Single-agent dynamic optimization models

In these lecture notes we consider specification and estimation of dynamic optimization models. 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.

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:

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

where

- 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"
- 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

$$x_{t+1} \begin{cases} \sim G(x'|x_t) & \text{if } i_t = 0 \text{ (don't replace engine in period } t) \\ \sim G(x'|0) & \text{if } i_t = 1\text{: once replaced, mileage gets reset to zero} \end{cases} \tag{2}$$

  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.[1]
- 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.

Define value function:

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

where maximum is over all possible sequences of $\{i_{t+1}, i_{t+2}, \ldots\}$. 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.[2]

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

[2] 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.

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

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

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

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

We make the following parametric assumptions on utility flow:

$$u(x_t, \epsilon_t, i; \theta) = -c\left((1 - i_t) * x_t; \theta\right) - i * RC + \epsilon_{it}$$

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_{it}, $i = 0, 1$ are structural errors, which represents factors which affect HZ's replacement choice i_t in period t, but are unobserved by the econometrician. Define $\epsilon_t \equiv (\epsilon_{0t}, \epsilon_{1t})$.

  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 observe as much as HZ does, and $i_t = i^*(x_t; \theta)$, so that replacement decision should be perfectly explained by mileage. Hence, model will not be able to explain situations where there are two periods with identical mileage, but in one period HZ replaced, and in the other HZ doesn't replace.

(Tension between this empirical practice and "falsifiability: of model)

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 $G(x'|x)$.

Remark: Distinguishing myopic from forward-looking behavior. 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

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

or a forward-looking model, which posits that

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

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 and $F(\epsilon)$, the distribution of the ϵ shocks. (Below, we show how knowledge of β and F, along with an additional normalization, permits nonparametric identification of the utility functions in this model.)

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. Examples: Chevalier and Goolsbee (2009) examine whether students' choices of purchasing a textbook now depend on the possibility that a new edition will be released soon. Becker, Grossman, and Murphy (1994) argue that cigarette addiction is "rational" by showing that cigarette consumption is response to permanent future changes in cigarette prices.

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 j 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:

Assumption 1 $(x_t, \vec{\epsilon}_t)$ *is a stationary controlled first-order Markov process, with transition*

$$\begin{aligned} p(x', \epsilon'|x, \epsilon, i) &= p(\epsilon'|x', x, \epsilon, i) \cdot p(x'|x, e, i) \\ &= p(\epsilon'|x') \cdot p(x'|x, i). \end{aligned} \tag{4}$$

The first line is just factoring the joint density into a conditional times a marginal. The second line shows the simplifications from Rust's assumptions. 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:

$$\begin{aligned} & l\left(x_1, \ldots, x_T, i_1, \ldots, i_T | x_0, i_0; \theta\right) \\ = & \prod_{t=1}^{T} Prob\left(i_t, x_t | x_0, i_0, \ldots, x_{t-1}, i_{t-1}; \theta\right) \\ = & \prod_{t=1}^{T} Prob\left(i_t, x_t | x_{t-1}, i_{t-1}; \theta\right) \\ = & \prod_{t=1}^{T} Prob\left(i_t | x_t; \theta\right) \times Prob\left(x_t | x_{t-1}, i_{t-1}; \theta_3\right). \end{aligned} \tag{5}$$

Both the third and fourth lines arise from the conditional independence assumption. Note that, in the dynamic optimization problem, the optimal choice of i_t depends on the state variables (x_t, ϵ_t). Hence the third line (implying that $\{x_t, i_t\}$ evolves as 1-order Markov) relies on the conditional serial independence of ϵ_t. The last equality also arises from this conditional serial independence assumption.

Hence, the log likelihood is additively separable in the two components:

$$\log l = \sum_{t=1}^{T} \log Prob\left(i_t | x_t; \theta\right) + \sum_{t=1}^{T} \log Prob\left(x_t | x_{t-1}, i_{t-1}; \theta_3\right).$$

Here $\theta_3 \subset \theta$ denotes the subset of parameters which enter G, the transition probability function for mileage. Because $\theta_3 \subset \theta$, we can maximize the likelihood function above in two steps.

