-time model the time path of the state is a sequence while in the continuous-time model that path is a piecewise-constant, rightcontinuous function of time. Jumps between states occur at random times according to a controlled nonhomogenous Poisson process. Thus we need to specify the hazard rates for each possible reason for a jump, and the probability of transitions conditional on the jump occurring. In a model with industry specific shocks there are four possible reasons for a jump; the investment of one of the incumbents produced a new outcome, the value of the outside alternative increased, an incumbent decided to exit, or a potential entrant decided to enter.

Setup. The horizon is infinite and the state of the game at time t is  $\omega_t = (\omega_{1t}, \ldots, \omega_{nt})$ , where  $\omega_{it}$  is the state of incumbent firm i at time t. Incumbent firm i has to sell its assets in order to exit the industry, and buyers for its assets arrive with hazard rate  $\lambda$ . Upon arrival the buyer makes a take-it-or-leave-it offer of a scrap value,  $\phi_{it}$ , drawn randomly from a distribution  $F(\cdot)$  with continuous and positive density. As in the discrete-time model, the offered scrap value is private information and hence is a random variable from the perspective of other firms. If it is above the incumbent firm's continuation value, the incumbent accepts it and leaves the industry. We let  $\psi(\omega_{it}, \omega_{-it})$  (sometimes abbreviated to  $\psi_{it}$ ) be the probability that incumbent firm *i* exits.

A continuing firm invests  $x_{it} \geq 0$  per unit of time and the probability of its  $\omega_{it}$  transiting to  $\omega_{it}+1$  in any time interval is increasing in this investment. If we use the familiar constant elasticity form that hazard becomes  $x_{i,t}^{\gamma}$ , where the parameter  $\gamma > 0$  is the elasticity of the success hazard with respect to investment expenditures or, equivalently, (the negative of) the elasticity of the expected time to an investment success. As in the discrete-time model we allow for an industry wide shock which has a probability of occurring of  $\delta$  per unit of time and decreases the state of all incumbents by one if it does occur.

We assume that there is always a finite number  $\mathcal{E}$  of potential entrants. Potential entrant i has to acquire assets to enter the industry. Sellers of assets arrive with hazard rate  $\lambda^e$  and make a take-it-or-leave-it offers of  $\phi^e_{it}$ , which is assumed to be drawn randomly from a distribution  $F^e(\cdot)$  with continuous and positive density. If the value of entering at state  $\omega^e$  exceeds this entry cost the potential entrant i accepts the offer and immediately begins to invest. We let  $r^e(\omega_t)$  denote the probability that potential entrant i accepts the offer and *enters* the industry.

At time t the hazard rate of a jump occurring depends on the firms' investment, entry, and exit policies at the given state; i.e. on  $x_t = (x_{1t}, \ldots, x_{nt}), \ \psi_t = (\psi_{1t}, \ldots, \psi_{nt}),$  and  $r^e(\omega_t)$ . If a jump occurs at time t then the change in the state of the industry depends on what generated the jump, and the transition rules given above. More specifically if we let  $\mu(\omega'|\omega_t, x_t, \psi_t, r_t^e)$  be the hazard rate of a jump to  $\omega'$  occurring conditional on current polices and  $\omega$  we have

$$\mu(\omega'|\omega_t, x_t, \psi_t, r_t^e) = \begin{cases} \lambda \psi_{it} & \text{if } \omega' = (\phi, \omega_{-it}), \text{ for } i = 1, \dots, n_t, \\ x_{it}^{\gamma} & \text{if } \omega' = (\omega_{it} + 1, \omega_{-it}), \text{ for } i = 1, \dots, n_t, \\ \mathcal{E}\lambda^e r_t^e & \text{if } \omega' = \omega_t \cup \omega^e, \\ \delta & \text{if } \omega' = (\omega_{1t} - 1, \dots, \omega_{nt} - 1), \end{cases}$$
(16)

where  $\omega^e \cup \omega_t$  is shorthand for the vector obtained by adding an  $\omega^e$  to the vector  $\omega_t$  and reorderring the result appropriately.<sup>22</sup>

The savings in the computation of the expectation over successor states results from the fact that to form that expectation we need only sum over the 2n+2 distinct  $\omega'$  with positive hazards. This occurs because the model implies that in a short time interval there will be (with probability arbitrarily close to one) at most one jump. In the discrete-time model we must keep track of all possible combinations of firms' transitions between time t and time t + 1. Thus the computational burden of computing the expectation over successor states in the continuous-time model grows linearly in the number of firms, thereby avoiding the curse of dimensionality in that expectation.

<sup>&</sup>lt;sup>22</sup>Were we to add the possibility of firm specific component, we would add a hazard for  $\omega' = (\omega_{it} - 1, \omega_{-it})$  for  $i = 1, \ldots, n_t$  to the hazards defined above. Note that  $\sum_{\omega'} \mu(\omega'|\omega_t, x_t, \psi_t, r_t^e)$  is the hazard rate of a change in the state of the game.

Equilibrium. At time t player i chooses an action  $x_{it}$  that depends solely on the current state  $\omega_t$ . As in the discrete-time model, given that all his rivals adopt a Markovian strategy, a player faces a dynamic programming problem and Bellman's (1957) principle of optimality implies that he can do no better than to also adopt a Markovian strategy. For a statement of the principle of optimality in a continuous-time setting see, e.g., p. 327 of Intriligator (1971). Furthermore, although the player gets to pick his action from scratch at each point in time, his optimal action changes only when the state of the game changes.

An incumbent's problem. Suppose that the industry is in state  $\omega$ . The Bellman equation for incumbent firm *i* is similar to the one in discrete time. To see this reinterpret  $\pi(\cdot)$  and  $x_i$  as flows per unit of time and note that over a short interval of length  $\Delta > 0$  incumbent firm *i* solves the dynamic programming problem given by

$$V(\omega_i, \omega_{-i}) \approx \max_{x_i \ge 0} (\pi(\omega_i, \omega_{-i}) - x_i) \Delta + (1 - \rho \Delta) E\left\{ \max\left\{ \chi \phi', V\left(\omega'_i, \omega'_{-i}\right) \right\} | \omega_i, \omega_{-i}, x_i \right\},\right.$$

where  $(1 - \rho \Delta) \approx e^{-\rho \Delta}$  is the analog to the discrete-time discount factor  $\beta$ ,  $\chi$  is an indicator function which takes the value of one if a potential purchaser of the firm's assets arrives at the end of the interval of length  $\Delta$ , and it is understood that the expectations presumes all other agents are making equilibrium choices. Evaluating the right hand side of this equation we get

$$V(\omega_{i},\omega_{-i}) \approx \max_{x_{i}\geq 0} (\pi(\omega_{i},\omega_{-i})-x_{i})\Delta + (1-\Delta\rho) \bigg[ \lambda\psi(\omega_{i},\omega_{-i})\Delta E[\phi|\phi\geq V(\omega_{i},\omega_{-i})] + \\ + \sum_{\omega_{i}',\omega_{-i}'} q^{c}(\omega_{i}',\omega_{-i}'|\omega_{i},\omega_{-i},x_{i})\Delta V(\omega_{i}',\omega_{-i}') \\ + \Big(1-\lambda\psi(\omega_{i},\omega_{-i})\Delta - \sum_{\omega_{i}',\omega_{-i}'} q^{c}(\omega_{i}',\omega_{-i}'|\omega_{i},\omega_{-i},x_{i})\Delta\Big)V(\omega_{i},\omega_{-i})\bigg],$$

where

$$\psi(\omega_i, \omega_{-i}) = 1 - F(V(\omega_i, \omega_{-i}))$$

is the probability of exit and

$$q^{c}(\omega_{i}',\omega_{-i}'|\omega_{i},\omega_{-i},x_{i}) = \mu(\omega_{i}',\omega_{-i}'|\omega_{i},\omega_{-i},x_{i},x_{-i}(\omega_{i},\omega_{-i}),0,\psi_{-i}(\omega_{i},\omega_{-i}),r^{e}(\omega_{i},\omega_{-i}))$$
(17)

are the equilibrium perceptions of incumbent firm i about state-to-state transitions. The value of zero for  $\psi_i$  reflects the fact that these perceptions are conditional on continuing on in the industry.

This last equation breaks up the hazard rate of a change in the state of the industry into a part that is due to exit,  $\lambda \psi_i(\cdot)$ , and a part that is due to all other possible change in
the state,  $\sum_{\omega'_i,\omega'_{-i}} q^c(\omega'_i,\omega'_{-i}|\omega_i,\omega_{-i},x_i)$ . The first term reflects the fact that the firm earns the flow profits minus investment over the  $\Delta$  period. Exit occurs with probability  $\lambda \psi_i(\cdot)\Delta$ and, similar to the discrete-time model, the expected scrap value conditional on exiting is  $E\{\phi|\phi \geq V(\omega_i,\omega_{-i})\}$  rather than its unconditional expectation  $E\{\phi\}$ . The next term covers the possibility of some other change in the state, including an investment success and a depreciation shock. The final term covers the possibility that nothing changes between time t and time  $t + \Delta$ .

As  $\Delta \to 0$ , the above equation simplifies to the Bellman equation

$$\rho V(\omega_{i}, \omega_{-i}) = \max_{x_{i} \ge 0} \pi(\omega_{i}, \omega_{-i}) - x_{i} + \lambda \psi(\omega_{i}, \omega_{-i}) \Big( E\left[\phi | \phi \ge V(\omega_{i}, \omega_{-i})\right] - V(\omega_{i}, \omega_{-i}) \Big) \\ + \sum_{\omega_{i}', \omega_{-i}'} q^{c}(\omega_{i}', \omega_{-i}' | \omega_{i}, \omega_{-i}, x_{i}) \Big( V\left(\omega_{i}', \omega_{-i}'\right) - V\left(\omega_{i}, \omega_{-i}\right) \Big).$$
(18)

Hence,  $V(\omega_i, \omega_{-i})$  can be interpreted as the asset value to incumbent firm *i* of participating in the game. This asset is priced by requiring that the opportunity cost of holding it,  $\rho V(\omega_i, \omega_{-i})$ , equals the current cash flow,  $\pi(\omega_i, \omega_{-i}) - x_i$ , plus the expected capital gain or loss due to exit or some other change in the state.

The strategy of incumbent firm i is found by carrying out the maximization problem on the RHS of the Bellman equation (18). In particular, the optimal investment decision is

$$x(\omega_i, \omega_{-i}) = \arg\max_{x_i} - x_i + \sum_{\omega'_i, \omega'_{-i}} q^c(\omega'_i, \omega'_{-i} | \omega_i, \omega_{-i}, x_i) \left( V\left(\omega'_i, \omega'_{-i}\right) - V\left(\omega_i, \omega_{-i}\right) \right).$$
(19)

Given a success hazard of  $x_i^{\gamma}$  with  $0 < \gamma < 1$ , this simplifies to

$$x(\omega_i, \omega_{-i}) = \max\left\{0, \left(\gamma(V(\omega_i + 1, \omega_{-i}) - V(\omega_i, \omega_{-i}))\right)^{\frac{1}{1-\gamma}}\right\}.$$

More generally, the optimal investment decision satisfies a Kuhn-Tucker condition similar to the discrete-time model. An incumbent exits if the optimized continuation value is less than the offered scrap value. This gives an exit probability of

$$\psi(\omega_i, \omega_{-i}) = 1 - F(V(\omega_i, \omega_{-i})).$$
(20)

An entrant's problem. Suppose that the industry is in state  $\omega$ . If  $V(\omega, \omega^e)$  is the value of entering the industry in state  $\omega^e$  (holding fixed the other firms in state  $\omega$ ) and continuing on as an incumbent, then an entrant will enter if and only if  $V(\omega, \omega^e) \ge \phi^e$ . Consequently the probability of entry conditional on a seller of assets appearing at a particular instant is

$$r^{e}(\omega) = F^{e}(V(\omega, \omega^{e})), \qquad (21)$$

and the hazard rate of entry is  $\mathcal{E}\lambda^e r^e(\omega)$ .

**Computation.** While Doraszelski & Judd (2004) use a Gauss-Seidel algorithm to compute the equilibrium to their continuous-time model, here we provide a Gauss-Jacobi algorithm that is analogous to Pakes & McGuire's (1994) algorithm for the discrete-time model. As usual we begin with an initial set of values and policies for all states  $\omega \in S^{\circ}$ . Iteration l then maps old guesses into new guesses as it cycles over the state space  $S^{\circ}$  in some predetermined order.

Specifically, to update the values and policies for an incumbent firm in state  $\omega$ , we proceed as follows. First, we use the optimality condition in equation (19) to update the investment decision as

$$x^{l}(\omega_{i},\omega_{-i}) = \arg\max_{x_{i}} - x_{i} + \sum_{\omega_{i}',\omega_{-i}'} q^{c,l-1}(\omega_{i}',\omega_{-i}'|\omega_{i},\omega_{-i},x_{i}) \left( V^{l-1}(\omega_{i}',\omega_{-i}') - V^{l-1}(\omega_{i},\omega_{-i}) \right),$$

where  $q^{c,l-1}(\omega'_i, \omega'_{-i}|\omega_i, \omega_{-i}, x_i)$  is constructed by substituting the policies of the other firms from the previous iteration into equation (17). Second, we use equation (20) to update the exit decision as

$$\psi^l(\omega_i, \omega_{-i}) = 1 - F(V^{l-1}(\omega_i, \omega_{-i}))$$

Third, we use  $x^{l}(\omega_{i}, \omega_{-i})$  and  $\psi^{l}(\omega_{i}, \omega_{-i})$  to update the value as

$$V^{l}(\omega_{i},\omega_{-i}) = \frac{1}{\rho + \lambda\psi^{l}(\omega_{i},\omega_{-i}) + \sum_{\omega'_{i},\omega'_{-i}}q^{c,l-1}(\omega'_{i},\omega'_{-i}|\omega_{i},\omega_{-i},x^{l}(\omega_{i},\omega_{-i})))} \times \left\{\pi(\omega_{i},\omega_{-i}) - x^{l}(\omega_{i},\omega_{-i}) + \lambda\psi^{l}(\omega_{i},\omega_{-i})E[\phi|\phi \ge V^{l-1}(\omega_{i},\omega_{-i})]\right\} + \sum_{\omega'_{i},\omega'_{-i}}q^{c,l-1}(\omega'_{i},\omega'_{-i}|\omega_{i},\omega_{-i},x^{l}(\omega_{i},\omega_{-i}))V^{l-1}(\omega'_{i},\omega'_{-i})\right\}.$$

$$(22)$$

By rearranging equation (18) and dividing through by the sum of the discount factor and the hazard rate of a change in the state of the industry, we ensure that equation (22) would be a contraction were we to iterate on it holding the policies of the other firms fixed (just as it is in the equation which updates values in the discrete-time model). Fourth, we update the entry probability conditional on the appearance of a seller of assets by setting

$$r^{e,l}(\omega) = F^e(V^{l-1}(\omega, \omega^e)).$$

The smaller number in successor states in the continuous-time model may allow us to speed up the computations in one other way. Prior to adding each of the terms that enter in the expectation over successor states, both the continuous- and the discrete-time algorithms need to look them up in computer memory. This requires the algorithms to compute the addresses of the successor states in computer memory and imposes a further cost. One way to speed up the computations is to compute these addresses once and then store them for future reference. Precomputed addresses decreases running times but increases memory requirements. When the number of successor states is small enough the savings from this prior computation outweigh the costs, so it is frequently useful in the continuous-time model, but not in the discrete-time model. We note also that the homotopy methods that we discuss in Section 6 can use the continuous-time models just as the Gaussian methods described here do, and that this would result in similar reductions of the computational burden.

**Example.** Doraszelski & Judd (2004) compare a modified version of the Pakes & McGuire (1994) quality ladder model to compare the continuous- and discrete-time approaches. In order to avoid the existence problems that may arise from the discrete nature of firms' entry and exit decision (see Section 3.1 for details), they abstract from entry and exit. Moreover, they differ from the original quality ladder model in that their shocks are independent across firms whereas Pakes & McGuire (1994) assume an industry-wide shock (see Section 2).

