Single-agent dynamic optimization models: estimation and identification

1 Alternative approaches to estimation: avoid numeric dynamic programming

1.1 Hotz-Miller approach

- 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 bus j, time t, observe:

  - $\{x_{jt}, i_{jt}\}$: observed state variables x_{jt} and discrete decision (control) variable i_{jt}.

    Let $j = 1, \ldots, N$ index the buses, $t = 1, \ldots, T$ index the time periods.

  - For Harold Zurcher model: x_{jt} is mileage since last replacement on bus i in period t, and i_{jt} is whether or not engine of bus j was replaced in period t.

  - Unobserved state variables: ϵ_{jt}, *i.i.d.* over j and t. Assume that distribution is known (Type 1 Extreme Value in Rust model)

1.2 Hats and Tildes

In the following, let quantities with hats ˆ's denote objects obtained just from data.

Objects with tildes ˜'s denote "predicted" quantities, obtained from both data and calculated from model given parameter values θ.

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

- Choice probabilities, conditional on state variable: Prob $(i = 1|x)$[1], estimated by

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

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

- Transition probabilities of observed state and control variables: $G(x'|x,i)$[2], estimated by conditional empirical distribution

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

- In practice, when x is continuous, we estimate smoothed version of these functions by introducing a "smoothing weight" $w_{jt} = w(x_{jt}; x)$ such that $\sum_j \sum_t w_{jt} = 1$. Then, for instance, the choice probability is approximated by

$$\hat{p}(i=1|x) = \sum_j \sum_t w_{jt} \mathbf{1}(i_{jt}=1).$$

  One possibility for the weights is a kernel-weighting function. Consider a kernel function $k(\cdot)$ which is symmetric around 0 and integrates to 1. Then

$$w_{jt} = \frac{k\left(\frac{x_{jt}-x}{h}\right)}{\sum_{j'} \sum_{t'} k\left(\frac{x_{j't'}-x}{h}\right)}.$$

  h is a bandwidth. Note that as $h \to 0$, then $w_{it} \to \frac{\mathbf{1}(x_{jt}=x)}{\sum_{j'} \sum_{t'} \mathbf{1}(x_{j't'}=x)}$.

Tildes and forward simulation. Let $\tilde{V}(x,i;\theta)$ denote the choice-specific value function, minus the error term ϵ_i.

With estimates of $\hat{G}(\cdot|\cdot)$ and $\hat{p}(\cdot|\cdot)$, as well as a parameter vector θ, you can "estimate" these choice-specific value functions by exploiting an alternative representation of value function: letting i^* denote the optimal sequence of decisions, we have:

$$V(x_t, \epsilon_t) = \mathbb{E}\left[\sum_{\tau=0}^{\infty} \beta_{\tau t} \left\{u(x_\tau, i_t^*) + \epsilon_{i_t^*}\right\} .|x_t, \epsilon_t\right]$$

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

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

This implies that the choice-specific value functions can be obtained by constructing the sum[3]

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

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

Here $u(x,i;\theta)$ denotes the per-period utility of taking choice i at state x, *without* the additive logit error. Note that the knowledge of $i'|x'$ is crucial to being able to forward-simulate the choice-specific value functions. Otherwise, $i'|x'$ is multinomial with probabilities given by Eq. (1) below, and is impossible to calculate without knowledge of the choice-specific value functions.

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

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

Also, the expectation $\mathbb{E}_{\epsilon|i,x}$ denotes the expectation of the ϵ_i conditional on choice i being taken, and current mileage x. For the logit case, there is a closed form:

$$\mathbb{E}[\epsilon_i|i,x] = \gamma - \log(Pr(i|x))$$

where γ is Euler's constant (0.577...) and $Pr(i|x)$ is the choice probability of action i 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 $i=1,2$):

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

[3]Note that the distribution $(x',i',\epsilon'|x,i)$ can be factored, via the conditional independence assumption, into $(\epsilon'|i',x')(i'|x')(x'|x,i)$.

where

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

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

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

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

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

1.3 Estimation: match hats to tildes

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

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

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. We can calculate directly from the data.

