Identifying Dynamic Discrete Choice Models off Short Panels*

Peter Arcidiacono
Duke University & NBER

Robert A. Miller
Carnegie Mellon University

September 6, 2013

Abstract

This paper analyzes identification in dynamic discrete choice models of single agents and noncooperative games. We extend previous work by providing new conditions for identifying flow payoffs in stationary settings, and in nonstationary settings when the data is sampled every period respondents make decisions. These results are also a benchmark for investigating identification when the relevant time horizon extends beyond the length of the data. We show that in short panels the utility flows of models with finite dependence are partially identified for particular normalizations. Finally, when finite dependence does not hold, or when the required normalizations are unattractive, we show how exclusion restrictions or stability of the flow payoff function over time can be used to recover flow payoffs.

1 Introduction

Dynamic discrete choice models in both single agent and games settings are increasingly used to explain panel data in labor economics, industrial organization and marketing.[1] It is widely recognized that the interpretation of structural models and their accuracy in predicting the effects of policy innovation

*We thank Jean-Marc Robin and seminar participants at Sciences Po and Toulouse for helpful comments. We acknowledge support from National Science Foundation Grant Awards SES0721059 and SES0721098.

[1]For surveys of this literature see Eckstein and Wolpin (1989), Rust (1994), Pakes (1994), Miller (1997), Aguirregabiria and Mira (2010) and Arcidiacono and Ellickson (2012).

depend critically on the assumptions needed to identify the models. The central role identification plays in determining the value of estimating structural models has motivated a small but growing literature on the identification of dynamic discrete choice models in both single-agent and multi-agent settings.

Research in this area dates back to Rust (1994), who showed that solutions to stationary infinite horizon dynamic discrete choice models are invariant to a broad class of utility transformations. Magnac and Thesmar (2002) later established that the flow payoffs for a two period model are identified in discrete choice optimization problems when the econometrician knows the joint probability distribution of the choice specific idiosyncratic disturbances and the discount fact, subject to a normalization on the flow payoffs in each of the periods. Pesendorfer and Schmidt-Dengler (2008) proved that their Magnac and Thesmar's result does not extend to games because the expected flow payoffs embed state-specific flow payoffs that depend on the actions of the other players, implying that the number of parameters exceeds the number of equations that characterize the empirical content of the game. Absent parametric restrictions, the state-specific utility flows in games are not identified. Still more recently, Norets and Tang (2012) provide conditions for identifying the probability distribution of the choice specific disturbance in stationary binary choice environments in the pretense of exclusion restrictions whereby a set of variables affects the transitions of the states but not the utility flows themselves.[2] In addition, empirical work using dynamic discrete choice models with fully parametric formulations often includes a discussion of identification as a preface to estimation.

This research has focused on cases where the model is either stationary or where the data covers the full time horizon. Yet many data sets are short panels: they do not cover the full lifetime of the sampled firms, individuals and products, and the sample respondents are often subjected to aggregate shocks that cannot be averaged out in the cross section. 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 very strong restrictions on the aggregate processes. But nonstationarity and aggregate shocks arise naturally in the human life cycle through aging, business cycles and the general equilibrium

[2] Most work in this area, including ours, focus on the case when all the unobserved variables are independently distributed over time, but Kasahara and Shimotsu (2009) and Hu and Shum (2012) relax this assumption in their analyses of identification.

effects of evolving demographics, in industries because of innovation and growth within and external to the market under consideration, and in marketing through the diffusion of new products and more generally over the product life cycle.

This paper extends previous work by deriving new conditions on identification for dynamic discrete choice models of individual optimization problems and multi-agent noncooperative games, applicable to both stationary and nonstationary settings. In the latter case we distinguish between panels that are short versus, for want of a better descriptor, long.

Our first set of results applies to stationary settings and nonstationary finite horizon settings where the data covers the full time horizon. We extend the results of Magnac and Thesmar (2002) to settings that last beyond two periods by building on Arcidiacono and Miller (2011), which provides a representation of the value function as a mapping of future streams of conditional choice probabilities and flow payoffs associated with any sequence of future choices. Under the standard assumptions that the distribution of the choice specific idiosyncratic disturbances and subjective discount factor are known, we show that the flow payoffs associated with the other choices are identified up to any normalization on the flow payoff for one of the choices. The normalized choice can vary by state and, in non stationary settings, by time period. A corollary of this result is that when the time horizon extends beyond the length of the data, in the absence of further restrictions, the degree of underidentification is directly comparable to the size of the state space upon which the continuation value is defined at the end of the sample. The corollary shows why the results of Magnac and Thesmar cannot be interpreted as applying to models lasting more than two periods without making assumptions about the value function, an endogenous mapping defined on the state space. This first set of results provides a launch pad for analyzing short panels.

For short panels, we provide three ways of restoring identification. The first, finite dependence, arises naturally out of the structure of many models. Finite dependence is a testable restriction on state transitions that is developed in Altug and Miller (1998), Arcidiacono and Miller (2011), Gayle (2013) and Aguirregabiria and Magesan (2013). Intuitively, finite dependence means that the long term effects of a current choice can be obliterated within a finite number of periods by following certain choice paths that may be optimal but typically are not. This property effectively bounds the number of future periods dynamic considerations are important for a current choice. We show that in the finite dependence case the utility flows a finite number of periods from the end of the sample models are identified up a normalization on the utility flows of the choices that establish finite dependence, normalizations that are consistent with normalizations that could be made to achieve identification in the long panel case. For example the utility flows from not renewing in a renewal problem or not exiting in a game with an exit decision are identified as mappings from the states and periods up until one period before the sample ends precisely because the utility from renewing or exiting in the last period is normalized.

This leaves three other reasons why a nonstationary discrete choice model cannot be identified off a short panel for a given disturbance structure. The normalization necessary for identifying a model with finite dependence might be at odds with knowledge emanating from beyond the data. The model might not exhibit finite dependence, a property that we show is testable. Finally, in multiagent games, the underlying payoff primitives are a multivalued correspondence from the expected flows identified from the equilibrium best response functions. In these three cases point identification requires exclusion restrictions that reduce the size of the parameter space, or functional form restrictions on the preferences embodied in the utility flows. In the latter parts of the paper we focus on two important cases in the literature by dropping the time dependence of utility flows, implying the nonstationarity is driven by the state transitions or the time horizon, and including a set of state variables that affect the state transitions but do not directly affect current utility flows.

The rest of the paper proceeds as follows. In the next section we lay out the dynamic optimization models and noncooperative games with private information to be studied. Then in Section 3 we analyze identification in stationary settings and nonstationary environment with complete histories. Section 4 provides a necessary and sufficient condition for restoring partial identification of flow payoffs, that is identifying some or all of the utility flows, when there is finite dependence. Then in Sections 5 and 6 we explore ways of achieving identification by replacing that assumption on the state transitions with assumptions on the functional form of utility and exclusion restrictions in we provide several new examples to illustrate when finite dependence holds, and adding further structure to a identify the primitives in games. We briefly conclude with some comments on the implications of this analysis for model specification and estimation methods.

2 Framework

This section lays out a general class of dynamic discrete choice models. Drawing upon our previous work in Arcidicacono and Miller (2011), we extend our representation of the conditional value functions which plays an overarching role in our analysis. Finally, we modify our framework to accommodate games with private information and explain the additional complications that identification in dynamic games pose.

2.1 Dynamic optimization discrete choice

In each period until T, for $T \leq \infty$, an individual chooses among J mutually exclusive actions. Let d_{jt} equal one if action $j \in \{1, \ldots, J\}$ is taken at time t and zero otherwise. The current period payoff for action j at time t depends on the state $x_t \in \{1, \ldots, X\}$. 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 individual's current period payoff from choosing j at time t is also affected by a choice-specific shock, ϵ_{jt}, which is revealed to the individual at the beginning of the period t. We assume the vector $\epsilon_t \equiv (\epsilon_{1t}, \ldots, \epsilon_{Jt})$ has continuous support and is drawn from a probability distribution that is independently and identically distributed over time with density function $g(\epsilon_t)$. We model the individual's current period payoff for action j at time t by $u_{jt}(x_t) + \epsilon_{jt}$.

The individual takes into account both the current period payoff as well as how his decision today will affect the future. Denoting the discount factor by $\beta \in (0, 1)$, the individual chooses the vector $d_t \equiv (d_{1t}, \ldots, d_{Jt})$ to sequentially maximize the discounted sum of payoffs:

$$E\left\{\sum_{t=1}^{T}\sum_{j=1}^{J}\beta^{t-1}d_{jt}\left[u_{jt}(x_t) + \epsilon_{jt}\right]\right\} \tag{1}$$

where at each period t the expectation is taken over the future values of $x_{t+1}, \ldots, x_T$ and $\epsilon_{t+1}, \ldots, \epsilon_T$. Expression (1) is maximized by a Markov decision rule which gives the optimal action conditional on t, x_t, and ϵ_t. We denote the optimal decision rule at t as $d_t^o(x_t, \epsilon_t)$, with jth element $d_{jt}^o(x_t, \epsilon_t)$. The probability of choosing j at time t conditional on x_t, $p_{jt}(x_t)$, is found by taking $d_{jt}^o(x_t, \epsilon_t)$ and integrating over ϵ_t:

$$p_{jt}(x_t) \equiv \int d_{jt}^o(x_t, \epsilon_t)\, g(\epsilon_t)\, d\epsilon_t \tag{2}$$

We then define $p_t(x_t) \equiv (p_{1t}(x_t), \ldots, p_{Jt}(x_t))$ as the vector of conditional choice probabilities.

Denote $V_t(x_t)$, the (ex-ante) value function in period t, as the discounted sum of expected future payoffs just before ϵ_t is revealed and conditional on behaving according to the optimal decision rule:

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

Given state variables x_t and choice j in period t, the expected value function in period $t+1$, discounted one period into the future is $\beta\sum_{x_{t+1}=1}^{X}V_{t+1}(x_{t+1})f_{jt}\left(x_{t+1}|x_t\right)$. Under standard conditions, Bellman's principle applies and $V_t(x_t)$ can be recursively expressed as:

$$\begin{aligned} V_t(x_t) &= E\left\{\sum_{j=1}^{J}d_{jt}^{o}\left(x_t,\epsilon_t\right)\left[u_{jt}(x_t)+\epsilon_{jt}+\beta\sum_{x_{t+1}=1}^{X}V_{t+1}(x_{t+1})f_{jt}\left(x_{t+1}|x_t\right)\right]\right\} \\ &= \sum_{j=1}^{J}\int d_{jt}^{o}\left(x_t,\epsilon_t\right)\left[u_{jt}(x_t)+\epsilon_{jt}+\beta\sum_{x_{t+1}=1}^{X}V_{t+1}(x_{t+1})f_{jt}\left(x_{t+1}|x_t\right)\right]g\left(\epsilon_t\right)d\epsilon_t \end{aligned} \tag{3}$$

where the second line integrates out the disturbance vector ϵ_t. We then define the choice-specific conditional value function, $v_{jt}(x_t)$, as the flow payoff of action j without ϵ_{jt} plus the expected future utility conditional on following the optimal decision rule from period $t+1$ on:[3]

$$v_{jt}(x_t) = u_{jt}(x_t)+\beta\sum_{x_{t+1}=1}^{X}V_{t+1}(x_{t+1})f_{jt}\left(x_{t+1}|x_t\right) \tag{4}$$

Our analysis of identification is based on a representation of $v_{jt}(x_t)$ that slightly generalizes Theorem 1 of Arcidiacono and Miller (2011). Both results are based on their Lemma 1, that for every $t\in\{1,\ldots,T\}$ and $p\in\Delta^J$, the J dimensional simplex, there exists a real-valued function $\psi_j\left(p\right)$ such that:

$$\psi_j[p_t(x)] \equiv V_t(x) - v_{jt}(x) \tag{5}$$

To interpret (5), note that the value of committing to action j before seeing ϵ_t is $v_{jt}(x_t)+E\left[\epsilon_{jt}\right]$. Therefore the expected loss from precommitting to j versus waiting until ϵ_t is observed and only then making an optimal choice, $V_t(x_t)$, is the constant $E\left[\epsilon_{jt}\right]$ plus $\psi_j[p_t(x_t)]$, a composite function that only depends x_t through the conditional choice probabilities. This result leads to the following theorem, proved using an induction.

Theorem 1 *For each choice $j\in\{1,\ldots,J\}$ and $\tau\in\{t,\ldots,T\}$, let any $d_\tau^*(x_\tau,j)$ denote any mapping from the state space $\{1,\ldots,X\}$ to R^J satisfying the constraints that i) $d_{kt}^*(x_t,j)=1$ and ii) $\sum_{k=1}^{J}d_{k\tau}^*(x_\tau,j) =$*

[3] For ease of exposition we refer to $v_{jt}(x_t)$ as the conditional value function in the remainder of the paper.

1. *Recursively define* $\kappa^*_\tau(x_{\tau+1}|x_t, j)$ *as:*

$$\kappa^*_\tau(x_{\tau+1}|x_t, j) \equiv \begin{cases} f_{jt}(x_{t+1}|x_t) & \text{for } \tau = t \\ \sum_{x_\tau=1}^{X} \sum_{k=1}^{J} d^*_{k\tau}(x_\tau, j)\, f_{k\tau}(x_{\tau+1}|x_\tau)\kappa^*_{\tau-1}(x_\tau|x_t, j) & \text{for } \tau = t+1, \dots, T \end{cases} \tag{6}$$

Then:

$$v_{jt}(x_t) = u_{jt}(x_t) + \sum_{\tau=t+1}^{T} \sum_{k=1}^{J} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \left[u_{k\tau}(x_\tau) + \psi_k[p_\tau(x_\tau)]\right] d^*_{k\tau}(x_\tau, j)\kappa^*_{\tau-1}(x_\tau|x_t, j) \tag{7}$$

Arcidiacono and Miller (2011) prove the theorem when $d^*_{k\tau}(x_\tau, j) \geq 0$ for all k and τ. In that case $\kappa^*_\tau(x_{\tau+1}|x_t, j)$ is the probability of reaching $x_{\tau+1}$ by following the sequence defined by $d^*_\tau(x_\tau, j)$,. Here we relax the constraint that $0 \leq d^*_{k\tau}(x_\tau, j) \leq 1$ in order to demonstrate that negative weights, and weights that exceed one, or pseudo-choices, can sometimes be useful in establishing identification.[4]

2.2 Extension to dynamic games

This framework extends naturally to dynamic game. In the games setting, we assume that there are I players making choices in periods $[1, \dots, T]$, $T \leq \infty$. The systematic part of payoffs to the i^{th} player not only depends on his own choice in period t, denoted by $d^{(i)}_t \equiv \left(d^{(i)}_{1t}, \dots, d^{(i)}_{Jt}\right)$, the state variables x_t, but also the choices of the other players, which we now denote by $d^{(\sim i)}_t \equiv \left(d^{(1)}_t, \dots, d^{(i-1)}_t, d^{(i+1)}_t, \dots, d^{(I)}_t\right)$. Denote by $U^{(i)}_{jt}\left(x_t, d^{(\sim i)}_t\right) + \epsilon^{(i)}_{jt}$ the current utility of agent i in period t, where $\epsilon^{(i)}_{jt}$ is an identically and independently distributed random variable that is private information to the firm. Although the players all face the same observed state variables, these state variables will affect each of the players in different ways. For example, a characteristic of player i may affect the payoff for player i differently than a characteristic of player i'. Hence, the payoff function is superscripted by i.

