The Marhsall Lectures

Cambridge University

January, 2004.

Models of Markets for Applied Work

by Ariel Pakes

(Harvard Univsersity).

Lecture 2: Dynamics.

We introduced the last lecture with a brief discussion of the primitives needed to analyze simple static equilibria. These primitives enable an analysis of how prices, quantities, and the distribution of profits and consumer surplus, are determined conditional on the state variables of the problem (or how they change in response to a changes in the environment). In this section we will move to dynamic analysis; that is the analysis of how the state variables that are subject to the firms' control to evolve.

There are additional parameters that need to be estimated; the parameters determining the distribution of both entry and exit costs, and of the impacts of investments. Though there are several standard ways of estimating how investment impacts on the firm's state variables, it has proven much more difficult to get a handle on relevant ranges of entry and exit fees.

Once we have all the parameters we can compute the Markov Perfect Equilibrium policies. The second part of the talk discusses ways of doing that. We begin by quickly reviewing a standard iterative algorithm due to Pakes and McGuire (1994). That algorithm is subject to a "curse of dimensionality" which is often telling, so we conclude with a description of artificial intelligence algorithms which can often be used to circumvent this problem.

Empirical Entry Models

Entry and exit are a key part of the dynamics of market adjustments. In particular we usually think that it is the size and nature of entry costs that allow one (or a group of) firms to maintain a dominant position in a profitable market. Unfortunately most cost data are proprietary (and hence difficult for researchers to access), and data on sunk costs as virtually nonexistent. The decision of a firm on whether to exit is determined by whether its continuation value is greater than its sell-off value, and the latter is often associated with factors as hard to measure as "goodwill" or the value of the firm's building and equipment in its "second best" alternative employment. The potential entrants' sunk costs can be largely determined by the time and effort required to formulate the idea to be marketed, or by an individual entrepreneur's cost in accessing startup capital and/or the requisite permissions from a local administration. As a result we have to infer the extent of sunk costs from other variables whose behavior depends on them.

The variable that seems most directly related to the costs of entry is

entry itself. However to make use of the connection between actual entry and the costs of entry we need a framework which allows us to compute the value of entering (similarly to make use of the relationship of sell-off values and exit we need to be able to calculate the cost of continuing). Though such frameworks have been available for some time (e.g., the Ericson and Pakes, 1995, framework which we come back to presently), their implications cannot be used directly in estimation without encountering substantial (in many cases insurmountable) computational problems.

Largely as a result of these computational problems the early empirical analysis of entry and exit used, as a modelling framework, special cases of two-period entry games. The earliest two-period models assumed that the continuation values of all potential entrants are identical (Bresnahan and Reiss, 1987, 1990, and Berry, 1992). More recent work by Siem (2003) and Mazzeo (2003) provides ways of extending the framework to allow for a finite number of types which differed in their continuation value. Allowing for different "types" is crucial for realistic analysis of most issues we want to use entry models to analyze; e.g. the analysis of location and size decisions in retail trade and their response to various policies.

The two-period game framework is, however, problematic. Most importantly it makes little sense unless sunk costs are absent (or at least are very small relative to single period profits). This is an assumption which is hard to believe, particularly since much of our interest in entry models stems from a desire to investigate the origins and implications of sunk costs.

Note that dynamic economic model that we want to base our analysis on is transparent; agents enter (continue) if the expected discounted value of entering (continuing) is greater than the cost of entry (the sell-off value). It is the computation of the expected discounted values that is complex. My goal here is to show how semiparametric methods can be used to provide a simple way of approximating those values that is grounded in what actually happened. Most of this part of the talk is taken from Pakes, Ostrovsky, and Berry (2003), a paper I will refer to simply as POB, and I thank my coauthors for permission to use this material.

A Semiparametric Entry Model

I will focus on the simple case where there is only one entry location the same number of potential entrants in each period, and no investment after entry. All the essential points can be made with this simple case and you will

see that, at least in principle, the argument is easy to generalize to a finite number of entry locations, a random number of potential entrants in every period (see POB), and post entry investment.

Let n_t be the number of agents active at the beginning of each period, z_t be a vector of exogenous profit shifters which evolve as a discrete state Markov process, and assume that there is a one-period profit function that is determined by these variables, say $\pi(n, z; \theta)$, where θ is a parameter vector to be estimated. An incumbent chooses to exit if current profits plus the discounted selloff value is greater than profits plus the discounted continuation value. So if ϕ is the sell-off (or exit) value and $0 < \delta < 1$ is the discount rate, the "Bellman" equation for the value of an incumbent is

$$V(n, z; \phi, \theta) = \max \left\{ \pi(n, z; \theta) + \delta\phi, \pi(n, z; \theta) + \delta VC(n, z; \theta) \right\}, \qquad (1)$$

where $VC(\cdot)$ is the continuation value. If the max is the first term inside the curly brackets, the incumbent exits.

If e is the number of entrants, x is the number of exitors (both of which are unknown at the time the incumbents decisions are made), and $p(\cdot)$ is notation for a probability distribution, then $VC(\cdot)$ is just the expectation (over the possible numbers of exitors, entrants, and values of the profit shifters) of the next period's realization of the value function (of future $V(\cdot)$), or

$$VC(n, z; \theta) \equiv \sum_{\phi', z', e, x} V(n+e-x, z', \phi'; \theta) p(\phi'|\theta) p(e, x|n, z, \chi = 1) p(z'|z). \quad (2)$$

Note that to form this expectation we need to form the incumbent's perceptions of the likely number of entrants and exitors *conditional on the incumbent itself continuing*, perceptions that we write as the probability distribution

$$p(e, x|z, n, \chi = 1)$$

where $\chi = 1$ is notation for the incumbent continuing. We need these perceptions because the incumbent cannot estimate his returns to continuing without an idea of how many other firms will be active. It is the requirement that these perceptions be consistent with behavior that will generate our equilibrium conditions.

