

Identifying Dynamic Discrete Choice Models off Short Panels*

Peter Arcidiacono

Robert A. Miller

Duke University & NBER Carnegie Mellon University

September 9, 2017

Abstract

This paper analyzes the identification of flow payoffs and counterfactual choice probabilities (CCPs) in single-agent dynamic discrete choice models. We develop new results on non-stationary models where the time horizon for the agent extends beyond the length of the data (short panels). We show that counterfactual CCPs and policy functions in short panels are identified when induced by temporary policy changes affecting payoffs, even though the utility flows are not. Counterfactual CCPs induced by innovations to state transitions are generally not identified unless the model exhibits single action finite dependence (which includes terminal or renewal states as special cases), and the payoffs of those actions establishing single action finite dependence (the renewal action for example) are known.

1 Introduction

Dynamic discrete choice models are increasingly used to explain panel data in labor economics, industrial organization and marketing.¹ It is widely recognized that interpreting the predictions of policy innovations from structural models critically depends on the assumptions used to identify the model. This paper

*Corresponding author: Robert Miller, Tepper Business School, Carnegie Mellon University, Pittsburgh, Pa, 15217; phone, 412-268-3701; email, ramiller@cmu.edu.

¹For surveys of this literature see Eckstein and Wolpin (1989), Pakes (1994), Rust (1994), Miller (1997), Aguirregabiria and Mira (2010), Keane, Todd and Wolpin (2011) and Arcidiacono and Ellickson (2011).

extends previous work on identifying dynamic discrete choice models of individual optimization problems off panel data. We focus on nonstationary data generating processes where the time horizon of the agent extends beyond the length of the data. For convenience we refer to data of this form as *short*. Short panels contrast with *long* ones: data generated from stationary processes, or panel data generated by nonstationary data generating processes that sample every event with strictly positive probability in a finite horizon model. We analyze the identification of policy functions, structural parameters, and counterfactual policies, highlighting results for short panels and contrasting their differences with long panels that have received much more attention from the literature.

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. Conventional wisdom holds that accommodating nonstationarities within dynamic structures complicates inference, explaining why most applied work in this area assumes the data generating process is stationary, or impose other strong restrictions in estimation.

Our analysis draws extensively upon previously published work: Rust's (1987) conditional independence assumption limiting the role of unobserved heterogeneity; Hotz and Millers' (1993) inversion theorem, relating conditional choice probabilities to differences in continuation values, that we show below identifies the policy function when the distribution of unobserved heterogeneity is known; the observational equivalence Rust (1994) highlighting links between payoffs occurring at different times; the identification theorem of Magnac and Thesmar (2001) for primitives in a finite horizon model; Aguirregabiria's (2005) extension to infinite horizon stationary models; and the representation of utility payoffs in Arcidiacono and Miller (2011). Our results on counterfactuals extends the prior work of Aguirregabiria (2005, 2010) and Norets and Tang (2014), who show that in long panels counterfactuals only affecting flow payoffs do not depend on the normalization selected, but that counterfactuals affecting the transitions of the state variables generally depend on which flow payoff is normalized.

Following the papers cited above, and many other besides, we assume throughout that the unobserved variables are independently distributed over time, that the distribution of the unobserved variables is

known, as is the discount factor. However the earliest work on estimating dynamic discrete choice models (Miller, 1984; Pakes, 1986; Wolpin, 1984) included unobserved heterogeneity, and the identification of the distribution of unobserved variables has been taken up in several recent studies (Kasahara and Shimotsu, 2009; Aguirregabiria, 2010; Hu and Shum, 2012; Norets and Tang, 2014). A literature on the identification and estimation of multi-agent models has also emerged (Aguirregabiria and Mira, 2007; Bajari, Benkard and Levin, 2007; Pakes, Ostrovsky and Berry, 2007; Pesendorfer and Schmidt-Dengler, 2008; Bajari, Chernozhukov, Hong, and Nekepelov, 2009; Aguirregabiria and Suzuki, 2014; Aguirregabiria and Mira, 2015). Several studies explore within specialized frameworks tradeoffs between imposing exclusion restrictions and functional forms assumptions, or adding information about continuous choices, in order to identify the discount factor, features of the disturbance distribution and counterfactual policies (Heckman and Navarro, 2007; Aguirregabiria, 2010; Blevin, 2014; Norets and Tang, 2014; Bajari, Chu, Nekipelov, and Park, 2016). Chou's (2016) recent work within a binary choice context on identifying counterfactual predictions without normalizing per period payoffs, is close to ours, because in considering the tradeoffs described above, Chou also distinguishes between short and long panels. But without further restrictions on the parameter space or information over and above the choices and state variables, that is beyond the assumptions made in this paper, the primitives of these models are underidentified.

The main differences emerging from this study between identification in long and short panels can be summarized as follows. 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. 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 decision making and payoffs that occur after the panel ends.

Predictions about the future can be made from long panels, but not in short ones. This difference highlights the assumption of studies that assume long panels embody the future within the past through an ergodicity assumption, whereas short panels formally accommodate nonstationary features of the data generating process. If researchers do not impose functional form assumptions and exclusion restrictions on future choice sets, payoffs and state transitions pertaining to parts of the population excluded from the data generating process, counterfactuals for models estimated off short panels must be restricted to

behavior that would have been observed if the counterfactual policies had been implemented during the time span of the panel.

We find that even 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. However, without making further assumptions on the payoffs, and in contrast to long panels, counterfactual choice probabilities for temporary policy changes affecting the state transitions are not generally identified off short panels, even if the flow value for one of the choices is known for the entire history.

There is, however, one important specialization, *single action finite dependence*, that partially restores results available for long panels to short panels. Single action finite dependence, defined formally in Section 3.3, arises when upon taking a particular action for a certain number of periods following an initial choice, the distribution of states no longer depends on that initial choice. This condition is stronger than *finite dependence*, but weaker than *terminating* or *renewal* actions, common assumptions in empirical applications of dynamic discrete choice. If the model exhibits single action finite dependence *and* the flow payoff for that particular action is known over the course of the sample period, the identification of (some of) the primitives (pertaining to the periods sampled in the short panel), and temporary counterfactual changes to state transitions, are restored.

Nevertheless useful policy advice can be gleaned from short panels without making additional assumptions beyond those necessary for identification in long panels, even if the single action finite dependence property does not hold. For example many panels on early lifecycle behavior do not sample many periods beyond the phase of interest (such as early child development and educational choices), and our results show that predictions about subsidy and tax policy can be inferred from such panels without making strong assumptions about payoffs that occur after that phase of the lifecycle is over.

The next section lays out the decision framework. Section 3 derives the observationally equivalent sets of primitives, and analyzes the identification of the flow payoffs. Section 4 turns to the identification of the conditional choice probabilities (CCPs) under counterfactual regimes. Section 5 provides an example illustrating the main results of our analysis, while Section 6 concludes.

2 Framework

The following notation and assumptions define the class of dynamic discrete choice Markov models we consider.