First step: Estimate θ_3, the parameters of the Markov transition probabilities for mileage. We assume a discrete distribution for mileage x, taking K distinct and equally-spaced values $\left\{x_{[1]}, x_{[2]}, \ldots, x_{[K]}\right\}$, in increasing order, where $x_{[k']} - x_{[k]} = \Delta \cdot (k' - k)$, where Δ is a mileage increment (Rust considers $\Delta = 5000$). Also assume that given the current state $x_t = x_{[k]}$, the mileage in the next period can more up to at most $x_{[k+J]}$. (When $i_t = 1$¡ so that engine is replaced, we reset $x_t = 0 = x_{[0]}$.) Then the mileage transition probabilities can be expressed as:

$$P(x_{[k+j]} | x_{[k]}, d = 0) = \begin{cases} p_j & \text{if } 0 \leq j \leq J \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

so that $\theta_3 \equiv \{p_0, p_1, \ldots, p_J\}$, with $0 < p_0, \ldots, p_J < 1$ and $\sum_{j=1}^{J} p_j = 1$.

This first step can be executed separately from the substantial second step. θ_3 estimated just by empirical frequencies: $\hat{p}_j = \text{freq}\{x_{t+1} - x_t = \Delta \cdot j\}$, for all $0 \leq j \leq J$.

Second step: Estimate the remaining parameters $\theta \backslash \theta_3$, parameters of maintenance cost function $c(\cdots)$ and engine replacement costs.

Here, we make a further assumption:

Assumption 2 *The ϵ's are distributed i.i.d. (across choices and periods), according to the Type I extreme value distribution. So this implies that in Eq. (4) above, $p(\epsilon'|x') = p(\epsilon')$, for all x'.*

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

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

Because of the logit assumptions on ϵ_t, the replacement probability simplifies to a multinomial logit-like expression:

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

This is called a "dynamic logit" model, in the literature.

Defining $\bar{u}(x,i;\theta) \equiv u(x,\epsilon,i;\theta) - \epsilon_i$ the choice probability takes the form

$$
Prob\left(i_t|x_t;\theta\right) = \frac{\exp\left(\bar{u}(x_t,i_t,\theta) + \beta E_{x',\epsilon'|x_t,i_t} V(x',\epsilon')\right)}{\sum_{i=0,1} \exp\left(\bar{u}(x_t,i,\theta) + \beta E_{x',\epsilon'|x_t,i} V(x',\epsilon')\right)}. \tag{7}
$$

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 $V(x,\epsilon;\hat{\theta})$. After $V(x,\epsilon;\hat{\theta})$ is obtained, we can compute the LL fxn in Eq. (7).

1.2.2 Computational details for inner loop

Compute value function $V(x, \epsilon; \hat{\theta})$ 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):

$$EV(x, i) = \int_y \log \left\{ \sum_{j \in C(y)} \exp \left[\bar{u}(y, j; \theta) + \beta EV(y, j) \right] \right\} p(dy|x, i) \tag{8}$$

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):

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

The next-to-last equality uses the closed-form expression for the expectation of the maximum, for extreme-value variates.[3]

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

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

[3] See Chiong, Galichon, and Shum (2013) for the most general treatment of this.

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.) Here Rust assumes that mileage is discrete- (finite-) valued, and takes K values, each spaced 5000 miles apart, consistently with earlier modeling of mileage transition function in Eq. (6). Let the values of the state variable x be discretized into a grid of points, which we denote $\vec{r}$.

Because of this assumption that x is discrete, the $EV(x,i)$ function is now finite dimensional, having $2 \times K$ elements.

- $\tau = 0$: 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. (8) and $EV^0(x;\theta)$ to calculate, at each $x \in \vec{r}$, and $i \in \{0,1\}$.

$$EV^1(x,i) = \sum_{y \in \vec{r}} \log \left\{ \sum_{j \in C(y)} \exp\left[\bar{u}(y,j;\theta) + \beta EV^0(y,j)\right] \right\} p(y|x,i)$$

  where the transition probabilities $p(y|x,i)$ are given by Eq. (6) above.

  Now check: is $EV^1(x,i)$ close to $EV^0(x,i)$? 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 go to next iteration $\tau = 2$.

References

Becker, G., M. Grossman, and K. Murphy (1994): "An Empirical Analysis of Cigarette Addiction," *American Economic Review*, 84, 396–418.

Chevalier, J., and A. Goolsbee (2009): "Are Durable Goods Consumers Forward Looking? Evidence from the College Textbook Market," *Quarterly Journal of Economics*, 124, 1853–1884.

Chiong, K., A. Galichon, and M. Shum (2013): "Estimating Dynamic Discrete Choice Models via Convex Analysis," mimeo, Caltech.

Magnac, T., and D. Thesmar (2002): "Identifying Dynamic Discrete Decision Processes," *Econometrica*, 70, 801–816.

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