

Identifying Discrete Choice Models

Robert A. Miller

University of Tokyo

October 2019

Introduction

Nonstationarity and short panels

- Dynamic discrete choice models are used to explain panel data in labor economics, industrial organization and marketing.
- Interpreting predictions of policy innovations from structural models critically depends on the identifying assumptions of the model.
- This lecture focuses on distinctions between *short* and *long* panel data sets.
- Short panels are samples from nonstationary data generating processes where the time horizon of the agent extends beyond the length of the data.
- Long panels are data generated from stationary processes, or by nonstationary data generating processes that sample every event with strictly positive probability in a finite horizon model.

Introduction

Short panels and nonstationarity are common

- Short panels are common: many panel data sets do not cover the full lifetime of the sampled firm, individual, or product.
- Nonstationarities arise naturally: in the human life cycle through aging, and the general equilibrium effects of evolving demographics; in industries because of innovation and growth; and in marketing through the diffusion of new products and over the product life cycle.
- These features pose serious challenges to inference.
- Yet most applied work in this area assumes the data generating process is stationary, or impose other strong restrictions in estimation.

Introductory Example

A model of health

- The following two-period, two-choice example illustrates the main results in a simple context.
- Consider a two period model, $T = 2$, of the decision to smoke, $d_{2t} = 1$, or not, $d_{1t} = 1$
- The state variable is: healthy, $x = 1$, or sick, $x = 2$.
- Individuals begin healthy and remain so if they do not smoke in the first period.
- If an individual smokes in the first period the probability of falling sick in the second period is π .
- The disturbances are distributed Type 1 Extreme Value.
- The true value of the systematic component from not smoking is 0 when healthy and c when sick.
- That is $u_{1t}(1) = 0$ and $u_{1t}(2) = c$ for $t \in \{1, 2\}$.

Introductory Example

Lifetime utility

- In this simple example realized lifetime utility is:

$$\begin{aligned} & \sum_{t=1}^2 \sum_{j=1}^2 \beta^{t-1} d_{jt} [u_{jt}(x_t) + \epsilon_{jt}] \\ = & d_{11} [u_{11}(1) + \epsilon_{11}] + d_{21} [u_{21}(1) + \epsilon_{21}] \\ & + d_{11}\beta \{d_{12} [u_{12}(1) + \epsilon_{12}] + d_{22} [u_{22}(1) + \epsilon_{22}]\} \\ & + d_{21}1 \{x = 1\} \beta \{d_{12} [u_{12}(1) + \epsilon_{12}] + d_{22} [u_{22}(1) + \epsilon_{22}]\} \\ & + d_{21}1 \{x = 2\} \beta \{d_{12} [u_{12}(2) + \epsilon_{12}] + d_{22} [u_{22}(2) + \epsilon_{22}]\} \\ = & d_{11}\epsilon_{11} + d_{21} [u_{21}(1) + \epsilon_{21}] \\ & + d_{11}\beta \{d_{12}\epsilon_{12} + d_{22} [u_{22}(1) + \epsilon_{22}]\} \\ & + d_{21}1 \{x = 1\} \beta \{d_{12}\epsilon_{12} + d_{22} [u_{22}(1) + \epsilon_{22}]\} \\ & + d_{21}1 \{x = 2\} \beta \{d_{12} [c + \epsilon_{12}] + d_{22} [u_{22}(2) + \epsilon_{22}]\} \end{aligned}$$

Introductory Example

Lacking knowledge about payoffs from outside the data

- First suppose the econometrician does not know the true payoff from either action.
- Instead he normalizes the flow payoff in all periods to 0 for not smoking, regardless of the individual's health state.
- That is $u_{1t}^*(x) = 0$ for $x \in \{1, 2\}$ and $t \in \{1, 2\}$. Then from (8) below:

$$u_{21}^*(1) - u_{11}^*(1) = u_{21}(1) - u_{11}(1) + \beta\pi c \quad (1)$$

- Equation (1) illustrates a general property.
- Differences relative to the normalized action are not identified, in this case because c is not identified.