As Table 2 shows, each iteration of the continuous-time algorithm is far faster than its discrete-time equivalent. Partly offsetting this is the fact that for comparable games the continuous-time algorithm needs more iterations to converge to the equilibrium. The reason that the continuous-time algorithm suffers an iteration penalty is that the continuous-time "contraction factor", i.e.

$$\frac{\sum_{\omega_i',\omega_{-i}'} q^{c,l-1}(\omega_i',\omega_{-i}'|\omega_i,\omega_{-i},x^l(\omega_i,\omega_{-i}))}{\rho + \lambda \psi^l(\omega_i,\omega_{-i}) + \sum_{\omega_i',\omega_{-i}'} q^{c,l-1}(\omega_i',\omega_{-i}'|\omega_i,\omega_{-i},x^l(\omega_i,\omega_{-i}))}$$

in equation (22) varies with firms' policies and substantially exceeds the discrete-time discount factor  $\beta$ . Even though the continuous-time algorithm needs more iterations, the loss in the number of iterations is small when compared to the gain from avoiding the curse of dimensionality and the (much smaller) gain from precomputed addresses. Table 2 illustrates this comparison and the total gain from continuous time. Continuous time beats discrete time by 60% if n = 3, a factor of 12 if n = 6, a factor of 209 if n = 9, a factor of 3,977 if n = 12, and a factor of 29,734 if n = 14.

**Software.** Doraszelski & Judd (2004) make a set of C and Matlab codes available upon request.

### 5.3 Stochastic Approximation Algorithm

The stochastic algorithm breaks the curses of dimensionality in the number of elements of the summand in calculating the expectation over future states *and* in the number of states. It is able to do so because

• it never attempts to obtain accurate policies on the entire state space, just on a

				$\operatorname{ratio}$	
#firms	discrete	$\operatorname{continuous}$	time per	number of	time to con-
	time	time	iteration	iterations	vergence
	mins.	mins.			
2	$1.80 \times 10^{-4}$	$1.12 \times 10^{-4}$	1.73	0.93	1.61
3	$1.42 \times 10^{-3}$	$8.83 \times 10^{-4}$	2.88	0.56	1.60
4	$1.13 \times 10^{-2}$	$4.43 \times 10^{-3}$	6.10	0.42	2.54
5	$8.78 \times 10^{-2}$	$1.70 \times 10^{-2}$	14.57	0.36	5.18
6	$6.42 \times 10^{-1}$	$5.34 \times 10^{-2}$	37.12	0.32	12.03
7	$4.44 \times 10^0$	$1.47  imes 10^{-1}$	98.26	0.31	30.19
8	$2.67  imes 10^1$	$3.56 imes10^{-1}$	249.38	0.30	74.94
9	$1.66  imes 10^2$	$7.95  imes 10^{-1}$	709.04	0.29	208.85
10	$9.28 \times 10^2$	$1.77 \times 10^0$	1800.00	0.29	523.72
11	$4.94 \times 10^{3}$	$3.30 \times 10^0$	5187.50	0.29	1498.33
12	$2.46 \times 10^{4}$	$6.18 \times 10^0$	13770.00	0.29	3977.26
13	$1.27 \times 10^{5}$	$1.13 \times 10^1$	39033.56	0.29	11246.96
14	$6.00 \times 10^5$	$2.02  imes 10^1$	103195.27	0.29	29734.23

Table 2: Time to convergence and ratio of discrete to continuous time in terms of time per iteration, number of iterations, and time to convergence. Stopping rule is "distance to truth  $< 10^{-4}$ ." Entries in italics are based on an estimated 119 iterations to convergence in discrete time. Source: Doraszelski & Judd (2004).

recurrent class of points, and the number of points in the recurrent class need not grow in any particular way with the dimension of the state space, and

• it never calculates integrals over possible future values, rather it approximates those integrals with an average of past outcomes.

To ease exposition of how this is done this section will modify the presentation in two ways. First we will initially assume that the exit and entry costs are fixed numbers, our  $(\phi, \phi^e)$ , rather than draws from distributions, and that there is only one potential entrant in each period (who will either enter or not as a deterministic function of equilibrium values). Accordingly since all incumbents have the same  $\phi$  we do not include  $\phi$  as an argument of the value function. We include a paragraph indicating how to modify the algorithm to revert to the initial specification below.

Second we will rewrite the value function in a way that makes it explicit that the equilibrium conditions can be expressed as a fixed point in the  $\{W(\cdot)\}$ , i.e. in the expected discounted value of future net cash flows conditional on the different possible outcomes of the firm's investment expenditures, rather than as a fixed point in the values per se. That is we rewrite the fixed point conditions at every  $\omega \in S^{\circ}$  so that incumbent values are given by

$$V(\omega_i, \omega_{-i}|W) = \pi(\omega_i, \omega_{-i}) + \max\left\{\phi, \max_{x_i}\left(-x_i + \beta \sum_{\nu} W(\nu|\omega_i, \omega_{-i})p(\nu|x_i)\right)\right\}$$
(23)

potential entrants values are given by

$$V(\omega^{e}, \omega | W) = \max\left\{0, -\phi^{e} + \max_{x^{e}}\left(-x^{e} + \beta \sum_{\nu} W(\nu | \omega^{e}, \omega) p(\nu | x^{e})\right)\right\}$$
(24)

and

$$W(\nu|\omega_i,\omega_{-i}) = \sum_{\omega'_{-i},\eta} V(\omega_i + \nu - \eta, \omega'_{-i}|W)q(\omega'_{-i},\omega,\eta)p(\eta),$$
(25)

where  $q(\omega'_{-i}, \omega, \eta)$  are derived from the policies generated by equations (23) and (24) and the primitives of the problem. Note that here and below when i = e it is understood that  $\omega_{-i} = \omega$ , the vector of all of the incumbents  $\omega$ 's.

Writing the value function in this way makes it clear that the  $W(\cdot)$  are sufficient statistics for agents' behavior. That is if we knew the  $W(\cdot)$ 's we would know equilibrium policies and values. The stochastic algorithm recasts the search for equilibrium as a search for a set of numbers, the  $\{W(\cdot)\}$ , that satisfy the fixed point implicit in the definition of the value function in terms of  $W(\cdot)$ , our  $V(\cdot|W)$  in equations (23) and (24), and the definition of  $W(\cdot)$ in terms of  $V(\cdot|W)$  in equation (25).

**Iterative search procedure.** The stochastic algorithm is also iterative. However in contrast to the other algorithms discussed here it is *asynchronous*; it only calculates policies for a single state at each iteration. Consequently it must hold in memory both

- a location, and
- estimates of the expected discounted value of the future net cash flows that would result from each possible outcome of each firm's actions (i.e. of the  $W(\cdot)$ 's in equation (23) and (24)).

Let  $W(\omega)$  refer to the vector of values of  $W(\nu; \omega_i, \omega_{-i})$  associated with the incumbents and potential entrant at  $\omega$ , W refer to the entire collection of  $W(\omega)$  (so  $W = \{W(\omega) : \omega \in S^{\circ}\}$ ), and l index iterations. Then an iteration must update both the location, say  $\omega^l$ , and the estimates of equilibrium  $W(\cdot)$ , or  $W^l$ .

In the first step of the update incumbents and the potential entrant at the current location chose the policies that maximizes the current iteration's estimate of expected discounted values of net cash flows. That is the first step substitutes components of  $W^l(\omega^l)$ for the equilibrium components of  $W(\omega)$  appearing in equations (23) and (24), and then determines the optimal policies implied by these estimates in a manner analogous to that described in equations (4), (6), and (9) above. These policies determine the distribution of the next location. To get the actual realization of that location we take computer generated pseudo random draws from the distributions determined by these policies and the primitives of the model. Then a potential entrants whose policy was to enter does so at the state determined by the simulated draws, incumbents whose policy were to continue update their state with their simulated draws, and incumbents whose policy was to exit do so. More precisely the location of the firms active in l+1 are given by  $\omega_i^l + \nu_i^{l+1} - \eta^{l+1}$  for continuing incumbents and  $\omega^e + \nu_e^{l+1} - \eta^{l+1}$ for a new entrant that enters. Here the  $\nu_i^{l+1}$  and  $\nu_e^{l+1}$  are draws from the distributions  $p(\nu|x_i^l)$  and  $p(\nu|x_e^l)$  where the x's are the policies determined above, while  $\eta^{l+1}$  is a draw from  $p(\eta)$  (the distribution of increments to the outside alternative). The vector of new states, re-ordered to be in their natural order, becomes  $\omega^{l+1}$ , the new location.

To complete the iteration we need also to update the estimates of the W. We only update the estimates of  $\{W(\omega^l)\}$  (the components of W associated with the location of the algorithm at iteration l). From equation (25), each equilibrium  $W(\nu; \omega_i, \omega_{-i})$  is an expectations of the discounted value of the firm's state in the next period conditional on a realization of its own investment outcome. The expectation is over the possible outcome of the competitors' states and of the outside alternative.

The update of the estimate of  $W(\omega^l)$  acts as if the random draws on the outcomes of the competitors' states and of the outside alternative are random draws from the equilibrium distribution of outcomes for the competitors' states and from the outside alternative; i.e. as if they are random draws from the equilibrium  $q(\omega'_{-i}, \omega, \eta)$  and  $p(\eta)$ . It then evaluates the integrand at the new state by acting as if the current iteration's estimate of  $W(\omega^l)$ , our  $W^l(\omega^l)$ , were the equilibrium values of  $W(\omega^l)$ . That is it acts as if  $V(\omega_i^l + \nu - \eta^{l+1}, \omega_{-i}^{l+1}, |W^l)$ , formed from the random draws and  $W^l(\omega^l)$  and computed as in equation (23), is the equilibrium value associated with that point.  $V(\omega_i^l + \nu - \eta^{l+1}, \omega_{-i}^{l+1}, |W^l)$  is then treated as a random draw from the integral determining W and averaged with the current iterations estimate of W, or  $W^l$ , to obtain  $W^{l+1}$ . That is to form its  $W^{l+1}(\cdot)$  the algorithm chooses an  $\alpha(\omega, l) > 0$  and sets

$$W^{l+1}(\nu;\omega_{i}^{l},\omega_{-i}^{l}) = \frac{1}{1+\alpha(\omega,l+1)}V(\omega_{i}^{l}+\nu-\eta^{l+1},\omega_{-i}^{l+1}|W^{l}) + \frac{\alpha(\omega,l)}{1+\alpha(\omega,l+1)}W^{l}(\nu;\omega_{i},\omega_{-i})$$
(26)

So  $W^{l+1}(\nu_i; \omega_i, \omega_{-i}) - W^l(\nu_i; \omega_i, \omega_{-i})$  is greater than, equal to, or less than zero according as  $V(\omega_i^l + \nu - \eta^{l+1}, \omega_{-i}^{l+1} | W^l) - W^l(\nu_i; \omega_i, \omega_{-i})$  is greater than, equal to, or less than zero.

If  $\alpha(\omega, l)$  equals the number of times the state  $\omega = \omega^l$  had been visited prior to iteration l, then the formulae in (26) produces the simple of average of the values obtained from those visits. In general there are "better" ways of choosing the  $\alpha(\cdot)$  then this. This because the earlier outcomes should be less precise estimates of the numbers we are after (and hence should receive less weight; we provide one example below). On the other hand all proposed

weighting schemes satisfy Robbins & Monro's (1951) convergence conditions; i.e. the sum of the weight increase without bound as the number of visits increase, but the sum of the squared weights remains bounded (conditions which are satisfied by simple averages).

Two conceptual points about the updates. First if  $W^{l}(\omega^{l})$  equals the equilibrium  $W(\omega^{l})$ then the expectation of the  $V(\omega_{i}^{l} + \nu - \eta^{l+1}, \omega_{-i}^{l+1} | W^{l})$  is also the equilibrium  $W(\omega^{l})$ . That is once we are at the equilibrium values of  $W(\omega^{l})$  we will tend to stay there. However as is explained below the algorithm only repeatedly visits points in a recurrent subset of the state space. As a result we should not expect the algorithm to generate precise estimates of the equilibrium  $W(\omega)$ 's associated with  $\omega$ 's that are outside the recurrent class. Consequently below we modify the conditions which define equilibrium policies and values so that they pertain only to a recurrent class of points, and show why they are relevant for subgames starting from an initial condition in the recurrent class. We then discuss tests of whether these equilibrium conditions are satisfied by the estimate of W outputted by the algorithm.

Second use of the Monte Carlo draw to both determine the evolution of the state, and to estimate the returns to alternative actions, mimics what would happen were agents actually implementing policies based on our procedure and then using the actual market outcomes to update their estimate of the implications of their actions. This provides both the link to the learning literature in economic theory, and the reason one might want to use the algorithm as a behavioral model in certain instances. That is if we knew the  $W(\cdot)$ 's that the agents' perceived at any point in time, and were willing to specify how those agents updated their perceptions of W as a result of information that accumulated over time, we could generate a probability distribution over sample paths for both the policies followed and the states achieved.

**Computational burden.** We now review the reasons for the stochastic algorithm's computational savings. First the location of the stochastic algorithm will eventually wander into a recurrent class of points (our R), and once within that R will stay within it forever (with probability one). That is after an initial run in period the stochastic algorithm only updates the values and policies associated with points in this recurrent class (the other algorithms update on the entire state space). In the examples we have computed the recurrent class tends to grow linearly in the number of state variables, thus overcoming one of the two aspects of the curse of dimensionality noted in the last section.

Second the stochastic algorithm does not require a summation over possible future states to update the information in memory at each point visited. The computational burden at a point involves only; determining the optimal policies given  $W^l(\omega^l)$ , drawing the random variables whose distributions are determined by these policies, and updating the estimates of  $W(\omega)$  with equation (26). The computational burden of these tasks grows linearly in the number of firms active at the state, thus eliminating the second aspect of the curse of dimensionality noted in the last section. However unlike the deterministic updates used in the other algorithms discussed here, the Monte Carlo estimate of  $W(\omega)$  contains the variance induced by the simulated draws on outcomes. If the estimates of  $W(\omega)$  settles down, the precision of the Monte Carlo estimate of the integral will become proportional to the number of times the state is visited. As a result the accuracy of the estimates at a point depend on the number of visits to the point, and sufficient accuracy can take a larger number of iterations per point than the other algorithms discussed here. On the other hand the precision of the Monte Carlo estimate of the integral does not (necessarily) depend on the dimension of the integral being estimated, while the cost of doing the summation explicitly does. So the number of visits needed for given accuracy need not (and in our experience does not) depend on the number of state variables.

Equilibrium policies and stopping rules. Let  $\tilde{W}$  be a particular value of W. As noted once we substitute  $\tilde{W}(\omega)$  into equations (23) and (24) we determines policies for all agents active at that  $\omega$  (though these policies will not, in general, be equilibrium policies). These policies determine the probabilities of transiting to any future state if the policies generated by  $\tilde{W}(\omega)$  are followed. Let those probabilities be given by  $q(\omega', \omega | \tilde{W}(\omega))$ ,  $\forall \omega' \in$  $S^{\circ}$ . Now order the states and arrange these probabilities into a row vector in that order, say  $q(\omega | \tilde{W}(\omega))$ . Do this for each  $\omega \in S^{\circ}$ , and combine the resultant rows into a matrix whose rows are ordered by the same order used to order the elements in each row. The result is a Markov matrix (or transition kernel) for the industry structures, say  $Q(\cdot, \cdot | \tilde{W})$ . This matrix defines the Markov process for industry structures that would be generated if all agents acted as if  $\tilde{W}$  were the true expected discounted values of alternative outcomes.

Any finite state Markov kernel generates a stochastic process with at least one recurrent class. Say  $\tilde{R} \subset S^{\circ}$  is a recurrent class of the Markov process defined by  $Q(\cdot, \cdot | \tilde{W})$ . I.e. if we start the process at an  $\omega^0 \in \tilde{R}$  and follow the policies generated by  $\tilde{W}$  then each element of all possible sequences,  $\{\omega^k\}_{k=0}^{\infty}$ , will be in  $\tilde{R}$  with probability one. Call the subvector of  $\tilde{W}$  which contains the elements of  $\tilde{W}$  associated with the points in  $\tilde{R}$ ,  $\tilde{W}|\tilde{R}$ . Since  $\tilde{R}$  is a recurrent class of  $Q(\cdot, \cdot | \tilde{W})$ ,  $\tilde{W} | \tilde{R}$  contains all information needed to analyze any game starting from any initial condition  $\omega^0 \in \tilde{R}$ . That is  $\tilde{W} | \tilde{R}$  generates policies "for subgames from  $\tilde{R}$ ".

Pakes & McGuire (2001) provide conditions which insure that W|R generates equilibrium policies for subgames from  $\tilde{R}$ . They then propose a procedure which identifies a recurrent class of  $Q(\cdot, \cdot|\tilde{W})$ , our  $\tilde{R}$ , and tests whether  $\tilde{W}|\tilde{R}$  satisfies the equilibrium conditions. We explain their procedure next. More details on the conceptual issues underlying these calculations are given below.

