


Dynamic Discrete Choice¹

Holger Sieg

¹The discussion of the dynamic logit model largely follows Rust (1994) as well as lecture notes that John Rust shared with me. 

Steps in a Dynamic Structural Model

The typical steps in a dynamic structural econometric model:

1. Specify the primitives of the model,
 - ▶ Single period agents' payoff functions (utility or profit)
 - ▶ Evolution of state variables (e.g. capital)
2. Solve for optimal dynamic behavior
 - ▶ We typically assume that agents maximize the present discounted value of future utilities or profits
3. Search for parameter values that result in the “best match” between our model predictions and observed behavior

State and Control Variables

- ▶ Two types of variables
 - ▶ s_t state variables
 - ▶ d_t control variables
- ▶ The set of all possible choices is D . The action space is assumed to have a finite number of elements.
- ▶ Choices are constrained. The constraints depend on the state of the world. The set of feasible choices at s_t is denoted by

$$D_t(s_t) \subseteq D$$

- ▶ Example: budget constraint, law of motion for the capital stock, etc.

Example: Rust's Optimal Bus Replacement

- ▶ Rust (1987) estimated an optimal replacement model.
- ▶ The choices are to replace or not to replace an engine:

$$d_{1,t} = \begin{cases} 1 & \text{no replacement} \\ 0 & \text{replacement} \end{cases}$$

- ▶ State variables:
 - ▶ miles of the engine x_t ;
 - ▶ $\varepsilon_t = (\varepsilon_{1,t}, \varepsilon_{2,t})$, which are idiosyncratic shocks.

Preferences & Beliefs

- ▶ Period utility function is given by:

$$u(s_t, d_t, \theta)$$

- ▶ θ is a vector of parameters to be estimated. Note that we often suppress θ to simplify the notation.
- ▶ Agents maximize expected intertemporal utility with discount factor β .
- ▶ Beliefs over future states are given by Markov transition probabilities:

$$\pi(s_{t+1}|s_t, d_t, \theta)$$

Bus Replacement Example

- Choice specific utility is given by:

$$u_{1,t} = -C(x_t, \theta) + \varepsilon_{1,t}$$

$$u_{2,t} = -C(0, \theta) - rc + \varepsilon_{2,t}$$

where operating costs are given by $C(x_t) = \theta_0 + \theta_1 x_t$, and replacement costs are rc .

- $\mathbb{P}[x_{t+1}|x_t, d_t, \theta] = \begin{cases} \theta_2 e^{-\theta_2(x_{t+1}-x_t)} & d_{1,t} = 1 \\ \theta_2 e^{-\theta_2 x_{t+1}} & d_{2,t} = 1 \end{cases}$
- $(\varepsilon_{1,t}, \varepsilon_{2,t}) \sim \text{Type I extreme value}$

Optimal Decision Rules

Individuals choose an optimal decision rule $\delta_t(s_t, \theta)$ such that

$$d_t = \delta_t(s_t, \theta)$$

maximizes expected lifetime utility given by

$$V_t(s_t) = \max_{\delta} \left\{ \mathbb{E} \left[\sum_{\tau=t}^T \beta^{\tau-t} u(s_{\tau}, d_{\tau}, \theta) \middle| s_t = s \right] \right\}$$

subject to the feasible constraints. $V_t(s)$ is the value function associated with the dynamic discrete choice problem.

Bellman Equation

We can also express the individuals optimization problem recursively using Bellman's equation:

$$V_t(s) = \max_{d \in D_t(s)} \left\{ u(s, d) + \beta \int V_{t+1}(s') \pi(s'|s, d, \theta) ds' \right\} \quad (1)$$

In many applications, we will discretize the state space, i.e. we will assume that the support of s has a finite number of elements. Let S the number of possible states. In that case, we have:

$$V_t(s) = \max_{d \in D_t(s)} \left\{ u(s, d) + \beta \sum_{s'=1}^S V_{t+1}(s') \pi(s'|s, d, \theta) \right\}$$

Solving DP Problems with a Finite State Space

We can use a technique called value function iteration. In the finite horizon case ($T < \infty$), this is equivalent by solving the model backward:

$$V_T(s) = \max_{d \in D_T(s)} u(s, d)$$
$$V_t(s) = \max_{d \in D_t(s)} \left\{ u(s, d) + \beta \sum_{s'} V_{t+1}(s') \pi(s'|s, d) \right\}$$

Finite time horizon problems are non-stationary, i.e. the value function and the policy function depend on time, t .