Introductory Example

Long versus short panels

- In a long panel data is collected on both periods.
- If the true payoffs from not smoking are known then the remaining utility parameters are identified.
- For example, applying (6) below:

$$u_{21}(1) = \ln p_{21}(1) - \ln p_{11}(1) + \beta\pi [\ln p_{12}(2) - \ln p_{12}(1) - c] \quad (2)$$

- In a short panel where there is only data on the first period, the parameters are not identified even if value of not smoking is known, as is evident from (2) which is constructed using CCPs for both periods.

Introductory Example

Counterfactual regime temporarily affecting payoffs

- Next consider a counterfactual regime that subsidizes sick people with a payment of Δ .
- Denote by $\tilde{p}_{11}(1)$ the CCP of not smoking in the first period in the counterfactual regime.
- This regime change does not affect second period choices.
- Applying (12) below and simplifying:

$$\tilde{p}_{11}(1) = \frac{p_{11}(1)}{p_{11}(1) + [1 - p_{11}(1)] \exp(\beta\Delta\pi)}$$

- This formula illustrates the basic idea that only CCPs used in the current regime are necessary to compute a counterfactual that has no effects on choices in periods beyond the end of the panel
- Consequently a short panel suffices to compute this counterfactual.

Introductory Example

Counterfactual regime temporarily affecting payoffs in a long panel

- Now consider a new regime changing the probability of falling sick, conditional on smoking, from π to $\tilde{\pi}$; this change has no effect on second period choices either.
- Forming analogous expressions to (2) and (1) for the counterfactual regime, we substitute out $u_{21}(1)$ and $u_{21}^*(1)$ to obtain the odds ratios:

$$\frac{\tilde{p}_{21}^*(1)}{\tilde{p}_{11}^*(1)} = \frac{p_{21}(1)}{p_{11}(1)} \times \left[\frac{p_{12}(1)}{p_{12}(2)} \right]^{\beta(\pi - \tilde{\pi})} = \frac{\tilde{p}_{21}(1)}{\tilde{p}_{11}(1)} \exp [\beta (\pi - \tilde{\pi}) c]$$

- The ratio of the nonsmoking probabilities for the two periods differ between the normalization and the true payoffs by the factor $\exp [\beta (\pi - \tilde{\pi}) c]$.
- Using an incorrect normalization leads to incorrect predictions of counterfactual choice probabilities induced by changes in transition probabilities.

Introductory Example

Counterfactual regime temporarily affecting payoffs in a short panel

- Now suppose the econometrician knows the true values of $u_{1t}(x)$ for each (t, x) , but data is only available on the first period smoking decisions.
- Hence the CCPs for the second period in the current regime, are not identified.
- The counterfactual CCPs in the new regime cannot be recovered even when the new regime only changes the first period transitions on the state variables because $p_{12}(x)$ is not identified.

A Class of Dynamic Discrete Choice Markov Models

Discrete time and finite choice sets

- 1 Let $T \in \{1, 2, \dots\}$ with $T \leq \infty$ denote the horizon of the optimization problem and $t \in \{1, \dots, T\}$ denote the time period.
- 2 Each period the individual chooses amongst J mutually exclusive actions.
- 3 Let $d_t \equiv (d_{1t}, \dots, d_{Jt})$ where $d_{jt} = 1$ if action $j \in \{1, \dots, J\}$ is taken at time t and $d_{jt} = 0$ if action j is not taken at t .
- 4 $x_t \in \{1, \dots, X\}$ for some finite positive integer X for each t .
- 5 $\epsilon_t \equiv (\epsilon_{1t}, \dots, \epsilon_{Jt})$ where $\epsilon_{jt} \in \mathfrak{R}$ for all (j, t) .
- 6 The joint mixed density function for the state in period $t + 1$ conditional on (x_t, ϵ_t) , denoted by $g_{t,j,x,\epsilon}(x_{t+1}, \epsilon_{t+1} | x_t, \epsilon_t)$, satisfies the conditional independence assumption:

$$g_{t,j,x,\epsilon}(x_{t+1}, \epsilon_{t+1} | x_t, \epsilon_t) = g_{t+1}(\epsilon_{t+1} | x_{t+1}) f_{jt}(x_{t+1} | x_t)$$

where $g_t(\epsilon_t | x_t)$ is a conditional density for the disturbances, and $f_{jt}(x_{t+1} | x_t)$ is a transition probability for x conditional on (j, t) .