Analogously we assume that the entrant must commit to entering one period before it earns any profit, so the value of entry is

$$VE(n, z; \theta) \equiv \sum_{e, x, z', \phi'} V(n + e - x, z', \phi') p(\phi'|\theta) p(x, e|n, z, \chi^e = 1) p(z'|z), \quad (3)$$

where

$$p(x, e|n, z, \chi^e = 1)$$

provides the potential entrant's perceptions of the likely number of entrants and exitors *conditional on it entering*, or conditional on $\chi^e = 1$.

The potential entrant enters if

$$\delta VE(n, z; \theta) \geq \kappa$$

where κ is its sunk cost of entry.

Assumptions and Their Implications

In addition to technical "regularity" conditions, it is assumed that entry and exit decisions are made simultaneously at the beginning of the period. The following assumptions are also made.

1. There are a fixed number of potential entrants in each period (denoted by $\mathcal{E}$), and the distribution over

   - the sunk costs of entry, say $F^{\kappa}(r|\beta)$, which has a lower bound of $\underline{\kappa} > 0$, and
   - the returns to exiting, say $F^{\phi}(\cdot|\beta)$, which are assumed nonnegative,

   are i.i.d. over time and across markets. Incumbents and entrants know these distribution and their own realizations, but do not know the realizations of their competitors (so there is asymmetric information, as in Siem, 2001).

2. Entrant's and incumbent's perceptions of the probabilities of exit and entry by their competitors in period t depend only on (n_t, z_t) (the publicly available information at that time).

These assumptions are *truly restrictive*. They imply, for example, that there are no serially correlated state variables that are observed to the agents and not to the econometrician. However they (or simple generalizations of them that allow for multiple locations) are less restrictive than the assumptions used in any of the two-period models that have been taken to data to date. Moreover, as we will show below, these assumptions lead to an estimator of continuation and entry values that has a transparent relationship to

objects in the data, and therefore has both a great deal of intuitive appeal and is easy to work with.

This model is a special case of the model in Ericson and Pakes (1995) and so has an equilibrium, but there may be more than one of them (for more detail see Doraszelski and Satterwaite, 2003). Each equilibrium generates a finite state Markov chain in (n, z) couples: i.e. the distribution of possible (n, z)'s in the next period depends only on the current (n, z) (and not on either prior history, or time itself). Indeed one can go a bit further and note that every possible sequence of $\{(n_t, z_t)\}$ will eventually wander into a recurrent subset of the possible (n, z) couples, say $\mathcal{R}$, and once (n_t, z_t) is in the set $\mathcal{R}$ it will stay in it forever (Freedman, 1983). Thus, for example, there will be an $\overline{n}$ such that, provided the current n is lower than it, we will never observe an $n > \overline{n}$. The market is simply not profitable enough to induce entry if there are $\overline{n}$ incumbents, so states with $n > \overline{n}$ are not in $\mathcal{R}$. However all states in $\mathcal{R}$ "communicate" with each other, and will eventually be visited many times.

I want to emphasize, however, that though our assumptions do not guarantee a unique equilibrium, they do insure that *there is only one equilibrium that is consistent with a given data generating process.* As a result we will be able to use the data itself to "pick out" the equilibrium that is played, and at least for large enough samples, we will pick out the correct one. This is all we require to develop consistent estimators for the parameters of the model.

To see that the data can be used to pick out the equilibrium, note that (i) the agents only condition their perceptions of the behavior of their competitors on the publicly available information (on (n, z)) and (ii) precisely the same information is available to the econometrician. Moreover in equilibrium the realized distribution of entrants and exitors from each state must be consistent with these perceived distributions.

Now recall that the data will eventually wander into the recurrent subset of points, and once in that subset will visit each point in it repeatedly. As the sample gets large we obtain an empirical distribution of entrants and exitors from each (n, z), and by the law of large numbers that distribution will converge to the distribution which generated it (almost surely). As noted this must be the distribution the agents use to form their perceptions, so we have just identified the perceived distributions needed for agents to make their decisions.

Given those perceived distributions equations (1) and (2) generate a unique best response for each incumbent and potential entrant. This is just

the familiar statement that reaction functions are generically unique, and can be proven using Blackwell's theorem for single agent dynamic programs. Since there is only one policy that is consistent with both the data and our equilibrium assumptions at each $(n, z) \in \mathcal{R}$, and once we are in the set $\mathcal{R}$ we stay there forever, there is a unique equilibrium for any subgame starting from any (n, z) couple in $\mathcal{R}$ (a set which can be identified from the data).[1] We now provide semiparametric estimators for the continuation and entry values generated by that equilibrium.

Equilibrium Perceptions and Continuation Values

In equilibrium the perceptions of potential entrants and incumbents of the likelihood of entry and exit by their competitors must be consistent with the distribution of entry and exit actually generated by incumbent and potential entrant behavior. This observation leads directly to two semiparametric estimates of continuation values at each (n, z) observed at least once.

The first estimate simply averages the realized continuation values of all firms who did continue when $(n_t, z_t) = (n, z)$. Since agents perceptions of probabilities of reaching different states must be consistent with the actual probabilities of reaching the different states, the sample average of discounted future values will converge to expected continuation value we are after. This estimator is particularly simple, and I will provide detail on it below.

The second estimator, though also intuitive, is computationally more complex, so the reader who wants more detail should see POB and the literature referred to there. Briefly, if we knew $p(e, x|n, z, \chi = 1)$ we could compute the fixed point which defines the continuation value in equations (1) and (2) for any value of θ. What the second procedure does is substitute the empirical distributions of entrants and exitors from each state for $p(e, x|n, z, \chi = 1)$ and compute the fixed point implied by these empirical distributions. Equations (1) and (2) define a contraction mapping, so solving the associated fixed point problem is not too difficult, but it must be solved every time we need

[1] There is a detail missing here. Though points in $\mathcal{R}$ can only communicate with other points in $\mathcal{R}$ if optimal policies are followed, there are some points, "boundary points" in the terminology of Pakes and McGuire (2001), that could communicate with points outside of $\mathcal{R}$ if feasible but nonoptimal policies were followed. To analyze equilibria for subgames in $\mathcal{R}$ fully, boundary points need to be treated separately (see Pakes and McGuire, 2001). In our case the only decisions that involve boundary points are the decisions of entrants at the maximum n observed for any given z; thus we can easily isolate them.

to evaluate θ in the estimation algorithm. In the terminology of Rust (1994) this is a "nested fixed point" estimation algorithm.[2]

The next subsection provides a brief description of how to obtain the estimates based on sample averages. The reader who is not interested in these details should be able to skip directly to the remark at the end of this subsection, and continue from there.

