Introduction to Dynamic Discrete Choice Models

Alejandro Robinson-Cortés
EMERG: March 22, 2023

Setup. Consider a **forward-looking agent** i who chooses an **action** $a_{it} \in \mathcal{A}$ on every **period** $t \in \{1, 2, \ldots, T\}$, with $|\mathcal{A}| < \infty$ and $T \leq \infty$. Let $u_a(x_{it}, \varepsilon_{it})$ denote the agent's **per-period utility** of choosing $a \in \mathcal{A}$, where $(x_{it}, \varepsilon_{it})$ are **state variables** observed by the agent, with $x_{it} \in \mathcal{X}$, $|\mathcal{X}| < \infty$, and $\varepsilon_{it} = (\varepsilon_{ait})_{a\in\mathcal{A}} \in \mathbb{R}^{|\mathcal{A}|}$. Econometrician observes **data** of N agents: $\{(a_{it}, x_{it})\}_{i=1,\ldots,N;t=1,\ldots,T}$.[1] State variables $(x_{it}, \varepsilon_{it})$ evolve according to **Markov process** with transition probability F, which satisfies **conditional independence (CI)**:

$$\text{(CI)} \qquad F(x_{it+1}, \varepsilon_{it+1}) = F(x_{it+1}\,|\,a_{it}, x_{it})\, G(\varepsilon_{it+1}).$$

Assume ε_{it+1} is i.i.d. across agents and time, and G is absolutely continuous. Choice data generated by **dynamic program (DP)**:

$$\text{(DP)} \qquad V(x_0, \varepsilon_0) = \max_{(a_{it})_{t=1}^{T}} \mathbb{E}\left[\sum_{t=1}^{T} \beta^{t-1} u_{a_{it}}(x_{it}, \varepsilon_{it})\,|\,x_0\right],$$

where $V(\,\cdot\,)$ is the **value function**, the expectation is taken with respect to the transition F (conditioning on the chosen a's), and $\beta \in [0, 1)$ is a **discount factor**. **Empirical exercise** consists in estimating utility function $u_a(\,\cdot\,)$ that rationalizes data.

Bus engines (Rust 1987). Model bus engine replacement choices of Harold Zurcher, superintendent of maintenance at the Madison (Wisconsin) Metropolitan Bus Company. Assume $N = 1$ and $T = \infty$; $\mathcal{A} = \{0, 1\}$ where $a = 1\{\text{replace engine}\}$; x_t is mileage with $|\mathcal{X}| = 3$ (discretise continuous data).[2] Per-period utility function:

$$(1) \qquad u_a(x, \varepsilon; \theta) = -c((1-a)x; \theta) - aRC + \varepsilon_a, \quad a \in \{0, 1\},$$

where $c(x; \theta)$ are maintenance costs given mileage x (increasing in x; known up to parameter vector θ), $RC \geq 0$ unknown costs of replacing the engine (to be estimated), and $\varepsilon = (\varepsilon_0, \varepsilon_1)$ is i.i.d. **Type 1 Extreme Value (T1EV)**. Replacing engine resets mileage to $x = 0$. Process is stationary (infinite horizon). Under (CI), the value function satisfies **Bellman equation**:

$$(2) \qquad V(x, \varepsilon) = \max_{a\in\{0,1\}} \left\{u_a(x, \varepsilon; \theta) + \beta\mathbb{E}\left[V(x', \varepsilon')\,|\,a, x; \theta\right]\right\},$$

where (x', ε') denotes the state variables in the next period. Intuitively, β is *not* identified in this model. Data with high discounting β and high replacement costs RC look similar as with low β and low RC. Hence, assume β is known. Rust considers two extreme cases: $\beta = 0.9999$ and $\beta = 0$ (**myopic**).

Let $EV(a, x; \theta) = \mathbb{E}[V(x', \varepsilon')\,|\,a, x; \theta]$ be the **expected value function**. Let $\bar{u}_a(x; \theta) \equiv u_a(x, \varepsilon; \theta) - \varepsilon_a$. From (2), the **conditional choice probability (CCP)** takes the logit-form:

$$(3) \qquad p(a\,|\,x; \theta) = \frac{\exp\left(\bar{u}_a(x; \theta) + \beta EV(a, x; \theta)\right)}{\sum_{\tilde{a}\in\{0,1\}} \exp\left(\bar{u}_{\tilde{a}}(x; \theta) + \beta EV(\tilde{a}, x; \theta)\right)}.$$

[1] In the data, T is always finite; however, the model admits an infinite horizon, i.e., $T = \infty$.