Players make simultaneous choices in each period. We denote $P_t\left(d^{(\sim i)}_t |x_t\right)$ as the joint conditional choice probability that the players aside from i collectively choose $d^{(\sim i)}_t$ at time t conditional on the state variables x_t. Since $\epsilon^{(i)}_t$ is independently distributed across all the players, $P_t\left(d^{(\sim i)}_t |x_t\right)$ has the product representation:

$$P_t\left(d^{(\sim i)}_t |x_t\right) = \prod_{\substack{i'=1 \\ i' \neq i}}^{I} \left(\sum_{j=1}^{J} d^{(i')}_{jt} p^{(i')}_{jt}(x_t)\right) \tag{8}$$

[4] This extension is also noted in Gayle (2013).

We assume each player acts like a Bayesian when forming his beliefs about the choices of the other players and that a Markov-perfect equilibrium is played. Hence, the beliefs of the players match the probabilities given in equation (8). Taking the expectation of $U_{jt}^{(i)}\left(x_t, d_t^{(\sim i)}\right)$ over $d_t^{(\sim i)}$, we define the systematic component of the current utility of player i as a function of the state variables as:

$$u_{jt}^{(i)}\left(x_t\right) = \sum_{d_t^{(\sim i)} \in J^{I-1}} P_t\left(d_t^{(\sim i)} \left| x_t\right.\right) U_{jt}^{(i)}\left(x_t, d_t^{(\sim i)}\right) \tag{9}$$

The values of the state variables at period $t+1$ are determined by the period t choices by all the players as well as the values of the period t state variables. Denote $F_{jt}\left(x_{t+1}\left| x_t, d_t^{(\sim i)}\right.\right)$ as the probability of x_{t+1} occurring given action j by player i in period t, when its state variables are x_t and the other players choose $d_t^{(\sim i)}$. From the perspective of player i the probability of transitioning from x_t to x_{t+1} given action j is:

$$f_{jt}^{(i)}\left(x_{t+1}\left| x_t\right.\right) = \sum_{d_t^{(\sim i)} \in J^{I-1}} P_t\left(d_t^{(\sim i)} \left| x_t\right.\right) F_{jt}\left(x_{t+1}\left| x_t, d_t^{(\sim i)}\right.\right) \tag{10}$$

The expressions for the conditional value functions for player i are the same as we described in Subsection 2.1. Loosely speaking, player i solves a dynamic optimization problem treating the other player's equilibrium actions as nature that help determine both the flow payoffs and the state transitions.

As in Subsection 2.1, consider for all $\tau \in \{t, \ldots, T\}$ any sequence of pseudo-choices:

$$d_\tau^{*(i)}(x_\tau, j) \equiv \left(d_{1\tau}^{*(i)}(x_\tau, j), \ldots, d_{J\tau}^{*(i)}(x_\tau, j)\right)$$

subject to the constraints $\sum_{k=1}^{J} d_{k\tau}^{*(i)}(x_\tau, j) = 1$ and starting value $d_{jt}^{*(i)}(x_t, j) = 1$. Using the pseudo-choices, and taking the equilibrium actions of the other players as given, recursively define $\kappa_\tau^{*(i)}(x_{\tau+1}|x_t, j)$ in a similar manner to equation (6) as:

$$\kappa_\tau^{*(i)}(x_{\tau+1}|x_t, j) \equiv \begin{cases} f_{jt}^{(i)}(x_{t+1}|x_t) & \text{for } \tau = t \\ \sum_{x_\tau=1}^{X} \sum_{k=1}^{J} d_{k\tau}^{*(i)}\left(x_\tau, j\right) f_{k\tau}^{(i)}(x_{\tau+1}|x_\tau) \kappa_{\tau-1}^{*(i)}(x_\tau|x_t, j) & \text{for } \tau = t+1, \ldots, T \end{cases} \tag{11}$$

Adding i superscripts to (7) , it now follows that Theorem 1 applies to this multiagent setting in exactly the same way as in a single agent setting.

Two critical differences distinguish noncooperative discrete choice dynamic games from their single agents counterparts, and both are relevant for studying identification. Whereas $u_{jt}(x_t)$ denotes the primitive utility flow term in single agent optimization problems, $u_{jt}^{(i)}(x_t)$ defined by (9) is a reduced form parameter that depends on the actions of the other players. In dynamic games, the flow payoff $u_{jt}^{(i)}(x_t)$

is not a primitive but an expected utility found by integrating $U_{jt}^{(i)}\left(x_t, d_t^{(\sim i)}\right)$ over the joint probability distribution $P_t\left(d_t^{(\sim i)} | x_t\right)$ induced by the current actions of the other players simultaneously making their equilibrium choices that are partly determined by their private information.

Similarly, $f_{jt}(x_{t+1} | x_t)$ is the primitive defining the state transition probabilities in single agent optimization problems, but $f_{jt}^{(i)}(x_{t+1} | x_t)$, defined by (10), is a reduced form parameter that depends on the conditional choice probabilities of the other players, $P_t\left(d_t^{(\sim i)} | x_t\right)$, as well as the primitive $F_{jt}\left(x_{t+1} \middle| x_t, d_t^{(\sim i)}\right)$. While it is easy to interpret restrictions placed directly on the primitive $f_{jt}(x_{t+1} | x_t)$ in single agent problems, placing restrictions on $F_{jt}\left(x_{t+1} \middle| x_t, d_t^{(\sim i)}\right)$ complicates matters in dynamic games because of the endogenous effects arising from $P_t\left(d_t^{(\sim i)} | x_t\right)$ on $f_{jt}^{(i)}(x_{t+1} | x_t)$.

3 Three Theorems on Identification

The objects of identification in the optimization model are the utility flows, the discount factor, the transition matrix of the observed state variables, and the distribution of the unobserved variables,[5] summarized with the notation (u, β, F, G). In this section we build upon Rust (1994) and Magnac and Thesmar (2002) in single agent settings and Pesendorfer Schmidt-Dengler (2008) in games settings, by considering identification when (β, F, G) are known.[6] First we show that u is identified up to a normalization on the flow payoffs for one of the choices in each state when either the environment is stationary or when $\mathcal{T} = T$. Finally, we analyze the degree of under-identification when $\mathcal{T} < T$.

3.1 Normalizing utility flows

Rust (1994, Lemma 3.2 on page 3127) showed that the solution to a stationary infinite horizon discrete choice optimization problem is invariant to a broad class of utility transformations. His result can be simply

[5] Often 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 (2013) have proposed simple estimators for a broad range of error distributions.

[6] The assumption that (β, F, G) is known is standard. Typically F is identified from the transitions alone by assuming that all the state variables are observed, estimates of β, that calibrate a person's subjective discount factor in a stationary model are obtained from other data, and G is selected largely on the basis of tractability.

extended to nonstationary optimization problems and dynamic games, but leaves unanswered the question of how to partition the identified sets of utility specifications, so that specifications belonging to different partitions are associated with different decision rules, while those in the same partition are observationally equivalent. As a first step towards deriving this partition, we now show there is an observationally equivalent dynamic optimization problem to (u, β, F, G), which we denote by (u^*, β, F, G), where for each (t, x) we arbitrarily select any one choice $l(x,t) \in \{1, \ldots, J\}$ and set the flow utility associated with that choice, $u^*_{l(x,t),t}(x)$, to an arbitrary real value we denote by $c_t(x)$. Similarly in the infinite horizon analogue we select for each x any one choice $l(x) \in \{1, \ldots, J\}$ and set the flow utility associated with that choice, $u^*_{l(x)}(x)$, to an arbitrary real value denoted by $c(x)$.[7]

Theorem 2 *In the finite horizon case let $l(x,t) \in \{1, \ldots, J\}$ and $c_t(x) \in \Re$ respectively denote any arbitrarily defined normalizing action and benchmark flow utility the associated with (x,t), and define for all $j \in \{1, \ldots, J\}$:*

$$u^*_{jT}(x) \equiv u_{jT}(x) - u_{l(x,t),T}(x) + c_T(x)$$

and:

$$u^*_{jt}(x) = u_{jt}(x) + c_t(x) - u_{l(x,t),t}(x) + \sum_{\tau=t+1}^{T} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \left[u^*_{1\tau}(x_\tau) - u_{1\tau}(x_\tau)\right] \left[\kappa^*_{\tau-1}(x_\tau | x_t, l(x,t)) - \kappa^*_{\tau-1}(x_\tau | x_t, j)\right]$$

When the environment is stationary, define:

$$u_j \equiv \begin{bmatrix} u_j(1) \\ \vdots \\ u_j(X) \end{bmatrix}, \quad u^*_j \equiv \begin{bmatrix} u^*_j(1) \\ \vdots \\ u^*_j(X) \end{bmatrix}, \quad \tilde{u} \equiv \begin{bmatrix} u_{l(1)}(1) \\ \vdots \\ u_{l(X)}(X) \end{bmatrix}, \quad F_j \equiv \begin{bmatrix} f_j(1|1) & \ldots & f_j(X|1) \\ \vdots & \ddots & \vdots \\ f_j(1|X) & \ldots & f_j(X|X) \end{bmatrix}, \quad c \equiv \begin{bmatrix} c(1) \\ \vdots \\ c(X) \end{bmatrix}$$

Then $[\mathcal{I} - \beta \mathcal{F}_1]$ is invertible. Define for all $j \in \{1, \ldots, J\}$:

$$u^*_j = u_j + c - \tilde{u} + \beta (F_1 - F_j) [\mathcal{I} - \beta F_1]^{-1} (u^*_1 - u_1)$$

Then (u^*, β, F, G) is observationally equivalent to (u, β, F, G).

A common normalization in empirical work is to set $u^*_{1t}(x) = 0$ for all (t, x) in the finite horizon case and $u^*_1(x) = 0$ for all x in the stationary case. Theorem 2 demonstrates that a normalization like that is

[7] Analogously for the dynamic games defined in Section 2, there is an observationally equivalent dynamic game in which $U^{(i)}_{jt}\left(x_t, d^{(\sim i)}_t\right) = c_t(x)$ for all $\left(l, t, x_t, d^{(\sim i)}_t\right)$, where again j can vary with the time period and the state. Note that while $c_t(x)$ varies with the state and the time period, it does not vary with the decisions of the other players.

necessary to identify the remaining utility parameters. The next section provides conditions under which it is sufficient.

3.2 Data is sampled from the whole population

Magnac and Thesmar (2002, Theorem 2 and Corollary 3 on pages 807 and 808) establish identification of the flow payoff for $T = 2$ finite when $G(\epsilon_t)$ and β are known, $u_1(x)$ is normalized for all x, and the continuation value for one of the actions is also normalized. We extend their results to the case where data on the full time horizon is observed as well as to stationary environments.

Let $d^*_{1\tau}(x_\tau) = 1$ for all τ in equation (7) and subtract $v_{1t}(x_t)$ from $v_{jt}(x_t)$ to obtain:

$$v_{jt}(x_t) - v_{1t}(x_t) = u_{jt}(x_t) + \sum_{\tau=t+1}^{T} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \psi_1[p_\tau(x_\tau)] \left[\kappa^*_{\tau-1}(x_\tau|x_t, j) - \kappa^*_{\tau-1}(x_\tau|x_t, 1)\right] \tag{12}$$

An alternative expression for this difference can be obtained by differencing the expressions for $\psi_1(x_t)$ and $\psi_t(x_t)$ given in equation (5):

$$v_{jt}(x_t) - v_{1t}(x_t) = \psi_1[p_t(x_t)] - \psi_j[p_t(x_t)] \tag{13}$$

As shown in Theorem 3 below, the two expressions for $v_{jt}(x_t) - v_{1t}(x_t)$ can then be used to form expressions for $u_{jt}(x_t)$ as a function of the transition probabilities, the conditional choice probabilities, and the discount factor. Further, Theorem 3 shows how the problem simplifies in the stationary case where the time subscripts are dropped from the flow payoffs and the transition functions and when the time horizon is infinite.[8]

Theorem 3 *For all j, t, and x_t, the flow payoff $u_{jt}(x_t)$ can be expressed as:*

$$u_{jt}(x_t) = \psi_1[p_t(x_t)] - \psi_j[p_t(x_t)] + \sum_{\tau=t+1}^{T} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \psi_1[p_\tau(x_\tau)] \left[\kappa^*_{\tau-1}(x_\tau|x_t, 1) - \kappa^*_{\tau-1}(x_\tau|x_t, j)\right] \tag{14}$$

When the environment is stationary, let $\mathcal{I}$ denote the X dimensional identity matrix and define:

$$u_j \equiv \begin{bmatrix} u_j(1) \\ \vdots \\ u_j(X) \end{bmatrix}, \quad \mathcal{F}_j \equiv \begin{bmatrix} f_j(1|1) & \dots & f_j(X|1) \\ \vdots & \ddots & \vdots \\ f_j(1|X) & \dots & f_j(X|X) \end{bmatrix}, \quad \Psi_j \equiv \begin{bmatrix} \psi_j[p(1)] \\ \vdots \\ \psi_j[p(X)] \end{bmatrix}$$

[8] This theorem, first presented in a section on identification in Arcidiacono and Miller (2010), was omitted from the published version (Arcidiacono and Miller, 2011), which focuses on estimation in the presense of unobserved heterogeneity.