Using Sample Averages to Form Continuation Values

To form this estimate we need estimates of (i) the distribution of future states given the current state and (ii) the realized value at those future states.

We begin with the realized values at the future states. An incumbent at t who continues to $t+1$ receives that period's profits plus a choice between the continuation value at $t+1$ and a draw from the distribution of sell-off values. The incumbent chooses the sell-off value if and only if it is greater than the continuation value.

Let functions indexed by t refer to the function evaluated at (n_t, z_t), and let the probability of exit be given by

$$g^x_{t+1} \equiv Pr\{\phi > VC_{t+1}(\theta)\}. \tag{4}$$

Then given how the incumbent chooses, the continuation value of firms who did continue from period t conditional on the (publicly available) information available in period $t+1$, say $\widetilde{VC}R_t(\theta)$, is given by

$$\widetilde{VC}R_t(\theta) = \pi_{t+1}(\theta) + \delta(1 - g^x_{t+1})VC_{t+1}(\theta) + \delta g^x_{t+1} E[\phi|\phi > VC_{t+1}(\theta)] \tag{5}$$

where $E[\phi|\phi > VC_{t+1}(\theta)]$ is the expected sell-off value conditional on exiting.

[2] Aguirreberia and Mira (2003), in independent work, provide a closely related estimator. Instead of using the empirical transition probabilities for (n, z) couples, they use the empirical exit (and for entrants, entry) probabilities to generate these transitions. That is, since all incumbents are identical, a consistent estimate of the probability that x of n incumbents exit is given by the binomial formula with an exit probability estimated by the fraction of incumbents that exit. A similar procedure can be used to generate entry probabilities. Aguirreberia and Mira (2003) construct a consistent estimate of $p(e, x|n, z, \chi = 1)$ from these binomial probabilities. They assume that the random entry and exit fees are i.i.d. draws from an extreme value distribution. This is not necessary but eases the computational burden of their nested fixed point estimator significantly. POB compare these and other alternative estimators.

The last term has to be made more explicit. The easiest case is when sell-off values are exponentially distributed with parameter σ, because a truncated exponential is just another exponential with its origin shifted to the truncation point.[3] In this case, then

$$E[\phi|\phi > VC_{t+1}(\theta)] = \sigma + VC_{t+1}(\theta). \tag{6}$$

Now make two substitutions in the equation for $\widetilde{VCR}_t(\theta)$: (i) substitute this expression for $E[\phi|\phi > V_{t+1}(\theta)]$ and (ii) substitute an estimate of the exit probability obtained by averaging over the fraction who exit at the different periods in the data when $(n, z) = (n_t, z_t)$, say $\tilde{g}_t^x$, for the unobserved g_t^x. Rearranging terms we obtain

$$\widetilde{VCR}_t(\theta) = \pi_{t+1}(\theta) + \delta VC_{t+1}(\theta) + \delta\sigma\tilde{g}_{t+1}^x. \tag{7}$$

Note that the expectation of $\widetilde{VCR}_t(\theta)$ is equal to the expected continuation value of all those firms who continue. That is not quite $VC_t(\theta)$, since in forming the expectation which determines whether it should continue the incumbent conditions on itself continuing (while in the realizations measured by $\widetilde{VCR}_t(\theta)$ only a fraction continue). As shown in POB the ratio of the probabilities observed in the data, to the probabilities that the incumbent uses is consistently estimated by $w_t = [1 - \tilde{g}_t^x]/[1 - (x_t/n_t)]$, where $\tilde{g}_t^x$ is, as before, the fraction of incumbents who exit when $(n, z) = (n_t, z_t)$.

So if we let

$$VCR_{t+1} = w_t\widetilde{VCR}_t(\theta)$$

then

$$E[VCR_t(\theta)] = VC_t(\theta)$$

the unknown continuation value as defined in equation (1).

The sample average of the observations on $VCR_t(\theta)$ computed as an average over those periods where (n_t, z_t) was a particular value of (n, z), say (n^*, z^*), converges to the population expectation of the continuation value for (n^*, z^*), that is to the $VC(n^*, z^*; \theta)$ we are after. More precisely if we

[3] Virtually any other distributional assumption could be used here, but most others (with the exception of the uniform) would result in a somewhat more computational complex estimation algorithm. On the other hand the shape restrictions of the exponential seems reasonable for sell-off values.

let $T(n,z)$ be the set of periods when there was n firms active and the state variable had value z, and $\#T(n,z)$ be the number of such periods, and define

$$\widetilde{VC}(n,z) \equiv \frac{1}{\#T(n,z)} \sum_{t \in T(n,z)} w_t \widetilde{VCR}_t(\theta), \tag{8}$$

then if $\rightarrow_P$ is read as converges in probability, we have

$$\widetilde{VC}(n,z) \rightarrow_P VC(n,z)$$

provided $\#T(n,z) \rightarrow \infty$.

