Identifying Dynamic Discrete Choice Models off Short Panels*

Peter Arcidiacono — Duke University & NBER

Robert A. Miller — Carnegie Mellon University

April 2, 2013

Abstract

This paper considers nonparametric identification of nonstationary dynamic discrete choice models when the agent's time horizon extends beyond the length of the data. We show conditions under which flow payoffs are identified subject to standard normalizations and, when the payoff function does not depend directly on time, show identification even when the time horizon extends beyond the length of the data. We further establish identification for a class of nonstationary dynamic discrete choice games and show how the nonstationarity of the problem can be helpful in unbundling the agent's state-specific payoffs from the expected payoffs where the expectation is taken over the actions of the agent's competitors.

*We acknowledge support from National Science Foundation Grant Awards SES0721059 and SES0721098. Preliminary and Incomplete. Comments welcome.

1 Introduction

In this paper we consider identification of dynamic discrete choice models when the environment is non-stationary, analyzing both single and multi-agent settings. Further, we consider cases where individual's are forming expectations about events far beyond what is seen in the data and the researcher is unwilling to make assumptions about how expectations are formed in time periods beyond what is seen in the data.

Our identification results build on the insight of Arcidiacono and Miller (2011) who note that future payoffs can be represented as flow payoffs for any decision rule over the full time horizon plus correction terms that only depend on the conditional choice probabilities. By exploiting the structure of the model, a clever choice of which decision rule to write the future value term with respect to can result in the future value term–or differences in future value terms across two choices–only depending on a few period ahead conditional choice probabilities and flow payoffs.

For a particular class of models, we show identification, subject to standard normalizations, of a set of non-stationary flow payoff functions (i.e. the flow payoff depends directly on t) even when the length of the panel is shorter than the time horizon. The class of models includes renewal problems such as in Rust (1987) or games where there is an exit decision. Hence, with a short panel we are able to obtain identification of a class of non-stationary dynamic games.

In cases where the flow payoff function does not depend on time directly but only through the state variables, we are able to show identification for a much larger class of models, again even when the time horizon extends beyond what is seen in the data. Here the non-stationarity occurs either because the transitions of the states depend on time or because the time horizon is finite.

Indeed, our results suggest that non-stationarity can serve as an aid in identification when the model obeys certain properties. We show, for example, that in some cases the flow payoff for one choice in one state needs to be normalized as opposed to the flow payoff for one choice in all states.

We further consider identification in non-stationary dynamic games, first considering identification of expected payoffs of a particular choice where the expectation is taken over the actions of the agent's competitors. Here too finite dependence is helpful in establishing identification and we show how to obtain finite dependence paths in games. Further, we show how nonstationarity aids in unbundling state-specific payoffs from the expected payoffs.

The rest of the paper proceeds as follows. First, we consider the single agent setting, covering both cases where the full time horizon is seen in the data and other cases where the time horizon goes beyond the observed data. We show identification results in this latter case when there is either a terminal or a renewal action or when the flow payoff function does not depend directly on time. Next, we cover games and show how to recover finite dependence paths and how these finite dependence paths facilitate identification of expected payoffs. Finally, we show nonstationarity can aid in the unbundling of state-specific payoffs from expected payoffs.

2 Framework

We first develop the framework for single agent case, following closely Arcidiacono and Miller (2011) section 3. 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:

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

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:

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

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:[1]

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

3 Identification in Single Agent Settings

Let $\mathcal{T} \leq T$ denoted the date the panel ends. We differentiate between two scenarios. When $\mathcal{T} = T < \infty$, or when $T = \infty$ and the environment is stationary, we extend the results of and Thesmar (2002) and Pesendorfer and Schmidt-Dengler (2008), who analyze identification when the conditional choice probabilities are known. Identification when $\mathcal{T} < T < \infty$, or $T = \infty$ and the environment is nonstationary, is trickier. In this case the method of solving and imposing the solution of the underlying discrete choice problem on the data generating process is not feasible without making strong parametric assumptions about features of the framework that affect the flow payoffs and state variable transitions between dates $\mathcal{T}$ and T. We show, however, that the model is partially identified under much weaker assumptions involving specializations of the finite dependence property.

We adopt several normalizations commonly made in the static and dynamic discrete choice literature. Since only choices are observed rather than utility levels, we normalize the flow payoff

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

function for one of the choices to zero in every time period, setting $u_{1t}(x_t) = 0$ for all x_t and t.[2] Absent observations on utility, we also assume the distribution of the transitory vector of shocks, $G(\epsilon_t)$, is known. Finally because the time-subscripted utility functions depend in an unrestricted way upon the time period, the subjective discount factor β must also be normalized.[3]

Our identification results follow from the representation theorems given in Arcidiacono and Miller (2011). Lemma 1 of Arcidiacono and Miller (2011) shows that the value function given in (3) can be expressed as a function of the conditional choice probabilities, $p_t(x_t)$ and one conditional value function, $v_{jt}(x_t)$. Specifically they show the there exists a real-valued function ψ_j for every $j \in \{1, \dots, J\}$ such that:

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

Substituting (5) into the right hand side of (4) we obtain:

$$v_{jt}(x_t) = u_{jt}(x_t) + \beta \sum_{x_{t+1}=1}^{X} [v_{kt+1}(x_{t+1}) + \psi_k [p_{t+1}(x_{t+1})]] f_{jt} (x_{t+1}|x_t) \tag{6}$$

Equation (6) shows that the conditional value function can be expressed as the flow payoff of the choice plus a function of the one period ahead conditional choice probabilities and the one period ahead conditional value function for any choice.

Arcidiacono and Miller use this result to show that the value function can be expressed as a function of the flow payoffs associated with any decision rule that covers period t to T plus functions of the conditional choice probabilities. Consider a sequence of decisions from t to T. The first choice in the sequence is the initial choice j which sets $d^*_{jt}(x_t, j) = 1$. For periods $\tau \in \{t+1, \dots, T\}$, the

[2]As noted by Pesendorfer and Schmidt-Dengler (2008, page 913), in some situations, such when a firm exits an industry, the future value of the choice to exit may be known, providing an empirical justification for the normalization.

[3]Later in the paper we show how exclusion restrictions can be used to recover β and components of the utility levels of the normalized choice.

choice sequence maps x_τ and the initial choice j into $d^*_\tau(x_\tau, j) \equiv \{d^*_{1\tau}(x_\tau, j), \ldots, d^*_{J\tau}(x_\tau, j)\}$. The choices in the sequence then must satisfy $d^*_{k\tau}(x_\tau, j) \geq 0$ and $\sum_{k=1}^{J} d^*_{k\tau}(x_\tau, j) = 1$. Note that the choice sequence can depend upon new realizations of the state and may also involve mixing over choices.

