In these lecture notes we consider specification and estimation of dynamic optimization model. Focus on single-agent models.

1 Rust (1987)

Rust (1987) is one of the first papers in this literature. Model is quite simple, but empirical framework introduced in this paper for dynamic discrete-choice (DDC) models is still widely applied.

Agent is Harold Zurcher, manager of bus depot in Madison, Wisconsin. Each week, HZ must decide whether to replace the bus engine, or keep it running for another week. This engine replacement problem is an example of an *optimal stopping* problem, which features the usual tradeoff: (i) there are large fixed costs associated with "stopping" (replacing the engine), but new engine has lower associated future maintenance costs; (ii) by not replacing the engine, you avoid the fixed replacement costs, but suffer higher future maintenance costs. Optimal solution is characterized by a threshold-type of rule: there is a "critical" cutoff mileage level x^* below which no replacement takes place, but above which replacement will take place.

(Another well-known example of optimal stopping problem in economics is job search model: each period, unemployed worker decides whether to accept a job offer, or continue searching. Optimal policy is characterized by "reservation wage": accept all job offers with wage above a certain threshold.)

1.1 Behavioral Model

At the end of each week t, HZ decides whether or not to replace engine. Control variable defined as:

$$i_t = \begin{cases} 1 & \text{if HZ replaces} \\ 0 & \text{otherwise.} \end{cases}$$

For simplicity. we describe the case where there is only one bus (in the paper, buses are treated as independent entities).

HZ chooses the (infinite) sequence $\{i_1, i_2, i_3, \ldots, i_t, i_{t+1}, \ldots\}$ to maximize discounted expected utility stream:

$$\max_{\{i_1,i_2,i_3,\dots,i_t,i_{t+1},\dots\}} E\sum_{t=1}^{\infty}\beta_{t-1}u\left(x_t,\epsilon_t,i_t;\theta\right) \tag{1}$$

where

- x_t is the mileage of the bus at the end of week t. Assume that evolution of mileage is stochastic (from HZ's point of view) and follows

$$x_{t+1}\begin{cases}\sim G(x'|x_t) & \text{if } i_t=0 \text{ (don't replace engine in period } t)\\ =0 & \text{if } i_t=1\text{: once replaced, bus is good as new}\end{cases} \tag{2}$$

  and $G(x'|x)$ is the conditional probability distribution of next period's mileage x' given that current mileage is x. HZ knows G; econometrician knows the form of G, up to a vector of parameters which are estimated.

- ϵ_t denotes shocks in period t, which affect HZ's choice of whether to replace the engine. These are the "structural errors" of the model (they are observed by HZ, but not by us), and we will discuss them in more detail below.

- Since mileage evolves randomly, this implies that even given a sequence of replacement choices $\{i_1,i_2,i_3,\dots,i_t,i_{t+1},\dots\}$, the corresponding sequence of mileages $\{x_1,x_2,x_3,\dots,x_t,x_{t+1},\dots\}$ is still random. The expection in Eq. (1) is over this stochastic sequence of mileages and over the shocks $\{\epsilon_1,\epsilon_2,\dots\}$.

- The state variables of this problem are:

  1. x_t: the mileage. Both HZ and the econometrician observe this, so we call this the "observed state variable"
  2. ϵ_t: the utility shocks. Econometrician does not observe this, so we call it the "unobserved state variable"

Define value function:

$$V(x_t,\epsilon_t)=\max_{i_\tau,\ \tau=t+1,t+2,\dots} E_t\left[\sum_{\tau=t+1}^{\infty}\beta_{\tau-t}u\left(x_t,\epsilon_t,i_t;\theta\right)|x_t\right]$$

where maximum is over all possible sequences of $\{i_{t+1}, i_{t+2}, \dots\}$. Note that we have imposed stationarity, so that the value function $V(\cdot)$ is a function of t only indirectly, through the value that the state variable x takes during period t.

(An important distinction between empirical papers with dynamic optimization models is whether agents have infinite-horizon, or finite-horizon. Stationarity (or time homogeneity) is assumed for infinite-horizon problems, and they are solved using value function iteration. Finite-horizon problems are non-stationary, and solved by backward induction starting from the final period. Most structural dynamic models used in labor economics are finite-horizon.)

Using the Bellman equation, we can break down the DO problem into an (infinite) sequence of single-period decisions:

$$i_t = i^*(x_t, \epsilon_t; \theta) = \text{argmax}_i \left\{ u(x_t, \epsilon_t, i; \theta) + \beta E_{x', \epsilon' | x_t, \epsilon_t, i_t} V(x', \epsilon') \right\}$$

where the value function is

$$\begin{aligned} V(x, \epsilon) &= \max_{i=1,0} \left\{ u(x, \epsilon, i; \theta) + \beta E_{x', \epsilon' | x_t, \epsilon_t, i_t} V(x', \epsilon') \right\} \\ &= \max \left\{ u(x, \epsilon, 0; \theta) + \beta E_{x', \epsilon' | x_t, \epsilon_t, i_t} V(x', \epsilon'), u(x, \epsilon, 1; \theta) + \beta V(0). \right\} \\ &= \max \left\{ \tilde{V}(x_t, \epsilon_t, 1), \tilde{V}(x_t, \epsilon_t, 0) \right\}. \end{aligned} \tag{3}$$

In the above, we define the choice-specific value function

$$\tilde{V}(x_t, \epsilon_t, i_t) = \begin{cases} u(x, \epsilon, 1; \theta) + \beta V(0) & \text{if } i_t = 1 \\ u(x, \epsilon, 0; \theta) + \beta E_{x', \epsilon' | x_t, \epsilon_t, i_t} V(x', \epsilon') & \text{if } i_t = 0. \end{cases}$$

We make the following parametric assumptions on utility flow:

$$u(x, \epsilon, i; \theta) = -c\left((1-i) * x; \theta\right) - i * RC + \epsilon_i$$

where

- $c(\cdots)$ is the maintenance cost function, which is presumably increasing in x (higher x means higher costs)
- RC denotes the "lumpy" fixed costs of adjustment. The presence of these costs implies that HZ won't want to replace the engine every period.

- ϵ_i, $i = 0, 1$ are structural errors, which represents factors which affect HZ's replacement choice i_t in period t, but are unobserved by the econometrician.

  As Rust remarks (bottom, pg. 1008), you need this in order to generate a positive likelihood for your observed data. Without these ϵ's, we observed as much as HZ does, and model will not be able to explain situations where (say) mileage was 20,000, but in one case HZ replaces, and in second case HZ doesn't replace.