The Stationary Infinite Horizon Case

In the stationary infinite horizon case, we have

$$V_t(s_t) = V(s_t) \text{ for all } t$$

To compute $V(s_t)$, we start with an initial guess of the value function, denoted by $V(s)$. We then iterate and update the value function:

$$V'(s) = \max_{d \in D_t(s)} \left\{ u(s, d) + \beta \sum_{s'} V(s') \pi(s'|s, d) \right\}$$

until convergence:

$$\|V(s) - V'(s)\| \leq c \text{ for all } s$$

for some convergence level c . Recall that this algorithm converges quickly since the algorithm is a contraction mapping. The algorithm computes a fixed point of the contraction mapping.

Policy Function Iteration

- ▶ Instead of iterating on the value function, we could also iterate on the policy function.
- ▶ Pick a policy function $\delta(s)$. The value function that is induced by this policy function is given by:

$$V_\delta(s) = u(s, \delta(s)) + \beta \sum_{s'} V_\delta(s') \pi(s'|s, \delta(s))$$

or in matrix algebra

$$\underbrace{V_\delta}_{[S \times 1]} = \underbrace{u}_{[S \times 1]} + \beta \underbrace{E_\delta}_{[S \times S]} \underbrace{V_\delta}_{[S \times 1]}$$

Hence we get

$$(I - \beta E_\delta) V_\delta = u$$

This is a linear system of equations which can be solved to obtain:

$$V_\delta = (I - \beta E_\delta)^{-1} u$$

Now we can update the policy function:

$$\delta'(s) = \operatorname{argmax}_{d \in D(s)} \left\{ u(s, d) + \beta \sum_{s'} V_\delta(s') \pi(s'|s, d) \right\}$$

And iterate on this until convergence:

$$\|\delta(s) - \delta'(s)\| \leq c$$

Observed and Unobserved State Variables

- ▶ If the goal of the analysis is to estimate the parameters of the model, we need to impose additional assumptions.
- ▶ We partition the state space into two components:
 - ▶ x observed by the econometrician
 - ▶ ε unobserved by the econometrician
- ▶ Note that we still assume that $s = (x, \varepsilon)$ is observed by the individual.

Deterministic Decision Rules

- ▶ Optimal decision rules of DP models are deterministic. (See Rust (1994) for a discussion.)
- ▶ We, therefore, need unobserved state variables to obtain a well-behaved econometric model.
- ▶ Let us rewrite the decision rule as

$$\delta(s, \theta) = \delta(x, \varepsilon, \theta)$$

- ▶ The idea is to obtain conditional choice probabilities by integrating out the state variables that are not observed by the econometrician.
- ▶ This is the same idea that McFadden (1974) used for static discrete choice models.

Conditional Independence

- ▶ To estimate the transition probabilities, it is useful to assume that:

$$\pi(x_{t+1}, \varepsilon_{t+1} | x_t, \varepsilon_t, d_t, \theta) = q(\varepsilon_{t+1} | x_{t+1}) f(x_{t+1} | x_t, d_t, \theta)$$

- ▶ x_{t+1} is a sufficient statistic for ε_{t+1}
- ▶ In practice, we often assume:

$$q(\varepsilon_{t+1} | x_{t+1}) = q(\varepsilon_{t+1})$$

- ▶ $f(x_{t+1} | x_t, d_t, \theta)$ does not depend on ε_t .

