

Dynamic Models: single agent problems*

Spring 2009

1 Rust (1987): “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher”

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

- 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:
 1. there are large fixed costs associated with “stopping” (replacing the engine), but new engine has lower associated future maintenance costs;
 2. by not replacing the engine, you avoid the fixed replacement costs, but suffer higher future maintenance costs.

*These notes rely on Matthew Shum’s lecture notes.

- Optimal solution is characterized by a threshold-type of rule: there is a “critical” cutoff mileage level \hat{x} below which no replacement takes place, but above which replacement will take place.
- Remark: 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 $\{i_1, i_2, i_3, \dots, i_t, i_{t+1}, \dots\}$ to maximize discounted expected utility stream:

$$\max_{\{i_1, i_2, \dots\}} E \sum_{t=1}^{\infty} \beta^{t-1} u(x_t, \epsilon_t, i_t; \theta) \quad (1)$$

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

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

where $G(x'|x)$ 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.

- Note: G only depends on last period's x as opposed to history of x (which is much more complicated).
 - ϵ 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)
 - 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 mileage $\{x_1, x_2, x_3, \dots, x_t, x_{t+1}, \dots\}$ is still random. The expectation in equation 1 is over this stochastic sequence of mileage and over the shocks.
- 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:

$$V(x_t, \epsilon_t) = \max_{\{i_\tau, i_{\tau+1}, \dots\}} E_t \left[\sum_{\tau=t+1}^{\infty} \beta^{\tau-t} u(x_\tau, \epsilon_\tau, i_\tau; \theta) | x_t \right] \quad (3)$$

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 .

- Remark: 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 and macro are finite-horizon.
- Using the Bellman equation, we can break down the DO problem into an (infinite) sequence of single-period decisions

$$i_t = i^*(x_t, \epsilon_t; \theta) = \operatorname{argmax}_i \{u(x_t, \epsilon_t; \theta) + \beta E_{x', \epsilon'} V(x', \epsilon')\} \quad (4)$$

where the value function is

$$V(x, \epsilon) = \max_{i \in \{0,1\}} \{u(x, \epsilon, i; \theta) + \beta E_{x', \epsilon'} V(x', \epsilon')\} \quad (5)$$

$$= \max \{u(x, \epsilon, 0; \theta) + \beta E_{x', \epsilon'} V(x', \epsilon'), u(x, \epsilon, 1; \theta) + \beta E_{\epsilon'} V(0, \epsilon)\} \quad (6)$$

$$= \max \left\{ \tilde{V}(x_t, \epsilon_t, 0; \theta), \tilde{V}(x_t, \epsilon_t, 1; \theta) \right\} \quad (7)$$

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

$$\tilde{V}(x_t, \epsilon_t, i_t) = \begin{cases} u(x_t, \epsilon_t, 1; \theta) + \beta E_{\epsilon'} V(0, \epsilon') & \text{if } i_t = 1 \\ u(x_t, \epsilon_t, 0; \theta) + \beta E_{x', \epsilon'} V(x', \epsilon') & \text{if } i_t = 0 \end{cases} \quad (8)$$

- We make the following parametric assumptions on utility

$$u(x_t, \epsilon_t, i; \theta) = -c((1 - i_t) \cdot x_t; \theta) - i * RC + \epsilon_t$$

where

1. $c(\cdot)$ is the maintenance cost function, increasing in x (higher mileage implies higher costs).
2. RC denotes the “lumpy” fixed costs of adjustment (completely re-hauling the engine). RC ensures that HZ will not replace the engine every period.

3. $\epsilon_{it}, i = \{0, 1\}$ are structural errors. These are factors that affect HZ's replacement choice i_t but are unobserved by the econometrician – let $\epsilon_t \equiv (\epsilon_{0t}, \epsilon_{1t})$
- Rust remarks (bottom, pg. 1008), you need ϵ in order to generate a positive likelihood for your observed data. Without them, we observe as much as HZ does, and $i_t = i^*(x_t; \theta)$. So the 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 (see figure 1).
 - Given additivity, we also define the choice-specific utility

$$u(x, i, \epsilon; \theta) = \mu(x, i; \theta) + \epsilon_i$$

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 the maintenance cost function $c(\cdot)$
 2. replacement cost RC
 3. parameters of the mileage transition function $G(x'|x)$.
- Remark: 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

$$i_t = \operatorname{argmax}_{i \in \{0, 1\}} \{u(x_t, \epsilon_t, 0), u(x_t, \epsilon_t, 1)\}, \quad (9)$$

or a forward-looking model which posits that

$$i_t = \operatorname{argmax}_{i \in \{0, 1\}} \left\{ \tilde{V}(x_t, \epsilon_t, 0), \tilde{V}(x_t, \epsilon_t, 1) \right\}. \quad (10)$$