Now consider the probability of being in state $x_{\tau+1}$ conditional on following the choices in the sequence. Denote this probability as $\kappa^*_\tau(x_{\tau+1}|x_t, j)$ which is recursively defined by:

$$\kappa^*_\tau(x_{\tau+1}|x_t, j) \equiv \begin{cases} f_{jt}(x_{t+1}|x_t) & \text{for } \tau = t \\ \sum_{x_\tau=1}^{X} \sum_{k=1}^{J} d^*_{k\tau}(x_\tau, j)\, f_{k\tau}(x_{\tau+1}|x_\tau)\kappa^*_{\tau-1}(x_\tau|x_t, j) & \text{for } \tau = t+1, \ldots, T \end{cases} \tag{7}$$

The future value term can now be expressed relative to the conditional value functions for the choices in the sequence. Theorem 1 of Arcidiacono and Miller (2011) shows that continuing to express the future value term relative to the value of the next choice in the sequence yields an alternative expression for $v_{jt}(x_t)$:

$$v_{jt}(x_t) = u_{jt}(x_t) + \sum_{\tau=t+1}^{T} \sum_{k=1}^{J} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \left[u_{k\tau}(x_\tau) + \psi_k[p_\tau(x_\tau)]\right] d^*_{k\tau}(x_\tau, j)\kappa^*_{\tau-1}(x_\tau|x_t, j) \tag{8}$$

Note this expression holds for any decision sequence. It is this representation that we exploit in showing identification.

Before doing so, however, we slightly relax Theorem 1 of Arcidiacono and Miller (2011). Namely, the only requirement we place on $d^*_{k\tau}(x_\tau, j)$ is that $\sum_{k=1}^{J} d^*_{k\tau}(x_\tau, j) = 1$. That is, we can place negative weights on particular choice paths so long as the sum of the weights equals one in every time period.

Lemma 1 *$v_{jt}(x_t)$ can be expressed as in (8) for all decision sequences such that $\sum_{k=1}^{J} d^*_{k\tau}(x_\tau, j) = 1$ for all τ.*

As shown below, applying negative weights can be useful in establishing identification for some problems.

3.1 When the setting is stationary or $\mathcal{T} = T$

Building on the results of Magnac and Thesmar (2002), we begin by showing how to recover flow payoff functions when the sampling period is the same as the horizon or when the problem is stationary. Theorem 2 and Corollary 3 of Magnac and Thesmar (2002, pages 807 and 808) establish exact identification of the differences in conditional value functions when T is finite. An equivalent result holds for the identification of flow utilities. Let $d_{1\tau}^{*}(x_\tau) = 1$ for all τ in equation (8) and subtract $v_{1t}(x_t)$ from $v_{jt}(x_t)$ to obtain:

$$v_{jt}(x_t) - v_{1t}(x_t) = u_{jt}(x_t) + \sum_{\tau=t+1}^{T} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \psi_1[p_\tau(x_\tau)] \left[\kappa_{\tau-1}^{*}(x_\tau|x_t, j) - \kappa_{\tau-1}^{*}(x_\tau|x_t, 1)\right] \tag{9}$$

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):

$$v_{jt}(x_t) - v_{1t}(x_t) = \psi_1[p_t(x_t)] - \psi_j[p_t(x_t)] \tag{10}$$

As shown in Theorem 1 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 1 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.

Theorem 1 *For all j, t, and x_t, the flow payoff $u_{jt}(x_t)$ can be expressed as.*

$$u_{jt}(x_t) = \psi_1[p_t(x_t)] - \psi_j[p_t(x_t)] + \sum_{\tau=t+1}^{T} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \psi_1[p_\tau(x_\tau)] \left[\kappa_{\tau-1}^{*}(x_\tau|x_t, 1) - \kappa_{\tau-1}^{*}(x_\tau|x_t, j)\right] \tag{11}$$

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

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

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

$$u_j = \Psi_j - \Psi_1 + \beta\left(\mathcal{F}_1 - \mathcal{F}_j\right)\left[\mathcal{I} - \beta\mathcal{F}_1\right]^{-1}\ \Psi_1 \tag{12}$$

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 (11) and (12) is known and, therefore, both systems are exactly identified. However, by putting further structure on the flow payoff function, the error distribution can be made more flexible and the discount factor may be identified.

3.2 When the setting is nonstationary and $\mathcal{T} < T$

When no data on choices or state transitions are available for the last part of the lifecycle, the expressions derived above cannot aid identification, because there are no conditional choice probabilities or transition functions past $\mathcal{T}$. In this section we show two cases where it is is possible to recover the same flow payoffs as in section 3.1 until time $\mathcal{T} - 1$. We then show for an expanded class of models that it is possible to recover flow payoffs when the flow payoff function does not depend on time directly, but only through the values of the states.

3.2.1 Renewal and terminal choices

The first case where identification can be restored for payoff functions until $\mathcal{T} - 1$ is when there is a terminal choice–a choice where no future decisions are made. Normalize the payoff of the terminal choice to zero. Since the value function can be expressed relative to the conditional value functions for one of the choices plus a function of the conditional choice probabilities, we can express the value function relative to the conditional value function of the terminal choice. But since the terminal choice has no future component, the conditional value function is zero.

A second case occurs when there is a renewal action, an action which, when taken, resets the states that were influenced by past actions. An example is the bus-engine replacement problem in Rust (1987). Replacing the engine resets the bus mileage regardless of when the engine was replaced previously. Labeling the renewal action as action 1, a renewal action at $t+1$ satisfies:

$$\sum_{x_{t+1}=1}^{X} f_{1,t+1}(x_{t+2}|x_{t+1})f_{jt}(x_{t+1}|x_t) = \sum_{x_{t+1}=1}^{X} f_{1,t+1}(x_{t+2}|x_{t+1})f_{j't}(x_{t+1}|x_t) \tag{13}$$

for all x_t and all $\{j, j'\} \in [1, \ldots, J]$.