Assumption 1 Time is discrete, and the choice set is finite:

1A 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.

1B Each period the individual chooses amongst J mutually exclusive actions. 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 . Thus $d_{jt} \in \{0, 1\}$ with $\sum_{j=1}^J d_{jt} = 1$ for all $j \in \{1, \dots, J\}$ and $t \in \{1, \dots, T\}$.

1C If action j is taken at time t , the probability of x_{t+1} occurring in period $t + 1$ is denoted by $f_{jt}(x_{t+1}|x_t)$. The current period payoff for action j at time t depends on the state.

Assumption 2 Denote the realization of the state at t by (x_t, ϵ_t) . We assume ϵ_t is a J dimensional vector of disturbances with continuous support, that x_t has finite support, and following Rust (1987), that the conditional independence assumption is satisfied:

2A $x_t \in \{1, \dots, X\}$ for some finite positive integer X for each $t \in \{1, \dots, T\}$.

2B $\epsilon_t \equiv (\epsilon_{1t}, \dots, \epsilon_{Jt})$ where $\epsilon_{jt} \in \Re$ for all $j \in \{1, \dots, J\}$ and $t \in \{1, \dots, T\}$.

2C The joint mixed density function for the state in period $t + 1$ conditional on (x_t, ϵ_t) , denoted by $g_{t,x,\epsilon}(x_{t+1}, \epsilon_{t+1} | x_t, \epsilon_t)$, satisfies the conditional independence assumption:

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

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

Assumption 3 The preferences of the optimizing agent are defined over states and actions by a utility function that is both additively separable over time, and between the contemporaneous disturbance and the Markovian state variables:

3A 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, we 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:

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

3B At the beginning of each period $t \in \{1, \dots, T\}$ the agent observes the realization (x_t, ϵ_t) chooses d_t to sequentially maximize the discounted sum of payoffs:

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

where at each period t 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) .

For comparison purposes, we nest both the stationary infinite horizon model, defined by setting $T = \infty$, $u_{jt}(x_t) = u_j(x_t)$ and $f_{jt}(x_{t+1}|x_t) = f_j(x_{t+1}|x_t)$, and the finite horizon model in which $T < \infty$. Our main focus is, however, on models where stationarity is not assumed and the horizon is infinite, or at least extends beyond the length of the data set, which is the case for many panels in labor economics.

Given the assumptions above, an optimal decision rule at t exists, which we now define as $d_t^o(x_t, \epsilon_t)$, with j^{th} element $d_{jt}^o(x_t, \epsilon_t)$. The conditional choice probability (CCP) of choosing j at time t conditional on x_t is found by integrating $d_{jt}^o(x_t, \epsilon_t)$ over ϵ_t :

$$p_{jt}(x_t) \equiv \int d_{jt}^o(x_t, \epsilon_t) g_t(\epsilon_t | x_t) d\epsilon_t \quad (2)$$

We define $p_t(x_t) \equiv (p_{1t}(x_t), \dots, p_{Jt}(x_t))$ as the CCP vector. Denote the ex-ante value function in period t by:

$$V_t(x_t) \equiv E \left\{ \sum_{\tau=t}^T \sum_{j=1}^J \beta^{\tau-t} d_{j\tau}^o(x_\tau, \epsilon_\tau) (u_{j\tau}(x_\tau) + \epsilon_{j\tau}) \right\}$$

Thus $V_t(x_t)$ is the discounted sum of expected future payoffs just before ϵ_t is revealed and conditional on behaving according to the optimal decision rule. Let $v_{jt}(x_t)$ denote the choice-specific conditional value function, the flow payoff of action j without ϵ_{jt} plus the expected future utility conditional on following the optimal decision rule from period $t + 1$ onwards:

$$v_{jt}(x_t) \equiv u_{jt}(x_t) + \beta \sum_{x_{t+1}=1}^X V_{t+1}(x_{t+1}) f_{jt}(x_{t+1}|x_t) \quad (3)$$

Also define $\psi_{jt}(x) \equiv V_t(x) - v_{jt}(x)$. Since the value of committing to action j before seeing ϵ_t is $v_{jt}(x) + E[\epsilon_{jt} | x]$, the expected loss from pre-committing to j versus waiting until ϵ_t is observed and only then making an optimal choice, $V_t(x)$, is $\psi_{jt}(x) + E[\epsilon_{jt} | x]$. Denoting the indicator function by $1\{\cdot\}$, the policy function can be expressed as:

$$d_{jt}^o(x_t, \epsilon_t) = \prod_{k=1}^J 1\{\epsilon_{kt} - \epsilon_{jt} \leq v_{jt}(x) - v_{kt}(x)\} = \prod_{k=1}^J 1\{\epsilon_{kt} - \epsilon_{jt} \leq \psi_{kt}(x) - \psi_{jt}(x)\} \quad (4)$$

3 Identifying the Primitives

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,² 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.³

3.1 Observational equivalence

Proposition 1 of Hotz and Miller (1993) and Lemma 1 of Arcidiacono and Miller (2011) together imply that $\psi_{jt}(x)$ is identified off the CCPs if G is known. That is for each (x, j, t) a mapping denoted by $\Psi_{jt}(p, x)$ is identified off G with the property that $\psi_{jt}(x) = \Psi_{jt}[p_t(x), x]$. For example if $g_t(\epsilon_t | x)$ does not depend on x , then $\psi_{jt}(x)$ only depends on x through $p_t(x)$; further specializing, it is well known that if $g_t(\epsilon_t | x_t)$ is a standard Type 1 Extreme Value then $\psi_{jt}(x) = -\ln[p_{jt}(x)]$. Thus the policy functions are also identified from (4); similarly the counterfactual policy functions are identified if the counterfactual

²Often the distribution of unobserved variables is assumed to be extreme value for tractability. However, Arcidiacono and Miller (2011) showed how generalized extreme value distributions can easily be accommodated within a CCP estimation framework, and recently Chiong, Galichon, and Shum (2016) have proposed simple estimators for a broad range of error distributions.

³There are exceptions; for example Miller (1984) and Miller and Sanders (1997) estimate the discount rate β . Alternatively if the optimizing agents are firms, or individuals with exponential utility (CARA), the estimation equations depend on an (observed) market interest rate, rather than a preference parameter; see Gayle, Golan and Miller (2015) or Khorunzhina and Miller (2017) for examples of the latter case.

CCPs are identified and the counterfactual distribution of the disturbances is known.

However it seems almost common knowledge that u is only identified relative to one choice per period for each state. For example Rust (1994, Lemma 3.2 page 3127) showed that the solution to a stationary infinite horizon discrete choice problems is invariant to a broad class of utility transformations substituting between current and future payoffs. Thus u is not point identified. But to the best of our knowledge there is no formalism fully characterizing set identification given (T, β, f, g) .