As remarked earlier, these assumption imply a very simple type of optimal decision rule $i^*(x, \epsilon; \theta)$: in any period t, you replace when $x_t \geq x^*(\epsilon_t)$, where $x^*(\epsilon_t)$ is some optimal cutoff mileage level, which depends on the value of the shocks ϵ_t.

Parameters to be estimated are:

1. parameters of maintenance cost function $c(\cdots)$;
2. replacement cost RC;
3. parameters of mileage transition function $G(x'|x)$.

Note: in these models, the discount factor β is typically not estimated. Essentially, the time series data on $\{i_t, x_t\}$ could be equally well explained by a myopic model, which posits that

$$i_t = \text{argmax}_{i \in \{0,1\}} \left\{ u(x_t, \epsilon_t, 0), u(x_t, \epsilon_t, 1) \right\},$$

or a forward-looking model, which posits that

$$i_t = \text{argmax}_{i \in \{0,1\}} \left\{ \tilde{V}(x_t, \epsilon_t, 0), \tilde{V}(x_t, \epsilon_t, 1) \right\}.$$

In both models, the choice i_t depends just on the current state variables x_t, ϵ_t. Indeed, Magnac and Thesmar (2002) shows that in general, DDC models are nonparametrically underidentified, without knowledge of β or $G(\epsilon)$, the distribution of the ϵ shocks.

Intuitively, in this model, it is difficult to identify β apart from fixed costs. In this model, if HZ were myopic (ie. β close to zero) and replacement costs RC were low, his decisions may look similar as when he were forward-looking (ie. β close to 1) and RC were large. Reduced-form tests for forward-looking behavior exploit scenarios in which some variables which affect future utility are known in period t: consumers are deemed forward-looking if their period t decisions depends on these variables. (Example: Chevalier and Goolsbee (2005) examine whether students' choices of purchasing a textbook now depend on the possibility that a new edition will be released soon.)

1.2 Econometric Model

Data: observe $\{i_t, x_t\}$, $t = 1, \ldots, T$ for 62 buses. Treat buses as homogeneous and independent (ie. replacement decision on bus i is not affected by replacement decision on bus j).

Rust makes the following conditional independence assumption, on the Markovian transition probabilities in the Bellman equation above:

$$p(x', \epsilon'|x, \epsilon, i) = p(\epsilon'|x') \cdot p(x'|x, i). \tag{4}$$

Namely, two types of conditional independence: (i) given x, ϵ's are independent over time; and (ii) conditional on x and i, x' is independent of ϵ.

Likelihood function for a single bus:

$$\begin{aligned} & l\left(x_1, \ldots, x_T, i_t, \ldots, i_T | x_0, i_0; \theta\right) \\ = & \prod_{t=1}^{T} Prob\left(i_t, x_t | x_0, i_0, \ldots, x_{t-1}, i_{t-1}; \theta\right) \\ = & \prod_{t=1}^{T} Prob\left(i_t, x_t | x_{t-1}, i_{t-1}; \theta\right) \\ = & \prod_{t=1}^{T} Prob\left(i_t | x_t; \theta\right) \times Prob\left(x_t | x_{t-1}, i_{t-1}; \theta_3\right). \end{aligned} \tag{5}$$

The third line arises from the Markovian feature of the problem, and the last equality arises due to the conditional independence assumption.

Given the factorization of the likelihood function above, we can estimate in two steps:

1. Estimate θ_3, the parameters of the Markov transition probabilities for mileage,

conditional on non-replacement of engine (i.e., $i_t = 0$). (Recall that $x_{t+1} = 0$ wp1 if $i_t = 1$.)

We assume a discrete distribution for $\Delta x_t \equiv x_{t+1} - x_t$, the incremental mileage between any two periods:

$$\Delta x_t = \begin{cases} [0, 5000) & \text{w/prob } p \\ [5000, 10000) & \text{w/prob } q \\ [10000, \infty) & \text{w/prob } 1 - p - q \end{cases}$$

so that $\theta_3 \equiv \{p, q\}$, with $0 < p, q < 1$ and $p + q < 1$.

This first step can be executed separately from the more substantial second step.

2. Estimate θ, parameters of maintenance cost function $c(\cdots)$ and engine replacement costs.

Here, we make a further assumption that the ϵ's are distributed i.i.d. (across choices and periods), according to the Type I extreme value distribution.

Expand the expression for $Prob(i_t = 1|x_t; \theta)$ equals

$$\begin{aligned} & Prob \left\{ -c(0;\theta) - RC + \epsilon_{1t} + \beta V(0) > -c(x_t;\theta) + \epsilon_{0t} + \beta E_{x',\epsilon'|x_t,\epsilon_t,i_t=0} V(x',\epsilon') \right\} \\ =& Prob \left\{ \epsilon_{1t} - \epsilon_{0t} > c(0;\theta) - c(x_t;\theta) + \beta \left[E_{x',\epsilon'|x_t,\epsilon_t,i_t=0} V(x,\epsilon) - V(0) \right] + RC \right\} \end{aligned}$$

Next, we make a further assumption that the ϵ's are distributed i.i.d. (across choices and periods), according to the Type I extreme value distribution. Then the replacement probability simplifies further to a multinomial logit-like expression:

$$= \frac{\exp\left(-c(0;\theta) - RC + \beta V(0)\right)}{\exp\left(-c(0;\theta) - RC + \beta V(0)\right) + \exp\left(-c(x_t;\theta) + \beta E_{x',\epsilon'|x_t,\epsilon_t,i_t=0} V(x',\epsilon')\right)}$$

where the last line follows if we assume that ϵ_{1t} and ϵ_{0t} are independent, and each is distributed iid TIEV, also independently over time. This is called a “dynamic logit” model, in the literature.

We use the notation $u(x, i; \theta)$ and $\tilde{V}(x, i)$ to denote, respectively, the per-period utility function, and choice-specific value function, minus the additive ϵ error. Then the choice probability takes the form

$$Prob\left(i_t|x_t;\theta\right) = \frac{\exp\left(u(x_t, i_t, \theta) + \beta E_{x',\epsilon'|x_t,\epsilon_t,i_t} V(x',\epsilon')\right)}{\sum_{i=0,1} \exp\left(u(x_t, i, \theta) + \beta E_{x',\epsilon'|x_t,\epsilon_t,i} V(x',\epsilon')\right)}. \quad (6)$$

1.2.1 *Estimation method for second step: Nested fixed-point algorithm*

The second-step of the estimation procedures is via a "nested fixed point algorithm".

Outer loop: search over different parameter values $\hat{\theta}$.

Inner loop: For $\hat{\theta}$, we need to compute the value function $V(x;\hat{\theta})$. After $V(x,\epsilon;\hat{\theta})$ is obtained, we can compute the LL fxn in Eq. (6).