Say we want to test whether a subset of the  $\tilde{W}$  outputted by the algorithm, and the associated values and policies, are equilibrium policies for a recurrent class of  $Q(\cdot, \cdot | \tilde{W})$ . First we need to obtain a candidate for a recurrent class of  $Q(\cdot, \cdot | \tilde{W})$ . To do so use the policies generated by  $\tilde{W}$  to simulate  $\{\omega^j\}_{j=1}^{J_1+J_2}$ . Let  $R(J_1, J_2)$  be the set of states visited at least once between  $j = J_1$  and  $j = J_2$ . Then, as both  $J_1$  and  $J_2 - J_1 \to \infty$ ,  $R(J_1, J_2)$ must converge to a recurrent class of the the process  $Q(\cdot, \cdot|\tilde{W})$ . As we shall see it typically does not take long to generate a million iterations of the stochastic algorithm. As a result it is easy to simulate several million draws, throw out a few million, and then consider the locations visited by the remainder as the recurrent class.

Let  $\tilde{W}|\tilde{R}(J_1, J_2)$  contain the components of the  $\tilde{W}$  vector outputted from the algorithm that are needed to obtain the policies for all the  $\omega \in R(J_1, J_2)$ . Pakes & McGuire's (2001) test of whether  $W|\tilde{R}(J_1, J_2)$  generates equilibrium policies for subgames starting in  $R(J_1, J_2)$ consists of checking if equations (23) and (24) are satisfied to sufficient accuracy with Wcomputed from (25) after substituting  $V(\cdot|\tilde{W})$  for  $V(\cdot|W)$  into that equation. Sufficient accuracy is defined in terms of a norm on the  $\omega \in R(J_1, J_2)$  (and the comparison is made for each incumbent and potential entrant at those  $\omega$ ). Since the stochastic algorithm is designed to obtain more accurate estimates of policies and values for points that are visited more frequently, Pakes & McGuire (2001) suggest using a weighted sum of squares norm with weights proportional to the number of visits to the point. Note that this test is the same as the test used to stop the Gaussian algorithm in Section (4) except that now we confine the test to points in  $R(J_1, J_2)$  and use a weighted norm.

If the researcher is particularly interested in policies and values from a given point, say a point which reflects the structure of a particular industry at a time of interest, the test can be altered to reflect that fact. More generally the algorithm, as programmed in Pakes & McGuire (2001), outputs a count of how many times each point has been visited. This provides a good indication of how precisely the values and policies at the given point have been estimated. Local restart procedures are advisable if, for some reason, the researcher is interested in a point rarely visited by the natural progression of the algorithm.

Note that the computational burden of the test in Pakes & McGuire (2001) goes up either geometrically or exponentially in the dimension of the state space (depending on the type of states). Moreover the test is the *only* aspect of the stochastic algorithm whose computational burden grows more than linearly in the dimension of the state space. So as we increase the number of state variables the computational burden of the test eventually outweighs the burden of the rest of the algorithm (see below). As a result Fershtman & Pakes (2005) developed an alternative test which both circumvents this curse, and is easy to apply. We come back to this alternative test below.

**Equilibrium: Conceptual issues.** We now provide more detail on the conceptual issues underlying the notion of equilibrium and the test (a reader not interested in these details should be able to skip this section and have no difficulty with the sections that follow).

Formally the conditions being tested above are not sufficient to guarantee that the policies and values outputted by the algorithm are Markov Perfect equilibrium policies and values for states in the recurrent class. To see this note that though all points in the recurrent class only communicate with other points in the recurrent class if optimal policies are followed, there are points in the recurrent class that could communicate with points outside the recurrent class if feasible (though in-optimal) policies are followed. When this occurs the  $V(\cdot, \cdot | \tilde{W})$  needed to check equation (25) will contain values which the test does not insure are equilibrium values (since they are points which are not in the recurrent class we do not test the condition in equation 23 for them). For an example of this recall that  $\overline{n}$  is the maximum number of agents ever active, and consider a state in  $R(J_1, J_2)$  with  $\overline{n}$  firms active. To determine whether the policies at that state are optimal we need to know what values would be were the potential entrant to enter and no incumbent exit; i.e. what the values would be if there were  $\overline{n} + 1$  firms active. We do not check whether the values at such states satisfy our equilibrium conditions<sup>23</sup>.

Pakes & McGuire (2001) call points that are in  $\tilde{R}$  but could communicate with points outside of  $\tilde{R}$  if feasible (though in-optimal) policies were followed, boundary points of  $\tilde{R}$ . Points in  $\tilde{R}$  that are not boundary points are called interior points of R. Interior points could not communicate with a point outside of R under any feasible policy. They then show that when the test needs to use a  $V(\cdot, \cdot | \tilde{W})$  associated with an  $\omega \notin \tilde{R}$  to compute a W (using equation (25)) to check equation (23) or (24) at a point in  $\tilde{R}$ , then that  $V(\cdot, \cdot|W)$  need not satisfy the equilibrium condition (equation (23) or (24)) exactly; it only needs to be larger than the true equilibrium value. The reason is that if the firm chooses not to undertake a policy that would lead them to a point outside of  $\hat{R}$  when those points are evaluated too highly, we know they would not chose such an action if those points were evaluated at their true values. They then consider ways of bringing the weaker condition into the test for equilibrium at boundary points, and report that boundary points are visited so infrequently that bringing in different reasonable ways of checking these conditions had no effect on the estimates<sup>24</sup>. We note, however, that as in single agent applications of stochastic algorithms (see the references below), to minimize the probability that the algorithm ignores relevant actions they initialized the algorithm with an initial estimate of W (a  $W^0$ ) thought to be much larger than the equilibrium W (two candidates are the value function from the one firm problem and  $\pi(\omega_i, \omega_{-i})/(1-\beta)$ ). This tends to make the program stop at a set of W which are, if anything, too large, and this in turn insures that the conditions required for boundary points are satisfied (though we can not prove this in the context of the dynamic games that we are considering here).

Fershtman & Pakes (2005) treat this issue differently. They consider a notion of equilibrium whose conditions can be consistently tested on actual data. They call their equilib-

<sup>&</sup>lt;sup>23</sup>Moroever since the algorithm does not visit points with  $\overline{n} + 1$  agents active repeatedly, it can not be expected to produce an accurate estimates of the values at that point.

<sup>&</sup>lt;sup>24</sup>Given this fact they found it computationally efficient to set the probability of events that would lead from an  $\omega \in R(J_1, J_2)$  to a point outside of  $R(J_1, J_2)$  to zero, and renormalize the remaining probabilities appropriately.

rium notion applied Markov perfect equilibrium; it is closely related to the notion of selfconfirming equilibrium in Fudenberg & Levine (1993). The conditions of applied Markov perfect equilibrium only place restrictions on policies and values on the recurrent class of points. Roughly those conditions are that: (i) in equilibrium policies at those points must be optimal with respect to the equilibrium W, and (ii) for every  $(\omega_i, \omega_{-i})$  associated with a point in the recurrent class and every  $\nu$ ,  $W(\nu; \omega_i, \omega_{-i})$  must satisfy the fixed point condition in equation (25) only if optimal policies indicate that the probability of that  $\nu$  is greater than zero. One test of the sufficient conditions for applied Markov perfect equilibrium is the test introduced above, but they introduce a test that is less computationally burdensome (see below).

**Experience with stochastic algorithms.** Most of the experience with stochastic algorithms is with single agent problems or zero sum games in which all points are recurrent (for an easy to read review see Barto, Bradtke & Singh (1995)). In these problems one can generally prove convergence (in probability, or almost surely, see Bertsekas & Tsitsiklis (1996)). These proofs require one to start with an *overestimate* of the estimated values (so all points are tried before they are discarded). The striking result from the numerical exercises is that the imprecision in their estimates at points with little weight in the ergodic distribution *does not* seem to impact adversely on the precision of the estimates of the values associated with points with heavy weight in the ergodic distribution (see also below).

Once we move to dynamic games we cannot guarantee convergence of the stochastic algorithm (just as we could not for the deterministic algorithms discussed above). On the other hand the two applications of the stochastic algorithm we know of have never experienced a convergence problem. Interestingly those models worked with the version of the EP model which assumes a single known entry fee and exit value, just as we have done in this section. This is precisely the version of the EP model which Pakes & McGuire (1994) found often generate convergence problems when computed with their Gaussian algorithm. This makes the absence of convergence problems in the stochastic algorithm particularly notable.

To allow for random entry and exit fees we need only modify the algorithm slightly. Equation (23) must be modified to read

$$V(\omega_i, \omega_{-i}|W) = \pi(\omega_i, \omega_{-i}) + [1 - r(\omega_i, \omega_{-i}|W)]\phi(\omega_i, \omega_{-i}|W)$$
$$+ r(\omega_i, \omega_{-i}|W) \left\{ -x(\omega_i, \omega_{-i}|W) + \beta \sum_{\nu} W(\nu|\omega_i, \omega_{-i})p(\nu|x(\omega_i, \omega_{-i}|W)) \right\},$$

where  $r(\omega_i, \omega_{-i}|W)$  is the probability of drawing a random exit fee less than the continuation value, and  $\phi(\omega_i, \omega_{-i}|W)$  is the expected value of the exit fee conditional on it being greater than the continuation value. These values are calculated from the exit policy at each  $(\omega_i, \omega_{-i})$  and held in memory. Similarly the entry policy becomes a value for  $\phi^e$  above which the potential entrant does not enter, and that value is also held in memory at  $(\omega_i, \omega_{-i})$ .

**Example.** To illustrate Pakes & McGuire (2001) computed the same equilibrium computed pointwise in Pakes & McGuire (1994), and showed that the stochastic algorithm computed estimates whose implications are indistinguishable from those of the Gaussian algorithm.

The model in Pakes & McGuire (1994) has only one state per firm, so the only dimension in which we can trace out the relationship between the components of computational burden, and the size of the problem, is as we increase  $\bar{n}$ .  $\bar{n}$  for Pakes & McGuire's (1994) problem is largely determined by their M (the number of consumers serviced by the market). So we will focus on how the components of the computational burden changes as we push Mup from Pakes & McGuire's (1994) initial M = 5 by units of 1 until M = 10. Recall that we know that the time per state in the stochastic algorithm scales linearly in the number of firms, but we do not know how the size of the recurrent class increases in the number of state variables, so this is the issue we are particularly interested in.

The version of the algorithm used in Pakes & McGuire (2001) downweighs the early estimates of continuation values (which, recall, are imprecisely estimated). It does so by restarting the algorithm after each of the first seven million iterations, using the terminal values from those iterations as starting values of the next million iterations. The algorithm then begins a long run which was interrupted every million iterations to run the test. The test compared the estimates of the values obtained from the current estimate of W and equations (23) and (24) to the estimates obtained from equations (23) and (24) when the W estimates were obtained from the calculation in equation (25) (the weights were set proportional to the number of visits to the location in the last million draws). If the weighted correlation between the two estimates of the value function was over .995 and the difference between their weighted means were less than 1%, the algorithm stopped. The bottom panel of the table provides statistics which enable us to assess how the computational burden changed with market size. The top panel provides the distribution of the number of firm's active from a 100,000 period simulation and our estimated policies.

The number of points in the last row refers to the number of points visited at least once in the last million iterations. This will be our approximation to the size of the recurrent class (#R). There were 21,300 such points when M = 5. As expected  $\bar{n}$  increases in M. However #R does not increase geometrically in  $\bar{n}$ . The top part of the panel makes it clear why; when we increase M the number of distinct points at which there are a large number of firms active does increase, but the larger market no longer supports configurations with a small number of firms. Indeed though the function relating #R to M seems initially convex, it then turns concave giving the impression that it may asymptote to a finite upper bound. The ratio of the number of states at which we compute policies and values in the stochastic algorithm to that in the Gaussian algorithms described above is #R/#S.

M =	5	6	7	8	9	10	
Percer	Percentage of equilibria with $n$ firms active						
n =	0	1					
3	58.3	00.8	00.0	00.0	00.0	00.0	
4	33.7	77.5	48.9	04.4	00.7	00.1	
5	06.3	16.8	41.4	62.3	33.0	07.2	
6	01.5	04.2	07.3	25.0	44.3	41.8	
7	00.2	00.6	02.2	06.5	15.3	34.3	
8	00.0	00.1	00.2	01.7	05.9	13.1	
9	00.0	00.0	00.0	00.0	00.8	03.5	
10	00.0	00.0	00.0	00.0	00.0	00.0	
Avera	Average $n$						
	3.43	4.26	4.64	5.39	5.95	6.64	
Minutes per Million Iterations							
	5.5	6.5	7.5	8.6	10	11	
Minutes per Test							
	3.6	8.15	17.1	42.8	100	120	
Number of Iterations (millions)							
	7 + 5	7 + 2	7 + 21	7 + 4	7 + 9	7 + 3	
Number of Points (thousands)							
	21.3	30.5	44.2	68.1	98.0	117.5	

Table 3: Comparisons for Increasing Market Size<sup>a</sup>.

<sup>a</sup>All runs were on a Sun SPARCStation 2. Source: Pakes & McGuire (2001).

 $\#R/\#S \approx 3.3\%$  when  $\overline{n} = 6$ , about .4% when  $\overline{n} = 10$  (then  $\#S \approx 3.2 \times 10^7$ ), and would decline further for larger  $\overline{n}$ . We note here that we have gotten reductions at least as large as this when we computed equilibria to models that had more than one state per firm. In particular in models where products had more than one dimension which could vary with the firms' investments, there were typically many combinations of characteristics that were not profitable to produce, and the stochastic algorithm would simply never visit them.

The other two components of computational time are the time needed to update values and compute policies at each state, and the number of iterations. As noted we expected the time per state to be a function of the number of firms active at the state, and what the table shows is that the ratio of the average number of active firms to the time per million iterations was essentially constant as we increased M (it varied between 1.53 and 1.65 minutes). The good news was that the number of iterations until our test criteria was satisfied did not tend to increase in M (though it did vary quite a bit between runs).

Thus *absent* test times, the average CPU time needed for our algorithm seems to grow *linearly* in the average number of firms active. Given the theoretical discussion this is as good as we could have expected. Comparing the level of our CPU times to those from backward solution algorithms we find that when  $\overline{n} = 6$  the stochastic algorithm took about a third of the CPU time, when  $\overline{n} = 10$  even the most optimistic projection for the backward techniques leads to a ratio of CPU times of .09%. If  $\overline{n}$  grew much beyond that, or if there were more than one state variable per firm, the backward solution algorithm simply could

not be used (even on the most powerful of modern computing equipment). In contrast, we have analyzed several such problems on our workstation.

As noted the *test* times do grow exponentially in the number of firms and in our runs they rose from about 3 minutes when M = 5 to over two hours when M = 10. By M = 10the algorithm spends ten times as much time computing the test statistic after each million iterations as it spends on the iterations themselves. So we now outline a modification to the test due Fershtman & Pakes (2005) which circumvents this problem.

A more computationally efficient test. Fershtman & Pakes (2005) provide a test of the equilibrium conditions on the recurrent class whose; (i) computational burden scales linearly in both the number of points in the recurrent class and the number of states at each point (and so is not subject to a curse of dimensionality), and (ii) has an interpretation as a norm in the percentage difference between the actual expected values generated by the estimate of W and the values implied by the Bellman equations (23) and (24).

Say we want to test whether a  $\tilde{W}$  generates equilibrium policies and values on the recurrent class. One way to view what we want a test to do is to measure the difference between the estimate of  $V(\omega_i, \omega_{-i}|W)$  from equations (23) and (24), and the expected discounted values of future net cash flows that each agent would obtain were all agents using the policies generated by  $\tilde{W}$ . Fershtman & Pakes (2005) suggest approximating the expected returns the agents would earn from following the policies generated by W with averages of the discounted value of net cash flows earned from different sample paths that are simulated using the policies generated by  $\tilde{W}$ , and comparing that to the values,  $V(\omega_i, \omega_{-i}|\tilde{W})$ , from equations (23) and (24). The squared differences between the  $V(\omega_i, \omega_{-i}|\tilde{W})$  from these equations and the average of the discounted value over the simulated sample paths is a sum of; (i) the sampling variance in the average of the discounted value of the simulated sample paths (we will refer to this as the sampling variance term), and (ii) the difference between the *expectation* of the discounted net cash flows from the simulated paths and the value functions estimates in the formula in (23) and (24) (we will refer to this as the "bias" term). Fershtman & Pakes (2005) subtract a consistent estimate of the sampling variance term from this squared difference to obtain a test statistic which, at least in the limit, will depend only on the bias term.