Finally we need to connect the estimates of the continuation values from the different (n,z) to each other. Here I will need matrix notation. Arrange the $\widetilde{VC}(n,z;\theta)$ of the (n,z) combinations observed in the data into the vector $\widetilde{VC}(\theta)$, let the matrix of empirical transition probabilities from one observed (n,z) combination to another (weighted as above) be $\tilde{M}$, and gather the average fraction of firms exiting from each state into the vector $\tilde{g}^x$. Then, if we substitute our expression for $\widetilde{VCR}(\theta)$ from equation (7) into the expression for $\widetilde{VC}(n,z)$ in equation (8) and rearrange we have

$$\widetilde{VC}(\theta) \equiv \tilde{M}\left[\pi(\theta) + \delta\sigma\tilde{g}^x\right] + \delta\tilde{M}\ VC(\theta).$$

This gives us our estimates of continuation values in terms of objects that we know (at least up to the parameters to be estimated), and the continuation values itself. We now iteratively substitute our estimate of continuation values for the continuation values that appear in this expression. Doing this once we obtain

$$\widetilde{VC}(\theta) \equiv \tilde{M}\left[\pi(\theta) + \delta\sigma\tilde{g}^x\right] + \delta\tilde{M}\left[\pi(\theta) + \delta\sigma\tilde{g}^x\right] + \delta^2\tilde{M}^2 VC(\theta),$$

and if we continue this process of substitution we get

$$\begin{aligned} \widetilde{VC}(\theta) &= \sum_{\tau=0}^{\infty} \delta^{\tau+1} M^{\tau}\left[\pi(\theta) + \delta\sigma\tilde{g}^x\right] \\ &= [I - \delta\tilde{M}]^{-1}\tilde{M}[\pi(\theta) + \delta\sigma\tilde{g}^x], \end{aligned}$$

as a consistent estimator for the unknown continuation values.

Note that $\widetilde{VC}(\theta)$ is just the discounted value of the returns of the incumbents who did continue (adjusted to account for the fact that the incumbent conditions on itself continuing). This is the sense in which our estimator of continuation values is transparent and likely to have empirical content. We expect the actual average of realized continuation values to be close to the true expected continuation value.

Remark on Ease of Computation

Note how easy it is to compute our estimates of continuation values, or $\widetilde{VC}(\theta)$. If δ is known (and we usually think that the prior information we have on δ is likely to swamp the information on δ available from estimating an entry model), then

$$\widetilde{VC}(\theta) = \tilde{A}\pi(\theta) + \sigma\tilde{a}$$

for a data matrix $\tilde{A} = [I - \delta\tilde{M}]^{-1}\tilde{M}$ and a data vector $\tilde{a} = \delta[I - \delta\tilde{M}]^{-1}\tilde{g}^x$. So if profits were linear functions of θ, the first stage estimates of continuation values are also.

POB shows (and it is easy to verify) that consistent estimates of entry values, say $\widetilde{VE}(\theta)$, can be obtained as

$$\widetilde{VE}(\theta) = \tilde{B}\pi(\theta) + \sigma\tilde{b},$$

where

$$\tilde{B} \equiv \tilde{M}^e + \delta\tilde{M}^e\tilde{A}, \qquad \tilde{b} \equiv \delta\tilde{M}^e\tilde{a} + \delta\tilde{M}^e\tilde{g}^x,$$

and $\tilde{M}^e$ is the Markov transition matrix formed after weighting the observed transitions with $w_t^e = (e_t/\mathcal{E})/\tilde{g}_t^e$, where $\tilde{g}_t^e$ is an estimate of the entry probability obtained by averaging over the fraction who enter at the different periods in the data when $(n, z) = (n_t, z_t)$.

Estimation Methods

We estimate models by finding the value of the model's parameter vector that makes the predictions of the model as close as possible to the observed data. In our case the model predicts that, conditional on (n, z), the number of entrants and exitors are independent draws from two binomial distributions. Moreover a consistent estimate of the model's binomial probability of exit in a given period is

$$1 - F^{\phi}(\tilde{A}_t\pi(\theta) + \sigma\tilde{a}_t|\theta)$$

and of the binomial entry probability is

$$F^{\kappa}(\tilde{B}_t\pi(\theta) + \sigma\tilde{b}_t|\theta)$$

where the t subscript on (A, a, B, b) picks out the row of each at which $(n, z) = (n_t, z_t)$.

POB considers different estimators for the models parameters obtained by using different measures of the distance between the theoretical predictions for the entry and exit probabilities conditional on θ, and the empirical observations. All of them will be consistent and have limiting normal distributions. They will, however, differ in both their variance, and in their computational complexity. Moreover though the computational issues might now be not nearly as crucial, the variance issue is likely to be of some importance. This because the first step estimates entry and exit values at each possible state, and the variances of those estimates effect the variances of the estimates of the parameters of interest[4].

For now all I want to do is note the following:

- There are good reasons to think that pseudo maximum likelihood estimator, the estimator obtained by substituting the estimates of the probabilities given above into the likelihood function and maximizing with respect to the parameter vector, will not do well in this context, so one should probably focus on alternative estimators that fit the expected number of entrants and exitors at each state.

- Second though our estimator is a "two-step" estimator with second step "regressors" that are estimated, there is a simple "parametric bootstrap" procedure which provides consistent estimators for the variances of the parameter estimates.[5] .

[4]Thus in the two location model that we used for Monte Carlo analysis, there were 404 observed states implying that our preliminary estimators estimated $4 \times 404 \approx 1600$ non parametric exit and entry values.

[5]The parametric bootstrap is obtained as follows. Let the estimate of the parameter vector be $\hat{\theta}$, and compute $\hat{VC}(\hat{\theta})$ and $\hat{VE}(\hat{\theta})$, the estimates of the continuation and entry values obtained by substituting the $\hat{\theta}$ into the equations for those values. Starting from the observed $(n_{t=0}, z_{t=0})$, randomly draw sell-off values from $F^{\phi}(\hat{\theta})$ and entry values from $F^{\kappa}(\hat{\theta})$ for each of the incumbents and potential entrants. Assume each incumbent exits if its sell-off value is greater than the component $\hat{VC}(\hat{\theta})$ that is associated with $(n_{t=0}, z_{t=0})$. Similarly the potential entrants enter if their randomly drawn entry cost was less than $\hat{VE}(\hat{\theta})$. Keep this predicted number of entrants and exitors in memory, and construct $n_{t=1}$ as the initial number of incumbents minus the exitors plus the entrants. Next randomly draw a $z_{t=1}$ from (a consistent estimate of) the Markov process generating z. This gives us $(n_{t=1}, z_{t=1})$, and we can determine exit and entry from this point in the same way as we did above. Continue in this way until we have a simulated sample with exactly the same number of periods (or of markets and periods) as in the original sample. Once we have this "pseudo random" sample, substitute it into the estimation algorithm and

I conclude this section by stressing two characteristics of our estimators. The first is their *simplicity*. I.e. to obtain them we need only minimize analytic functions of the parameters. Indeed the only aspect of the algorithm that is at all computationally burdensome are the one time costs of organizing the data into the Markov chain M, and then computing the matrix inverse needed for

$$\widetilde{VC}(\theta) = [I - \delta\tilde{M}]^{-1}\tilde{M}[\pi(\theta) + \delta\sigma\tilde{g^x}].$$

This makes the computational burden of these estimators as easy (or easier) than the estimators for any of the earlier two-period entry models. Moreover the model is easily generalized to allow entry into multiple locations (the number of entrants to the different locations are then draws from a multinominal rather than that of a binomial distribution), and to models with a random number of entrants (the number of entrants into the different locations are then a draw from a mixture of multinominals, so the estimating equation is still an analytic function of the parameters of the problem). Our Monte Carlo work clearly bears this out; we obtain estimates of two location models on panels of five years of observations on five hundred markets in under two minutes on my desk top.