Then $[\mathcal{I} - \beta \mathcal{F}_1]$ *is invertible and for all* j*:*

$$u_j = \Psi_j - \Psi_1 + \beta \left(\mathcal{F}_1 - \mathcal{F}_j\right) \left[\mathcal{I} - \beta \mathcal{F}_1\right]^{-1} \ \Psi_1 \tag{15}$$

Given the assumptions made at the beginning of this section regarding the state transitions, conditional choice probabilities, the discount factor, and the distribution of the structural errors, everything on the right hand side of both (14) and (15) is known and, therefore, both systems are exactly identified. These equations yield asymptotically efficient estimators of the unrestricted utility flows. They are defined by substituting sample analogues for the conditional choice probabilities into the closed form 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.

Note that in the finite horizon case, intertemporal preferences are fully accounted for by subscripting the utility flow terms by t, and the appearance of β is purely cosmetic, which without loss of generality can be set to one. Therefore the intertemporal features of finite models are identifed by the data generating process in long panels. In contrast the trade off between current and future preferences in the stationary model is not identified at all, and the interpretation of the parameters hinges critically on the value of β assigned from outside the data generating process.

3.3 Observational equivalence in short panels

Given that we do not see state transitions and conditional choice probabilities after $\mathcal{T}$, we express u_{jt} as in (14) relative to choice 1 (the normalized choice) for the first $\mathcal{T}$ periods and then use the value function at $\mathcal{T}+1$. This leads to the following expression for u_{jt}:

$$\begin{aligned} u_{jt}(x_t) \quad = \quad & \psi_1[p_t(x_t)] - \psi_j[p_t(x_t)] + \sum_{\tau=t+1}^{\mathcal{T}} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \psi_1[p_\tau(x_\tau)] \left[\kappa^*_{\tau-1}(x_\tau|x_t,1) - \kappa^*_{\tau-1}(x_\tau|x_t,j)\right] \\ & + \sum_{x_{\mathcal{T}+1}=1}^{X-1} \beta^{\mathcal{T}-t} V_{\mathcal{T}+1}(x_{\mathcal{T}+1})] \left[\kappa^*_{\mathcal{T}}(x_{\mathcal{T}+1}|x_t,1) - \kappa^*_{\mathcal{T}}(x_{\mathcal{T}+1}|x_t,j)\right] \end{aligned} \tag{16}$$

This expression provides the basis for the following theorem giving the degree of under-identification. Note that is is the last term that leads to under-identification, the degree of which is specified in Theorem 4 below. Given the choice probabilities, $u_{jt}(x_t)$ is a linear mapping of $V_{\mathcal{T}+1}(x_{\mathcal{T}+1})$. Let $X_{\mathcal{T}+1}$ denote the number states that can be reached at time $\mathcal{T}$ +1.

Theorem 4 *Given β, $G(\epsilon)$ and $u_{1t}(x_t) = 0$ for all t and x_t, the degree of under-identification for the first $\mathcal{T}$ flow payoffs is at most $X_{\mathcal{T}+1} - 1$.*

Theorem 4 shows that the degree of under-identification is less than or equal to the number of different states that can be reached in the last period of the data. In order to achieve identification, we could normalize the value functions in the last period to zero. At that point we would treat the sample as if the time horizon was $\mathcal{T}$ rather than T. The crippling limitation of this approach is that all the flow utilities can only be interpreted relative to the arbitrary normalization adopted in the final period of an endogenous mapping, a mapping that would change with the length of the panel. However, progress can be made when the transitions of the state variables satisfy certain properties (finite dependence), when there are exclusion restrictions, or when the flow payoffs are assumed to be independent of time. In the last case, the non-stationarity would then be driven either by the finite time horizon or non-stationarity in the state transitions.

4 Finite dependence

In this section we show that if the state transitions satisfy a finite dependence property defined below, then a subset of the flow payoffs are identified even when the time horizon extends beyond the sample period with no normalizations made on the continuation values. Identification of this subset, however, does come at a cost. As Theorem 2 showed, there are a set of normalizations that are observationally equivalent. Identification in the short panel case is restored under particular normalizations but not others.

To fix ideas we first consider dynamic optimization problems with terminal choices (that end the problem) and renewal problems (that restart a cycle) to show how the inversion that recovers utility flows from the conditional choice probabilities simplifies under certain normalizations but not under others. Then after defining the finite dependence property, we formally state the identification result, further illustrate its application in single agent problems and dynamic games, and provide an algorithm for showing how finite dependence is established.

4.1 Terminal and renewal choices

To see how finite dependence aids identification we first consider two special cases: when there is a terminal or renewal choice. Terminal choices are named that way because they end the optimization problem or game by preventing any future decisions; irreversible sterilization against future fertility, (Hotz and Miller, 1993) and firm exit from an industry (Aguirregabiria and Mira 2007) are examples. The defining feature of a renewal choice is that it resets the states that were influenced by past actions. Turnover and job matching (Miller, 1984), or replacing a bus engine (Rust,1987), are illustrative of renewal actions. Let the first choice denote the terminal or renewal choice. In such models, following any choice $j \in \{1, \ldots, J\}$ with the first one leads to same value of state variables after two periods. Thus for all $t < T$ and x_t the probability distribution of x_{t+2} conditional on x_t does not depend on the choice made in period t if the terminal or renewal choice is taken in period $t+1$:

$$\sum_{x_{t+1}=1}^{X} f_{1,t+1}(x_{t+2}|x_{t+1}) f_{jt}(x_{t+1}|x_t) = \sum_{x_{t+1}=1}^{X} f_{1,t+1}(x_{t+2}|x_{t+1}) f_{1t}(x_{t+1}|x_t) \tag{17}$$

Normalizing the utility for the terminal or renewal choice to zero in all states greatly simplifies the inversion embedded in Theorem 2. The formulas in the theorem give the utility from taking any other action $j \in \{2, \ldots, J\}$ in period $t < \mathcal{T} - 1$ as:

$$u_{jt}(x_t) = \psi_1[p_t(x_t)] - \psi_j[p_1(x_t)] - \sum_{x_{t+1}=1}^{X} \beta\psi_1[p_{t+1}(x_{t+1})] f_{jt}(x_{t+1}|x_t) \tag{18}$$

when there is a terminal choice and as:

$$u_{jt}(x_t) = \psi_1[p_t(x_t)] - \psi_j[p_1(x_t)] + \sum_{x_{t+1}=1}^{X} \beta\psi_1[p_{t+1}(x_{t+1})] \left[f_{1t}(x_{t+1}|x_t) - f_{jt}(x_{t+1}|x_t)\right] \tag{19}$$

when there is a renewal choice.

Notice that in either case only the current conditional choice probabilities and those one period into the future are used to identify current utility under this normalization. Now suppose there are three periods of data in a model with terminal or renewal choices, and select any normalization for choices made in the first two periods, but retain the normalization of the first choice in the third. Appealing to Theorem 2, we write:

$$u_{j1}^*(x) = u_{j1}(x) + c_1(x) - u_{l(x,1)1}(x) + \sum_{\tau=2}^{3} \sum_{x_\tau=1}^{X} \beta^{\tau-1} \left[u_{1\tau}^*(x_\tau) - u_{1\tau}(x_\tau)\right] \left[\kappa_{\tau-1}^*(x_\tau|x_1, l(x,1)) - \kappa_{\tau-1}^*(x_\tau|x_1, j)\right]$$

To amplify, every normalization of the utility flows is identified, except for the utilities in last sampled period, subject to a proviso that in the second last period of the sample the utility of the terminal or renewal action is the normalization. Finally we note that the condition of normalizing the terminating or renewal action in the last period for which data is available is necessary.

4.2 Defining finite dependence

Consider two sequences of decision weights that begin at date t in state x_t, one with choice j and the other with choice k. We say that the pair of choices $\{j, j'\}$ exhibits ρ-period dependence if there exists sequences of decision weights from j and j' such that:

$$\kappa^*_{t+\rho}(x_{t+\rho+1}|j) = \kappa^*_{t+\rho}(x_{t+\rho+1}|j') \tag{20}$$

for all $x_{t+\rho+1}$. That is, the weights associated with each state are the same across the two paths after ρ periods.

The equation defining ρ-period finite dependence is then a multi-period generalization of (17) . More generally ρ-period finite dependence requires ρ future periods of future conditional choice probabilities to identify the utility flows in the first period of the data. If there are more than ρ periods of data, identification requires the final ρ periods to be normalized in a way dictated by the finite dependence property. Under finite dependence, differences in current utility $u_{jt}(x_t) - u_{kt}(x_t)$ can be expressed as:

$$\begin{aligned} u_{jt}(x_t) - u_{j't}(x_t) &= \psi_{j'}[p_t(x_t)] - \psi_j[p_t(x_t)] \\ &+ \sum_{\tau=t+1}^{t+\rho} \sum_{k=1}^{J} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \left\{u_{k\tau}(x_\tau) + \psi_k[p_\tau(x_\tau)]\right\} \left[d^*_{k\tau}(x_\tau|j)\kappa^*_{\tau-1}(x_\tau|j) - d^*_{k\tau}(x_\tau|j')\kappa^*_{\tau-1}(x_\tau|j')\right] \end{aligned} \tag{21}$$

This result follows directly from Theorem 1. For the first ρ periods, we se the decision sequences such that finite dependence holds. After ρ periods we set the two decision sequences to be the same, implying the differenced term in the second line is zero at $t + 1 + \rho$.

4.3 Finite dependence and identification

Theorem 4 established the degree of under-identification when the panel does not cover the relevant time horizon. The under-identification resulted from decisions made in the observed panel depending on expectations beyond the sample period. Finite dependence offers some respite. When this property holds, differences in the distribution of state variables that arise from making different current choices can be obliterated within a finite number of periods through a judicious weighting of future decisions in the interim phase. While the choice probabilities associated with the interim decisions do not correspond to the CCPs generated by optimal behavior, and not all of them are necessarily positive, the correction factor between the conditional valuation function and the ex-ante continuation value rebalances differences in lifetime utility arising from non-optimizing and possibly infeasible behavior. In this way the role of future consequences from current choices is captured by a relatively small number of future utility flow and correction factors even though the decision horizon of the individual is much longer.

A caveat limits the role of finite dependence in restoring identification. As is evident from Equation (21), future utility flow terms of the form $u_{i\tau}(x_\tau)$ appear in the finite dependence representation of the current utility difference between $u_{jt}(x_t)$ and $u_{kt}(x_t)$. Thus none of the utility flow terms for the periods immediately preceding the last sampling period are identified. It now follows by an induction that the assumption of finite dependence does not suffice to identify any of the utility flow terms when the paths defining finite dependence that determine Equation (21) involve flow utilities. At first glance this argument seems to doom finite dependence as a device for achieving identification. However whether future utility flow terms explicitly appear in the finite dependence representation implied by (20) depends on which normalization is adopted. Theorem demonstrates that all normalizations benchmarking one choice for each period and state are observationally equivalent. In other words the theorem gives provisional licence to judiciously pick a normalization that makes the flow utilities in the interim periods on the second line of the right side of (21) drop out, simplifying the equation to:

$$\begin{aligned} u_{jt}(x_t) - u_{kt}(x_t) &= \psi_k[p_t(x_t)] - \psi_j[p_t(x_t)] \\ &+ \sum_{\tau=t+1}^{t+\rho} \sum_{i=1}^{J} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \psi_i[p_\tau(x_\tau)] \left[d_{i\tau}^*(x_\tau|j)\kappa_{\tau-1}^*(x_\tau|j) - d_{i\tau}^*(x_\tau|k)\kappa_{\tau-1}^*(x_\tau|k) \right] \end{aligned} \tag{22}$$

Theorem 5 below states the required set of normalizations for identifying utility flows in a nonstationary environment with panels of limited length when there is finite dependence.

Theorem 5 *If there exist decision sequences for a pair of choices* $\{1, j\}$ *such that* (20) *holds for some* $\rho < \mathcal{T}$ *and* x_t *, then there exists a set of normalizations consistent with Theorem 2 for the flow payoffs between* t *and* $t + \rho$ *so that* (21) *simplifies to* (22).

One can prove by induction that no restrictions on the normalization adopted in periods preceding $\mathcal{T}$ $-\rho-1$ are necessary to identify the utility flows in those periods. In some settings, the normalizations along the finite dependence path are in concordance with prior beliefs formed from outside sources about the current value of taking certain benchmark actions, that is natural. In other cases, the normalizations given in Theorem 5 are unnatural even though they are observationally equivalent to those preferred by the researcher. In these cases, discussed in Sections 5 and 6, identification can be achieved under more plausible normalizations by imposing exclusion or functional form restrictions on the utility flows that effectively shrink the parameter space.

4.4 Establishing finite dependence in single agent settings

In the previous section we established identification conditional on finite dependence holding. Here we establish conditions under which finite dependence holds for a pair of choices $\{j, j'\}$. In the process, we also show simple ways of checking for finite dependence. We derive sufficient conditions for one-period-ahead finite dependence to hold with two choices and then prove the more general case for ρ-period dependence.

For finite dependence to hold, there must exist weights $d^*_{2t+1}(x_{t+1}, 1)$ and $d^*_{1t+1}(x_{t+1}, 1)$ such that the probability of being in each state at $t+2$ is the same given the two different initial decisions. Using (6), the probability of state x_{t+2} resulting from choice $j \in \{1, 2\}$ in state x_t at time t and a decision rule $d^*_{t+1}(x_t, j)$ can be written as:

$$\begin{aligned} \kappa^*_{t+1}(x_{t+2}|x_t, j) &= \sum_{x_{t+1}} d^*_{2t+1}(x_{t+1}, j) \left[f_{2t+1}(x_{t+2}|x_{t+1}) - f_{1t+1}(x_{t+2}|x_{t+1})\right] f_{jt}(x_{t+1}|x_t) \\ &\quad + \sum_{x_{t+1}} f_{1t+1}(x_{t+2}|x_{t+1}) f_{jt}(x_{t+1}|x_t) \end{aligned} \tag{23}$$

Given choice j at time t only a subset of the state space may be reached from state x_t where we refer to this subset as the attainable states. Denote $N^*_{t+1}(j, x_t)$ as the number of attainable states in period $t+1$ given the choice j in state x_t at time t. Now consider all the possible states that could be reached two-periods-ahead given *either* choice j or j' which we denote as $N^*_{t+2}(x_t)$. It the weights on each of these states that need to be aligned across the two choice paths.