We normalize the flow payoff of the renewal action in each time period to zero. Expressing the future value term for the flow payoff at time t for action j relative to the renewal choice implies (4) can be written as:

$$\begin{aligned} v_{jt}(x_t) &= u_{jt}(x_t) + \beta \sum_{x_{t+1}=1}^{X} \left[v_{1t+1}(x_{t+1}) + \psi_1\left[p_{t+1}(x_{t+1})\right]\right] f_{jt}\left(x_{t+1}|x_t\right) \qquad (14) \\ &= u_{jt}(x_t) + \beta \sum_{x_{t+1}=1}^{X} \left[\psi_1\left[p_{t+1}(x_{t+1})\right] + \sum_{x_{t+2}=1}^{X} \beta V_{t+2}(x_{t+2}) f_{1t+1}(x_{t+2}|x_{t+1})\right] f_{jt}\left(x_{t+1}|x_t\right) \end{aligned}$$

Expressing $v_{1t}(x_t)$ in a similar way implies that the period $t+2$ value function on the right hand side of (14) will enter in the same way in the two expressions. Differencing the two expressions, rearranging terms, and substituting in for $v_{jt}(x_t) - v_{1t}(x_t)$ with (10) yields:

$$u_{jt}(x_t) = \psi_1[p_t(x_t)] - \psi_j[p_1(x_t)] + \sum_{x_{t+1}=1}^{X} \beta\psi_1[p_{t+1}(x_{t+1})]\left[f_{1t}(x_{t+1}|x_t) - f_{jt}(x_{t+1}|x_t)\right] \tag{15}$$

Given known functional forms for ψ, the conditional choice probabilities through $\mathcal{T}$, and β, the flow payoffs up until $\mathcal{T}-1$ are identified. The reason the flow payoffs for $\mathcal{T}$ are not identified is that we need the one-period-ahead conditional choice probabilities and the data only extend until $\mathcal{T}$.

Example 1: Exclusion restrictions and identification of β

Note that identification at this point has not relied on the presence of exclusion restrictions: the same variables that affect the transitions of the states also affect the flow payoffs. Suppose we can partition the state vector x into a set of variables that affect both the flow payoffs, x_1, and a set that affect both the flow payoffs and the transitions, x_2. In this case, we can use exclusion restrictions to potentially identify discount factor or the normalized flow payoffs as in Norets and Tang (2012). Here we show that when an exclusion restriction is present and there is a renewal choice, there is a closed form expression for β and β is over-identified.

Consider the set of individuals whose state variables at time t have the same values of x_{1t}. Now consider two different values of x_{2t}, x_{2t}^A and x_{2t}^B. Differencing (15) across the two different values of x_{2t} yields:

$$\begin{aligned} 0 &= \psi_1[p_t(x_{1t},x_{2t}^A)] - \psi_j[p_t(x_{1t},x_{2t}^A)] + \sum_{x_{t+1}=1}^{X} \beta\psi_1[p_{t+1}(x_{t+1})]\left[f_{1t}(x_{t+1}|x_{1t},x_{2t}^A) - f_{jt}(x_{t+1}|x_{1t},x_{2t}^A)\right] \\ &+ \psi_j[p_t(x_{1t},x_{2t}^B)] - \psi_1[p_t(x_{1t},x_{2t}^B)] + \sum_{x_{t+1}=1}^{X} \beta\psi_1[p_{t+1}(x_{t+1})]\left[f_{jt}(x_{t+1}|x_{1t},x_{2t}^B) - f_{1t}(x_{t+1}|x_{1t},x_{2t}^B)\right] \end{aligned}$$

Solving for β yields:

$$\beta = \frac{\psi_1[p_t(x_{1t},x_{2t}^A)] - \psi_j[p_t(x_{1t},x_{2t}^A)] + \psi_j[p_t(x_{1t},x_{2t}^B)] - \psi_1[p_t(x_{1t},x_{2t}^B)]}{\sum_{x_{t+1}=1}^{X}\psi_1[p_{t+1}(x_{t+1})]\left[f_{jt}(x_{t+1}|x_{1t},x_{2t}^A) - f_{1t}(x_{t+1}|x_{1t},x_{2t}^A) + f_{1t}(x_{t+1}|x_{1t},x_{2t}^B) - f_{jt}(x_{t+1}|x_{1t},x_{2t}^B)\right]} \tag{16}$$

3.2.2 Under-identification and stability of the payoff function

Absent a renewal or terminal choice, we cannot establish identification of the flow payoffs for the first $\mathcal{T}-1$ periods. However, we can describe the degree of under-identification. In particular, similar to Magnac and Thesmar (2002), we show under-identification relative to the value function at $\mathcal{T}+1$.

Given that we do not see state transitions and conditional choice probabilities after $\mathcal{T}$, we express u_{jt} as in (11) 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}:

$$\begin{aligned} u_{jt}(x_t) &= \psi_1[p_t(x_t)] - \psi_j[p_t(x_t)] + \sum_{\tau=t+1}^{\mathcal{T}} \sum_{x_\tau=1}^{X} \beta^{\tau-t}\psi_1[p_\tau(x_\tau)] \left[\kappa^*_{\tau-1}(x_\tau|x_t,1) - \kappa^*_{\tau-1}(x_\tau|x_t,j)\right] \\ &\quad + \sum_{x_{\mathcal{T}+1}=1}^{X} V_{\mathcal{T}+1}(x_{\mathcal{T}+1})] \left[\kappa^*_{\mathcal{T}}(x_{\mathcal{T}+1}|x_t,1) - \kappa^*_{\mathcal{T}}(x_{\mathcal{T}+1}|x_t,j)\right] \end{aligned} \tag{17}$$

Note that is is the last term that leads to under-identification, the degree of which is specified in theorem 2.

Theorem 2 *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 $X-1$.*

Progress can be made, however, when the payoff function is stable over time, $u_{jt}(x) = u_{jt'}(x)$ 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.

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 $u_{1t}(x) = u_{1t+1}(x) = 0$. Applying Theorem 1 to the second period:

$$\begin{aligned} u_2(x) &= \psi_1[p_{t+1}(x)] - \psi_2[p_{t+1}(x)] + \sum_{x'=1}^{X} \beta V_3(x') \left[f_{1t+1}(x'|x) - f_{2t+1}(x'|x)\right] \\ &= \psi_1[p_{t+1}(x)] - \psi_2[p_{t+1}(x)] + \sum_{x'=1}^{X-1} \beta V_{t+2}(x') \left[f_{1t+1}(x'|x) - f_{2t+1}(x'|x)\right] \end{aligned} \tag{18}$$

Applying Theorem 1 to the first and second periods:

$$\begin{aligned}
u(x) &= \psi_1[p_t(x)] - \psi_2[p_t(x)] \\
&\quad + \sum_{x'=1}^{X} \beta\psi_1[p_{t+1}(x')]\left[f_{1t}(x'|x) - f_{2t}(x'|x)\right] + \sum_{x'=1}^{X} \beta^2 V_{t+2}(x')\left[\kappa^*_{t+1}(x'|x,1) - \kappa^*_{t+1}(x'|x,2)\right] \\
&= \psi_1[p_t(x)] - \psi_2[p_t(x)] \qquad (19) \\
&\quad + \sum_{x'=1}^{X-1} \beta\psi_1[p_{t+1}(x')]\left[f_{1t}(x'|x) - f_{2t}(x'|x)\right] + \sum_{x'=1}^{X-1} \beta^2 V_{t+2}(x')\left[\kappa^*_{t+1}(x'|x,1) - \kappa^*_{t+1}(x'|x,2)\right]
\end{aligned}$$