More precisely for each state in the recurrent class use the policies generated by W to simulate the algorithm T iterations into the future and keep track of the discounted value of net cash flows generated over those T iterations. Add this sum to the discounted value equation (23) assigns to the terminal state; i.e. if the terminal state is  $(\omega_i^T, \omega_{-i}^T)$  add  $\beta^T V(\omega_i^T, \omega_{-i}^T | \tilde{W})$  as computed from equation (23). Call this sum an estimated value of the simulated path. Compute several independent estimates of the value of the simulated path from each state, and compute the average and variance of those values.

If we let the average of the values of the simulated paths be  $VS(\omega_i, \omega_{-i}|\tilde{W}, T)$ , and E

be the expectation operator over the simulated random draws then

$$E\left(\frac{VS(\omega_i,\omega_{-i}|\tilde{W},T)}{V(\omega_i,\omega_{-i}|\tilde{W})}-1\right)^2 =$$
(27)

$$E\Big(\frac{VS(\omega_i,\omega_{-i}|\tilde{W},T)-E[VS(\omega_i,\omega_{-i}|\tilde{W})]}{V(\omega_i,\omega_{-i}|\tilde{W})}\Big)^2 + \Big(\frac{E[VS(\omega_i,\omega_{-i}|\tilde{W},T)]-V(\omega_i,\omega_{-i}|\tilde{W})}{V(\omega_i,\omega_{-i}|\tilde{W})}\Big)^2.$$

The first term in this expression is the variance term, and the second term is the bias. Moreover the variance term can be unbiasedly estimated from the variance across the simulated sample paths; i.e. we can form  $\hat{Var}\left(\frac{VS(\omega_i,\omega_{-i})}{V(\omega_i,\omega_{-i}|\tilde{W})}\right)$  such that

$$E\Big[\hat{Var}\Big(\frac{VS(\omega_i,\omega_{-i})}{V(\omega_i,\omega_{-i}|\tilde{W})}\Big)\Big] = E\Big(\frac{VS(\omega_i,\omega_{-i}|\tilde{W},T) - E[VS(\omega_i,\omega_{-i}|\tilde{W})]}{V(\omega_i,\omega_{-i}|\tilde{W})}\Big)^2.$$

Consequently

$$E\left[\left(\frac{VS(\omega_i,\omega_{-i}|\tilde{W},T)}{V(\omega_i,\omega_{-i}|\tilde{W})}-1\right)^2-\hat{Var}\left(\frac{VS(\omega_i,\omega_{-i})}{V(\omega_i,\omega_{-i}|\tilde{W})}\right)\right] = \left(\frac{E[VS(\omega_i,\omega_{-i}|\tilde{W},T)]-V(\omega_i,\omega_{-i}|\tilde{W})}{V(\omega_i,\omega_{-i}|\tilde{W})}\right)^2.$$

This provides us with an unbiased estimate of the bias term. Since the variance (and higher order moments) of this estimate of the bias is finite, any weighted average of independent estimates of the bias terms over the recurrent class of points will converge to the same weighted average of the true bias term across these points (a.s.). It is weighted averages of this form that Fershtman & Pakes (2005) use as a test statistic.

**Potential usefulness.** The usefulness of the stochastic algorithm is likely to vary from problem to problem. There are some problems which it cannot be applied to without further adaptation; a good example being problems in which we want to constrain "off the equilibrium path" behavior in a more detailed way (as in many collusion models). Where we can use it, the tradeoff between using it and the other algorithms can be thought of as a tradeoff between number of iterations, number of states, and the time per state at each iteration. The stochastic algorithm can be expected to do particularly well when; (i) the dimension of the recurrent class is small relative to the dimension of the state space, and (ii) the formulae for computing the probability distribution over future possible outcomes is complicated. This last point is particularly important in computing equilibria to models with asymmetric information, as the stochastic algorithm can compute those equilibria without ever actually computing posterior distributions at any given state (and the posteriors are often horribly complex; for more detail see Fershtman & Pakes (2005)).

### 5.4 Function Approximation Methods

Function approximation techniques attempt to approximate either the value or the policy function (or both) by a sufficiently rich set of bases functions (e.g. polynomials or splines). Judd (1998) discusses these methods in some detail in the context of dynamic programming models.

In the context of dynamic games function approximation methods have been used mostly for dynamic games with a continuum of states and without entry and exit. These games pose both theoretical (see, e.g., Whitt 1980) and computational challenges. In particular, they require an entirely different set of techniques to compute an equilibrium since players' values and policies can no longer be represented as (finite-dimensional) vectors and instead have to be represented as (infinite-dimensional) functions. Judd (1998) gives an introduction to these so-called projection techniques and Rui & Miranda (1996, 2001), Doraszelski (2003), and Hall, Royer & Van Audenrode (2003) apply them to dynamic games with a continuum of states.

There has been some attempt to apply these techniques to games with discrete state spaces. Pakes & McGuire (1994) provide one example and some theoretical results on the dimensionality reductions available from these techniques. Briefly they consider a polynomial basis, and project an initial estimate of the value functions at a small number of points down onto this basis to obtain initial coefficients of the basis functions. These coefficients and the basis functions allow one to construct continuation values for each firm operating at the basis points. These continuation values, in turn, can be used to obtain policies and new estimates of the values for all firms active at those points. The new value functions are projected down against the basis functions to obtain new coefficients and the process is repeated.

Pakes & McGuire (1994) show that if one restricts the basis functions to functions which preserve symmetry and anonymity (what they call an "exchangeable" basis), the number of basis functions needed for any order of polynomial, and hence the number of points we need to keep track of, is bounded by a finite constant (i.e. does not grow in  $\bar{n}$  at all). Moreover moment generating functions can be used to reduce the computational burden from computing continuation values at a given point. So the potential for decreasing the computational burden of computing equilibrium for large markets using this technique is dramatic. On the other hand they also report that in their particular example with  $\bar{n} = 6$ there were extensive convergence problems.

The convergence problems were likely to result from the fact that in their model the firm's value function can be a discontinuous function of its competitor's policies (entry and exit by competitors can cause jumps in the value function), and the polynomials they used have problems approximating discontinuities. So these problems may well disappear if either different forms of approximations were used (e.g. splines), or if the number of firms grew larger (in which case exit by any one firm should have a smaller effect on the value of its

competitors). This suggests that function approximation might still be a useful technique for problems with a large number of firms (as if often the case in the dynamic models used in trade, development, and productivity), though such applications might require some further technical developments.

### 5.5 Oblivious Equilibrium

Weintraub et al. (2005) propose an approximation method for analyzing EP-style dynamic models of imperfect competition. They derive an algorithm for computing an "oblivious" equilibrium in which each firm is assumed to make decisions based only on its own state, an aggregate shock, and knowledge of the long run average industry state. When firms chose their controls they ignore the influence of current and past values of the aggregate shock, and current information about rivals' states, on likely future states.

They prove that if the distribution of firms obeys a certain "light-tail" condition, then as the market size becomes large oblivious equilibrium closely approximates a Markov perfect equilibrium in a certain sense. The light tail condition is more likely to be satisfied in industries where there are no firms that are clearly dominant.

They also derive a performance bound that can be used to assess how well the approximation performs in any given applied problem. Finally, Weintraub et al. (2005) apply these methods to a dynamic investment game and find that the approximation typically works well for markets with hundreds of firms, and in some cases works well even for markets with only tens of firms.

## 6 Computing Multiple Equilibria

Below we review the homotopy method to computing multiple equilibria in the EP framework due to Besanko et al. (2004). At the end of this section, we briefly discuss how the problem of multiple equilibria has been handled in other settings.

A drawback of all the computational methods discussed above is that they offer no systematic approach to computing multiple equilibria. Homotopy or path-following methods partially resolve this issue.<sup>25</sup> Starting from a single equilibrium computed for a particular value of the parameter vector, the homotopy algorithm traces out an entire path of equilibria obtained by varying the parameter vector. Whenever we can find such a path and the path folds back on itself, then the homotopy algorithm will identify multiple equilibria. We note at the outset that there is no guarantee that any given path computes all possible

<sup>&</sup>lt;sup>25</sup>Recall that two functions  $f: X \to Y$  and  $g: X \to Y$  from one topological space X to another Y are called homotopic if one can be continuously deformed into the other, i.e., if there exists a continuous function  $H: X \times [0,1]$  such that H(x,0) = g(x) and H(x,1) = f(x) for all  $x \in X$ . Such a deformation is called a homotopy. See Zangwill & Garcia (1981) for an introduction to homotopy methods, Schmedders (1998, 1999) for an application to general equilibrium models with incomplete asset markets, and Berry & Pakes (2006) for an application to estimating demand systems.



Figure 3: Homotopy example. Source: Besanko et al. (2004).

equilibria at a given value of the parameter vector (unless the system of equations that defines the equilibrium happens to be polynomial, see Judd & Schmedders (2004)). Moreover, homotopy methods are computationally demanding.

We begin with a simple example which explains how the homotopy algorithm works. Consider the single nonlinear equation  $H(x, \tau) = 0$ , where

$$H(x,\tau) = -15.289 - \frac{\tau}{1+\tau^4} + 67.500x - 96.923x^2 + 46.154x^3.$$
(28)

Equation (28) implicitly defines the relationship between an endogenous variable x and an exogenous parameter  $\tau$  (in our case  $\tau$  would be one of the parameters of the model). The object of interest is the set of solutions  $H^{-1} = \{(x,\tau) | H(x,\tau) = 0\}$  as graphed in Figure 3. Inspection of Figure 3 shows that there are multiple solutions to equation (28), e.g., at  $\tau = 0.3$  there are three solutions: x = 0.610, x = 0.707, and x = 0.783. Finding these solutions is trivial once the graph is drawn, but producing the graph is less than straightforward even in this very simple case. Whether one solves  $H(x,\tau) = 0$  for x taking  $\tau$  as given or for  $\tau$  taking x as given, the result is a multi-valued correspondence, not a single-valued function.

To apply the homotopy method, we introduce an auxiliary variable s that indexes each point on the graph starting at point A for s = 0 and ending at point D for  $s = \bar{s}$ . The graph is just the parametric path given by a function pair  $(x(s), \tau(s))$  that satisfies  $H(x(s), \tau(s)) = 0$  or, equivalently,  $(x(s), \tau(s)) \in H^{-1}$ . Since there are infinitely many such function pairs, we need a simple way to pick out a member of this family. To do this, we differentiate  $H(x(s), \tau(s)) = 0$  with respect to s to obtain

$$\frac{\partial H(x(s),\tau(s))}{\partial x}x'(s) + \frac{\partial H(x(s),\tau(s))}{\partial \tau}\tau'(s) = 0.$$
(29)

This single differential equation in two unknowns, x'(s) and  $\tau'(s)$ , captures the conditions that are required to remain "on path." One possible approach for tracing out a path in  $H^{-1}$  is thus to solve equation (29) for the ratio  $\frac{x'(s)}{\tau'(s)}$ , i.e., the direction of the next step along the path from s to s + ds. This approach, however, creates difficulties because the ratio may switch from  $+\infty$  to  $-\infty$ , e.g., at point B in Figure 3. So instead of solving for the ratio, we solve for each term of the ratio. To do this, we use the system of differential equations given by

$$x'(s) = \frac{\partial H(x(s), \tau(s))}{\partial \tau}, \qquad (30)$$

$$\tau'(s) = -\frac{\partial H(x(s), \tau(s))}{\partial x}$$
(31)

to determine the next step along the path. Now we have avoided the problem of dividing by zero when the path folds back on itself.

These are the so-called basic differential equations for our simple example. Their significance is that they reduce the task of tracing out the set of solutions to solving a system of differential equations. Given an initial condition this is can be done using a variety of methods (see, e.g., Chapter 10 of Judd 1998). In our example, note that if  $\tau = 0$ , then  $H(x,\tau) = 0$  is easily solved for x = 0.5. This provides the initial condition (point A in Figure 3). From there the homotopy algorithm follows the path until it reaches  $\tau = 1$ (point D). Whenever  $\tau'(s)$  switches sign from positive to negative (point B), the path is bending backward and there are multiple solutions. Conversely, whenever the sign of  $\tau'(s)$ switches back from negative to positive (point C), the path is bending forward.<sup>26</sup>

Besanko et al. (2004) apply the homotopy method to computing equilibria in their model of learning-by-doing and organizational forgetting. They write the system of nonlinear equations that defines an equilibrium as  $H(x, \tau) = 0$ . The system contains the Bellman equation and optimality conditions for each firm's policy at each state in  $S^{\circ}$ .  $\tau$  is a parameter of the model; in their case the rate of depreciation  $\delta$ . The object of interest is the set of equilibria  $H^{-1} = \{(x, \tau) | H(x, \tau) = 0\}.$ 

<sup>&</sup>lt;sup>26</sup>In computing equilibria to their learning-by-doing model Besanko et al. (2004) prove that on the portion of the path that is backward bending (in Figure 3 this is the part of the path between points B and C) the mapping that provides the iterations for the Pakes & McGuire (1994) algorithm has an eigenvalue outside of the complex unit circle and is therefore unstable. It follows that, holding fixed the homotopy parameter, "inbetween" two equilibria with  $\tau'(s) > 0$ , there is one equilibrium with  $\tau'(s) \leq 0$  that cannot be computed using the Pakes & McGuire (1994) algorithm.

Proceeding as in our simple example, they define a parametric path to be a set of functions  $y(s) = (x(s), \tau(s))$  such that  $y(s) \in H^{-1}$ . Here x(s) is the vector of value functions, policies, and exit or entry probabilities associated with each firm at each state. As s varies y(s) describes a path in  $H^{-1}$ , that satisfies

$$H(y(s)) = 0.$$

Differentiating both sides with respect to s yields the conditions that are required to remain on path

$$\sum_{i} \frac{\partial H(y(s))}{\partial y_i} y'_i(s) = 0.$$
(32)

This is a system differential equations in the same number of unknowns as number of equations.

As before there is at least one solution and it obeys the basic differential equations

$$y'_{i}(s) = (-1)^{i+1} \det\left(\left(\frac{\partial H(y(s))}{\partial y}\right)_{-i}\right)$$
(33)

for all *i*, where the notation  $(\cdot)_{-i}$  is used to indicate that the *i*th column is removed from the Jacobian  $\frac{\partial H(y(s))}{\partial y}$  (see pp. 27-28 in Zangwill & Garcia (1981) for the proof). Note that equation (33) reduces to equations (30) and (31) if *x* is a scalar instead of a vector. Given an initial condition the basic differential equations (33) can be solved using numerical methods.

A few comments are in order. First, if the Jacobian has less than full rank, then the determinants of all its submatrices are zero. Thus,  $y'_i(s) = 0$  for all *i* and the homotopy algorithm stalls. Indeed, a central condition in the mathematical literature on homotopy methods is that the Jacobian has full rank (see, e.g., Zangwill & Garcia 1981). If so, the homotopy is called regular and the algorithm is guaranteed to trace out a path.<sup>27</sup> Unfortunately, proving that the Jacobian has full rank is frequently difficult when the goal is to follow an economically meaningful parameter such as the rate of depreciation  $\delta$  or the discount factor  $\beta$ .

Second,  $H(\cdot)$  must be continuously differentiable. This means that we need the distribution of scrap values  $F(\cdot)$  and setup costs  $F^e(\cdot)$  to be differentiable. Also the investment decisions which, since zero investment is a possible outcome, are typically characterized by Kuhn-Tucker conditions should be reformulated as continuously differentiable equations (Judd, Kübler & Schmedders (2003) show how to do this).<sup>28</sup>

Finally we note that the homotopy algorithm is computationally demanding. Comput-

<sup>&</sup>lt;sup>27</sup>The mathematical literature calls this the continuation problems and has developed methods for dealing with the bifurcations that may arise if the Jacobian has less than full rank (see, e.g., Tesfatsion & Kabala (1991)).

<sup>&</sup>lt;sup>28</sup>The mathematical literature has developed piecewise linear homotopy methods that can handle problems with discontinuous equations (see, e.g., Eaves (1972) and Chapter 5 of Judd (1998)).

ing the Jacobian is a major burden. Numeric derivatives are slow to compute and though analytic derivatives speed up the computations, deriving and coding them by hand is time consuming and error prone. One alternative is to use automatic differentiation software that takes a procedure for computing a function as input and produces a procedure for computing its derivatives as output. However even with analytic derivatives the computational demands of homotopy algorithm that have been used to date are such they have been restricted to relatively small systems of nonlinear equations. On the other hand there has been very little exploration of methods which might reduce its computational burden. Some initial experience suggests that exploiting sparsity patterns in the Jacobian yields substantial gains.