A Class of Dynamic Discrete Choice Markov Models

Bounded additively separable preferences

- Denote the discount factor by $\beta \in (0, 1)$ and the current payoff from taking action j at t given (x_t, ϵ_t) by $u_{jt}(x_t) + \epsilon_{jt}$.
- To ensure a transversality condition is satisfied, assume $\{u_{jt}(x)\}_{t=1}^T$ is a bounded sequence for each $(j, x) \in \{1, \dots, J\} \times \{1, \dots, X\}$, and so is:

$$\left\{ \int \max \{|\epsilon_{1t}|, \dots, |\epsilon_{Jt}|\} g_t(\epsilon_t | x_t) d\epsilon_t \right\}_{t=1}^T$$

- At the beginning of each period t the agent observes the realization (x_t, ϵ_t) chooses d_t to sequentially maximize:

$$E \left\{ \sum_{\tau=t}^T \sum_{j=1}^J \beta^{\tau-t} d_{j\tau} [u_{j\tau}(x_\tau) + \epsilon_{j\tau}] \mid x_t, \epsilon_t \right\} \quad (3)$$

where the expectation is taken over future realized values x_{t+1}, \dots, x_T and $\epsilon_{t+1}, \dots, \epsilon_T$ conditional on (x_t, ϵ_t) .

Identifying the Primitives

Identifying assumptions and data generating process

- The optimization model is fully characterized by the time horizon, the utility flows, the discount factor, the transition matrix of the observed state variables, and the distribution of the unobserved variables, summarized with the notation (T, β, f, g, u) .
- The data comprise observations for a real or synthetic panel on the observed part of the state variable, x_t , and decision outcomes, d_t .
- In our analysis, let $S \leq T$ denote the last date for which data is available (for a real or synthetic cohort).
- Following most of the empirical work in this area we consider identification when (T, β, f, g) are assumed to be known.
- Thus the goal is to identify u from (x_t, d_t) when (T, β, f, g) is known.

Identifying the Primitives

Observational Equivalence

- It is common knowledge that u is only identified relative to one choice per period for each state.
- Can we say more than that?
- For each (x, t) let $l(x, t) \in \{1, \dots, J\}$ denote any arbitrarily defined normalizing action and $u_{l(x,t),t}^*(x)$ its associated flow utility.
- Assume $\left\{ u_{l(x,t),t}^*(x) \right\}_{t=1}^T$ is bounded for each $x \in \{1, \dots, X\}$.
- Let $\kappa_\tau^*(x_{\tau+1}|x_t, j)$ denote the probability distribution of $x_{\tau+1}$, given a state of x_t taking action j at t , and then repeatedly taking the normalized action from period $t+1$ through to period τ .
- Thus $\kappa_t^*(x_{t+1}|x_t, j) \equiv f_{jt}(x_{t+1}|x_t)$ and for $\tau \in \{t+1, \dots, T\}$:

$$\kappa_\tau^*(x_{\tau+1}|x_t, j) \equiv \sum_{x=1}^X f_{l(x,\tau),\tau}(x_{\tau+1}|x) \kappa_{\tau-1}^*(x|x_t, j) \quad (4)$$

Theorem

For each $R \in \{1, 2, \dots\}$, define for all (x, j, t) :

$$u_{jR}^*(x) \equiv u_{jR}(x) + u_{l(x,R),R}^*(x) - u_{l(x,R),R}(x) \quad (5)$$

and $u_{jt}^*(x) \equiv$

$$u_{jt}(x) + u_{l(x,t),t}^*(x) - u_{l(x,t),t}(x) + \lim_{R \rightarrow T} \sum_{\tau=t+1}^R \sum_{x'=1}^X \beta^{\tau-t} \left\{ \begin{array}{l} \left[u_{l(x',t),t}^*(x') - u_{l(x,\tau),\tau}(x') \right] \times \\ \left[\kappa_{\tau-1}^*(x'|x_t, l(x,t)) - \kappa_{\tau-1}^*(x'|x_t, j) \right] \end{array} \right\} \quad (6)$$

- Then (T, β, f, g, u^*) is observationally equivalent to (T, β, f, g, u) .
- Conversely suppose (T, β, f, g, u^*) is observationally equivalent to (T, β, f, g, u) . For each date and state select any action $l(x, t) \in \{1, \dots, J\}$ with payoff $u_{l(x,t),t}^*(x)$ implied by u^* . Then (5) and (6) hold for all (t, x, j) .