Subtracting the first equation (18) from the second (19) yields:

$$\begin{aligned}
&\psi_1[p_{t+1}(x)] - \psi_2[p_{t+1}(x)] + \sum_{x'=1}^{X-1} \beta V_{t+2}(x')\left[f_{1t+1}(x'|x) - f_{2t+1}(x'|x)\right] \\
= \;&\psi_1[p_t(x)] - \psi_2[p_t(x)] + \sum_{x'=1}^{X-1} \beta\psi_1[p_{t+1}(x')]\left[f_{1t}(x'|x) - f_{2t}(x'|x)\right] \\
&+ \sum_{x'=1}^{X-1} \beta^2 V_{t+2}(x')\left[\kappa^*_{t+1}(x'|x,1) - \kappa^*_{t+1}(x'|x,2)\right]
\end{aligned}$$

for each $x \in \{1, \ldots, X\}$, and collecting terms:

$$\begin{aligned}
&\sum_{x'=1}^{X-1} V_{t+2}(x')\left[\beta f_{1t+1}(x'|x) - \beta f_{2t+1}(x'|x) - \beta^2\kappa^*_{t+1}(x'|x,1) + \beta^2\kappa^*_{t+1}(x'|x,2)\right] \\
= \;&\psi_1[p_t(x)] - \psi_2[p_t(x)] - \psi_1[p_{t+1}(x)] + \psi_2[p_{t+1}(x)] \\
&+ \sum_{x'=1}^{X-1} \beta\psi_1[p_2(x')]\left[f_{1t}(x'|x) - f_{2t}(x'|x)\right]
\end{aligned}$$

We write the $X-1$ dimensional vectors:

$$V_{t+2} \equiv \begin{bmatrix} V_{t+2}(1) \\ \vdots \\ V_{t+2}(X-1) \end{bmatrix}$$

$$\Psi \equiv \begin{bmatrix} \psi_1[p_t(1)] - \psi_2[p_t(1)] - \psi_1[p_{t+1}(1)] + \psi_2[p_{t+1}(1)] \\ +\sum_{x'=1}^{X-1} \beta\psi_1[p_{t+1}(x')]\left[f_{1t}(x'|1) - f_{2t}(x'|1)\right] \\ \vdots \\ \psi_1[p_t(X-1)] - \psi_2[p_t(X-1)] - \psi_1[p_{t+1}(X-1)] + \psi_2[p_{t+1}(X-1)] \\ +\sum_{x'=1}^{X} \beta\psi_1[p_{t+1}(x')]\left[f_{1t}(x'|X-1) - f_{2t}(x'|X-1)\right] \end{bmatrix}$$

and the X dimensional square matrix:

$$A \equiv \begin{bmatrix} A^{(1,1)} & \dots & A^{(1,X-1)} \\ \vdots & \ddots & \vdots \\ A^{(X-1,1)} & \dots & A^{(X-1,X-1)} \end{bmatrix}$$

where:

$$A^{(x,x')} \equiv \left[\beta f_{1t+1}(x'|x) - \beta f_{2t+1}(x'|x) - \beta^2\kappa^*_{t+1}(x'|x,1) + \beta^2\kappa^*_{t+1}(x'|x,2)\right]$$

In matrix form now:

$$\Psi = AV_{t+2}$$

so if A is invertible

$$V_{t+2} = A^{-1}\Psi$$

and hence

$$\begin{aligned} u_2(x) &= \psi_1[p_{t+1}(x)] - \psi_2[p_{t+1}(x)] + \sum_{x'=1}^{X-1} \beta V_{t+2}(x')\left[f_{1t+1}(x'|x) - f_{2t+1}(x'|x)\right] \\ &= \psi_1[p_{t+1}(x)] - \psi_2[p_{t+1}(x)] + \sum_{x'=1}^{X-1} \beta\left(F_{1x} - F_{2x}\right)A^{-1}\Psi \end{aligned}$$

where

$$F_{jx} = (f_{jt+1}(1|x) \dots f_{jt+1}(X-1|x))$$

So the question of identification turns on the invertibility of A. Note the elements of A are solely mappings from the transitions, model primitives, not conditional choice probabilities.

3.2.3 Finite dependence

While identification may be achieved in non stationary models where the flow payoffs are stable, establishing identification may be impractical due to having to invert a matrix with dimension equal to the size of the state space. We next consider a class of models that satisfy an expanded version of the finite dependence property of Arcidiacono and Miller (2011). We show that exploiting the finite dependence property can often make it much easier to establish identification as well as showing how to construct finite dependence paths.

Recall that Arcidiacono and Miller (2011) Theorem 1 established that the future utility term could be expressed as a sequence of flow payoffs from any decision rule plus a function of the conditional choice probabilities. Lemma 1 expanded the decision rules such that the weight on a particular choice in a particular state could be negative as long as the sum of weights on all choices conditional on a particular state sum to one.

Consider two sequences of decision weights, one that begins with choice j and the other with choice j'. We say that the pair of choices $\{j, j'\}$ exhibits ρ-period dependence at state x_t if there exists sequences of decision weights from j and j' such that:

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

for all $x_{t+\rho+1}$. That is, the probabilities of being in each state are the same across the two paths after ρ periods. Under finite dependence, and assuming stability of the payoff function over time,

$u_j(x_t)$ can be expressed as:

$$u_j(x_t) = \psi_1[p_t(x_t)] - \psi_j[p_t(x_t)] + \sum_{\tau=t+1}^{t+\rho} \sum_{k=1}^{J} \sum_{x_\tau=1}^{X} \beta^{\tau-t} \{u_k(x_\tau) + \psi_k[p_\tau(x_\tau)]\} [\kappa_\tau^*(x_\tau|x_t, j) - \kappa_\tau^*(x_\tau|x_t, 1)] \tag{21}$$

We then have the following identification results:

Theorem 3 *If there there exists decision sequences for all pairs choices $\{1, j\}$, $j \in [2, \ldots, J]$ and for all x_t, such that for some ρ (20) holds, then $u_j(x_t)$ is identified for all $j > 1$ and all x_t subject to a rank condition when $\mathcal{T} > \rho$.*

Corollary 4 *If there exists decision sequences for some pairs of choices $\{1, j\}$, $j \in [2, \ldots, J]$ and for all x_t such that for some ρ (20) holds and the decision sequences only include choices where there exists sequences such that (20) holds for ρ, then the flow payoffs for those choices are identified for all x_t subject to a rank condition when $\mathcal{T} > \rho$.*

The first corollary establishes that when all choices exhibit finite dependence relative to choice 1 then all flow payoff terms can be covered if a rank condition is satisfied. The second corollary states that identification of some of the flow payoff terms can be achieved even if all choices do not exhibit finite dependence relative to choice 1. Namely, if the finite dependence paths for the choices do not involve the choice where there is no finite dependence path, identification of the flow payoffs for choices that *do* have finite dependence paths is preserved.