Second, the actual estimates of the entry and exit values are grounded in reality in the sense that they equal the average of the true discounted value of continuing and entering from the different states. As a result the entry and exit cost distributions we estimate are those that rationalize the behavior of the incumbents and potential entrants in the face of what, on average, *actually happened*. Unless the incumbents and potential entrants were basing decisions on information that was systematically biased, the results from this procedure should both make "sense" and be robust to small perturbations of the model.

Computing Markov Perfect Equilibrium.

This section assumes that we have estimates of all the parameters of the problem and considers computing the dynamic equilibria that those estimates imply.

---

generate a new estimate of θ. Do this for a number of simulated samples, and calculate the variance-covariance of the estimates of θ obtained from the alternative pseudo random samples. This is a consistent estimate of the variance covariance matrix of $\hat{\theta}$.

We go back to the first lecture's examples where firms can differ in the products they are marketing and/or their costs of production. Recall that their price (or quantity) decisions are made to maximize current profits and the static analysis can output profits as a function of the firm's own characteristics, the characteristics of its competitors, and common state variables. For simplicity we work with unidimensional "types", and think of i as indexing either the quality of the product the firm produces, or the efficiency of its production technology. We compute equilibria to get an idea of sensible investment, entry, and exit strategies.

More formally, we let

- $i \in \mathcal{Z}^+$.
- s_i be the number of firms with efficiency level i,
- $s = [s_i;\ i \in \mathcal{Z}^+]$ be the "industry structure" (the number of firms at each different efficiency level),

so the function $\pi(i, s)$ provides the firm's profits when its state is i and the industry structure is s.

Bellman Equation for Incumbent Behavior.

The Bellman equation for the incumbent is similar to that in the entry model except that now we allow for investment, so

$$V(i,s) = \max\{\phi, \pi(i,s) + \sup_{(x \geq 0)} [-cx + \beta \sum V(i',s') pr(i',s'|x,i,s)]\}.$$

The only term which is not a primitive in this expression, and hence which needs to be calculated before we can determine $V(i, s)$, is the term representing the firm's perceptions about the likelihood of the alternative states that the firm is likely to face in future periods, or $pr(i', s'|x, i, s)$.

We assume that its own future states depends only on its own current state and its investments, and that decisions on investment are made simultaneously so that

$$pr(i', s'|x, i, s) = pr(i'|i, x) pr(s' - e(i')|i, s)$$

where $e(i)$ is a vector which puts one in the i^{th} slot and zero elsewhere, so $s' - e(i') \equiv \hat{s}_i$ is the vector consisting of the location of all of the competitor's of a firm at i.

Here $pr(i'|i,x)$ is a primitive reflecting how a firm's investment is likely to effect its state. On the other hand if we stick with standard subgame perfect notions of equilibrium, $pr(s' - e(i')|i,s)$, the firm's perceptions on the possible locations of its future competitors, has to be consistent with the behavior of those competitors. It is therefore an equilibrium construct.

Starting with the specification of investment we assume that

$$i_{t+1} - i_t \equiv \tau_{t+1}$$

where

$$\tau_t \equiv \nu_t - \zeta_t$$

ν is a firm's investment outcome, and ζ is a common industry shock representing either the state of demand or factor costs, depending on the model. It is the presence of ζ that causes positive correlation in the profits of the firms in a given industry.

We assume that the distribution of ν depends on the amount invested through the family

$$\mathcal{P} = \{p_\nu(\cdot|x), x \in \mathcal{R}^+\},$$

which is stochastically increasing in x, and has the property that $Pr\{\nu > 0|x = 0\} = 0$, so agents can not improve their state without investing. ζ has density $\mu(\zeta)$.

Recall that $\hat{s}_i = s - e_i$ provide the states of the competitors of a firm at state i for a particular s. A firm can not make rational investment decision without some perception of the likelihood of different numbers of competitors in the future. Let $q[\hat{s}_i{}'|i,s,\zeta]$ provide the firm's perceived probability of its competitors future states conditional on a particular value of ζ. Then given the above

$$pr(i' = i^*, s' = s^*|x,i,s) =$$
$$\Sigma_\zeta p(\nu = i^* - i - \zeta|x) q[\hat{s}_i{}' = s^* - e(i^*)|i,s,\zeta]\mu(\zeta).$$

where $e(i)$ has a one in the i^{th} slot and zero elsewhere.

Given any $q(\cdot|\cdot)$ standard single agent dynamic programming results allow us to solve for the policies of each agent. Of course the $q(\cdot|\cdot)$ actually chosen must "make sense"; and this is the equilibrium condition we come to presently.

Note also that $q[\cdot|i,s,\zeta]$ embodies the incumbent's beliefs about entry and exit. We take the simplest entry model possible. A potential entrant pays an amount x_e ($> \beta\phi$) to enter, and if it does enter, it enters one period later at state $\imath^e \in \Omega^e \subset \mathcal{Z}^+$ with probability $p^e(\cdot)$. The potential entrant only enters if the expected discounted value of future net cash flows from entering is greater than the cost of entry. The cost of entry can be specified as either x_e, or as a random variable which distributes uniformly on $[x_{e,l}, x_{e,u}]$[6].

Dynamic Equilibrium:Characterization Results.