We now look to express the finite dependence condition in matrix form. Define $F^*_t(j, x_t)$ as a vector containing the probabilities of transitioning to each of the $N^*_{t+1}(j, x_t)$ attainable states given the choice sequence beginning with j and state x_t. Note that by definition all the elements of $F_t^*(j, x_t)$ are greater than zero. Denote $D_{2t+1}^*(j, x_t)$ as a vector giving the weight placed on choice 2 for each of the $N_{t+1}^*(j, x_t)$ attainable states at $t+1$. Finally, denote $F_{kt+1}(j, x_t)$ as an $N_{t+1}^*(j, x_t) \times (N_{t+2}^*(x_t) - 1)$ which gives the probability of transitioning from each of the $N_{t+1}^*(j, x_t)$ attainable states given initial choice j to $N_{t+2}^*(x_t) - 1$ of the attainable states at $t+1$ given choice k at $t+1$. The finite dependence condition is then:

$$\begin{aligned} & \left[F_{2t+1}^T(2, x_t) - F_{1t+1}^T(2, x_t)\right]\left[D_{2t+1}^*(2, x_t) \circ F_t^*(2, x_t)\right] + F_{1t+1}^T(2, x_t) F_t^*(2, x_t) \\ = \quad & \left[F_{2t+1}^T(1, x_t) - F_{1t+1}^T(1, x_t)\right]\left[D_{1t+1}^*(2, x_t) \circ F_t^*(1, x_t)\right] + F_{1t+1}^T(1, x_t) F_t^*(1, x_t) \end{aligned} \tag{24}$$

where $\circ$ refers to element-by-element multiplication.

There are then $N_{t+2}^*(x) - 1$ equations with $N_{t+1}^*(1, x_t) + N_{t+1}^*(2, x_t)$ unknowns. We can rewrite this system as:

$$\begin{aligned} & \begin{bmatrix} F_{2t+1}^T(2, x_t) - F_{1t+1}^T(2, x_t) & F_{1t+1}^T(1, x_t) - F_{2t+1}^T(1, x_t) \end{bmatrix} \begin{bmatrix} D_{2t+1}^*(2, x_t) \circ F_t^*(2, x_t) \\ D_{1t+1}^*(2, x_t) \circ F_t^*(1, x_t) \end{bmatrix} \\ & = F_{1t+1}^T(1, x_t) F_t^*(1, x_t) - F_{1t+1}^T(2, x_t) F_t^*(2, x_t) \end{aligned} \tag{25}$$

implying that if the rank of

$$\begin{bmatrix} F_{2t+1}^T(2, x_t) - F_{1t+1}^T(2, x_t) & F_{1t+1}^T(1, x_t) - F_{2t+1}^T(1, x_t) \end{bmatrix} \tag{26}$$

is $N_{t+2}^*(x) - 1$ then finite dependence holds. Note that the size of the matrix for which we need to check the rank condition may be quite small in practice. Because the decision weights can depend on the initial state, the size of the matrix is determined by the number of states that are attainable in two periods from a particular initial state.

For the general case, where there are more than two choices and when finite dependence cannot be achieved in one period, we can derive a similar result. Namely, suppose we wanted to see if the problem exhibited ρ period dependence given prescribed decision sequences up through $\rho - 1$ ($\rho > 1$). Let $\tau = t + \rho$. As before, denote $N_{\tau+1}^*(j, x_t)$ as the number of attainable states given the prescribed decision sequence up through τ that begins with choice j and denote $N_{\tau+2}^*(x_t)$ given either prescribed decision sequence. Denote $F_{k\tau+1}(j, x_t)$ as an $N_{\tau+1}^*(j, x_t) \times (N_{\tau+2}^*(x_t) - 1)$ which gives the probability of transitioning from each of the $N^*_{\tau+1}(j, x_t)$ attainable states given initial choice j to the $N^*_{\tau+2}(x_t) - 1$ attainable states at $\tau + 2$ given *either* initial choice j or j'. Define $\mathcal{F}_{\tau+1}(j, x_t)$ as an $(N^*_{\tau+2}(x_t) - 1) \times ((J-1)N^*_{\tau+1}(j, x_t))$ matrix given by:

$$\mathcal{F}_{\tau+1}(j, x_t) = \begin{bmatrix} F_{2\tau+1}(j, x_t) - F_{1\tau+1}(j, x_t) \\ \vdots \\ F_{k\tau+1}(j, x_t) - F_{1\tau+1}(j, x_t) \\ \vdots \\ F_{J\tau+1}(j, x_t) - F_{1\tau+1}(j, x_t) \end{bmatrix}^T$$

which is the analog to (26). We then have the following result:

Theorem 6 *If the rank of* $\begin{bmatrix} \mathcal{F}_{\tau+1}(j, x_t) & -\mathcal{F}_{\tau+1}(j', x_t) \end{bmatrix}$ *is* $N^*_{\tau+2}(x+1) - 1$ *then finite dependence can be achieved in* $\tau - t + 1$ *periods.*

Example 1: A nonstationary search model

To illustrate how negative weights are useful in obtaining finite dependence, we consider a simple search model in which jobs last only one period.[9] Each period $t \in \{1, \ldots, T\}$ an individual may stay home by setting $d_{1t} = 1$, or apply for temporary employment setting $d_{2t} = 1$. Job applicants are successful with probability λ_t, and the value of the position depends on the experience of the individual denoted by $x \in \{1, \ldots, X\}$. If the individual works his experience increases by one unit, and remains at the current level otherwise. The preference primitives are given by the current utility from staying home, denoted by $U_{1t}(x_t)$, and the utility from working, $U_{2t}(x_t)$. Thus the dynamics of the model come strictly through experience, and the nonstationarities arising through offer arrival weights, λ_t, and through wages varying over time, as indicated by subscripting utilities with the time period.

To demonstrate this model satisfies finite dependence with $\rho = 1$, we construct two paths, one starting with stay home decision, $d_{1t} = 1$, and the other beginning with an employment application, $d_{2t} = 1$. These two paths have differently weighted decisions in period $t + 1$ that generate the same probability distribution of x_{t+2} conditional on x_t. Following either the two period sequence work, $d_{2t} = 1$, home,

[9] The example can easily be extended to the case where the individual can choose to stay with his current job. We focus on the simpler case here for ease of exposition, at the same time noting that a sizable fraction of the Spanish workforce gain experience through a sequence of temporary jobs before finding a permanent position.

$d_{1t+1} = 1$, or the two period sequence home, $d_{1t} = 1$, work with weight $d_{2t+1} = \lambda_t/\lambda_{t+1}$, home with weight $1-\lambda_t/\lambda_{t+1}$ generates the same distribution for x_{t+2}. Namely $x_{t+2} = x_t$ with probability $f_{2t}(x_t|x_t) = 1-\lambda_t$, and $x_{t+2} = x_t + 1$ with probability $f_{2t}(x_t + 1|x_t) = \lambda_t$. Notice that if $\lambda_t < \lambda_{t+1}$ then $d_{2,t+1} > 1$ and $d_{2,t+1} = 1 - \lambda_t/\lambda_{t+1} < 0$.

This search model has finite dependence, and is therefore identified by Theorem 5. It is nevertheless instructive to directly establish identification, because it yields an expression for the utility primitives in terms of the conditional choice, job offer and experience transition probabilities in terms of the normalization. The current utility from staying home is just $u_{1t}(x_t) \equiv U_{1t}(x_t)$, and the (expected) current utility from applying for a temporary position in period t is:

$$u_{2t}(x_t) \equiv (1 - \lambda_t) U_{1t}(x_t) + \lambda_t U_{2t}(x_t) \tag{27}$$

Since staying at home does not increase working experience $f_{1t}(x_t|x_t) = 1$, and conditional on applying for a position $f_{2t}(x_t + 1|x_t) = \lambda_t$, Equation (21) simplifies to:

$$\begin{aligned} u_{2t}(x) - u_{1t}(x) \quad = \quad & \psi_1[p_t(x)] - \psi_2[p_t(x)] - \beta \{u_{1t+1}(x) + \psi_1[p_{t+1}(x)]\} (1 - \lambda_{t+1}) \frac{\lambda_t}{\lambda_{t+1}} \\ & -\beta \{u_{1t+1}(x + 1) + \psi_1[p_{t+1}(x + 1)]\} \lambda_t + \beta \{u_{2t+1}(x) + \psi_2[p_{t+1}(x)]\} \frac{\lambda_t}{\lambda_{t+1}} \end{aligned} \tag{28}$$

for all $x \in \{1, \ldots, t + 1\}$. To solve for the normalization, we require all the utility flow terms on the right side of (28) to cancel. After factoring out $\beta\lambda_t/\lambda_{t+1}$ we obtain:

$$\begin{aligned} 0 \quad &= \quad u_{2t+1}(x) - u_{1t+1}(x) (1 - \lambda_{t+1}) - u_{1t+1}(x + 1)\lambda_{t+1} \\ &\Longrightarrow \quad u_{2t+1}(x) - u_{1t+1}(x) = \lambda_{t+1} [u_{1t+1}(x + 1) - u_{1t+1}(x)] \end{aligned} \tag{29}$$

Appealing to (29) we substitute $\lambda_t [u_{1t}(x + 1) - u_{1t}(x)]$ for $u_{2t}(x) - u_{1t}(x)$ in (28) to obtain a recursive expression for $u_{1t}(x + 1)$ that implies:

$$u_{1t}(x+1) = u_{1t}(1) + \sum_{y=1}^{x} \frac{\psi_1[p_t(y)] - \psi_2[p_t(y)]}{\lambda_{t+1}} - \frac{\beta\psi_1[p_{t+1}(y)] (1 - \lambda_{t+1})}{\lambda_{t+1}} - \beta\psi_1[p_{t+1}(y+1)] + \frac{\beta\psi_2[p_{t+1}(y)]}{\lambda_{t+1}}$$

where $u_{1t}(1)$ can take any real value because expected utility is invariant to additive constants applied to the kernel functions $u_{1t}(x)$. Hence from (28):

$$\begin{aligned} u_{2t}(x) \quad = \quad & u_{1t}(x) + \psi_1[p_t(x)] - \psi_2[p_t(x)] - \beta\psi_1[p_{t+1}(x)] (1 - \lambda_{t+1}) \frac{\lambda_t}{\lambda_{t+1}} \\ & -\beta\psi_1[p_{t+1}(x + 1)]\lambda_t + \beta\psi_2[p_{t+1}(x)] \frac{\lambda_t}{\lambda_{t+1}} \end{aligned} \tag{30}$$

from which expressions for $U_{1t}(x_t) = u_{1t}(x_t)$ and $U_{2t}(x_t) = \lambda_t^{-1} [u_{2t}(x_t) - (1 - \lambda_t) u_{1t}(x_t)]$ directly follow.

4.5 Establishing finite dependence in games

Showing a model exhibits finite dependence becomes complicated in games because the decisions of a player today affects the decisions of the other players tomorrow. However, for many games settings the structure of the game gives a natural way of obtaining finite dependence. The structure we have considered is one in which the current action of the player does not affect the choices of the other players. The basic idea for games is to first line up the states of the other players through the period $t+1$ action and then line up the agent's state at $t+2$, assuming the agent can line up his own state in one period.

Note that $\mathcal{F}^{(i)}_{t+1}(j)$ contains transition probabilities from $t+1$ to $t+2$ given initial choice j by player i. Note also that the choice of one's competitors at $t+2$ does not depend on the player's choice at $t+2$ except through expectations over the choice conditional on the state. What we would like is that the choice at $t+2$ of one's competitors lines up the competitors' states at $t+3$. Denote $N^{\sim i}_{t+3}$ as all possible competitor states that can result from choice sequences beginning with j or j'. Denote $\mathcal{P}^{\sim i}_{t+2}$ as the transpose of the transition matrix from N^{*}_{t+2} feasible period 2 states to the $N^{\sim i}_{t+3}-1$ competitor states at $t+3$.

$$\mathcal{P}^{\sim i}_{t+2}\left[\begin{array}{cc}\mathcal{F}^{(i)}_{t+1}(j) & -\mathcal{F}^{(i)}_{t+1}(j')\end{array}\right]\left[\begin{array}{c}\mathcal{D}^{(i)}_{t+1}(j)\\ \\ \mathcal{D}^{(i)}_{t+1}(j')\end{array}\right]=\mathcal{P}^{\sim i}_{t+2}\left[F^{(i)}_{1t+1}(j'^{T}F^{(i)}_{j't}(x_t)-F^{(i)}_{1t+1}(j)^{T}F^{(i)}_{jt}(x_t)\right] \tag{31}$$

This leaves us with an $N^{\sim i}_{t+3}-1$ system of equations. If the rank of $\mathcal{P}^{\sim i}_{t+2}\left[\begin{array}{cc}\mathcal{F}^{(i)}_{t+1}(j) & -\mathcal{F}^{(i)}_{t+1}(j')\end{array}\right]=$ $N^{\sim i}_{t+3}-1$, then we have a sufficient condition for competitor states lining up at $t+2$. If we further assume that one's own state can be lined up with the period $t+2$ decision, we are done.

Example 2: A coordination game

To illustrate how finite dependence can be applied in games, we consider a simple coordination game. Each player $i\in\{1,2\}$ choose to whether or not to compete in a market at time t by setting $d^{(i)}_{2t}=1$ if competing and setting $d^{(i)}_{1t}=1$ if not. The dynamics of the game arise purely from the effect of decisions made by both players in the previous period on current payoffs; in this model $x_t=\{d^{(1)}_{2t-1},d^{(2)}_{2t-1}\}$. The non-stationarity occurs through the flow payoffs and corresponding choice probabilities rather than through the transitions on the state variables.

This model exhibits two period finite dependence. To prove this claim we find two sequences of choices by the first player, which differ in their initial choice at t, such that when the second player makes his equilibrium choices, the joint distribution of $\left(d_{t+2}^{(1)}, d_{t+2}^{(2)}\right)$ is the same for both sequences.

Per the discussion above, the first step in establishing finite dependence is choosing weights on the decisions at $t+1$ such that after the $t+2$ distribution of competitor's states will be the same across the two choice paths. In this case, there is only one competitor state variable which is whether or not the competitor will be in the market at $t+2$. Hence, the number of rows in $\mathcal{P}_{t+2}^{(2)}$ is one, implying that as long as the one of the columns of $\mathcal{P}_{t+2}^{(2)}\left[\begin{array}{cc} \mathcal{F}_{t+1}(1) & -\mathcal{F}_{t+1}(2) \end{array}\right]$ is not zero, there exists a choice path such that the expected probability of the competitor being in the market after the period $t+2$ decision in the same across the initial choice of being in or out of the market at period t. Further, we can ensure that player 1's state is the same after the $t+2$ decision by setting the $t+2$ choice for player 1 to be the same across the two paths. This has no effect on player 2's choice at $t+2$ since it is not one of player 2's state variables at $t+2$. The following theorem then establishes that a finite dependence path does indeed exist.