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 non-parametrically underidentified, without knowledge of β and $F(\epsilon)$, 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 (i.e. β close to zero) and replacement costs RC were low, his decisions may look similar as when he were forward-looking (i.e. β close to 1) and RC were large.
- Reduced-form tests for forward-looking behavior exploit scenarios in which some variables which affected 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 Data

- Observe $\{i_t, x_t\}, t = 1, \dots, T$ for 104 buses.
- Treat buses as homogenous and independent (i.e. replacement of bus engine i is not affected by the replacement decision on bus j).
- Buses are actually quite heterogenous – see Table IIa and IIb (and figure 1)
- While the model can control for variation across groups of buses (by estimating different parameters), can't control for variation within group. So selects groups where buses look homogenous (started with 162, ended up with 104)
- Also worry about replacement vs utilization – older buses may be kept as backups for busy times. RC for older buses is smaller than for newer buses, but in the data see newer bus's engines replaced more frequently. Model doesn't address this “bigger” decision.

1.3 Econometric Model

- Assume the following conditional independence assumption, on the Markovian transition probabilities:

$$p(x', \epsilon' | x, \epsilon, i) = p(\epsilon' | x', x, \epsilon, i) \cdot p(x' | x, \epsilon, i) \quad (11)$$

$$= p(\epsilon' | x') \cdot p(x' | x, i) \quad (12)$$

- The second line shows the simplifications due to Rust's assumptions: There are 2 types of conditional independence
 1. given x , ϵ 's are independent over time, and
 2. conditional on (x, i) , x' is independent of ϵ .
- Likelihood function for a single bus is:

$$L(x_1, \dots, x_T, i_1, \dots, i_T; x_0, i_0, \theta) = \prod_{t=1}^T \text{Prob}(i_t, x_t | (x_0, i_0), \dots, (x_{t-1}, i_{t-1}); \theta) \quad (13)$$

$$= \prod_{t=1}^T \text{Prob}(i_t, x_t | (x_{t-1}, i_{t-1}); \theta) \quad (14)$$

$$= \prod_{t=1}^T \text{Prob}(i_t | x_t; \theta_1) \text{Prob}(x_t | (x_{t-1}, i_{t-1}); \theta_2) \quad (15)$$

where the Markovian feature of the problem is exposed along with the conditional independence assumptions.

- The log likelihood is additively separable in the two components

$$\mathcal{L} = \sum_{t=1}^T \log \text{Prob}(i_t | x_t; \theta_1) + \sum_{t=1}^T \log \text{Prob}(x_t | (x_{t-1}, i_{t-1}); \theta_2). \quad (16)$$

- This can estimate the likelihood in two steps
 1. Estimate θ_2 , the parameters of the transition probabilities for mileage conditional on non-replacement.

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

$$\Delta x_t = \begin{cases} [0, 5000) & \text{w/prob } p, \\ [5, 000, 10, 000) & \text{w/prob } q, \\ [10, 000, \infty) & \text{w/prob } 1 - p - q \end{cases} \quad (17)$$

so that $\theta_2 = \{p, q\}$ with $0 < p, q < 1$ and $p + q < 1$.

This first step can be executed separately from the (substantial) second step

2. Estimate θ_1 – parameters of the maintenance cost function and engine replacement costs

- Make a further assumption on the distribution of ϵ 's: there are distributed i.i.d. (across choices and periods) according to the Type 1 extreme value distribution. So, $p(\epsilon'|x') = p(\epsilon')$ for all x' .
- So $\text{Prob}(i_t = 1|x_t; \theta_1)$ equals

$$\text{Prob}(-c(0; \theta_1) - RC + \epsilon_{1t} + \beta V(0) > -c(x_t; \theta_1) + \epsilon_{0t} + \beta E_{x', \epsilon'} V(x', \epsilon')), \quad (18)$$

$$= \text{Prob}(\epsilon_{1t} - \epsilon_{0t} > c(0; \theta_1) - c(x_t; \theta_1) + \beta[E_{x', \epsilon'} V(x', \epsilon') - V(0) + RC]) \quad (19)$$

because of the logit assumptions, the replacement probability simplifies to a (usual) MNL expression:

$$\frac{\exp(-c(0; \theta_1) - RC + \beta V(0))}{\exp(-c(0; \theta_1) - RC + \beta V(0)) + \exp(-c(x_t; \theta_1) + \beta E_{x', \epsilon'} V(x', \epsilon'))}. \quad (20)$$