[2] We omit discussion of the finite horizon case ($T < \infty$). In this case, the period-specific value functions can be computed by **backward induction**.

Under conditional independence, the log-likelihood of data $(\mathbf{a}, \mathbf{x}) = (a_t, x_t)_{t=1}^T$ is:

$$\text{(4)} \qquad \log \mathcal{L}(\mathbf{a}, \mathbf{x}; \theta) = \sum_{t=1}^{T} \log p(a_t \,|\, x_t; \theta) + \sum_{t=1}^{T} \log F(x_{t+1} \,|\, a_t, x_t).$$

One can estimate θ and F separately (F can be estimated nonparametrically or one can assume a parametric model). To evaluate $\log \mathcal{L}(\mathbf{a}, \mathbf{x}; \theta)$ at θ, we need to compute $EV(a, x; \theta)$ at every point in $\mathcal{A} \times \mathcal{X}$. In Rust's case, it has $|\mathcal{A}| \times |\mathcal{X}| = 2 \times 3 = 6$ points, so $EV(\,\cdot\,)$ is a vector in $\mathbb{R}^6$. Rust shows that the expected value function satisfies the Bellman equation[3]:

$$\text{(5)} \qquad EV(a, x; \theta) = \sum_{\tilde{x} \in \mathcal{X}} \log \left\{ \sum_{\tilde{a} \in \mathcal{A}} \exp \Big(\bar{u}_{\tilde{a}}(\tilde{x}; \theta) + \beta EV(\tilde{a}, \tilde{x}; \theta) \Big) \right\} F(\tilde{x} \,|\, a, x).$$

Equation (5) is a **contraction mapping**, so we can find its **fixed-point**—the conditional value function $EV(\,\cdot\,)$—through iteration. Very useful computationally; otherwise, would need to compute value function $V(x, \varepsilon)$, which is infinite-dimensional. This results in the **nested fixed-point (NFXP)** algorithm: given θ, in the **inner loop** iterate over (5) to find $EV(\,\cdot\,)$; in the **outer loop**, look for θ to maximize $\log \mathcal{L}(\mathbf{a}, \mathbf{x}; \theta)$.

CCP estimation (Hotz and Miller 1993). A drawback of the NFXP algorithm is that it requires computation of the expected value function at each iteration of the parameter vector θ. Hotz and Miller (HM) propose an estimator that does involve solving a DP. Their approach exploits a relationship between the CCPs and the **choice-specific (aka. conditional) value function** $V_a(x; \theta)$, which is defined by:

$$\text{(6)} \qquad V_a(x; \theta) = \bar{u}_a(x; \theta) + \beta \mathbb{E} \left[V(x', \varepsilon') \,|\, a, x; \theta \right], \quad a \in \mathcal{A}.$$

Let $p(x; \theta) = \big(p(a \,|\, x; \theta) \big)_{a \in \mathcal{A}}$. For any distribution $G(\varepsilon')$, HM show that under (CI):

$$\text{(7)} \qquad V_a(x; \theta) - V_{a'}(x; \theta) = \phi_{a,a'}(p(x; \theta)), \quad a, a' \in \mathcal{A}$$

where $\phi_{a,a'}(\,\cdot\,)$ is an invertible real-valued function which only depends on G. This is known as the **Hotz-Miller inversion**. In the T1EV case, $\phi_{a,a'}(p) = \log p_a - \log p_{a'}$. The key of the HM-estimator is to note that $p(x; \theta)$ can be estimated directly from the data nonparametrically, and we can use these estimated CCPs, $\hat{p}(x)$, to **forward-simulate** $V_a(x; \theta)$.[4] First, use (DP) to write (6):

$$\text{(8)} \qquad V_a(x; \theta) = \bar{u}_a(x; \theta) + \beta \mathbb{E} \Big[\bar{u}_{a'}(x'; \theta) + \varepsilon'_{a'} + \beta \mathbb{E} \Big[\bar{u}_{a''}(x''; \theta) + \varepsilon''_{a''} + \beta \mathbb{E} [\,\cdots] \,|\, a', x'; \theta \Big] \,|\, a, x; \theta \Big]$$