The following notation is used to derive the identified set of observationally equivalent primitives in the absence of any further information about payoffs. For each (x, t) let $l(x, t) \in \{1, \dots, J\}$ denote any arbitrarily defined normalizing action and $c_t(x) \in \mathfrak{R}$ its associated benchmark flow utility, meaning $u_{l(x,t),t}^*(x) \equiv c_t(x)$. Assume $\{c_t(x)\}_{t=1}^T$ 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 . Formally:

$$\kappa_\tau^*(x_{\tau+1}|x_t, j) \equiv \begin{cases} f_{jt}(x_{t+1}|x_t) & \text{for } \tau = t \\ \sum_{x=1}^X f_{l(x,\tau),\tau}(x_{\tau+1}|x) \kappa_{\tau-1}^*(x|x_t, j) & \text{for } \tau = t+1, \dots, T \end{cases} \quad (5)$$

For a given (T, β, f, g) Theorem 1 exhibits all the observational equivalent dynamic optimization problems to (T, β, f, g, u) , which we denote by (T, β, f, g, u^*) .

Theorem 1 For each $R \in \{1, 2, \dots\}$, define for all $x \in \{1, \dots, X\}$, $j \in \{1, \dots, J\}$ and $t \in \{1, \dots, R\}$:

$$u_{jR}^*(x) \equiv u_{jR}(x) + c_R(x) - u_{l(x,R),R}(x) \quad (6)$$

$$u_{jt}^*(x) \equiv u_{jt}(x) + c_t(x) - u_{l(x,t),t}(x) + \lim_{R \rightarrow T} \left\{ \sum_{\tau=t+1}^R \sum_{x'=1}^X \beta^{\tau-t} [c_\tau(x') - u_{l(x,\tau),\tau}(x')] [\kappa_{\tau-1}^*(x'|x_t, l(x,t)) - \kappa_{\tau-1}^*(x'|x_t, j)] \right\} \quad (7)$$

The model defined by (6) and (7), denoted by (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 $u_{l(x,t),t}^*(x) \equiv c_t(x) \in \mathfrak{R}$, where $\{c_t(x)\}_{t=1}^T$ is bounded for each $x \in \{1, \dots, X\}$. Then (6) and (7) hold for all (t, x, j) .

Previous work on identification focused on two cases, when the horizon of the underlying optimization problem is finite, and when it is stationary, simplifying the statement of the theorem. When $T < \infty$ we

set $R = T$ in (6) and (7), as well as dropping the limit operator in (7). In stationary environments (7) has the following matrix representation.

Corollary 2 *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 $c(x) \equiv (c(1), \dots, c(X))'$ is bounded for each $x \in \{1, 2, \dots\}$. Then (7) reduces to:*

$$u_j^* = u_j + [I - \beta F_j][I - \beta F_l]^{-1}(c - u_l) \quad (8)$$

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

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

A common normalization is to let $l(x, \tau) = 1$ and $c_t(x) = 0$ for all (t, x) , normalizing the payoff from the first choice to zero by defining $u_{1t}^*(x) \equiv 0$, and interpreting the payoffs for other actions as net of, or relative to, the current payoff for the first choice. The theorem shows that with the important exception of the static model (when $T = 1$), this interpretation is grossly misleading, if not false. For example if $1 < T < \infty$ then (6) and (7) simplify to:

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

where $\kappa_\tau(x_{\tau+1}|x_t, j)$ is defined by setting $f_{l(x, \tau), \tau}(x_{\tau+1}|x) = f_{1\tau}(x_{\tau+1}|x)$ in (5). Note that $u_{jt}^*(x) - u_{1t}(x) \neq u_{jt}(x)$ for all $j \in \{2, \dots, J\}$ and $t \in \{1, \dots, T-1\}$, and depends on the true unknown value of the $u_{1\tau}(x_\tau)$ payoffs, rendering the interpretation of $u_{jt}^*(x)$ problematic.

3.2 Identifying the primitives off long panels

We now assume that one of the payoffs is known for every state and time. Without loss of generality and to simplify the notation we reorder the actions so that the known payoff corresponds to the first action. While the reordering of actions might be time and state dependent, this does not affect the generality

of the argument, and is notationally less burdensome than retaining the original ordering of actions. A consequence of Theorem 1 of Arcidiacono and Miller (2011) is that $v_{jt}(x)$ may be expressed as:

$$v_{jt}(x_t) = u_{jt}(x_t) + \sum_{\tau=t+1}^R \sum_{x_\tau=1}^X \beta^{\tau-t} [u_{1\tau}(x_\tau) + \psi_{1\tau}(x_\tau)] \kappa_{\tau-1}(x_\tau|x_t, j) \quad (9)$$

Manipulating (9) and using the definition of $\psi_{jt}(x)$ we obtain a set of necessary and sufficient conditions for identifying u when (T, β, f, g) is known.

Theorem 3 *For all j, t , and x :*

$$\begin{aligned} u_{jt}(x) &= u_{1t}(x) + \psi_{1t}(x) - \psi_{jt}(x) \\ &+ \sum_{\tau=t+1}^T \sum_{x_\tau=1}^X \beta^{\tau-t} [u_{1\tau}(x_\tau) + \psi_{1t}(x_\tau)] [\kappa_{\tau-1}(x_\tau|x, 1) - \kappa_{\tau-1}(x_\tau|x, j)] \end{aligned} \quad (10)$$

For stationary processes, define the $X \times 1$ column vector $\Psi_j \equiv [\psi_j(1) \dots \psi_j(X)]'$. Then all j :

$$u_j = \Psi_j - \Psi_1 - u_1 + \beta (F_1 - F_j) [I - \beta F_1]^{-1} (\Psi_1 + u_1) \quad (11)$$

Magnac and Thesmar (2002, Theorem 2 and Corollary 3 on pages 807 and 808) establish identification of the flow payoff for $T = 2$ finite; Equation (11) is almost identical to Aguirregabiria (2005, Proposition 1 on pages 395 and 396) who identifies the structural parameters up to a particular normalization in the stationary settings; Norets and Tang (2014, Lemma 1 on page 1234) characterize the binary choice stationary environment. This theorem unifies their previous results and provides a springboard for contrasting the results for short panels.

Everything on the right hand side of both (10) and (11) is known; since there are as many equations as unknowns, the system is exactly identified. These equations therefore yield asymptotically efficient estimators of the unrestricted utility flows. They are defined by substituting sample analogues for the conditional choice probabilities into the mappings that represent the utility flows; they are efficient because the mapping of the conditional choice probabilities on to the current utility flows is the one to one, and the relative frequencies observed in the data are the maximum estimates of the conditional choice probabilities. Asymptotic precision can only be increased by exploiting information outside the data set about true restrictions on the utility flows; false restrictions, such as adopting convenient functional forms for the payoffs, typically create misspecifications.