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

Given the logit assumption, from equation (1), we know that

$$\log \tilde{p}(i = 1|x) - \log \tilde{p}(i = 0|x) = \left[\tilde{V}(x, i = 1) - \tilde{V}(x, i = 0)\right].$$

Hence, by equating $\tilde{V}(x, i = 1) - \tilde{V}(x, i = 0)$ to $\hat{\delta}_x$, we obtain an alternative estimator for θ:

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

1.4 A further shortcut in the discrete state case

In this section, for convenience, we will use Y instead of i to denote the action.

For the case when the state variables X are discrete, it turns out that, given knowledge of the CCP's $P(Y|X)$, solving for the value function is just equivalent to solving a system of linear equations. This was pointed out in Pesendorfer and Schmidt-Dengler (2008) and Aguirregabiria and Mira (2007). Specifically:

- Assume that choices Y and state variables X are all *discrete* (ie. finite-valued). $|X|$ is cardinality of state space X. Here X includes just the observed state variables (not including the unobserved shocks ϵ)

- Per-period utilities:

  $$u(Y, X, \epsilon_Y; \Theta) = \bar{u}(Y, X; \Theta) + \epsilon_Y$$

  where ϵ_Y, for $y = 1, \dots Y$, are i.i.d. extreme value distributed with unit variance.

- Parameters Θ. The discount rate β is treated as known and fixed.

- Introduce some more specific notation. Define the integrated or ex-ante value function (before ϵ observed, and hence before the action Y is chosen):[4]

  $$W(X) = \mathbb{E}[V(X, \epsilon)|X].$$

  Along the optimal dynamic path, at state X and optimal action Y, the continuation utility is

  $$\bar{u}(Y, X) + \epsilon_Y + \beta \sum_{X'} P(X'|X, Y) W(X').$$

[4] Similar to Rust's $EV(\cdots)$ function, but *not the same*. See appendix.

This integrated value function satisfies a Bellman equation:

$$
\begin{aligned}
W(X) &= \sum_Y \left[P(Y|X)\left\{\bar{u}(Y,X) + \mathbb{E}(\epsilon_Y|Y,X)\right\}\right] + \beta \sum_Y \sum_{X'} P(Y|X)P(X'|X,Y)W(X') \\
&= \sum_Y \left[P(Y|X)\left\{\bar{u}(Y,X) + \mathbb{E}(\epsilon_Y|Y,X)\right\}\right] + \beta \sum_{X'} P(X'|X)W(X')
\end{aligned}
\tag{2}
$$

- To derive the above, start with "real" Bellman equation:

$$
\begin{aligned}
V(X,\epsilon) &= \bar{u}(Y^*,X) + \sum_Y \epsilon_Y \mathbf{1}(Y=Y^*) + \beta \mathbb{E}_{X'|X,Y}\mathbb{E}_{\epsilon'|X'}V(X',\epsilon') \\
&= \bar{u}(Y^*,X) + \sum_Y \epsilon_Y \mathbf{1}(Y=Y^*) + \beta \mathbb{E}_{X'|X,Y}W(X') \\
\Rightarrow W(X) = \mathbb{E}_{\epsilon|X}V(X,\epsilon) &= \mathbb{E}_{Y^*,\epsilon|X}\left\{\bar{u}(Y^*,X) + \sum_Y \epsilon_Y \mathbf{1}(Y=Y^*) + \beta \mathbb{E}_{X'|X,Y}W(X')\right\} \\
&= \mathbb{E}_{Y^*|X}\mathbb{E}_{\epsilon|Y^*,X}\left\{\cdots\right\} \\
&= \mathbb{E}_{Y^*|X}\left[\bar{u}(Y^*,X) + E[\epsilon_{Y^*}|Y^*,X] + \beta \mathbb{E}_{X'|X,Y}W(X')\right] \\
&= \sum_Y P(Y=Y^*|X)\left[\cdots\right].
\end{aligned}
$$

(Note: in the 4th line above, we first condition on the optimal choice Y^*, and take expectation of ϵ conditional on Y^*. The other way will not work.)

