14.385 Nonlinear Econometric Analysis

Dynamic discrete choice

Maximilian Kasy

Department of Economics, MIT

Fall 2022

Outline

- Goal:
  - From observation of choices, infer preferences,
  - which allow to predict choice probabilities in counterfactual settings.
- Static multinomial choice:
  - Cross-section of individuals who choose between multiple alternatives.
  - Choices maximize utility.
  - Choice probabilities might depend on choice characteristics and individual characteristics.
- Dynamic discrete choice:
  - Decisions affect both utility now, and future states.
  - States evolve according to a Markov transition function.
  - Choices maximize expected discounted utility (value function).

Takeaways for this part of class

- With functional form assumptions and restrictions of heterogeneity, we can infer preferences from choices.
- Functional form assumptions often have strong implications for counterfactual behavior.
- Dynamic discrete choice models assume
  - Agents have correct knowledge of state transition probabilities,
  - and solve the full dynamic optimization problem.
  - Actions are observed, utilities are not.
- Reinforcement learning (next part of class) considers the reverse problem:
  - We have to learn transition probabilities from data,
  - and want to construct an agent maximizing expected discounted utility.
  - Rewards (utilities) are observed, actions are to be chosen.

Recap: Static multinomial choice

Dynamic discrete choice

References

Multinomial Choice Models

- Choices from a set of unordered alternatives.
  Examples:
  - Commuting mode: car, train, walk, ...
  - Mobile phone network.
- For a random sample of N individuals, we observe the choice

$$Y_i \in \{0, 1, \ldots, m\}.$$

- Let $Y_{ij} = 1$ if $Y_i = j$, and $Y_{ij} = 0$ otherwise, for $i = 1, \ldots, N$ and $j = 0, \ldots, m$.
- We observe characteristics of the individual,

$$X_i \quad \text{for} \quad i = 1, \ldots, N \quad \text{(e.g., income)}.$$

- We may observe covariates with values depending on the alternative,

$$X_{ij} \quad \text{for} \quad i = 1, \ldots, N \quad \text{and} \quad j = 0, \ldots, m \quad \text{(e.g., prices)}.$$

Multinomial Logit Model

- Suppose first that the covariates only vary across individuals, X_i.
- The **multinomial logit model** is then an extension of the basic logit model.
- Let $p_{ij} = \Pr(Y_i = j | X_i)$. The multinomial logit model postulates:

$$p_{ij} = \frac{e^{X_i' \beta_j}}{\sum_{k=0}^{m} e^{X_i' \beta_k}}.$$

- Equal translation of all vectors β_k to $\beta_k + \alpha$
  leaves all choice probabilities, p_{ij}, unchanged.
- Therefore, we typically normalize $\beta_0 = 0$. This implies $p_{ij}/p_{i0} = e^{X_i' \beta_j}$.

Multinomial Logit Model

- The log-likelihood is:

$$\sum_{i=1}^{N}\sum_{j=0}^{m} Y_{ij} \ln p_{ij} = \sum_{i=1}^{N}\sum_{j=0}^{m} Y_{ij} \ln \left(\frac{e^{X_i' \beta_j}}{\sum_{j=0}^{m} e^{X_i' \beta_j}} \right).$$

- We maximize this log-likelihood with respect to $\beta_1, \ldots, \beta_m$ (remember that we set $\beta_0 = 0$).

- The implied marginal effects is

$$\frac{\partial p_{ij}}{\partial X_i} = p_{ij} \left(\beta_j - \sum_{k=0}^{m} p_{ik} \beta_k \right).$$

- $\beta_j = 0$ does not imply $\partial p_{ij} / \partial X_i = 0$!

- The average marginal effect is:

$$\frac{1}{N} \sum_{i=1}^{N} p_{ij} \left(\beta_j - \sum_{k=0}^{m} p_{ik} \beta_k \right).$$

Conditional Logit Model

- Regressors vary by alternative (and possibly by individual), X_{ij}.
- This model postulates:

$$p_{ij} = \frac{e^{X'_{ij}\beta}}{\sum_{k=0}^{m} e^{X'_{ik}\beta}}.$$

- The model cannot contain variables that do not vary by alternative (like a constant, or individual income).
- Imagine that we try to introduce such a variable, Z_i, with coefficient α. The probability of choice j for individual i is now:

$$p_{ij} = \frac{e^{X'_{ij}\beta + Z_i\alpha}}{\sum_{k=0}^{m} e^{X'_{ik}\beta + Z_i\alpha}} = \frac{e^{Z_i\alpha} e^{X'_{ij}\beta}}{e^{Z_i\alpha}\sum_{k=0}^{m} e^{X'_{ik}\beta}} = \frac{e^{X'_{ij}\beta}}{\sum_{k=0}^{m} e^{X'_{ik}\beta}}.$$

  So, α is not identified.

