University of Toronto
Department of Economics

Working Paper 297

Dynamic Discrete Choice Structural Models: A Survey

By Victor Aguirregabiria and Pedro mira

July 27, 2007

Dynamic Discrete Choice Structural Models: A Survey

Victor Aguirregabiria
University of Toronto

Pedro Mira
CEMFI, Madrid

July 26, 2007

Abstract

This paper reviews methods for the estimation of dynamic discrete choice structural models and discusses related econometric issues. We consider single agent models, competitive equilibrium models and dynamic games. The methods are illustrated with descriptions of empirical studies which have applied these techniques to problems in different areas of economics. Programming codes for the estimation methods are available in a companion web page.

Keywords: Dynamic structural models; Discrete choice; Estimation methods.

JEL: C14, C25, C61.

1 Introduction

This paper reviews recent developments in the literature on the estimation of discrete choice dynamic programming models of individual behavior. The goal of this paper is to provide an update of existing surveys of this literature, e.g. Eckstein and Wolpin (1989) and Rust (1994a).[1] In order to avoid repetition, we emphasize the methodological contributions during the last decade. Thus, some of the major themes of the survey are: the extension of methods which avoid repeated full solution of the structural model in estimation; the development and increased use of simulation and approximation methods; and the exploration of techniques that allow researchers to estimate dynamic equilibrium models, both strategic and competitive. This paper tries to make the reader familiar with these recent developments. With that purpose in mind, this survey is complemented with programs that implement the estimation methods we describe. These programs are available at the journal's web site.

In dynamic discrete choice structural models, agents are forward looking and maximize expected intertemporal payoffs. The parameters to be estimated are *structural* in the sense that they describe agents' preferences and beliefs about technological and institutional constraints. Under the principle of *revealed preference*, these parameters are estimated using micro data on individuals' choices and outcomes. Thus an attractive feature of this literature is that structural parameters have a

[1] Other excellent surveys are Rust (1994b), Pakes (1994) and Miller (1997).

transparent interpretation within the theoretical model that frames the empirical investigation. Moreover, econometric models in this class are useful tools for the evaluation of new (counterfactual) policies.[2] Seminal papers include Wolpin (1984) on fertility and child mortality, Miller (1984) on occupational choice, Pakes (1986) on patent renewal, and Rust (1987) on machine replacement. A well known impediment to the development of this literature has been the computational complexity of estimation. Solving the structural model or evaluating an estimation criterion such as the likelihood can *both* be non-trivial numerical tasks, and estimation in this context typically requires the use of algorithms in which a dynamic programming solution procedure is nested in the optimization of the estimation criterion. In spite of this, over the last twenty years there has been a significant number of interesting applications of these models to different areas in economics. In this paper, we will select a few of these applications as examples to illustrate estimation methods and econometric issues. Many of the problems encountered by applied researchers are the same as in other discrete choice microeconometric models, e.g., permanent unobserved heterogeneity, initial conditions, censored outcomes and sample selection, measurement error, endogeneity, identification, etc. Having to consider explicit solutions to the dynamic optimization problem that is postulated to describe individual behavior adds another layer of complexity. But such a close link between economic theory and the econometric model can also provide more insight into the econometric problems.

Our discussion of estimation methods will follow a classification based on two criteria. The first criterion is the type of interactions between agents which the structural model explicitly takes into account. According to this criterion we distinguish single-agent models, dynamic games, and competitive equilibrium models. The estimation of equilibrium models and dynamic strategic games involves specific econometric issues which are not present in single agent models. In particular, multiple equilibria, the endogeneity of other players' actions and prices, and the curse of dimensionality associated with the number of heterogeneous players and the state of the economy.

A second natural classification criterion is the structure of the unobservables. Different assumptions on the structure of the unobservables lead to very different estimation methods, as already noted by Eckstein and Wolpin (1989). Multi-dimensional numerical integration and the so called *initial conditions problem* are important issues that arise in dynamic discrete choice models (structural or not) in which unobservables are correlated across choices or over time. Dealing with these issues can add considerable complexity to the solution and estimation of dynamic discrete choice structural models.

[2] See Wolpin (1996) for a review of some uses of these models for public policy anlaysis.

Our discussion of estimation methods for single agent problems proceeds in three steps. First, we review methods for *Rust's model* with additively separable, conditionally independent and extreme value distributed unobservables. Second, we consider several important departures from Rust's framework, such as allowing for permanent unobserved heterogeneity, non additive shocks, correlation across choices and observable but choice-censored state-variables or payoffs. We group all such models under the label *Eckstein-Keane-Wolpin models* after the authors who have contributed many applications and methodological advances. Third, we discuss estimation of models with serially correlated unobservables and continuous state variables. While we can compute exactly the solution of a discrete, finite horizon dynamic decision model, the solution to models with infinite horizons and continuous observable state variables is always an approximation and most of the methods and applications in the literature have considered specifications where continuous state variables are discretized. Approximation errors may have an effect on inferences and conclusions in applied work. We include a brief discussion of this issue.

For the sake of space, this paper does not cover several topics in this literature which have received attention in recent years. Some omissions that we are aware of are nonparametric and semiparametric identification and estimation (see Magnac and Thesmar, 2002, Aguirregabiria, 2007, Bajari and Hong, 2005, and Heckman and Navarro, 2007), and application of parallel computing (see Ferrall, 2005).

Section 2 introduces the notation and basic assumptions, illustrates the main issues that arise in estimation of dynamic discrete choice structural models and presents four examples of applications. Each of these examples deals with one of the four classes of models that we examine in the survey: single-agent models under Rust's framework; single-agent models under the Eckstein-Keane-Wolpin framework; dynamic general equilibrium models; and dynamic strategic games. This section is self-contained and tries to provide an introduction to this literature for a second year PhD student. Sections 3, 4 and 5 deal with the detail of methods for the estimation of single-agent models, dynamic games and general equilibrium models, respectively. The idea is that, after reading section 2, the reader can go to either of these sections to learn about the details of specific methods.

2 Models and examples

2.1 Single agent models

Time is discrete an indexed by t. We index agents by i. Agents have preferences defined over a sequence of states of the world from period $t = 0$ until period $t = T$. The time horizon T can be

either finite or infinite. The state of the world at period t for individual i has two components: a vector of state variables s_{it} that is known at period t; and a decision a_{it} chosen at period t that belongs to the discrete set $A = \{0, 1, ..., J\}$. The time index t can be a component of the state vector s_{it}, which may also contain time-invariant individual characteristics. Agents' preferences over possible sequences of states of the world can be represented by a utility function $\sum_{j=0}^{T} \beta^j \, U(a_{i,t+j}, s_{i,t+j})$, where $\beta \in (0,1)$ is the discount factor and $U(a_{it}, s_{it})$ is the current utility function.[3] The decision at period t affects the evolution of future values of the state variables, but the agent faces uncertainty about these future values. The agent's beliefs about future states can be represented by a Markov transition distribution function $F(s_{i,t+1}|a_{it}, s_{it})$.[4] These beliefs are *rational* in the sense that they are the true transition probabilities of the state variables.[5] Every period t the agent observes the vector of state variables s_{it} and chooses his action $a_{it} \in A$ to maximize the expected utility

$$E\left(\sum_{j=0}^{T-t} \beta^j \, U(a_{i,t+j}, s_{i,t+j}) \mid a_{it}, s_{it}\right). \tag{1}$$

This is the agent's dynamic programming (DP) problem. Let $\alpha(s_{it})$ and $V(s_{it})$ be the optimal decision rule and the value function of the DP problem, respectively.[6] By Bellman's principle of optimality the value function can be obtained using the recursive expression:

$$V(s_{it}) = \max_{a \in A} \left\{ U(a, s_{it}) + \beta \int V(s_{i,t+1}) \, dF(s_{i,t+1} \mid a, s_{it}) \right\} \tag{2}$$

and the optimal decision rule is then $\alpha(s_{it}) = \arg\max_{a \in A} \{v(a, s_{it})\}$ where, for every $a \in A$,

$$v(a, s_{it}) \equiv U(a, s_{it}) + \beta \int V(s_{i,t+1}) \, dF(s_{i,t+1} \mid a, s_{it}) \tag{3}$$

is a choice-specific value function.

We are interested in the estimation of the structural parameters in preferences, transition probabilities, and the discount factor β.[7] Suppose that a researcher has panel data for N individuals

[3] We do not consider deviations from this expected utility, time-separable framework. Note that β is constant over time, e.g., hyperbolic discounting is ruled out.

[4] The utility function and the transition probability functions are not indexed by time or individual. This is without loss of generality because time and individual heterogeneity can be arguments in the state vector s_{it}.

[5] If the only data available are (longitudinal) data on choices and states, then preferences, beliefs and the actual transition probabilities cannot be separately identified in general. In this sense, rational expectations is an identification assumption. It can be relaxed if the researcher has data on elicited beliefs and is able to use that data in estimation. See Manski (2004) and Delavande (2006) for a discussion and example.

[6] Again, and without loss of generality, we omit time and individual subindexes from the arguments of these functions because they are implicit in the state vector s_{it}.

[7] The discount factor is assumed constant across agents. In most applications this parameter is not estimated because it is poorly identified (e.g., see Rust, 1987). Also note that the decision horizon T is the same for all agents and known by the econometrician.

who behave according to this decision model. For every observation (i,t) in this panel dataset, the researcher observes the individual's action a_{it} and a subvector x_{it} of the state vector s_{it}. Therefore, from an econometric point of view, we can distinguish two subsets of state variables: $s_{it} = (x_{it}, \varepsilon_{it})$, where the subvector ε_{it} is observed by the agent but not by the researcher. Note that ε_{it} is a source of variation in the decisions of agents conditional on the variables observed by the researcher. It is the model's 'econometric error', which is given a structural interpretation as an unobserved state variable.[8] In some applications the researcher also observes one or more *payoff variables*. We define a *payoff variable* as a variable y_{it} which contains information about utility but is not one of the model's actions or state variables.[9] Payoff variables depend on current action and state variables. We specify this relationship as $y_{it} = Y(a_{it}, x_{it}, \varepsilon_{it})$, where $Y(.)$ is the *payoff function*, e.g., an earnings function, a production function, etc.[10] In summary, the researcher's dataset is:

$$Data = \{\ a_{it}\ ,\ x_{it}\ ,\ y_{it} : i = 1, 2, ..., N\ ;\ t = 1, 2, ..., T_i\} \tag{4}$$

where T_i is the number of periods over which we observe individual i.[11] In microeconometric applications of single-agent models, we typically have that N is large and T_i is small.

Let θ be the vector of structural parameters and let $g_N(\theta)$ be an estimation criterion for this model and data, such as a likelihood or a GMM criterion. For instance, if the data are a random sample over individuals and the criterion is a log-likelihood, then $g_N(\theta) = \sum_{i=1}^{N} l_i(\theta)$, where $l_i(\theta)$ is the contribution to the log-likelihood function of individual i's history:

$$\begin{aligned} l_i(\theta) &= \log \Pr\left(a_{it}, y_{it}, x_{it} : t = 1, 2, ..., T_i \mid \theta\right) \\ &= \log \Pr\left(\alpha\left(x_{it}, \varepsilon_{it}, \theta\right) = a_{it},\ Y\left(a_{it}, x_{it}, \varepsilon_{it}, \theta\right) = y_{it},\ x_{it} : t = 1, 2, ..., T_i \mid \theta\right) \end{aligned} \tag{5}$$

Whatever the estimation criterion, in order to evaluate it for a particular value of θ it is necessary to know the optimal decision rules $\alpha(x_{it}, \varepsilon_{it}, \theta)$. Therefore, for each trial value of θ the DP problem needs to be solved exactly, or its solution approximated in some way.

So far we have not made any assumption on the relationship between observable and unobservable variables. These are key modelling decisions in the econometrics of dynamic discrete structural

[8] See Rust (1994) for a discussion of alternative interpretations of the econometric error.

[9] That is, we can write $U(a_{it}, s_{it})$ as $\tilde{U}(y_{it}, a_{it}, s_{it})$. For instance, in a model of firm behavior the researcher may observe firms' output, revenue or the wage bill; or in a model of individual behavior the econometrician may observe individual earnings.

[10] The payoff function can also incorporate stochastic components such as measurement error in payoff variables or structural innovations which are not state variables because they are iid over time and unknown to the agent when he makes his decision.

[11] Here the index t sequences each individual's observations. Strictly speaking, it is not the same as the time period index in the description of the structural model, but distinguishing between the two at this stage would make the notation unnecessarily cumbersome.

models. The form of $l_i(\theta)$ and the choice of the appropriate solution and estimation methods crucially depend on these assumptions. We first introduce six assumptions that describe what we define as *Rust's model*. This is the simplest framework for estimation and it has been used in many applications beginning with the bus engine replacement model in Rust (1987).

ASSUMPTION AS (Additive separability): The one-period utility function is additively separable in the observable and unobservable components: $U(a, x_{it}, \varepsilon_{it}) = u(a, x_{it}) + \varepsilon_{it}(a)$, where $\varepsilon_{it}(a)$ is a zero mean random variable with support the real line.[12] That is, there is one unobservable state variable for each choice alternative, so the dimension of ε_{it} is $(J+1) \times 1$.

ASSUMPTION IID (iid unobservables): The unobserved state variables in ε_{it} are independently and identically distributed over agents and over time with CDF $G_\varepsilon(\varepsilon_{it})$ which has finite first moments and is continuous and twice differentiable in ε_{it}.[13]

ASSUMPTION CI-X (Conditional independence of future x): Conditional on the current values of the decision and the observable state variables, next period observable state variables do not depend on current ε: i.e., $CDF(x_{i,t+1}|a_{it}, x_{it}, \varepsilon_{it}) = F_x(x_{i,t+1}|\ a_{it}, x_{it})$. We use θ_f to represent the vector of parameters that describe the transition probability function F_x.

ASSUMPTION CI-Y (Conditional independence of y): Conditional on the values of the decision and the observable state variables, the value of the payoff variable y is independent of ε: i.e., $Y(a_{it}, x_{it}, \varepsilon_{it}) = Y(a_{it}, x_{it})$. The vector of parameters that describe Y is θ_Y.

ASSUMPTION CLOGIT: The unobserved state variables $\{\varepsilon_{it}(a) : a = 0, 1, ..., J\}$ are independent across alternatives and have an extreme value type 1 distribution.

ASSUMPTION DIS (Discrete support of x): The support of x_{it} is discrete and finite: $x_{it} \in X = \{x^{(1)}, x^{(2)}, ..., x^{(|X|)}\}$ with $|X| < \infty$.