Note that, in contrast to previous examples, flow payoff terms are on both sides of (21). Note further that there are $(J-1) \times (\mathcal{T} - \rho) \times X$ equations but only $(J-1) \times X$ unknowns. Hence, if a rank condition on the equations is met and conditional on one flow payoff being normalized, the remaining flow payoff terms may be recovered. Further, since the number of equations exceed the number of unknowns, it may also possible to parameters such as the flow payoffs of the normalized choice for some states, the discount factor, or the correlation patterns of the structural errors.

Example 2: Identification in a non-stationary search model

To illustrate how non-stationarity helps with identification as well as showing how negative weights are useful in obtaining finite dependence, we consider a simple search model. In each period t, $t \in [1, \ldots, T]$ an individual receives a job offer with probability λ. When the individual receives an offer, he can either stay home or accept the offer, $d_{1t} = 1$ and $d_{2t} = 1$ respectively. The value of the offer depends on the experience of the individual, $x \in [0, \ldots, X]$. If the individual accepts the offer, his experiences increases by one unit with probability $\pi_t > 0$ and remains at the current level otherwise. Jobs last only one period so the dynamics of the model come strictly through experience.[4] An individual who does not receive an offer must set $d_{1t} = 1$. As above, we assume that $u_{jt}(x_t) = u_j(x_t)$ for all t.

This example satisfies the finite dependence property but not through a renewal action. Consider the sequence of choices $d_{2t} = 1$, $d_{1t+1} = 1$. Note that this sequence of choices is feasible because the option of staying home is available regardless of whether the individual receives an offer. Now consider the sequence beginning with $d_{1t} = 1$. With probability $(1 - \lambda)$ the individual must choose $d_{1t+1} = 1$. However, we can line up the probabilities of each experience level across the two paths by weighting $d_{2t+1} = 1$ by $\pi_1/(\lambda\pi_{t+1})$ along the path where the offer occurs. Note that this weight can be greater than one, implying a negative weight on $d_{1t+1} = 1$ along the path where an offer arrives at $t + 1$. Using these finite dependence paths, we show that the flow payoff for only one

[4]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.

state-choice combination needs to be normalized. Hence, not only does a finite dependence path exist, but the non-stationarity results in baseline utilities being recovered for all states but one.

Theorem 5 *Given conditional choice probabilities $p_{jt}(x_t)$ for all j, $t \in [1, \ldots, \mathcal{T}]$ and $x_t \in [0, \ldots, X]$, given the distribution of the ϵ's and β, and setting $u_1(0) = 0$, $u_j(x)$ is identified for all j, x, as long as $\mathcal{T} \geq 2$ and $\pi_t \neq \pi_{t'}$ for some pair $\{t, t'\} \in [1, \ldots, \mathcal{T}]$.*

The intuition behind the proof is that we can consider two individuals with the same values of x but in different time periods. By writing out the finite dependence paths for the two time periods, we can see that the two equations are not redundant and can solve for closed form expressions of the payoffs for every state-choice combination subject to normalizing the payoffs for one state-choice combination to zero.

3.2.4 Establishing Finite Dependence

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.

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:

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

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* $N^*_{\tau+2}-1$ *then finite dependence can be achieved in* $\tau - t + 1$ *periods.*

The proof for Theorem 6 shows that the $N^*_{\tau+2}-1$ system of equations we need to solve can be expressed as:

$$\begin{bmatrix} \mathcal{F}_{\tau+1}(j,x_t) & -\mathcal{F}_{\tau+1}(j',x_t) \end{bmatrix} \begin{bmatrix} \mathcal{D}_{\tau+1}(j,x_t) \\ \mathcal{D}_{\tau+1}(j',x_t) \end{bmatrix} = F_{1\tau+1}(j',x_t)^T \mathcal{K}_\tau(j',x_t) - F_{1\tau+1}(j,x_t)^T \mathcal{K}_\tau(j,x_t) \tag{22}$$

4 Games

4.1 Multi-agent framework

In the games setting, we assume that there are I players making choices in periods $[1, \ldots, T]$, $T \leq \infty$. The systematic part of payoffs to the i^{th} player not only depends on his own choice in period t, denoted by $d_t^{(i)} \equiv \left(d_{1t}^{(i)}, \ldots, d_{Jt}^{(i)}\right)$, the state variables x_t, but also the choices of the other players, which we now denote by $d_t^{(-i)} \equiv \left(d_t^{(1)}, \ldots, d_t^{(i-1)}, d_t^{(i+1)}, \ldots, d_t^{(I)}\right)$. Denote by $U_{jt}^{(i)}\left(x_t, d_t^{(-i)}\right) + \epsilon_{jt}^{(i)}$ the current utility of agent i in period t, where $\epsilon_{jt}^{(i)}$ 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 $P_t\left(d_t^{(-i)} | x_t\right)$ as the probability firm i's competitors choose $d_t^{(-i)}$ at time t conditional on the state variables x_t. Since $\epsilon_t^{(i)}$ is independently distributed across all the firms, $P_t\left(d_t^{(-i)} | x_t\right)$ has the product representation:

$$P_t\left(d_t^{(-i)} | x_t\right) = \prod_{\substack{i'=1 \\ i' \neq i}}^{I} \left(\sum_{j=1}^{J} d_{jt}^{(i')} p_{jt}^{(i')}(x_t)\right) \tag{23}$$

We impose rational expectations on the player's beliefs about the choices of its competitors and assume a Markov-perfect equilibrium is played. Hence, the beliefs of the firm match the probabilities given in equation (23). Taking the expectation of $U_{jt}^{(i)}\left(x_t, d_t^{(-i)}\right)$ over $d_t^{(-i)}$, we define the systematic component of the current utility of firm i as a function of the firm's state variables as

$$u_{jt}^{(i)}(x_t) = \sum_{d_t^{(-i)} \in J^{I-1}} P_t\left(d_t^{(-i)} | x_t\right) U_{jt}^{(i)}\left(x_t, d_t^{(-i)}\right) \tag{24}$$

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 period t state variables. Denote $F_{jt}\left(x_{t+1}\middle|x_t, d_t^{(-i)}\right)$ as the probability of x_{t+1} occurring given action j by firm i in period t, when its state variables are x_t and the other firms in its markets choose $d_t^{(-i)}$. The probability of transitioning from x_t to x_{t+1} given action j by firm i in then given by:

$$f_{jt}^{(i)}\left(x_{t+1}\left|x_t\right.\right) = \sum_{d_t^{(-i)}\in J^{I-1}} P_t\left(d_t^{(-i)}\left|x_t\right.\right) F_{jt}\left(x_{t+1}\middle|x_t, d_t^{(-i)}\right) \tag{25}$$

The expressions for the conditional value functions for player i are then no different than what was described in Section 3 subject to the fact that we are now working with expected flow payoffs where the expectation is taken over the decisions of the other palavers. Equation (6) is modified in the games environment to:

$$v_{jt}^{(i)}(x_t) = u_{jt}^{(i)}(x_t) + \beta \sum_{x_{t+1}=1}^{X}\left[v_{kt}^{(i)}(x_{t+1}) + \psi_k\left[p_t^{(i)}(x_{t+1})\right]\right] f_{jt}^{(i)}\left(x_{t+1}|x_t\right) \tag{26}$$

where k indexes any feasible choice in period $t+1$.