3.3 Single action finite dependence and short panels

Next we consider cases where the sampling period, S , falls short of the time horizon T . In one specialization some of the primitives can be identified off short panels without resorting to further restrictions on the payoffs. This happens when the probability transitions $f_{jt}(x_{t+1}|x_t)$ exhibit a special form of finite dependence, called single action finite dependence, and if the current payoff associated with that particular action is known. Formally, single action ρ -dependence holds for action one (the choice for which payoffs are observed) if for some $t < T - \rho$ and all j :

$$\kappa_{\rho-1}(x_{t+\rho}|x_t, 1) = \kappa_{\rho-1}(x_{t+\rho}|x_t, j) \quad (12)$$

More specialized than finite dependence (Arcidiacono and Miller 2011, 2015), single action finite dependence nevertheless encompasses many applications; it includes terminal choices that end the optimization problem or prevent any future decisionmaking; irreversible sterilization against future fertility, (Hotz and Miller, 1993), firm exit from an industry (Aguirregabiria and Mira, 2007; Pakes, Ostrovsky, and Berry, 2007) and retirement (Gayle, Golan and Miller, 2015) are examples. Single action finite dependence also includes renewal choices that resets the state next period to a value which is (deterministically or stochastically) independent of the current state. Turnover and job matching (Miller, 1984), or replacing a bus engine (Rust, 1987), are illustrative of renewal actions. Multiperiod renewal, such as Altug and Miller (1998), where repeatedly taking an action for a finite number of periods obliterates the effects of all previous actions, is yet another example of single action finite dependence.⁴ Note however, the primitives in these applications are only identified if the payoff to the single action is actually known, not just normalized to some notationally convenient arbitrary value.

Appealing to Corollary 4 it now follows that for all $t < S - \rho$:

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

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 in (10) that occur

⁴Similarly in Gayle and Miller (2006) and Khorunzhina and Miller (2017) single action finite dependence applies because older offspring do not directly affect the current birth choices of their mother.

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. Formally if (12) holds and $u_{1t}(x_t)$ is known for all $t \leq S$, then $u_{jt}(x_t)$ is identified for all $t \leq S - \rho$.

3.4 Lack of identification off nonstationary short panels

Since choices and state transitions are not observed after period S , the corresponding conditional choice probabilities and state transition matrices are not identified beyond that period either. Rather than express $u_{jt}(x)$ relative to the known payoff for first choice for the full horizon as in (10), we express u_{jt} relative to the known u_{0t} until period S and then use the value function at $S + 1$. This yields an expression for $u_{jt}(x_t)$ that provides the basis for the following corollary, which illuminates the extent of underidentification.

Corollary 4 *For all j, t , and x :*

$$\begin{aligned}
 u_{jt}(x) = & u_{1t}(x) + \psi_{1t}(x) - \psi_{jt}(x) + \sum_{\tau=t+1}^S \sum_{x_\tau=1}^X \beta^{\tau-t} [u_{1\tau}(x_\tau) + \psi_{1t}(x_\tau)] [\kappa_{\tau-1}(x_\tau|x, 1) - \kappa_{\tau-1}(x_\tau|x, j)] \\
 & + \beta^{S-t} V_{S+1}(x_{S+1}) [\kappa(x_{S+1}|x, 1) - \kappa(x_{S+1}|x, j)]
 \end{aligned} \tag{13}$$

The last expression in (13) 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.

4 Identifying the Effects of Policy Innovations

An important rationale for estimating structural models is their policy invariance; they yield robust predictions about the effects of changes in the primitives on equilibrium in different regimes. Aguirregabiria (2005, 2010) and Norets and Tang (2014) established two key results for stationary environments where there are no aggregate shocks: the CCPs for counterfactual regimes involving only payoff innovations are identified from the data generating process for the current regime, but to predict the effects of changing

a state transition it is also necessary to identify the primitives, not just the observationally equivalent set of primitives. This section investigates the extension of their results to short panels.

Denote the true payoffs in the sampled regime by $u_{jt}(x)$, the true payoffs in the counterfactual regime by $\tilde{u}_{jt}(x)$, and a payoff innovation by $\Delta_{jt}(x) \equiv \tilde{u}_{jt}(x) - u_{jt}(x)$. Let $u_{jt}^*(x)$ denote any normalization that is observationally equivalent to $u_{jt}(x)$ in the current regime, and $\tilde{u}_{jt}^*(x)$ any normalization that is observationally equivalent to $u_{jt}^*(x)$ in the counterfactual regime. Similarly, transition innovations are denoted by $\Lambda_{jt}(x'|x) \equiv \tilde{f}_{jt}(x'|x) - f_{jt}(x'|x)$, where $f_{jt}(x'|x)$ is the observed transition for the sampled regime. Since $f_{jt}(x'|x)$ and $\tilde{f}_{jt}(x'|x)$ are both probability transitions, $-f_{jt}(x'|x) \leq \Lambda_{jt}(x'|x) \leq 1$ for all (j, t, x) and $\sum_{x'} \Lambda_{jt}(x'|x) = 0$ for all (t, x) . Finally let the vector functions $\tilde{d}_t^o(x, \epsilon_t)$ and $\tilde{p}_t(x)$ respectively denote the optimal decision rule and the CCP associated with (x, t) and define:

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

where $\tilde{V}_t(x)$ denotes the ex ante value function associated with the counterfactual regime, and $\tilde{v}_{jt}(x)$ the conditional value function for the j^{th} action.

We limit our analysis to temporary policy innovations that expire before the sample ends at S . Even if $V_S(x)$ is identified from restrictions placed on the functional form of $u_t(x)$, and even if the policy changes are perfectly foreseen right until the end of the horizon at T , this is not sufficient to recover $p_t(x)$ and hence $u_t(x)$ for $t > S$. Since the primitives for the agent are not identified off short panels, the solution to the counterfactual regime cannot be computed and therefore $\tilde{p}_t(x)$ is not identified for any t (either before or after S). This contrasts with stationary environments, where forecasting the future is resolved by fiat, since current utilities estimated in periods before T are identical to those which nobody has observed; whether stationarity is a reasonable assumption depends on the application.

4.1 Counterfactual policies that affect payoffs

The starting point for investigating temporary payoff innovations is to note that by construction $\tilde{p}_{S+1}(x) = p_{S+1}(x)$. Solving the backwards recursion optimization problem we thus obtain the CCPs for the counterfactual regime. Theorem 5 shows that if the CCPs for the current regime cover the periods during which payoffs could have been changed, then the counterfactual CCPs can be computed.⁵

⁵A similar result also applies to temporary changes in G , the proof following the same logic.