Note that assumptions IID and CI-X together imply that $F(x_{i,t+1}, \varepsilon_{i,t+1}|a_{it}, x_{it}, \varepsilon_{it}) = G_\varepsilon(\varepsilon_{i,t+1})$ $F_x(x_{i,t+1}|\ a_{it}, x_{it})$, what corresponds to the conditional independence assumption in Rust (1987). It is convenient to distinguish several components in the vector of structural parameters: $\theta = \{\theta_u, \theta_Y, \theta_f\}$, where θ_Y and θ_f have been defined above and θ_u represents all the parameters in the utility function which are not in θ_Y as well as the parameters in the distribution of G_ε. In

[12] Unbounded support is as important as additive separability in assumption AS. To see this, suppose that $U(a, x_{it}, \varepsilon_{it}) = u(a, x_{it}) \exp\{\varepsilon_{it}(a)\}$ with $E(\exp\{\varepsilon_{it}(a)\}|x_{it}) = 1$. It is clear that we can rewrite this utility function as $u(a, x_{it}) + \varepsilon^*_{it}(a)$, where $\varepsilon^*_{it}(a)$ is equal to $u(a, x_{it}) [\exp\{\varepsilon_{it}(a)\} - 1]$. Note that $\varepsilon^*_{it}(a)$ is additive in the utility and it is a zero mean random variable. However, its range of variation is not the real line but only the positive real numbers.

[13] Rust (1994) presents a weaker version of this assumption where the second and higher moments of ε_{it} may depend on x_{it}. However, this weaker version of the assumption is hardly ever used in practive.

order to illustrate this framework, we present an example of a model of retirement from the labor force which is based on the models of Rust and Phelan (1997) and Karlstrom, Palme and Svensson (2004).

EXAMPLE 1. (Retirement from the labor force). Every period the individual decides whether to continue working ($a_{it} = 1$) or to retire and start collecting social security pension benefits ($a_{it} = 0$). This is an optimal stopping problem with a finite horizon, where $t = 1$ is the earliest age at which the individual can retire and T is age at death.[14] Let ra_{it} denote the individual's retirement status. If he has not retired yet, $ra_{it} = 0$. If retired, ra_{it} is the age at which he retired. The utility function is additively separable in consumption and leisure. More specifically,

$$u(a_{it}, x_{it}) = E\left(c_{it}^{\theta_{u1}} | a_{it}, x_{it}\right) \ \exp\left\{\theta_{u2} + \theta_{u3} h_{it} + \theta_{u4} m_{it} + \theta_{u5} \frac{t_{it}}{1 + t_{it}}\right\} - \theta_{u6} \ a_{it} \tag{6}$$

c_{it} is current consumption. θ_{u1} is the coefficient of relative risk aversion. The expression in the exponential term captures individual heterogeneity in the marginal utility of consumption. In particular, h_{it} is an indicator of good health status, m_{it} is an indicator of marital status, and t_{it} is age. Finally, the last term is associated with the utility of leisure. θ_{u6} is the disutility of working. In our notation, $\theta_u = (\theta_{u1}, \theta_{u2}, ..., \theta_{u6})$. The unobservable state variables $\varepsilon_{it}(1)$ and $\varepsilon_{it}(0)$ enter additively in the utilities of working and not working, respectively, and they can be interpreted as transitory and idiosyncratic shocks to the utility from leisure. These random variables are independently distributed over time and over individuals with an extreme value distribution.

Consumption is equal to current income (y_{it}) minus health care expenditures net of insurance reimbursements (hc_{it}): $c_{it} = y_{it} - hc_{it}$. If the individual works, his income is equal to stochastic labor earnings (w_{it}), hence the expectation $E\left(c_{it}^{\theta_{u1}} | a_{it}, x_{it}\right)$. If the individual decides to retire, then his earnings are equal to social security pension benefits b_{it}. The econometrician observes earnings y_{it} and $y_{it} = a_{it} w_{it} + (1 - a_{it}) b_{it}$. Labor earnings depend on age, health status, marital status, past earnings history through pension points and an unobservable shock ξ_{it}. In particular,

$$w_{it} = \exp\left\{\theta_{w1} + \theta_{w2} h_{it} + \theta_{w3} m_{it} + \theta_{w4} \frac{t_{it}}{1 + t_{it}} + \theta_{w5} pp_{it} + \xi_{it}\right\} \tag{7}$$

Earnings y_{it} are a payoff variable and $\theta_Y = (\theta_{w1}, \theta_{w2}, ..., \theta_{w5}, \sigma_\xi^2)$ is the vector of parameters in the corresponding payoff function, where σ_ξ^2 is the variance of ξ_{it}. Retirement benefits depend on retirement age (ra_{it}) and on *pension points or social security wealth* pp_{it}: $b_{it} = b(ra_{it}, pp_{it})$. The

[14] Rust and Phelan allow for uncertainty about the age at death. In their specification, the hazard of death varies with age and T is a terminal age at which the probability of death is 1.

form of the function $b(.,.)$ depends on the rules of the pension system, which are known to the econometrician. Health care expenditures are stochastic, with a Pareto distribution conditional on health and marital status.

The vector of observable state variables is $x_{it} = \{h_{it}, hc_{it}, m_{it}, t_{it}, ra_{it}, pp_{it}\}$. Age and retirement age have obvious deterministic transition rules. Pension points are a deterministic function of past earnings history. However, as argued by Rust and Phelan (1997) and Rust et al (2000), for many pension systems, including the US system, the transition rule of pension points can be very closely approximated by a Markov process with transition probability function $F_{pp}(pp_{i,t+1}|w_{it}, pp_{it})$. This transition probability is nonparametrically specified. Health status and marital status follow first order Markov processes with transition probabilities $F_h(h_{i,t+1}|h_{it})$ and $F_m(m_{i,t+1}|m_{it})$ which are nonparametrically specified. Though health expenditures and pension points are continuous variables, Rust and Phelan discretize these variables. An important feature of this model is that the shock to wages ξ_{it} is assumed serially uncorrelated, independent of the state variables x_{it} and ε_{it}, and unknown to the individual at the time he makes his period t decision. Therefore, ξ_{it} determines the transition of labor earnings but it is not a state variable and assumption CI-Y holds: i.e., conditional of (a_{it}, x_{it}) observed earnings are independent of the unobserved state variables in ε_{it}.

Rust and Phelan use a richer specification of this model which allows for part time work and post-retirement work and includes a detailed description of Social Security and Medicare benefits to help explain several aspects of retirement behavior in the US. Counterfactual experiments based on their estimated model suggest that the peak in retirement behavior at age 65 is largely due to the fact that Social Security benefits are actuarially unfair after age 65 and to the fact that 'health insurance constrained' individuals have to wait to age 65 in order to apply for Social Security benefits and qualify for Medicare. Karlstrom, Palme and Svensson (2004) use their model to simulate the effect of a three year delay in retirement benefits.■

Imposing additive separability implies that the marginal utility of observable state variables, today and in future periods, does not depend on unobservables. For instance, in Example 1 the decision to retire or to continue working determines the current and future level of consumption. However, additive separability in the econometric model implies that observed variation in retirement choices cannot be linked to unobserved heterogeneity in the marginal utility of consumption. Relatedly, an individual considering early retirement in this model would presumably weigh uncertainty about the value of her marginal utility of consumption in the future. Additive separability

implies that unobserved state variables induce no uncertainty about the marginal utility of consumption, which may have an effect on the patterns of behavior the model can explain and on estimation of structural parameters such as the coefficient of relative risk aversion. As in the case of assumption CI-X which we discuss next, assumption AS may not be too restrictive if the model specification and the data are sufficiently rich in *observable* explanatory variables; i.e., observable state variables which are unconstrained by assumptions AS, IID and CI-X, vary across individuals and time and correlate with behavior.

Assumptions IID and CI-X restrict the joint transition probability of state variables. These restrictions have two main implications which can be illustrated in Example 1. First, the unobserved shocks to the utility of leisure are serially uncorrelated, i.e., transitory. And second, the probability that a person's health or marital status will change between the current period and the next one could depend on whether this person is currently working, but once we take this into account it does not depend on the current value of the shocks to the utility of leisure.[15] Suppose we interpret these unobservable shocks as health shocks not included in the observable h_{it}. The shocks are important enough because they can explain individuals' changes in labor supply decisions, but they should be 'transitory' since according to CI-X they cannot have any direct impact on next period's health.

An important implication of assumptions IID and CI-X is that the solution to the DP problem is fully characterized by the *integrated value function or E*max function, $\bar{V}(x_{it})$, which is the expectation of the value function over the distribution of unobservable state variables, conditional on the observable state variables: $\bar{V}(x_{it}) \equiv \int V(x_{it}, \varepsilon_{it})\ dG_\varepsilon(\varepsilon_{it})$. This function is the unique solution to the *integrated Bellman equation*:

$$\bar{V}(x_{it}) = \int \max_{a \in A} \left\{ u(a, x_{it}) + \varepsilon_{it}(a) + \beta \sum_{x_{i,t+1}} \bar{V}(x_{i,t+1})\ f_x(x_{i,t+1}|a, x_{it}) \right\} dG_\varepsilon(\varepsilon_{it}) \tag{8}$$

Under these assumptions, the size of the state space X is the relevant measure of computational complexity, and given that X is discrete and finite the DP problem can be solved exactly. The choice-specific value function in (3) can be decomposed as $v(a, x_{it}) + \varepsilon_{it}(a)$ as in static random utility models, where:[16]

$$v(a, x_{it}) = u(a, x_{it}) + \beta \sum_{x_{i,t+1}} \bar{V}(x_{it+1})\ f_x(x_{i,t+1}|a, x_{it}) \tag{9}$$

[15] In Rust and Phelan's specification the probability that a person's health or marital status will change between the current period and the next one does not depend on whether this person is currently working.

[16] There is a slight abuse of notation here because we use the same $v()$ function as in equation (3), but the alternative-specific value function here does not depend on $\varepsilon_{it}(a)$.

Another important implication of assumptions IID and CI-X is that the observable state vector x_{it} is a sufficient statistic for the current choice. The contribution of individual i to the log-likelihood function can be factored as follows:[17]

$$\begin{aligned} l_i(\theta) &= \sum_{t=1}^{T_i} \log P(a_{it}|x_{it},\theta) + \sum_{t=1}^{T_i} \log f_Y(y_{it}|a_{it},x_{it},\theta_Y) \\ &+ \sum_{t=1}^{T_i-1} \log f_x(x_{i,t+1}|a_{it},x_{it},\theta_f) + \log \Pr(x_{i1}|\theta) \end{aligned} \tag{10}$$

f_Y is the density of the payoff variable conditional on (a_{it}, x_{it}). In example 1, under assumption CI-Y, this density is quite straightforward: $I\{a_{it}=1\}\phi([\log y_{it} - \theta_{w1} - \theta_{w2} h_{it} - \theta_{w3} m_{it} - \theta_{w4} \frac{t_{it}}{1+t_{it}} - \theta_{w5} pp_{it}]/\sigma_\xi)$, where $I\{.\}$ is the indicator function and $\phi(.)$ is the density of a standard normal. This can greatly simplify the estimation of θ_Y.[18] f_x is the transition density function associated with F_x. The term $\log \Pr(x_{i1}|\theta)$ is the contribution of initial conditions to the likelihood of individual i. In most applications a conditional likelihood approach is followed and this term is ignored. No bias is incurred as long as unobservables are serially independent. Even when there is a loss of efficiency, the conditional likelihood is simpler to compute. Accordingly, hereafter the term 'likelihood' by default refers to the conditional likelihood. The term $P(a_{it}|x_{it},\theta)$ is the *Conditional Choice Probability* (CCP) obtained by integrating the optimal decision rule over the unobservable state variables. The optimal decision rule is $\alpha(x_{it},\varepsilon_{it}) = \arg\max_{a\in A} \{v(a,x_{it}) + \varepsilon_{it}(a)\}$. Therefore, for any $(a,x) \in A \times X$ and $\theta \in \Theta$, the conditional choice probability is:

$$\begin{aligned} P(a|x,\theta) &\equiv \int I\{\alpha(x,\varepsilon;\theta) = a\}\, dG_\varepsilon(\varepsilon) \\ &= \int I\{v(a,x_{it}) + \varepsilon_{it}(a) > v(a',x_{it}) + \varepsilon_{it}(a') \text{ for all } a'\}\, dG_\varepsilon(\varepsilon_{it}) \end{aligned} \tag{11}$$

When $\{\varepsilon_{it}(a)\}$ are iid type 1 extreme value random variables, as in example 1, the multi-dimension integrals in the integrated Bellman equation and in the definition of CCPs have closed form analytical expressions. This is the DP conditional logit model with Bellman equation[19]

$$\bar{V}(x_{it}) = \log\left(\sum_{a=0}^{J} \exp\left\{u(a,x_{it}) + \beta \sum_{x_{i,t+1}} \bar{V}(x_{i,t+1})\ f_x(x_{i,t+1}|a,x_{it})\right\}\right) \tag{12}$$

[17] To see this, define $\tilde{a}_{it}$, $\tilde{x}_{it}$ and $\tilde{y}_{it}$ as the vectors with the histories of the individual's decisions, states and outcomes, respectively, from period 1 until t. Assumptions IID, CI-X and CI-Y imply that $\Pr(a_{it}, y_{it}, x_{it} \mid \tilde{a}_{i,t-1}, \tilde{y}_{i,t-1}, \tilde{x}_{i,t-1}) = f_Y(y_{it} \mid a_{it}, x_{it}) \Pr(a_{it} \mid x_{it}) f_x(x_{it} \mid a_{i,t-1}, x_{i,t-1})$.

[18] Note that ξ_{it} is censored for the econometrician since wages are observed only when the individual works. However, under assumption CI-Y this censoring does not introduce any bias in the estimation of the wage equation by OLS. We discuss this issue in more detail in section 3.