Corollary

Suppose $u_{jt}(x) = u_j(x)$ and let $u_j \equiv (u_j(1), \dots, u_j(X))'$. Similarly suppose $f_{jt}(x_{t+1}|x_t) = f_j(x_{t+1}|x_t)$ for all $t \in \{1, 2, \dots\}$. Denote by $l(x)$ the normalizing action for that state, with true payoff vector $u_l = (u_{l(1)}(1), \dots, u_{l(X)}(X))'$, and assume $u_l^* \equiv (u_l^*(1), \dots, u_l^*(X))'$ is bounded for each $x \in \{1, 2, \dots\}$. Then (6) reduces to:

$$u_j^* = u_j + [I - \beta F_j] [I - \beta F_l]^{-1} (u_l^* - u_l) \quad (7)$$

where $u_j^* \equiv (u_j^*(1), \dots, u_j^*(X))'$, the X dimensional identity matrix is denoted by I , and:

$$F_j \equiv \begin{bmatrix} f_j(1|1) & \dots & f_j(X|1) \\ \vdots & \ddots & \vdots \\ f_j(1|X) & \dots & f_j(X|X) \end{bmatrix}, \quad F_l \equiv \begin{bmatrix} f_{l(1)}(1|1) & \dots & f_{l(1)}(X|1) \\ \vdots & \ddots & \vdots \\ f_{l(X)}(1|X) & \dots & f_{l(X)}(X|X) \end{bmatrix}$$

Identifying the Primitives

Observational Equivalence

- A common normalization is to set $u_{1t}^*(x) = 0$ for all (t, x) .
- The theorem shows that with the important exception of the static model (when $T = 1$), we should not interpret the payoffs for the other actions as net of, or relative to, the current payoff for the first choice.
- Intuitively, this does not account for the effect of imposing an incorrect normalization on future utilities that affect the continuation value.
- Define $\kappa_\tau(x_{\tau+1}|x_t, j)$ by setting $f_{l(x, \tau), \tau}(x_{\tau+1}|x) = f_{1\tau}(x_{\tau+1}|x)$ in (4), if $T < \infty$ then (5) and (6) simplify to:

$$u_{jT}^*(x) = u_{jT}(x) - u_{1T}(x)$$

and:

$$u_{jt}^*(x) = u_{jt}(x) - u_{1t}(x) - \sum_{\tau=t+1}^T \sum_{x_\tau=1}^X \beta^{\tau-t} u_{1\tau}(x_\tau) [\kappa_{\tau-1}(x_\tau|x_t, 1) - \kappa_{\tau-1}(x_\tau|x_t, j)]$$

Identifying the Primitives

Identification off long panels

Theorem

For all j , t , and x :

$$u_{jt}(x) = u_{1t}(x) + \psi_{1t}(x) - \psi_{jt}(x) + \sum_{\tau=t+1}^T \sum_{x_{\tau}=1}^X \beta^{\tau-t} \left\{ \begin{array}{l} [u_{1\tau}(x_{\tau}) + \psi_{1t}(x_{\tau})] \times \\ [\kappa_{\tau-1}(x_{\tau}|x, 1) - \kappa_{\tau-1}(x_{\tau}|x, j)] \end{array} \right\} \quad (8)$$

In stationary models, define $\Psi_j \equiv [\psi_j(1) \dots \psi_j(X)]'$, and for all j :

$$u_j = \Psi_1 - \Psi_j - u_1 + \beta (F_1 - F_j) [I - \beta F_1]^{-1} (\Psi_1 + u_1) \quad (9)$$

- If (T, β, f, g) is known, and if a payoff, say the first, is also known for every state and time, then u is identified.