1.2.2 *Computational details for inner loop*

Compute value function $V(x;\hat{\theta})$ by iterating over Bellman's equation (3).

A clever and computationally convenient feature in Rust's paper is that he iterates over the *expected* value function $EV(x,i) \equiv E_{x',\epsilon'|x,i}V(x',\epsilon';\theta)$. The reason for this is that you avoid having to calculate the value function at values of ϵ_0 and ϵ_1, which are additional state variables. He iterates over the following equation (which is Eq. 4.14 in his paper):

$$EV(x,i) = \int_y \log\left\{\sum_{j\in C(y)} \exp\left[u(y,j;\theta) + \beta EV(y,j)\right]\right\} p(dy|x,i) \tag{7}$$

Somewhat awkward notation: here "EV" denotes a function. Here x,i denotes the *previous* period's mileage and replacement choice, and y,j denote the *current* period's mileage and choice (as will be clear below).

This equation can be derived from Bellman's equation (3):

$$\begin{aligned}
V(y,\epsilon;\theta) =& \max_{j\in 0,1}\left[u(y,j;\theta) + \epsilon + \beta EV(y,j)\right] \\
\Rightarrow E_{y,\epsilon}\left[V(y,\epsilon;\theta) \mid x,i\right] \equiv EV(x,i;\theta) =& E_{y,\epsilon|x,i}\left\{\max_{j\in 0,1}\left[u(y,j;\theta) + \epsilon + \beta EV(y,j)\right]\right\} \\
=& E_{y|x,i}E_{\epsilon|y,x,i}\left\{\max_{j\in 0,1}\left[u(y,j;\theta) + \epsilon + \beta EV(y,j)\right]\right\} \\
=& E_{y|x,i}\log\left\{\sum_{j=0,1}\left[u(y,j;\theta) + \beta EV(y,j)\right]\right\} \\
=& \int_y \log\left\{\sum_{j=0,1}\left[u(y,j;\theta) + \beta EV(y,j)\right]\right\} p(dy|x,i).
\end{aligned}$$

The next-to-last equality uses the closed-form expression for the expectation of the maximum, for extreme-value variates.

Once the $EV(x, i; \theta)$ function is computed for θ, the choice probabilities $p(i_t|x_t)$ can be constructed as

$$\frac{\exp\left(u(x_t, i_t; \theta) + \beta EV(x_t, i_t; \theta)\right)}{\sum_{i=0,1} \exp\left(u(x_t, i; \theta) + \beta EV(x_t, i; \theta)\right)}.$$

The value iteration procedure: The expected value function $EV(\cdots; \theta)$ will be computed for each value of the parameters θ. The computational procedure is iterative.

Let τ index the iterations. Let $EV^\tau(x, i)$ denote the expected value function during the τ-th iteration. (We suppress the functional dependence of EV on θ for convenience.) Let the values of the state variable x be discretized into a grid of points, which we denote $\vec{r}$.

- $\tau = 0$: Start from an initial guess of the expected value function $EV(x, i)$. Common way is to start with $EV(x, i) = 0$, for all $x \in \vec{r}$, and $i = 0, 1$.
- $\tau = 1$: Use Eq. (7) and $EV^0(x; \theta)$ to calculate, at each $x \in \vec{r}$, and $i \in \{0, 1\}$.

$$\begin{aligned} V^1(x, i) =& \int_y \log\left\{\sum_{j \in C(y)} \exp\left[u(y, j; \theta) + \beta EV^0(y, j)\right]\right\} p(dy|x, i) \\ =& p \cdot \int_x^{x+5000} \log\left\{\sum_{j \in C(y)} \exp\left[u(y, j; \theta) + \beta EV^0(y, j)\right]\right\} dy + \\ & q \cdot \int_{x+5000}^{x+10000} \log\{\cdots\}\, dy + (1 - p - q) \cdot \int_{x+10000}^{\infty} \log\{\cdots\}\, dy. \end{aligned}$$

Now check: is $EV^1(x, i)$ close to $EV^0(x, i)$? One way is to check whether

$$sup_{x,i}|EV^1(x, i) - EV^0(x, i)| < \eta$$

where η is some very small number (eg. 0.0001). If so, then you are done. If not, then

  - Interpolate to get $V^1(\cdot, i)$ at all points $x \notin \vec{r}$.
  - Go to next iteration $\tau = 2$.

2 Hotz-Miller approach: avoid numeric dynamic programming

- One problem with Rust approach to estimating dynamic discrete-choice model very computer intensive. Requires using numeric dynamic programming to compute the value function(s) for every parameter vector θ.

- Alternative method of estimation, which avoids explicit DP. Present main ideas and motivation using a simplified version of Hotz and Miller (1993), Hotz, Miller, Sanders, and Smith (1994).

- For simplicity, think about Harold Zurcher model.

- What do we observe in data from DDC framework? For agent i, time t, observe:

  - $\{\tilde{x}_{it}, d_{it}\}$: observed state variables $\tilde{x}_{it}$ and discrete decision (control) variable d_{it}. For simplicity, assume d_{it} is binary, $\in \{0, 1\}$

    Let $i = 1, \ldots, N$ index the buses, $t = 1, \ldots, T$ index the time periods.
  - For Harold Zurcher model: $\tilde{x}_{it}$ is mileage on bus i in period t, and d_{it} is whether or not engine of bus i was replaced in period t.
  - Given renewal assumptions (that engine, once repaired, is good as new), define transformed state variable x_{it}: mileage since last engine change.
  - Unobserved state variables: ϵ_{it}, $i.i.d.$ over i and t. Assume that distribution is known (Type 1 Extreme Value in Rust model)

- In the following, let quantities with hats $\hat{}$'s denote objects obtained just from data.

  Objects with tildes $\tilde{}$'s denote "predicted" quantities, obtained from both data and calculated from model given parameter values θ.

- From this data alone, we can estimate (or "identify"):

– Transition probabilities of observed state and control variables: $G(x'|x,d)$[1], estimated by conditional empirical distribution

$$\hat{G}(x'|x,d) \equiv \begin{cases} \sum_{i=1}^{N}\sum_{t=1}^{T-1} \frac{1}{\sum_i \sum_t \mathbf{1}(x_{it}=x,d_{it}=0)} \cdot \mathbf{1}\left(x_{i,t+1} \leq x', x_{it}=x, d_{it}=0\right), & \text{if } d=0 \\ \sum_{i=1}^{N}\sum_{t=1}^{T-1} \frac{1}{\sum_i \sum_t \mathbf{1}(d_{it}=1)} \cdot \mathbf{1}\left(x_{i,t+1} \leq x', d_{it}=1\right), & \text{if } d=1. \end{cases}$$

– Choice probabilities, conditional on state variable: $\text{Prob}\,(d=1|x)$[2], estimated by

$$\hat{P}(d=1|x) \equiv \sum_{i=1}^{N}\sum_{t=1}^{T-1} \frac{1}{\sum_i \sum_t \mathbf{1}\left(x_{it}=x\right)} \cdot \mathbf{1}\left(d_{it}=1, x_{it}=x\right).$$

Since $\text{Prob}\,(d=0|x) = 1 - \text{Prob}\,(d=1|x)$, we have $\hat{P}(d=0|x) = 1 - \hat{P}(d=1|x)$.