As in Section 3, we can then consider a sequence of decisions by player i from t to T and the corresponding probabilities of being in each state given this sequence. Denote the first choice in the sequence as j and denote $d_{k\tau}^{*(i)}(x_\tau, j)$ as the mixing weight placed on choice k at time τ, given that the sequence began with choice j and the state is x_τ. The probability of being in state $x_{\tau+1}$ conditional on following the choices in the sequence can then be defined recursively in a similar manner to equation (7):

$$\kappa_\tau^{*(i)}(x_{\tau+1}|x_t, j) \equiv \begin{cases} f_{jt}^{(i)}(x_{t+1}|x_t) & \text{for } \tau = t \\ \sum_{x_\tau=1}^{X}\sum_{k=1}^{J} d_{k\tau}^{*(i)}\left(x_\tau, j\right) f_{k\tau}^{(i)}(x_{\tau+1}|x_\tau)\kappa_{\tau-1}^{*(i)}(x_\tau|x_t, j) & \text{for } \tau = t+1, \ldots, T \end{cases} \tag{27}$$

Our discussion of identification in multi-agent settings is then broken down into two parts. First, is non-parametric identification of the expected payoffs. Here, the same results as in the single agent setting apply, though satisfying finite dependence may be more difficult. Second is disentangling the state-specific payoffs from the expected payoffs of a particular action.

4.2 Identification of expected payoffs when the environment is stationary

When the environment is stationary, identification of expected utility can be shown as a simple extension of Theorem 1 , assuming that only one equilibrium is played in the data. As in the single agent setting, we normalize the expected payoff for one of the choices to be zero.[5] Without loss of generality, label this assumption as choice 1. Given this normalization, we can express the conditional value function for any choice as the flow payoff for that choice plus functions of the conditional choice probabilities and state transitions, yielding the following theorem:

Theorem 7 *Let $\mathcal{I}$ denote the X dimensional identity matrix and define*

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

Then $\left[\mathcal{I} - \beta \mathcal{F}_1^{(i)}\right]$ is invertible and for all j:

$$u_j^{(i)} = \Psi_j^{(i)} - \Psi_1^{(i)} + \beta \left(\mathcal{F}_1^{(i)} - \mathcal{F}_j^{(i)}\right) \left[\mathcal{I} - \beta \mathcal{F}_1^{(i)}\right]^{-1} \Psi_1^{(i)} \tag{28}$$

where $\mathcal{F}^{(i)}$ and $\Psi^{(i)}$ are evaluated using the choice probabilities for the equilibrium played in the data.

[5]We relax this assumption in some of the non-stationary cases we describe below.

4.3 Terminal and renewal actions

When the environment is non-stationary, the infinite horizon implies immediately that the problem is severely under-identified. However, as in the single agent case, there are two scenarios under which a subset of the expected payoff functions can be identified: where there is a terminal or a renewal action. Many problems in industrial organization have a terminal choice such as games where there is an exit decision. For renewal, there must be a way for one of the agents to reset the state for all players. As in the single agent setting, label the renewal action as action 1 where a renewal action at time $t+1$ satisfies:

$$\sum_{x_{t+1}=1}^{X} f_{1,t+1}^{(i)}(x_{t+2}|x_{t+1}) f_{jt}^{(i)}(x_{t+1}|x_t) = \sum_{x_{t+1}=1}^{X} f_{1,t+1}^{(i)}(x_{t+2}|x_{t+1}) f_{j't}^{(i)}(x_{t+1}|x_t) \tag{29}$$

for all x_t and all $\{j, j'\} \in [1, \ldots, J]$. When there is either a terminal or renewal choice, a subset of the expected payoff functions are identified, subject to standard normalizations.

Theorem 8 *When (29) holds or when there is a choice such that no further decisions are made, $J-1 \times \mathcal{T}-1 \times X$ expected payoffs are identified.*

Note that embedded in $f_{1,t+1}^{(i)}(x_{t+2}|x_{t+1})$ are how the state transitions are affected by the decisions by the other players. Renewal can still occur, however, when a decision by the agent makes the past decisions of the other players irrelevant. An example is making a decision to adopt a frontier technology, rendering the technologies of the other firms obsolete.

4.4 Finite dependence and games

We now apply our finite dependence results to games. Similar to the single agent case, a pair of choices $\{j, j'\}$ exhibits ρ-period dependence at state x_t if there exists a sequence of choices from j

and j' such that:

$$\kappa^{*(i)}_{t+\rho}(x_{t+\rho+1}|x_t, j) = \kappa^{*(i)}_{t+\rho}(x_{t+\rho+1}|x_t, j') \tag{30}$$

Showing a model exhibits finite dependence becomes complicated 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 4: Finite dependence in a coordination game

We now give an example of how to construct finite dependence sequences in a games environment. The game we consider has two players, each of whom can decide whether or not to compete in the market. When $d_{1t}^{(i)} = 1$ player i does not compete at time t and the flow payoff i receives is $\epsilon_{1t}^{(i)}$. Note that this payoff does not depend on the other player's choice, nor does it depend on past choices of player i. The expected payoff for entering at time t, $d_{2t}^{(i)} = 2$, is given by $u_{2t}^{(i)}(x_t)$ where x_t are the relevant state variables. In this case, the dynamics come through the expected payoffs depending on the decisions to compete in the previous period, $x_t = \{d_{2t-1}^{(1)}, d_{2t-1}^{(2)}\}$. The expected payoffs for competing are given by:

$$u_{2t}^{(i)}(x_t) = \sum_{j=1}^{2} p_{jt}^{(i)}(x_t) U_{2t}^{(i)} \left(d_{2t-1}^{(i)}, j \right) + \epsilon_{2t}^{(i)} \tag{32}$$

where the flow payoff for entering is allowed to vary over time.

Per the discussion above, the first step in establishing finite dependence is choosing weights on the decisions at $t+1$ such that after the $t+2$ decision the 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 begin 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 normalizing the $t+2$ choice 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 9 *Finite dependence can be achieved after two periods for all x_t*

4.5 Recovery of flow payoffs

Example 5: Unbundling state-specific payoffs in an entry/exit game

To illustrate how non-stationarity aids in the recovery of flow payoffs, we consider an entry/exit game. 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:

1. whether there is another firm in the market $d_{2t}^{\sim i}$,
2. whether the firm is an incumbent and therefore does not have to pay the entry cost, d_{2t}^{i},
3. and 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.