Identifying the Primitives

Proving the theorem

- Appealing to lecture 3 we specialize the mixed decision rule to taking the first action to obtain:

$$v_{jt}(x) = u_{jt}(x) + \sum_{\tau=t+1}^T \sum_{x_{\tau}=1}^X \beta^{\tau-t} [u_{1\tau}(x_{\tau}) + \psi_{1\tau}(x_{\tau})] \kappa_{\tau-1}(x_{\tau}|x, j)$$

- Subtract from the expression above the corresponding expression for $v_{1t}(x_t)$ yielding:

$$\begin{aligned} & v_{jt}(x) - u_{jt}(x) - [v_{1t}(x) - u_{1t}(x)] \\ &= \psi_{1t}(x) - \psi_{jt}(x) - u_{jt}(x) + u_{1t}(x) \\ &= \sum_{\tau=t+1}^T \sum_{x_{\tau}=1}^X \beta^{\tau-t} \left\{ \begin{array}{l} [u_{1\tau}(x_{\tau}) + \psi_{1t}(x_{\tau})] \times \\ [\kappa_{\tau-1}(x_{\tau}|x, 1) - \kappa_{\tau-1}(x_{\tau}|x, j)] \end{array} \right\} \end{aligned}$$

- The theorem follows from rearrangement.

Identifying the Primitives

Lack of identification off short panels

- Alternatively suppose the sampling period, S , falls short of the time horizon T .
- Then choices and state transitions are not observed after period S .
- Rather than express $u_{jt}(x)$ relative to the known payoff for first choice for the full horizon as in (8), we express u_{jt} relative to the known u_{0t} until period S and then use the value function at $S + 1$.

Corollary

For all j , t , and x :

$$\begin{aligned} u_{jt}(x) &= u_{1t}(x) + \psi_{1t}(x) - \psi_{jt}(x) & (10) \\ &+ \sum_{\tau=t+1}^S \sum_{x_{\tau}=1}^X \beta^{\tau-t} \left\{ \begin{array}{l} [u_{1\tau}(x_{\tau}) + \psi_{1\tau}(x_{\tau})] \times \\ [\kappa_{\tau-1}(x_{\tau}|x, 1) - \kappa_{\tau-1}(x_{\tau}|x, j)] \end{array} \right\} \\ &+ \sum_{x_{S+1}=1}^X \beta^{S-t} V_{S+1}(x_{S+1}) [\kappa(x_{S+1}|x, 1) - \kappa(x_{S+1}|x, j)] \end{aligned}$$

Identifying the Primitives

Lack of identification off short panels

- The last expression in (10) gives the underidentification result.
- Since the choice probabilities and state transition matrices are identified from the data up to S , and $u_{jt}(x_t)$ is a linear mapping of $V_{S+1}(x)$, the utility flows would be exactly identified if $V_{S+1}(x)$ was known.
- However $V_{S+1}(x)$ is endogenous and depends on CCPs that occur after the sample ends.
- In general the primitives are not identified off a short panel without imposing X further restrictions.

Identifying the Primitives

Single action finite dependence defined

- In one specialization, however, some of the primitives can be identified off short panels without resorting to further restrictions on the payoffs.
- Single action ρ -dependence holds for an action, again say the first, if for some $t < T - \rho$ and for all j :

$$\kappa_{\rho-1}(x_{t+\rho}|x_t, 1) = \kappa_{\rho-1}(x_{t+\rho}|x_t, j) \quad (11)$$

- More specialized than finite dependence (Arcidiacono and Miller 2011, 2015), single action finite dependence includes:
 - terminal choices . . . irreversible sterilization (Hotz and Miller, 1993); firm exit (Aguirregabiria and Mira, 2007; Pakes, Ostrovsky, and Berry, 2007); retirement (Gayle, Golan and Miller, 2015).
 - renewal choices . . . job turnover (Miller, 1984); replacing a bus engine (Rust, 1987).
 - multiperiod renewal . . . capital depletion (Altug and Miller, 1998).