Theorem 7 *Finite dependence can be achieved after two periods for all x_t,*

5 Exclusion Restrictions and Functional Form Restrictions

The previous section showed that, when the problem is non-stationary and the sampling period is shorter than the time horizon, identification of some of the flow payoffs can be achieved when the problem exhibits finite dependence. In this section we consider how to restore identification in cases where finite dependence does not hold within the sample period or when the finite dependence normalizations are unattractive. The next two subsections show identification conditions when there are exclusions restrictions, variables that affect the state transitions but do not enter the flow payoffs, or when the the flow payoffs are stable. By stability of the flow payoffs we mean that the flow payoffs depend on on the state and the choice but not time itself.

5.1 Exclusion restrictions and identification of flow payoffs

When finite dependence does not hold or the number of periods necessary to achieve finite dependence is shorter than $\mathcal{T}-1$, the model is under-identified. One means of restoring identification is to impose exclusion restrictions Assuming that some variables affect the state transitions but do not enter current utility provides a means of restoring identification of the flow payoffs as well as the ex-ante value function at $\mathcal{T}$. Indeed, with exclusion restrictions, only one period of data followed by one set of transitions may result in identification of the flow payoffs in that period.

For example in the non-stationary case with three choices, we have the following system of equations:

$$\begin{bmatrix} u_{2t} \\ u_{3t} \end{bmatrix} = \begin{bmatrix} \Psi_{1t} - \Psi_{2t} \\ \Psi_{1t} - \Psi_{3t} \end{bmatrix} + \beta \begin{bmatrix} \mathcal{F}_{1t} - \mathcal{F}_{2t} \\ \mathcal{F}_{1t} - \mathcal{F}_{3t} \end{bmatrix} \begin{bmatrix} \mathcal{V}_{t+1} \\ \mathcal{V}_{t+1} \end{bmatrix} \tag{32}$$

Absent exclusion restrictions, we have a system of $2X$ equations with $3X - 1$ unknowns. Now suppose X can be partitioned into X_1 and X_2 where X_2 contains two values and where x_2 affects the state transitions but not the flow payoffs:

$$u_{jt}(x_1, x_2) = u_{jt}(x_1)$$

Letting superscripts indicate the value of x_2 at time t, we have the following system of equations:

$$\begin{bmatrix} u_{2t}^{(1)} \\ u_{2t}^{(2)} \\ u_{3t}^{(1)} \\ u_{3t}^{(2)} \end{bmatrix} = \begin{bmatrix} \Psi_{1t}^{(1)} - \Psi_{2t}^{(1)} \\ \Psi_{1t}^{(2)} - \Psi_{2t}^{(2)} \\ \Psi_{1t}^{(1)} - \Psi_{3t}^{(1)} \\ \Psi_{1t}^{(2)} - \Psi_{3t}^{(2)} \end{bmatrix} + \beta \begin{bmatrix} \mathcal{F}_{1t}^{(1)} - \mathcal{F}_{2t}^{(1)} \\ \mathcal{F}_{1t}^{(2)} - \mathcal{F}_{2t}^{(2)} \\ \mathcal{F}_{1t}^{(1)} - \mathcal{F}_{3t}^{(1)} \\ \mathcal{F}_{1t}^{(2)} - \mathcal{F}_{3t}^{(2)} \end{bmatrix} \begin{bmatrix} \mathcal{V}_{t+1} \\ \mathcal{V}_{t+1} \end{bmatrix} \tag{33}$$

which now has $4X_1$ equations and $4X_1 - 1$ unknowns. Since $u_{jt}^{(1)} = u_{jt}^{(2)}$, taking differences and rearranging terms yields:

$$\begin{bmatrix} \Psi_{2t}^{(1)} - \Psi_{1t}^{(1)} + \Psi_{1t}^{(2)} - \Psi_{2t}^{(2)} \\ \Psi_{3t}^{(1)} - \Psi_{1t}^{(1)} + \Psi_{1t}^{(2)} - \Psi_{3t}^{(2)} \end{bmatrix} = \beta \begin{bmatrix} \mathcal{F}_{1t}^{(1)} - \mathcal{F}_{2t}^{(1)} - \mathcal{F}_{1t}^{(2)} + \mathcal{F}_{2t}^{(2)} \\ \mathcal{F}_{1t}^{(1)} - \mathcal{F}_{3t}^{(1)} - \mathcal{F}_{1t}^{(2)} + \mathcal{F}_{3t}^{(2)} \end{bmatrix} \mathcal{V}_{t+1} \tag{34}$$

Hence if the rank of the $2X_1 \times 2X_1$ matrix:

$$\begin{bmatrix} \mathcal{F}_{1t}^{(1)} - \mathcal{F}_{2t}^{(1)} - \mathcal{F}_{1t}^{(2)} + \mathcal{F}_{2t}^{(2)} \\ \mathcal{F}_{1t}^{(1)} - \mathcal{F}_{3t}^{(1)} - \mathcal{F}_{1t}^{(2)} + \mathcal{F}_{3t}^{(2)} \end{bmatrix}$$

is at least $2X_1 - 1$, we can solve for the value functions. Given the value functions, we obtain the flow payoffs using (1) and (2).

5.2 Stable utility functions and identification of flow payoffs

When the panel is sufficiently short such that finite dependence cannot be achieved, stability of the flow payoffs can also restore identification. We define stability of the flow payoffs as $u_{jt}(x) = u_{jt'}(x)$ for all

$\{t, t'\}$ and for all $j \in [1, \ldots, J]$. In this case the non-stationarity comes from either the state transitions or the time horizon. Identification is achieved by solving for both the flow payoffs and the value functions in the last period, similar to Section 5.1.

To illustrate the nature of identification in the case where there are incomplete histories, suppose there are only two choices each period, and the data covers two periods, t and $t+1$. We now assume $u_{2t}(x) = u_{2t+1}(x) = u_2(x)$ for all $x \in \{1, \ldots, X\}$ and adopt the normalization $u_{1t}(x) = u_{1t+1}(x) = 0$.

We can express u_2, the vector of flow payoff for action 2 in every state, relative to choosing action 1 in the next period. Similarly, we can also express u_2 relative to choosing action 1 in the next period and in the period after that. Hence, given conditional choice probabilities in two periods, t and $t+1$, we can express u_2 as:

$$u_2 = \Psi_{1t+1} - \Psi_{2t+1} + \beta(\mathcal{F}_{1t+1} - \mathcal{F}_{2t+1})\mathcal{V}_{t+2} \tag{35}$$

$$u_2 = \Psi_{1t} - \Psi_{2t} + \beta(\mathcal{F}_{1t} - \mathcal{F}_{2t})\Psi_{1t+1} + \beta^2(\mathcal{F}_{1t} - \mathcal{F}_{2t})\mathcal{F}_{1t+1}\mathcal{V}_{t+2} \tag{36}$$

Taking differences and rearranging terms yields:

$$\Psi_{1t} - \Psi_{2t} - \Psi_{1t+1} + \Psi_{2t+1} + \beta(\mathcal{F}_{1t} - \mathcal{F}_{2t})\Psi_{1t+1} = \beta(\mathcal{F}_{1t+1} - \mathcal{F}_{2t+1} - \beta\mathcal{F}_{1t}\mathcal{F}_{1t+1} + \beta\mathcal{F}_{2t}\mathcal{F}_{1t+1})\mathcal{V}_{t+2} \tag{37}$$

Since adding a constant to the future value terms does not affect choice probabilities, we only need to identify $\mathcal{V}_{t+2}$ up to a constant, implying we need the rank of

$$(\mathcal{F}_{1t+1} - \mathcal{F}_{2t+1} - \beta\mathcal{F}_{1t}\mathcal{F}_{1t+1} + \beta\mathcal{F}_{2t}\mathcal{F}_{1t+1})$$

to be $X-1$ to identify the $X-1$ differenced value functions. Note that the rank condition is on differences in state transition matrices–it does not depend on the conditional choice probabilities or the flow payoffs–implying that the rank condition is straightforward to check given reasonably-sized problems.

5.3 Alternative normalizations

Even when the problem satisfies finite dependence, the normalizations required to achieve identification may be unattractive. Clearly imposing exclusion restrictions or stability of the flow payoffs shrinks the number of parameters that need to be identified. In the latter case, standard normalizations such as normalizing the payoffs for one of the choices to be zero in all states can be used to achieve identification of $(J-1)\times X$ flow payoffs. Note that the problem may still be non-stationary through the state transitions or through the time horizon. Here we are working in a middle ground between the fully non-stationary case, where both the state transitions and the flow payoffs vary with time, and the infinite horizon stationary case, where both the state transitions and flow payoffs are constant over time.

Example 3: Revisiting non-stationary search

To illustrate how stability allows us to use normalizations on the flow payoffs besides those dictated by the finite dependence path, we return to the non-stationary search example in Section 4.4. Recall that one of the finite dependence paths entailed searching for work in the first period and not searching in the second, while the other entailed not searching for work in the first period and mixing in the second such that the expected accumulation of human capital was the same across the two paths. Substituting in for $u_{2t}(x)$ with (27) in (28) and simplifying yields:

$$\begin{aligned}\lambda_t[U_{2t}(x)-U_{1t}(x)] \quad = \quad & \psi_1[p_t(x)]-\psi_2[p_t(x)]+\beta\lambda_t\left[U_{2t+1}(x)-U_{1t+1}(x+1)\right] \\ & +\frac{\beta\lambda_t}{\lambda_{t+1}}\left\{\psi_2[p_{t+1}(x)]-\psi_1[p_{t+1}(x)]\right\}+\beta\lambda_t\left\{\psi_1[p_{t+1}(x)]-\psi_1[p_{t+1}(x+1)]\right\}\end{aligned} \tag{38}$$

If we assume that i) the flow payoffs do not depend on t and that ii) the flow payoff for not working is zero in all states, $U_1(x)=0$ for all x, then the (38) simplifies to:

$$\begin{aligned}U_{2t}(x) \quad = \quad & \frac{\psi_1[p_t(x)]-\psi_2[p_t(x)]}{\lambda_t(1-\beta)}+\frac{\beta}{(1-\beta)\lambda_{t+1}}\left\{\psi_2[p_{t+1}(x)]-\psi_1[p_{t+1}(x)]\right\} \\ & +\frac{\beta}{(1-\beta)}\left\{\psi_1[p_{t+1}(x)]-\psi_1[p_{t+1}(x+1)]\right\}\end{aligned} \tag{39}$$

Note that $U_2(x)$ is heavily over-identified as (39) must hold for all t.

6 Unbundling

In the games context, the previous sections focused on the assumptions necessary to identify ex-ante flow payoffs. But these are not the primitives of the game and hence the assumptions made in order to achieve identification may also not be on the primitives. As in Bajari et al. (2009), we focus on recovering ex-post flow payoffs, the primitives of the game. Here, however, we focus on how non-stationarity aids in the recovery of the ex-post flow payoffs. We first consider a case where non-stationarity occurs through the state transitions but the flow payoff function is stable. Next, we provide an example where the flow payoff function is non-stationary but where unbundling can still be achieved.

6.1 Stable utility functions and unbundling in games

By definition utility functions are stable over time in stationary environments, so assuming a stationary environment automatically eliminates the prospect of identifying game primitives through technological change. An alternative approach is to maintain the assumptions that utility functions are stable over time, but relax the assumption that state transition probabilities are stationary. For example innovation and growth in technology might reveal features of firms' payoff functions through their responses to evolving business conditions.

Example 4: An entry/exit game

To illustrate how nonstationarity aids in the recovery of flow payoffs, we consider an entry/exit game. The example exploits the terminal state property to recover $u_{jt}^{(i)}(x_t)$, normalize the utility from the exit choice as $U_1^{(i)}\left(x_t, d_t^{(\sim i)}\right) = 0$ and identify utility from the entry choice $U_{2t}^{(i)}\left(x_t, d_t^{(\sim i)}\right)$ by assuming the flow utility from entry (and remaining in the industry) does not depend on time; formally $U_{2t}^{(i)}\left(x_t, d_t^{(\sim i)}\right) = U_2^{(i)}\left(x_t, d_t^{(\sim i)}\right)$ for $t < \mathcal{T}$. Markets can have at most two firms. An incumbent firm can choose to remain in the market or exit. Exit is a terminal choice. An exiting firm is replaced by a potential entrant in the next period who faces the choices: remain in the market (enter) or exit. Let $d_{jt}^{(i)} = 1$ if action j is taken by player i at time t and is zero otherwise. Label exit as action 1 and entry as action 2. The time horizon is infinite.

The flow payoff of exiting is normalized to $\epsilon_{1t}^{(i)}$, a transitory shock that is private information to player i. Since it is a terminal choice, there are no future payoffs for exiting. Current period payoffs for entering or remaining in the market depend on three state variables: (i) whether there is another firm in the market $d_{2t}^{(\sim i)}$, (ii) whether the firm is an incumbent and therefore does not have to pay the entry cost, d_{2t}^{i}, and (iii) a discrete market state variable $x_{1t} \in X$ with state transitions given by $f_t(x_{1t+1}|x_{1t}){>}0$ for all $x_{1t+1} \in X$. Note that the transitions on the market state variable depends on time which is what makes the model non-stationary.

Conditional on the other player's action, the flow payoff for i at time t for entering the market is $U_2^{(i)}(d_{2t}^{(\sim i)}, d_{t-1}^{(i)}, x_{1t})$. The expected payoffs of entering depends on the $x_t \equiv \left\{d_{2t-1}^{(\sim i)}, d_{2t-1}^{(i)}, x_{1t}\right\}$. It is then defined as:

$$u_{2t}^{(i)}(x_t) = \sum_j p_{jt}^{(\sim i)}(x_t) U_2^{(i)}(j, d_{2t-1}^{(i)}, x_{1t}) \tag{40}$$

The total expected payoff for taking action 2 are then given by $u_{2t}^{(i)}(x_t) + \epsilon_{2t}^{(i)}$ where $\epsilon_{2t}^{(i)}$ is a transitory shock to the payoff for action 2 that is private information to player i.