Given θ, knowledge of F and G, and $\hat{p}(x)$, we can simulate the expectations in the previous equation with respect to $(a' \,|\, x)$, $(x' \,|\, a, x)$, $(a'' \,|\, x')$, $(x'' \,|\, a', x')$, and so on. Moreover, the conditional expectation $\mathbb{E}[\varepsilon'_a \mid a, x]$ has a closed-form solution in the logit case: $\mathbb{E}[\varepsilon'_a \mid a, x] = \gamma - \log p(a \mid x)$. Simulate S sequences of values $x' \sim F(\,\cdot\,|\, a, x)$, $a' \sim \hat{p}(x')$, $x'' \sim F(\,\cdot\,|\, a', x')$, etc., then approximate $V_a(x; \theta)$ taking averages across simulations:

[3] Equation (5) uses the expectation of the maximum of T1EV variables, i.e., if $(\varepsilon_a)_{a \in \mathcal{A}}$ are i.i.d. T1EV, then $\mathbb{E} \left[\max_{a \in \mathcal{A}} \{ u_a + \varepsilon_a \} \right] = \log \left(\sum_{a \in \mathcal{A}} \exp(u_a) \right) + \gamma$, where $\gamma \approx 0.57$ is Euler's constant.

[4] The nonparametric estimates $\hat{p}(x)$ depend directly on the data; they are *not* a function of the parameter vector θ, which we are trying to estimate. This estimation can be done by sample averages of kernel regression, which in practice requires $|\mathcal{X}|$ to not be too large.

$$(9) \qquad V_a(x;\theta) \approx \frac{1}{S}\sum_{s=1}^{S} \bar{u}_a(x;\theta) + \beta\Big[\bar{u}_{a'}(x';\theta) + \gamma - \log\hat{p}(a'\,|\,x') + \beta\Big[\bar{u}_{a''}(x'';\theta) + \gamma - \log\hat{p}(a''\,|\,x'') + \beta[\cdots]\Big]\Big].$$

In practice, truncate simulated series at some finite period T_0. Let $\tilde{V}_a(x;\theta)$ denote the estimated choice-specific value functions (i.e., the RHS of (9)). The **predicted conditional choice probabilities** are then:

$$(10) \qquad \tilde{p}(a\,|\,x;\theta) = \frac{\exp\left(\tilde{V}_a(x;\theta)\right)}{\sum_{\tilde{a}\in\{0,1\}}\exp\left(\tilde{V}_{\tilde{a}}(x;\theta)\right)}.$$

Let $\tilde{p}(x;\theta) = \left(\tilde{p}(a\,|\,x;\theta)\right)_{a\in\mathcal{A}}$. The predicted CCPs $\tilde{p}(x;\theta)$ depend on the specific parameter θ used in the forward-simulation and, in general, are different to the actual CCPs estimated from the data $\hat{p}(x)$. HM propose two estimators for θ. The first one matches the predicted CCPs to the ones estimated directly from the data:

$$(11) \qquad \hat{\theta} = \arg\min_{\theta} \|\tilde{\mathbf{p}}(\theta) - \hat{\mathbf{p}}\|,$$

where $\tilde{\mathbf{p}}(\theta) = \left(\tilde{p}(x;\theta)\right)_{x\in\mathcal{X}}$ and $\hat{\mathbf{p}} = \left(\hat{p}(x)\right)_{x\in\mathcal{X}}$. Alternatively, one can also use the HM inversion and exploit the T1EV assumption to estimate θ via:

$$(12) \qquad \hat{\theta} = \arg\min_{\theta} \left\|\left[\tilde{\mathbf{V}}_1(\theta) - \tilde{\mathbf{V}}_0(\theta)\right] - \left[\log\hat{\mathbf{p}}_1 - \log\hat{\mathbf{p}}_0\right]\right\|,$$

where $\tilde{\mathbf{V}}_a(\theta) = \left(\tilde{V}_a(x;\theta)\right)_{x\in\mathcal{X}}$ and $\hat{\mathbf{p}}_a = \left(\hat{p}(a \mid x)\right)_{x\in\mathcal{X}}$. A key advantage of the HM approach to **estimation via CCPs** is that the simulation draws in $\tilde{V}_a(x;\theta)$ stay fixed throughout the estimation (i.e., as we vary θ to find the optimum). HM study the large sample properties of both estimators.