[19] Since X is discrete and finite, we can represent Bellman equation as a system of equations in the Euclidean space $\mathbb{R}^{|X|}$. Let $\bar{\mathbf{V}}$ be the $|X| \times 1$ vector with the values $\bar{V}(x^{(1)})$, $\bar{V}(x^{(2)})$, ..., $\bar{V}(x^{(|X|)})$. Then, $\bar{\mathbf{V}}$ is the unique solution to: $\bar{\mathbf{V}} = \log\left(\sum_{a=0}^{J} \exp\left\{\mathbf{u}(a) + \beta\, \mathbf{F}(a)\, \bar{\mathbf{V}}\right\}\right)$, where $\mathbf{u}(a)$ is the $|X| \times 1$ vector of current utilities $(u(a,x^{(1)}), u(a,x^{(2)}), ..., u(a,x^{(|X|)}))'$ and $\mathbf{F}(a)$ is the $|X| \times |X|$ matrix with the transition probabilities $f_x(x_{t+1}|a,x_t)$.

and choice probabilities:

$$P(a|x_{it},\theta) = \frac{\exp\{v(a,x_{it})\}}{\sum_{j=0}^{J}\exp\{v(j,x_{it})\}} \tag{13}$$

If $v(a,x_{it})$ were a linear function of the parameters θ, these expressions would be familiar as the choice probability of binary probit, logit or multinomial logit models. In general, $v(a,x_{it})$ is a complex non-linear function of θ which has to be computed from the Bellman equation in (12).

The relative simplicity of dynamic discrete choice models under Rust's assumptions and their similarity with static discrete choice econometric models, has contributed to more extensive development of this framework. The econometric theory is better understood, and some general results on identification are available (see Rust, 1994, Magnac and Thesmar, 2002, and Aguirregabiria, 2007). Factorization of the likelihood in (10) allows for a computationally advantageous estimation approach. Our review of estimation methods begins in section 3.1 with methods for this class of models. One line of research has developed estimation methods which avoid repeated solution of the DP problem (see sections 3.1.2-3.1.4), and this computational advantage has made Rust's model the framework of choice in recent research on models with strategic interactions (see section 4). If the choice is not binary, an important restriction is involved in the CLOGIT assumption, i.e., that the unobservable state variables are independent across alternatives.

Our next example, based on Keane and Wolpin (1997), serves as an illustration of models which relax some of the assumptions in Rust's framework. More specifically, we highlight four departures from the previous model: (1) unobservables which do not satisfy assumption AS; (2) observable payoff variables which are choice-censored and do not satisfy assumption CI-Y; (3) permanent unobserved heterogeneity (a departure of assumption IID); and (4) unobservables which are correlated across choice alternatives (i.e., no CLOGIT assumption). Many applications have included at least one of these four features, and in this survey we group all of them under the label *Eckstein-Keane-Wolpin (EKW)* after the authors who are the main contributors. We review estimation methods for EKW models in section 3.2.

EXAMPLE 2. (Occupational choice and the career decisions of young men). Each period (year), starting at age 16 through a maximum age T, an individual chooses between staying at home ($a_{it}=0$), attending school ($a_{it}=4$), or working at one of three occupations: white collar ($a_{it}=1$), blue collar ($a_{it}=2$), or the military ($a_{it}=3$). The specification of the one-period utility

function is:

$$\begin{aligned} U(0, s_{it}) &= \omega_i(0) + \varepsilon_{it}(0) \\ U(4, s_{it}) &= \omega_i(4) - \theta_{tc1}\ I(h_{it} \geq 12) - \theta_{tc2}\ I(h_{it} \geq 16) + \varepsilon_{it}(4) \\ U(a, s_{it}) &= W_{it}(a) \qquad \text{for } a = 1, 2, 3 \end{aligned} \tag{14}$$

h_{it} is schooling (in years). θ_{tc1} and θ_{tc2} are parameters that represent tuition costs in college and graduate school, respectively. $W_{it}(a)$ is the wage of individual i at period t in occupation a. Wages in occupation a are the product of the skill price in that occupation, r_a, and the individual's skill level for that occupation, which is an exponential function of an individual-specific endowment at age 16, schooling, experience and a transitory shock. That is:

$$W_{it}(a) = r_a\ \exp\left\{\omega_i(a) + \theta_{a1}\ h_{it} + \theta_{a2}\ k_{it}(a) - \theta_{a3}\ \left(k_{it}(a)\right)^2 + \varepsilon_{it}(a)\right\} \tag{15}$$

where $k_{it}(a)$ is cumulated work experience (in years) in occupation a. The vector $\varepsilon_{it} = \{\varepsilon_{it}(a) : a \in A\}$ contains choice-specific shocks to skill levels and to the monetary values of a year at school or at home. They are assumed to be serially uncorrelated with a joint normal distribution with zero means and unrestricted variance matrix. The vector $\omega_i = \{\omega_i(a) : a \in A\}$ contains occupation and individual-specific endowments which are fixed from age 16. This vector has a discrete support, and its probability distribution is nonparametrically specified. Both ε_{it} and ω_i are unobservable to the econometrician but observable to the individual when he makes his decision at period t. The vector of observable state variables is $x_{it} = \{h_{it}, t_{it}, k_{it}(a) : a = 1, 2, 3\}$, where t_{it} represents age. All the variables in x_{it} have a discrete and finite support. Labor earnings are observable. This payoff variable is equal to zero when $a_{it} \in \{0, 4\}$ and equal to $W_{it}(a_{it})$ when $a_{it} \in \{1, 2, 3\}$. Therefore, the payoff function is $I\{a_{it} \in \{1, 2, 3\}\} r_a \exp\{\omega_i(a) + \theta_{a1}\ h_{it} + \theta_{a2}\ k_{it}(a) - \theta_{a3}\ \left(k_{it}(a)\right)^2 + \varepsilon_{it}(a)\}$. Note that assumption CI-Y does not hold because the unobserved state variables $\varepsilon_{it}(1)$, $\varepsilon_{it}(2)$ and $\varepsilon_{it}(3)$ have a direct effect on observed labor earnings.

The opportunity cost of investing in human capital by attending school is the value of foregone earnings and work experience, or the utility of staying home. Working also has an investment value since it increases occupation-specific skills and future earnings. An individual's optimal career path is partly determined by comparative advantage embedded in endowments at age 16. Keane and Wolpin estimated this model on NLSY data.[20] They found that unobserved skill endowments at age 16 are a very important source of inequality in lifetime career paths, earnings and utility. To the

[20] Keane and Wolpin report estimates of two versions of this model. The first 'bare bones' version is the one described here. The second one includes additional features such as disutilities of switching between choices and permament unobserved heterogeneity in preferences, which are introduced in order to help the model fit the degree of persistence in choices observed in the data.

extent that they are the root source of inequality, their counterfactual policy experiments suggest that the impact of (large) college tuition subsidies on college attainment and income distribution would be rather small. ■

We now use example 2 to illustrate the practical implications for estimation of relaxing some of Rust's assumptions. The sum of the permanent and transitory unobserved skill components, $\varepsilon_{it}(a) + \omega_i(a)$, is a serially correlated state variable. The presence of autocorrelated unobservables has important implications for the estimation of structural parameters. Individuals self-select into occupations and schooling classes on the basis of persistent differences in skills which are unobserved by the researcher. Ignoring this source of self-selection can result into an overestimation of the returns to schooling and occupation-specific experience. Also, persistence in occupational choices in the data may arise because there is a disutility of switching occupations (state dependence), or because there are persistent differences in skills (unobserved heterogeneity). Failure to control for the latter, if present, would lead to biased estimates of switching disutilities.[21] With serially correlated unobservables, the probability of an individual's sequence of choices cannot be factored into a product of conditional choice probabilities as in (10). The observable state x_{it} is not a sufficient statistic for a_{it} because lagged choices contain information about the permanent components ω_i. However, conditional on ω_i the transitory components $\{\varepsilon_i(a)\}$ do satisfy assumption IID. Since ω_i has discrete support, each individual's likelihood contribution can be obtained as a finite mixture of likelihoods, each of which has the same form as in (10). If the support of ω_i is $\Omega = \{\omega^1, \omega^2, ..., \omega^L\} \subset R^{JL}$, then we have that

$$l_i(\theta, \Omega, \pi) = \log \left(\sum\nolimits_{\ell=1}^{L} L_i(\theta, \omega^\ell)\ \pi_{\ell|x_{i1}} \right) \tag{16}$$

where $\pi_{\ell|x} \equiv \Pr(\omega_i = \omega^\ell | x_{i1} = x)$; π is the vector of parameters $\{\pi_{\ell|x} : \ell = 1, 2, ..., L; x \in X\}$; and $L_i(\theta, \omega^\ell)$ is $\prod_{t=1}^{T_i} P(a_{it}|x_{it}, \theta, \omega^\ell)\ f_Y(y_{it}|a_{it}, x_{it}, \theta, \omega^\ell)\ \prod_{t=1}^{T_i-1} f_x(x_{i,t+1}|a_{it}, x_{it}, \theta_f, \omega^\ell)$.[22] Note that the conditional density of the payoff variable y_{it} depends not only on θ_Y, as in Rust's model, but on the whole vector of structural parameters θ. Conditioning on the current action a_{it} introduces selection or censoring in the payoff variable and this selection effect depends on all the parameters of the model. An important implication of this feature of the model is that (even if there is not permanent unobserved heterogeneity) θ_Y and θ_u cannot be estimated separately.

Several issues are worth highlighting here. First, in general permanent unobserved heterogene-

[21] See Heckman (1981) for the first discussion of this issue.

[22] In Example 2, the transitions of the observable state variables are all deterministic and do not depend on any structural parameter. Therefore, in this example θ_f is an empty vector and the likelihood $L_i(\theta, \omega^\ell)$ does not include the term $\prod_{t=1}^{T_i-1} f_x(x_{i,t+1}|a_{it}, x_{it}, \omega^\ell, \theta_f)$.

ity poses an initial conditions problem because the initial state x_{i1} is correlated with permanent unobserved components. This problem is avoided if the structural model has an initial period in which all the individuals have the same value of the vector x, and if this initial period is observed in the sample. Under these conditions $x_{i1} = x_1$ for every i and then $\Pr(\omega^\ell | x_{i1})$ is constant over individuals and equal to π_ℓ, i.e., the unconditional mass probability of type ℓ individuals which is a primitive structural parameter. Keane and Wolpin's model is an example of this in that all individuals start 'life' in the model at age 16 with zero experience in all occupations and the NLSY data provide histories as of age 16. If x can take different values in the behavioral model's first decision period, then the researcher needs to specify how the distribution of permanent unobserved heterogeneity varies with x_{i1}. In Keane and Wolpin's data there is variation in schooling measured at age 16. It seems likely that initial schooling is correlated with other age 16 endowments, and Keane and Wolpin allow for this as our mixture likelihood in (16) suggests. A more difficult case arises if individual histories are left-censored. We return to this issue in our review of estimation methods for finite mixture models in Section 3.2.1. The focus there is on a structure with (discrete) permanent unobserved heterogeneity and (continuous) iid transitory components, which is the simplest way of allowing for autocorrelation in unobservable state variables.

Second, in order to evaluate the mixture of likelihoods the DP problem needs to be solved as many times as the number of components in the mixture. This is the reason why permanent unobserved heterogeneity is almost always introduced with a discrete and finite support. The integrated value function or Emax function (conditional on individual's type) $\bar{V}_\ell(x_{it})$ still fully characterizes the solution to the DP problem. The integrated Bellman equation for individual type ℓ is:

$$\bar{V}_\ell(x_{it}) = \int \max_{a \in A} \left\{ U(a, x_{it}, \omega^\ell, \varepsilon_{it}) + \beta \sum_{x_{i,t+1}} \bar{V}_\ell(x_{i,t+1}) \; f_x(x_{i,t+1} | a, x_{it}, \omega^\ell) \right\} dG_\varepsilon(\varepsilon_{it}) \qquad (17)$$

The optimal decision rule is $\alpha(x_{it}, \omega^\ell, \varepsilon_{it}) = \arg\max_{a \in A} \left\{ v(a, x_{it}, \omega^\ell, \varepsilon_{it}) \right\}$, where $v(a, x_{it}, \omega^\ell, \varepsilon_{it})$ is the term in brackets $\{\}$ in equation (17). When $\varepsilon' s$ are additively separable and extreme value distributed, then we still have closed form expressions for the integrated Bellman equation and for CCPs in terms of choice-specific value functions. In the occupational choice model of example 2 the transitory shocks to skills are not additively separable and, more important, they are correlated across choices or occupations. This seems like a realistic (and even necessary) assumption to make in this application; e.g., if the transitory component of labor market skills reflects an unobservable health shock, the shock would likely have an effect on both white collar and blue collar skills. But

correlation across choices implies that integrated value functions and conditional choice probabilities do not have the convenient closed forms of McFadden's conditional logit. In this case, as well as in models with more general autocorrelation structures, repeated calculation of multidimensional integrals in the solution and/or estimation of the model is unavoidable. Section 3.2.2 reviews Keane and Wolpin's simulation and interpolation method, developed for the occupational choice model. Simulation is used in order to compute multidimensional integrals, interpolation in order to handle problems with very large state spaces.

Third, in example 2 transitory shocks to wages are observable to the decision maker and determine the optimal occupational choice. This feature of the model implies that, in contrast with example 1, there is a self-selection bias if we estimate the wage equation for an occupation by OLS using the sub-sample of individuals whose wages are observed because they chose that occupation. And fourth, in example 2 the shocks $\varepsilon_{it}(2)$, $\varepsilon_{it}(3)$ and $\varepsilon_{it}(4)$ do not satisfy assumption AS, i.e., one shock per alternative and additive separability combined with unbounded support. This implies that the model is *saturated* (Rust 1994). That is, optimal choice probabilities $P(a|x, \omega^\ell, \theta)$ are not strictly positive for every value of $(a, x, \omega^\ell, \theta)$ and this leads to well-known complications in estimation.

So far, we have maintained the assumption that X is a discrete and finite set. However, in many applications some state variables are continuous. DP problems with continuous state variables cannot be solved exactly and the solution need to be approximated using discretization or interpolation methods. These approximation methods introduce an additional error in the estimation of the model. The implications of this error for the properties of estimators is a complicated issue because of the nonlinearity of the structural model. Section 3.3 discusses some recent developments in this area. In that section we also discuss methods for models with (time-variant) serially correlated unobservables, and Bayesian methods.