Additive Separability

Furthermore, we assume that the utility function is additively separable in ε_t :

$$\begin{aligned} u(s_t, d_t, \theta) &= u(x_t, d_t, \theta) + \varepsilon_t(d_t) \\ &= \sum_i d_{i,t} (u_i(x_t, \theta) + \varepsilon_{i,t}) \end{aligned}$$

This assumption makes it easier to integrate out the error terms.

Conditional Value Function

Under these assumptions, the Bellman equation takes the form:

$$V(x, \varepsilon) = \max_{d \in D(x)} \{v(x, d) + \varepsilon(d)\},$$

where $v(x, d)$ is called the conditional value function, which is given by:

$$v(x, d) = u(x, d) + \beta \int \left[\int V(x', \varepsilon') q(\varepsilon' | x') d\varepsilon' \right] f(x' | x, d) dx' \quad (2)$$

In some papers $v_j(x_t)$ is used to denote the conditional value function:

$$V_t(x, \varepsilon) = \max \left\{ \sum_{j=1}^J d_{jt} (v_{jt}(x) + \varepsilon_{jt}) \right\}$$

Conditional Choice Probabilities

By integrating out the unobserved state variables, we obtain well-behaved conditional choice probabilities:

$$\begin{aligned} P[d|x, \theta] &= \int 1[d = \delta(x, \varepsilon, \theta)] q(\varepsilon|x) d\varepsilon \\ &= \int 1[d = \operatorname{argmax} \{v(x, d) + \varepsilon(d)\}] q(\varepsilon|x) d\varepsilon \quad (3) \end{aligned}$$

Note that the ε 's need to be continuous with unbounded support for the choice probabilities to be well-defined.

The Likelihood Function

Given a random sample $\{x_t^i, d_t^i\}_{i=1, t=1}^{N, T}$, the likelihood function is given by:

$$L(\theta) = \prod_{i=1}^N \prod_{t=1}^T P_t(d_t^i | x_t^i, \theta) f(x_{t+1}^i | x_t^i, d_t^i, \theta)$$

The likelihood function has two components:

- a) the conditional choice probabilities
- b) the transition probabilities

A Nested Fixed Point Algorithm

Estimation consists of:

- ▶ an outer loop that searches over the parameter space;
- ▶ an inner loop that evaluates the conditional value functions of each parameter vector.

Unobserved Types

- ▶ It is often useful to control for unobserved heterogeneity among agents.
- ▶ Following Heckman and Singer (1984), define $m = 1, \dots, M$ different types and assume that preferences and beliefs are type-specific.
- ▶ Assuming that each type occurs with probability q_m , we have

$$L(\theta) = \prod_{i=1}^N \sum_{m=1}^M \left(q_m \prod_{t=1}^T P_t^m(d_t^i | x_t^i, \theta) f_t^m(x_{t+1}^i, x_t^i, d_t^i, \theta) \right)$$

- ▶ Arcidiacono and Jones (2003) discuss how to use an EM algorithm to sequentially estimate this model.

Example: Keane & Wolpin

Let us assume that errors are normally distributed.

1. Pick parameters vector.
2. Solve equation (1) given by $V(s) = \Gamma(V(s))$ using value function iteration for the finite time horizon.
3. Compute $v(x, d)$ using (2).
4. Compute conditional choice probabilities using (3).
5. Evaluate likelihood.
6. Check for convergence.
7. Repeat until convergence of likelihood function.

Discussion

This approach is time-intensive since it requires:

- ▶ numerical integration,
- ▶ repeated optimization,
- ▶ a relatively large state space.

The advantage is that we can deal with non-stationary problems and can easily extend the method to deal with partially unobserved state variables (such as wages.)

Type I Extreme Value Errors (McFadden & Rust)

The model above can be simplified by assuming that ε follow a type I extreme value distribution. Setting $\sigma = 1$, we obtain

$$P[d|x, \theta] = \frac{\exp v(x, d, \theta)}{\sum_{d'} \exp v(x, d', \theta)}$$

Recall that $v(x, d, \theta)$ is the conditional value function associated with choice d .