Theorem 5 *Given any temporary payoff innovation in which $\Delta_{jt}(x) = 0$ for all $t > S$ then:*

$$\tilde{p}_{jS}(x) = \int \prod_{k=1}^J 1 \{ \epsilon_{kS} - \epsilon_{jS} + \Delta_{kS}(x) - \Delta_{jS}(x) \leq \psi_{kS}(x) - \psi_{jS}(x) \} g(\epsilon_S | x) d\epsilon_S$$

For all $t < S$ the CCPs for the counterfactual regime can be recursively expressed as:

$$\tilde{p}_{jt}(x) = \int \prod_{k=1}^J 1 \{ \epsilon_{kt} - \epsilon_{jt} \leq \tilde{\psi}_{jt}(x) - \tilde{\psi}_{kt}(x) \} g(\epsilon_t | x) d\epsilon_t$$

where $\tilde{\psi}_{jt}(x) \equiv \Psi_{jt}[\tilde{p}_t(x), x]$ and:

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

4.2 Counterfactual policies that affect state transitions

Identifying counterfactual CCPs that result from changes in the state transitions requires more information. From (4) and (9):

$$\tilde{p}_{jt}(x) = \int \prod_{k=1}^J 1 \left\{ \begin{aligned} &\epsilon_{kt} - \epsilon_{jt} + u_{kt}(x) - u_{jt}(x) \\ &\leq \sum_{\tau=t+1}^T \sum_{x_\tau=1}^X \beta^{\tau-t} \left[\tilde{\psi}_{1\tau}(x_\tau) + u_{1\tau}(x_\tau) \right] [\tilde{\kappa}_{\tau-1}(x_\tau | x, k) - \tilde{\kappa}_{\tau-1}(x_\tau | x, j)] \end{aligned} \right\} g(\epsilon_t | x) d\epsilon_t \quad (14)$$

where $\tilde{\kappa}_{\tau-1}(x_\tau | x, k)$ and $\tilde{\kappa}_{\tau-1}(x_\tau | x, j)$ are defined analogously to $\kappa_{\tau-1}(x_\tau | x, k)$ and $\kappa_{\tau-1}(x_\tau | x, j)$ by replacing $f_{j,t+1}(x' | x)$ with $\tilde{f}_{j,t+1}(x' | x)$ as appropriate (and then repeating the first action). 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. Supposing $u_{1\tau}(x_\tau)$ is known for all (τ, x) , the previous section implies $u_{jt}(x)$ is identified for all (j, t, x) . Consequently $\tilde{p}_{jt}(x)$ can be recursively recovered: $\tilde{p}_T(x) = p_T(x)$ which implies $\tilde{\psi}_{1T}(x_T) = \Psi_{1T}[p_T(x), x]$; successive substitutions into $\tilde{\psi}_{1s}(x_\tau) = \Psi_{1s}[\tilde{p}_s(x), x]$ for $s \in \{\tau + 1, \dots, T\}$ solve for $\tilde{p}_{j\tau}(x)$ in (14), thus establishing identification.

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 $\tilde{\kappa}_{\tau-1}(x_\tau | x, k) = \tilde{\kappa}_{\tau-1}(x_\tau | x, j)$, the recursive procedure described above applies.

In general the case of short panels is more dire, because knowing the true systematic payoff for one of the choices is generally not sufficient to identify the effects of even a temporary innovation. Since $p_\tau(x_\tau)$

is not identified for $\tau > S$, neither is $\psi_{1\tau}(x_\tau) = \Psi_{1\tau}[p_\tau(x), x]$. From (14) it now follows that $\tilde{p}_{jt}(x)$ is not identified for any t . Intuitively as a function of the states, the continuation values at S are unknown, and the counterfactual regime redistributes the weights of reaching the various states at S ; since the earlier choices are partly determined by the unknown continuation values, it is not possible to solve for the CCPs.

5 Example

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, $d_{1t} = 1$, where the relevant state variable is whether the individual is healthy, $x = 1$, or sick, $x = 2$. All individuals begin healthy and remain so if they do not smoke in the first period, but should the individual smoke in the first period the probability of falling sick in the second period is π . The disturbances are distributed Type 1 Extreme Value, implying $\psi_{jt}(x) = -\ln[p_{jt}(x)]$. 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 $t \in \{1, 2\}$.

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 (7) in Theorem 2:

$$u_{21}(1) = \ln p_{21}(1) - \ln p_{11}(1) + \beta\pi [\ln p_{12}(2) - \ln p_{12}(1) - c] \quad (15)$$

Suppose, however, the econometrician does not know the true payoff from either action, and 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 (10) in Theorem 1 and (15):

$$u_{21}^*(1) - u_{11}^*(1) = u_{21}(1) - u_{11}(1) + \beta\pi c \quad (16)$$

Equation (16) illustrates a general property: differences relative to the normalized action are not identified, in this case because c is not identified. 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 (15) which is constructed using CCPs for both periods.

Next consider a counterfactual regime that subsidizes sick people with a payment of Δ , a regime change

that does not affect second period choices. Applying Theorem 4 and simplifying:

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

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, so in this case a short panel suffices to compute the counterfactual.

Finally, 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 (15) and (16) 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 $\exp[\beta(\pi - \tilde{\pi})c]$. Therefore using a normalization that does not correspond to information about the true value of a payoff leads to incorrect predictions of counterfactual choice probability that are induced by changes in transition probabilities. Lastly, 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. It is not possible to recover any of the counterfactual CCPs in the new regime even when the new regime only changes the first period transitions on the state variables because $p_{12}(x)$, the CCPs for the second period in the current regime, are not identified.

6 Conclusion

Previous work shows current flow payoffs are exactly identified off long panels from the conditional choice probabilities when the payoffs for one of the choices is known, along with the discount factor, and the distribution of the unobservables. This paper shows these assumptions are not sufficient to identify the remaining parameters off nonstationary short panels. In contrast to nonstationary short panels, inference from long panels can be used to predict future events and permanent policy shifts. Such inference is possible because in long panels, but not short, choice probabilities from the past fully capture anything that might happen in the future. This is the reason for focusing on temporary shifts when working

with short panels. Although the primitives are not identified off nonstationary short panels, knowing the discount factor, the distribution of the unobserved variables and the choice probabilities suffices to identify the behavioral effects of temporary changes in flow payoffs. Knowing the payoffs for one of the choices suffices to identify the temporary shifts in the transition function off long panels, but not off nonstationary short panels, except in special cases discussed in this paper. 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.

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 parameter; 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.

Acknowledgement *We thank Shakeeb Khan, Jean-Marc Robin, the editor, four referees and seminar participants at Duke, Sciences Po, Toulouse, and Toronto for their helpful comments. We acknowledge support from National Science Foundation Grant Awards SES0721059 and SES0721098.*

7 Appendix

Proof of Theorem 1.