2.2 Competitive equilibrium models

The single-agent models of Examples 1 and 2 are partial equilibrium models. They study one side of the market (i.e., labor supply) taking prices and aggregate quantities as exogenously given. Though useful for policy evaluation, it is well known that partial equilibrium analysis can give misleading results if the assumption that prices are invariant to changes in the policy variables of interest is not a good approximation. Equilibrium models are also better suited to improve our understanding of economy-wide trends. Despite the limitations of partial equilibrium, there have been very few studies that specify and estimate dynamic general equilibrium (GE) models with

heterogeneous agents using micro data. An important exception is the study by Heckman, Lochner and Taber (1998) who estimate and calibrate a heterogeneous-agent dynamic GE model of human capital accumulation and earnings. They use their estimated model to study the sources of rising wage inequality in the US economy. More recently, Lee (2005) and Lee and Wolpin (2006) have estimated dynamic GE models of human capital accumulation in the same spirit as the study in Heckman, Lochner and Taber, relaxing some of their assumptions and more fully incorporating into the estimation process the equilibrium restrictions embodied in the theory.[23]

EXAMPLE 3. (Occupational choice in equilibrium). In order to make it more suitable for the analysis of economy-wide aggregate data, consider the following changes in the model of Example 2: a) Exclude the 'military' occupational choice. b) Allow the utility parameters to vary with gender and make the utility of the 'home' alternative dependent on the number of children of pre-school age n_{it}: $U(0, s_{it}) = \omega_i(0) + \theta_c\, n_{it} + \varepsilon_{it}(0)$. The observable state variable n_{it} which is now part of x_{it} follows an exogenous Markov process conditional on age, gender, education and cohort. In Example 2, the skill rental prices r_a in the wage equations (15) were assumed constants. In this example the state vector is augmented to include skill prices and, most important, the law of motion of skill prices is endogenously determined in the model as an equilibrium outcome. The labor market is assumed to be competitive. The supply side consists of overlapping generations of individuals aged 16 through 65, whose behavior corresponds to the occupational choice model. The demand side can be characterized by the following Cobb-Douglas, constant returns to scale aggregate production function:

$$Y_t = z_t\ S_{1t}^{\alpha_{1t}}\ S_{2t}^{\alpha_{2t}}\ K_t^{1-\alpha_{1t}-\alpha_{2t}} \tag{18}$$

where Y_t is aggregate output, S_{1t} and S_{2t} are the aggregate quantities of white collar and blue collar skills, and z_t and K_t are total factor productivity and the aggregate capital stock, which follow exogenously determined processes.[24] Technical change is captured (deterministically) in time-varying factor shares $\{\alpha_{1t}, \alpha_{2t}\}$ as well as in total factor productivity. Demand side competitive behavior implies that skill rental prices satisfy the value-of-marginal-product conditions:

$$r_{at} = (\alpha_{at}/S_{at})\left(z_t\ S_{1t}^{\alpha_{1t}}\ S_{2t}^{\alpha_{2t}}\ K_t^{1-\alpha_{1t}-\alpha_{2t}}\right) \qquad \text{for } a = 1, 2 \tag{19}$$

[23] This emergent microeconometric literature on estimation of dynamic GE models is related to macro literature on GE models with heterogeneous agents. See Krusell and Smith (1998) and Rios-Rull (1999). There is also an important and large literature in labor economics on the estimation of equilibrium search models with heterogeneous workers and firms. See Mortensen and Pissarides (1999) and Eckstein and van den Berg (2007) for surveys.

[24] It is implicitly assumed that the distribution of capital rental income among individuals has no effect on labor supply decisions.

Let $\widetilde{X}_t$ denote aggregate state variables relevant to the individual's occupational choice problem, i.e. current skill rental prices and other aggregate variables which agents use to predict future skill rental prices. The state vector of the occupational choice model is augmented with $\widetilde{X}_t$ and individuals are assumed to know its law of motion. The aggregate supplies of skills are obtained by adding individual skill supplies over all individuals in the economy:

$$S_{at} = \int_{x,\omega,\epsilon} k_{it}(a)\ I\left\{a = \alpha(x_{it}, \omega_i, \varepsilon_{it}, \widetilde{X}_t)\right\} \qquad \text{for } a = 1, 2 \tag{20}$$

where $k_{it}(a)$ is individual i's stock of skill in occupation a, $\alpha()$ is the optimal decision rule, and the integral represents the appropriately weighted sum over the distribution of individual state variables.

In a (rational expectations) competitive equilibrium, the sequence of skill rental prices $\{r_{1t}, r_{2t} : t = 1, 2, ...\}$ and the law of motion of $\widetilde{X}_t$ satisfy the following conditions. First, individuals solve the occupational choice problem taking the law of motion of $\widetilde{X}_t$ as given. Second, labor markets for both white and blue collar skills clear, i.e., the resulting aggregate skill supply functions, together with skill prices, satisfy the labor demand conditions (19). And, third, the actual law of motion of $\widetilde{X}_t$ is consistent with individuals' beliefs. The equilibrium need not be stationary: besides dependence on initial conditions $\widetilde{X}_1$, there are other sources of non-stationarity such as time variation in factor shares in the aggregate production function, changes in cohort sizes, cohort effects in fertility, etc.

An important question the equilibrium model can address is the extent to which 'feedback' effects offset the impact of policies such as the college tuition subsidy evaluated by Keane and Wolpin in partial equilibrium. That is, the subsidy may increase the supply of white collar skills to the extent that it induces higher college enrollment. If the relative price of white collar skills falls as a result, this reduces the returns to schooling and college enrollment. Lee (2005) estimates the equilibrium model on CPS data and concludes that feedback effects in the US were quite small. Lee and Wolpin (2006) extend this framework, allowing for three occupations in an economy with two sectors, goods and services. They estimate the model on CPS and NLSY data and use it to study the determinants of the large growth of employment in the service sector in the US over the last 50 years. They conclude that demand factors (technological change and movements in product and capital prices) were much more important than supply factors such as changes in cohort sizes or the decline in fertility. ■

The estimation of this model is considerably more demanding than that of its partial equilibrium version in Example 2 for two reasons. First, imposing the equilibrium restrictions increases the computational burden of estimation by an order of magnitude. Second, estimation of the equilib-

rium model requires additional data which can only be obtained from different sources and having to combine multiple data sources poses some complications for estimation and inference.

Recall that the state space of the individual agent's occupational choice problem is augmented with aggregate variables $\widetilde{X_t}$, e.g. current and past values of skill prices, total factor productivity, cohort sizes, the distributions of schooling and occupation-specific experience, and other variables which predict future skill prices. Note that any such variable omitted by the econometrician is a potential source of endogeneity bias because it is likely to be correlated with current skill prices. The dimensionality of $\widetilde{X}_t$ is potentially so large as to make solution - let alone estimation - infeasible without some further simplification. Lee (2005) assumes that skill price sequences are deterministic and individual agents have perfect foresight; in this case, the state space is augmented with the sequence of current and future deterministic skill prices. Lee and Wolpin (2006) assume that stochastic skill prices are linear functions of a small number of state variables and search for equilibria within this class of pricing functions. The equilibria they consider are thus approximations to the 'full' stochastic rational expectations equilibria implied by the model. For any given values of the parameters to be estimated, the equilibrium laws of motion solve a fixed point problem which uses as input the solutions to the individual DP problems of all agents who interact in the economy. This implies that a single evaluation of the estimation criterion nests two layers of fixed point problems.[25] The estimation criterion is based on a set of moment conditions which summarize occupational choices and wages obtained from micro survey data. We review the estimation of these models in more detail in section 5.

2.3 Dynamic discrete games

The analysis of many economic and social phenomena requires the consideration of dynamic strategic interactions between a relatively small number of agents. The study of the dynamics of oligopoly industries is perhaps the most notorious example of these dynamic strategic interactions. Competition in oligopoly industries involves important investment decisions which are irreversible to a certain extent. Market entry in the presence of sunk entry costs is a simple example of this type of investment. Dynamic games are powerful tools for the analysis of these dynamic strategic interactions. Until very recently, econometric models of discrete games had been limited to relatively simple static games. Two main econometric issues explain this limited range of applications: the

[25] Multiplicity of equilibria in these models cannot be ruled out in general, but computing multiple equilibria - if they exist - is a difficult task. Because of this, this issue tends to be ignored in empirical work. Estimation is carried out under the assumption that the equilibrium is unique, and the validity of the assumption is scrutinized at the final parameter estimates. The recent literature on estimation of dynamic discrete games, which we review in sections 2.3 and 4, has been more concerned with the issue of multiple equilibria.

computational burden in the solution of dynamic discrete games, and the indeterminacy problem associated with the existence of multiple equilibria. The existence of multiple equilibria is a prevalent feature in most empirical games where best response functions are non-linear in other players' actions. Models with multiple equilibria do not have a unique reduced form and this incompleteness may pose practical and theoretical problems in the estimation of structural parameters. The computational burden in the structural estimation of games is specially severe. The dimension of the state space, and the cost of computing an equilibrium, increases exponentially with the number of heterogeneous players. An equilibrium is a fixed point of a system of best response operators and each player's best response is itself the solution to a dynamic programming problem.

In section 4 we review recent papers which have taken Rust's framework of single agent problems, combined with Hotz-Miller's estimation approach, as a starting point in the development of estimable dynamic discrete games of incomplete information. Private information state variables are a convenient way of introducing unobservables in the econometric model.[26] Furthermore, under certain regularity conditions dynamic games of incomplete information have at least one equilibrium while that is not case in dynamic games of complete information (see Doraszelski and Satterthwaite, 2003). We provide here a description of the basic framework with no payoff variables, based on suitably extended versions of the AS, IID and CI-X assumptions (also see Rust (1994) pp. 154-158).

Consider a game that is played by N players that we index by $i \in I = \{1, 2, ..., N\}$. Every period t these players decide simultaneously a discrete action. Let $a_{it} \in A = \{0, 1, ..., J\}$ be the action of player i at period t. At the beginning of period t a player is characterized by two vectors of state variables which affect her current utility: x_{it} and ε_{it}. Variables in x_{it} are common knowledge for all players in the game, but the vector ε_{it} is private information of player i. Let $x_t \equiv (x_{1t}, x_{2t}, ..., x_{Nt})$ be the vector of common knowledge state variables, and similarly define $a_t \equiv (a_{1t}, a_{2t}, ..., a_{Nt})$, and let ε_t be the vector with all players' private information. Let $U_i(a_t, x_t, \varepsilon_{it})$ be player i's current payoff function, that depends on the actions of all the players, the common knowledge state variables, and his own private information ε_{it}. A player chooses his action to maximize expected discounted intertemporal utility $E_t[\sum_{j=0}^{T-t} \beta^j U_i(a_{t+j}, x_{t+j}, \varepsilon_{i,t+j})]$, where $\beta \in (0, 1)$ is the discount factor. Players have uncertainty about other players' current and future actions, about future common knowledge state variables, and about their own future private

[26] As far as we know, there are no estimable dynamic games of complete information. The estimation of this class of games involves non trivial complications. For instance, other players' current actions are not independent of common knowledge unobservables. In contrast, unobservables which are private information state variables, independently distributed across players, can explain at least part of the heterogeneity in players' actions without generating this endogeneity problem.

information shocks. We assume that $\{x_t,\varepsilon_t\}$ follows a controlled Markov process with transition probability function $F(x_{t+1},\varepsilon_{t+1}|a_t,x_t,\varepsilon_t)$. This transition probability is common knowledge.

Players' strategies depend only on payoff relevant state variables. Let $\boldsymbol{\alpha} = \{\alpha_i(x_t, \varepsilon_{it}) : i \in I\}$ be a set of strategy functions, one for each player. Taking as given the strategies of all players other than i, the decision problem of player i is just a standard single-agent DP problem. Let $V_i^{\boldsymbol{\alpha}}(x_t, \varepsilon_{it})$ be the value function of this DP problem. The Bellman equation is $V_i^{\boldsymbol{\alpha}}(x_t, \varepsilon_{it}) = \max_{a_i \in A} \{v_i^{\boldsymbol{\alpha}}(a_i, x_t, \varepsilon_{it})\}$ where, for every $a_i \in A$,

$$\begin{aligned} v_i^{\boldsymbol{\alpha}}(a_i, x_t, \varepsilon_{it}) &\equiv E_{\varepsilon_{-it}} \{U_i(a_i, \alpha_{-i}(x_t, \varepsilon_{-it}), x_t, \varepsilon_{it})\} \\ &+ \beta\ E_{\varepsilon_{-it}} \left\{ \int V_i^{\boldsymbol{\alpha}}(x_{t+1}, \varepsilon_{t+1})\ dF(x_{t+1},\varepsilon_{t+1}|a_i, \alpha_{-i}(x_t, \varepsilon_{-it}),x_t, \varepsilon_t) \right\} \end{aligned} \tag{21}$$

and $E_{\varepsilon_{-it}}$ represents the expectation over other players' private information shocks. The best response function of player i is $b_i(x_t, \varepsilon_{it}, \alpha_{-i}) = \arg\max_{a_i \in A} \{v_i^{\boldsymbol{\alpha}}(a_i, x_t, \varepsilon_{it})\}$. This best response function gives the optimal strategy of player i if the other players behave, now and in the future, according to their respective strategies in α_{-i}. A Markov perfect equilibrium (MPE) in this game is a set of strategy functions $\boldsymbol{\alpha}^*$ such that for any player i and for any (x_t, ε_{it}) we have that $\alpha_i^*(x_t, \varepsilon_{it}) = b_i(x_t, \varepsilon_{it}, \alpha_{-i}^*)$.

We now formulate the AS, IID and CI-X assumptions in the context of this game.

ASSUMPTION AS-Game: The one-period utility function is additively separable in common knowledge and private information components: $U_i(a_t, x_t, \varepsilon_{it}) = u_i(a_t, x_t) + \varepsilon_{it}(a_{it})$, where $\varepsilon_{it}(a)$ is the a-th component of vector ε_{it}. The support of $\varepsilon_{it}(a)$ is the real line for all a.

ASSUMPTION IID-Game: Private information shocks ε_{it} are independently and identically distributed over agents and over time with CDF $G_\varepsilon(\varepsilon_{it})$ which has finite first moments and is continuous and twice differentiable in ε_{it}.

ASSUMPTION CI-X-Game: Conditional on the current values of players' actions and common knowledge state variables, next period common knowledge state variables do not depend on current private information shocks: i.e., $CDF(x_{t+1}|a_t, x_t, \varepsilon_t) = F_x(x_{t+1}|\ a_t, x_t)$.