- This is sometimes referred to as a “dynamic logit.” More generally we have

$$\text{Prob}(i_t|x_t; \theta) = \frac{\exp(u(x_t, i_t, \theta) + \beta E_{x', \epsilon'} V(x', \epsilon'))}{\sum_{i=0,1} \exp(u(x_t, i_t, \theta) + \beta E_{x', \epsilon'} V(x', \epsilon'))} \quad (21)$$

1.4 Estimation method for the 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 $\hat{\theta}_1$

Inner loop: For $\hat{\theta}_1$, we need to compute the value function $V(x, \epsilon; \hat{\theta}_1)$. After V is obtained, we can compute the log likelihood function.

1.5 Computational details for inner loop

- Compute the value function $V(x, \epsilon; \hat{\theta})$ by iterating over Bellman’s equation, see equation 5.
- A clever and computationally convenient feature in Rust’s paper is that he iterates over the *expected* value function $EV(x; i) \equiv E_{x', \epsilon'} V(x', \epsilon'; \theta)$. The reason for this is that you avoid having to calculate the value function at values ϵ_0 and ϵ_1 , which are additional state variables.
- He iterates over the following equation (which is Eq. 4.14 in his paper)

$$EV(x, i) = \int_y \log \left\{ \sum_{j \in C(y)} \exp[u(y, j; \theta) + \beta EV(y, j)] \right\} p(dy|x, i) \quad (22)$$

- lets explain the notation: “EV” denotes a function. Also (x, i) denotes the previous period’s mileage and replacement choice, while (y, j) denote the current period’s mileage and choice So the bellman equation is

$$V(y, \epsilon; \theta) = \max_{j \in \{0,1\}} [u(y, j; \theta) + \epsilon + \beta EV(y, j)] \quad (23)$$

where

$$E_{y,\epsilon}[V(y, \epsilon; \theta)|x, i] \equiv EV(x, i; \theta) = E_{y,\epsilon|x,i} \left\{ \max_{j \in \{0,1\}} [u(y, j; \theta) + \epsilon + \beta EV(y, j)] \right\} \quad (24)$$

$$= E_{y|x,i} E_{\epsilon|y,x,i} \left\{ \max_{j \in \{0,1\}} [u(y, j; \theta) + \epsilon + \beta EV(y, j)] \right\} \quad (25)$$

$$= E_{y|x,i} \log \left\{ \sum_{j=0,1} \exp [u(y, j; \theta) + \beta EV(y, j)] \right\} \quad (26)$$

$$= \int_y \log \left\{ \sum_{j=0,1} \exp [u(y, j; \theta) + \beta EV(y, j)] \right\} p(dy|x, i) \quad (27)$$

The next-to-last inequality uses the closed-form expression for the expectation of the maximum, for extreme-value variates (a.k.a. the log of sums)

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

$$\frac{\exp(u(x_t, i_t; \theta) + \beta EV(x_t, i_t; \theta))}{\sum_{i=0,1} \exp(u(x_t, i_t; \theta) + \beta EV(x_t, i_t; \theta))} \quad (28)$$

The value iteration procedure

- The expected value function $EV(\cdot; \theta)$ 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 (suppressing the dependence on θ).

- Let the values of the state variable x be discretized into a grid of points, which we denote \bar{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 $x \in \bar{r}$, and $i = 0, 1$.
- $\tau = 1$: Use equation 22 and $EV^0(x, i)$ to calculate, at each $x \in \bar{r}$ and $i = 0, 1$

$$EV^1(x, i) = \int_y \log \left\{ \sum_{j \in C(y)} \exp[u(y, j; \theta) + \beta EV^0(y, j)] \right\} p(dy|x, i) \quad (29)$$

$$= p \cdot \int_x^{x+5000} \log \left\{ \sum_{j \in C(y)} \exp[u(y, j; \theta) + \beta EV^0(y, j)] \right\} dy + \quad (30)$$

$$q \cdot \int_{x+5000}^{x+10,000} \log \left\{ \sum_{j \in C(y)} \exp[u(y, j; \theta) + \beta EV^0(y, j)] \right\} dy + \quad (31)$$

$$(1 - p - q) \cdot \int_{x+10,000}^{\infty} \log \left\{ \sum_{j \in C(y)} \exp[u(y, j; \theta) + \beta EV^0(y, j)] \right\} dy \quad (32)$$

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

$$\sup_{x, i} |EV^1(x, i) - EV^0(x, i)| < \eta \quad (33)$$

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

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

1.6 Summary/Review of Rust (1987)

Content question: Can we do a good job explaining the aggregate time series of capital replacement by modeling its micro-foundations?