Identifying the Primitives

Single action finite dependence and short panels

- Appealing to the corollary above it now follows that for all $t < S - \rho$:

$$u_{jt}(x_t) = u_{1t}(x) + \psi_{1t}(x) - \psi_{jt}(x) + \sum_{\tau=t+1}^{t+\rho} \sum_{x_\tau=1}^X \beta^{\tau-t} \left\{ \begin{array}{l} [u_{1\tau}(x_\tau) + \psi_{1\tau}(x_\tau)] \times \\ [\kappa_{\tau-1}(x_\tau|x, 1) - \kappa_{\tau-1}(x_\tau|x, j)] \end{array} \right\}$$

- Intuitively $\kappa_{\tau-1}(x_\tau|x_t, 1)$ and $\kappa_{\tau-1}(x_\tau|x_t, j)$, the sequence of state probabilities from following the two paths $(1, 1, 1, \dots)$ and $(j, 1, 1, \dots)$ respectively, merge after ρ periods, obliterating terms occurring after $t + \rho$.
- Thus if the payoffs for the choices that establish single action finite dependence are known, then identification of the primitives up until period $S - \rho$ are identified.

Counterfactual Policy Innovations

Comparing long with short panels

- A rationale for estimating structural models is policy invariance.
- Structural models yield robust predictions about the effects of changes in the primitives on equilibrium in different regimes.
- Long panels, but not short panels, embody the future within the past through an ergodicity assumption.
- So long, but not short, panels are amenable to predicting the future.
- Models estimated off short panels can be used to reconstruct counterfactual histories.
- Whether a panel is long or short is determined by the data generating process of the underlying model.
- Our analysis highlights a trade-off between committing specification errors by treating data as a long panel, or by accepting the limitations that accompany nonstationary short panels.

Counterfactual Policy Innovations

Limiting the scope of counterfactuals addressed with short panels

- We limit our analysis to temporary policy innovations that expire before the sample ends at S .
- Useful policy advice can be gleaned from short panels that do not sample many periods beyond the phase of interest:
 - clinical trials
 - public policy experiments
 - early child development
 - education choices
 - medical innovations curing disease
 - venture capital funding.

Counterfactual Policy Innovations

Notation for changing payoffs and transitions

- Denote by:
 - $u_{jt}(x)$ the true payoffs in the current regime
 - $\tilde{u}_{jt}(x)$ the true payoffs in the counterfactual regime
 - $\Delta_{jt}(x) \equiv \tilde{u}_{jt}(x) - u_{jt}(x)$ the payoff innovation
 - $u_{jt}^*(x)$ an observationally equivalent normalization to $u_{jt}(x)$
- Similarly denote by:
 - $f_{jt}(x'|x)$ the transition in the current regime
 - $\tilde{f}_{jt}(x'|x)$ the transition in the counterfactual regime
 - $\Lambda_{jt}(x'|x) \equiv \tilde{f}_{jt}(x'|x) - f_{jt}(x'|x)$ a transition innovation.

Counterfactual Policy Innovations

Back to fundamentals

- Recall defining $Q_t^{-1}(p, x)$ as the inverse of $Q_t(\delta, x) \equiv (Q_{1t}(\delta, x), \dots, Q_{J-1,t}(\delta, x))'$ where:

$$Q_{jt}(\delta, x) \equiv \int_{-\infty}^{\infty} G_{jt}(\epsilon_j + \delta_j - \delta_1, \dots, \epsilon_j, \dots, \epsilon_j + \delta_j | x) d\epsilon_j$$

- With this notation we identified $\psi_{it}(x) \equiv V_t(x) - v_{it}(x)$

$$\begin{aligned} &= \sum_{j=1}^J \left\{ p_{jt}(x) [v_{jt}(x) - v_{it}(x)] + \int \epsilon_{jt} d_{jt}^0(x_t, \epsilon_t) g_t(\epsilon_t | x) d\epsilon_t \right\} \\ &= \sum_{j=1}^J p_{jt}(x) \left[Q_{jt}^{-1}[p_t(x), x] - Q_{it}^{-1}[p_t(x), x] \right] + \\ &\quad \sum_{j=1}^J \int \prod_{k=1}^J 1 \left\{ \begin{array}{l} \epsilon_{kt} - \epsilon_{jt} \leq \\ Q_{jt}^{-1}[p_t(x), x] - Q_{kt}^{-1}[p_t(x), x] \end{array} \right\} \epsilon_{jt} g_t(\epsilon_t | x) d\epsilon_t \end{aligned}$$