(i) First we show that if (6) and (7) hold, observational equivalence follows. As a starting point consider the finite horizon case in which (6) and (7) reduce to:

$$\begin{aligned}
 u_{jt}^*(x_t) &= u_{jt}(x_t) + c_t(x_t) - u_{l(x,t)t}(x_t) \\
 &+ \sum_{\tau=t+1}^T \sum_{x=1}^X \beta^{\tau-t} [c_\tau(x) - u_{l(x,\tau),\tau}(x)] [\kappa_{\tau-1}^*(x|x_t, l(x_t, t)) - \kappa_{\tau-1}^*(x|x_t, j)]
 \end{aligned} \tag{17}$$

where we set $t = T$ and drop all the terms involving τ for the last period and the static case $T = 1$. Given

the representation of $v_{jt}(x_t)$ provided by (9), it is optimal to set $d_{jt} = 1$ if:

$$j = \arg \max_{k \in \{1, \dots, J\}} \left\{ u_{kt}(x_t) + \epsilon_{kt} + \sum_{\tau=t+1}^T \sum_{x=1}^X \beta^{\tau-t} \left[u_{l(x,\tau),\tau}(x) + \psi_{l(x,\tau),\tau}(x) \right] \kappa_{\tau-1}^*(x|x_t, k) \right\}$$

Subtracting the constant:

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

does not change the optimal choice, so $d_{jt} = 1$ is optimal if $j \in \{1, \dots, J\}$ maximizes

$$\begin{aligned} & u_{kt}(x_t) - u_{l(x,t)t}(x_t) + \epsilon_{kt} \\ & + \sum_{\tau=t+1}^T \sum_{x=1}^X \beta^{\tau-t} \left\{ u_{l(x,\tau),\tau}(x) \left[\kappa_{\tau-1}^*(x|x_t, k) - \kappa_{\tau-1}^*(x|x_t, l(x_t, t)) \right] + \psi_{l(x,\tau),\tau}(x) \kappa_{\tau-1}^*(x|x_t, k) \right\} \end{aligned} \quad (18)$$

over $k \in \{1, \dots, J\}$. From (17):

$$\begin{aligned} & u_{jt}(x_t) - u_{l(x,t)t}(x_t) - \sum_{\tau=t+1}^T \sum_{x=1}^X \beta^{\tau-t} u_{l(x,\tau),\tau}(x) \left[\kappa_{\tau-1}^*(x|x_t, l(x_t, t)) - \kappa_{\tau-1}^*(x|x_t, j) \right] \\ & = u_{j_t}^*(x_t) - c_t(x_t) - \sum_{\tau=t+1}^T \sum_{x=1}^X \beta^{\tau-t} c_\tau(x) \left[\kappa_{\tau-1}^*(x|x_t, l(x_t, t)) - \kappa_{\tau-1}^*(x|x_t, j) \right] \end{aligned}$$

Substitute the second line into the maximand of (18). Then $d_{jt} = 1$ is optimal if:

$$\begin{aligned} j & = \arg \max_{k \in \{1, \dots, J\}} \left\{ \begin{aligned} & u_{kt}^*(x_t) - c_t(x_t) + \epsilon_{kt} \\ & + \sum_{\tau=t+1}^T \sum_{x=1}^X \beta^{\tau-t} \left\{ c_\tau(x) \left[\kappa_{\tau-1}^*(x|x_t, k) - \kappa_{\tau-1}^*(x|x_t, l(x_t, t)) \right] + \psi_{l(x,\tau),\tau}(x) \kappa_{\tau-1}^*(x|x_t, k) \right\} \end{aligned} \right\} \\ & = \arg \max_{k \in \{1, \dots, J\}} \left\{ u_{kt}^*(x) + \epsilon_{kt} + \sum_{\tau=t+1}^T \sum_{x=1}^X \beta^{\tau-t} \left[c_\tau(x) + \psi_{l(x,\tau),\tau}(x) \right] \kappa_{\tau-1}^*(x|x_t, k) \right\} \end{aligned}$$

where the second line follows because the dropped terms do not depend on the choice. Therefore the optimal choices are unaffected so the CCPs are the same, proving observational equivalence for $T < \infty$.

(ii) A prerequisite for proving the infinite horizon extension is to show that a finite limit in (7) exists.

Since $\{u_{l(x,\tau),\tau}(x)\}_{\tau=1}^\infty$ and $\{c_\tau(x_\tau)\}_{\tau=1}^\infty$ are bounded sequences, by M say, for all $R \leq \infty$:

$$\begin{aligned} & \left| \sum_{\tau=t+1}^R \sum_{x_\tau=1}^X \beta^{\tau-t} \left[c_\tau(x_\tau) - u_{l(x,\tau),\tau}(x_\tau) \right] \left[\kappa_{\tau-1}^*(x_\tau|x_t, l(x, t)) - \kappa_{\tau-1}^*(x_\tau|x_t, j) \right] \right| \\ & \leq \sum_{\tau=t+1}^R \sum_{x_\tau=1}^X \beta^{\tau-t} \left[|c_\tau(x_\tau)| + |u_{l(x,\tau),\tau}(x_\tau)| \right] \left[\left| \kappa_{\tau-1}^*(x_\tau|x_t, l(x, t)) \right| + \left| \kappa_{\tau-1}^*(x_\tau|x_t, j) \right| \right] \\ & \leq 4M / (1 - \beta) \end{aligned} \quad (19)$$

Referring to (7), this proves $\{u_{j_t}^*(x)\}_{\tau=1}^\infty$ is bounded for each $j \in \{1, \dots, J\}$ and $x \in \{1, \dots, X\}$.

We now notate the horizon of the problem in the value functions, conditional value functions, and the CCPs by $V_{t,T}(x)$, $v_{jt,T}(x)$ and $p_{jt,T}(x)$ respectively for the (T, β, f, g, u) framework; similarly we write $V_{t,T}^*(x)$, $v_{jt,T}^*(x)$ and $p_{jt,T}^*(x)$ for those functions when the framework is (T, β, f, g, u^*) . The proof for finite horizon problems above shows $p_{jtT}(x) = p_{jtT}^*(x)$ for all $t < T$, and hence $\psi_{jt,T}(x) = \psi_{jt,T}^*(x)$. The inequalities in (19) imply $\lim V_{tT}^*(x) = V_{t,\infty}^*(x)$ and $\lim v_{jtT}^*(x) = v_{jt,\infty}^*(x)$, with both limits finite. Hence $\lim \psi_{jt,T}^*(x) = V_{t,\infty}^*(x) - v_{jt,\infty}^*(x) \equiv \psi_{jt,\infty}^*(x)$ exists and is finite too. Since $\psi_{jt,T}^*(x) = \psi_{jt,T}(x)$ for all $T < \infty$, and $\lim \psi_{jt,T}(x) = \psi_{jt,\infty}(x)$, it now follows that $\psi_{jt,\infty}^*(x) = \psi_{jt,\infty}(x)$. Appealing to (??), observational equivalence follows because:

$$p_{jt,\infty}(x) = \int \prod_{k=1}^J 1 \{ \epsilon_{kt} - \epsilon_{jt} \leq \psi_{kt,\infty}^*(x) - \psi_{jt,\infty}^*(x) \} g_t(\epsilon_t | x) = p_{jt,\infty}^*(x).$$