- In matrix notation, this is:

$$
\begin{aligned}
\bar{W}(\Theta) &= \sum_{Y\in(0,1)} P(Y) * \left[\bar{u}(Y;\Theta) + \epsilon(Y)\right] + \beta \cdot F \cdot \bar{W}(\Theta) \\
\Leftrightarrow \bar{W}(\Theta) &= (I-\beta F)^{-1}\left\{\sum_{Y\in(0,1)} P(Y) * \left[\bar{u}(Y;\Theta) + \epsilon(Y)\right]\right\}
\end{aligned}
\tag{3}
$$

where

  - $\bar{W}(\Theta)$ is the vector (each element denotes a different value of X) for the integrated value function at the parameter Θ
  - '*' denotes elementwise multiplication

- F is the $|X|$-dimensional square matrix with (i,j)-element equal to $Pr(X' = j|X = i)$.
- $P(Y)$ is the $|X|$-vector consisting of elements $Pr(Y|X)$.
- $\bar{u}(Y)$ is the $|X|$-vector of per-period utilities $\bar{u}(Y;X)$.
- $\epsilon(Y)$ is an $|X|$-vector where each element is $E[\epsilon_Y|Y,X]$. For the logit assumptions, the closed-form is

$$E[\epsilon_Y|Y,X] = \text{Euler's constant} - \log(P(Y|X)).$$

  Euler's constant is 0.57721.

Based on this representation, P/S-D propose a class of "least-squares" estimators, which are similar to HM-type estimators, except now we don't need to "forward-simulate" the value function. For instance:

- Let $\hat{P}(\bar{Y})$ denote the estimated vector of conditional choice probabilities, and $\hat{F}$ be the estimated transition matrix. Both of these can be estimated directly from the data.

- For each posited parameter value Θ, and given $(\hat{F},\ \hat{P}(\bar{Y}))$ use equation (3) to evaluate the integrated value function $\bar{W}(X,\Theta)$, and derive the vector $\tilde{P}(\bar{Y};\Theta)$ of implied choice probabilities at Θ, which has elements

$$\tilde{P}(Y|X;\Theta) = \frac{\exp\left[\bar{u}(Y,X;\Theta) + \mathbb{E}_{X'|X,Y}W(X';\Theta)\right]}{\sum_Y \exp\left[\bar{u}(Y,X;\Theta) + \mathbb{E}_{X'|X,Y}W(X';\Theta)\right]}.$$

- Hence, Θ can be estimated as the parameter value minimizing the norm $||\hat{P}(\bar{Y}) - \tilde{P}(Y;\Theta)||$.

2 Semiparametric identification of DDC Models

We can also use the Hotz-Miller estimation scheme as a basis for an argument regarding the identification of the underlying DDC model. In Markovian DDC models, without unobserved state variables, the Hotz-Miller routine exploits the fact that the

Markov probabilities $x', d'|x, d$ are identiified directly from the data, which can be factorized into

$$f(x', d'|x, d) = \underbrace{f(d'|x')}_{\text{CCP}} \cdot \underbrace{f(x'|x, d)}_{\text{state law of motion}}. \tag{4}$$

In this section, we argue that once these "reduced form" components of the model are identified, the remaining parts of the models – particularly, the per-period utility functions – can be identified without any further parametric assumptions. These arguments are drawn from Magnac and Thesmar (2002) and Bajari, Chernozhukov, Hong, and Nekipelov (2007).

We make the following assumptions, which are standard in this literature:

1. Agents are optimizing in an infinite-horizon, stationary setting. Therefore, in the rest of this section, we use primes $'$'s to denote next-period values.

2. Actions D are chosen from the set $\mathcal{D} = \{0, 1, \ldots, K\}$.

3. The state variables are X.

4. The per-period utility from taking action $d \in \mathcal{D}$ in period t is:

   $$u_d(X_t) + \epsilon_{d,t}, \ \forall d \in \mathcal{D}.$$

   The $\epsilon_{d,t}$'s are utility shocks which are independent of X_t, and distributed i.i.d with known distribution $F(\epsilon)$ across periods t and actions d. Let $\vec{\epsilon}_t \equiv (\epsilon_{0,1}, \epsilon_{1,t}, \ldots, \epsilon_{K,t})$.

5. From the data, the "conditional choice probabilities" CCP's

   $$p_d(X) \equiv \text{Prob}(D = d|X),$$

   and the Markov transition kernel for X, denoted $p(X'|D, X)$, are identified.

6. $u_0(X)$, the per-period utility from $D = 0$, is normalized to zero, for all X.

7. β, the discount factor, is known.[5]

[5]Magnac and Thesmar (2002) discuss the possibility of identifying β via exclusion restrictions, but we do not pursue that here.