The Social Surplus Function

The expected value function or the social surplus function (McFadden '81) is defined as

$$\begin{aligned} V(x) &= \int V(x, \varepsilon) q(\varepsilon|x) d\varepsilon \\ &= \int \max_d \left\{ u(x, d) + \varepsilon(d) + \beta \int V(x') f(x'|x, d) dx' \right\} q(\varepsilon|x) d\varepsilon \\ &= \Gamma(V(x)) \end{aligned} \tag{4}$$

Thus $V(x)$ is a fixed point of the mapping above.

Expected Value Function and Conditional Value Function

Using our definitions, we have

$$v(x, d) = u(x, d) + \beta \int V(x') f(x'|x, d) dx'$$

and

$$V(x) = \int \max_d \{v(x, d) + \varepsilon(d)\} q(\varepsilon|x) d\varepsilon$$

Properties of Extreme Value Distributions I

Let (Y_1, \dots, Y_D) be INID (independent, non-identically distributed) extreme value random variables with location parameters $(\alpha_1, \dots, \alpha_D)$ and common scale parameter σ , i.e. the distribution of Y_d is given by:

$$\begin{aligned} F(y | \alpha_d, \sigma) &= P\{Y_d \leq y | \alpha_d, \sigma\} \\ &= \exp \left\{ - \exp \left\{ \frac{-(x - \alpha_d)}{\sigma} \right\} \right\}. \end{aligned} \tag{5}$$

Properties of Extreme Value Distributions II

We would like to show that this family is max-stable by proving that $\max(Y_1, \dots, Y_D)$ is an extreme value random variable with scale parameter σ and location parameter

$$\alpha = \sigma \ln \left[\sum_{d=1}^D \exp\{\alpha_d/\sigma\} \right] \quad (6)$$

Properties of Extreme Value Distributions III

$$\begin{aligned}P(\max_d Y_d \leq x) &= \prod_{d=1}^D P(Y_d \leq x) \\&= \prod_{d=1}^D \exp \left\{ -\exp \left\{ \frac{-(x - \alpha_d)}{\sigma} \right\} \right\} \\&= \exp \left\{ \sum_d -\exp \left\{ \frac{-(x - \alpha_d)}{\sigma} \right\} \right\} \\&= \exp \left\{ \exp \left\{ -\frac{x}{\sigma} \right\} \sum_d \exp \left\{ \frac{\alpha_d}{\sigma} \right\} \right\} \\&= \exp \left\{ -\exp \left\{ \frac{-(x - \sigma \ln \sum \exp \frac{\alpha_d}{\sigma})}{\sigma} \right\} \right\}\end{aligned}$$

Properties of Social Surplus Function I

We would like to show that the Social Surplus function or expected value function has the form given by

$$V(X) = \sigma\gamma + \sigma \ln \left[\sum_{d=1}^D \exp\{v_d/\sigma\} \right].$$

we use the fact that if $\{\epsilon_d\}$ are independent random variables, we have following formula for the probability distribution of the random variable $\max_{d=1,\dots,D}[v_d + \epsilon_d]$:

$$\Pr \left\{ \max_{d=1,\dots,D}[v_d + \epsilon_d] \leq x \right\} = \prod_{d=1}^D \Pr \{v_d + \epsilon_d \leq x\}.$$

Now, let ϵ_d have an extreme value distribution with location parameter $\alpha_d = 0$ and scale parameter $\sigma > 0$. Then it is easy to see that $v_d + \epsilon_d$ is also an extreme value random variate with location parameter v_d and scale parameter σ .

Properties of Social Surplus Function II

As we showed above, $\max[v_d + \epsilon_d]$ is also extreme value. Hence

$$\Pr \left\{ \max_{d=1,\dots,D} [v_d + \epsilon_d] \leq x \right\} = \exp \left\{ - \exp \left\{ \frac{-(x - \alpha)}{\sigma} \right\} \right\},$$

where the location parameter is given by $\ln \left[\sum_{d=1}^D \exp\{v_d/\sigma\} \right]$.