As in the case of single-agent models, under assumptions AS, IID and CI-X the integrated value function $\bar{V}_i^{\boldsymbol{\alpha}}(x_t) \equiv \int V_i^{\boldsymbol{\alpha}}(x_t, \varepsilon_{it}) dG_\varepsilon(\varepsilon_{it})$ fully characterizes player i's DP problem. The integrated Bellman equation is $\bar{V}_i^{\boldsymbol{\alpha}}(x_t) = \int \max_{a_i \in A} \{v_i^{\boldsymbol{\alpha}}(a_i, x_t) + \varepsilon_{it}(a_i)\} dG_\varepsilon(\varepsilon_{it})$ where

$$v_i^{\boldsymbol{\alpha}}(a_i, x_t) = E_{\varepsilon_{-it}} \left\{ u_i(a_i, \alpha_{-i}(x_t, \varepsilon_{-it}), x_t) + \beta \int \bar{V}_i^{\boldsymbol{\alpha}}(x_{t+1})\ dF_x(x_{t+1}|a_i, \alpha_{-i}(x_t, \varepsilon_{-it}),x_t) \right\} \tag{22}$$

Another implication of these assumptions is that a MPE can be described as a fixed point of a mapping in the space of CCPs which 'integrate out' players' private information variables. Given a set of strategy functions $\boldsymbol{\alpha} = \{\alpha_i(x_t, \varepsilon_{it}) : i \in I\}$ we define a set of *conditional choice probabilities* $\mathbf{P}^{\boldsymbol{\alpha}} = \{P_i^{\boldsymbol{\alpha}}(a_i|x) : (i, a_i, x) \in I \times A \times X\}$ such that,

$$P_i^{\boldsymbol{\alpha}}(a_i|x) \equiv \Pr\left(\alpha_i(x_t, \varepsilon_{it}) = a_i \ | x_t = x\right) = \int I\left\{\alpha_i(x_t, \varepsilon_{it}) = a_i\right\} \ dG_\varepsilon(\varepsilon_{it}) \tag{23}$$

These probabilities represent the expected behavior of firm i from the point of view of the rest of the firms when firm i follows its strategy in $\boldsymbol{\alpha}$. Similarly, we can define player i's best response probability function as:

$$\Lambda(a_i|v_i^{\boldsymbol{\alpha}}(., x_t)) = \int I\left\{a_i = \arg\max_{j \in A}\{v_i^{\boldsymbol{\alpha}}(j, x_t) + \varepsilon_{it}(j)\}\right\} \ dG_\varepsilon(\varepsilon_{it}) \tag{24}$$

Note that the value functions $v_i^{\boldsymbol{\alpha}}(a_i, x_t)$ depend on other players' strategies only through other players' choice probabilities. That is, in equation (22) we can replace the expectation $E_{\varepsilon_{-it}}\{...\alpha_{-i}(x_t, \varepsilon_{-it})...\}$ by $\sum_{a_{-i}} \Pr(a_{-i}|x_t; \boldsymbol{\alpha})\{...a_{-i}...\}$ where $\sum_{a_{-i}}$ represents the sum over all possible values in A^{N-1}, and $\Pr(a_{-i}|x_t; \boldsymbol{\alpha}) = \prod_{j \neq i} P_i^{\boldsymbol{\alpha}}(a_j|x_t)$. To emphasize this point we will use the notation $v_i^{\mathbf{P}}$ instead $v_i^{\boldsymbol{\alpha}}$ to represent these value functions.

Let $\boldsymbol{\alpha}^*$ be a set of MPE strategies and let $\mathbf{P}^* \equiv \{P_i^*(a_i|x) : \text{for every } (i, a_i, x)\}$ be the corresponding CCPs. Then, it is straightforward to show that for every (i, a_i, x), $P_i^*(a_i|x) = \Lambda(a_i|v_i^{\mathbf{P}^*}(., x))$. That is, $\mathbf{P}^*$ is a fixed point of $\mathbf{P} = \boldsymbol{\Lambda}(v^{\mathbf{P}})$, where $\boldsymbol{\Lambda}(v^{\mathbf{P}}) \equiv \{\Lambda(a_i|v_i^{\mathbf{P}}(., x)) : \text{for every } (i, a_i, x)\}$. Given the assumptions on the distribution of private information, we can define best response probability functions which are continuous in the compact set of players' choice probabilities. By Brower's theorem, there exists at least one equilibrium. In general, the equilibrium is not unique.[27]

We now distinguish between observable and unobservable variables from the point of view of the econometrician. In principle, we could distinguish observable and unobservable components both in common knowledge and in private information state variables. We start with a more restrictive and simpler case.

ASSUMPTION OC: All the common knowledge state variables in x_t are observable to the econometrician, and all the private information variables in ε_t are unobservable.

To complete this description of the econometric model we should comment on the sampling framework. In the case of single agent models, we assumed that the econometrician has a random sample of many agents (e.g., firms, households) behaving according to the model. That is not

[27] See Doraszelski and Satterthwaite (2003) amd Aguirregabiria and Mira (2007) for more details.

the case in applications of dynamic strategic games where typically the number of players is quite small, e.g., firms in an oligopoly market, members of a family, political parties, etc. We assume that the game is played independently at different locations, indexed by m, and that we have a random sample of M of these locations. Therefore, taking into account Assumption OC, the data consist of:

$$Data = \{\ a_{mt},\ x_{mt} : m = 1, 2, ..., M\ ;\ t = 1, 2, ..., T_m\} \tag{25}$$

Note that an assumption that is implicit in this description of the data is that we observe the actions of *all* the players in the game. Our next example illustrates models with strategic interactions and is based on Aguirregabiria and Mira (2007).

EXAMPLE 4. (An entry-exit game of incomplete information). The players are firms making decisions on whether to enter, continuing to operate in, or exit from a market. The market is a small local retail market and each active firm operates at one location or store. We observe a random sample of markets, $m = 1, \ldots, M$. In each market there are N potentially active and infinitely lived firms. Every period all firms decide simultaneously whether to operate their store or not. If firm i in market m operates its store at time t ($a_{imt} = 1$), its variable profits depend on the number of firms that choose to be active and on market demand conditions, as follows: $\theta_{RS} \log(S_{mt}) - \theta_{RN} \log\left(1 + \sum_{j \neq i} a_{jmt}\right)$, where θ_{RS}, θ_{RN} are parameters. Market demand conditions are represented by market (population) size, S_{mt}. Market size is common knowledge to firms and observable to the econometrician, and it follows a first order Markov process. The parameters θ_{RS} and θ_{RN} measure the sensitivity of variable profits to market size and to the number of active competitors, respectively.[28] Total current profits of an active firm are:

$$U_{imt}(1) = \theta_{RS} \log(S_{mt}) - \theta_{RN} \log\left(1 + \sum_{j \neq i} a_{jmt}\right) - \theta_{FC,i} - \theta_{EC,i}(1 - a_{im,t-1}) + \varepsilon_{imt}(1) \tag{26}$$

where $\theta_{FC,i} - \varepsilon_{imt}$ is firm i's fixed operating cost which has two components: $\theta_{FC,i}$ is time-invariant and common knowledge, and $\varepsilon_{imt}(1)$ is private information of firm i and time-varying. The term $(1 - a_{i,t-1})\,\theta_{EC,i}$ is an entry cost, where the entry cost parameter $\theta_{EC,i}$ is multiplied by $(1 - a_{i,t-1})$ since this cost is paid only by entrants. If a firm does not operate its store, it can put its capital to other uses. Current profits of a non-active firm, $U_{imt}(0)$, are equal to the value of the best outside opportunity. We assume that $U_{imt}(0) = \varepsilon_{imt}(0)$, which is private information of firm i. The choice-specific private information variables, $\varepsilon_{imt}(0)$ and $\varepsilon_{imt}(1)$, are transitory normally

[28] One may interpret variable profits as the equilibrium payoffs of a one-period static game in which firms with identical products and variable costs compete in quantities (i.e., Cournot).

distributed shocks, iid across firms, markets and time with zero mean.[29]

The vector of state variables of the game which are observable to the researcher is $x_{mt} = (S_{mt}, a_{mt-1})$ where $a_{mt-1} = \{a_{mi,t-1} : i = 1, 2, ..., N\}$ are indicators of incumbency status. Incumbency status matters because incumbent firms do not pay entry costs and are thus more likely to operate their store than non-incumbents. The unobservable state variables are the private information shocks ε_{imt}. In this game, a firm's strategy function, $\alpha_i(x_{mt}, \varepsilon_{imt})$, is a binary decision rule. The associated conditional choice probability, $P_i^{\alpha}(1|x_{mt})$, is the probability that the firm operates in the market.

Aguirregabiria and Mira estimate this model using panel data of Chilean local retail markets. They assume that all potential entrants in each market are identical and consider symmetric equilibria.[30] They estimate the model separately for five different retail industries and analyze how economies of scale, the sensitivity of profits to the number of active firms and the magnitude of sunk costs contribute to explain differences in the number of active firms across industries. As we describe in section 4, Aguirregabiria and Mira relax Assumption OC by including a time-invariant and market-specific component of market profitability that is common knowledge for firms but unobserved to the researcher They find that this unobserved market heterogeneity is important. Failure to account for it may lead to implausible estimates of the parameter θ_{RN} which measures the sensitivity of profits to the number of active competitors, because more profitable markets tend to have a larger number of active firms in equilibrium. ■

The estimation of dynamic games, relative to single agent models, poses several specific problems. Simple as this example is, it illustrates how the size of the game's state space tends to grow exponentially with the number of players if it includes player-specific state variables such as incumbency status. As in the competitive equilibrium example, solving the model nests two levels of fixed point problems, i.e., the best response equilibrium condition and the individual player's dynamic programming problem. In order to alleviate the computational complexity, one might attempt to estimate structural parameters from each player's individual decision problem and actions. But the decision rule depends on the (lagged) actions of other players, such as their incumbency status. If some aggregate state variables are unobservable (e.g., the permanent component of market profitability), other player's actions are endogenous.[31] For given parameter values, multiple equi-

[29] We might write the value of the best outside oppotunity as $\mu_i + \varepsilon_{imt}(0)$, but the parameter μ_i cannot be identified separately from the average fixed cost $\theta_{FC,i}$ so we normalize it to zero.

[30] Therefore, $\theta_{FC,i}$ and $\theta_{EC,i}$ is the same for all firms. Aguirregabiria, Mira and Roman (2007) relax this assumption. We discuss the issue of firm permanent unobserved heterogeneity in section 4.

[31] And, furthermore, an initial conditions problem has to be taken into account.

libria are a (likely) possibility which will make the empirical model indeterminate and call for some additional structure.

On a different level, note that example 4 is different from the first three in that only choice data are used in estimation. In the first three examples, data on payoffs (i.e., wages) was also available. 'Pure choice' data is a fairly common situation which tends to make estimation simpler because payoff data are often choice-censored and corrections for selection can be more difficult to implement in structural models, as we will see for example 2. On the other hand, it is clear that availability of payoff data is an important source of identification. Furthermore, the basic framework that we have described for empirical discrete games is easily extended to the case in which payoff variables satisfy assumption CI-Y.

3 Estimation methods for single agent models

3.1 Methods for *Rust's model*

We introduced *Rust's model* as a single agent model with the conditional independence and additive separability assumptions. We describe four methods/algorithms which have been applied to the estimation of this class of models: (1) the nested fixed point algorithm; (2) Hotz-Miller's CCP method; (3) recursive CCP method; and (4) Hotz-Miller with simulation. The first of these is a full solution method, i.e., the DP-problem is solved for every trial value of the parameters. Methods (2)-(4) avoid repeated full solutions of the DP problem, taking advantage of the existence of an invertible mapping between conditional choice probabilities and differences in choice-specific value functions, a result due to Hotz and Miller (1993). We have noted that almost all applications in models with more than two choice alternatives have also imposed the extreme value assumption on unobservables. Although the extreme value assumption per se is not essential for the statistical properties of the model and the methods which we review here, it leads to closed form expressions for several of the econometric model's key objects. This is an advantage which these methods fully exploit, specially (2)-(4). Extending the range of applicability of these methods to models which do not impose the extreme value assumption is a topic for further research.

3.1.1 Nested fixed point algorithm

The nested fixed point algorithm (NFXP) is a gradient iterative search method to obtain the maximum likelihood estimator of the structural parameters. More specifically, this algorithm combines a BHHH method (outer algorithm), that searches for a root of the likelihood equations, with a

value function or policy iteration method (inner algorithm), that solves the dynamic programming problem for each trial value of the structural parameters. The algorithm is initialized with an arbitrary vector of structural parameters, say $\hat{\theta}_0$. A BHHH iteration is defined as:

$$\hat{\theta}_{k+1} = \hat{\theta}_k + \left(\sum_{i=1}^{N} \nabla l_i(\hat{\theta}_k) \nabla l_i(\hat{\theta}_k)' \right) \left(\sum_{i=1}^{N} \nabla l_i(\hat{\theta}_k) \right) \tag{27}$$

where $\nabla l_i(\theta)$ is the gradient in θ of the log-likelihood function for individual i. Given the form of the likelihood in equation (10), and that $\theta = (\theta_u', \theta_Y', \theta_f')'$, we have:

$$\nabla l_i(\theta) = \left[\begin{array}{l} \sum_{t=1}^{T_i} \nabla_{\theta_u} \log P(a_{it}|x_{it}, \theta) \\ \sum_{t=1}^{T_i} \nabla_{\theta_Y} \log P(a_{it}|x_{it}, \theta) + \sum_{t=1}^{T_i} \nabla_{\theta_Y} \log f_Y(y_{it}|a_{it}, x_{it}, \theta_Y) \\ \sum_{t=1}^{T_i} \nabla_{\theta_f} \log P(a_{it}|x_{it}, \theta) + \sum_{t=1}^{T_i - 1} \nabla_{\theta_f} \log f_x(x_{i,t+1}|a_{it}, x_{it}, \theta_f) \end{array} \right] \tag{28}$$

The terms $\nabla_{\theta_Y} \log f_Y$ and $\nabla_{\theta_f} \log f_x$ are standard because the transition probability function and the payoff function are primitives of the model. However, to obtain $\nabla_\theta \log P$ we need to solve the DP problem for $\theta = \hat{\theta}_k$ in order to compute the conditional choice probabilities and their derivatives with respect to the components of $\hat{\theta}_k$. There are different ways to solve the DP problem. When the model has finite horizon (i.e., T is finite) the standard approach is to use backward induction. For infinite horizon models, one can use either value function iterations (described below) or policy function iterations, or an hybrid of both.