Given exit is a terminal choice, we can express the conditional value function for entering the market as:

$$\begin{aligned} v_{2t}^{(i)}(x_t) &= u_{2t}^{(i)}(x_t) + \beta \sum_j \sum_{x_{1t+1}} p_{jt}^{(\sim i)} V_{t+1}^{(i)}(j, 1, x_{1t+1}) f_t(x_{1t+1}|x_{1t}) \\ &= u_{2t}^{(i)}(x_t) + \beta \sum_j \sum_{x_{1t+1}} p_{jt}^{(\sim i)} \psi_2 \left[p_{t+1}^{(i)}(j, 1, x_{1t+1})\right] f_t(x_{1t+1}|x_{1t}) \end{aligned} \tag{41}$$

implying we can express $u_{2t}^{(i)}(x_t)$ as:

$$u_{2t}^{(i)}(x_t) = \psi_{1t}\left[p_t^{(i)}(x_t)\right] - \psi_{2t}\left[p_t^{(i)}(x_t)\right] - \beta \sum_j \sum_{x_{1t+1}} p_{jt}^{(\sim i)} \psi_2 \left[p_{t+1}^{(i)}(j, 1, x_{1t+1})\right] f_t(x_{1t+1}|x_{1t}) \tag{42}$$

Note that under our assumptions everything on the right hand side of (42) is known. Substituting in on the left hand side with (40) yields:

$$\sum_j p_{jt}^{(\sim i)}(x_t) U_2^{(i)}(j, d_{2t-1}^{(i)}, x_{1t}) = \psi_{1t}\left[p_t^{(i)}(x_t)\right] - \psi_{2t}\left[p_t^{(i)}(x_t)\right] - \beta \sum_j \sum_{x_{1t+1}} p_{jt}^{(\sim i)} \psi_2 \left[p_{t+1}^{(i)}(j, 1, x_{1t+1})\right] f_t(x_{1t+1}|x_{1t}) \tag{43}$$

There are then two unknowns on the left hand side of equation (43). By evaluating this expression at a particular value of x_t and then using those same values just in a different time period, we obtain two equations and two unknowns. The following theorem then establishes identification of the $U_2^{(i)}$'s:

Theorem 8 *Given a known distribution for ϵ where ϵ is independent across players and time, β, $u_1^{(i)}(x) = \epsilon_{1t}^{(i)}$, and $p_{2t}^{(i)}(x_t) \neq p_{2t+1}^{(i)}(x_t)$, then $U_2^{(i)}\left(j, d_{2t-1}^{(i)}, x_{1t}\right)$ is identified for all j if $\mathcal{T} \geq 2$.*

6.2 Exclusion restrictions and alternative normalizations

While allowing complete flexibility regarding the role of non-stationarity in ex-post payoffs makes unbundling impossible, some structure on how time affects ex-post flow payoffs can be incorporated. By placing structure on the model, we are also able to relax normalizations that were required in the finite dependence case but might seem undesirable. For example, if we have information from another source on the value of particular alternatives at particular points in time and this information may not correspond to the finite dependence normalizations needed for identification. The next example illustrates how exclusion restrictions permit identification of non-stationary flow payoffs as well as achieving identification of the flow payoffs under an alternative normalization.

Example 5: Revisiting the coordination game

To see how exclusion restrictions can be used in conjunction with finite dependence to recover ex-post payoffs, we return to the coordination game described in Example 2. We showed that this model exhibited finite dependence so by Theorem 5 we can recover ex-ante flow payoffs given a particular normalization using three periods of data. Here we work with a different set of normalizations and use exclusion restrictions to unbundle the ex-ante payoffs to recover the ex-post payoffs.

We normalize the systematic flow payoff from not competing to zero. Thus $U_{1t}^{(i)}(x_t, d_t^{(\sim i)}) = 0$ for all four values of x_t, both choices of the rival player $d_t^{(\sim i)}$, and all time periods t. We also impose the exclusion restriction:

$$U_{2t}^{(i)}(x_t, d_t^{(\sim i)}) \equiv U^{(i)}(x_t, d_t^{(\sim i)}) + \theta_t \tag{44}$$

such that the effect of time on the flow payoffs of action 2 does not depend on the state or the other player's action. The ex-ante payoff is then given by:

$$u_{2t}^{(1)}(x) \equiv \sum_{j=0}^{1} p_{jt}^{(2)}(x) U_2^{(1)}(x, j) + \theta_t \tag{45}$$

with $u_{1t}^{(i)}(x) \equiv 0$.

Given values of $p_{jt}^{(i)}(x)$ and $u_{2t}^{(i)}(x)$ observed over $\mathcal{T}$ different time periods for the 4 values of the state variables $\left(d_{t-1}^{(1)}, d_{t-1}^{(2)}\right)$, Equation (45) represents a a linear system of $4\mathcal{T}$ equations for each player, with $\mathcal{T}+8$ unknown variables comprising the time effects θ_t plus the 8 values $U_2^{(1)}(x, j)$ can take for different combinations of (x, j). We then have the following result:

Theorem 9 *Subject to a rank condition, the ex-post payoffs absent the θ_t's can be recovered for $\mathcal{T} \geq 3$. Further, the θ_t's can be recovered up until period $\mathcal{T} - 1$.*

7 Conclusion

This paper establishes conditions for identifying dynamic discrete choice models, both for long panels where the sample period covers the full time horizon or the model is stationary, and for short panels where the sample period is shorter than the time horizon. For a known disturbance structure and discount factor, dynamic discrete choice models of individual optimization are identified up to any normalization of one choice-specific flow payoff for each period in each state when the model is non-stationary and the panel is long.[10] We also show that, with one significant exception, the same models are not identified from short panels without imposing exclusion and functional form restrictions on preferences that reduce the size of the parameter space. The exception to this rule is that a subset of the flow payoffs may be identified in the class of models that exhibit finite dependence. This class is, however, identified only for particular normalizations, not all normalizations. When the panel is short and finite dependence fails, or is associated with a unacceptable normalization, the models of individual optimization are under-identified by the number of elements in the state space. Our analysis carries over directly to ex-ante payoffs in noncooperative games.

Identification precedes consistency. Thus our results on identification provide a set of general conditions that fill a gap in the literature for proving maximum likelihood estimators of dynamic discrete choice models are consistent.[11] As we have explained in the text, the equations characterizing identification also yield asymptotically efficient estimators of the unrestricted utility flows in models exhibiting finite dependence that are surprisingly easy to compute.

Our analysis has ramifications for the empirical implementation of structural models of discrete choice dynamic optimization and equilibrium strategies. Our results suggest that in nonstationary settings relatively few restrictions on the role of time dependence, such as assuming just two periods in the life-cycle have the same utility flow as a mapping from the state variables, suffice to identify the remaining flows

[10] In the stationary case, flow payoffs are identified given a normalization of one choice-specific flow payoff where the normalization can vary by state.

[11] The conditions also apply to simulation and indirect inference estimators that are based on the score, Aguirregaberia and Mira (2002) swapping nested fixed point estimators that are asymptotically equivalent to maximum likelihood, as well as the estimator of Keane and Wolpin (1994), subject to the proviso that approximation errors introduced for computational tractability are eliminated as the sample grows.

in the sample as well as the continuation value for sample respondents at the end of the panel. Equivalently finite dependence coupled to an appropriate normalization can achieve the same ends. The upshot is that almost all empirical models of life-cycle behavior in labor economics are overidentified by many degrees of freedom, because of strong stability assumptions made on utility both during and after the sample ends. For example, all full solution and nested estimation methods assume the same functional form for the utility flows in every period. The overidentifying restrictions offer scope for specification testing and relaxing the assumption of a constant subjective discount factor as well as permitting flexibility of the disturbance structure.

The implications for model specification in industrial organization are just as striking. Structural empirical work in this field leans heavily on the assumption of stationarity, primarily because the panel is necessarily short with infinitely-lived firms. By exploiting finite dependence, however, much can still be recovered. Time dependent state transition processes arise very naturally in industrial organization and, rather than being a barrier to identification and estimation, may actually be useful in unbundling ex-post flow payoffs.

A Proofs

Proof of Theorem 1. The proof follows directly from Arcidiacono and Miller (2011) since there is no use of the fact that the weights are probabilities. ■

Proof of Theorem 2.

It is convenient to prove the finite horizon and stationary cases separately, the nonstationary case first. Let $l(x,t) \in \{1,\ldots,J\}$ and $c_t(x)$ respectively denote the normalizing action and benchmark flow utility the associated with (t,x). We set $u^*_{l(x,t),t}(x) = c_t(x)$ and for all $j \neq l(x,t)$ define:

$$u^*_{jT}(x) \equiv u_{jT}(x) - u_{l(x,T),T}(x) + c_T(x) \tag{46}$$

and:

$$u^*_{jt}(x) = u_{jt}(x) + c_t(x) - u_{l(x,t)t}(x) + \sum_{\tau=t+1}^{T} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \left[u^*_{1\tau}(x_\tau) - u_{1\tau}(x_\tau)\right] \left[\kappa^*_{\tau-1}(x_\tau|x_t, l(x,t)) - \kappa^*_{\tau-1}(x_\tau|x_t, j)\right] \tag{47}$$

In the final period T, supposing $x_T = x$ the agent optimally sets $d_{jT} = 1$ if:

$$u_{jT}(x) + \epsilon_{jT} \geq \max_{k \in \{1,\ldots,J\}} \{u_{kT}(x) + \epsilon_{kT}\}$$

inequalities that are satisfied if and only if:

$$u^*_{jT}(x) + \epsilon_{jT} \geq \max_{k \in \{1,\ldots,J\}} \{u^*_{kT}(x) + \epsilon_{kT}\}$$

as required by the theorem, and establishing the result for $T = 1$.

For the representation of $v_{jt}(x_t)$ provided by (7), set $d^*_{1\tau}(x_\tau, k) = 1$ for all $\tau = \{t+1, \ldots, T\}$ $k \in \{1, \ldots, J\}$ and $x_\tau \in \{1, \ldots, X\}$. Supposing $x_t = x$ in period t, the decision maker optimally sets $d_{jt} = 1$ if:

$$j = \underset{k \in \{1,\ldots,J\}}{\arg\max} \left\{ u_{kt}(x) + \epsilon_{kt} + \sum_{\tau=t+1}^{T} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \left[u_{1\tau}(x_\tau) + \psi_1[p_\tau(x_\tau)] \right] \kappa^*_{\tau-1}(x_\tau | x_t, k) \right\}$$

Subtracting the constant:

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

does not change the optimal choice, so $d_{jt} = 1$ is optimal if:

$$j = \underset{k \in \{1,\ldots,J\}}{\arg\max} \left\{ u_{kt}(x) - u_{l(x,t)t}(x) + \epsilon_{kt} + \sum_{\tau=t+1}^{T} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \left\{ \begin{array}{c} u_{1\tau}(x_\tau) \left[\kappa^*_{\tau-1}(x_\tau | x_t, k) - \kappa^*_{\tau-1}(x_\tau | x_t, l(x,t)) \right] \\ + \psi_k[p_\tau(x_\tau)] \kappa^*_{\tau-1}(x_\tau | x_t, k) \end{array} \right\} \right\} \tag{48}$$

From (47):

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

Substitute the second line into the maximand of (48). Then $d_{jt} = 1$ is optimal if:

$$\begin{aligned} j &= \underset{k \in \{1,\ldots,J\}}{\arg\max} \left\{ u^*_{kt}(x) - c_t(x) + \epsilon_{kt} + \sum_{\tau=t+1}^{T} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \left\{ \begin{array}{c} u^*_{1\tau}(x_\tau) \left[\kappa^*_{\tau-1}(x_\tau | x_t, k) - \kappa^*_{\tau-1}(x_\tau | x_t, l(x,t)) \right] \\ + \psi_k[p_\tau(x_\tau)] \kappa^*_{\tau-1}(x_\tau | x_t, k) \end{array} \right\} \right\} \\ &= \underset{k \in \{1,\ldots,J\}}{\arg\max} \left\{ u^*_{kt}(x) + \epsilon_{kt} + \sum_{\tau=t+1}^{T} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \left[u^*_{1\tau}(x_\tau + \psi_k[p_\tau(x_\tau)] \right] \kappa^*_{\tau-1}(x_\tau | x_t, k) \right\} \end{aligned}$$

as required, where the last line follows because the dropped terms do not depend on the choice. This proves the result for all finite T.

We now turn to infinite horizon stationary models. We start by defining a c_x an i_x and a u_{i_x} for each x analogously to the finite horizon case and set:

$$u_j^*(x) = u_j(x) + c_x - u_{i_x}(x) + \sum_{\tau=1}^{\infty} \sum_{x_\tau=1}^{X} \beta^\tau \left[u_1^*(x_\tau) - u_1(x_\tau)\right] \left[\kappa_{\tau-1}^*(x_\tau|x_t, l(x,t)) - \kappa_{\tau-1}^*(x_\tau|x_t, j)\right]$$

or in matrix notation:

$$u_j^* = u_j + c - \tilde{u} + \beta \left(F_1 - F_j\right) \left[\mathcal{I} - \beta F_1\right]^{-1} \left(u_1^* - u_1\right)$$

which is the result in the text. ■

Proof of Theorem 3. Substituting in for $v_{jt}(z_t) - v_{1t}(z_t)$ in (12) with the corresponding expression in (13) implies:

$$\psi_1[p_t(z_t)] - \psi_j[p_t(z_t)] = u_{jt}(z_t) + \sum_{\tau=t+1}^{T} \sum_{z_\tau=1}^{Z} \beta^{\tau-t} \psi_1[p_\tau(z_\tau)] \left[\kappa_{\tau-1}^*(z_\tau|z_t, j) - \kappa_{\tau-1}^*(z_\tau|z_t, 1)\right]$$

Solving for $u_{jt}(z_t)$ completes the first part of the theorem:

$$u_{jt}(z_t) = \psi_1[p_t(z_t)] - \psi_j[p_t(z_t)] + \sum_{\tau=t+1}^{T} \sum_{z_\tau=1}^{Z} \beta^{\tau-t} \psi_1[p_\tau(z_\tau)] \left[\kappa_{\tau-1}^*(z_\tau|z_t, 1) - \kappa_{\tau-1}^*(z_\tau|z_t, j)\right] \tag{49}$$

To prove the second part, note that the two decision sequences set the initial choices such that $d_{jt} = 1$ or $d_{1t} = 1$ and then both decision sequences set $d_{1t'} = 1$ for all $t' > t$. From the definition of F_1, the columns of F_1^τ gives the probabilities of being in each state after τ periods conditional choosing alternative 1 in each of those periods. The rows indicate how these probabilities differ given the initial state. Hence, for $\tau \geq 1$, the (z, z') element of F_1^τ is $\kappa_{t+\tau-1}^*(z'|z, 1)$. Similarly, the (z, z') element of $F_j F^\tau$ is $\kappa_{t+\tau-1}^*(z'|z, j)$.