Recall that the expectation of a single extreme value random variable, $\tilde{\epsilon}$, with location parameter α and scale parameter σ is given by:

$$E(\tilde{\epsilon}) = \alpha + \sigma \gamma, \tag{7}$$

where $\gamma = .577216\dots$ is Euler's constant.

The form of the Social Surplus Function then follows from the formula for the expectation of an extreme value random variable.

Deriving the Conditional Choice Probabilities I

We want to show that the partial derivative of $\max_{d=1,\dots,D}[v_d + \epsilon_d]$ equals the indicator function $I\{d = \delta(\epsilon)\}$. Consider first the case that alternative d yields the highest utility, i.e. $d = \delta(\epsilon)$. We then have

$$v_d + \epsilon_d > v_{d'} + \epsilon_{d'} \quad \forall d' \neq d$$

Thus

$$v_d + \epsilon_d = \max_{d'=1,\dots,D} [v_{d'} + \epsilon_{d'}]$$

and hence we have

$$\frac{\partial \max_{d=1,\dots,D}[v_d + \epsilon_d]}{\partial v_d} = 1$$

Deriving the Conditional Choice Probabilities II

Next consider the case that alternative d is not the utility maximizing choice, i.e. $d \neq \delta(\epsilon)$. In that case we have:

$$\max_{d'=1,\dots,D} [v_{d'} + \epsilon_{d'}] > v_d + \epsilon_d$$

It follows that we have

$$\frac{\partial v_d \max_{d=1,\dots,D} [v_d + \epsilon_d]}{\partial v_d} = 0$$

Thus far we have assumed that the optimal choice is unique. For the sake of completeness consider the case in which the optimal choice is not unique. However, based on our assumption on the distribution of error terms, this case has probability zero. Thus we conclude that the identity above holds with probability 1.

Deriving the Conditional Choice Probabilities III

We want to appeal to the Lebesgue Dominated Convergence Theorem to justify the interchange of integration and differentiation operators. As long as the distribution of the $\{\epsilon_d\}$'s has a density, the derivative

$$\frac{\partial \max_{d=1,\dots,D}[v_d + \epsilon_d]}{\partial v_d} = I\{d = \delta(\epsilon)\}.$$

exists almost everywhere with respect to this density and is bounded by 1, so that the Lebesgue Dominated Convergence Theorem applies.

Deriving the Conditional Choice Probabilities IV

As a consequence, we obtain the Williams-Daly-Zachary Theorem:

$$\begin{aligned}\frac{\partial V(X)}{\partial v_d} &= \frac{\partial}{\partial v_d} \int \dots \int \max_{d=1,\dots,D} [v_d + \epsilon_d] f(\epsilon_1, \dots, \epsilon_D | X) d\epsilon_1 \dots d\epsilon_D \\ &= \int \dots \int \frac{\partial}{\partial v_d} \max_{d=1,\dots,D} [v_d + \epsilon_d] f(\epsilon_1, \dots, \epsilon_D | X) d\epsilon_1 \dots d\epsilon_D \\ &= \int \dots \int I\{d = \delta(\epsilon_1, \dots, \epsilon_D)\} f(\epsilon_1, \dots, \epsilon_D | X) d\epsilon_1 \dots d\epsilon_D \\ &= P\{d|X\}\end{aligned}$$

Deriving the Conditional Choice Probabilities V

Then by the Williams-Daly-Zachary Theorem we have

$$\begin{aligned} P\{d|X\} &= \frac{\partial}{\partial v_d} \left[\sigma \gamma + \sigma \ln \left(\sum_{d=1}^D \exp(v_d/\sigma) \right) \right] \\ &= \frac{\exp v_d/\sigma}{\sum_{d'=1}^D \exp v_{d'}/\sigma}. \end{aligned}$$

Setting $\sigma = 1$ gives the familiar result.