To illustrate this algorithm in more detail, consider a version of Rust's DP conditional logit model with infinite horizon, and discrete observable state variables x. Let $\bar{\mathbf{V}}(\theta)$ be the column vector of values $\{\bar{V}(x, \theta) : x \in X\}$. Following equation (12), this vector of values can be obtained as the unique fixed point in $\bar{\mathbf{V}}$ of the following Bellman equation in vector form:

$$\bar{\mathbf{V}} = \log \left(\sum_{a=0}^{J} \exp \left\{ \mathbf{u}(a, \theta) + \beta \ \mathbf{F}_x(a) \ \bar{\mathbf{V}} \right\} \right) \tag{29}$$

where $\mathbf{u}(a, \theta)$ is the vector of utilities $\{u(a, x, \theta) : x \in X\}$ and $\mathbf{F}_x(a)$ is the transition probability matrix with elements $f_x(x'|a, x)$. The choice probabilities $P(a|x, \theta)$ have the conditional logit form:

$$P(a|x, \theta) = \frac{\exp \left\{ u(a, x_{it}, \theta) + \beta \ \mathbf{F}_x(a, x)' \bar{\mathbf{V}}(\theta) \right\}}{\sum_{j=0}^{J} \exp \left\{ u(j, x_{it}, \theta) + \beta \ \mathbf{F}_x(j, x)' \bar{\mathbf{V}}(\theta) \right\}} \tag{30}$$

where $\mathbf{F}_x(a, x)$ is the column vector $\{f_x(x'|a, x) : x' \in X\}$. To obtain the gradient $\nabla_\theta \log P_{it}$ of this DP-conditional logit model, it is useful to take into account that the denominator in equation (30) is equal to the exponential of $\bar{V}(x_{it}, \theta)$. It can be shown that this gradient has the following analytic form:

$$\begin{aligned} \nabla_{\theta_u} \log P_{it} &= \frac{\partial u(a_{it}, x_{it})}{\partial \theta_u} + \beta \frac{\partial \bar{\mathbf{V}}'}{\partial \theta_u} \mathbf{F}_x(a_{it}, x_{it}) - \frac{\partial \bar{V}(x_{it})}{\partial \theta_u} \\ \nabla_{\theta_f} \log P_{it} &= \beta \left(\frac{\partial \mathbf{F}_x(a_{it}, x_{it})'}{\partial \theta_f} \bar{\mathbf{V}} + \frac{\partial \bar{\mathbf{V}}'}{\partial \theta_f} \mathbf{F}_x(a_{it}, x_{it}) \right) - \frac{\partial \bar{V}(x_{it})}{\partial \theta_f} \end{aligned} \tag{31}$$

The expression for $\nabla_{\theta_Y} \log P_{it}$ is equivalent to that of $\nabla_{\theta_u} \log P_{it}$. And given the Bellman equation in (29), the Jacobian matrix $\partial \bar{\mathbf{V}}(\theta)/\partial \theta'$ is:

$$\begin{aligned} \frac{\partial \bar{\mathbf{V}}(\theta)}{\partial \theta'_u} &= \left(I - \beta \sum_{a=0}^{J} \mathbf{P}(a|\theta) * \mathbf{F}_x(a) \right)^{-1} \left(\sum_{a=0}^{J} \mathbf{P}(a|\theta) * \frac{\partial \mathbf{u}(a,\theta)}{\partial \theta'_u} \right) \\ \frac{\partial \bar{\mathbf{V}}(\theta)}{\partial \theta'_f} &= \beta \left(I - \beta \sum_{a=0}^{J} \mathbf{P}(a|\theta) * \mathbf{F}_x(a) \right)^{-1} \left(\sum_{a=0}^{J} \mathbf{P}(a|\theta) * \frac{\partial \mathbf{F}_x(a)}{\partial \theta'_f} \bar{\mathbf{V}}(\theta) \right) \end{aligned} \tag{32}$$

where $\mathbf{P}(a|\theta)$ is the column vector of choice probabilities $\{P(a|x,\theta) : x \in X\}$, and $*$ represents the element-by-element product. Note that we obtain the gradient of the value functions in a relatively simple manner as a by-product of the iterative DP solution method. This is an important additional advantage of the DP-conditional logit specification. Without it, it would be necessary to 'perturb' each element of θ and obtain new solutions of the DP model for each perturbation in order to compute numerical derivatives, which is much more costly.[32]

The NFXP algorithm proceeds as follows. We start with an arbitrary value of θ, say $\hat{\theta}_0$. Given $\hat{\theta}_0$, in the 'inner' algorithm we obtain the vector $\bar{\mathbf{V}}(\hat{\theta}_0)$ by successive iterations in the Bellman equation (29): starting with some guess $\bar{\mathbf{V}}_0$, we iterate in $\bar{\mathbf{V}}_{h+1} = \log(\sum_{a=0}^{J} \exp\{\mathbf{u}(a,\hat{\theta}_0) + \beta\, \mathbf{F}_x(a) \bar{\mathbf{V}}_h\})$ until convergence. Then, given $\hat{\theta}_0$ and $\bar{\mathbf{V}}(\hat{\theta}_0)$ we construct the choice probabilities $P(a|x,\hat{\theta}_0)$ using the formula in (30), the matrix $\partial \bar{\mathbf{V}}(\hat{\theta}_0)/\partial \theta'$using (32), and the gradient $\nabla l_i(\hat{\theta}_0)$ using equation (31). Finally, in the 'outer' algorithm we use the gradient $\nabla l_i(\hat{\theta}_0)$ to make a new BHHH iteration to obtain $\hat{\theta}_1$. We proceed in this way until the distance between $\hat{\theta}_{k+1}$ and $\hat{\theta}_k$ or the difference in the likelihoods is smaller than a pre-specified convergence constant. When the model has finite horizon, we can solve for the value function, its gradient and choice probabilities using backward induction in the inner algorithm of the NFXP. That is, the sequence of value vectors at ages T, $T-1$, etc, can be obtained starting with $\bar{\mathbf{V}}_T(\hat{\theta}) = \log(\sum_{a=0}^{J} \exp\{\mathbf{u}_T(a,\hat{\theta})\})$, and then using the recursive formula $\bar{\mathbf{V}}_t(\hat{\theta}) = \log(\sum_{a=0}^{J} \exp\{\mathbf{u}_t(a,\hat{\theta}) + \beta \mathbf{F}_{x,t}(a) \bar{\mathbf{V}}_{t+1}(\hat{\theta})\})$ for $t \leq T-1$. At each iteration the choice probabilities are $P_t(a|\hat{\theta}) = \exp\{u_t(a,\hat{\theta}) + \beta\, \mathbf{F}_{x,t}(a)' \bar{\mathbf{V}}_{t+1}(\hat{\theta})\} \;/\; [\sum_{j=0}^{J} \exp\{u_t(j,\hat{\theta}) + \beta\, \mathbf{F}_{x,t}(j)' \bar{\mathbf{V}}_{t+1}(\hat{\theta})\}]$, and the gradients have the following recursive forms $\partial \bar{\mathbf{V}}_t(\hat{\theta})/\partial \theta'_u = \sum_{a=0}^{J} \mathbf{P}_t(a|\hat{\theta}) * \{\partial \mathbf{u}_t(a,\hat{\theta})/\partial \theta'_u + \beta \mathbf{F}_{x,t}(a)\; \partial \bar{\mathbf{V}}_{t+1}(\hat{\theta})/\partial \theta'_u\}$ and $\partial \bar{\mathbf{V}}_t(\hat{\theta})/\partial \theta'_f = \beta \sum_{a=0}^{J} \mathbf{P}_t(a|\hat{\theta}) * \{\partial \mathbf{F}_{x,t}(a)/\partial \theta'_f \bar{\mathbf{V}}_{t+1}(\hat{\theta}) + \beta \widehat{\mathbf{F}}_{x,t}(a)\; \partial \bar{\mathbf{V}}_{t+1}(\hat{\theta})/\partial \theta'_f\}$.

As any other gradient method, the NFXP algorithm returns a solution to the likelihood equations. In general, the likelihood function of this class of models is not globally concave. Therefore,

[32] If policy iteration is used to solve the infinite horizon DP model, the gradient of the value function can also be obtained as a by-product. The computational cost of computing numerical derivatives is the main reason why BHHH, which avoids second derivatives, is paricularly useful in the estimation of structural models.

some global search is necessary to check whether the root of the likelihood equations that has been found is actually the global maximum and not just a local optimum.

We have described the NFXP algorithm in the context of full information MLE (FIML). However, most applications of this algorithm have considered a sequential partial likelihood approach first advocated by Rust. In the first step, the partial likelihoods $\sum_{i,t} \log f_Y(y_{it}|a_{it}, x_{it}, \theta_Y)$ and $\sum_{i,t} \log f_x(x_{i,t+1}|a_{it}, x_{it}, \theta_f)$ are maximized to obtain estimates of parameters θ_Y and θ_f, respectively. This step does not involve solving the DP problem. Given these estimators, in the second step the parameters in θ_u are estimated using the NFXP algorithm and the partial likelihood $\sum_{i,t} \log P(a_{it}|x_{it}, \theta_u, \hat{\theta}_Y, \hat{\theta}_f)$. This two-step approach can greatly simplify the estimation problem in models with many parameters in the transition probabilities. It was used by Rust and Phelan (1997) to estimate the model that we described in Example 1. In that application, the state variables with stochastic transitions were health status, health expenses, marital status and public pension points. The payoff function is the labor earnings equation. There are two main reasons why this sequential approach reduces the cost of estimating these and other models. First, the number of BHHH iterations needed to reach convergence typically increases with the number of parameters that are estimated. Note that BHHH iterations here are particularly costly because they involve the full solution of the dynamic programming problem. And second, the computational cost associated with a global search increase as well with the dimension of the parameter space. Nevertheless, it should be noted that in the sequential partial likelihood approach the calculation of standard errors in the second step correcting for estimation error in the first step is not a trivial task. It is often no more costly to consider a third estimation step consisting of a single FIML-BHHH iteration, which as is well known delivers an asymptotically efficient estimator as well as the correct standard errors.

There is a long list of applications which have used the NFXP algorithm to estimate models in Rust's class. For instance: investment models of machine replacement, as in Rust (1987), Sturm (1991), Das (1992), Kennet (1993 and 1994), and Rust and Rothwell (1995); or models of retirement from the labor force, as in Rust and Phelan (1997), and Karlstrom, Palme and Svensson (2004).

3.1.2 Hotz-Miller's CCP estimator

The main advantages of the NFXP algorithm are its conceptual simplicity and, more importantly, that it gives the MLE which is asymptotically efficient under the assumptions of the model. The main limitation of this algorithm is its computational cost. In particular, the DP problem has to be solved exactly for each trial value of the structural parameters. Given the cost of solving some DP

problems, this characteristic of the algorithm limits the range of applications that are feasible, even in the DP-conditional logit case. Hotz and Miller (1993) observed that, under the assumptions of Rust model, it is not necessary to solve the DP problem even once in order to estimate the structural parameters A key idea in their method is that, using nonparametric estimates of choice and transition probabilities, it is possible to obtain a simple representation of the choice-specific value functions $v(a, x, \theta)$ for values of θ around the true vector of structural parameters (see also Manski, 1993, for a method that exploits a similar idea). This representation is particularly simple and useful for estimation in the DP-conditional logit model with linear-in-parameters utility. We assume that $u(a, x_{it}, \theta_u) = z(a, x_{it})'\theta_u$, where $z(a, x_{it})$ is a vector of known functions. We describe Hotz-Miller's method in detail for this case but we also include a brief discussion of alternative specifications, including the nonlinear DP-conditional logit case which corresponds to the model in Example 1.[33] Furthermore, we assume that the dataset does not include payoff variables or, if it does, assumption CI-Y holds and θ_Y has been estimated and subsumed into the known z functions.

In Hotz and Miller's representation the choice-specific value functions are written as follows:[34]

$$v(a, x, \theta) = \tilde{z}(a, x, \theta)\ \theta_u + \tilde{e}(a, x, \theta) \tag{33}$$

where $\tilde{z}(a, x, \theta)$ is the expected and discounted sum of current and future z vectors $\{z(a_{t+j}, x_{t+j}) : j = 0, 1, 2, ...\}$ which may occur along all possible histories originating from $(a_t, x_t) = (a, x)$; and likewise $\tilde{e}(a, x, \theta)$ is the expected and discounted sum of the stream $\{\varepsilon(a_{t+j}) : j = 1, 2, ...\}$. More formally,

$$\begin{aligned} \tilde{z}(a, x_t, \theta) &= z(a, x_t) + \sum_{j=1}^{T-t} \beta^j\ E_{x_{t+j}|a_t=a,x_t} \left[\sum_{a'=0}^{J} P(a'|x_{t+j}, \theta)\ z(a', x_{t+j}) \right] \\ \tilde{e}(a, x_t, \theta) &= \sum_{j=1}^{T-t} \beta^j\ E_{x_{t+j}|a_t=a,x_t} \left[\sum_{a'=0}^{J} P(a'|x_{t+j}, \theta)\ e(a', x_{t+j}) \right] \end{aligned} \tag{34}$$

where $E_{x_{t+j}|a_t=a,x_t}(.)$ represents the expectation over the distribution of x_{t+j} conditional to $(a_t = a, x_t)$. This expectation is calculated under the assumption that the individual behaves optimally in the future and, therefore, state variables evolve as a controlled stochastic process defined by optimal choice probabilities and transition probability functions. $e(a, x_t)$ is the expectation of $\varepsilon_t(a)$

[33] Alternatively, a linear utility $z(a, x_{it})'\boldsymbol{\theta}_u$ can also be interpreted as a first order Taylor approximation to a nonlinear utility.