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:

$$u_{2t}^{(i)}(x_t) = \sum_j p_{jt}^{(\sim i)}(x_t) U_2^{(i)}(j, d_{2t-1}^{(i)}, x_{1t}) \tag{33}$$

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:

$$\begin{aligned} v_{2t}^{(i)}(x_t) &= u_{2t}^{(i)}(x_t) + \beta \sum_j \sum_{x_{1t+1}} p_{jt}^{(\sim i)} V_{t+1}^{(i)}(j,1,x_{1t+1}) f_t(x_{1t+1}|x_{1t}) \\ &= u_{2t}^{(i)}(x_t) + \beta \sum_j \sum_{x_{1t+1}} p_{jt}^{(\sim i)} \psi_2 \left[p_{t+1}^{(i)}(j,1,x_{1t+1})\right] f_t(x_{1t+1}|x_{1t}) \end{aligned} \tag{34}$$

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

$$u_{2t}^{(i)}(x_t) = \psi_{1t}\left[p_t^{(i)}(x_t)\right] - \psi_{2t}\left[p_t^{(i)}(x_t)\right] - \beta \sum_j \sum_{x_{1t+1}} p_{jt}^{(\sim i)} \psi_2 \left[p_{t+1}^{(i)}(j,1,x_{1t+1})\right] f_t(x_{1t+1}|x_{1t}) \tag{35}$$

Note that under our assumptions everything on the right hand side of (35) is known. Substituting in on the left hand side with (33) yields:

$$\sum_j p_{jt}^{(\sim i)}(x_t) U_2^{(i)}(j, d_{2t-1}^{(i)}, x_{1t}) = \psi_{1t}\left[p_t^{(i)}(x_t)\right] - \psi_{2t}\left[p_t^{(i)}(x_t)\right] - \beta \sum_j \sum_{x_{1t+1}} p_{jt}^{(\sim i)} \psi_2 \left[p_{t+1}^{(i)}(j,1,x_{1t+1})\right] f_t(x_{1t+1}|x_{1t}) \tag{36}$$

There are then two unknowns on the left hand side of equation (36). 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 10 *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 $\mathcal{T} \geq 2$.*

5 Conclusion

A Proofs

Proof of Theorem 1. Substituting in for $v_{jt}(z_t) - v_{1t}(z_t)$ in (9) with the corresponding expression in (10) 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:

$$u_{jt}(z_t) = \psi_1[p_t(z_t)] - \psi_j[p_t(z_t)] + \sum_{\tau=t+1}^{T} \sum_{z_\tau=1}^{Z} \beta^{\tau-t} \psi_1[p_\tau(z_\tau)] \left[\kappa^*_{\tau-1}(z_\tau|z_t, 1) - \kappa^*_{\tau-1}(z_\tau|z_t, j)\right] \tag{37}$$

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:

$$u_j = \Psi_j - \Psi_1 + \sum_{\tau=1}^{\infty} \beta^\tau \left(F_1 - F_j\right) F_1^{\tau-1} \Psi_1 = \Psi_j - \Psi_1 + \beta \left(F_1 - F_j\right) \left(\sum_{\tau=0}^{\infty} \beta^\tau F_1^\tau\right) \Psi_1 \tag{38}$$

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 $[\mathcal{I} - \beta F_1]^{-1}$ follows from Hadley (page 118, 1961) with:

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

Substituting the expression for Q into (38) we obtain:

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

which proves the theorem. ■

Proof of Theorem 5.

We first write down the conditional value function for working at time t with zero years of work experience, $v_{2t}(0)$, normalizing the flow payoff for not working with zero years of work experience to zero.

$$\begin{aligned}
v_{2t}(0) &= u_2(0) + \beta\pi_t\left[\lambda V_{t+1}(1) + (1-\lambda)v_{1t+1}(1)\right] + \beta(1-\pi_t)\left[\lambda V_{t+1}(0) + (1-\lambda)v_{1t+1}(0)\right] \\
&= u_2(0) + \beta\pi_t\left[\lambda\psi_1(p_{t+1}(1)) + v_{1t+1}(1)\right] + \beta(1-\pi_t)\left[\lambda\psi_1(p_{t+1}(0)) + v_{1t+1}(0)\right] \\
&= u_2(0) + \beta\pi_t u_1(1) + \beta(1-\pi_t)u_1(0) + \beta\pi_t\lambda\psi_1(p_{t+1}(1)) + \beta(1-\pi_t)\psi_1(p_{t+1}(0)) \qquad (39) \\
&\quad +\beta^2\pi_1\left[\lambda V_{t+2}(1) + (1-\lambda)v_{1t+2}(1)\right] + \beta^2(1-\pi_1)\left[\lambda V_{t+2}(0) + (1-\lambda)v_{1t+2}(0)\right]
\end{aligned}$$

After expressing the corresponding finite dependence path for an initial decision of staying home, differencing the two expressions will result in the bottom line of (39) canceling out.

We can express $v_{1t}(0)$ as follows, again treating $u_1(0)$ as zero:

$$\begin{aligned}
v_{1t}(0) &= \beta\lambda\left(\frac{\pi_t}{\pi_{t+1}\lambda}\left[\psi_2(p_{t+1}(0)) + v_{2t+1}(0)\right] + \left(1 - \frac{\pi_t}{\pi_{t+1}\lambda}\right)\left[\psi_1(p_{t+1}(0)) + v_{1t+1}(0)\right]\right) + \beta(1-\lambda)v_{1t+1}(0) \\
&= \frac{\beta\pi_t}{\pi_{t+1}}\left[\psi_2(p_{t+1}(0)) + u_2(0)\right] + \beta^2(1-\lambda)\left[\lambda V_{t+2}(0) + (1-\lambda)v_{1t+2}(0)\right] \\
&\quad +\beta^2\pi_t\left[\lambda V_{t+2}(1) + (1-\lambda)v_{1t+2}(1)\right] + \frac{\beta^2\pi_t(1-\pi_{t+1})}{\pi_{t+1}}\left[\lambda V_{t+2}(0) + (1-\lambda)v_{1t+2}(0)\right] \\
&\quad +\beta\lambda\left(1 - \frac{\pi_t}{\pi_{t+1}\lambda}\right)\left[\psi_1(p_{t+1}(0)) + \lambda\beta V_{t+2}(0) + (1-\lambda)\beta v_{1t+2}(0)\right] \\
&= \frac{\beta\pi_t u_2(0)}{\pi_{t+1}} + \frac{\beta\pi_t\psi_2(p_{t+1}(0))}{\pi_{t+1}} + \beta\left(\lambda - \frac{\pi_t}{\pi_{t+1}}\right)\psi_1(p_{t+1}(0)) \qquad (40) \\
&\quad +\beta^2\pi_t\left[\lambda V_{t+2}(1) + (1-\lambda)v_{1t+2}(1)\right] + \beta^2(1-\pi_t)\left[\lambda V_{t+2}(0) + (1-\lambda)v_{1t+2}(0)\right]
\end{aligned}$$