**Software.** A number of software packages implement homotopy methods. Besanko et al.'s (2004) programs are based on Hompack (Watson, Billups & Morgan 1987), which is written in Fortran 77, and are available upon request. Hompack 90 is able to exploit sparsity patterns and is described in Watson, Sosonkina, Melville & Morgan (1997). There are numerous software packages available for automatic differentiation, e.g., ADIFOR for Fortran 77 and ADIC for C.

**Summary.** The problem of finding all solutions to a system of nonlinear equations is largely unresolved in the mathematics literature. Indeed, as already noted, there is no guarantee that the homotopy algorithm finds all the equilibria of our dynamic stochastic game. In some cases it is possible to exploit the structure of the system of equations. For example, the system of equations that characterizes the set of Nash equilibria in static games is polynomial (see, e.g., McKelvey & McLennan 1996). For polynomial systems, in turn, there are methods that are guaranteed to find all solutions. These all-solutions homotopies have been implemented in the freely-available software package Gambit (McKelvey, McLennan & Turocy 2006) and used by Bajari, Hong & Ryan (2004) in the context of static games and by Judd & Schmedders (2004) to construct a computational uniqueness proof for dynamic games in which movements through the state space are unidirectional. Jia (2006) imposes supermodularity conditions on a two-stage game and uses the fact that in supermodular games there is a best and a worst equilibrium to delineate the set of subgame perfect equilibria.

Finally, there are methods for computing the entire set of subgame perfect equilibria in repeated games (see, e.g., Judd, Yeltekin & Conklin 2003). Recently, these methods have been extended to game with state variables (Conklin & Judd 1996, Judd & Yeltekin 2001). Unfortunately, however, none of the available methods is anywhere near capable of handling the large games that are typical for applications of EP's framework.

## 7 Applications and Extensions

There is a large and active literature using the EP framework in IO, and, most recently, it has also been used in other fields such as international trade (Erdem & Tybout 2003) and finance (Goettler, Parlour & Rajan 2004, Kadyrzhanova 2006). As Table 4 shows, there are too many applications to discuss all of them in detail. Instead we focus on some of the larger themes in this literature. They allow us to both; (i) illustrate how the framework has been extended by various authors to approximate alternative institutional structures, and (ii) point out some of the problems that need to be overcome before the framework can be applied to particular topics of obvious interest.

We begin with a brief review of the empirical work using the EP framework. This includes a small literature on testing the EP model against known alternatives, and a small but growing literature on using the framework to analyze the impacts of different policy or environmental changes. Notably absent from this subsection is a review of the recent literature which uses the EP framework to structure estimation algorithms. The relevant literature here includes Aguirregabiria & Mira (2006), Bajari, Benkard & Levin (2006), Pakes, Ostrovsky & Berry (2006), Pesendorfer & Schmidt-Dengler (2003) and has been recently reviewed in Ackerberg, Benkard, Berry & Pakes (2005).

We then turn to applied theory papers which use the framework. Our starting point are models of capacity and advertising dynamics. These models preserve the traditional "static-dynamic" breakdown we used to introduce the framework in prior sections. They illustrate nicely just how rich the industry dynamics generated by the framework can be even in this "textbook" setting. This fact is accentuated by comparing the results using the EP framework to other—typically analytically tractable—dynamic models. Next we discuss models of mergers. The EP framework in its basic form is well-suited to analyze the effect of a "one-time" exogenously specified merger on the future evolution of an industry. The papers which do this analysis emphasize the importance of explicitly modeling how a merger affects firms' incentives for investment, entry, and exit. By taking into account how these incentives evolve after a merger has taken place, these papers often arrive at welfare implications that are quite different from those obtained in simpler models. There has also been some work on endogenizing merger activity. This requires a much richer stage game to be played in each period; firms not only choose quanities/prices but also whether and with whom to merge.

Next we turn to models of learning-by-doing, network effects, and collusion. These types of models depart from the traditional "static-dynamic" breakdown in teaching IO., i.e. in these models the price or quantity that a firm sets in the product market has a direct effect on the dynamics of the industry. Consequently, the profit function can no longer be computed "off line." After providing a brief review of how to modify the framework to allow for this possibility we review some of the results from this literature. Again the emphasis is on dynamic effects that have not been captured by simpler models. As we shall see they often shed new light on both policy issues and our understanding of the implications of different environments.

application	$\operatorname{paper}(\mathrm{s})$
advertising	Doraszelski & Markovich (2006), Dube, Hitsch
	& Manchanda (2005)
capacity accumulation	Besanko & Doraszelski (2004), Ryan (2005),
	Beresteanu & Ellickson (2005)
collusion	Fershtman & Pakes (2000, 2005), de Roos (2004)
competitive convergence	Langohr (2003)
consumer learning	Ching (2003)
firm size and growth	Laincz & Rodrigues (2004)
learning-by-doing	Benkard $(2004)$ , Besanko et al. $(2004)$
mergers	Berry & Pakes (1993), Gowrisankaran (1999),
	Chen (2004)
network effects	Jenkins, Liu, Matzkin & McFadden (2004),
	Markovich (2004), Markovich & Moenius
	(2005), Chen, Doraszelski & Harrington (2004)
productivity growth	Laincz $(2005)$
R&D	Gowrisankaran & Town (1997), Auerswald
	(2001), Song (2002), Judd et al. (2002), Fersht-
	man & Markovich (2006)
technology adoption	Schivardi & Schneider (2005)
international trade	Erdem & Tybout (2003)
finance	Goettler et al. (2004), Kadyrzhanova (2006)

Table 4: Applications of EP's framework.

### 7.1 Empirics

First a brief note on testing. Pakes & Ericson (1998) show that the Markov nature of the equilibrium process distinguishes the EP model from industry equilibrium models based on learning about a time-invariant parameter, such as the dynamic equilibrium model of Jovanovic (1982). They formalize this distinction by developing a non-parametric test, versions of which have been used to determine the appropriateness of the different types of models for different data sets (see, e.g., Bhattacharjee (2005), Klette & Raknerud (2002), and Abbring & Campbell (2003)).

The test assumes there is an observable variable which orders the  $\omega_i$  of active firms (in their application the variable used was sales), and follows the values of this variable in cohorts of surviving firms over time. If the EP framework is appropriate, the dependence of the conditional distribution of current size, conditional on a few past sizes, on the initial size should disappear as we increase the age of the firm (formally a  $\phi$ -mixing condition is satisfied). If the learning model is appropriate that dependence never disappears. Interestingly Pakes & Ericson (1998) find a sharp distinction between manufacturing and retail firms in their Wisconsin data set; it is clear that manufacturing firms satisfy the phi-mixing condition while retail firms do not. That is at least in their data the EP model seems more appropriate for manufacturing firms while a model based on learning about time invariant parameters seems more appropriate for retail trade.

In addition, a number of papers apply variants of the EP framework to analyze policy or environmental changes in a variety of industries. In the first effort along these lines, Gowrisankaran & Town (1997) present a dynamic model of the hospital industry in which nonprofit and for-profit hospitals coexist and compete. They are differentiated by their objective functions, investment technologies, and taxation rates. In their model patients, who differ by income and type of insurance coverage, chose admission to their preferred hospital. Hospitals choose investment, entry, exit, and pricing strategies. Gowrisankaran & Town (1997) first estimate the parameters of the model and then use it to examine the effects of changes in the Medicare reimbursement system, universal health-care coverage, and taxation of nonprofits. The results were striking; the effect of the policy changes on entry, investment and exit policies often implied that the original goals of the policy changes would not only not be achieved, but in some cases the dynamic implications would quickly result in an environment that was worse than the original situation with respect to the stated policy goal.

Ryan (2005) evaluates the welfare costs of the 1990 Amendments to the Clean Air Act on the US Portland cement industry. The typical cost analysis of an environmental regulation consists of an engineering estimate of the compliance costs. In industries where fixed costs are an important determinant of market structure this static analysis ignores the dynamic effects of the regulation on entry, investment, and market power. Ryan (2005) accounts for these effects through a dynamic model of oligopoly in the tradition of EP. He recovers the entire cost structure of the industry, including the distribution of sunk entry costs and adjustment costs of investment, and finds that the Amendments have significantly increased the sunk cost of entry. Simulating the welfare effects of the Amendments, he shows that a static analysis misses the welfare penalty on consumers, and obtains the wrong sign on the welfare effects on incumbent firms.

The EP framework has also been applied to particular industries in order to shed light on the institutional arrangements in those industries. Song (2002) analyzes a research joint venture (RJV) led by SEMATECH in the semiconductor industry using the dynamic model of oligopoly. The paper first estimates firms' states, using product level data, and solves for the equilibrium research expenditures. Results show that a firm's research expenditure in SEMATECH is one fifth of what it would be in competitive research. Lower research expenditures result in higher net profits in RJVs, although variable profits are similar in both regimes. RJVs are also more likely to generate higher consumer surplus than competitive research. The paper also shows that firms react differently for the same changes in the product market, depending on whether they cooperate or compete in research. Beresteanu & Ellickson (2005) examine competition between supermarket chains using a dynamic model of strategic investment. Employing a eight year panel dataset of store level observations that includes every supermarket operating in the United States, they develop and estimate a fully dynamic model of chain level competition. The estimated parameters of the structural model are then used to evaluate the competitive impact from the introduction of superstores.

There are a number of other papers which use variants of the framework to analyze particular issues in different industries, (Benkard 2004, Ching 2003, de Roos 2004, Dube et al. 2005, Jenkins et al. 2004), some of which we will return to after we introduce more complex variants of the EP framework. Overall these papers have shown both that; (i) the framework is at least reasonably well-suited for applied dynamic analysis in a variety of industries, and (ii) that the dynamic effects of policy changes should not be ignored and often outweigh (sometimes reversing) the static effects.

### 7.2 Capacity and Advertising Dynamics

Empirical evidence suggests that there are substantial and persistent differences in the sizes of firms in most industries. For example, Mueller (1986) found that in 44 percent of 350 U.S. manufacturing industries, the identity of the industry leader remained unchanged over a twenty-two year period, and that the correlation between market shares over this period was 0.66 (see also Gort (1963) among others).

Asymmetric industry structures can, of course, arise as the outcome of a game in which firms differ in their economic fundamentals (e.g., cost structures) or their strategic positions at the outset of the game (e.g., first versus later mover). However, this begs the question: How do such differences in initial conditions arise in the first place? A number of papers has shown that asymmetric industry structures can arise as the outcome of a capacity accumulation game played by ex ante identical firms (e.g., Saloner (1987), Maggi (1996), and Reynolds & Wilson (2000)). Though valuable in highlighting that substantial differences in firm size can arise endogenously for strategic reasons, these models are less satisfactory in explaining the persistence of these differences over time which is highlighted by the evidence. This is because these papers consider an unchanging competitive environment. Once an equilibrium has been reached, nothing further happens to upset the positions of firms. That is, asymmetric industry structures persist in these models by default. Ideally, however, one would like to understand whether there are circumstances under which asymmetric industry structures persist in a competitive environment that changes over time, for example, because of firm-specific shocks. The competitive environment can also change due to feedback effects. For example, if large firms invest more aggressively than small firms, then a small initial asymmetry may become larger over time, but it may vanish otherwise. In general, one would expect feedback effects to play a role whenever the time horizon under consideration is long enough to allow firms to interact repeatedly.

For these reasons, an important question is when substantial and persistent size differences can arise endogenously in equilibrium in a market in which *ex ante* identical firms interact repeatedly and are subject to firm-specific shocks that continuously alter their positions. To address this question, Besanko & Doraszelski (2004) use the EP framework to track the evolution of an oligopolistic industry. Consistent with the empirical evidence, their dynamic model of capacity accumulation is able to explain the substantial and persistent differences in the sizes of firms.

Besanko & Doraszelski (2004) demonstrate that industry dynamics depend critically on the mode of product market competition and on the degree to which investment is reversible, as measured by the rate of depreciation. Under quantity competition, each firm accumulates enough capacity to supply the Cournot quantities, leading to an industry structure of equalsized firms independent of whether investment is irreversible (zero depreciation) or reversible (positive depreciation). With positive depreciation, firms tend to hold idle capacity out of a precautionary motive. By contrast, under price competition, there are forces that propel the industry towards asymmetric structures. In particular, if investment is reversible, then the industry evolves towards an outcome with one dominant firm and one small firm. Industry dynamics in this latter case resemble a rather brutal preemption race. During this race, firms invest heavily as long as they are of equal size even though this leads to substantial industry-wide overcapacity. Once one of the firms manages to pull slightly ahead in this race, however, the smaller firm "gives up", thereby propelling the larger firm eventually into a position of dominance.

Their paper sheds new light on the relationship between preemption and reversibility. Besanko & Doraszelski (2004) show that, under price competition, the preemption race between contending firms becomes more brutal as investment becomes more reversible. This stands in marked contrast to the usual intuition that depreciation reduces the commitment power of capacity and that capacity accumulation can therefore only lead to a temporary advantage. The key insight underlying their result is that with reversible investment the consequences to a firm of falling behind its rival are not fatal: the firm can allow its capacity to depreciate and assume the more profitable posture of a "puppy dog." The higher is the rate of depreciation, the easier it is for a lagging firm to "disengage" from the preemption race and hence the more attractive it is for firms to enter in such a race in the first place.

Laincz & Rodrigues (2004) further pursue the idea of using the EP framework to explain stylized facts about industry structure and dynamics. To this end, they develop a dynamic model of the firm size distribution. Their model allows for continually falling marginal costs through process R&D (as in Laincz (2005)). Empirical studies of the firm size distribution often compare the moments to a log-normal distribution as implied by Gibrat's Law and note important deviations. Thus, the first, and basic questions they ask are how well does the dynamic industry model reproduce Gibrat's Law and how well does it match the deviations uncovered in the literature. They show that the model reproduces these results when testing the simulated output using the techniques of the empirical literature. They then use the model to study how structural parameters affect the firm size distribution. They find that, among other things, fixed and sunk costs increase both the mean and variance of the firm size distribution while generally decreasing the skewness and kurtosis. The rate of growth in an industry also raises the mean and variance, but has non-monotonic effects on the higher moments.

Turning from capacity to advertising dynamics, it is worth noting that during 2003 close to 250 billion dollar was spent on advertising in the U.S., well above 2% of GDP. Practitioners know very well the value of advertising to achieving their long-term market share and profitability goals and presume that advertising is capable of giving them a sustainable competitive advantage over their rivals. The existing dynamic models of advertising competition, however, suggest quite the opposite. In these models there is a globally stable symmetric steady state (see e.g. Friedman 1983, Fershtman 1984, Chintagunta 1993, Cellini & Lambertini 2003). Consequently, any differences among firms are bound to vanish over time, and there is no room for a sustainable competitive advantage, not even if firms enter the market one by one and thus differ in their strategic positions at the outset of the game (Fershtman, Mahajan & Muller 1990).

Doraszelski & Markovich (2006) attempt to reconcile theory and observation by showing that advertising can indeed have a lasting effect on the structure of an industry. To this end, they propose a dynamic model of advertising competition based on the EP framework to track the evolution of an industry. Within this dynamic framework, they study two different models of advertising: In the first model, advertising influences the goodwill consumers extend towards a firm ("goodwill advertising"), whereas in the second model it influences the share of consumers who are aware of the firm ("awareness advertising"). They show that asymmetries may arise and persist under goodwill as well as awareness advertising. The basis for a strategic advantage, however, differs greatly in the two models of advertising.

Under goodwill advertising, the size of the market and the cost of advertising are key determinants of industry structure and dynamics. In particular, goodwill advertising leads to an extremely asymmetric industry structure with a large and a small firm if the market is small or if advertising is expensive. Because the marginal benefit of advertising is small relative to its cost, a small firm has only a weak incentive to advertise when competing against a large firm and, in fact, may choose not to advertise at all. If the market is large or if advertising is cheap, on the other hand, even a small firm has a fairly strong incentive to advertise. In this case we obtain a symmetric industry structure with two large firms.