Conditional Logit Model

- The marginal effects for the conditional logit model are:

$$\partial p_{ij}/\partial X_{ij} = p_{ij}(1-p_{ij})\beta, \quad \partial p_{ij}/\partial X_{ik} = -p_{ij}p_{ik}\beta.$$

- McFadden (1974) showed that the conditional logit model can be derived as the solution of a utility maximizing agent with utility from choice j given by

$$U_{ij} = X_{ij}'\beta + u_{ij}$$

where u_{ij} are independent and have a type I extreme value distribution.

- By choosing X_{ij} and β appropriately, it can be shown that the multinomial logit model is a particular case of the conditional logit model. To simplify, suppose that there are three alternatives $j = 0, 1, 2$, then we can make:

$$X_{i0} = \begin{pmatrix} 0 \\ 0 \end{pmatrix} \quad X_{i1} = \begin{pmatrix} X_i \\ 0 \end{pmatrix} \quad X_{i2} = \begin{pmatrix} 0 \\ X_i \end{pmatrix} \quad \beta = \begin{pmatrix} \beta_1 \\ \beta_2 \end{pmatrix},$$

obtaining the multinomial logit model.

Independence of Irrelevant Alternatives (IIA)

- Conditional logit model $\Rightarrow$ The conditional probability of $Y_i = j$ given $Y_i = j$ or k is:

$$\frac{p_{ij}}{p_{ij} + p_{ik}} = \frac{1}{1 + e^{-(X_{ij} - X_{ik})'\beta}}.$$

- That is, this conditional probability depends only on $X_{ij} - X_{ik}$ and not on the characteristics of other alternatives.

- It does not change even if other alternatives become available.
  This is called **independence of irrelevant alternatives** (IIA).

- But differently:

$$\frac{p_{ij}}{p_{ik}} = e^{(X_{ij} - X_{ik})'\beta}.$$

The Red Bus/Blue Bus example

- Suppose commuters choose between two modes of transportation: **car** and **red bus**.
- Initially, $p_{\texttt{car}}/p_{\texttt{red bus}} = 1$; the choice probability for each mode is 1/2:
$$p_{\texttt{car}} = p_{\texttt{red bus}} = 1/2.$$
- Suppose now that a new mode of transportation is introduced: The **blue bus**.
- If consumers do not value the color of the bus (the color of the bus is not in X_{ij}) then $p_{\texttt{red bus}}/p_{\texttt{blue bus}} = 1$.
- By the IIA property, it is still true that $p_{\texttt{car}}/p_{\texttt{red bus}} = 1$.
Therefore, we obtain:
$$p_{\texttt{car}} = p_{\texttt{red bus}} = p_{\texttt{blue bus}} = 1/3.$$
- Just by introducing the blue bus, the choice probability of car went from 1/2 to 1/3. However, because blue and red buses are close substitutes, we would expect that $p_{\texttt{car}}$ stays unchanged.

IIA and the Random Utility Model

- The conditional logit model can be derived as the solution of an agent maximizing utility over choices, where utility from choice j is given by

$$U_{ij} = X_{ij}'\beta + u_{ij}$$

  and u_{ij} are independent and have a type I extreme value distribution.

- The IIA property arises from the combination of two assumptions:
  (1) independence between the utilities of different alternatives (given the covariates),
  (2) type I extreme value distribution for u_{ij}.

- There exist extensions of the conditional logit model that dispose of the IIA property by relaxing either (1) or (2). Two models that relax (1):
  - nested logit
  - random coefficient logit

Recap: Static multinomial choice

Dynamic discrete choice

References

Dynamic Discrete Choice

- Consider an agent or set of agents
- choosing actions, a_{it} from a discrete set $A = \{0, 1, \ldots, J\}$
- over an infinite horizon.
- Agents observe state variables $S_{it} = (X_{it}, \varepsilon_{it})$.
- X_{it} is observed by the agent and by the econometrician.
- ε_{it} is observed by the agent but not observed by the econometrician.

Example: Rust (1987)

- The agent is Harold Zurcher, superintendent of maintenance at the Madison (WI) Metropolitan Bus Company
- X_{it} is engine mileage for bus i at month t
- ε_{it} are other characteristics of bus i at month t, which affect Zurcher's decisions, but unobserved by the econometrician.
- $a_{it} \in \{0, 1\}$ codes Zurcher's bus engine replacement decision

Other applications:

- Retirement decisions,
- occupational choice,
- dynamic discrete games.