Another Contraction Mapping

We have shown that the type I extreme value assumption then implies the following result:

$$V(x) = \gamma + \ln \left(\sum_{d'} \exp v(x, d') \right)$$

Hence $v(x, d)$ is the unique fixed point of the mapping:

$$v(x, d) = u(x, d) + \beta \int \left(\gamma + \ln \left(\sum_{d'} \exp v(x', d') \right) \right) f(x'|x, d) dx'$$

The Finite Horizon Case

If T is finite, then $v_t(x, d)$ is recursively defined as

$$v_T(x_t, d_t, \theta) = u_T(x_t, d_t, \theta)$$

$$v_t(x_t, d_t, \theta) = u(x_t, d_t, \theta) + \beta \int \left(\gamma + \ln \left(\sum_{d_{t+1}} \exp v_{t+1}(x_{t+1}, d_{t+1}, \theta) \right) \right) f(x_{t+1} | x_t, d_t, \theta) dx_{t+1}$$

In the infinite horizon case, $v(x, d, \theta)$ is just the fixed point of the contraction mapping above.

The Rust Algorithm

The Rust algorithm the consists of the following steps:

1. pick parameter vector θ
2. compute $v(x, d)$ by solving (4)
3. evaluate $P(d|x)$ (closed form solution)
4. evaluate likelihood
5. check for convergence
6. repeat

Discussion

The type I extreme value assumption facilitates the analysis because

- ▶ it provides a closed form solution of the conditional choice probabilities;
- ▶ reduces the curse of dimensionality by focusing on $v(x, d)$.

The main drawback of this approach is that we still need to compute $v(x, d)$ numerically, which is computationally expensive if x is large.

CCP Estimation (Hotz & Miller)

- ▶ We can further simplify the computation of the conditional value functions by conditioning on observed choices.
- ▶ Hotz and Miller's (1993) Conditional Choice Probability estimator (CCP) accomplishes this.
- ▶ The key insight of HM is that we can invert equation (3) and express (the normalized differences of) $v(x, d)$ as functions of the conditional choice probabilities $P(d|x)$.
- ▶ Hotz, Miller, Sanders and Smith (1994) extend these ideas to formulate an estimator of DDC model that relies on forward simulation to compute the conditional value functions, instead of backward induction.

Swapping the Nested Fixed Point Algorithm

To see how this works, we follow Aguirregabiria and Mira (2002). We treat the state space as discrete and rewrite (4) as

$$V(x) = \sum_d P(d|x) \left(u(x, d) + E[\varepsilon|d, x] + \beta \sum_{x'} V(x') f(x'|x, d) \right) \quad (8)$$

where

$$E[\varepsilon|x, d] = \frac{1}{P(d|x)} \int \varepsilon(d) 1[\delta(x, \varepsilon) = d] q(\varepsilon|x) d\varepsilon \quad (9)$$

Another Nice Property of Type I Extreme Value Errors

If ε is distributed as type I extreme value, then it can be shown that:

$$E[\varepsilon|x, d] = \gamma - \ln(P(d|x)) = E[\varepsilon|P, d]$$

where γ is Euler's constant.

This result is important since we can now express the expected value function as a function of the structural parameters and the conditional choice probabilities.

Consider, for simplicity, the case where $D = 2$. We want to show that:

$$E(\tilde{\epsilon}_d | d, X) = \gamma - \ln(P\{d|X\})$$

To simplify the notation, define $p_d = P\{d|X\}$. First, note that:

$$\begin{aligned} v_1 - v_2 &= \ln(p_1) - \ln(p_2) \\ &= \ln(p_1/(1 - p_1)) \end{aligned}$$

Second, $d = 1$ if and only if:

$$\begin{aligned}\epsilon_2 &< \epsilon_1 + v_1 - v_2 \\ &= \epsilon_1 + \ln(p_1/(1 - p_1))\end{aligned}$$

Third, notice that:

$$\begin{aligned}E(\epsilon_1 | d = 1, X) &= \int_{-\infty}^{\infty} \int_{-\infty}^{\epsilon_1 + \ln(p_1/(1-p_1))} \epsilon_1 f(\epsilon_1, \epsilon_2) d\epsilon_2 d\epsilon_1 / p_1 \\ &= \int_{-\infty}^{\infty} \epsilon_1 F(\epsilon_1 + \ln(p_1/(1 - p_1))) f(\epsilon_1) d\epsilon_1 / p_1\end{aligned}$$