[34] Hotz-Miller's representation can be derived either for alternative-specific value functions or for the integrated value function (Emax function). We have preferred to use the first version here. Also note that for problems with terminal actions Hotz and Miller derived a different representation of alternative-specific value functions which is simpler than the general representation that we present here. The estimator based on this alternative representation has different asymptotic variance and it has been used less often.

conditional on x_t and on alternative a being optimal, i.e., $e(a,x_t) \equiv E(\varepsilon_t(a)|x_t, \alpha(x_t,\varepsilon_t)=a)$. Hotz and Miller showed that $e(a,x_t)$ is a function of a, the probabilities $P(.|x_t,\theta)$, and the distribution G_ε only.[35] The particular functional form of $e(a,x)$ depends on the probability distribution G_ε. When the $\varepsilon's$ are iid extreme value, $e(a,x)$ has a closed-form and is equal to Euler's constant minus $\log P(a|x,\theta)$. Note that $\tilde{z}(a,x_t,\theta)$ and $\tilde{e}(a,x_t,\theta)$ depend on θ only through the parameters in the transition probabilities, θ_f, and the vector of CCPs $\mathbf{P}(\theta) = \{P(a|x,\theta) : (a,x) \in A \times X\}$. In fact, we can define $\tilde{z}$ and $\tilde{e}$ for an arbitrary vector of CCPs (optimal or not). To emphasize this point we use the notation $\tilde{z}^{\mathbf{P}}(a,x_t)$ to represent $\sum_{j=1}^{T-t} \beta^j E_{x_{t+j}|a_t=a,x_t}[\sum_{a'=0}^{J} P(a'|x_{t+j})\ z(a',x_{t+j})]$ and $\tilde{e}^{\mathbf{P}}(a,x_t)$ to represent $\sum_{j=1}^{T-t} \beta^j E_{x_{t+j}|a_t=a,x_t}[\sum_{a'=0}^{J} P(a'|x_{t+j})\ e(a',x_{t+j})]$.[36]

Let $\theta^0 = (\theta_u^0, \theta_f^0)$ be the true value of θ in the population of individuals under study. And let $\mathbf{P}^0$ be the conditional choice probabilities in the population. If we knew $(\mathbf{P}^0, \theta_f^0)$, we could construct the values $\tilde{z}^{\mathbf{P}^0}(a,x)$ and $\tilde{e}^{\mathbf{P}^0}(a,x)$ and then use Hotz and Miller's representation to approximate the choice-specific values $v(a,x,\theta)$ as simple linear functions of θ_u close to its true value. Although we do not know $(\mathbf{P}^0, \theta_f^0)$, we can estimate them consistently without having to solve the DP problem. Consistent estimates of transition probabilities can be obtained using a (partial) MLE of θ_f^0 that maximizes the (partial) likelihood $\sum_{i,t} \log f_x(x_{i,t+1}|a_{it},x_{it},\theta_f)$. Conditional choice probabilities can be estimated using nonparametric regression methods (i.e., $P^0(a|x) = E(I\{a_{it}=a\}|x_{it}=x)$) such as a Nadaraya-Watson kernel estimator or a simple frequency estimator. Let $\hat{\mathbf{P}}$ and $\hat{\theta}_f$ be the estimators of $\mathbf{P}^0$ and θ_f^0, respectively. Based on these estimates, Hotz and Miller's idea is to approximate $v(a,x_{it},\theta)$ by $\tilde{z}^{\hat{\mathbf{P}}}(a,x_{it})\theta_u + \tilde{e}^{\hat{\mathbf{P}}}(a,x_{it})$ in order to obtain the CCP's in equation (30). They propose the GMM estimator that solves in θ_u the sample moment conditions:

$$\sum_{i=1}^{N}\sum_{t=1}^{T_i} H(x_{it}) \begin{bmatrix} I\{a_{it}=1\} - \dfrac{\exp\left\{\tilde{z}^{\hat{\mathbf{P}}}(1,x_{it})\theta_u + \tilde{e}^{\hat{\mathbf{P}}}(1,x_{it})\right\}}{\sum_{a=0}^{J}\exp\left\{\tilde{z}^{\hat{\mathbf{P}}}(a,x_{it})\theta_u + \tilde{e}^{\hat{\mathbf{P}}}(a,x_{it})\right\}} \\ \vdots \\ I\{a_{it}=J\} - \dfrac{\exp\left\{\tilde{z}^{\hat{\mathbf{P}}}(J,x_{it})\theta_u + \tilde{e}^{\hat{\mathbf{P}}}(J,x_{it})\right\}}{\sum_{a=0}^{J}\exp\left\{\tilde{z}^{\hat{\mathbf{P}}}(a,x_{it})\theta_u + \tilde{e}^{\hat{\mathbf{P}}}(a,x_{it})\right\}} \end{bmatrix} = 0 \qquad (35)$$

[35] To see this, note that the event $\{\alpha(x,\varepsilon)=a\}$ is equivalent to $\{v(a,x)+\varepsilon(a) \geq v(a',x)+\varepsilon(a')$ for any $a' \neq a\}$. Then, $e(a,x) = \frac{1}{P(a|x)} \int \varepsilon(a) I\{\varepsilon(a') - \varepsilon(a) \leq v(a,x) - v(a',x)\ \forall a' \neq a\}\, dG(\varepsilon)$. The last expression shows that $e(a,x)$ depends on the primitives of the model only through the probability distribution of ε and the vector of value differences $\tilde{v}(x) \equiv \{v(a,x) - v(0,x) : a \in A\}$. The vector of choice probabilities $\{P(a|x) : a \in A\}$ is also a function of G_ε and $\tilde{v}(x)$, i.e., $P(a|x) = \Pr(\varepsilon(a') - \varepsilon(a) \leq v(a,x) - v(a',x)\ \forall a' \neq a|x)$. Hotz and Miller showed that this mapping that relates choice probabilities and value differences is invertible (see Proposition 1, page 501, in Hotz and Miller, 1993).

[36] For notational simplicity we omit θ_f as an argument, though it should be clear that $\tilde{z}^{\mathbf{P}}(a,x_t)$ and $\tilde{e}^{\mathbf{P}}(a,x_t)$ depend on $\mathbf{P}$ and θ_f.

where $H(x_{it})$ is a matrix with dimension $\dim(\theta_u) \times J$ with functions of x_{it} which are used as instruments.

The main advantage of this estimator is its computational simplicity. Nonparametric estimation of choice probabilities is a (relatively) simple task. The main task is the computation of the values $\tilde{z}^{\hat{\mathbf{P}}}(a, x_{it})$ and $\tilde{e}^{\hat{\mathbf{P}}}(a, x_{it})$. We provide more details on the calculation of these expected and discounted values below. However, these values are calculated just once and remain fixed in the search for the Hotz-Miller estimator. In contrast, in a full solution method such as the NFXP algorithm these values are recomputed exactly for each trial value of θ. Thus Hotz-Miller's method greatly reduces the computational burden of NFXP's 'inner' algorithm. Another important advantage of Hotz-Miller's CCP estimator is that, for the DP conditional logit model with linear-in-parameters utility, the system of equations (35) that defines the estimator has a unique solution. Therefore, a global search is not needed.

Previous conventional wisdom was that Hotz-Miller's estimator achieved a significant computational gain at the expense of efficiency, both in finite samples and asymptotically. Thus, researchers had the choice between two extremes: a full solution NFXP-ML estimator with the attendant computational burden, or the much faster but less efficient Hotz-Miller estimator. However, Aguirregabiria and Mira (2002) showed that a pseudo maximum likelihood version of Hotz-Miller's estimator is asymptotically equivalent to partial MLE. The 'two-step' pseudo maximum likelihood (PML) estimator is defined as the value of θ_u that maximizes the pseudo likelihood function:[37]

$$Q(\theta_u, \hat{\mathbf{P}}, \hat{\theta}_f) = \sum_{i=1}^{N} \sum_{t=1}^{T_i} \log \frac{\exp\left\{\tilde{z}^{\hat{\mathbf{P}}}(a_{it}, x_{it})\theta_u + \tilde{e}^{\hat{\mathbf{P}}}(a_{it}, x_{it})\right\}}{\sum_{a=0}^{J} \exp\left\{\tilde{z}^{\hat{\mathbf{P}}}(a, x_{it})\theta_u + \tilde{e}^{\hat{\mathbf{P}}}(a, x_{it})\right\}} \tag{36}$$

The asymptotic variance of this two-step PML estimator is just equal to the variance of the partial MLE of θ_u described at the end of section 3.1.1. That is, the initial nonparametric estimator of $\mathbf{P}^0$ and the PML estimator of θ_u^0 are asymptotically independent and therefore there is no asymptotic efficiency loss from using an inefficient initial estimator of $\mathbf{P}^0$. Nevertheless, although the two-step PML estimator is asymptotically equivalent to partial MLE, Monte Carlo experiments show that its finite sample bias can be much larger.[38] Imprecise initial estimates of choice probabilities do not affect the asymptotic properties of the estimator, but they can generate serious small sample biases in the two-step PML estimator and, more generally, in the whole class of Hotz-Miller's CCP

[37] It is well known that the PML estimator belongs to the class of GMM estimators defined in equation (35). More specifically, the PML estimator is the GMM estimator with a matrix of instruments $H(x_{it})$ equal to the pseudo scores.

[38] See the Monte Carlo experiments in Aguirregabiria and Mira (2002 and 2007), Pesendorfer and Schmidt-Dengler (2007), and Kasahara and Shimotsu (2006).

estimators.[39] The source of this bias is well understood in two-step methods: $\hat{\mathbf{P}}$ enters nonlinearly in the sample moment conditions that define the estimator, and the expected value of a nonlinear function of $\hat{\mathbf{P}}$ is not equal to that function evaluated at the expected value of $\hat{\mathbf{P}}$. The larger the variance of $\hat{\mathbf{P}}$, the larger the bias of $\hat{\theta}_u$. The variance of the nonparametric estimator of $P^0(a|x) = E(I\{a_{it} = a\}|x_{it} = x)$ increases with the number of cells in the set X. In applications with millions of cells in X and a few thousand observations, the variance of $\hat{\mathbf{P}}$ and the bias of $\hat{\theta}_u$ can be very large. A recursive extension of the two-step method, which we describe in the next subsection, deals with this problem.

We now describe in some detail on the computation of the values $\tilde{z}^{\mathbf{P}}(a, x_t)$ and $\tilde{e}^{\mathbf{P}}(a, x_t)$. We start with the case of a finite horizon model, i.e., $T < \infty$. We can obtain the sequence of values $\{\tilde{z}_t^{\mathbf{P}}(a, x_t)$ and $\tilde{e}_t^{\mathbf{P}}(a, x_t) : t = 1, 2, ..., T\}$ using backwards induction. Starting at the last period, we have $\tilde{z}_T^{\hat{\mathbf{P}}}(a, x_T) = z(a, x_T)$ and $\tilde{e}_T^{\hat{\mathbf{P}}}(a, x_T) = 0$. Then, for any period $t < T$, we have the following recursive expressions:

$$\begin{aligned}
\tilde{z}_t^{\mathbf{P}}(a, x_t) &= z(a, x_t) + \beta \sum_{x_{t+1} \in X} f_x(x_{t+1}|a, x_t) \left[\sum_{a'=0}^{J} P_{t+1}(a'|x_{t+1}) \ \tilde{z}_{t+1}^{\mathbf{P}}(a', x_{t+1}) \right] \\
\tilde{e}_t^{\mathbf{P}}(a, x_t) &= \beta \sum_{x_{t+1} \in X} f_x(x_{t+1}|a, x_t) \left[\sum_{a'=0}^{J} P_{t+1}(a'|x_{t+1}) \ \left(e_{t+1}(a', x_{t+1}) + \tilde{e}_{t+1}^{\mathbf{P}}(a', x_{t+1}) \right) \right]
\end{aligned} \tag{37}$$

The computational burden incurred to obtain these values is equivalent to that of solving the finite horizon DP problem. For models with infinite horizon, these values can be obtained by using successive approximations. It is possible to show that $\tilde{z}^{\mathbf{P}}(a, x) = z(a, x) + \beta \sum_{x' \in X} f_x(x'|a, x)$ $W_z^{\mathbf{P}}(x')$ and $\tilde{e}^{\mathbf{P}}(a, x) = \beta \sum_{x' \in X} f_x(x'|a, x)$ $W_e^{\mathbf{P}}(x')$, where $W_z^{\mathbf{P}}(.)$ is a $1 \times \dim(\theta_u)$ vector and $W_e^{\mathbf{P}}(.)$ is a scalar and both are (basis functions for) valuation operators. Define the matrix $\mathbf{W}^{\mathbf{P}} \equiv \{[W_z^{\mathbf{P}}(x), W_e^{\mathbf{P}}(x)] : x \in X\}$. Then, the valuation basis $\mathbf{W}^{\mathbf{P}}$ is defined as the unique solution in $\mathbf{W}$ to the following contraction mapping:

$$\mathbf{W} = \sum_{a=0}^{J} \mathbf{P}(a) * \{[\mathbf{z}(a), \mathbf{e}(a)] + \beta \ \mathbf{F}_x(a) \ \mathbf{W}\} \tag{38}$$

where $\mathbf{P}(a)$ is the column vector of choice probabilities $\{P(a|x) : x \in X\}$; $\mathbf{z}(a) = \{z(a, x) : x \in X\}$; and $\mathbf{e}(a) \equiv \{e(a, x) : x \in X\}$.[40] Again, the computational cost to obtain these values is equivalent to solving once the infinite horizon DP problem. Note that (38) is a linear system, so there is a

[39] As pointed out by Pakes, Ostrovsky and Berry (2007), the bias of CCP estimators can be smaller when the instruments $H(x_{it})$ do not depend on the first step estimator $\hat{\mathbf{P}}$. We discuss this issue in section 4.

[40] $W_z^{\mathbf{P}}(x)\theta_u + W_e^{\mathbf{P}}(x)$ is the expected discounted utility of behaving according to choice probabilities $\mathbf{P}$ from current period t and into the infinite future when $x_t = x$. See Aguirregabiria and Mira (2002) for more detail.

closed form expression for $\mathbf{W}^{\mathbf{P}}$, i.e., $\mathbf{W}^{\mathbf{P}} = (I - \beta \sum_{a=0}^{J} \mathbf{P}(a) * \mathbf{F}_x(a))^{-1} \sum_{a=0}^{J} \mathbf{P}(a) * [\mathbf{z}(a), \mathbf{e}(a)]$. When the number of cells in X is small enough, matrix inversion algorithms may be preferable to successive approximations.[41]

Our description of the CCP method has sofar assumed that the utility function is linear in parameters and that there is a closed form expression for the $e()$ function. We now discuss briefly the role of these assumptions. First, if the utility is not linear-in-parameters, we can represent the choice-specific value $v(a, x, \theta)$ as $\tilde{z}^{\mathbf{P}}(a, x, \theta_u) + \tilde{e}^{\mathbf{P}}(a, x)$, where θ_u is now an argument in $\tilde{z}^{\mathbf{P}}()$. All the previous expression for $\tilde{z}^{\mathbf{P}}$ and $\tilde{e}^{\mathbf{P}}$ still apply if we replace $z(a, x)$ by $z(a, x, \theta_u)$ and $W_z^{\mathbf{P}}(x)$ by $W_z^{\mathbf{P}}(x, \theta_u)$. The matrix $\mathbf{W}^{\mathbf{P}}$ is now $\{[W_z^{\mathbf{P}}(x, \theta_u), W_e^{\mathbf{P}}(x)] : x \in X\}$ which uniquely solves the system $\mathbf{W} = \sum_{a=0}^{J} \mathbf{P}(a) * \{[\mathbf{z}(a, \theta_u), \mathbf{e}(a)] + \beta \ \mathbf{F}_x(a) \ \mathbf{W}\}$. For each trial value of θ_u the terms $\sum_{a=0}^{J} \mathbf{P}(a) * \mathbf{z}(a, \boldsymbol{\theta}_u)$ do have to be recomputed and premultiplied by rows of the 'weighting' inverse matrix. This increases significantly the computational cost relative to a model with a linear-in-parameters utility. However, the inverse matrix $(I - \beta \sum_{a=0}^{J} \mathbf{P}(a) * \mathbf{F}_x(a))^{-1}$ only has to be computed once and collects a large part of the calculations involved in valuation.[42] Second, the $e(a, x)$ function has a straightforward expression when the $\varepsilon' s$ have independent extreme value distributions, as well as in binary choice models when unobservables have normal or exponential distributions (see Aguirregabiria and Mira (2007) and Pakes et al (2007) for examples). However, in multinomial models without extreme value unobservables, $e()$ does not have a closed form and would have to be computed numerically or by simulation. Furthermore, in general $e()$ might depend nonlinearly on unknown parameters and would have to be recomputed for different trial values of the parameters. Therefore, relaxing the logistic assumption in the multinomial case represents an important complication for all methods which rely on Hotz-Miller's invertibility result and their usefulness in that setting remains an open question. Finally, it should also be noted that all the methods that use Hotz-Miller's representation of value functions are based on the two-step partial likelihood approach and do not estimate the discount factor directly. To see why, recall that the valuation operator relies on previously obtained consistent estimates of the parameters θ_f of transition probability functions. Without these, the present values $\tilde{z}^{\mathbf{P}}$ and $\tilde{e}^{\mathbf{P}}$ would have to be recomputed repeatedly for different values of θ_f and β in the in the second step, rather than just once, and most of the computational advantage of the Hotz-Miller approach would be lost.

