Industrial Organization

- This note briefly introduces a Dynamic Single Agent example of investment and then motivates extensions. The objective is to connect some of our work in the course to an example from Industrial Organization. The material here follows Pakes (1994, and 1996) rather closely.

- The example is a standard investment problem modified to allow for stochastic outcomes to the investment process, and to allow for exit.

  We first need to define some primitives:

- Call $\pi(\omega, z) : \ \Omega \times Z \ \rightarrow R$, provides the profits accrue given an (ω, z) couple. $s \equiv (\omega, z)$ is called the state of the system. It determines profits.

- $\omega \in \Omega$ is the state variable whose evolution we can affect using investment. Think of it as knowledge, or advertising, physical capital, or the stock of the resource we are exploring.

- The stochastic process generating the sequence of realization $\{\omega_t\}$ is a "controlled" Markov process. By this we mean that the distribution of ω_{t+1} is determined by ω_t, the Markov component, and the control (investment) x_t.

- We need a set of distribution functions, one for each possible (ω_t, x_t). This determines the distribution of ω_{t+1} given the (ω_t, x_t):

$$\Lambda_\omega = \{P(\cdot|\omega, x);\ \omega \in \Omega,\ x \in X \subset R^+\}. \tag{1}$$

  The family Λ_ω is assumed to be stochastically increasing in the natural order of both Ω and X. Meaning, that ceteris paribus, the higher current ω the higher future ω is likely to be, and the higher current x the higher future ω is likely to be.

- We will need a cost of investment that we will designate as $c(x)$ or just x.

- z designates an exogenous variable that helps determine profits. Examples include factor costs, the general state of demand, etc. The difference between ω and z is that we have no control over the sequence $\{z_t\}$. It evolves as an exogenous stochastic process with family

$$\Lambda_z = \{P(\cdot|z),\ z \in Z \in R\}, \tag{2}$$

  assumed stochastically increasing in z.

- Let superscript t denote the history of a variable until t, and subscript t denote an actual year, so the history of states is

$$s_t \equiv \{s_1, \cdots, s_t\} \in S^t \equiv (S \times S \times \cdots \times S). \tag{3}$$

- Call $\psi_t(S^t)\ :\ S^t \rightarrow X$ for $t = 1, 2, \cdots$ a policy function for period t. This maps the observed history until that t into a choice for investment at that t. It is a function because it is defined for all possible histories.

- The policy is a sequence of such policy functions, one for each possible future t, and will be denoted by ψ. It is an infinite sequence of functions.

- We say that the policy ψ is feasible from s_1 iff $\psi_t(s^t) \in X, \forall t$. And let $\Psi(s_1)$ be the set of all feasible policies.

- We will say that the objective of the decision maker is to choose a $\psi \in \Psi(s_1)$ to maximize the expected discounted value of future net cash flows. To do this we have to calculate the EDV of each ψ and that is given by

$$V_\psi(s_1) \equiv E_\psi\{\sum_{t=0}^{\infty} \beta^t(\pi(s_t) - c(\psi_t(s^t)))|s_1\}, \tag{4}$$

  where it is understood that E_ψ refers to the expectation given that policy ψ is followed. But this is equal to

$$\sum_{t=0}^{\infty} \beta^t E_\psi\{\pi(s_t) - c(\psi_t(s^t))|s_1\} = \sum_{t=0}^{\infty} \beta^t \sum_{s_t} \left[\pi(s_t) - c(\psi_t(s^t))\right]\ p(s_t|\psi, s_1). \tag{5}$$

- Now define $V^*(s_1) = \sup_{\psi \in \Psi(s_1)} V_\psi(s_1)$ and $\psi^*(s_1)$ to be the corresponding policy. We want to find these two, they are functions, because they take values for each $s \in S$.

And then there was Bellman

- We can rewrite the problem in terms of a single unknown function:

$$v(s) = \sup_{x \in X}\{\pi(s) - c(x) + \beta \int_{\omega', z'} v(\omega', z') p(\omega'|\omega, z) p(z'|z).\} \tag{6}$$

  This function solves for a function in terms of itself, or specifies a fixed point in a space of functions. It says that the value today is current returns plus the expected discounted value of tomorrow given optimal behavior in the interim.