Methodological question: Here is a way to estimate the primitive (i.e. structural) parameters of a stochastic dynamic programming problem with observed and unobserved state variables.

Results:

- Methodology: Big leap. Lots of papers now use this technique to estimate dynamic models. In combination with other computational innovations (i.e. simulation), able to estimate lots of complicated dynamic models. All fields have been influenced.
- Still face computational difficulties with: distributions are state variables, strategic dynamic games, etc.
- One of the latest things is to bypass solving the model for every parameter vector in the estimation process (e.g. estimating static games with private information).
- Content: Model generates implied demand for engine replacement (or investment replacement). Can answer a lot of “what if” questions. Change the value of RC for example, via investment tax credits. Or changes in scrappage rules, etc. Performing counterfactuals is one of the main advantages of structural work over reduced-form.

Identification

- First best: talk about identification in the usual way. Especially focus on how variation in the data pins down the value of parameters in the model. But this can be difficult to do with complicated dynamic models.
- Often helps to vary a parameter and see how the objective criterion changes (or how the model’s predictions of certain key variables change).
- Second best: Estimate the model and run a battery of tests. Look at model predictions and see how compare with the data (or independent work). Basically, do goodness-of-fit and be able to explain why the model fits certain patterns well and not other patterns.

- There is a lot of discussion between reduced-form and structural economists about the importance of identification and how to show a model is identified.
- For some insight (at least from the structural side) go to Prof John Rust's website:
<http://gemini.econ.umd.edu/jrust/papers.html>
and look for
Comments on Michael Keane's paper, "Structural vs. Atheoretic Approaches to Econometrics" Keane's paper

2 Hotz and Miller (1993) “Conditional Choice Probabilities and the Estimation of Dynamic Models”

- One problem with Rust’s approach to estimating dynamic discrete-choice model is that it is very computer intensive. Requires using numeric dynamic programming to compute the value function(s) for every parameter vector.
- There is an alternative method of estimation, which avoids explicit DP. We’ll 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 bus 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, and so $d_{it} \in \{0, 1\}$. Let $i = 1, \dots, N$ index the buses, $t = 1, \dots, T$ 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, $\hat{\cdot}$, denote objects obtained just from data. Objects with tildes, $\tilde{\cdot}$, denote “predicted” quantities, obtained from both data and calculated from model given parameter values θ .
- From the data alone, we can estimate (or “identify”):
 - Transition probabilities of observed state and control variables: $G(x'|x, d)$,

estimated by conditional empirical distribution

$$\hat{G}(x'|x, d) \equiv \begin{cases} \sum_{i=1}^N \sum_{t=1}^{T-1} \frac{1}{\sum_i \sum_t \mathbf{1}(x_{it}=x, d_{it}=0)} \cdot \mathbf{1}(x_{i,t+1} \leq x', x_{it} = x, d_{it} = 0), & \text{if } d=0 \\ \sum_{i=1}^N \sum_{t=1}^{T-1} \frac{1}{\sum_i \sum_t \mathbf{1}(d_{it}=1)} \cdot \mathbf{1}(x_{i,t+1} \leq x', d_{it} = 1), & \text{if } d=1 \end{cases} \quad (34)$$

Choice Probabilities, conditional on state variable $\text{Prob}(d = 1|x)$, estimated by

$$\hat{P}(d = 1|x) \equiv \sum_{i=1}^N \sum_{t=1}^{T-1} \frac{1}{\sum_i \sum_t \mathbf{1}(x_{it} = x,)} \cdot \mathbf{1}(x_{i,t} = x, d_{it} = 1) \quad (35)$$

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

– Note: there are stationarity assumptions made: the G and $\text{Prob}(d = 1|x)$ functions are explicitly not indexed with time t .