- With estimates of $\hat{G}(\cdot|\cdot)$ and $\hat{p}(\cdot|\cdot)$, as well as a parameter vector θ, you can "estimate" the choice-specific value functions by constructing the sum

$$\begin{aligned}\tilde{V}(x,d=1;\theta) =& u(x,d=1;\theta) + \beta E_{x'|x,d=1} E_{d'|x'} E_{\epsilon'|d',x'}\left[u(x',d';\theta) + \epsilon'\right. \\ & \left. + \beta E_{x''|x',d'} E_{d''|x''} E_{\epsilon'|d'',x''}\left[u(x'',d'';\theta) + \epsilon'' + \beta \cdots\right]\right]\end{aligned}$$

$$\begin{aligned}\tilde{V}(x,d=0;\theta) =& u(x,d=0;\theta) + \beta E_{x'|x,d=0} E_{d'|x'} E_{\epsilon'|d',x'}\left[u(x',d';\theta) + \epsilon'\right. \\ & \left. + \beta E_{x''|x',d'} E_{d''|x''} E_{\epsilon'|d'',x''}\left[u(x'',d'';\theta) + \epsilon'' + \beta \cdots\right]\right].\end{aligned}$$

Here $u(x,d;\theta)$ denotes the per-period utility of taking choice d at state x, *without* the additive logit error. Note that the observation of $d'|x'$ is crucial to being able to forward-simulate the choice-specific value functions. Otherwise, $d'|x'$ is multinomial with probabilities given by Eq. (8) below, and is impossible to calculate without knowledge of the choice-specific value functions.

- In practice, "truncate" the infinite sum at some period T:

$$\begin{aligned}\tilde{V}(x,d=1;\theta) =& u(x,d=1;\theta) + \beta E_{x'|x,d=1} E_{d'|x'} E_{\epsilon''|d',x'}\left[u(x',d';\theta) + \epsilon'\right. \\ & + \beta E_{x''|x',d''} E_{d''|x''} E_{\epsilon'|d'',x''}\left[u(x'',d'';\theta) + \epsilon'' + \cdots\right. \\ & \left.\left.\beta E_{x^T|x^{T-1},d^{T-1}} E_{d^T|x^T} E_{\epsilon^T|d^T,x^T}\left[u(x^T,d^T;\theta) + \epsilon^T\right]\right]\right]\end{aligned}$$

[1]By stationarity, note we do not index the G function explicitly with time t.
[2]By stationarity, note we do not index this probability explicitly with time t.

Also, the expectation $E_{\epsilon|d,x}$ denotes the expectation of the ϵ conditional on choice d being taken, and current mileage x. For the logit case, there is a closed form:

$$E[\epsilon|d,x] = \gamma - \log(Pr(d|x))$$

where γ is Euler's constant (0.577...) and $Pr(d|x)$ is the choice probability of action d at state x.

Both of the other expectations in the above expressions are observed directly from the data.

- Both choice-specific value functions can be simulated by (for $d = 1, 2$):

$$\begin{aligned}\tilde{V}(x,d;\theta) \approx = \frac{1}{S}\sum_s \Big[& u(x,d;\theta) + \gamma - \log(\hat{P}(d|x)) + \beta \Big[u(x'^s, d'^s;\theta) + \gamma - \log(\hat{P}(d'^s|x'^s)) \\ & +\beta \Big[u(x''^s, d''^s;\theta) + \gamma - \log(\hat{P}(d''^s|x''^s)) + \beta \cdots\Big]\Big]\Big]\end{aligned}$$

  where

  - $x'^s \sim \hat{G}(\cdot|x,d)$
  - $d'^s \sim \hat{p}(\cdot|x'^s)$, $x''^s \sim \hat{G}(\cdot|x'^s, d'^s)$
  - &etc.

  In short, you simulate $\tilde{V}(x,d;\theta)$ by drawing S "sequences" of (d_t, x_t) with a initial value of (d,x), and computing the present-discounted utility correspond to each sequence. Then the simulation estimate of $\tilde{V}(x,d;\theta)$ is obtained as the sample average.

- Given an estimate of $\tilde{V}(\cdot,d;\theta)$, you can get the *predicted choice probabilities*:

$$\tilde{p}(d=1|x;\theta) \equiv \frac{exp\left(\tilde{V}(x,d=1;\theta)\right)}{\exp\left(\tilde{V}(x,d=1;\theta)\right) + \exp\left(\tilde{V}(x,d=0;\theta)\right)} \tag{8}$$

  and analogously for $\tilde{p}(d=0|x;\theta)$. Note that the predicted choice probabilities are different from $\hat{p}(d|x)$, which are the actual choice probabilities computed from the actual data. The predicted choice probabilities depend on the parameters θ, whereas $\hat{p}(d|x)$ depend solely on the data.

- One way to estimate θ is to minimize the distance between the predicted conditional choice probabilities, and the actual conditional choice probabilities:

$$\hat{\theta} = \text{argmin}_{\theta} || \hat{\mathbf{p}}(d=1|x) - \tilde{\mathbf{p}}\left(d=1|x;\theta\right) ||$$

where $\mathbf{p}$ denotes a vector of probabilities, at various values of x.

- Another way to estimate θ is very similar to the Berry/BLP method. Given the logit assumption, we can equate the actual conditional choice probabilities $\hat{p}(d|x)$ to the model's predicted choice probabilities $\tilde{p}(d|x;\theta)$ to obtain that

$$\hat{\delta}_x \equiv \log \hat{p}(d=1|x) - \log \hat{p}(d=1|x) = \left[\hat{V}(x,d=1) - \hat{V}(x,d=0)\right].$$

An alternative estimator could proceed by doing

$$\bar{\theta} = \text{argmin}_{\theta} || \hat{\delta}_x - \left[\tilde{V}(x,d=1;\theta) - \tilde{V}(x,d=0;\theta)\right] ||.$$

3 Introduction to structure of dynamic oligopoly models

- Consider a simple two-firm model, and assume that all the dynamics are deterministic.

- Let x_{1t}, x_{2t}, denote the state variables for each firm in each period. Let q_{1t}, q_{2t} denote the control variables. Example: x's are capacity levels, and q's are incremental changes to capacity in each period.