(iii) To prove the converse, first note that since u^* and u are observationally equivalent then they generate the same set of CCPs, implying from the identification of $\psi_{jt}(x)$ that:

$$\psi_{jt}(x) = \Psi_{jt}[p(x), x] = \Psi_{jt}[p^*(x), x] = \psi_{jt}^*(x) \quad (20)$$

and hence from (9) that:

$$\begin{aligned} u_{jt}^*(x) - u_{jt}(x) &= v_{jt}^*(x) - v_{jt}(x) \\ &+ \sum_{\tau=t+1}^T \sum_{x_\tau=1}^X \beta^{\tau-t} [u_{1\tau}^*(x_\tau) - u_{1\tau}(x_\tau)] \kappa_{\tau-1}(x(ii_\tau|x, j)) \end{aligned} \quad (21)$$

Once again with reference to (20) and (9):

$$\begin{aligned} v_{jt}^*(x) - v_{jt}(x) &= v_{l(x,t),t}^*(x) - v_{l(x,t),t}(x) \\ &= c_t(x) - u_{l(x,t),t}(x) + \sum_{\tau=t+1}^T \sum_{x_\tau=1}^X \beta^{\tau-t} [u_{1\tau}^*(x_\tau) - u_{1\tau}(x_\tau)] \kappa_{\tau-1}(x_\tau | x_t, l(x, t)) \end{aligned}$$

Substituting the expression for $v_{jt}^*(x) - v_{jt}(x)$ obtained above into (21) proves the converse. ■

Proof of Corollary 2.

Using the matrix notation defined in the theorem, express u_j as:

$$u_j^* = u_j + c - u_l + \sum_{\tau=1}^{\infty} \beta^\tau (F_l - F_j) F_l^{\tau-1} (c - u_l) = u_j + c - u_l + \beta (F_l - F_j) \sum_{\tau=0}^{\infty} \beta^\tau F_l^\tau (c - u_l) \quad (22)$$

Since $\beta f_j(x'|x) > 0$ for all (j, x, x') and $\beta \sum_{z'=1}^Z f_j(x'|x) = \beta < 1$ for all (j, x) , the existence of $[I - \beta F_l]^{-1}$ now follows from Hadley (page 118, 1961), where:

$$Q_l \equiv \sum_{\tau=0}^{\infty} \beta^\tau F_l^\tau = I + \beta Q_l F_l = [I - \beta F_l]^{-1}$$

Substituting for the expressions for Q_l in (22):

$$\begin{aligned} u_j^* &= u_j + c - u_l + \beta (F_l - F_j) [I - \beta F_l]^{-1} (c - u_l) \\ &= u_j + [I - \beta F_l + \beta (F_l - F_j)] [I - \beta F_l]^{-1} (c - u_l) \end{aligned}$$

yielding (8). ■

Proof of Theorem 3. From (9) and (4):

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

Solving for $u_{jt}(x_t)$ completes the first part of the theorem. For the stationary case, we use the matrix notation defined in Theorems 1 and 2 to express u_j as:

$$\begin{aligned} u_j &= u_1 + \Psi_1 - \Psi_j + \sum_{\tau=1}^{\infty} \beta^\tau (F_1 - F_j) F_1^{\tau-1} (u_1 + \Psi_1) \\ &= u_1 + \Psi_1 - \Psi_j + \beta (F_1 - F_j) \left(\sum_{\tau=0}^{\infty} \beta^\tau F_1^\tau \right) (u_1 + \Psi_1) \end{aligned}$$

Then following arguments used in the proof of Theorem 1 we substitute $[I - \beta F_1]^{-1}$ for $\sum_{\tau=0}^{\infty} \beta^\tau F_1^\tau$ in the equation above to obtain (11). ■

Proof of Theorem 5. In the counterfactual regime, dynamic optimization requires the agent to choose the action that maximizes $\epsilon_{jt} + \tilde{v}_{jt}(x)$ over $j \in \{1, \dots, J\}$ which implies:

$$\tilde{p}_{jt}(x) = \int \prod_{k=1}^J 1 \left\{ \epsilon_{kt} - \epsilon_{jt} \leq \tilde{\psi}_{jt}(x) - \tilde{\psi}_{kt}(x) \right\} g(\epsilon_t | x) d\epsilon_t \quad (23)$$

But:

$$\begin{aligned} \tilde{v}_{jS}(x) - \tilde{v}_{kS}(x) &= u_{jS}(x) - u_{kS}(x) + \Delta_{jS}(x) - \Delta_{kS}(x) + \sum_{x'=1}^{X-1} \beta V_{S+1}(x') [f_{jS}(x'|x) - f_{kS}(x'|x)] \\ &= \Delta_{jS}(x) - \Delta_{kS}(x) + v_{jS}(x) - v_{kS}(x) \\ &= \Delta_{jS}(x) - \Delta_{kS}(x) + \psi_{kS}(x) - \psi_{jS}(x) \end{aligned} \quad (24)$$

Substituting (24) into (23) yields:

$$\tilde{p}_{jS}(x) = \int \prod_{k=1}^J 1 \left\{ \epsilon_{kS} - \epsilon_{jS} \leq \Delta_{jS}(x) - \Delta_{kS}(x) + \psi_{kS}(x) - \psi_{jS}(x) \right\} g(\epsilon_S) d\epsilon_S$$

But:

$$\begin{aligned}
\tilde{\psi}_{jt}(x) - \tilde{\psi}_{kt}(x) &= u_{jt}(x) - u_{kt}(x) + \Delta_{jt}(x) - \Delta_{kt}(x) + \sum_{x'=1}^{X-1} \beta \tilde{V}_{t+1}(x') [f_{jt}(x'|x) - f_{kt}(x'|x)] \\
&= u_{jt}(x) - u_{kt}(x) + \Delta_{jt}(x) - \Delta_{kt}(x) \\
&\quad + \sum_{x'=1}^{X-1} \beta V_{t+1}(x') [f_{jt}(x'|x) - f_{kt}(x'|x)] + \sum_{x'=1}^{X-1} \beta [\tilde{V}_{t+1}(x') - V_{t+1}(x')] [f_{jt}(x'|x) - f_{kt}(x'|x)] \\
&= \psi_{jt}(x) - \psi_{kt}(x) + \Delta_{jt}(x) - \Delta_{kt}(x) + \sum_{x'=1}^{X-1} \beta [\tilde{V}_{t+1}(x') - V_{t+1}(x')] [f_{jt}(x'|x) - f_{kt}(x'|x)]
\end{aligned}$$

Now we exploit the fact from (9) that for all t :

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

with an analogous expression for $\tilde{V}_t(x)$ which implies:

$$\begin{aligned}
\tilde{V}_t(x) - V_t(x) &= \Delta_{1t}(x_t) + \tilde{\psi}_{1t}(x_t) - \psi_{1t}(x_t) \\
&\quad + \sum_{\tau=t+1}^T \sum_{x_\tau=1}^X \beta^{\tau-t} [\Delta_{1\tau}(x_\tau) + \tilde{\psi}_{1\tau}(x_\tau) - \psi_{1\tau}(x_\tau)] \kappa_{\tau-1}(x_\tau|x_t, 1)
\end{aligned}$$