Bellman equation

- Agents' beliefs about future states follow a Markov transition process with transition probability function

$$P(S_{it+1}|a_{it}, S_{it}, \theta_p).$$

- The value function $V_\theta(X_{it}, \varepsilon_{it})$ is the solution to the Bellman equation

$$V_\theta(S_{it}) = \max_{a\in A}\left[U(a, S_{it}, \theta_u) + \beta \int V_\theta(S_{it+1}) dP(S_{it+1}|a, S_{it}, \theta_p)\right]$$

  where

  - U is the instantaneous utility function,
  - β is the discount factor, (typically imputed, not estimated)
  - $\theta = (\theta_p, \theta_u)$.

- The optimal decision rule solves the Bellman equation.

Functional form assumptions

Recall $S_{it} = (X_{it}, \varepsilon_{it})$.

- **Additive separability + Logit:**

$$U(a, S_{it}, \theta_u) = u(a, X_{it}, \theta_u) + \varepsilon_{it}(a),$$

  and

  - $\varepsilon_{it} = (\varepsilon_{it}(0), \varepsilon_{it}(1), \ldots, \varepsilon_{it}(J))$ is i.i.d. across i and t,
  - with mutually independent (centered) type I extreme value components.

- **Conditional independence:**

$$P(X_{it+1} | X_{it}, \varepsilon_{it}, a_{it}, \theta_p) = P(X_{it+1} | X_{it}, a_{it}, \theta_p)$$

- **Discrete support:** X_{it} has discrete and finite support $\mathcal{X}$.

Rust (1987)

For the engine replacement application in Rust (1987):

- Instantaneous utility is

$$u(a, X_{it}, \theta_u) + \varepsilon(a) = \begin{cases} -c(X_{it}, \theta_{u1}) + \varepsilon(0) & \text{if } a_{it} = 0 \\ -\theta_{u2} - c(0, \theta_{u1}) + \varepsilon(1) & \text{if } a_{it} = 1 \end{cases}$$

  where

  - $c(X_{it}, \theta_{u1})$: operating cost of a bus with X_{it} mileage (could be, e.g., polynomial), normalize $c(0, \theta_{u1}) = 0$,
  - θ_{u2}: engine replacement cost.

- Mileage (X_{it}) is discretized in 90 intervals of length 5000.
- Transition probabilities $P(X_{it+1}|X_{it}, a_{it}, \theta_p)$:
  Multinomial with three values corresponding to $[0, 5000)$, $[5000, 10000)$, $[10000, \infty)$ for mileage between t and $t+1$.

Integrated Bellman equation

- Additive separability and Logit and Conditional Independence assumptions $\Rightarrow$ integrated version of the Bellman equation with closed form:

$$\bar{V}_\theta(X_{it}) = \ln\left(\sum_{a=0}^{J} \exp\left\{u(a, X_{it}, \theta_u) + \beta \sum_{x\in\mathcal{X}} \bar{V}_\theta(x) P(X_{it+1} = x | a, X_{it}, \theta_p)\right\}\right).$$

- The conditional choice probabilities are

$$P(a_{it} = a | X_{it}, \theta) = \frac{\exp\{\bar{v}_\theta(a, X_{it})\}}{\sum_{j=0}^{J} \exp\{\bar{v}_\theta(j, X_{it})\}}$$

where

$$\bar{v}_\theta(a, X_{it}) = u(a, X_{it}, \theta_u) + \beta \sum_{x\in\mathcal{X}} \bar{V}_\theta(x) P(X_{it+1} = x | a, X_{it}, \theta_p).$$

Log-likelihood

The log-likelihood is

$$\sum_{i=1}^{N}\sum_{t=1}^{T}\ln P(a_{it}|X_{it},\theta)+\sum_{i=1}^{N}\sum_{t=2}^{T}\ln P(X_{it}|a_{it-1},X_{it-1},\theta_p)$$

- Transition probabilities are specified as primitives of the model, which makes it easy to evaluate the second term of the log-likelihood.
- Typically, θ_p is estimated separately in a first step by maximizing that term.
- Evaluating the first term of the likelihood is more difficult because it involves the integrated value function, $\bar{V}_\theta(x)$.
- Rust (1987) proposes a *nested fixed point algorithm* (NFXP):
  - **Outer loop:** Iterates over θ to maximize the likelihood
  - **Inner loop:** Inside each iteration of the outer loop, iterates the Bellman equation until convergence to find $\bar{V}_\theta(x)$ (this is facilitated by the discrete nature of x)

Discrete states

- Suppose $\mathcal{X} = \{x_1, \ldots, x_k\}$ (in Rust's paper, $k = 90$).
- Let
$$\boldsymbol{u}(a, \theta_u) = \begin{pmatrix} u(a, x_1, \theta_u) \\ \vdots \\ u(a, x_k, \theta_u) \end{pmatrix}$$
and
$$\boldsymbol{F}(a) = \begin{pmatrix} p(x_1|a, x_1, \theta_p) & \cdots & p(x_k|a, x_1, \theta_p) \\ \vdots & \ddots & \vdots \\ p(x_1|a, x_k, \theta_p) & \cdots & p(x_k|a, x_k, \theta_p)) \end{pmatrix}$$
where
$$p(x'|a, x, \theta_p) = P(X_{it+1} = x'|a_{it} = a, X_{it} = x, \theta_p).$$
- The matrix $\boldsymbol{F}(a)$ can be estimated in the first step from the second term in the log-likelihood function. So consider it as known for the rest of the argument.