- Assume (for now) that $x_{it+1} = g\left(x_{it}, q_{it}\right)$, $i = 1, 2$, so that next period's state is a *deterministic* function of this period's state and control variable. (Can allow for cross-effects with no problem.)

- Firm i (=1,2) chooses a sequence $q_{i1}, q_{i2}, q_{i3}, \ldots$ to maximize its discounted profits:

$$\sum_{t=0}^{\infty} \beta^t \Pi\left(x_{1t}, x_{2t}, q_{1t}, q_{2t}\right)$$

where $\Pi(\cdots)$ denotes single-period profits.

- Because the two firms are duopolists, and they must make these choices recognizing that their choices can affect their rival's choices. We want to consider a dynamic equilibrium of such a model, when (roughly speaking) each firm's sequence of q's is a "best-response" to its rival's sequence.

- A firm's strategy in period t, q_{it}, can potentially depend on the whole "history" of the game ($\mathcal{H}_{t-1} \equiv \{x_{1t'}, x_{2t'}, q_{1t'}, q_{2t'}\}_{t'=0,\dots,t-1}$), and well as on the time period t itself. This becomes quickly intractable, so we usually make some simplifying regularity conditions:

  - Firms employ *stationary* strategies: so that strategies are not explicitly a function of time t (i.e. they depend on time only indirectly, through the history $\mathcal{H}_{t-1}$). Given stationarity, we will drop the t subscript, and use primes $'$ to denote next-period values.
  - A dimension-reducing assumption is usually made: for example, we might assume that q_{it} depends only on x_{1t}, x_{2t}, which are the "payoff-relevant" state variables which directly affect firm i's profits in period i. This is usually called a "Markov" assumption. With this assumption $q_{it} = q_i(x_{1t}, x_{2t})$, for all t.
  - Furthermore, we usually make a *symmetry* assumption, that each firm employs an identical strategy assumption. This implies that $q_1(x_{1t}, x_{2t}) = q_2(x_{2t}, x_{1t})$.

- To characterize the equilibrium further, assume we have an equilibrium strategy function $q^*(\cdot,\cdot)$. For each firm i, then, and at each state vector x_1, x_2, this optimal policy must satisfy Bellman's equation, in order for the strategy to constitute subgame-perfect behavior:

$$q^*(x_1, x_2) = \text{argmax}_q \left\{\Pi(x_1, x_2, q, q^*(x_2, x_1)) + \beta V(x_1' = g(x_1, q), x_2' = g(x_2, q^*(x_2, x_1)))\right\} \tag{9}$$

  from firm 1's perspective, and similarly for firm 2. $V(\cdot,\cdot)$ is the value function, defined recursively at all possible state vectors x_1, x_2 via the Bellman equation:

$$V(x_1, x_2) = \max_q \left\{\Pi(x_1, x_2, q, q^*(x_2, x_1)) + \beta V(x_1' = g(x_1, q), x_2' = g(x_2, q^*(x_2, x_1)))\right\}. \tag{10}$$

- I have described the simplest case; given this structure, it is clear that the following extensions are straightforward:
  - Cross-effects: $x_i' = g\left(x_i, x_{-i}, q_i, q_{-i}\right)$
  - Stochastic evolution: $x_i'|x_i, q_i$ is a random variable. In this case, replace last term of Bellman eq. by $E\left[V\left(x_1', x_2'\right)|x_1, x_2, q, q_2 = q^*\left(x_2, x_1\right)\right]$.

    This expectation denotes player 1's *equilibrium beliefs* about the evolution of x_1 and x_2 (equilibrium in the sense that he assumes that player 2 plays the equilibrium strategy $q^*(x_2, x_1)$).
  - > 2 firms
  - Firms employ asymmetric strategies, so that $q_1\left(x_1, x_2\right) \neq q_2\left(x_2, x_1\right)$
  - ...
- Computing the equilibrium strategy $q^*(\cdots)$ consists in iterating over the Bellman equation (9). However, the problem is more complicated than the single-agent case for several reasons:
  - The value function itself depends on the optimal strategy function $q^*(\cdots)$, via the assumption that the rival firm is always using the optimal strategy. So value iteration procedure is more complicated:
    1. Start with initial guess $V^0(x_1, x_2)$
    2. If q's are continuous controls, then when strategies are symmetric, then $q^0(x_1, x_2) = q^0(x_2, x_1) \equiv q^0$, and first-order condition defines the unknown q^0:

    $$0 = \Pi_3\left(x_1, x_2, q^0, q^0\right) + \beta V_1^0\left(g\left(x_1, q^0\right), g\left(x_2, q^0\right)\right) \cdot g_2\left(x_1, q^0\right). \quad (11)$$

    3. If q's are discrete, taking values $\in \mathcal{Q}$, then in the symmetric case:

    $$q^0 = \text{argmax}_{q \in \mathcal{Q}} \left\{\Pi\left(x_1, x_2, q, q^0\right) + \beta V^0\left(g\left(x_1, q\right), g\left(x_2, q^0\right)\right)\right\} \quad (12)$$

    4. Symmetry assumption helps a lot: computational problem is essentially the same as single-agent problem (except state space is expanded to include state variables of both firms).

When strategies are asymmetric, then for continuous controls, we must solve for $q_1^0 \equiv q^0(x_1, x_2)$ and $q_2^0 \equiv q^0(x_2, x_1)$ to satisfy the system of first-order conditions (here subscripts denotes partial derivatives)

$$\begin{aligned} 0 &= \Pi_3\left(x_1, x_2, q_1^0, q_2^0\right) + \beta V_1^0\left(g\left(x_1, q_1^0\right), g\left(x_2, q_2^0\right)\right) \cdot g_2\left(x_1, q_1^0\right) \\ 0 &= \Pi_3\left(x_2, x_1, q_2^0, q_1^0\right) + \beta V_1^0\left(g\left(x_2, q_2^0\right), g\left(x_1, q_1^0\right)\right) \cdot g_2\left(x_2, q_2^0\right). \end{aligned} \tag{13}$$

For the discrete control case:

$$\begin{aligned} q^0 &= \text{argmax}_{q\in\mathcal{Q}}\left\{\Pi\left(x_1, x_2, q, q_2^0\right) + \beta V^0\left(g\left(x_1, q\right), g\left(x_2, q_2^0\right)\right)\right\} \\ q_2^0 &= \text{argmax}_{q\in\mathcal{Q}}\left\{\Pi\left(x_2, x_1, q, q_1^0\right) + \beta V^0\left(g\left(x_2, q\right), g\left(x_1, q_1^0\right)\right)\right\}. \end{aligned} \tag{14}$$

5. Update the next iteration of the value function:

$$V^1(x_1, x_2) = \left\{\Pi\left(x_1, x_2, q_1^0, q_2^0\right) + \beta V^0\left(g\left(x_1, q_1^0\right), g\left(x_2, q_2^0\right)\right)\right\}. \tag{15}$$

Note: this and the previous step must be done at all points (x_1, x_2) in the discretized grid. As usual, use interpolation or approximation to obtain $V^1(\cdots)$ at points not on the grid.