Informally, an equilibrium for this model consists of

- a set of policies that maximizes each agent's value given its perceptions on the evolution of the likely states of its competitors, and
- a set of perceptions on competitor outcomes that are consistent with these policies.

Ericson and Pakes (1995) provide a formal characterization of the equilibrium, prove existency, and note that any equilibrium has the following characteristics

1. It is "computable". I.e. we will only ever observe values of "i" on the finite set $\Omega = \{1, ..., K\}$, and there will never be more than $\overline{n}$ firms active. Note that this implies that we need only compute equilibria for $(i,s) \in \Omega \times S$, where

   $$S \equiv \{s = [s_1, ..., s_k] : \sum s_j \leq \overline{n} < \infty\}$$

   so that the number of elements in S or $\#S \leq K^N$.

2. It is Markov. Equilibrium policies generate a homogeneous Markov chain for industry structures [for $\{s_t\}$], i.e.

   $$Pr[s_{t+1} = s' | s_t, ..., s_1] = Pr[s_{t+1} = s' | s_t] \equiv Q[s'|s_t].$$

   with the Markov transition "kernel" $Q(\cdot,\cdot)$ on $S \times S$.

[6] When random entry costs are used only the potential entrant knows the realization of the entry costs, the other incumbents know only that entry costs will be a random draw from this uniform distribution.

3. Any equilibrium $Q[\cdot|\cdot]$ is *ergodic*.

The lower bound to i comes from the fact that at any $\hat{s}_i$ there is an i low enough to induce exit, and there is only a finite number of $\hat{s}_i$. The upper bound to i comes from the assumption that the profits are bounded, which implies the value is bounded, which in turn implies that at some point it no longer pays to invest (the benefit of investment is it gives the firm a higher probability of increasing its i which, for high enough i is not worth much). The bound on n comes from regularity conditions on the profit function.

The ergodicity just says that s_t will eventually wander into a subset of states, the recurrent class or $\mathcal{R}$ and once in $\mathcal{R}$ will stay within it forever. These are just the states that the incentives generated by your parameter values can support in any lasting way (thus if the market is large enough we will never get too low a value for n, since even if a sequence of bad draws went to all of the incumbents, before we got to a very low n there would be entry). It is frequently the case that $\mathcal{R}$ is much smaller than S, indeed it tends to grow linearly, instead of exponentially, in the number of state variables, a fact which we will use intensively below.

Computation of Policies.

Here I want to briefly outline two computational algorithms, one the "dynamic programming" like iterative algorithm provided in Pakes and McGuire (1994) and the other an artificial intelligence algorithm in Pakes and McGuire (2001).

To understand how these algorithm's work it is useful to rewrite the Bellman equation as

$$V(i,s) = \max_{\chi\in\{0,1\}}\{[1-\chi]\phi + \tag{9}$$
$$\chi\{\pi(i,s) - \sup_{x\geq 0}[-cx + \beta\sum_{\nu} w(\nu;i,s)p(\nu|x_1)]\},$$

where

$$w(\nu;\ i,s) \equiv$$
$$\sum_{(\hat{s}_i',\zeta)} V(i+\nu-\zeta, \hat{s}_i' + e(i+\nu-\zeta)|w)q[\hat{s}_i'|i,s,\zeta]\mu(\zeta), \quad (1a)$$

and

$$q[\hat{s}_i^{'} = s_i^* | i, s, \zeta] \equiv$$

$$Pr\{\hat{s}_i^{'} = \hat{s}_i^* | i, s, \zeta, \text{equilibrium policies } \} \quad (1b).$$

Here $w(\nu; i, s)$ is the expected discounted value of future net cash flow conditional on the current year's investment resulting in a particular value of ν, and the current state being (i, s) (it integrates out over the possible outcomes of both the investment strategies of competitors (the $\hat{s}_i^{'}$), and over the outside alternative (the ζ)).

Note that $\{w(\nu; i, s)\}$ is all the firm needs to know in order to make optimal decisions. Moreover were it to learn what the $\{w(\nu; i, s)\}$ are the firm could solve its problem by behaving like it is facing a single agent, static, maximizing, problem.

Of course to be in a Markov Perfect Equilibrium, the $\{w(\nu; i, s)\}$ would have to satisfy our equilibrium conditions. On the other hand if one wanted to impose rules directly on how the firm learns about $\{w(\nu; i, s)\}$, one could also determine the play, and the nature of industry evolution, that is likely to emanate from those learning rules.

Fixed Point Computation of $\{w(\nu; i, s)\}$.

Pakes and McGuire's (1994) procedure for calculating the equilibrium is a standard iterative procedure.

- In memory we have estimates of the value function and policies associated with each $(i, s) \in \Omega \times S$.
- The updating *Synchronous*; i.e. it circles through the points in S in some fixed order and updates all estimates associated with *every* $s \in S$ at each iteration.
- Convergence occurs when the values and policies from successive iterations are the same. Converged policies and values satisfy all the properties of equilibrium values and policies.

The updating rules are straightforward. At iteration k the algorithm updates each agents policies at each s as follows:

- calculate $w^{k-1}(\cdot|i,s)$ from the information in memory (as in 1a);
- substitute $w^{k-1}(\cdot)$ for $w(\cdot)$ in (1) and then solve the resultant *single agent* optimization problem for the k^{th} iteration's entry, exit and investment polices at (i,s), and put these policies in memory;
- substitutes the new policies and the w^{k-1} for the w, x and the max operator in (1), and label the result $V^k(\cdot)$;
- calculates $V^k(\cdot) - V^{k-1}(\cdot)$ and then substitutes $V^k(\cdot)$ and the policies from iteration k, for the iteration $k-1$ values that were in memory.

This type of algorithm, together with extensions to it, have been used extensively both in applied theory and for empirical work (for a dated review see Pakes, 2000). The extensions, which were largely done by others, were generally begun as modifications to the model to fit the details of the industry or the analytic question of interest, but often resulted in conceptual contributions. Still the computational burden of the algorithm can be limiting, a fact which is not only a practical limitation to using the algorithm, but is also at least somewhat problematic, as it assumes the agents can find answers that we are having trouble computing.