In contrast to the cost/benefit considerations that give rise to a strategic advantage under goodwill advertising, whether or not asymmetries arise and persist under awareness advertising depends on the intensity of product market competition. If competition is soft, the industry evolves towards a symmetric structure, but it evolves towards an asymmetric structure if competition is fierce. Industry dynamics in this latter case resemble a preemption race. In this race, both firms start off advertising heavily and continue to do so as long as they are neck-and-neck. Once one firm gains a slight edge over its competitor, however, there is a marked change in advertising activity. While the smaller firm scales back the larger firm ratchets up its advertising, thus eventually securing itself a position of dominance. The ensuing asymmetric industry structure persists because it is in the self-interest of the smaller firm to stay behind. In fact, the nature of product market competition is such that once the smaller firm tries to grow, the larger firm responds aggressively by triggering a "price war," thereby pushing prices and hence profits down. This gives the smaller firm an overwhelming incentive to remain inconspicuous. In sum, the central idea of their model of awareness advertising is that "more is less." This is a rationale for persistent asymmetries that has mostly been ignored in the literature on dynamic games.

The benefits from improving upon earlier work are most obvious in comparison to linearquadratic games (Friedman 1983, Cellini & Lambertini 2003). Since the law of motion in such a game is given by a system of linear difference (or differential) equation, the dynamics are generically either explosive and thus inconsistent with equilibrium or there is a globally stable symmetric steady state. Hence, the very nature of a linear-quadratic game goes against the notion of a sustainable competitive advantage. Taken together, Doraszelski & Markovich's (2006) departures from earlier work lead to a model of advertising competition that exhibits much richer dynamics. On the other hand, they force them to leave analytically tractable modeling frameworks such as linear-quadratic games behind.

Dube et al. (2005) use the EP framework to study a higher-frequency feature of firms' advertising policies in more detail. More specifically, they develop a model of dynamic advertising competition in order to explain "pulsing", a widely observed phenomenon in many consumer-goods industries whereby firms systematically switch advertising on and off at a high-frequency. Hence, we observe periods of zero and non-zero advertising, as opposed to a steady level of positive advertising. Using an estimated demand system for the frozen entree product category, they verify whether the use of pulsing can be justified as an equilibrium advertising practice. They find evidence for a threshold effect, which is qualitatively similar to the S-shaped advertising response that the theoretical literature has put forth as an explanation for pulsing. Their estimates imply that firms should indeed pulse in equilibrium. Predicted advertising in equilibrium is higher, on average, than observed advertising. On average, the optimal advertising policies yield a moderate profit improvement over the profits under observed advertising.

### 7.3 Mergers

Almost all of the formal models of merger activity condition on the cost, qualities, and variety of products sold in the market. These models hold the distribution of characteristics of the products being marketed (as well as the nature of competition) fixed, and analyze the impact of the ownership change on producer and consumer surplus. The producer surplus analysis provides a vehicle for analyzing the incentives to merge, while either the consumer surplus or the sum of the two provides the vehicle for analyzing whether the merger might be socially beneficial.

As noted in Stigler's (1968) investigation of the US Steel mergers, the results from such a "static" analysis of mergers can easily be overturned by simple dynamic considerations (his discussion allowed for adjustment costs in an analysis of mergers in a homogenous homogeneous goods industry). The first attempts to build a model to analyze the dynamic effects of mergers are Berry & Pakes (1993) and Cheong & Judd (2006). Both papers analyze the impact of a "one-time" exogenously specified merger in a dynamic model which allows for investment but does not allow for any further mergers. These papers show that mergers can be beneficial to *both* the firms merging and to society, even if the profits of the merging firms *and* consumer surplus falls at the time of the merger. The predominant reason is that there is less of an incentive to invest in the merged industry, and the Markov perfect equilibrium generates more investment than a social planner would (see Mankiw & Whinston (1986) for the intuition underlying these arguments).

Chen (2004) extends the dynamic model of capacity accumulation due to Besanko & Doraszelski (2004) in which firms produce nearly homogeneous products and compete in prices to investigates the price and welfare effects of mergers. In contrast to the results stressed in the results cited in the last paragraph, Chen (2004) finds that mergers reduce welfare. This negative effect results from the fact that certain firms in the postmerger industry optimally choose to let their capacities shrink, resulting in higher prices and lower consumer surplus. This divergence in results emphasizes both just how rich a set of outcomes the framework can generate, and the importance of programming an appropriate institutional structure into the model before coming to any policy conclusions.

To be more realistic a model which investigated the dynamic impacts of mergers would want to allow mergers to arise endogenously, and not just investigate the impacts of a "onetime" exogenously specified merger. In particular one might think that a merger by one firm might lead to further mergers of competitors, as it has often been noted that there tend to be "merger waves". There are many unsolved problems here, not least among them being the diversity of views on the factors motivating merger activity in different industries. In addition to specifying the possible sources of gains from mergers, a merger model must also specify a market mechanism for determining which among the possible profitable mergers at any point of time are in fact consummated.

One such mechanism is provided in Gowrisankaran (1999). He takes the capacityconstrained quantity-setting game with homogeneous goods and adds to it a merger game. The merger game occurs at the beginning of each period and proceeds in the following sequential manner. The largest firm is allowed to choose a merger partner first. All other firms present the largest firm with a "take it or leave it" price at which they are willing to be bought. Information is symmetric except that the largest firm draws a "synergy" value for each merger which is known only to it; i.e. the synergy value for a given firm is not known to any of the firms that might be acquired. The largest firm chooses to merge with the firm which generates the highest net merger value provided that value is positive. The net merger value is the expected discounted value of future net cash flow if a merger would take place net of the price of the acquisition and what the value of the firm would be if the merger did not take place. If a merger takes place the process restarts (there are new take it or leave if offers, and new synergy values), and the (new) largest firm can choose another merger partner. When the largest firm chooses not to merge further, the second largest firm gets to choose a merger partner in the same way. This process continues until the smallest active firm chooses not to merge further. At that point production, investment, and then exit followed by entry and investment decisions are made. All offers and actions are made to maximize the expected discounted value of future net cash flows given the agents' information sets, and the equilibrium is Markov perfect.

Perhaps the most striking part of this analysis is that it can be done. Given the quantitative magnitude of the merger phenomena in recent years and the extent that it can be impacted by policy, any step in developing a usable models of mergers is welcome. Still, it is clear that we are only at the beginnings of developing realistic dynamic models that allow for mergers; a lot of work remains to be done.

#### 7.4 Learning-by-Doing and Network Effects

In the simple special case of the framework described above the distribution of the next period's state, conditional on today's states and firms' investment, entry, and exit decisions, is independent of the prices or quantities that the firms set in the product market. This is the assumption which allows us to compute the profit function "off line" and to study static equilibrium without considering the impact of pricing or quantity decisions on future profits. The assumption is, however, inappropriate whenever either; (i) future cost (ii) future demand or (iii) the choice of equilibrium prices or quantities in future periods, depends on the prices or the quantities set in the current period. Relevant cases when this assumption is inappropriate include models with learning-by-doing, adjustment costs, durable goods, experience goods, collusion, and network effects.

In all these cases the profit function can no longer be computed "off line" and fed into the algorithm for computing the equilibrium of the dynamic stochastic game. The reason is that the choice rule for prices or quantities will now have to account for the impact of the choice on future cash flows as well as on current profits. As a result we cannot solve for prices or quantities without knowing the value function, and this makes for a more difficult computational problem. The collusion case has a slightly different structure then the others and we discuss it in a separate section below. Here we consider only the cases where the price or quantity set today have an independent effect on future demand or cost conditions.

The modification to the first-order condition needed to accommodate these cases differs

depending on whether the firm's control is price or quantity. To see this consider a setting with learning-by-doing. Let  $\omega_i$  be the stock of experience or know-how of firm *i* that determines its cost of production. The law of motion for firm *i*'s state is

$$\omega_i' = \omega_i + \nu_i - \eta_i$$

where the random variable  $\nu_i$  indicates additions to the stock of experience and the random variable  $\eta_i$  represents firm-specific depreciation of experience (also called organizational forgetting in the literature). The distribution of  $\nu_i$  is stochastically increasing in the sales  $q_i$  of firm *i*.

In games of quantity competition a firm can change its control without directly affecting the quantities, and therefore state-to-state transitions, of its competitors. Then the Nash first-order condition only involves the derivative of the firm's current profit and transition probabilities for its own state with respect to the control. In contrast, in games of price competition any change in the firm's control induces a change in the quantities, and therefore in the probabilities of the state-to-state transitions, of the firm's competitors. This adds further derivatives to the Nash first-order condition.<sup>29</sup> Benkard (2004) analyzes a learningby-doing model with quantity competition and since his work has the added realism of a model built up from estimated parameters, we begin with an outline of it. Later we discuss the model of learning-by-doing and organizational forgetting of Besanko et al. (2004) as an example of a price setting game.

Benkard's (2004) goal is to analyze competition in the market for wide bodied commercial aircraft. Benkard adds estimates of demand parameters to the cost functions he estimates in Benkard (2000), and then computes and analyzes a model of dynamic quantity competition among producers which is a reasonably realistic approximation to the competition that occurred in the commercial aircraft market at the time the Lockheed Tristar was being introduced. Benkard (2004) finds that the effect of current quantity on future costs, and through future costs on its future competitive status, will induce the firm to produce large quantities in the early production years (the experience curve is steep early on in the production process). In fact production in the early years is pushed so far that price falls well below marginal cost in those years. This implication is clearly borne out by the price and cost data, and is inconsistent with a static quantity setting model. Benkard (2004) then proceeds to an analysis of the producer and consumer surplus generated by the outcomes of the interactions in this market, and a series of counterfactuals allow him to analyze what would have been likely to happen if we had imposed other institutional constraints.

As pointed out by Benkard (2004), organizational forgetting, i.e., depreciation of experience, is needed to explain the dynamics in the market for wide-bodied airframes in the 1970s and 1980s. It is often said that learning-by-doing promotes market dominance because it

 $<sup>^{29}</sup>$ This also happens when the control is a bid in a repeated auction with capacity constraints, see Jofre-Bonet & Pesendorfer (2003).

gives a more experienced firm the ability to profitably underprice its less experienced rival. However if learning-by-doing can be "undone" by organizational forgetting, then there is a question of whether organizational forgetting can reverse the market dominance effects of learning-by-doing. Besanko et al. (2004) find that this is not necessarily the case; on the contrary, over a wide range of parameterizations, organizational forgetting tends to make firms more instead of less aggressive. This aggressive pricing behavior, in turn, puts the industry on a path towards market dominance. They extend Cabral & Riordan's (1994) model of learning-by-doing by allowing for organizational forgetting. In the absence of organizational forgetting, the price that a firm sets reflects two goals. First, by winning a sale, the firm moves down its learning curve. Second, the firm prevents its rival from moving down its learning curve. Organizational forgetting accentuates these possibilities; by winning a sale, a firm makes itself less vulnerable to future losses from organizational forgetting and, at the same time, it makes its rival more vulnerable. This creates strong incentives to cut prices.

Incorporating organizational forgetting in the model as suggested by the empirical studies of Argote, Beckman & Epple (1990), Darr, Argote & Epple (1995), Benkard (2000), and Thompson (2003) into the Cabral & Riordan (1994) model of learning-by-doing makes the model much less analytically tractable. As a result Besanko et al. (2004) move from the elegance of the analytic results in Cabral & Riordan (1994) to numerical analysis. On the other hand, adding organizational forgetting to a model of learning-by-doing leads to a rich array of pricing behaviors and industry dynamics that the existing literature never considered.

In particular, Besanko et al. (2004) show that the model with both learning-by-doing and organizational forgetting can give rise to multiple equilibria, whereas a model with learning-by-doing alone cannot. They show that with their parameterization these equilibria range from "peaceful coexistence" to "trench warfare." If the inflow of know-how into the industry due to learning-by-doing is substantially smaller than the outflow of know-how due to organizational forgetting, then it is virtually impossible that both firms reach the bottom of their learning curves. Conversely, if the inflow is substantially greater than the outflow, then it is virtually inevitable that they do. In both cases, the primitives of the model tie down the equilibrium. This is no longer the case if the inflow roughly balances the outflow, and the stage is set for multiple equilibria. If firms believe that they cannot profitably coexist at the bottom of their learning curves and that instead one firm comes to dominate the market, then both firms cut their prices in the hope of acquiring a competitive advantage early on and maintaining it throughout. This aggressive pricing behavior, in turn, leads to market dominance. However, if firms believe that they can profitably coexist, then neither firm cuts its price, thereby ensuring that the anticipated symmetric industry structure actually emerges. Consequently, in addition to the degree of organizational forgetting, the equilibrium by itself is an important determinant of pricing behavior and industry dynamics.

Just as learning-by-doing can lead to decreasing cost over time, network effects can lead to increasing utility over time. Markovich (2004) and Markovich & Moenius (2005) model the dynamics caused by the interactions between hardware and software choices; that is by "indirect" network effects. Consumers make a hardware choice that lasts two periods. Software is designed to run on one, and only one, of the two types of hardware. Software firms must commit to one of the two types of hardware when they enter, and then can invest to improve the quality of their product, or exit, just as in the core version of our model. The demand for a given software product depends not only on the qualities of software products available for each of the two hardware types, but also on the number of consumers who have purchased the different types of hardware in the past. Thus consumer demand for hardware products depends on beliefs about the likelihood of future software products available for each hardware type, while the entry exit and investment decisions of software firm's depends on their beliefs on the future hardware purchases of consumers.

Markovich (2004) shows that with her parameterization excess inertia does not occur. A platform would be the standard in the market only if it is better than the competing platforms. Furthermore, if the industry's competitors (the outside alternative) are not progressing too quickly the equilibrium is one where both types of hardware are produced (the "variety" equilibrium), while if the competitors to the industry are growing quickly we see an equilibrium with only a single type of hardware produced (the "standardization" equilibrium). Markovich & Moenius (2005) find that network effects tie together the performance of firms using the same platform. This is driven by two effects: a successful competitor increases its platform's market share, and that in turn increases the incentives to invest in quality for all firms on this platform. So a firm may even enjoy a windfall increase in its market value resulting when a competitor on the same platform has a success.

Jenkins et al. (2004) apply a dynamic model with network externalities to a stylized description of the browser war between Netscape and Microsoft (see also Hall et al. 2003). They show that network effects may be a substantial barrier to entry, giving both entrants and incumbents powerful strategic incentives to tip the market. An incumbent that dominates a market through share and scope has an incentive to maintain share as a barrier that makes entry costly and risky. An entrant has an incentive to grab market share to overcome the incumbent's advantage. In the browser war between Netscape and Microsoft, network effects appear to have sharpened strategic incentives and driven a "no-holds-barred" battle for market share. Counterfactual experiments that compare "as is" market trajectories with "but for" trajectories suggest Microsoft's "bad acts" may have been decisive in tipping the market.

### 7.5 Collusion

Most of the theoretical work on collusion and price wars assumes a fixed or exogenously changing environment. Though these assumptions help clarify what determines when a collusive agreement can be enforced and hence when it breaks down (see, in particular, the classic work of Green & Porter (1984) and Abreu, Pearce & Stacchetti (1986)), they limit the investigation of the implications of collusion to its impact on prices; ignoring the (possibly equally important) effects of collusion on the costs, qualities, and varieties of the products marketed.

If we are willing to give up on the elegance of analytic results and rely instead on numerical analysis, it is not difficult to analyze collusive models that allow for heterogeneity among firms who invest to develop their products, and can enter and exit. Fershtman & Pakes (2000) assume that firms either collude to set prices or set prices as in a static Nash pricing equilibrium. Collusive prices and profits are determined by the outcomes of a Nash bargaining game in which the threat value is the profits from the static, Nash in prices, equilibrium. The choice of which price vector to play depends on whether any incumbent has deviated from collusive prices in the past, and on whether the punishments currently available are sufficient to insure no firm has an incentive to deviate in the current period. If there is an incumbent who has deviated the static Nash in price solution is played as long as that incumbent remains active. As in much of the repeated game literature no incumbent ever deviates.<sup>30</sup> However there are tuples of states for which the punishment of reverting to non-collusive prices is not sufficient to support collusion, and this generates "price wars" (reversions to a Nash pricing equilibrium).

Fershtman & Pakes (2000) find that collusion is hard to sustain when either one of the firms does not keep up with the advances of its competitors, or a "low quality" entrant enters. In either case there will be an active firm that is quite likely to exit in the near future. Not only is it hard to punish a firm who is likely to exit after it deviates, but if one of the competitors is near an exit state the other incumbent(s) has an incentive to price predatorily (that is to deviate themselves) in order to hasten that exit.