Therefore:

$$\begin{aligned}
\tilde{\psi}_{jt}(x) - \tilde{\psi}_{kt}(x) &= \psi_{jt}(x) - \psi_{kt}(x) + \Delta_{jt}(x) - \Delta_{kt}(x) + \sum_{x'=1}^{X-1} \beta [\tilde{V}_{t+1}(x') - V_{t+1}(x')] [f_{jt}(x'|x) - f_{kt}(x'|x)] \\
&= \psi_{jt}(x) - \psi_{kt}(x) + \Delta_{jt}(x) - \Delta_{kt}(x) \\
&\quad + \sum_{\tau=t+1}^T \sum_{x_\tau=1}^X \beta^{\tau-t} [\Delta_{1\tau}(x_\tau) + \tilde{\psi}_{1\tau}(x_\tau) - \psi_{1\tau}(x_\tau)] [\kappa_{\tau-1}(x_\tau|x, j) - \kappa_{\tau-1}(x_\tau|x, k)]
\end{aligned}$$

as required. ■

References

- [1] Aguirregabiria, V. (2005): "Nonparametric Identification of Behavioral Responses to Counterfactual Policy Interventions in Dynamic Discrete Choice Processes," *Economic Letters* 87, 393-398.
- [2] Aguirregabiria, V. (2010): "Another Look at the Identification of Dynamic Discrete Decision Processes, with an Application to Retirement Behavior", *Journal of Business and Economic Statistics* 28, 201-218.
- [3] Aguirregabiria, V., and P. Mira (2010): "Sequential Estimation of Dynamic Discrete Games," *Econometrica* 75, 1-53.

- [4] Aguirregabiria, V., and P. Mira (2010): "Dynamic Discrete Choice Structural Models: A Survey," *Journal of Econometrics* 156, 38-67.
- [5] Aguirregabiria, V., and P. Mira (2015): "Identification of Games of Incomplete Information with Multiple Equilibria and Unobserved Heterogeneity", working paper.
- [6] Aguirregabiria, V. and J. Suzuki (2014): "Identification and Counterfactuals in Dynamic Models of Market Entry and Exit", *Quantitative Marketing and Economics* 12, 267-304.
- [7] Altug, S., and R. Miller (1998): "The Effect of Work Experience on Female Wages and Labour Supply," *Review of Economic Studies* 62, 45-85.
- [8] Arcidiacono, P., and P. Ellickson (2011): "Practical Methods for Estimation of Dynamic Discrete Choice Models", *Annual Review of Economics* 3, 363-394.
- [9] Arcidiacono, P., and R. Miller (2011): "Conditional Choice Probability Estimation of Dynamic Discrete Choice Model with Unobserved Heterogeneity", *Econometrica* 79, 1823-1867.
- [10] Arcidiacono, P., and R. Miller (2015): "Nonstationary Dynamic Models with Finite Dependence", working paper.
- [11] Bajari, P., L. Benkard and J. Levin (2007): "Estimating Dynamic Models of Imperfect Competition", *Econometrica*, 75, 1331-1370.
- [12] Bajari, P., V. Chernozhukov, H. Hong, and D. Nekipelov (2009): "Nonparametric and Semiparametric Analysis of a Dynamic Discrete Game", working paper.
- [13] Bajari, P., C. Chu, D. Nekipelov, and M. Park (2016): Identification and Semiparametric Estimation of a Finite Horizon Dynamic Discrete Choice Model with a Terminating Action", *Quantitative Marketing and Economics* 14, 271-323.
- [14] Blevins, J. (2014): "Nonparametric Identification of Dynamic Decision Processes with Discrete and Continuous Choices," *Quantitative Economics* 5, 531-554.
- [15] Chiong, K., A. Galichon, and M. Shum (2013): "Duality in Dynamic Discrete-Choice Models", *Quantitative Economics* 7, 83-115.

- [16] Chou, C. (2016): "Identification and Linear Estimation of General Dynamic Programming Discrete Choice Models", working paper.
- [17] Eckstein, Z. and K. Wolpin (1989): "The Specification and Estimation of Dynamic Stochastic Discrete Choice Models: A Survey", *The Journal of Human Resources* 24, 562-598.
- [18] Gayle, G. and R. Miller (2006): "Life-Cycle Fertility and Human Capital Accumulation", working paper.
- [19] Gayle, G., L. Golan and R. Miller (2015): "Promotion, Turnover and Compensation in the Executive Labor Market," *Econometrica* 83, 2293-2369.
- [20] Heckman, J., and S. Navarro (2007): "Dynamic Discrete Choice and Dynamic Treatment Effects", *Journal of Econometrics* 136, 341-396.
- [21] Hotz, V., and R. Miller (1993): "Conditional Choice Probabilities and Estimation of Dynamic Models", *Review of Economic Studies* 60, 497-529.
- [22] Hu, Y., and M. Shum (2012): "Nonparametric Identification of Dynamic Models with Unobserved State Variables", *Journal of Econometrics* 171, 32-44.
- [23] Kasahara, H., and K. Shimotsu (2009): "Nonparametric Identification and Estimation of Finite Mixture Models of Dynamic Discrete Choices", *Econometrica* 77, 135-175.
- [24] Keane M., P. Todd and K. Wolpin (2011): "The Structural Estimation of Behavioral Models: Discrete Choice Dynamic Programming Methods and Applications", *Handbook of Labor Economics*, Volume 4A, edited by O. Ashenfelter and D. Card, North Holland, Chapter 4, 331-461.
- [25] Khorunzhina N., and R. Miller (2017): "American Dream Delayed: Shifting Determinants of Home Ownership," working paper.
- [26] Magnac T., and D. Thesmar (2002): "Identifying Dynamic Discrete Decision Processes", *Econometrica* 70, 801-816.
- [27] Miller, R. (1984): "Job Matching and Occupational Choice", *Journal of Political Economy* 92, 1086-1120.

- [28] Miller, R. (1997): "Estimating Models of Dynamic Optimization with Microeconomic data", in *Handbook of Applied Econometrics*, edited by M. Pesaran and P. Schmidt, Basil Blackwell, Vol. 2, pp. 246-299.
- [29] Norets, A., and X. Tang (2014): "Semiparametric Inference in Dynamic Binary Choice Models", *Review of Economic Studies* 81, 1229-1262.
- [30] Pakes, A. (1994): "The Estimation of Dynamic Structural Models: Problems and Prospects, Part II. Mixed Continuous-Discrete Control Models and Market Interactions", in *Advances in Econometrics: Proceedings of the 6th World Congress of the Econometric Society*, edited by J. Laffont and C. Sims, 171-259.
- [31] Pesendorfer, M., and P. Schmidt-Dengler (2008): "Asymptotic Least Squares Estimators for Dynamic Games", *The Review of Economic Studies* 75, 901-908.
- [32] Rust, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher", *Econometrica* 55, 999-1033.
- [33] Rust, J. (1994): "Structural Estimation of Markov Decision Processes", in *Handbook of Econometrics, Volume 4*, edited by R. Engle and D. McFadden, Elsevier-North Holland, 3081-3143.