- Let $\tilde{V}(x, d; \theta)$ denote the choice-specific value function, minus the error term ϵ_d . That is,

$$\tilde{V}(x, d; \theta) \equiv \tilde{V}(x, \epsilon, d; \theta) - \epsilon_d \quad (36)$$

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

$$\tilde{V}(x, d = 1; \theta) = u(x, d = 1; \theta) + \beta E_{x'|x, d=1} E_{d'|x'} E_{\epsilon'|d', x'} \left[u(x', d'; \theta) + \epsilon' + \right. \quad (37)$$

$$\left. \beta E_{x''|x', d'} E_{d''|x'', d'} E_{\epsilon''|d'', x''} [u(x'', d''; \theta) + \epsilon'' + \beta \dots] \right] \quad (38)$$

$$\tilde{V}(x, d = 0; \theta) = u(x, d = 0; \theta) + \beta E_{x'|x, d=0} E_{d'|x'} E_{\epsilon'|d', x'} \left[u(x', d'; \theta) + \epsilon' + \right. \quad (39)$$

$$\left. \beta E_{x''|x', d'} E_{d''|x'', d'} E_{\epsilon''|d'', x''} [u(x'', d''; \theta) + \epsilon'' + \beta \dots] \right] \quad (40)$$

- Above $u(x, d; \theta)$ 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.
- In practice, “truncate” the infinite sum at some period T (far enough in the future such that the discount factor reduces its impact).
- Note the expectation of $E_{\epsilon|d,x}$ denotes the expectation of ϵ conditional on (x, d) . For the logit case there is closed form:

$$E[\epsilon|d, x] = \gamma - \log(\Pr(d|x)) \quad (41)$$

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

- Both of the other expectations in the above expressions are observed directly from the data.
- Alternatively, both choice-specific value functions can be **simulated** for $d = 1, 2$ by

$$\tilde{V}(x, d; \theta) \approx \frac{1}{S} \sum_{s=1}^S u(x, d; \theta) + \beta \left[u(x'^s, d^s; \theta) + \gamma - \log(\hat{P}(d^s|x'^s)) + \right. \quad (42)$$

$$\left. \beta [u(x''^s, d''^s; \theta) + \gamma + \hat{P}(d''^s|x''^s) + \beta \dots] \right] \quad (43)$$

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.

- So simulate $\tilde{V}(x, d; \theta)$ by drawing S “sequences” of (d_t, x_t) with an initial value of (d, x) and computing the present-discounted utility corresponding to each sequence. Then the simulation average of $\tilde{V}(x, d; \theta)$ is obtained as the sample average.
- When do you use simulation versus “infinite sum”?
- Given the estimated $\tilde{V}(x, d; \theta)$ you can compute the *predicted choice probabilities*

$$\tilde{p}(d = 1|x; \theta) \equiv \frac{\exp\left(\tilde{V}(x, d = 1; \theta)\right)}{\exp\left(\tilde{V}(x, d = 1; \theta)\right) + \exp\left(\tilde{V}(x, d = 0; \theta)\right)} \quad (44)$$

and analogously $\tilde{p}(d = 0|x; \theta)$.

- Note that $\tilde{p}(d|x; \theta)$ and $\hat{p}(d|x)$ differ because one depends on parameters and one is based solely on the data.
- One way to estimate θ is to minimize the distance between predicted and actual conditional choice probabilities

$$\hat{\theta} = \operatorname{argmin}_{\theta} \|\hat{\mathbf{p}}(d = 1|x) - \tilde{\mathbf{p}}(d = 1|x; \theta)\| \quad (45)$$

where \mathbf{p} denotes a vector of probabilities.

- Another way to estimate θ is very similar to the Berry/BLP method. We can calculate directly from the data.

$$\hat{\delta}_x \equiv \log \hat{p}(d = 1|x) - \log \hat{p}(d = 0|x) \quad (46)$$

Given the logit assumption, from equation 44 we know that

$$\log \tilde{p}(d = 1|x; \theta) - \log \tilde{p}(d = 0|x; \theta) = \left[\tilde{V}(x, d = 1; \theta) - \tilde{V}(x, d = 0; \theta) \right] \quad (47)$$

Hence we obtain an alternative estimator

$$\bar{\theta} = \operatorname{argmin}_{\theta} \|\hat{\delta}_x - \left[\tilde{V}(x, d = 1; \theta) - \tilde{V}(x, d = 0; \theta) \right]\| \quad (48)$$