Counterfactual Policy Innovations

An extra piece of notation

- Now let $\tilde{p}_t(x)$ is the CCP associated $\tilde{d}_t^o(x, \epsilon_t)$
- For $\epsilon \equiv (\epsilon_1, \dots, \epsilon_J)$ and $p \equiv (p_1, \dots, p_J)$ we now define:

$$Y_{it}(p, x) \equiv \sum_{j=1}^J p_j \left[Q_{jt}^{-1}(p, x) - Q_{it}^{-1}(p, x) \right] \\ + \sum_{j=1}^J \int \prod_{k=1}^J 1 \left\{ \begin{matrix} \epsilon_k - \epsilon_j \leq \\ Q_{jt}^{-1}(p, x) - Q_{kt}^{-1}(p, x) \end{matrix} \right\} \epsilon_j dG_t(\epsilon | x)$$

- Note that $\psi_{jt}(x) = Y_{jt}[\tilde{p}_t(x), x]$ for all (j, t, x) .
- Moreover writing $\tilde{V}_t(x)$ and $\tilde{v}_{jt}(x)$ respectively as the ex ante and conditional value functions associated with the counterfactual regime, it follows that:

$$\tilde{\psi}_{jt}(x) \equiv Y_{jt}[\tilde{p}_t(x), x] = \tilde{V}_t(x) - \tilde{v}_{jt}(x)$$

- From the definition of $Q_t^{-1}(p, x)$ it follows that if $\tilde{p}_t(x)$ is identified, then $Y_{jt}[\tilde{p}_t(x), x]$ and hence $\tilde{\psi}_{jt}(x)$ is identified when G is known.

Identifying the Primitives

Temporary counterfactuals involving payoffs are identified off short panels

Theorem

Supposing $\Delta_{jt}(x) = 0$ for all $t \geq S$ then $\tilde{p}_{jS}(x) = p_{jS}(x)$ and for all $t < S$:

$$\tilde{p}_{jt}(x) = \int \prod_{k=1}^J 1 \left\{ \begin{array}{l} \epsilon_{kt} - \epsilon_{jt} \leq \\ \tilde{\psi}_{jt}(x) - \tilde{\psi}_{kt}(x) \end{array} \right\} dG_t(\epsilon_t | x) \quad (12)$$

where:

$$\begin{aligned} \tilde{\psi}_{jt}(x) - \tilde{\psi}_{kt}(x) &= \psi_{jt}(x) - \psi_{kt}(x) + \Delta_{jt}(x) - \Delta_{kt}(x) \\ &\quad + \sum_{\tau=t+1}^S \sum_{x_\tau=1}^X \beta^{\tau-t} \left\{ \begin{array}{l} [\Delta_{1\tau}(x_\tau) + \tilde{\psi}_{1\tau}(x_\tau) - \psi_{1\tau}(x_\tau)] \\ \times [\kappa_{\tau-1}(x_\tau | x, j) - \kappa_{\tau-1}(x_\tau | x, k)] \end{array} \right\} \end{aligned}$$

Counterfactual Policy Innovations

Computing counterfactual state transitions

- Identifying counterfactual CCPs that result from changes in the state transitions requires more information, because in this case:

$$\tilde{p}_{jt}(x) = \int \prod_{k=1}^J \mathbf{1} \left\{ \begin{array}{l} \epsilon_{kt} - \epsilon_{jt} + u_{kt}(x) - u_{jt}(x) \\ \leq \sum_{\tau=t+1}^T \sum_{x_{\tau}=1}^X \beta^{\tau-t} \left\{ \begin{array}{l} [\tilde{\psi}_{1\tau}(x_{\tau}) + u_{1\tau}(x_{\tau})] \times \\ [\tilde{\kappa}_{\tau-1}(x_{\tau}|x, k) - \tilde{\kappa}_{\tau-1}(x_{\tau}|x_t, j)] \end{array} \right\} \end{array} \right\}$$

where $\tilde{\kappa}_{\tau-1}(x_{\tau}|x, k)$ is defined analogously to $\kappa_{\tau-1}(x_{\tau}|x, k)$:

- by replacing $f_{j,t+1}(x'|x)$ with $\tilde{f}_{j,t+1}(x'|x)$
- and then repeating the first action.