More on estimation. While the Hotz-Miller's CCP estimator is computationally less burdensome than NFXP, it is less efficient due to simulations (in finite samples and asymptotically). Aguirregabiria and Mira (2002) show that a pseudo-maximum likelihood version of Hotz-Miller's CCP estimator is asymptotically equivalent to partial MLE. For an estimation method that exploits the discreteness of $\mathcal{X}$, and allows computation of an integrated value function via matrix inversion, see Aguirregabiria and Mira (2007), and Pesendorfer and Schmidt-Dengler (2008). For estimation based on mathematical programming with equilibrium constraints (MPEC), see Su and Judd (2012).

Identification. In general, dynamic discrete choice (DDC) models are *not* identified nonparametrically (Magnac and Thesmar 2002). It is common to impose parametric assumptions along the lines of Rust (1987):

$$(13) \qquad u_a(x,\varepsilon) = u(a,x;\theta) + \varepsilon_a,$$

where $u(a,x;\theta)$ is known up to the parameter vector θ, and the distribution of $\varepsilon = (\varepsilon_a)_{a\in\mathcal{A}}$ is assumed to be known (usually assumed to be T1EV). Similarly, it is also common to assume β is known. Formally, Magnac and Thesmar (2002) show that $u_a(\,\cdot\,)$ is nonparametrically identified under knowledge of β, G (the distribution of ε), and a utility normalization. Without knowledge of β and G, the model is nonparametrically underidentified. For some recent results on nonparametric identification of the discount factor β and counterfactuals in DDC models, see Abbring and Daljord (2020), and Kalouptsidi, Scott, and Souza-Rodrigues (2021), respectively.

Applications. DDC models are used in multiple fields in economics, e.g., labor, industrial organization, education, and health. Examples of applications in the literature include labor market transitions, education choices, retirement decisions, fertility and contraceptive choices, search models, dynamic games, demand for durable goods, auctions, etc. Much of the literature has focused on relaxing some of the assumptions in Rust (1987). For methodological reviews and examples of applications, see Keane and Wolpin (2009), Aguirregabiria and Mira (2010), Arcidiacono and Ellickson (2011), and Shum (2017).

References

Abbring, Jaap H., and Øystein Daljord. 2020. "Identifying the discount factor in dynamic discrete choice models." *Quantitative Economics* 11 (2): 471–501.

Aguirregabiria, Victor, and Pedro Mira. 2002. "Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models." *Econometrica* 70 (4): 1519–1543.

———. 2007. "Sequential Estimation of Dynamic Discrete Games." *Econometrica* 75 (1): 1–53.

———. 2010. "Dynamic discrete choice structural models: A survey." Structural Models of Optimization Behavior in Labor, Aging, and Health, *Journal of Econometrics* 156 (1): 38–67.

Arcidiacono, Peter, and Paul B. Ellickson. 2011. "Practical Methods for Estimation of Dynamic Discrete Choice Models." *Annual Review of Economics* 3 (1): 363–394.

Hotz, V. Joseph, and Robert A. Miller. 1993. "Conditional Choice Probabilities and the Estimation of Dynamic Models." *The Review of Economic Studies* 60 (3): 497–529.

Kalouptsidi, Myrto, Paul T. Scott, and Eduardo Souza-Rodrigues. 2021. "Identification of counterfactuals in dynamic discrete choice models." *Quantitative Economics* 12 (2): 351–403.

Keane, Michael P., and Kenneth I. Wolpin. 2009. "Empirical applications of discrete choice dynamic programming models." *Review of Economic Dynamics* 12 (1): 1–22.

Magnac, Thierry, and David Thesmar. 2002. "Identifying Dynamic Discrete Decision Processes." *Econometrica* 70 (2): 801–816.

Pesendorfer, Martin, and Philipp Schmidt-Dengler. 2008. "Asymptotic Least Squares Estimators for Dynamic Games1." *The Review of Economic Studies* 75 (3): 901–928.

Rust, John. 1987. "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher." *Econometrica* 55 (5): 999–1033.

Shum, Matthew. 2017. *Econometric Models for Industrial Organization.* World Scientific.

Su, Che-Lin, and Kenneth L. Judd. 2012. "Constrained Optimization Approaches to Estimation of Structural Models." *Econometrica* 80 (5): 2213–2230.