The Computational Burden of "Pointwise" Fixed Point Algorithms.

The computational burden is (essentially) the product of three factors,

- the number of points evaluated at each iteration;
- the time per point evaluated;
- the number of iterations.

Since each of the $\overline{n}$ active firms can only be at K distinct states, the number of points we need to evaluate at each iteration, or

$$\#S \leq K^{\overline{n}}.$$

Exchangeability, of the value and the policy functions in the state variables of a firm's competitors implies that we do not need to differentiate between two vectors of competitors that are permutations of one another. As a result Pakes (1993) has shown an upper bound for $\#S$ is given by the combinatoric

$$\left(\tfrac{K+\overline{n}-1}{\overline{n}}\right) << K^{\overline{n}},$$

but for $\overline{n}$ large enough this bound is tight.

The time per point is determined primarily by the cost of calculating the expected value of future states conditional on outcomes (of obtaining the $w^j(\cdot; i, s)$ from the information in memory) (the cost of obtaining the optimal polices and the new value function given $w^j(\cdot; i, s)$ is small and does not increase in the number of state variables).

Little is known about the relationship between the number of iterations and the number of state variables of the problem; but it doesn't seem to grow large. However, both the number of states, and the cost of calculating expected future values, grow at least geometrically in the number of firms, and exponentially in the states per firm. These two aspects of the problem are jointly known as the "curse of dimensionality", and clearly limit the applicability of the algorithm (especially for empirical work).

Stochastic Algorithms ("Learning Theory").

The PM (2001) algorithm circumvents the two aspects of the curse of dimensionality for a subclass of the models of interest. It is able to do this because:

1. It never attempts to calculate the needed policies on the entire state space, just on the recurrent class, or R, whose size bears no necessary relationship to the number of state variables.

2. The integral over possible future values needed to determine outcomes is never calculated; rather it is estimated by an average of past outcomes. This generates a tradeoff between the computational burden per point, and the number of iterations needed for a given level of precision. Since the precision of the Monte Carlo estimate does not (necessarily) depend on the dimension of the integral being estimated, while the cost of doing the summation explicitly does, the larger is n the more we expect the tradeoff to favor the Monte Carlo procedure.

The algorithm is *asynchronous* (it only calculates policies for a single location at each iteration). Averages of past outcomes are used to form estimates of the expected discounted value of the future net cash flows that would result from the different possible actions of each of the agents active at that location. The policy with the highest estimated expected discounted value is chosen as the agent's policy for this iteration (hence policies are determined

by a simple, single agent, maximization problem). Jointly, the policies the active agents chose determine the distribution of the next iteration's state. A random draw from that distribution is then used to update both the location of the algorithm, and the estimate of the integral determining the expected discounted value of future net cash flows from that location.

Note on Learning. 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.

Details.

Go back to the Bellman equation written as

$$V(i,s) = \max\{\phi, \pi(i,s) - \sup_{x \geq 0}[-cx + \beta \sum_{\nu} w(\nu; i, s) p(\nu|x_1)]\}.$$

where

$$w(\nu; i, s) \equiv \sum_{(\hat{s}_i', \zeta)} V(i + \nu - \zeta, \hat{s}_i' + e(i + \nu - \zeta)) q[\hat{s}_i'|i, s, \zeta] p(\zeta)$$

Recall that because the values of w determine the values of possible investment outcomes, they determine x. Further x and $w(\cdot)$ determine whether a firm should exit. So if either the agent, or the computer, knew the value of the $w(\cdot)$ associated with a location, it could make optimal policy choices *without* either

- specifying a $q[\cdot|\cdot]$ (i.e. forming a perceived distribution of its competitors future states),
- explicitly calculating the summation in the Bellman equation (which was one of the sources of the curse of dimensionality).

As a result the search for equilibrium policies can be recast as a search for the $w(\cdot)$ that satisfy the fixed point implicit in the definition of the value function in terms of $w(\cdot)$, say $V(\cdot, \cdot|w)$, and the definition of $w(\cdot)$ in terms of the value function.

Note also that

- Any set of $w(\cdot;\cdot,\cdot)$ can be used to determine a set of policies,
- these policies will define a "de facto" $q^w[\cdot|\cdot]$
- if these policies were followed, either "nature" or a computer could choose the market outcome (using the implied probabilities)
- this outcome can be viewed as a random realization from the integral defining components of w,
- as a result the agents might use that outcome to update their estimates of w.

This is precisely how the iterations of the artificial intelligence algorithms work. Each iteration is defined by its estimate of w, say w^j, and its location, an s^j. To iterate we must update both of these.

The location is updated by assuming that each firm believes that the current w^j provides the true conditional continuation values. Then the optimal policies are given by

$$\max_{\chi\in\{0,1\}}\{[1-\chi]\phi$$
$$+\chi sup_x[\pi(i,s^j)-cx+\beta\sum_{\nu_1}w^j(\nu;i,s^j)p(\nu|x_1)]\}.$$

The solutions to this problem determine a distribution of transitions for s. s^{j+1} is determined by taking a random draw from that distribution. I.e. first determine who remains active, then choose the next location for each firm that remains active, then reorder the data into a new s vector.

We now need to update the $\{w\}$. We only update the w of location s^j. Hence if s wanders into $\mathcal{R}$ and stays there, we eventually only update on $\mathcal{R}$. For an agent at location (i,s^j), compute

$$V(i+\nu-\zeta^{j+1},\hat{s}_i^{j+1}+e(i+\nu-\zeta^{j+1})|w^j),$$

for *all possible* values of ν (this is computed as the solution to the single agent problem from the w^j in memory). This the j^{th} period evaluation of being in location $(i+\nu-\zeta^{j+1})$ when all the other competitors states are determined by their simulated draws. To update, we treat this value as a random realization from the integral determining $w^j(\cdot)$. More precisely

$$w^{j+1}(\nu_1; i, s^j) - w^j(\nu_1; i, s^j) = \alpha(j, s^j) \times$$
$$[V(i + \nu - \zeta^{j+1}, \hat{s}_i^{j+1} + e(i + \nu - \zeta^{j+1})|w^j)$$
$$-w^j(\nu; i, s^j)].$$

Note that if $w^j = w^*$, then the expectation of

$$V(i + \nu - \zeta^{j+1}, \hat{s}_i^{j+1} + e(i + \nu - \zeta^{j+1})|w^j)$$

is $w_j^*(\cdot)$. So if $w^j = w^*$ then the expectation of $w^{j+1}(\cdot)$ is $w^j(\cdot)$, and there is a tendency to stay at w^* (a fixed point).