Counterfactual Policy Innovations

Identifying state transitions in long panels

- The presence of the $u_{1\tau}(x_\tau)$ terms show that they cannot be derived without knowing the true systematic payoff for one of the choices, regardless of the sample length.
- Suppose $u_{1t}(x_t)$ is known for all t .
- This extra knowledge identifies $u_{jt}(x)$ for all (j, t, x) in a long panel.
- Then we could recursively recover $\tilde{p}_{jt}(x)$ from:
 - $\tilde{p}_T(x) = p_T(x)$
 - this implies $\tilde{\psi}_{1T}(x_T) = Y_{1T}[p_T(x), x]$
 - use formula above to recover $\tilde{p}_{T-1}(x)$
 - successively substitute into $\tilde{\psi}_{1s}(x_\tau) = Y_{1s}[\tilde{p}_s(x), x]$ for $s \in \{\tau + 1, \dots, T\}$.

Counterfactual Policy Innovations

Counterfactual state transitions in short panels

- This argument extends to temporary changes in state transitions when single action ρ -dependence holds.
- In this special case $u_{jt}(x)$ is identified for all (j, x) and $t < S - \rho$.
- Since $\tilde{\kappa}_{\tau-1}(x_\tau|x, k) = \tilde{\kappa}_{\tau-1}(x_\tau|x_t, j)$, the recursive procedure described above applies.
- In general the case of short panels is more problematic.
- Knowing the true systematic payoff for one of the choices is generally not sufficient to identify the effects of even a temporary innovation.
- Note that $p_\tau(x_\tau)$ is not identified for $\tau > S$.
- Hence neither is $\psi_{1\tau}(x_\tau) = Y_{1\tau}[p_\tau(x), x]$.
- From (13) it now follows that $\tilde{p}_{jt}(x)$ is not identified for any t .
- Intuitively the continuation value at S is unknown even a primitive flow utility is known.

Conclusion

Identifying the primitives

- Flow payoffs are exactly identified off long panels from the CCPs when the researcher knows the payoffs for one of the choices, the discount factor, and the distribution of the unobservables.
- These assumptions are not sufficient to identify the remaining parameters off nonstationary short panels.
- In contrast to long panels, knowing the flow payoff for one of the actions over the course of the sample period is not generally enough to restore identification of the model primitives in short panels.
- An important exception is the special case of single action finite dependence when the payoff from that action is known.
- Loosely speaking this is because behavior observed during a short panel is not solely attributable to payoffs that occur during the panel, but partially reflects payoffs that occur after the panel ends.

Conclusion

Identifying counterfactuals

- Long panels, but not short, can be used to predict the future.
- Short panels can however be used to construct counterfactual histories.
- If none of the payoffs to any the actions are known, the effects of counterfactual temporary policy changes are identified in short panels if the policy change only affects the flow payoffs, a result that mimics the long panel analogue.
- Predictions from counterfactuals affecting state transitions can only be identified off long panels if one choice specific payoff for each state is known.
- This assumption is not sufficient for identification off short panels even if the counterfactual is temporary except in the special case of single action finite dependence when the payoff from that action is known.

Conclusion

A summary

- The following table summarizes our results on identifying u when (T, β, f, g) are known:

Identification	primitives	counterfactual payoffs	counterfactual transitions
long panel with one known payoff per state	Y	Y	Y
long panel lacking a known payoff	N	Y	N
short panel with one known payoff per state	N	Y	N
short panel lacking a known payoff	N	Y	N

Conclusion

One final speculative remark

- Finally, the case for estimating utility functions purely as a vehicle for making counterfactual predictions is not compelling unless the researcher has reason to impose restrictions on the utility functions because of knowledge outside the data.
- To compute behavior induced by changing payoffs off panels either short or long, it is not necessary to know the values of a choice specific payoff, but it is a requirement for estimating the remaining utility parameters
- To compute behavior induced by changing the transition function off long panels and short panels with the single action finite dependence property, aside from the CCPs, only data from outside the sample on the true value of a choice-specific payoff is necessary.