Formally the difference between the Fershtman & Pakes's (2000) model and the original EP framework is that they introduce a second state variable for each firm and let price choices depend on it. The second state variable, borrowed from the repeated game literature, is an indicator function which is one if the given firm has ever deviated from the collusive agreement in the past.<sup>31</sup> The Bellman equation for an incumbent firm is then more complicated because now at each state vector for which no one has deviated in the past, we must check to see that no one has an incentive to deviate in the current period. If either someone has deviated in the past, or someone has an incentive to deviate in the current period, then a static Nash pricing ensues. If neither of these conditions are satisfied, then the collusive prices are played. With this modification to the Bellman equation one

 $<sup>^{30}</sup>$ See Green & Porter (1984) and Abreu et al. (1986) and Judd, Yeltekin & Conklin (2003) for a method to compute the set of subgame perfect equilibria in repeated games.

<sup>&</sup>lt;sup>31</sup>This had been used previously to allow for collusion in Markov perfect settings See Haurie & Tolwinski (1986), Tolwinski, Haurie & Leitmann (1986), and Reynolds (1991). It is sometimes also possible to construct collusive equilibria without adding another state variable, see Fudenberg & Tirole (1983) and Nocke (2000).

can compute equilibrium values iteratively, using techniques analogous to those described in Section 4.1. Note that this implies that policies are computed for values of the state vector in which each incumbent firm has deviated in the past, as well as for cases when none have ever deviated. Since, in the equilibria that Fershtman & Pakes (2000) compute, no firm ever deviates, this implies that they need to compute values and policies which should never actually be observed (states that are "off the equilibrium path").

The results illustrate the potential importance of dynamic considerations in evaluating the benefits and costs of collusion. In particular for the parameter values they chose consumers prefer the collusive equilibria to the equilibria with no collusion; i.e. the greater quality and variety of goods that are marketed in the collusive equilibria more than compensates consumers for the higher prices they have to pay when there is collusion.

As Fershtman & Pakes (2000) note there are many ways to model collusion in a numerically tractable way and the choice among the possibilities should probably be made with a particular industry in mind. de Roos (2004) computes a model designed to match the available data on the the lysine cartel that operated in the 1990s. In his model firms' profits in the collusive regime are based on their market shares at the time the collusion agreement is struck and there is a punishment regime in addition to a non-cooperative and a collusive regime.

# 8 Topics for Further Study

There are a number of topics of potential importance to the successful use of the EP framework which have not yet been studied. Some of them involve a more in depth examination of the basic behavioral assumptions used in that framework. Others involve extensions to the framework which would be required to apply it to settings of obvious empirical importance. We now briefly outline four among those topics that we think may well be particularly important.

**Entry.** There has been a proliferation of entry models in which the potential entrants are short lived (i.e. entrants either enter this period or disappear, but cannot delay entry until the next period). EP assume that there is a finite number of potential entrants each period whose entry decisions are made sequentially. Letting n be the number of currently active firms, Doraszelski & Satterthwaite (2003) assume that every period there are  $\bar{n}-n$  potential entrants that make simultaneous decisions on whether to enter. Pakes et al. (2006) assume that the number of potential entrants is a random draw from a known distribution, and the potential entrants who do appear make their entry decisions simultaneously. As noted part of the reason for this proliferation is an absence of empirical facts on entry.

In any given applied setting it may be preferable to tie the number of potential entrants to particular exogenous state variables. As long as the exogenous state variable(s) evolve as a Markov process, this just requires us to add a state variable(s) to the model. It is also relatively easy to allow potential entrants to delay entry provided the number of active potential entrants in each period are observed by the other agents in the industry and the next period entry costs of the agents who do delay are independent of their current entry cost. This just requires us to add a state variable which specifies the number of potential entrants in each period.

It is more difficult to allow potential entrants in a given year to wait and enter in a future period when either; i) the number of available potential entrants is unknown to the other agents in the market, and/or (ii) when a potential entrant who delays in a given period has a setup cost in a following period which is correlated with their current setup cost. The difficulty does not lie in formulating the Bellman equation for the potential entrant; it is easy enough to add to that equation an option of remaining in the industry as a potential entrant in the next period. What is harder is to keep track of the number of potential entrants and the distribution of their setup costs when either the potential entrant's setup costs are correlated over periods, or when past potential entrants may (but need not) still be present (as they may have left to engage in some other activity). The reason is that then the number of potential entrants and the distribution of their setup costs depends on past states. I.e then the incumbent firms would use the observation that there were potential entrants who had not entered in prior periods to infer the number and potential setup costs of the potential entrants in the current period.

There are a number of possible workable alternatives here. For example we might allow potential entrants to delay by a small number of periods and then keep track of their number and the distribution of their setup costs over those periods. Alternatively we could keep track of sufficient statistics for the distribution of the number of potential entrants and their setup costs (as is done in the asymetric information model discussed below). Yet another alternative is to assume that the setup cost is composed of a deterministic part that is publicly known and a random part that is privately observed and then allow for a number of different types of entrants according to the deterministic part of setup costs (assuming the privately observed costs are independent over time). We do not know of anyone who has explored these (or any other) alternatives to date.

**Timing.** There are at least two aspects of timing that we have not delved into. One is the timing of when decisions are made. The other is the timing of the realizations of random variables and any lags between their updated values and the values of the "payoff relevant" random variables that affect the profit function. We begin with the timing of decisions.

The continuous-time model assumes that controls can be changed continuously while the discrete time version assumes that they can only be changed at fixed intervals. Though it may be the case that one (or both) of these alternatives provide an adequate approximation to a given setting, there may well be other cases in which the timing of decisions is crucial
to an understanding of outcomes, and the simple alternatives we have put forth are not rich enough to capture essential features of the environment. For example it is likely that firms have a planned schedule of meetings of the decision-making staff, but that extra meetings can be called when a change in the state of the industry warrants it. This makes the timing of decisions endogenous, a feature which is not allowed in (but could be added to) the framework. In cases where there are sharp changes in values resulting from being a first mover, a model with the endogenous timing of decisions may be necessary.

Relatedly we note that we are currently constraining decisions on our two instruments, price and investment, to occur at the same time. This may not lead to an adequate approximation for some industries. I.e. in some environments it may be more appropriate to have different (possibly endogenous) lengths of time for price commitment than for changes in investment policies. Moreover as we add further controls we should keep in mind that decisions on different controls need not be perfectly aligned.<sup>32</sup>

The continuous- and discrete-time models also treat realizations of the investment process differently. Starting at any initial state the continuous-time model will generate a change to one firm's state at some future time. The probability of the change occurring to any particular firm depends on the investments of all firms. After that change every agent will face the new industry state which consists of an updated state for the firm whose state experienced the change, and exactly the same state for the n-1 firms that did not. As in all Markov perfect models the game then continues from this new state. As a result to compute the continuation value in the continuous-time model we need only compute probabilities for which agent's state is the next to change and an integral over where it is likely to change to should it change; a computation whose burden is linear in the number of firms. Note that this implies that the investments made by n-1 firms in the interval of time between the initial state and the changed state vector will, with probability one, have no impact on the sample path of any of the agents. The discrete-time model assumes that realizations of all firms random outcomes occur between the instants when successive decisions are made. This forces the agent to take account of the possibility of simultaneous changes in all firms' states in the time interval between its current and future decisions, and hence in calculating continuation values. However even in the discrete-time model once a period has elapsed the investment made prior to that period is irrelevant for future sample paths.

Note that both specifications of investment assume that either the investment is successful, and leads to an immediate change in a "payoff relevant" random variable, or it is "lost." It may be that there are additional ordering, delivery and/or installation lags before the output of the investment activity can be embodied in marketable goods or services. Alternatively the whole idea that each new piece of information leads to an increment in profits may be too narrow a conceptualization to adequately characterize the investment process

 $<sup>^{32}</sup>$ We thank Joe Harrington for pointing this out to us.

in some industries. Consider, for example, pharmaceutical research. Even if investments do not lead to the discovery of a new drug in the current period, they may still yield information which changes the firm's perceptions of its (or of other firms') likely state in the future. If the information produced but not embodied in current output is known to all we could accommodate it by adding additional state variables to the model.<sup>33</sup> Alternatively it may well be important to account for the possibility that not all the information the firm discovers is known to the firm's competitors. In that case we would have to move to a model with private information which is correlated over time, like the model of asymmetric information model discussed below. In fact there has been very little investigation of the relationship of these timing issues to either the implications of the model, to its computational burden, or to the ability of the model to adequately mimic real-world environments.

Asymmetric information. In many situations it is natural to think that firms know more about their own state (their cost structure, the status of their research programs, etc.) then about their rivals' states. This leads us to the theoretical literature on asymmetric information (see Green & Porter (1984), Abreu et al. (1986), Athey & Bagwell (2001), Athey, Bagwell & Sanchirico (2004)). In these models firms observe the actions of (and/or signals sent by) their competitors and use them to formulate posterior distributions on their competitors' states. At least from a computational point of view constructing these posteriors is difficult. They depend on all the variables that contain information on a firm's competitors, and on the equilibrium conditions which tell the firm how to extract information on the competitors' states from whatever they observe. As a result the integral over successor states needed to determine policies is typically quite complicated.

Fershtman & Pakes (2005) present a different approach to computing equilibria in dynamic games with asymmetric information which circumvents the problem of computing posteriors. The idea is to obtain estimates of the expected future discounted value of net cash flows conditional on everything the firm knows. To do this they enlarge the state space to allow for informationally relevant variables, that is variables that contain information on the states of its competitors, as well as payoff relevant random variables. The estimates of the expected discounted values conditional on all observables are obtained from a "learning process" similar to that used in stochastic algorithm (see Section 5.3). Therefore there is no curse of dimensionality in the computation of the integral determining the value of successor states, and the algorithm converges to a recurrent class of the game in these states (which again need not grow in any particular way as a function of the number of state variables). This is a literature which is still is in its infancy. On the other hand, as noted above, there are a number of environments in which asymmetric information is likely to play a central

 $<sup>^{33}</sup>$ If it is one to one with investment expenditures, then we would add a state that measures the dollar value of the relevant investments. Alternatively there may be intermediate stages a firm must pass before it can embody the output of the research in a marketable good or service and we would add the state achieved; see, e.g., Doraszelski (2003) discussion of this point in the context of R&D race models.

role, so it is probably worth developing further.

**Dynamic consumers.** There are many situations in which it seems appropriate to allow consumers as well as producers to be forward looking. Analyzing markets with either durable or experience goods are two examples which, in and of themselves, are important enough to warrant detailed consideration of the relevant issues. When consumers, as well as producers, are forward looking, the fixed point that defines equilibrium values and policies requires consumers to maximize expectations based on consistent perceptions of the likelihood of future producer states, and producers to maximize expectations based on consistent perceptions on the likelihood of future consumer states. I.e. the conditions we used to define equilibrium in section (3) must be augmented to insure the optimality of consumer decisions. Though this is not conceptually difficult, it does imply that the state variables determining behavior include the the distribution of consumer, as well as that of producer, characteristics; a fact which, without further developments, is likely to increase the burden of computing equilibria significantly.

There have been some attempts to allow for dynamic consumers. Markovich (2004) and Markovich & Moenius (2005) have consumers that look two periods into the future in an oligopoly model and Nair (2005) has two types of forward-looking consumers in a monopoly model. Ching (2003) incorporates consumers learning into a dynamic model of the market for clonidine, a prescription drug, using estimated parameters from Ching (2002). However, he assumes that consumers base their purchase decision solely on their utility in the current period. Since most of manufacturing produces durable goods, and markets with experience goods typically raise a number of important policy issues (think of the market for perscription pharmaceuticals), the fact that we know of so few applications of equilibrium modeling when both consumers and producers are forward looking testifies to the difficulty of computing the equilibria from such models. Advances in computing hardware will no doubt help here, but this is definitely an area where algorithmic advances would be most welcome.

## 9 Conclusions

Dynamic analysis of imperfectly competitive markets is still in its infancy. Indeed there remain open questions in just about every dimension of the analysis. We have pointed out a number of ways one could make our assumptions richer. However even if we accept the modeling assumptions there remains an open and important theoretical question about how to select among multiple equilibria and about the relationship of the output of the algorithm to that selection mechanism. On a more applied front the extensions of the framework needed to analyze topics of obvious importance to the economy, such as the analysis of markets for durable, experience, or (most types of) network goods, are not yet available. Moreover the burden of currently available techniques for computing the equilibria to the models we do know how to analyze is still large enough to be a limiting factor in the analysis of many empirical and theoretical issues of interest. Finally we have adhered throughout to the notion of Markov perfect equilibria. Weaker notions of equilibria may be all that one can justify when describing the evolution of the industries we want to study (see e.g. Fudenberg & Levine's (1993) notion of self-confirming equilibrium). The weaker notions may also be easier to use in applied work (see Fershtman & Pakes (2005)), and have different implications then the stronger Markov perfect notion (see e.g. Esponda (2005)).

On the other hand we have to start somewhere. Moreover perhaps the most notable aspect of the applied results thus far is that even given all their assumptions and limitations they have reproduced phenomena that we observe in actual industries rather closely (see the discussion in the last section and the literature cited there). At the very least this has generated a deeper understanding of those phenomena. Of course we want to push the tools reviewed here further, possibly far enough to use to form predictions that would be accurate enough to use in policy analysis. The predictive accuracy of the current tools has not really been analyzed in any formal way, and is likely to vary with both the institutional setting and the issues of interest. Still one has to keep in mind that policy choices are often based on predictions of some form, and a computational tool would at least have the advantage of enabling an internally consistent quantititative assessment of alternatives.

## References

- Abbring, J. & Campbell, J. (2003), The entry and survival thresholds of a dynamic oligopoly, Working paper, University of Amsterdam, Amsterdam.
- Abreu, D., Pearce, D. & Stacchetti, E. (1986), 'Optimal cartel equilibria with imperfect monitoring', Journal of Economic Theory 39(1), 251–269.
- Ackerberg, D., Benkard, L., Berry, S. & Pakes, A. (2005), Econometric tools for analyzing market outcomes, in J. Heckman, ed., 'Handbook of Econometrics', Vol. 6, North-Holland, Amsterdam, p. forthcoming.
- Aguirregabiria, V. & Mira, P. (2006), 'Sequential estimation of dynamic discrete games', *Econometrica* forthcoming.
- Argote, L., Beckman, S. & Epple, D. (1990), 'The persistence and transfer of learning in an industrial setting', *Management Science* 36(2), 140–154.
- Athey, S. & Bagwell, K. (2001), 'Optimal collusion with private information', *Rand Journal* of Economics **32**(3), 428–465.
- Athey, S., Bagwell, K. & Sanchirico, C. (2004), 'Collusion and price rigidity', Review of Economic Studies 72(2), 317–349.

- Auerswald, P. (2001), The complexity of production, technological volatility and interindustry differences in the persistence of profits above the norm, Working paper, Harvard University.
- Bajari, P., Benkard, L. & Levin, J. (2006), 'Estimating dynamic models of imperfect competition', *Econometrica* forthcoming.
- Bajari, P., Hong, H. & Ryan, S. (2004), Identification and estimation of discrete games of complete information, Working paper, Duke University, Durham.
- Barto, A., Bradtke, S. & Singh, S. (1995), 'Learning to act using real-time dynamic programming', *Artifical Intelligenze* **72**(1), 81–138.
- Basar, T. & Olsder, J. (1999), *Dynamic noncooperative game theory*, 2nd edn, Society for Industrial and Applied Mathematics, Philadelphia.
- Bellman, R. (1957), Dynamic programming, Princeton University Press, Princeton.
- Benkard, L. (2000), 'Learning and forgetting: The dynamics of aircraft production', American Economic Review 90(4), 1034–1054.
- Benkard, L. (2004), 'A dynamic analysis of the market for wide-bodied commercial aircraft', *Review of Economic Studies* **71**(3), 581–611.
- Beresteanu, A. & Ellickson, P. (2005), The dynamics of retail oligopoly, Working paper, Duke University, Durham.
- Berry, S. & Pakes, A. (1993), 'Some applications and limitations of recent advances in empirical industrial organization: Merger analysis', American Economic Review 83(2), 247– 252.
- Berry, S. & Pakes, A. (2006), 'The pure characteristics demand model', *International Economic Review* forthcoming.
- Bertsekas, D. & Tsitsiklis, J. (1996), *Neuro-dynamic programming*, Athena Scientific, Belmont.
- Bertsekas, D. & Tsitsiklis, J. (1997), Parallel and distributed computation: Numerical methods, Athena Scientific, Belmont.
- Besanko, D. & Doraszelski, U. (2004), 'Capacity dynamics and endogenous asymmetries in firm size', *Rand Journal of Economics* **35**(1), 23–49.
- Besanko, D., Doraszelski, U., Kryukov, Y. & Satterthwaite, M. (2004), Learning-by-doing, organizational forgetting, and industry dynamics, Working paper, Northwestern University, Evanston.
- Bhattacharjee, A. (2005), Models of firms dynamics and the hazard rate of exits: Reconciling theory and evidence using regression models, Working paper, University of St. Andrews, St. Andrews.
- Cabral, L. & Riordan, M. (1994), 'The learning curve, market dominance, and predatory pricing', *Econometrica* 62(5), 1115–1140.