Following the arguments in Magnac and Thesmar (2002) and Bajari, Chernozhukov, Hong, and Nekipelov (2007), we will show the nonparametric identification of $u_d(\cdot),\ d = 1, \ldots, K$, the per-period utility functions for all actions except $D = 0$.

The Bellman equation for this dynamic optimization problem is

$$V(X, \vec{\epsilon}) = \max_{d \in \mathcal{D}} \left(u_d(X) + \epsilon_d + \beta \mathbb{E}_{X', \vec{\epsilon}' | D, X} V(X', \vec{\epsilon}') \right)$$

where $V(X, \vec{\epsilon})$ denotes the value function. We define the choice-specific value function as

$$V_d(X) \equiv u_d(X) + \beta \mathbb{E}_{X', \vec{\epsilon}' | D, X} V(X', \vec{\epsilon}').$$

Given these definitions, an agent's optimal choice when the state is X is given by

$$y^*(X) = \operatorname{argmax}_{y \in \mathcal{D}} \left(V_d(X) + \epsilon_d \right).$$

Hotz and Miller (1993) and Magnac and Thesmar (2002) show that in this setting, there is a known one-to-one mapping, $q(X) : \mathbb{R}^K \to \mathbb{R}^K$, which maps the K-vector of choice probabilities $(p_1(X), \ldots, p_K(X))$ to the K-vector $(\Delta_1(X), \ldots, \Delta_K(X))$, where $\Delta_d(X)$ denotes the difference in choice-specific value functions

$$\Delta_d(X) \equiv V_d(X) - V_0(X).$$

Let the i-th element of $q(p_1(X), \ldots, p_K(X))$, denoted $q_i(X)$, be equal to $\Delta_i(X)$. The known mapping q derives just from $F(\epsilon)$, the known distribution of the utility shocks.

Hence, since the choice probabilities can be identified from the data, and the mapping q is known, the value function differences $\Delta_1(X), \ldots, \Delta_K(X)$ is also known.

Next, we note that the choice-specific value function also satisfies a Bellman-like equation:

$$\begin{aligned} V_d(X) &= u_d(X) + \beta \mathbb{E}_{X'|X,D} \left[\mathbb{E}_{\vec{\epsilon}'} \max_{d' \in \mathcal{D}} (V_{d'}(X') + \epsilon'_{d'}) \right] \\ &= u_d(X) + \beta \mathbb{E}_{X'|X,D} \left\{ V_0(X') + \left[\mathbb{E}_{\vec{\epsilon}'} \max_{d' \in \mathcal{D}} (\Delta_{d'}(X') + \epsilon'_{d'}) \right] \right\} \\ &= u_d(X) + \beta \mathbb{E}_{X'|X,D} \left[H(\Delta_1(X'), \ldots, \Delta_K(X')) + V_0(X') \right] \end{aligned} \tag{5}$$

where $H(\cdots)$ denotes McFadden's "social surplus" function, for random utility models (cf. Rust (1994, pp. 3104ff)). Like the q mapping, H is a known function, which depends just on $F(\epsilon)$, the known distribution of the utility shocks.

Using the assumption that $u_0(X) = 0,\ \forall X$, the Bellman equation for $V_0(X)$ is

$$V_0(X) = \beta \mathbb{E}_{X'|X,D} \left[H(\Delta_1(X'), \ldots, \Delta_K(X')) + V_0(X') \right]. \tag{6}$$

In this equation, everything is known (including, importantly, the distribution of $X'|X, D$), except the $V_0(\cdot)$ function. Hence, by iterating over Eq. (6), we can recover the $V_0(X)$ function. Once $V_0(\cdot)$ is known, the other choice-specific value functions can be recovered as

$$V_d(X) = \Delta_d(X) + V_0(X),\ \forall y \in \mathcal{D},\ \forall X.$$

Finally, the per-period utility functions $u_d(X)$ can be recovered from the choice-specific value functions as

$$u_d(X) = V_d(X) - \beta \mathbb{E}_{X'|X,D} \left[H(\Delta_1(X'), \ldots, \Delta_K(X')) + V_0(X') \right],\ \forall y \in \mathcal{D},\ \forall X,$$

where everything on the right-hand side is known.