Differencing (40) from (39) and recognizing that $v_{2t}(0) - v_{1t}(0) = \psi_1(p_t(0)) - \psi_2(p_t(0))$ yields:

$$\begin{aligned}
\psi_1(p_t(0)) - \psi_2(p_t(0)) &= u_2(0)\left(1 - \frac{\beta\pi_t}{\pi_{t+1}}\right) + \beta\pi_t u_1(1) + \beta\pi_t\lambda\psi_1(p_{t+1}(1)) \qquad (41) \\
&\quad +\beta\left(1 - \lambda - \pi_t + \frac{\pi_t}{\pi_{t+1}}\right)\psi_1(p_{t+1}(0)) - \frac{\beta\pi_t\psi_2(p_{t+1}(0))}{\pi_{t+1}}
\end{aligned}$$

Forming the similar difference at t' yields:

$$\begin{aligned}\psi_1(p_{t'}(0))-\psi_2(p_{t'}(0)) &= u_2(0)\left(1-\frac{\beta\pi_{t'}}{\pi_{t'+1}}\right)+\beta\pi_{t'}u_1(1)+\beta\pi_{t'}\lambda\psi_1(p_{t'+1}(1)) \\ &\quad +\beta\left(1-\lambda-\pi_{t'}+\frac{\pi_{t'}}{\pi_{t'+1}}\right)\psi_1(p_{t'+1}(0))-\frac{\beta\pi_{t'}\psi_2(p_{t'+1}(0))}{\pi_{t'+1}}\end{aligned} \tag{42}$$

Since by assumption $\pi_t > 0$ for all t and $\pi_{t'} \neq \pi_t$, the system of equations defined by (41) and (42) is of full rank and both $u_2(0)$ and $u_1(1)$ can be recovered. Then, proceeding by induction we can recover the remaining flow payoff for experience levels of one and above. ■

Proof of Theorem 9. We can establish that a finite dependence path exists by showing that the rank of:

$$\mathcal{P}^2_{t+2}\left[\begin{array}{cc}\mathcal{F}^{(1)}_{t+1}(2) & -\mathcal{F}^{(1)}_{t+1}(1)\end{array}\right]$$

is one.

We begin by defining the terms in the above expression:

$$\mathcal{P}^{(2)}_{t+2}=\left[\begin{array}{cccc}p^{(2)}_{2t+2}(2,2) & p^{(2)}_{2t+2}(2,1) & p^{(2)}_{2t+2}(1,2) & p^{(2)}_{2t+2}(1,1)\end{array}\right] \tag{43}$$

$$\left[\begin{array}{cc}\mathcal{F}^{(1)}_{t+1}(2) & -\mathcal{F}^{(1)}_{t+1}(1)\end{array}\right]=\left[\begin{array}{cccc}p^{(2)}_{2t+1}(2,2) & p^{(2)}_{2t+1}(2,1) & -p^{(2)}_{2t+1}(1,2) & -p^{(2)}_{2t+1}(1,1)\\ p^{(2)}_{1t+1}(2,2) & p^{(2)}_{1t+1}(2,1) & -p^{(2)}_{1t+1}(1,2) & -p^{(2)}_{1t+1}(1,1)\\ -p^{(2)}_{2t+1}(2,2) & -p^{(2)}_{2t+1}(2,1) & p^{(2)}_{2t+1}(1,2) & p^{(2)}_{2t+1}(1,1)\\ -p^{(2)}_{1t+1}(2,2) & -p^{(2)}_{1t+1}(2,1) & p^{(2)}_{1t+1}(1,2) & p^{(2)}_{1t+1}(1,1)\end{array}\right] \tag{44}$$

These terms will then multiply:

$$\left[\begin{array}{c}D^{(1)}_{t+1}(2,x_t)\\ D^{(1)}_{t+1}(1,x_t)\end{array}\right]=\left[\begin{array}{c}D^*_{2t+1}(2,2)p^{(2)}_{2t}(x_t)\\ D^*_{2t+1}(2,1)p^{(2)}_{1t}(x_t)\\ D^*_{2t+1}(1,2)p^{(2)}_{2t}(x_t)\\ D^*_{2t+1}(1,1)p^{(2)}_{1t}(x_t)\end{array}\right] \tag{45}$$

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:

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

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 10. Denote $P^{(\tilde{i})}$ as a 2×2 matrix given by:

$$P^{(\tilde{i})}=\begin{bmatrix} p^{(\tilde{i})}_{1t}(x) & p^{(\tilde{i})}_{2t}(x) \\ p^{(\tilde{i})}_{1t+1}(x) & p^{(\tilde{i})}_{2t+1}(x) \end{bmatrix} \tag{46}$$

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

$$U^{(i)}_2=\begin{bmatrix} U^{(i)}_2(1,x) \\ U^{(i)}_2(2,x) \end{bmatrix} \tag{47}$$

Finally, define A as:

$$A=\begin{bmatrix} \psi_{1t}\left[p^{(i)}_t(x)\right]-\psi_{2t}\left[p^{(i)}_t(x)\right]-\beta\sum_j\sum_{x_{1t+1}}p^{(\sim i)}_{jt}\psi_2\left[p^{(i)}_{t+1}(j,1,x_{1t+1})\right]f_t(x_{1t+1}|x) \\ \psi_1\left[p^{(i)}_{t+1}(x)\right]-\psi_2\left[p^{(i)}_{t+1}(x)\right]-\beta\sum_j\sum_{x_{1t+2}}p^{(\sim i)}_{jt+1}\psi_2\left[p^{(i)}_{t+2}(j,1,x_{1t+2})\right]f_{t+1}(x_{1t+2}|x) \end{bmatrix} \tag{48}$$

The system of equation is then:

$$P^{(\tilde{i})}U^{(i)}_2=A \tag{49}$$

Since by assumption the choice probabilities vary between t and $t+1$, the rank of $P^{(\tilde{i})}$, implying we can invert $P^{(\tilde{i})}$ and solve for $U_2^{(i)}$. ■

References

[1] 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.

[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] 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.

[6] 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).

[7] Bajari, P., L. Benkard, and J. Levin (2007): " Estimating Dynamic Models of Imperfect Competition," *Econometrica,* 75, 1331-1371.

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

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

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

[11] 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.

[12] Norets, A. and X. Tang (2010): "Semiparametric Inference in Dynamic Binary Choice Models", working paper.

[13] 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.

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

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

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

[17] 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.