6. Stop when $\sup_{x_1,x_2} ||V^{i+1}(x_1, x_2) - V^i(x_1, x_2)|| \leq \epsilon$.

– In principle, one could estimate a dynamic games model, given time series of $\{x_t, q_t\}$ for both firms, by using a nested fixed-point estimation algorithm. In the outer loop, loop over different values of the parameters θ, and then in the inner loop, you compute the equilibrium of the dynamic game (in the way described above) for each value of θ.

– However, there is an inherent "Curse of dimensionality" with dynamic games, because the dimensionality of the state vector (x_1, x_2) is equal to the number of firms. (For instance, if you want to discretize 1000 pts in one dimension, you have to discretize at 1,000,000 pts to maintain the same fineness in two dimensions!)

Some papers provide computational methods to circumvent this problem (Keane and Wolpin (1994), Pakes and McGuire (2001), Imai, Jain, and Ching (2005)). Generally, these papers advocate only computing the value function at a (small) subset of the state points each iteration, and then approximating the value function at the rest of the state points using values calculated during previous iterations.

- Clearly, it is possible to extend the Hotz-Miller insights to facilitate estimation of dynamic oligopoly models, in the case where q is a discrete control. Advantage, as before, is that you can avoid numerically solving for the value function.

  Data directly tell you: the choice probabilities (distribution of $q|x_1, x_2$); state transitions: (joint distribution of $x_1' x_2' | x_1, x_2, q_2, q_2$). This will be a topic in Han Hong's lectures.

■■■

EXTRA TOPIC

4 An example of dynamic oligopoly: automobile market with secondary markets

We go over Esteban and Shum (2004).

In durable goods industries (like car market), secondary markets leads to intertemporal linkages between primary adn secondary markets. Used goods of today were new goods of yesterday.

Interesting economic question: Does this harm or benefit producers?

- Intuition different from static markets:

  Benchmark in DG setting is Coase outcome (firm's inability to commit to low levels of production can erode market power)

- Vs. this benchmark, 2-mkts can benefit producers:

  1. 2-mkt offers substitutes for firms' new production $\Rightarrow$ curtails Coasian tendency to overproduce ("*commitment benefit*")
  2. With heterogeneous consumers, 2-mkts segment market, allow firms to target new goods to high-valuation consumers ("*sorting benefit*")

■■■

Economic Model: Car market

Multiproduct firms producing cars which differ in quality, durability and depreciation schedule.

Empirical model accommodates cost/demand shocks; for simplicity, describe deterministic model.

- Firms $j = 1, \dots, N$. (e.g. Ford, GM, Honda)
- L is total number of brands/models (e.g. Taurus, Accord, Escort).
- Firm j produces L_j models; set of products denoted $\mathcal{L}_j$.
- Model i lasts T_i periods. There are $K \equiv \sum_{i=1}^{L} T_i$ "model-years" actively traded during any given period.
- Each model year differs in one-dimensional quality $\Rightarrow$ quality ladder

  $$[\alpha_1,\ \alpha_2,\ \ \dots,\ \alpha_K, \alpha_{K+1} = 0]$$

  where α_{K+1} is quality of outside option.
- Notation: depreciation schedules for different models
  - Define: $\eta(i)$ is ranking of model i when new.
  - Define: $v(\eta(i))$ is ranking of 1-yr old; $v^2(\eta(i)) \equiv v\left(v(\eta(i)\right)$ is ranking of 2-yr old, etc.

  $\Rightarrow$ Depreciation schedule of model i described by sequence

  $$\left\{\eta\left(i\right), v\left(\eta\left(i\right)\right), \dots, v^{T_i-1}\left(\eta\left(i\right)\right)\right\}.$$

  Note: each model has its own depreciation schedule.
  - Note that firms are *asymmetric*. Hence equilibrium is characterized by set of L Bellman equations (as will be derived below).

■■■

Economic model: Dynamic demand

Derive from individual-level optimizing behavior.

- A continuum of infinitely-lived consumers who differ in their preference for quality θ (one dimension)

- Quasilinear per-period utility: $U_t = \theta\alpha_k + m - p_t^k$, where m is total income. Assume no liquidity constraints.

- Choice set: model-years $k = 1, \ldots, K$, plus outside option (utility normalized =0)

- Consumers incur no **transactions costs**: abstract away from timing issues.

  Implies simple form of dynamic decision rule: in period t, consumer θ chooses model-year k yielding maximal "rental utility":

$$k_t = \text{argmax}_k \left\{ 0, \alpha_k\theta - p_t^k + \delta E_t p_{t+1}^{v(k)},\ k = 1, \ldots, K \right\}$$

  where δ is discount factor and expected rental price is $p_t^k - \delta E_t p_{t+1}^{v(k)}$.

  (Drop E_t for convenience: assume perfect foresight, so $E_t p_{t+1} = p_{t+1}$.)

■■■

Demand functions

- Given prices p_t^k, $p_{t+1}^{v(k)}$ for all $k = 1, \ldots, K$, each period there are K indifferent consumers $\bar{\theta} \geq \tilde{\theta}_t^1 \geq \tilde{\theta}_t^2 \geq \tilde{\theta}_t^3 \geq \ldots \geq \tilde{\theta}_t^K \geq 0$, s.t.

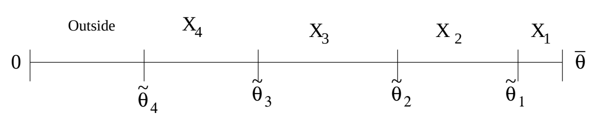

- The indifferent consumers solve

$$\alpha_k\tilde{\theta}_t^k - p_t^k + \delta p_{t+1}^{v(k)} = \alpha_{k+1}\tilde{\theta}_t^k - p_t^{k+1} + \delta p_{t+1}^{v(k+1)}, \quad \text{for } k = 1, \ldots, K-1$$

- Consumer heterogeneity: θ is uniformly distributed.