Also if $\alpha(j, s^j)$ equals the inverse of the number of times the point being updated has been visited (say $1/h^j(i, s)$), then

$$w^j(\nu_1; i, s) =$$
$$h^j(i, s)^{-1} \sum_{\{\cdot\}} V(i + \nu - \zeta^{q+1}, \hat{s}_i^{q+1} + e(i + \nu - \zeta^{q+1})|w^q),$$

where it is understood that the summation is over all draws from the point (i, s). I.e. w^j is just an average of past random draws on the conditional continuation values at that point[7].

There is no guarantee that this algorithm will converge to an MPE, but PM (2001) provide a test of whether it has converged. The test essentially computes the fixed point on the recurrent class, and is therefore inefficient (it requires us to compute the integral determining w). Hence work is now being done on constructing more efficient tests. Nonetheless it is instructive to look at the results from the computations which use the "inefficient test".

The relevant comparisons are in terms of memory requirements and CPU time. The ratio of memory requirements in the two algorithms is effectively $\#R/\#S$. $\#R/\#S \approx 3.3\%$ when $\overline{n} = 6$, about .4% when $\overline{n} = 10$ (then $\#S \approx 3.2 \times 10^7$), and would decline further for larger $\overline{n}$.

[7]In general there are "better" ways of choosing the $\alpha(j, s^j)$ than this (since the earlier outcomes should be less precise estimates of the numbers we are after, the better ways usually downweight them). On the other hand all proposed weighting schemes should satisfy Robbins and Monroe's (1956) convergence conditions (the sum of the weights increase without bound as the number of hits increase, but the sum of the squared weights remains bounded).

Table 1: *Comparisons for Increasing Market Size*[a].

M =	5	6	7	8	9	10
Percentage of equilibria with n firms active						
n=						
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
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

[a] All runs were on a Sun SPARCStation 2.

The CPU time of the stochastic algorithm is determined by the time per iteration, the number of iterations, and the test time. The time per iteration is essentially a function of the number of firms active so that the ratio of the average number of active firms to the time per million iterations only varied between 1.53 and 1.65 minutes. The number of iterations until our test criteria was satisfied varied quite a bit between runs; but did not tend to increase in M. Thus *absent* test times, the average CPU time needed for our algorithm seems to grow *linearly* in the average number of firms active.

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 $M = 10$ the algorithm spends ten times as much time computing the test statistic after each million iterations as it spends on the iterations themselves.

Comparing to CPU times from the backward solution algorithms we find

that when $\overline{n} = 6$ the stochastic algorithm took about half the CPU time (about a third if one ignores test time). When $\overline{n} = 10$ even the most optimistic projection for the backward techniques leads to a ratio of CPU times of .4% (.09% without test times). 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.

References

Aguirreberia, V., and P. Mira (2002) "Identification and Estimation of Dynamic Discrete Games", Boson University mimeo.

Berry, S. (1992) "Estimation of a Model of Entry in the Airline Industry", *Econometrica*, **60**, 889–917.

Brenahan, T., and P. Reiss (1987) "Do Entry Conditions Vary Across Markets?", *Brookings Papers on Economic Activity, Microeconomics Annual*, **1**, 833–882.

Bresnahan, T. and P. Reiss (1990) "Entry in Monopoly Markets", *Review of Economic Studies* **57**, 531–553.

Doraszelski, U. and M. Satterwaite (2003)"Foundations of Markov-Perfect Dynamics: Existence, Purification, and Multiplicity", *mimeo* Northwestern University.

Ericson, R., and A. Pakes (1995) "Markov-Perfect Industry Dynamics: A Framework for Empirical Work", *Review of Economic Studies*, **62**, 53–82.

Freedman, D. (1983) *Markov Chains.* Springer Verlag.

Mazzeo, M. (2002) "Product Choice and Oligopoly Market Structure", *RAND Journal of Economics*, **33**, 221–242.

Pakes, A (1993) "The Estimation of Dynamic Structural Models; Problems and Prospects" in *Advances in Economic Theory: Proceedings of the 6^{t}h World Congress of the Econometric Society*, edited by J.J.Laffont and C. Sims.

Pakes, A (2000) "A Framework for Applied Dynamic Analysis in I.O." N.B.E.R. discussion paper 8024, forthcoming in the *Handbook of Industrial Organization.*

Pakes, A., and P. McGuire (1994) "Computing Markov-Perfect Nash Equilibria: Numerical Implications of a Dynamic Differentiated Product Model", *RAND Journal of Economics*, **25**, 555–589.

Pakes, A., and P. McGuire (2001) "Stochastic Algorithms, Symmetric Markov Perfect Equilibria and The 'Curse' of Dimensionality", *Econometrica*, **69**, 1261–1281.

Pakes, A., M. Ostrovsky, and S. Berry (2003) "Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples)", Harvard University mimeo.

Rust, J. (1994) "Estimation of Dynamic Structural Models: Problems and Prospects", in C. Sims, ed. *Advances in Econometrics. Sixth World Congress of the Econometric Society.* Cambridge University Press.

Robbins, H., and S. Monroe, (1951), "A Stochastic Approximation Method", *Annals of Mathematics and Statistics*, 22, pp.400-407.

Seim, K. (2002) "Geographic Differentiation and Firms' Entry Decisions: The Video Retail Industry", Stanford University Graduate School of Business mimeo.

Sutton, J. (1991) *Sunk Costs and Market Structure.* M.I.T. Press.

Tamer, E. (2003) "Incomplete Simultaneous Discrete Response Model with Multiple Equilibria", *Review of Economic Studies* **70**, 147–165.