Remark: For the case where $F(\epsilon)$ is the Type 1 Extreme Value distribution, the social surplus function is

$$H(\Delta_1(X), \ldots, \Delta_K(X)) = \log \left[1 + \sum_{d=1}^{K} \exp(\Delta_d(X)) \right]$$

and the mapping q is such that

$$q_d(X) = \Delta_d(X) = \log(p_d(X)) - \log(p_0(X)),\ \forall d = 1, \ldots K,$$

where $p_0(X) \equiv 1 - \sum_{d=1}^{K} p_d(X)$. $\blacksquare$

Remark: The above argument also holds if ϵ_d is not independent of $\epsilon_{d'}$, and also if the joint distribution of $(\epsilon_0, \epsilon_1, \ldots, \epsilon_K)$ is explicitly dependent on X. However, in that case, the mappings q_X and H_X will depend explicitly on X, and typically not be available in closed form, as in the multinomial logit case. For this reason, practically all applications of this machinery maintain the multinomial logit assumption.

A A result for multinomial logit model

Show: for the multinomial logit case, we have $E[\epsilon_j|\text{choice } j \text{ is chosen}] = \gamma - \log(P_j)$ where γ is Euler's constant (0.577...) and $Pr(d|x)$ is the choice probability of action j.

This closed-form expression has been used much in the literature on estimating dynamic models: eg. Eq. (12) in Aguirregabiria and Mira (2007) or Eq. (2.22) in Hotz, Miller, Sanders, and Smith (1994).

Use the fact: for a univariate extreme value variate with parameter a, CDF $F(\epsilon) = \exp(-ae^{-\epsilon})$, and density $f(\epsilon) = \exp(-ae^{-\epsilon})(ae^{-\epsilon})$, we have

$$E(\epsilon) = \log a + \gamma, \quad \gamma = 0.577..$$

Also use McFadden's (1978) results for generalized extreme value distribution:

- For a function $G(e^{V_0}, \ldots, e^{V_J})$, we define the generalized extreme value distribution of $(\epsilon_0, \ldots, \epsilon_j)$ with joint CDF $F(\epsilon_0, \ldots, \epsilon_J) = \exp\{-G(e^{\epsilon_0}, \ldots, e^{\epsilon_J})\}$.
- $G(\ldots)$ is a homogeneous-degree-1 function, with nonnegative odd partial derivatives and nonpositive even partial derivatives.
- Theorem 1: For a random utility model where agent chooses according to $j = \text{argmax}_{j' \in \{0,1,\ldots,J\}} U_j = V_j + \epsilon_j$, the choice probabilities are given by

$$\begin{aligned} P(j) &= \int_{-\infty} F_j(V_j + \epsilon_j - V_0, V_j + \epsilon_j - V_1, \ldots, V_j + \epsilon_j - V_J) d\epsilon_j \\ &= \frac{e^{V_j} G_j(e^{V_0}, \ldots, e^{V_J})}{G(e^{V_0}, \ldots, e^{V_J})} \end{aligned}$$

- Corollary: Total expected surplus is given by

$$\bar{U} = \mathbb{E} \max_j (V_j + \epsilon_j) = \gamma + \log(G(e^{V_0}, \ldots, e^{V_J}))$$

  and choice probabilities by $P_j = \frac{\partial \bar{U}}{\partial V_j}$. For this reason, $G(\ldots)$ is called the "social surplus function"

In what follows, we use McFadden's shorthand of $< V_{j'} >$ to denote a $J+1$ vector with $j'-th$ component equal to $V_{j'-1}$ for $j' = 1, \ldots, J$.

Imitating the proof for the corollary above, we can derive that (defining $a = G(< e^{V_{j'}} >)$)

$$
\begin{aligned}
&\mathbb{E}(V_j + \epsilon_j | j \text{ is chosen}) \\
=&\frac{1}{P_j} \int_{-\infty}^{+\infty} (V_j + \epsilon_j) F_j(< V_j + \epsilon_j - V_{j'} >) d\epsilon_j \\
=&\frac{1}{P_j} \int_{-\infty}^{+\infty} (V_j + \epsilon_j) \exp(-G(< e^{-V_j - \epsilon_j + V_{j'}} >)) G_j(< e^{-V_j - \epsilon_j + V_{j'}} >) e^{-\epsilon_j} d\epsilon_j \\
=&\frac{a}{e^{V_j} G_j(< e^{V_{j'}} >)} \int_{-\infty}^{+\infty} (V_j + \epsilon_j) \exp(-a e^{-V_j - \epsilon_j}) G_j(< e^{V_{j'}} >) e^{-\epsilon_j} d\epsilon_j \qquad \text{(by props. of } G\text{)} \\
=&\int_{-\infty}^{+\infty} (V_j + \epsilon_j) \exp(-a e^{V_j - \epsilon_j}) a e^{-V_j} e^{-\epsilon_j} d\epsilon_j \\
=&\int_{-\infty}^{+\infty} w \exp(-a e^{-w}) a e^{w} dw \qquad (V_j + \epsilon_j \to w) \\
=&\log(a) + \gamma.
\end{aligned}
$$