- Derive inverse demand function (subject to non-negativity constraints)

$$p_t^k = (\alpha_k - \alpha_{k+1})\,\bar{\theta}\left(1 - \frac{1}{M}\sum_{r=1}^{k} x_t^r\right) + \delta p_{t+1}^{\upsilon(k)} + p_t^{k+1} - \delta p_{t+1}^{\upsilon(k+1)}$$

■■■

Supply side

- $\boldsymbol{y}_t = [1, x_t^1, \ldots, x_t^K]'$: vector of all cars *transacted* in period t.
- $\boldsymbol{d}_t \equiv \left[x_t^{\eta(1)}, x_t^{\eta(2)}, \ldots, x_t^{\eta(L)}\right]'$: vector of all cars *produced* in period t.
- Define matrices A and B, to get law of motion for y_t:

$$\boldsymbol{y}_t = \boldsymbol{A}\boldsymbol{y}_{t-1} + \boldsymbol{B}\boldsymbol{d}_t.$$

- Marginal costs constant; no (dis-)economies of scope: $C_{jt} = \sum_{i\in\mathcal{L}} c_i \cdot x_t^{\eta(i)}$.
- Period t profits for car i is

$$\Pi_t^i\left(\boldsymbol{y}_t, \boldsymbol{y}_{t+1}, \ldots, \boldsymbol{y}_{t+T_i-1}\right) = \left(p_t^i(\boldsymbol{y}_t, \boldsymbol{y}_{t+1}, \ldots, \boldsymbol{y}_{t+T_i-1}) - c_t^i\right) \cdot q_t^1 \;:$$

  depends on past, current, future prod'n of car i.

  Important: dependence of current profits on future actions leads to a *time-consistency* problem, which is absent from "usual" dynamic problems. Very roughly, time-inconsistency implies that an agent's optimal action in period t differs depending on whether the agent is deciding in period t, or period $t-1$, or period $t-2$, etc.

  Think of durable goods monopoly: in period 1, his optimal period 2 price is the monopoly price (because that would raise his profits in period 1). But when period 2 comes, his optimal period 2 price is actually a lower price (since he wants to sell to people who did not buy in period 1).

- For individual firm: $\forall t,\ \forall j \in \mathcal{N},\ \forall i \in \mathcal{L}_j$, period-$t$ production $x_t^{\eta(i)}$ maximizes

$$(*) \qquad \max_{x_t^{\eta(i)}, \forall i\in\mathcal{L}_j} \sum_{\tau=0}^{\infty}\sum_{i\in\mathcal{L}_j} \delta^\tau \underbrace{\left[\Pi_{t+\tau}^i\left(\boldsymbol{y}_{t+\tau}, \boldsymbol{y}_{t+\tau+1}, \ldots, \boldsymbol{y}_{t+\tau+T_i-1}\right)\right]}_{\text{period } t+\tau \text{ profits}},\ \ s.t.$$

$$\boldsymbol{y}_{t+\tau} \quad = \boldsymbol{A}\boldsymbol{y}_{t+\tau-1} + \boldsymbol{B}\boldsymbol{d}_{t+\tau} \ , \quad \text{for } \tau = 1, \dots, \infty.$$

- Note: obj fxn different in each period t: usual problem is

$$\max_{\left\{x_t^{\eta(i)},\ i \in \mathcal{L}_j\right\}_{t=0}^{\infty}} \sum_{t=0}^{\infty} \sum_{i \in \mathcal{L}_j} \delta^t \Pi_t^i \left(\boldsymbol{y}_t, \boldsymbol{y}_{t+1}, \dots, \boldsymbol{y}_{t+T_i-1}\right).$$

FOC for $x_t^{\eta(i)}$ contains derivative (say) $\frac{\partial \Pi_{t-1}^i}{\partial x_t^{\eta(i)}}$: in choosing period-$t$ prodn, recognize that it affects period-$(t+1)$ profits $\Rightarrow$ time-inconsistent.

■■■

Time-consistent Equilibrium production

- Restrict attention to Markov strategies: $\boldsymbol{A}\boldsymbol{y}_{t-1}$ is the "payoff-relevant state vector" for period t (stocks of cars produced prior to period t which are still actively traded in period t) $\Longrightarrow$

  Therefore consider production rules $x_t^{\eta(i)} = g_i(\boldsymbol{A}\boldsymbol{y}_{t-1}), \forall i \in \mathcal{L}_j,\ \forall j \in \mathcal{N}$.

- In Markov-perfect equilibrium, $g_1(.), \dots, g_L(.)$ satisfy Bellman equation

$$V_j(\boldsymbol{A}\boldsymbol{y}_{t-1}) = \max_{x_t^{\eta(i)},\ i \in \mathcal{L}_j} \sum_{i \in \mathcal{L}_j} \Pi_j(\boldsymbol{y}_t, \boldsymbol{y}_{t+1}, \dots, \boldsymbol{y}_{t+T_i-1}) + \delta V_j(\boldsymbol{A}\boldsymbol{y}_t)$$

  for all firms j, and the Markov decision rules

$$x_t^i = g_i(\boldsymbol{A}\boldsymbol{y}_{t-1}), \text{ for all } i \in \mathcal{L}_j.$$

  Value function $V_j(\boldsymbol{A}\boldsymbol{y}_{t-1})$=(*) (optimal profits from t onwards).

■■■

Linear Quadratic (LQ) Specification

- We focus on *linear* equilibrium decision rule $x_{t+h} = \boldsymbol{G}\boldsymbol{A}\boldsymbol{y}_{t+h-1}$

  No "trigger" strategies (step functions)

- $\Rightarrow$ Quadratic value function $V(\boldsymbol{A}\boldsymbol{y}_{t-1}) = y_t' \boldsymbol{A}' \boldsymbol{S} \boldsymbol{A} y_t$.

- Bellman equation can be rewritten in matrix notation

$$\boldsymbol{y}'_{t-1}\boldsymbol{A}'\boldsymbol{S}\boldsymbol{A}\boldsymbol{y}_{t-1} = \max_{x_t^i,\ i\in\mathcal{J}} \left\{ \sum_{i\in\mathcal{J}} \left[\sum_{z=0}^{T_i-1} \boldsymbol{y}'_{t+z}\delta^z \boldsymbol{R}_{\upsilon^z(i)}\boldsymbol{y}_t \right] \right\} - \boldsymbol{y}'_t\boldsymbol{C}_j\boldsymbol{y}_t + \boldsymbol{y}'_t\delta\left[\boldsymbol{A}'\boldsymbol{S}_j\boldsymbol{A}\right]\boldsymbol{y}_t$$