- Caplin, A. & Nalebuff, B. (1991), 'Aggregation and imperfect competition: On the existence of equilibrium', *Econometrica* **59**(1), 26–59.
- Cellini, R. & Lambertini, L. (2003), 'Advertising with spillover effects in a differential oligopoly game with differentiated goods', *Central European Journal of Operations Research* **11**, 409–423.
- Chen, J. (2004), Bias in merger evaluation due to cost misspecification, Working paper, UC Irvine, Irvine.
- Chen, J., Doraszelski, U. & Harrington, J. (2004), Network externalities, compatibility choice, and industry dynamics, Working paper, Hoover Institution, Stanford.
- Cheong, K. & Judd, K. (2006), 'Mergers and dynamic oligopoly', *Journal of Economic Dynamics and Control* forthcoming.
- Ching, A. (2002), Consumer learning and heterogeneity: Dynamics of demand for prescription drugs after patent expiration, Working paper, Ohio State University, Columbus.
- Ching, A. (2003), A dynamic oligopoly structural model for the prescription drug markt after patent expiration, Working paper, University of Wisconsin, Madison.
- Chintagunta, P. (1993), 'Investigating the sensitivity of equilibrium profits to advertising dynamics and competitive effects', *Management Science* **39**(9), 1146–1162.
- Conklin, J. & Judd, K. (1996), Computing value correspondences for repeated games with state variables, Working paper, Hoover Institution, Stanford.
- Darr, E., Argote, L. & Epple, D. (1995), 'The acquisition, transfer, and depreciation of knowledge in service organizations: Productivity in franchises', *Management Science* 41(11), 1750–1762.
- Davis, S. & Haltiwanger, J. (1992), 'Gross job creation, gross job destruction, and employment reallocation', Quarterly Journal of Economics 107(3), 819–863.
- de Roos, N. (2004), 'A model of collusion timing', International Journal of Industrial Organization 22, 351–387.
- Dockner, E., Jorgensen, S., Van Long, N. & Sorger, G. (2000), *Differential games in eco*nomics and management science, Cambridge University Press, Cambridge.
- Doraszelski, U. (2003), 'An R&D race with knowledge accumulation', Rand Journal of Economics **34**(1), 19–41.
- Doraszelski, U. (2005), A note on dynamic stochastic games with sequential moves, Working paper, Harvard University, Cambridge.
- Doraszelski, U. & Judd, K. (2004), Avoiding the curse of dimensionality in dynamic stochastic games, Technical working paper no. 304, NBER, Cambridge.
- Doraszelski, U. & Markovich, S. (2006), 'Advertising dynamics and competitive advantage', Rand Journal of Economics forthcoming.

- Doraszelski, U. & Satterthwaite, M. (2003), Foundations of Markov-perfect industry dynamics: Existence, purification, and multiplicity, Working paper, Hoover Institution, Stanford.
- Dube, J., Hitsch, G. & Manchanda, P. (2005), 'An empirical model of advertising dynamics', Quantitative Marketing and Economics 3, 107–144.
- Dunne, T., Roberts, M. & Samuelson, L. (1988), 'Patterns of firm entry and exit in U.S. manufacturing', Rand Journal of Economics 19(4), 495–515.
- Eaves, C. (1972), 'Homotopies for computation of fixed points', *Mathematical Programming* **3**(1), 1–22.
- Erdem, E. & Tybout, J. (2003), Trade policy and industrial sector reponses: Using evolutionary models to interpret the evidence, Working paper no. 9947, NBER, Cambridge.
- Ericson, R. & Pakes, A. (1995), 'Markov-perfect industry dynamics: A framework for empirical work', *Review of Economic Studies* 62, 53–82.
- Escobar, J. (2006), Time-homogenous Markov equilibrium in dynamic stochastic games, Working paper, Stanford University, Stanford.
- Esponda, I. (2005), Behavioral equilibrium in economies with adverse selection, Working paper, Stanford University, Stanford.
- Fershtman, C. (1984), 'Goodwill and market shares in oligopoly', *Economica* **51**(203), 271–281.
- Fershtman, C., Mahajan, V. & Muller, E. (1990), 'Market share pioneering advantage: A theoretical approach', *Management Science* 36(8), 900–918.
- Fershtman, C. & Markovich, S. (2006), Patent regimes in an asymmetric dynamic R&D race, Working paper, Northwestern University, Evanston.
- Fershtman, C. & Pakes, A. (2000), 'A dynamic oligopoly with collusion and price wars', Rand Journal of Economics 31, 294–326.
- Fershtman, C. & Pakes, A. (2005), Finite state dynamic games with asymmetric information: A framework for applied work, Working paper, Harvard University, Cambridge.
- Filar, J. & Vrieze, K. (1997), Competitive Markov decision processes, Springer, New York.
- Freedman, D. (1983), Markov chains, Springer, Berlin.
- Friedman, J. (1983), 'Advertising and oligopolistic equilibrium', Bell Journal of Economics 14(2), 464–473.
- Fudenberg, D. & Levine, D. (1993), 'Self-confirming equilibrium', *Econometrica* 61(3), 523– 545.
- Fudenberg, D. & Levine, D. (1998), The theory of learning in games, MIT Press, Cambridge.
- Fudenberg, D. & Tirole, J. (1983), 'Capital as commitment: Strategic investment to deter mobility', Journal of Economic Theory 31(2), 227–250.

Ghemawat, P. (1997), Games businesses play, MIT Press, Cambridge.

- Goettler, R., Parlour, C. & Rajan, U. (2004), 'Equilibrium in a dynamic limit order market', Journal of Finance forthcoming.
- Gort, M. (1963), 'Analysis of stability and changes in market shares', Journal of Political Economy 71, 51–63.
- Gowrisankaran, G. (1995), A dynamic analysis of mergers, Ph.D. thesis, Yale University, New Haven.
- Gowrisankaran, G. (1999), 'A dynamic model of endogenous horizontal mergers', Rand Journal of Economics **30**(1), 56–83.
- Gowrisankaran, G. & Town, R. (1997), 'Dynamic equilibrium in the hospital industry', Journal of Economics and Management Strategy 6(1), 45–74.
- Green, E. & Porter, R. (1984), 'Noncooperative collusion under imperfect price information', *Econometrica* **52**(1), 87–100.
- Hall, R., Royer, J. & Van Audenrode, M. (2003), Potential competition and the prices of network goods: Desktop software, Working paper, Hoover Institution, Stanford.
- Haurie, A. & Tolwinski, B. (1986), 'Definition and properties of cooperative equilibria in a two-player game of infinite duration', *Journal of Optimization Theory and Applications* 46, 525–534.
- Intriligator, M. (1971), *Mathematical optimization and economic theory*, Prentice Hall, Englewood Cliffs.
- Isaacs, R. (1954), Differential games, John Wiley & Sons, New York.
- Jenkins, M., Liu, P., Matzkin, R. & McFadden, D. (2004), The browser war: Econometric analysis of Markov perfect equilibrium in markets with network effects, Working paper, Stanford University, Stanford.
- Jia, P. (2006), What happens when Wal-Mart comes to town: An empirical analysis of the discount retailing industry., Working paper, Yale University, New Haven.
- Jofre-Bonet, M. & Pesendorfer, M. (2003), 'Estimation of a dynamic auction game', Econometrica 71(5), 1143–1489.
- Jovanovic, B. (1982), 'Selection and the evolution of industry', *Econometrica* **50**(3), 649–670.
- Judd, K. (1998), Numerical methods in economics, MIT Press, Cambridge.
- Judd, K., Kübler, F. & Schmedders, K. (2003), Computational methods for dynamic equilibria with heterogeneous agents, in M. Dewatripont, L. Hansen & S. Turnovsky, eds, 'Advances in economics and econometrics: Theory and applications, Eighth World Congress', Vol. 3, Cambridge University Press, Cambridge.
- Judd, K. & Schmedders, K. (2004), A computational approach to proving uniqueness in dynamic games, Working paper, Hoover Institution, Stanford.

- Judd, K., Schmedders, K. & Yeltekin, S. (2002), Optimal rules for patent races, Working paper, Hoover Institution, Stanford.
- Judd, K. & Yeltekin, S. (2001), Computing supergame equilibria with state variables, Working paper, Hoover Institution, Stanford.
- Judd, K., Yeltekin, S. & Conklin, J. (2003), 'Computing supergame equilibria', Econometrica 71(4), 1239–1254.
- Kadyrzhanova, D. (2006), The leader-bias hypothesis: Corporate control dynamics in industry equilibrium, Working paper, Columbia University, New York.
- Klette, T. & Raknerud, A. (2002), How and why do firms differ?, Working paper, University of Oslo, Oslo.
- Laincz, C. (2005), 'Market structure and endogenous productivity growth: How do R&D subsidies affect market structure?', Journal of Economic Dynamics and Control 29, 187–223.
- Laincz, C. & Rodrigues, A. (2004), A theoretical foundation for understanding firm size distributions and Gibrat's law, Working paper, Drexel University, Philadelphia.
- Langohr, P. (2003), Competitive convergence and divergence: Capability and position dynamics, Working paper, Northwestern University, Evanston.
- Lettau, M. & Uhlig, H. (1999), 'Rules of thumb versus dynamic programming', American Economic Review 89(1), 148–174.
- Maggi, G. (1996), 'Endogenous leadership in a new market', *Rand Journal of Economics* **27**(4), 641–659.
- Mankiw, G. & Whinston, M. (1986), 'Free entry and social inefficiency', Rand Journal of Economics 17(1), 48–58.
- Markovich, S. (2004), Snowball: A dynamic oligopoly model with network externalities, Working paper, Northwestern University, Evanston.
- Markovich, S. & Moenius, J. (2005), Winning while losing: Competition dynamics in the presence of indirect network effects, Working paper, Northwestern University, Evanston.
- Maskin, E. & Tirole, J. (1987), 'A theory of dynamic oligopoly, III: Cournot competition', European Economic Review **31**(4), 947–968.
- Maskin, E. & Tirole, J. (1988*a*), 'A theory of dynamic oligopoly, I: Overview and quantity competition with large fixed costs', *Econometrica* **56**(3), 549–569.
- Maskin, E. & Tirole, J. (1988b), 'A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and Edgeworth cycles', *Econometrica* 56(3), 571–599.
- McFadden, D. (1974), Conditional logit analysis of qualitative choice behavior, *in* P. Zarembka, ed., 'Frontiers of Econometrics', Academic Press, New York, pp. 105–142.

- McKelvey, R. & McLennan, A. (1996), Computation of equilibria in finite games, in H. Amman, D. Kendrick & J. Rust, eds, 'Handbook of Computational Economics', North-Holland, Amsterdam, pp. 87–142.
- McKelvey, R., McLennan, A. & Turocy, T. (2006), Gambit: Software tools for game theory, Technical report, California Institute of Technology, Pasadena. \*http://econweb.tamu.edu/gambit
- Mueller, D. (1986), Profits in the long run, Cambridge University Press, Cambridge.
- Nair, H. (2005), Intertemporal price discrimination and forward-lloking consumers: Application to the US market for console video-games, Working paper, Stanford University, Stanford.
- Nocke, V. (2000), Underinvestment and market structure, Working paper, University of Pennsylvania, Philadelphia.
- Noel, M. (2004), Edgeworth cycles and focal prices: Computational dynamic Markov equilibria, Working paper, UC San Diego, San Diego.
- Ortega, J. & Rheinboldt, W. (1970), Iterative solution of nonlinear equations in several variables, Academic Press, New York.
- Pakes, A. (2000), A framework for applied dynamic analysis in I.O., Working paper no. 8024, NBER, Cambridge.
- Pakes, A. (2006), Econometrics and theory in empirical io, Fisher schulz lecture, Harvard University, Cambridge.
- Pakes, A. & Ericson, R. (1998), 'Empirical implications of alternative models of firm dynamics', Journal of Economic Theory 79, 1–45.
- Pakes, A., Gowrisankaran, G. & McGuire, P. (1993), Implementing the Pakes-McGuire algorithm for computing Markov perfect equilibria in Gauss, Working paper, Yale University, New Haven.
- Pakes, A. & McGuire, P. (1994), 'Computing Markov-perfect Nash equilibria: Numerical implications of a dynamic differentiated product model', *Rand Journal of Economics* 25(4), 555–589.
- Pakes, A. & McGuire, P. (2001), 'Stochastic algorithms, symmetric Markov perfect equilibrium, and the "curse" of dimensionality', *Econometrica* **69**(5), 1261–1281.
- Pakes, A., Ostrovsky, M. & Berry, S. (2006), 'Simple estimators for the parameters of discrete dynamic games (with entry/exit examples)', Rand Journal of Economics forthcoming.
- Pesendorfer, M. & Schmidt-Dengler, P. (2003), Identification and estimation of dynamic games, Working paper no. 9726, NBER, Cambridge.
- Reynolds, S. (1991), 'Dynamic oligopoly with capacity adjustment costs', Journal of Economic Dynamics and Control 15(3), 491–514.

- Reynolds, S. & Wilson, B. (2000), 'Bertrand-Edgeworth competition, demand uncertainty, and asymmetric outcomes', *Journal of Economic Theory* **92**, 122–141.
- Robbins, H. & Monro, S. (1951), 'A stochastic approximation technique', Annals of Mathematical Statistics 22, 400–407.
- Rui, X. & Miranda, M. (1996), 'Solving nonlinear dynamic games via orthogonal collocation: An application to international commodity markets', Annals of Operations Research 68, 89–108.
- Ryan, S. (2005), The costs of environmental regulation in a concentrated industry, Working paper, Duke University, Durham.
- Saloner, G. (1987), 'Cournot duopoly with two production periods', *Journal of Economic Theory* **42**(1), 183–187.
- Sargent, T. (1993), Bounded rationality in macroeconomics: The Arne Ryde memorial lectures, Clarendon Press, Cambridge.
- Schivardi, F. & Schneider, M. (2005), Strategic experimentation and disruptive technological change, Working paper, New York University, New York.
- Schmedders, K. (1998), 'Computing equilibria in the general equilibrium model with incomplete asset markets', Journal of Economic Dynamics and Control 22, 1375–1401.
- Schmedders, K. (1999), 'A homotopy algorithm and an index theorem for the general equilibrium model with incomplete asset markets', *Journal of Mathematical Economics* 32(2), 225–241.
- Shapley, L. (1953), 'Stochastic games', Proceedings of the National Academy of Sciences 39, 1095–1100.
- Song, M. (2002), Competition vs. cooperation: A dynamic model of investment in the semiconductor industry, Working paper, Harvard University, Cambridge.
- Starr, A. & Ho, Y. (1969), 'Nonzero-sum differential games', Journal of Optimization Theory and Applications 3(3), 184–206.
- Stigler, G. (1968), The organization of industry, University of Chicago Press, Homewood.
- Tesfatsion, L. & Kabala, R. (1991), 'Solving nonlinear equations by adaptive homotopy continuation', Applied Mathematics and Computation 41(2), 99–115.
- Thompson, P. (2003), How much did the Liberty shipbuilders forget?, Working paper, Florida International University, Miami.
- Tolwinski, B., Haurie, A. & Leitmann, G. (1986), 'Cooperative equilibria in differential games', *Journal of Mathematical Analysis and Applications* **119**, 182–202.
- Vedenov, D. & Miranda, M. (2001), 'Numerial solution of dynamic oligopoly games with capital investment', *Economic Theory* 18, 237–261.
- Watson, L., Billups, S. & Morgan, A. (1987), 'HOMPACK: A suite of codes for globally convergent homotopy algorithms', ACM Transcations on Mathematical Software 13(3), 281–310.

- Watson, L., Sosonkina, M., Melville, R. & Morgan, A. (1997), 'HOMPACK90: A suite of Fortran 90 codes for globally convergent homotopy algorithms', ACM Transcations on Mathematical Software 23(4), 514–549.
- Weintraub, G., Benkard, L. & Van Roy, B. (2005), Markov perfect industry dynamics with many firms, Working paper, Stanford University, Stanford.
- Whitt, W. (1980), 'Representation and approximation of noncooperative sequential games', SIAM Journal of Control and Optimization 18(1), 33–48.
- Zangwill, W. & Garcia, C. (1981), *Pathways to solutions, fixed points, and equilibria*, Prentice Hall, Englewood Cliffs.