Substituting in the density and distribution function of an extreme value distribution and doing some algebra yields:

$$E(\epsilon_1 | d = 1, X) = \int_{-\infty}^{\infty} \epsilon_1 \exp(-(\epsilon_1 + \ln p_1)) \exp(-\exp(-(\epsilon_1 + \ln p_1))) d\epsilon_1$$

The result then follows from the formula for the expectation of an extreme value random variable.

Using Matrix Algebra

Stacking the M equations, we can rewrite (5) using matrix algebra as:

$$\underbrace{V}_{[M \times 1]} = \sum_d \underbrace{P(d)}_{[M \times 1]} .* \left(\underbrace{u(d)}_{[M \times 1]} + \underbrace{e(d, P)}_{[M \times 1]} + \beta \underbrace{F(d)}_{[M \times M]} V \right)$$

where $.*$ represents element-by-element multiplication

An Example

$X \in \{1, 2\}$, and $d \in \{a, b\}$. Then we have:

$$\begin{aligned} V(1) &= \mathbb{P}[a|1] (u(1, a) + e(1, a) + \beta (f_{11}^a V(1) + f_{12}^a V(2))) \\ &\quad + \mathbb{P}[b|1] (u(1, b) + e(1, b) + \beta (f_{11}^b V(1) + f_{12}^b V(2))) \end{aligned}$$

$$\begin{aligned} V(2) &= \mathbb{P}[a|2] (u(2, a) + e(2, a) + \beta (f_{21}^a V(1) + f_{22}^a V(2))) \\ &\quad + \mathbb{P}[b|2] (u(2, b) + e(2, b) + \beta (f_{21}^b V(1) + f_{22}^b V(2))) \end{aligned}$$

$$V(1) = \sum_d \mathbb{P}[d|1] \left(u(1, d) + e(1, d) + \beta \sum_i f_{1i}^d V(i) \right)$$

$$V(2) = \sum_d \mathbb{P}[d|2] \left(u(2, d) + e(2, d) + \beta \sum_i f_{2i}^d V(i) \right)$$

$$V(x) = \sum_d \mathbb{P}[d|x] \left(u(x, d) + e(x, d) + \beta \sum_i f_{x,i}^d V(i) \right)$$

Hence $V(x)$ is defined by a linear system of M equations in M unknowns. This system can be solved to obtain:

$$\begin{aligned} V &= \varphi(P) \\ &= (I_M - \beta F^u(P))^{-1} \left(\sum_d P(d) \cdot (u(d) + e(d, P)) \right) \end{aligned} \quad (10)$$

where

$$F^u(P) = \sum_d P(d) F(d)$$

is the $M \times M$ transition matrix induced by P . Note that the solution of the expected value function depends on θ since u and F depend on θ .

The HM-AM algorithm:

1. Obtain a nonparametric estimate of $P^N(d)$.
2. Pick θ
3. Compute $V^N(x, \theta) = \varphi(P^N, \theta)$
4. Compute $v^N(x, d, \theta)$.
5. Compute $P(d|x, \theta)$.
6. Compute likelihood.
7. Iterate until convergence of the likelihood function.

Some Additional Discussion

- ▶ The main advantage of the HM-AM algorithm is we only need to solve a system of linear equations to obtain the expected value functions.
- ▶ The main disadvantage is that we need to be able to estimate the full set of conditional choice probabilities which places high demands on the available data.
- ▶ AM (2002) show that we can generalize the HM approach by iterating on the choice probabilities. They also show that this estimator is asymptotically efficient.
- ▶ Arcidiacono and Miller (2012) show how to extend the CCP estimators to allow for unobserved heterogeneity of the Heckman-Singer type.
- ▶ Magnac and Thesmar (2002) use the HM inversion result to show that DDC models are non-parametrically not identified, extending an earlier result by Rust (1994).