- Recursive substitution yields

$$\begin{aligned} &\boldsymbol{y}'_{t-1}\boldsymbol{A}'\boldsymbol{S}\boldsymbol{A}\boldsymbol{y}_{t-1} = \\ &\max_{x_t^i,\ i\in\mathcal{J}} \boldsymbol{y}'_t \left\{ \left[\sum_{i\in\mathcal{J}} \sum_{z=0}^{T_i-1} (\boldsymbol{A}')^z[(\boldsymbol{I}+\boldsymbol{B}\boldsymbol{G})']^z\delta^z\boldsymbol{R}_{\upsilon^z(i)} \right] - \boldsymbol{C}_j + \delta\left[\boldsymbol{A}'\boldsymbol{S}_j\boldsymbol{A}\right] \right\} \boldsymbol{y}_t \\ &\equiv \max_{x_t^{\eta(i)},\ i\in\mathcal{J}} \boldsymbol{y}'_t\boldsymbol{Q}_j\boldsymbol{y}_t. \end{aligned}$$

■■■

Deriving equilibrium production rules

- Value iteration: solve for $\boldsymbol{S}$ and $\boldsymbol{G}$ by iterating over Bellman equation.
- For each $\boldsymbol{S}$, derive corresponding $\boldsymbol{G}$ via FOC of right-hand side:

$$\boldsymbol{B}'_j\left(\boldsymbol{Q}_j+\boldsymbol{Q}'_j\right)\boldsymbol{y}_t = \boldsymbol{B}'_j\left(\boldsymbol{Q}_j+\boldsymbol{Q}'_j\right)\boldsymbol{A}\boldsymbol{y}_{t-1} + \boldsymbol{B}'_j\left(\boldsymbol{Q}_j+\boldsymbol{Q}'_j\right)\boldsymbol{B}\boldsymbol{d}_t = 0.$$

  This system of FOC's corresponds to Eq. (3) in dynamic games handout.

  Rearranging, we get:

$$\boldsymbol{d}_{t+h} = -\left(\boldsymbol{W}\boldsymbol{B}\right)^{-1}\left(\boldsymbol{W}\boldsymbol{A}\right)\boldsymbol{y}_{t-1},$$

  where $\boldsymbol{W}_j \equiv \boldsymbol{B}'_j\left(\boldsymbol{Q}_j+\boldsymbol{Q}'_j\right)$ for each firm j and $\boldsymbol{W} \equiv [\boldsymbol{W}_1,\ldots,\boldsymbol{W}_N]'$.

  Basis for estimating supply side of model.

■■■

Estimation

θ is not identified, set to constant.

To generate estimating equations, introduce shocks to firms' marginal costs:

$$C(x_t^{\eta(i)}) = x_t^{\eta(i)} \left(c_i + \epsilon_{it}\right).$$

Assumptions: $\boldsymbol{\epsilon}_t \equiv [\epsilon_{1t}, \ldots, \epsilon_{Lt}]'$ has zero-mean, *i.i.d.* across t. The vector $\boldsymbol{\epsilon}_t$ is known to all firms when they make their period t choices (no asymmetric information).

From "certainty equivalence" properties of linear-quadratic model, optimal firm strategies are the same as in deterministic model, but with an additive shock:

$$\boldsymbol{d}_t = \boldsymbol{G}\boldsymbol{A}\boldsymbol{y}_{t-1} + \boldsymbol{w}_t, \text{where } E(\boldsymbol{w}_t) = \boldsymbol{0}.$$

where $E(\boldsymbol{w}_t) = \boldsymbol{0}$ and independent of $\boldsymbol{y}_{t-1}$.

With the cost shocks, then, production (and also prices) will be random over time. However, due to independence of shocks across time, the innovations in prices will have mean zero at time t:

$$p_{t+1} = E_t p_{t+1} + \omega_{t+1}$$

where $E_t \omega_{t+1} = 0$.

This implies that the realized "demand function residual" will also have mean zero, conditional on formation known at time t:

$$\begin{aligned} 0 &= E\left[p_t^{\eta(i)} - \left(\alpha_{\eta(i)} - \alpha_{\eta(i)+1}\right)\left(1 - \frac{1}{M}\sum_{r=1}^{\eta(i)} x_t^r\right) - \delta p_{t+1}^{\upsilon(\eta(i))} - p_t^{\eta(i)+1} + \delta p_{t+1}^{\upsilon(\eta(i)+1)} \middle| \Omega_t \right] \\ &= E\left[\left(1 - \delta L^{-1}\right)\left(p_t^{\eta(i)} - p_t^{\eta(i)+1}\right) - \left(\alpha_{\eta(i)} - \alpha_{\eta(i)+1}\right)\left(1 - \frac{1}{M}\sum_{r=1}^{\eta(i)} x_t^r\right) \middle| \Omega_t \right] \end{aligned}$$

These form the basis for the moment conditions which we use for estimation, which we denote $\boldsymbol{\mu}_T(\psi)$.

■■■

GMM Estimation

- Source of identification: co-movements time series of prices and production.

-

- GMM estimator $\boldsymbol{\psi} \equiv \text{argmin}_{\boldsymbol{\psi}} \boldsymbol{\mu}_T(\boldsymbol{\psi})' \boldsymbol{\Omega}_T^{-1} \boldsymbol{\mu}_T(\boldsymbol{\psi})$.

- Nested GMM procedure: for each value of parameters $\boldsymbol{\psi}$, solve LQ dynamic programming problem for coefficients $\boldsymbol{G}(\boldsymbol{\psi})$ of optimal production rules. LQ dynamic programming problem is very easy and quick to solve.

- Results

References

CHEVALIER, J., AND A. GOOLSBEE (2005): "Are Durable Goods Consumers Forward Looking?," NBER working paper 11421.

ESTEBAN, S., AND M. SHUM (2004): "Durable Goods Oligopoly with Secondary Markets: the Case of Automobiles," Forthcoming, *RAND Journal of Economics*.

HOTZ, J., AND R. MILLER (1993): "Conditional Choice Probabilties and the Estimation of Dynamic Models," *Review of Economic Studies*, 60, 497–529.

HOTZ, J., R. MILLER, S. SANDERS, AND J. SMITH (1994): "A Simulation Estimator for Dynamic Models of Discrete Choice," *Review of Economic Studies*, 61, 265–289.

IMAI, S., N. JAIN, AND A. CHING (2005): "Bayesian Estimation of Dynamic Discrete Choice Models," Queen's University, mimeo.

KEANE, M., AND K. WOLPIN (1994): "The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence," *Review of Economics and Statistics*, 76, 648–672.

MAGNAC, T., AND D. THESMAR (2002): "Identifying Dynamic Discrete Decision Processes," *Econometrica*, 70, 801–816.

PAKES, A., AND P. MCGUIRE (2001): "Stochastic Algorithms, Symmetric Markov Perfect Equilibrium, and the 'Curse' of Dimensionality," *Econometrica*, 69, 1261–1282.

RUST, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55, 999–1033.