For the multinomial logit model, we have $G(< e^{V_{j'}} >) = \sum_{j'} e^{V_{j'}}$. For this case $P_j = \exp(V_j)/G(< e^{V_{j'}} >)$, and $G_j(\cdots) = 1$ for all j. Then

$$
\begin{aligned}
\mathbb{E}[\epsilon_j | j \text{ is chosen}] &= \log(a) + \gamma - (V_j - V_0) - V_0 \\
&= \log(G(< e^{V_{j'}} >)) + \gamma - \log(P_j) + \log(P_0) - V_0 \quad (\text{using } V_j - V_0 = \log(P_j/P_0)) \\
&= \log(G(< e^{V_{j'}} >)) + \gamma - \log(P_j) + V_0 - \log(G(< e^{V_{j'}} >)) - V_0 \\
&= \gamma - \log(P_j)
\end{aligned}
$$

B Relations between different value function notions

Here we delve into the differences between the "real" value function $V(x, \epsilon)$, the $EV(x, y)$ function from Rust (1987), and the integrated or ex-ante value function $W(x)$ from Aguirregabiria and Mira (2007) and Pesendorfer and Schmidt-Dengler (2008).

By definition, Rust's EV function is defined as:

$$EV(x, y) = \mathbb{E}_{x', \epsilon' | x, y} V(x', \epsilon').$$

By definition, the integrated value function is defined as

$$W(x) = \mathbb{E}[V(X, \epsilon)|X = x].$$

By iterated expectations,

$$\begin{aligned} EV(x, y) &= \mathbb{E}_{x',\epsilon'|x,y}V(x', \epsilon') \\ &= \mathbb{E}_{x'|x,y}\mathbb{E}_{\epsilon'|x'}V(x', \epsilon') \\ &= \mathbb{E}_{x'|x,y}W(x') \end{aligned}$$

given the relationship between the EV and integrated value functions.

Hence, the "choice-specific value function" (without the ϵ) is defined as:

$$\begin{aligned} \bar{v}(x, y) &\equiv u(x, y) + \beta\mathbb{E}_{x',\epsilon'|x,y}V(x', \epsilon') \\ &= u(x, y) + \beta EV(x, y) \\ &= u(x, y) + \beta\mathbb{E}_{x'|x,y}W(x'). \end{aligned}$$

Also note

$$\begin{aligned} V(x, \epsilon) &= \max_y \{\bar{v}(x, y) + \epsilon_y\} \\ \Rightarrow W(x) &= \mathbb{E}\left[\max_y \{\bar{v}(x, y) + \epsilon_y\} |x\right] \end{aligned}$$

which corresponds to the *social surplus function* of this dynamic discrete choice model.

References

AGUIRREGABIRIA, V., AND P. MIRA (2007): "Sequential Estimation of Dynamic Discrete Games," *Econometrica*, 75, 1–53.

BAJARI, P., V. CHERNOZHUKOV, H. HONG, AND D. NEKIPELOV (2007): "Nonparametric and Semiparametric Analysis of a Dynamic Game Model," Manuscript, University of Minnesota.

GELFAND, I., AND S. FOMIN (1965): *Calculus of Variations*. Dover.

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.

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

MCFADDEN, D. (1978): "Modelling the Choice of Residential Location," in *Spatial Interaction Theory and Residential Location*, ed. by A. K. et. al. North Holland Pub. Co.

PESENDORFER, M., AND P. SCHMIDT-DENGLER (2008): "Asymptotic Least Squares Estimators for Dynamic Games," *Review of Economic Studies*, 75, 901–928.

RUST, J. (1994): "Structural Estimation of Markov Decision Processes," in *Handbook of Econometrics, Vol. 4*, ed. by R. Engle, and D. McFadden, pp. 3082–146. North Holland.