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

$$u_j = \Psi_j - \Psi_1 + \sum_{\tau=1}^{\infty} \beta^\tau \left(F_1 - F_j\right) F_1^{\tau-1} \Psi_1 = \Psi_j - \Psi_1 + \beta \left(F_1 - F_j\right) \left(\sum_{\tau=0}^{\infty} \beta^\tau F_1^\tau\right) \Psi_1 \tag{50}$$

Noting that $\beta f_j(z'|z) > 0$ for all (j, z, z') and $\beta \sum_{z'=1}^{Z} f_j(z'|z) = \beta < 1$ for all (j, z), the existence of $\left[\mathcal{I} - \beta F_1\right]^{-1}$ follows from Hadley (page 118, 1961) with:

$$Q \equiv \sum_{\tau=0}^{\infty} \beta^\tau F_1^\tau = \mathcal{I} + \beta Q F_1 = \left[\mathcal{I} - \beta F_1\right]^{-1}$$

Substituting the expression for Q into (50) we obtain:

$$u_j = \Psi_j - \Psi_1 + \beta \left(F_1 - F_j\right) \left[\mathcal{I} - \beta F_1\right]^{-1} \Psi_1$$

which proves the theorem. ■

Proof of Theorem 4. Note that if the V's were known we would be exactly identified. ■

Proof of Theorem 5. First we demonstrate by backwards induction that the equations characterizing finite dependence, (20) and (21) , fully capture the empirical content of the model. That is, if the sample was drawn from whole population then, given F and G repeatedly applying (21) can be used to identify u up to a normalization.

Next we define $\widehat{u}$, an element in the parameter space, which induces the same behavior as the data generating process from u as:

$$\widehat{u}_{1t}(x) = \left\{ \begin{array}{c} \psi_1[p_t(x)] - \psi_j[p_t(x)] \\ -\sum_{\tau=t+1}^{t+\rho} \sum_{i=1}^{J} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \psi_i[p_\tau(x_\tau)] \Delta d_{i\tau}^*(x_\tau|j) \kappa_{\tau-1}^*(x_\tau|j) \end{array} \right\} \frac{\Delta d_{j\tau}^*(x|j)\kappa_{\tau-1}^*(x|j)}{\Delta d_{1\tau}^*(x|j)\kappa_{\tau-1}^*(x|j) + \Delta d_{j\tau}^*(x|j)\kappa_{\tau-1}^*(x|j)} \tag{51}$$

and:

$$\begin{aligned} \widehat{u}_{jt}^*(x) &= u_{1t}^*(x) + \psi_1[p_t(x_t)] - \psi_j[p_t(x_t)] \\ &\quad + \sum_{\tau=t+1}^{t+\rho} \sum_{i=1}^{J} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \psi_i[p_\tau(x_\tau)] \Delta d_{i\tau}^*(x_\tau|j) \kappa_{\tau-1}^*(x_\tau|j) \end{aligned} \tag{52}$$

where:

$$\Delta d_{i\tau}^*(x_\tau|j)\kappa_{\tau-1}^*(x_\tau|j) = d_{i\tau}^*(x_\tau|j)\kappa_{\tau-1}^*(x_\tau|j) - d_{i\tau}^*(x_\tau|k)\kappa_{\tau-1}^*(x_\tau|1)$$

First we show that $\widehat{u}$ satisfies the equations characterizing finite dependence.

Substituting (52) into (51) we obtain:

$$\widehat{u}_{1t}(x) = \frac{\left[u_{jt}^*(x) - u_{1t}^*(x)\right] \Delta d_{j\tau}^*(x|j)\kappa_{\tau-1}^*(x|j)}{\Delta d_{1\tau}^*(x|j)\kappa_{\tau-1}^*(x|j) + \Delta d_{j\tau}^*(x|j)\kappa_{\tau-1}^*(x|j)}$$

which upon rearrangement, and then summing over τ and x_τ yields:

$$\begin{aligned} &\sum_{\tau=t+1}^{t+\rho} \sum_{i=2}^{J} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \left[u_{i\tau}(x_\tau) - u_{1\tau}(x_\tau)\right] \Delta d_{i\tau}^*(x_\tau|j)\kappa_{\tau-1}^*(x_\tau|j) = \\ &\sum_{\tau=t+1}^{t+\rho} \sum_{x_\tau=1}^{X} \beta^{\tau-t} u_{1\tau}(x_\tau) \left[\Delta d_{1\tau}^*(x_\tau|j)\kappa_{\tau-1}^*(x_\tau|j) + \Delta d_{2\tau}^*(x_\tau|j)\kappa_{\tau-1}^*(x_\tau|j)\right] \end{aligned}$$

Appealing to the definition of finite dependence given by (20) it follows from (21) that for all parameterizations u generating the data:

$$\begin{aligned}
u_{jt}(x_t) - u_{1t}(x_t) &= \psi_1[p_t(x_t)] - \psi_j[p_t(x_t)] + \sum_{\tau=t+1}^{t+\rho} \sum_{i=1}^{J} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \psi_i[p_\tau(x_\tau)] \Delta d^*_{i\tau}(x_\tau|j) \kappa^*_{\tau-1}(x_\tau|j) \\
&+ \sum_{\tau=t+1}^{t+\rho} \sum_{i=2}^{J} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \left[u_{i\tau}(x_\tau) - u_{1\tau}(x_\tau)\right] \Delta d^*_{i\tau}(x_\tau|j) \kappa^*_{\tau-1}(x_\tau|j) \\
&+ \sum_{\tau=t+1}^{t+\rho} \sum_{x_\tau=1}^{X} \beta^{\tau-t} u_{1\tau}(x_\tau) \left[\Delta d^*_{1\tau}(x_\tau|j) \kappa^*_{\tau-1}(x_\tau|j) + \Delta d^*_{2\tau}(x_\tau|j) \kappa^*_{\tau-1}(x_\tau|j)\right] \\
&= \psi_1[p_t(x_t)] - \psi_j[p_t(x_t)] \\
&+ \sum_{\tau=t+1}^{t+\rho} \sum_{i=1}^{J} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \left\{u_{i\tau}(x_\tau) + \psi_i[p_\tau(x_\tau)]\right\} \Delta d^*_{i\tau}(x_\tau|j) \kappa^*_{\tau-1}(x_\tau|j)
\end{aligned}$$

This demonstrates the existence of a normalization $\widehat{u}$ satisfying the equations that characterize finite dependence for the data generating process. ■

Proof of Theorem 6. Define $\mathcal{K}_\tau(j, x_t)$ as an $N^*_{\tau+1}(j, x_t)$ vector containing the probabilities of transitioning to each of the $N^*_{\tau+1}(j, x_t)$ attainable states given the choice sequence beginning with j and state x_t. Denote $D^*_{k\tau+1}(j)$ as a vector giving the weight placed on choice $k \in [1, \ldots, J]$ for each of the $N_{\tau+1}(j)$ possible states at $t+1$. Let $\mathcal{D}_{\tau+1}(j)$ be a $(J-1)N_{\tau+1}(j, x_t)$ vector defined by:

$$\mathcal{D}_{\tau+1}(j) = \begin{bmatrix} D^*_{2\tau+1}(j, x_t) \circ \mathcal{K}_\tau(j, x_t) \\ \vdots \\ D^*_{k\tau+1}(j, x_t) \circ \mathcal{K}_\tau(j, x_t) \\ \vdots \\ D^*_{J\tau+1}(j) \circ \mathcal{K}_\tau(j, x_t) \end{bmatrix}$$

where $\circ$ refers to element-by-element multiplication.

Denote $F_{k\tau+1}(j)$ as an $N_{\tau+1}(j) \times (N^*_{\tau+2} - 1)$ which gives the probability of transitioning from each of the $N_{\tau+1}(j)$ attainable states given initial choice j to the $N^*_{\tau+2} - 1$ attainable states at $\tau + 2$ given *either* initial choice j or j'. Define $\mathcal{F}_{\tau+1}(j)$ as an $(N^*_{\tau+2} - 1) \times ((J-1)N_{\tau+1}(j))$ matrix given by:

$$\mathcal{F}_{\tau+1}(j) = \begin{bmatrix} F_{2\tau+1}(j) - F_{1\tau+1}(j) \\ \vdots \\ F_{k\tau+1}(j) - F_{1\tau+1}(j) \\ \vdots \\ F_{J\tau+1}(j) - F_{1\tau+1}(j) \end{bmatrix}^T$$

The $N^*_{\tau+2}-1$ system of equations we need to solve can be expressed as:

$$\begin{bmatrix} \mathcal{F}_{\tau+1}(j,x_t) & -\mathcal{F}_{\tau+1}(j',x_t) \end{bmatrix} \begin{bmatrix} \mathcal{D}_{\tau+1}(j,x_t) \\ \mathcal{D}_{\tau+1}(j',x_t) \end{bmatrix} = F_{1\tau+1}(j',x_t)^T \mathcal{K}_\tau(j',x_t) - F_{1\tau+1}(j,x_t)^T \mathcal{K}_\tau(j,x_t) \tag{53}$$

implying that if the rank of $\begin{bmatrix} \mathcal{F}_{\tau+1}(j,x_t) & -\mathcal{F}_{\tau+1}(j',x_t) \end{bmatrix}$ is $N^*_{\tau+2}-1$ then finite dependence holds at period $\tau-t+1$. ■

Proof of Theorem 7. We can establish that a finite dependence path exists by showing that the rank of:

$$\mathcal{P}^2_{t+2} \begin{bmatrix} \mathcal{F}^{(1)}_{t+1}(2) & -\mathcal{F}^{(1)}_{t+1}(1) \end{bmatrix}$$

is one.

We begin by defining the terms in the above expression:

$$\mathcal{P}^{(2)}_{t+2} = \begin{bmatrix} p^{(2)}_{2t+2}(2,2) & p^{(2)}_{2t+2}(2,1) & p^{(2)}_{2t+2}(1,2) & p^{(2)}_{2t+2}(1,1) \end{bmatrix} \tag{54}$$

$$\begin{bmatrix} \mathcal{F}^{(1)}_{t+1}(2) & -\mathcal{F}^{(1)}_{t+1}(1) \end{bmatrix} = \begin{bmatrix} p^{(2)}_{2t+1}(2,2) & p^{(2)}_{2t+1}(2,1) & -p^{(2)}_{2t+1}(1,2) & -p^{(2)}_{2t+1}(1,1) \\ p^{(2)}_{1t+1}(2,2) & p^{(2)}_{1t+1}(2,1) & -p^{(2)}_{1t+1}(1,2) & -p^{(2)}_{1t+1}(1,1) \\ -p^{(2)}_{2t+1}(2,2) & -p^{(2)}_{2t+1}(2,1) & p^{(2)}_{2t+1}(1,2) & p^{(2)}_{2t+1}(1,1) \\ -p^{(2)}_{1t+1}(2,2) & -p^{(2)}_{1t+1}(2,1) & p^{(2)}_{1t+1}(1,2) & p^{(2)}_{1t+1}(1,1) \end{bmatrix} \tag{55}$$

These terms will then multiply:

$$\begin{bmatrix} D^{(1)}_{t+1}(2,x_t) \\ D^{(1)}_{t+1}(1,x_t) \end{bmatrix} = \begin{bmatrix} D^*_{2t+1}(2,2)p^{(2)}_{2t}(x_t) \\ D^*_{2t+1}(2,1)p^{(2)}_{1t}(x_t) \\ D^*_{2t+1}(1,2)p^{(2)}_{2t}(x_t) \\ D^*_{2t+1}(1,1)p^{(2)}_{1t}(x_t) \end{bmatrix} \tag{56}$$

Since the D^*_{2t+1}'s are weights on choices, we can set the weights on $D^*_{2t+1}(1,2)$ and $D^*_{2t+1}(1,1)$ to zero. Now consider the other two weights. Multiplying the matrices and rearranging terms yields the following expression:

$$\begin{aligned} &D^*_{2t+1}(2,2)p^{(2)}_{2t}(x_t)\left(p^{(2)}_{2t+2}(2,1) - p^{(2)}_{2t+2}(1,1) + p^{(2)}_{2t+1}(2,2)\left[p^{(2)}_{2t+2}(2,2) + p^{(2)}_{2t+2}(1,1) - p^{(2)}_{2t+2}(2,1) - p^{(2)}_{2t+2}(1,2)\right]\right) \\ &+D^*_{2t+1}(2,1)p^{(2)}_{2t}(x_t)\left(p^{(2)}_{2t+2}(2,1) - p^{(2)}_{2t+2}(1,1) + p^{(2)}_{2t+1}(2,1)\left[p^{(2)}_{2t+2}(2,2) + p^{(2)}_{2t+2}(1,1) - p^{(2)}_{2t+2}(2,1) - p^{(2)}_{2t+2}(1,2)\right]\right) \end{aligned}$$

Note that the expression multiplying each of the D^*_{2t+1}'s are the same except for the weights on the terms in brackets. Since we have assumed all the states are relevant for the decision, then the term multiplying $D^*_{2t+1}(2,2)$ and the term multiplying $D^*_{2t+1}(2,1)$ cannot both be zero. Hence, there exist decision weights at $t+1$ such that the probability of each of player 2's states is the same on both choice paths. Since player 1's state will be the same if the same action is chosen on each path at period $t+2$, the theorem is proved.