Value function as fixed point

- Let
$$\bar{\boldsymbol{V}}_\theta = \begin{pmatrix} \bar{V}_\theta(x_1) \\ \vdots \\ \bar{V}_\theta(x_k) \end{pmatrix}.$$
be the vectorized version of the value function, $\bar{V}_\theta(x)$.
- Then, for any given θ_u, $\bar{\boldsymbol{V}}_\theta$ is given by the fixed point
$$\bar{\boldsymbol{V}}_\theta = \ln\left(\sum_{a=0}^{J} \exp\left\{\boldsymbol{u}(a,\theta_u) + \beta \boldsymbol{F}(a)\bar{\boldsymbol{V}}_\theta\right\}\right).$$
- And the conditional choice probabilities are given by
$$P(a_{it} = a | X_{it} = x, \theta) = \frac{\exp\left(u(a,x,\theta_u) + \beta \boldsymbol{F}(a,x)\bar{\boldsymbol{V}}_\theta\right)}{\sum_{j=0}^{J} \exp\left(u(j,x,\theta_u) + \beta \boldsymbol{F}(j,x)\bar{\boldsymbol{V}}_\theta\right)},$$
where $\boldsymbol{F}(a,x) = (p(x_1|a,x,\theta_p), \ldots, p(x_k|a,x,\theta_p))$.

Two-step GMM estimation

- A problem with the NFXP is its computational cost.
  The Bellman equation is solved iteratively in each optimization step.
- The **conditional choice probability** (CCP) algorithm of Hotz and Miller (1993) does not require solving the Bellman equation, reducing computational cost drastically.
- Hotz and Miller notice that, for given values of θ_u, it is often possible to obtain estimates $\widehat{v}(a, X_{it}, \theta_u)$ of $\bar{v}_\theta(a, X_{it})$ from preliminary estimates of θ_p and $P(a_{it} = a|X_{it})$.
- Then, θ_u can be estimated in a second step by GMM. For example, MLE maximizes

$$\sum_{i=1}^{N}\sum_{t=1}^{T} \ln\left(\frac{\exp(\widehat{v}(a_{it}, X_{it}, \theta_u))}{\sum_{j=0}^{J}\exp(\widehat{v}(j, X_{it}, \theta_u))}\right).$$

Two-step estimation example

- Utility function linear in θ_u (e.g., polynomial):

$$u(a,x,\theta_u) = z(a,x)'\theta_u + \varepsilon(a).$$

- Then, it can be shown

$$\bar{v}_\theta(a,x) = \tilde{z}(a,x,\theta)'\theta_u + \tilde{e}(a,x,\theta),$$

  where $\tilde{z}(a,x,\theta)$ and $\tilde{e}(a,x,\theta)$ depend on θ only through θ_p, which can be estimated in a first step, and $P(a_{it}|X_{it}=x)$, which can be estimated non-parametrically.

- Therefore, we can obtain estimates $\widehat{z}(a,x)$ and $\widehat{e}(a,x)$ of $\tilde{z}(a,x,\theta)$ and $\tilde{e}(a,x,\theta)$, so $\widehat{v}(a,x,\theta_u) = \widehat{z}(a,x)'\theta_u + \widehat{e}(a,x)$.

- Then, the conditional choice probabilities are approximated as

$$P(a_{it}=a|X_{it}=x,\theta_u) = \frac{\exp(\widehat{z}(a,x)'\theta_u + \widehat{e}(a,x))}{\sum_{j=0}^{J}\exp(\widehat{z}(j,x)'\theta_u + \widehat{e}(j,x))}.$$

References

Seminal contributions to the dynamic discrete choice literature are:

Rust (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," Econometrica

Hotz and Miller (1993): "Conditional Choice Probabilities and the Estimation of Dynamic Models," Review of Economic Studies

These notes draw heavily from:

Aguirregabiria and Mira (2010): "Dynamic Discrete Choice Structural Models: A Survey," Journal of Econometrics

and from the slides by **Alberto Abadie** for previous iterations of 14.385.