- Note that we have transformed the problem from a problem of finding an infinite sequence of functions of increasing complexity to a problem of finding a single function. Given the value function we can directly calculate optimal policies.

- To make the problem a bit more realistic let the firm have the option of exiting the market. If it exits then it receives a one time payoff of ϕ dollars and never reappears again.

- The functional equation for our problem then becomes:

$$v(s) = \max\{\phi, \sup_{x \in X}\{\pi(s) - c(x) + \beta \int_{\omega', z'} v(\omega', z') p(\omega'|\omega, z) p(z'|z)\}\}. \qquad (7)$$

- It is possible to show using the properties of Dynamic Programming that the value function is weakly increasing in ω, and increasing in z.

- Notice that once we have the problem in this framework we can use the techniques we know to solve this problem. We can use value function iteration, policy function iteration, parametric policy iteration, smooth approximations to the value function or the decision rule, etc.

- Notice that in this problem the controls are investment and the exiting decision. Given this setup it can be shown (Pakes 1994) that the exiting rule is a stopping rule that depends on the evolution of the investment

Market Interactions

- One of the main distinctions between how Industrial Organization approaches most problems and how we have been doing so, is that agents (market) interactions are not only important but the key to the interesting features of most models, and the source of complications and computational burdens.

- In our simple framework we can incorporate market interactions allowing the returns an agent earns in a given period to depend not only on the value of the agent's own state vector (s), but also on the vector of state variables of the other agents active in the market $(\hat{y})$. Consider $y = (\hat{y}, s)$ to be the list of state variables of all active agents.

- We will assume that there is a finite upper bound to the number of agents active in the market in a given period (usually this will be a consequence of the primitives of the model).

- A particular value of y is a finite list of the state vectors of the firms currently active in the industry, and will be called an industry structure.
- This type of market interaction already makes us realize that a firm's investment profitability depends on the investments of its (potential and actual) competitors.
- As before we are going to assume that all decisions are made to maximize the expected discounted value of the future net cash flow conditional on the current information set. That information set includes a distribution for the counting measure of possible industry structures in future years, conditional on the current structure.
- The equilibrium notion used to close the model insists that this distribution is in fact consistent with optimal behavior by all incumbents and potential entrants.
- To simplify this potentially very complex problem in most cases practitioners restrict the state variables from their competitors to include either current production costs or current demand conditions. This is for example relaxed when presenting models with collusion.
- Also, it is assumed that decisions are not made simultaneously by all agents, but rather there are alternating moves. In period one the first agent chooses its strategy, and then the second agent is constrained to the implications of history. In period two, the second agent moves but not the first, etc.
- The agents current strategy then will, in addition to having an effect on current net cash flow, determine the state which the competitor will have to deal with in the next period. The control is chosen strategically, with the response of the competitor in mind. If on top of this assumption we assume a finite horizon it is possible to show that we have a generic unique equilibrium and we can develop an algorithm to compute the equilibrium.
- All these assumptions basically mean that we are going to restrict attention to Markov Perfect Nash Equilibrium in investment strategies (see Maskin and Tirole 1987 in EER, and two other papers in 1988 in Econometrica).
- Note that here we are not assuming for important issues like non-pecuniary spillovers among firms, or asymmetric information.

- We will not get into estimation issues here, but in Pakes' opinion the strategy of estimating the model's parameters by solving for the complete set of dynamic equilibrium responses for different candidate values of the parameter vector, and then fitting these into an iterative Maximum Likelihood or Minimum Distance search procedure has computational and data requirements unlikely to be satisfied.

- The strategy is then to break the estimation problem in smaller parts, allowing the researcher to obtain an estimator of a sub-vector of the total vector of the model's parameters. This estimator should be consistent and asymptotically normal under the complete set of equilibrium assumptions.

- Thanks to concentrating on MPNE it is easy to break the problem into estimating a static return function and then estimating the impact of investment on the changes in the state variables.

- On the other hand notice that once we have our estimated parameters, we will still want to use them to compute the equilibrium they imply, and then investigate how they change with policy and environmental changes. So we still need the algorithm to compute equilibrium responses, but it does not have to be as fast as it would be required in order to embed it into a Maximum Likelihood routine. This algorithms, however, are evolving and improving all the time (see Pakes and McGuire 2001).