■

Proof of Theorem 8. Denote $P^{(\sim i)}$ as a 2×2 matrix given by:

$$P^{(\sim i)} = \begin{bmatrix} p^{(\sim i)}_{1t}(x) & p^{(\sim i)}_{2t}(x) \\ p^{(\sim i)}_{1t+1}(x) & p^{(\sim i)}_{2t+1}(x) \end{bmatrix} \tag{57}$$

Noting that x_t provides all the relevant state variables expect for the choice of the competitors, define $U^{(i)}_2$ as:

$$U^{(i)}_2 = \begin{bmatrix} U^{(i)}_2(1,x) \\ U^{(i)}_2(2,x) \end{bmatrix} \tag{58}$$

Finally, define A as:

$$A = \begin{bmatrix} \psi_{1t}\left[p^{(i)}_t(x)\right] - \psi_{2t}\left[p^{(i)}_t(x)\right] - \beta\sum_j\sum_{x_{1t+1}} p^{(\sim i)}_{jt}\psi_2\left[p^{(i)}_{t+1}(j,1,x_{1t+1})\right] f_t(x_{1t+1}|x) \\ \psi_1\left[p^{(i)}_{t+1}(x)\right] - \psi_2\left[p^{(i)}_{t+1}(x)\right] - \beta\sum_j\sum_{x_{1t+2}} p^{(\sim i)}_{jt+1}\psi_2\left[p^{(i)}_{t+2}(j,1,x_{1t+2})\right] f_{t+1}(x_{1t+2}|x) \end{bmatrix} \tag{59}$$

The system of equation is then:

$$P^{(\sim i)}U^{(i)}_2 = A \tag{60}$$

Since by assumption the choice probabilities vary between t and $t+1$, the rank of $P^{(\sim i)}$ is two, implying we can invert $P^{(\sim i)}$ and solve for $U^{(i)}_2$. ■

Proof of Theorem 9. Set all D^*_{2t+1}'s to zero except for $D^*_{2t+1}(2,2)$. Using the proof for Theorem 9, we can solve for $D^*_{2t+1}(2,2)$ such that finite dependence results under a certain restriction. If this restriction is violated, then we can solve for $D^*_{2t+1}(2,1)$ and proceed accordingly.

$D^*_{2t+1}(2,2|x_t)$ is then given by:

$$\textbf{??}D^*_{2t+1}(2,2|x_t) = \frac{\mathcal{P}^{(2)}_{2t+2}(F_{1t+1}(1)F_{1t}(x_t) - F_{1t+1}(2)F_{2t}(x_t))}{p^{(2)}_{2t}(x_t)\left(p^{(2)}_{2t+2}(2,1) - p^{(2)}_{2t+2}(1,1) + p^{(2)}_{2t+1}\left[p^{(2)}_{2t+2}(2,2) + p^{(2)}_{2t+2}(1,1) - p^{(2)}_{2t+2}(2,1) - p^{(2)}_{2t+2}(1,2)\right]\right)} \tag{61}$$

The restriction is then that the denominator in (**??**) is non-zero. But if it is zero, then it will not be zero if we instead set all D^*_{2t+1}'s to zero except for $D^*_{2t+1}(2,1)$.

The numerator of (**??**) is:

$$
\begin{aligned}
& \begin{bmatrix} p^{(2)}_{2t+2}(2,2) & p^{(2)}_{2t+2}(2,1) & p^{(2)}_{2t+2}(1,2) & p^{(2)}_{2t+2}(1,1) \end{bmatrix} \\
& \times \left(\begin{bmatrix} 0 & 0 \\ 0 & 0 \\ p^{(2)}_{2t+1}(1,2) & p^{(2)}_{2t+1}(1,1) \\ p^{(2)}_{1t+1}(1,2) & p^{(2)}_{2t+1}(1,1) \end{bmatrix} \begin{bmatrix} p^{(2)}_{2t}(x_t) \\ p^{(2)}_{1t}(x_t) \end{bmatrix} - \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ p^{(2)}_{2t+1}(2,2) & p^{(2)}_{2t+1}(2,1) \\ p^{(2)}_{1t+1}(2,2) & p^{(2)}_{2t+1}(2,1) \end{bmatrix} \begin{bmatrix} p^{(2)}_{2t}(x_t) \\ p^{(2)}_{1t}(x_t) \end{bmatrix} \right) \\
= & \sum_{j=1}^{2} p^{(2)}_{2t+2}(1,j) \left(p^{(2)}_{2t}(x_t) \left[p^{(2)}_{jt+1}(1,2) - p^{(2)}_{jt+1}(2,2) \right] + p^{(2)}_{1t}(x_t) \left[p^{(2)}_{jt+1}(1,1) - p^{(2)}_{jt+1}(2,1) \right] \right)
\end{aligned}
$$

Now consider the expression for $u^{(1)}_{2t}(2,2)$ which can be written as:

$$
\begin{aligned}
u^{(1)}_{2t}(2,2) \quad = \quad & \psi_1(p_t(2,2)) + \beta p^{(2)}_{1t} \left[\psi_1(p_{t+1}(1,1)) - \psi_1(p_{t+1}(2,1)) \right] + \beta p_{2t}(2,2)\psi_1(p_{t+1}(1,2)) \\
& + \beta p^{(2)}_{2t}(2,2) D^*_{2t+1}(2,2|2,2) \left[\psi_2(p_{t+1}(2,2)) + u^{(1)}_{2t+1}(2,2) \right] + \beta p^{(2)}_{2t}(2,2)(1 - D^*_{2t+1}(2,2|2,2))\psi_1(p_{t+1}(2,2))
\end{aligned}
$$

Note that we can normalize one of the θ's to zero. Normalizing θ_{t+1} to zero, substituting in for u_2's, and rearranging terms yields:

$$
\begin{aligned}
& \sum_{j=1}^{2} \left(p^{(2)}_{jt}(2,2) - \beta p^{(2)}_{2t}(2,2) D^*_{2t+1}(2,2|2,2) p^{(2)}_{jt}(2,2) \right) U^{(1)}_2(2,j) + \theta_t = \\
& \psi_1(p^{(1)}_t(2,2)) + \beta p^{(2)}_{1t}(2,2) \left[\psi_1(p^{(1)}_{t+1}(1,1)) - \psi_1(p^{(1)}_{t+1}(2,1)) \right] + \beta p^{(2)}_{2t}(2,2)\psi_1(p^{(1)}_{t+1}(1,2)) \\
& + \beta p^{(2)}_{2t}(2,2) D^*_{2t+1}(2,2|2,2)\psi_2(p^{(1)}_{t+1}(2,2)) + \beta p^{(2)}_{2t}(2,2)(1 - D^*_{2t+1}(2,2|2,2))\psi_1(p^{(1)}_{t+1}(2,2))
\end{aligned}
$$

The analogous expression when the initial state is (2,1) is:

$$
\begin{aligned}
& \sum_{j=1}^{2} \left(p^{(2)}_{jt}(2,1) - \beta p_{2t}(2,1) D^*_{2t+1}(2,2|2,2) p^{(2)}_{jt}(2,2) \right) U^{(1)}_2(2,j) + \theta_t = \\
& \psi_1(p^{(1)}_t(2,1)) + \beta p^{(2)}_{1t}(2,1) \left[\psi_1(p^{(1)}_{t+1}(1,1)) - \psi_1(p^{(1)}_{t+1}(2,1)) \right] + \beta p^{(2)}_{2t}(2,1)\psi_1(p^{(1)}_{t+1}(1,2)) \\
& + \beta p^{(2)}_{2t}(2,1) D^*_{2t+1}(2,2|2,1)\psi_2(p^{(1)}_{t+1}(2,2)) + \beta p^{(2)}_{2t}(2,1)(1 - D^*_{2t+1}(2,2|2,1))\psi_1(p^{(1)}_{t+1}(2,2))
\end{aligned}
$$

The system above has two equations and three unknowns. Rolling back one period gives us two more equations with one more unknown, θ_{t-1}:

$$
\begin{aligned}
& \sum_{j=1}^{2} \left(p^{(2)}_{jt-1}(2,2) - \beta p_{2t-1}(2,2) D^*_{2t}(2,2|2,2) p^{(2)}_{jt-1}(2,2) \right) U^{(1)}_2(2,j) + \theta_{t-1} = \\
& \psi_1(p^{(1)}_{t-1}(2,2)) + \beta p^{(2)}_{1t-1}(2,2) \left[\psi_1(p^{(1)}_t(1,1)) - \psi_1(p^{(1)}_t(2,1)) \right] + \beta p_{2t-1}(2,2)\psi_1(p^{(1)}_t(1,2)) \\
& + \beta p^{(2)}_{2t-1}(2,2) D^*_{2t}(2,2|2,2) \left[\psi_2(p^{(1)}_t(2,2)) + \theta_t \right] + \beta p^{(2)}_{2t-1}(2,2)(1 - D^*_{2t}(2,2|2,2))\psi_1(p^{(1)}_t(2,2))
\end{aligned}
$$

$$
\sum_{j=1}^{2}\left(p^{(2)}_{jt-1}(2,1)-\beta p_{2t-1}(2,1)D^{*}_{2t}(2,2|2,1)p^{(2)}_{jt-1}(2,1)\right)U^{(1)}_{2}(2,j)+\theta_{t-1}=
$$

$$
\psi_1(p^{(1)}_{t-1}(2,1))+\beta p^{(2)}_{1t-1}(2,1)\left[\psi_1(p^{(1)}_{t}(1,1))-\psi_1(p^{(1)}_{t}(2,1))\right]+\beta p_{2t-1}(2,1)\psi_1(p^{(1)}_{t}(1,2))
$$

$$
+\beta p^{(2)}_{2t-1}(2,1)D^{*}_{2t}(2,2|2,1)\left[\psi_2(p^{(1)}_{t}(2,2))+\theta_t\right]+\beta p^{(2)}_{2t-1}(2,1)(1-D^{*}_{2t}(2,2|2,1))\psi_1(p^{(1)}_{t}(2,2))
$$

implying we now have four equations and four unknowns. ■

References

[1] Aguirregabiria, V. and A. Magesan (2013): “Euler Equations for the Estimation of Dynamic Discrete Choice Structural Models”, working paper.

[2] Aguirregabiria, V., and P. Mira (2002):“Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models,” *Econometrica*, 70, 1519-1543.

[3] Aguirregabiria, V., and P. Mira (2007): “ Sequential Estimation of Dynamic Discrete Games,” *Econometrica,* 75, 1-54.

[4] Aguirregabiria, V., and P. Mira (2010): “Dynamic Discrete Choice Structural Models: A Survey," *Journal of Econometrics*, 156(1), 38-67.

[5] Altug, S., and R. A. Miller (1998): “The Effect of Work Experience on Female Wages and Labour Supply,” *Review of Economic Studies*, 62, 45-85.

[6] Arcidiacono, P. and P. Ellickson (2011):“Practical Methods for Estimation of Dynamic Discrete Choice Models", *Annual Review of Economics* Vol. 3, 363-394.

[7] Arcidiacono, P., and R.A. Miller (2011):“Conditional Choice Probability Estimation of Dynamic Discrete Choice Model with Unobserved Heterogeneity”, *Econometrica*, Vol. 7, No. 6, 1823-1868.

[8] Arcidiacono, P., H. Sieg, and F. Sloan (2007):“Living Rationally Under the Volcano? An Empirical Analysis of Heavy Drinking and Smoking.” *International Economic Review*, 48(1).

[9] Bajari, P., V. Chernozhukov, H. Hong, and D. Nekipelov (2009):“Identification and Efficient Semiparametric Estimation of a Dynamic Discrete Game,” working paper.

[10] Chiong, K.X., A. Galichon, and M. Shum (2013): "Estimating Dynamic Discrete Choice Models via Convex Analysis," working paper.

[11] Eckstein, Z. and K.I. Wolpin (1989):"The Specification and Estimation of Dynamic Stochastic Discrete Choice Models: A Survey," *The Journal of Human Resources*, 24(4), 562-598.

[12] Gayle, W. (2013): "CCP Estimation of Dynamic Discrete/Continuous Choice Models with Generalized Finite Dependence and Correlated Unobserved Heterogeneity", working paper.

[13] Hotz, V. J., and R. A. Miller (1993): " Conditional Choice Probabilities and Estimation of Dynamic Models," *Review of Economic Studies*, 61, 265-289.

[14] Hu, Y. and M. Shum (2012):"Nonparametric identification of dynamic models with unobserved state variables," *Journal of Econometrics*, 171, 32-44.

[15] Joefre-Bonet, M. and M. Pesendorfer (2003):"Estimation of a Dynamic Auction Game," *Econometrica*, 71(5), pp. 1443-1489.

[16] Kasahara, H., and K. Shimotsu (2009): "Nonparametric Identification and Estimation of Finite Mixture Models of Dynamic Discrete Choices," *Econometrica*, 77(1), 135-175.

[17] Keane, M. P., and K. I. Wolpin (1994): "The solution and Estimation of Discrete Choice Dynamic Models by Simulation and Interpolation: Monte Carlo Evidence," *The Review of Economic Statistics*, 76, pp. 648 - 672.

[18] Magnac T. and D. Thesmar (2002):"Identifying Dynamic Discrete Decision Processes," *Econometrica*, 70(2), 801-816.

[19] McFadden, D. (1978): "Modelling the Choice of Residential Location," *Spatial Interaction and Planning Models,* editors F. Snickars and J. Weibull, North Holland Amsterdam, 75-96.

[20] Miller, R.A. (1984): "Job Matching and Occupational Choice," *Journal of Policitical Economy*, 92(6), pp. 1086 -1120.

[21] Miller, R.A. (1997) "Estimating models of dynamic optimization with microeconomic data," in *Handbook of Applied Econometrics*, M. Pesaran and P. Schmidt, editors, Basil Blackwell, Vol. 2, pp. 246-299.

[22] Norets, A. and X. Tang (2012): "Semiparametric Inference in Dynamic Binary Choice Models", working paper.

[23] Pakes, A. (1994), "The Estimation of Dynamic Structural Models: Problems and Prospects, Part II. Mixed Continuous-Discrete Control Models and Market Interactions," Chapter 5 of *Advances in Econometrics: Proceedings of the 6th World Congress of the Econometric Society*, edited by J.J. Laffont and C. Sims., pp. 171-259,

[24] Pakes, A., M. Ostrovsky, and S. Berry (2007): " Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry/Exit Examples)," *RAND Journal of Economics*, 38(2), 373-399.

[25] Pesendorfer, M., and P. Schmidt-Dengler (2008): " Asymptotic Least Square Estimators for Dynamic Games," *Review of Economic Studies*, 75, 901-908.

[26] Pesendorfer, M. and P. Schmidt-Dengler (2010): "Sequential Estimation of Dynamic Discrete Games: A Comment", *Econometrica*, Vol. 78, No. 2, 833-842.

[27] Rust, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica*, 55, 999-1033.

[28] Rust, J. (1994): "Structural Estimation of Markov Decision Processes," in *Handbook of Econometrics, Volume 4*, ed. by R.E. Engle and D. McFadden. Amsterdam: Elsevier-North Holland, Chapter 14.