[41] The matrix $(I - \beta F)^{-1}$ can be also approximated using the series $I + \beta F + \beta^2 F^2 + ... + \beta^K F^K$, with K large enough. This can be easier than matrix inversion. More generally, this inverse matrix can be obtained iterating in A (succesive approximations) in the mapping $A = I + \beta F \ A$.

[42] This inverse matrix computes the expected number of times each state will be visited in the future, with each visit weighted by the corresponding discount factor.

An important limitation of Hotz-Miller's CCP estimator and most of its extensions is that consistency depends crucially on Assumption IID. If unobservables are serially correlated, or if there is permanent unobserved heterogeneity as in the finite mixture model, consistent non parametric estimates of CCP's, an essential element of Hotz-Miller's 2-step approach, cannot be obtained from choice data.[43] However, Aguirregabiria and Mira (2007), Aguirregabiria, Mira and Roman (2007) and Arcidiacono and Miller (2006) have recently proposed and applied recursive versions of the CCP estimator which provide consistent estimates for finite mixture models. Also, Kasahara and Shimotsu (2006b) have shown than under certain conditions it is possible to obtain consistent nonparametric estimators of CCPs in finite mixture models, which can be used to construct a root-N consistent CCP estimator.

Some applications which have used the CCP estimator are: contraceptive choice, Hotz and Miller (1993); price adjustment and inventories in retail firms, Slade (1998), Aguirregabiria (1999), Kano (2006); and firms' investment and labor demand, Sanchez-Mangas (2002) and Rota (2004).

3.1.3 Recursive CCP estimation (NPL)

Let $\hat{\theta}_u$ be the two-step PML estimator of θ_u^0 that we have described above. Given this estimator, one can obtain new estimates of the choice probabilities, $\hat{\mathbf{P}}_1 = \{\hat{\mathbf{P}}_1(a|x)\}$, as $\hat{\mathbf{P}}_1(a|x) = \exp\left\{\tilde{z}^{\hat{\mathbf{P}}}(a,x)\ \hat{\theta}_u + \tilde{e}^{\hat{\mathbf{P}}}(a,x)\right\} / \sum_{j=0}^{J} \exp\left\{\tilde{z}^{\hat{\mathbf{P}}}(j,x)\ \hat{\theta}_u + \tilde{e}^{\hat{\mathbf{P}}}(j,x)\right\}$.[44] Given the new estimates $\hat{\mathbf{P}}_1$ we can compute new values $\tilde{z}^{\hat{\mathbf{P}}_1}(a,x_{it})$ and $\tilde{e}^{\hat{\mathbf{P}}_1}(a,x_{it})$, a new pseudo likelihood function $Q(\theta_u, \hat{\mathbf{P}}_1, \hat{\theta}_f)$, and a new PML estimator that maximizes this function. We can iterate in this way to generate a sequence of estimators of structural parameters and conditional choice probabilities $\{\hat{\theta}_{u,K}, \hat{\mathbf{P}}_K : K = 1, 2, ...\}$ such that for any $K \geq 1$:

$$\hat{\theta}_{u,K} = \arg\max_{\theta_u \in \mathbf{\Theta}}\ Q(\theta_u, \hat{\mathbf{P}}_{K-1}, \hat{\theta}_f) \tag{39}$$

and

$$\hat{\mathbf{P}}_K(a|x) = \frac{\exp\left\{\tilde{z}^{\hat{\mathbf{P}}_{K-1}}(a,x)\ \hat{\theta}_{u,K} + \tilde{e}^{\hat{\mathbf{P}}_{K-1}}(a,x)\right\}}{\sum_{j=0}^{J} \exp\left\{\tilde{z}^{\hat{\mathbf{P}}_{K-1}}(j,x)\ \hat{\theta}_{u,K} + \tilde{e}^{\hat{\mathbf{P}}_{K-1}}(j,x)\right\}} \tag{40}$$

All the estimators in this sequence are asymptotically equivalent to partial MLE and to the two-step PML (Aguirregabiria and Mira, 2002, Proposition 4). Therefore, iterating in this procedure does

[43] Relatedly, note that Hotz-Miller's approach also relies on a DGP such that all states are visited with positive probability. This assumption may be problematic in life-cycle applications.

[44] Given parameter values, this expression defines an operator mapping from CCP's into CCP's. Aguirregabiria and Mira (2002) show that this is a policy iteration operator and use it to characterize the solution of the DP problem in the space of conditional choice probabilities. Also note that if unobserved state variables do not have extreme value distribution, the mapping from choice-specific value functions to CCPs which is used in the policy iteration step in (40) need not have a closed form.

not give any asymptotic gain. However, it seems intuitive that if the pseudo likelihood is built from estimates of choice probabilities that exploit the structure of the model one may get estimates of structural parameters with smaller finite sample bias an variance. Aguirregabiria and Mira (2002) present Monte Carlo experiments that illustrate how iterating in this procedure does in fact produce significant reductions in finite sample bias. Kasahara and Shimotsu (2006a) provide a proof of this result using higher order expansions for the bias and variance of the sequence of PML estimators. Aguirregabiria and Mira also show that upon convergence the recursive procedure gives, exactly, a root of the likelihood equations. This result holds regardless of whether the initial estimator of $\mathbf{P}^0$ is consistent or not, and the procedure is called the *nested pseudo likelihood* algorithm (NPL).

Therefore, the NPL procedure can be seen both as a method to reduce the finite sample bias of Hotz-Miller's CCP estimator and as an algorithm to obtain the MLE. As a bias reduction method, we do not have to iterate until convergence and the computational cost is clearly smaller than NFXP. As an algorithm to obtain the MLE, it can also be computationally much cheaper than NFXP. The example in Aguirregabiria and Mira (2002) suggests that this is likely to be the case in infinite horizon models when maximization in θ_u of the pseudo likelihood function is a simple task, such as Rust's DP-conditional logit model with a linear-in-parameters utility where the pseudo likelihood is globally concave in θ_u.[45] Applications of this method include: Aguirregabiria and Alonso-Borrego (1999) on labor demand; Sanchez-Mangas (2002) and Lorincz (2005) on machine replacement and firms' investment; and Kano (2006) on price adjustments with menu costs.

3.1.4 Simulation-based Hotz-Miller estimator

Though Hotz-Miller's CCP estimator is computationally much cheaper than NFXP, it is still impractical for applications where the dimension of the state space X is very large, e.g., a discrete state space with millions of points or a model in which some of the observable state variables are continuous. To deal with this problem, Hotz, Miller, Sanders and Smith (1994) propose an extension of the Hotz-Miller estimator that uses simulation techniques to approximate the values $\tilde{z}^{\hat{\mathbf{P}}}(a,x)$ and $\tilde{e}^{\hat{\mathbf{P}}}(a,x)$. For every value of x_{it} in the sample and every choice alternative $a \in A$ (in the sample or not), we consider (a, x_{it}) as the initial state and generate R simulated paths of future actions and state variables from period $t+1$ to $t+T^*$ (i.e., T^* periods ahead). We index simulated paths by $r \in \{1, 2, ..., R\}$. The $r-th$ simulated path associated with the initial state (a, x_{it}) is

[45] Computational savings will be larger the smaller the number of NPL iterations relative to the number of trial values of θ_u required by NFXP's outer algorithm. Little is known about the relative merits of NPL and NFXP in other contexts, e.g., finite horizon models.

$\{a_{i,t+j}^{(r,a)}, x_{i,t+j}^{(r,a)} : j = 1, 2, ..., T^*\}$. Then, Hotz et al. consider the following simulators:

$$\begin{aligned} \tilde{z}_R^{\hat{\mathbf{P}}}(a, x_{it}) &= z(a, x_{it}) + \frac{1}{R} \sum_{r=1}^{R} \left[\sum_{j=1}^{T^*} \beta^j \; z\left(a_{i,t+j}^{(r,a)}, x_{i,t+j}^{(r,a)} \right) \right] \\ \tilde{e}_R^{\hat{\mathbf{P}}}(a, x_{it}) &= \frac{1}{R} \sum_{r=1}^{R} \left[\sum_{j=1}^{T^*} \beta^j \; e\left(a_{i,t+j}^{(r,a)}, x_{i,t+j}^{(r,a)} \right) \right] \end{aligned} \tag{41}$$

Simulated paths are obtained using the initial estimates of choice and transition probabilities, $\hat{\mathbf{P}}$ and $\hat{\theta}_f$. The path is generated sequentially. Starting at the observed state x_{it} and given the hypothetical action a, the state next period, $x_{i,t+1}^{(r,a)}$, is a random draw from the distribution $f_x(.|a, x_{it}, \hat{\theta}_f)$. Then, the action $a_{i,t+1}^{(r,a)}$ is a random draw from the distribution $\hat{P}(.|x_{i,t+1}^{(r,a)})$. Given $(a_{i,t+1}^{(r,a)}, x_{i,t+1}^{(r,a)})$, then the state $x_{i,t+2}^{(r,a)}$ is a random draw from $f_x(.|a_{i,t+1}^{(r,a)}, x_{i,t+1}^{(r,a)}, \hat{\theta}_f)$, and $a_{i,t+2}^{(r,a)}$ is drawn from $\hat{P}(.|x_{i,t+2}^{(r,a)})$. And so on. Simulations are independent across the R paths. If the DP problem has finite horizon, or if T^* is large enough such that the approximation error associated with the truncation of paths is negligible, then these simulators are unbiased. That is, for any number of simulations R we have that $E_R(\tilde{z}_R^{\hat{\mathbf{P}}}(a, x_{it}))) = \tilde{z}^{\hat{\mathbf{P}}}(a, x_{it})$ and $E_R(\tilde{e}_R^{\hat{\mathbf{P}}}(a, x_{it})) = \tilde{e}^{\hat{\mathbf{P}}}(a, x_{it})$, where $E_R(.)$ is the expectation over the simulation draws.

Hotz et al. propose an estimator that is root-N consistent for any number of simulations, even with $R = 1$. This property of simulation-based estimators obtains when, in the system of equations that define the estimator, the unbiased simulator enters linearly and averaged over sample observations. This is not satisfied by the simulation versions of the GMM estimator in (35) or of the PML estimator. Hotz et al consider a GMM estimator that exploits moment conditions for the choice-specific value functions. Given that the mapping that relates choice-specific value functions and choice probabilities is invertible (see Proposition 1 in Hotz and Miller, 1993), we can represent the value differences $v(a, x, \theta) - v(0, x, \theta)$ as functions of choice probabilities. For the DP-conditional logit model, this inverse function has a very simple closed-form expression: i.e., $v(a, x, \theta) - v(0, x, \theta) = \log(P(a|x, \theta)) - \log(P(0|x, \theta))$. Based on this representation, we can construct the following moment conditions:

$$E\left(h(x_{it}) \left[\log\left(\frac{P^0(a_{it}|x_{it})}{P^0(0|x_{it})} \right) - \left\{ \tilde{z}^{\mathbf{P}^0}(a_{it}, x_{it}) - \tilde{z}^{\mathbf{P}^0}(0, x_{it}) \right\} \theta_u^0 - \left\{ \tilde{e}^{\mathbf{P}^0}(a_{it}, x_{it}) - \tilde{e}^{\mathbf{P}^0}(0, x_{it}) \right\} \right] \right) = 0 \tag{42}$$

where $h(x_{it})$ is a vector of instruments with the same dimension as θ_u. These moment conditions still hold if we replace the population parameters $(\mathbf{P}^0, \theta_f^0)$ by consistent estimates, and the values $\tilde{z}$ and $\tilde{e}$ by the unbiased simulators $\tilde{z}_R$ and $\tilde{e}_R$. Then, for the DP-conditional logit model, the

simulation-based estimator is defined as the value of θ_u that solves the sample moment conditions:

$$\sum_{i=1}^{N}\sum_{t=1}^{T_i} h(x_{it}) \left[\log\left(\frac{\hat{P}_0(a_{it}|x_{it})}{\hat{P}_0(0|x_{it})} \right) - \left\{ \tilde{z}_R^{\hat{\mathbf{P}}}(a_{it}, x_{it}) - \tilde{z}_R^{\hat{\mathbf{P}}}(0, x_{it}) \right\} \theta_u - \left\{ \tilde{e}_R^{\hat{\mathbf{P}}}(a_{it}, x_{it}) - \tilde{e}_R^{\hat{\mathbf{P}}}(0, x_{it}) \right\} \right] = 0 \tag{43}$$

This estimator has a closed form expression. In fact, the expression is the one of an IV estimator in a linear regression model. It is clear that the simulation error averages out over the sample and does not have any influence on the consistency or the rate of convergence of the estimator. However, the simulation error affects the variance of the estimator: as R goes to infinity, the asymptotic variance of this estimator converges to the variance of Hotz et al estimator without simulation. The later is larger than the variance of the one-step PML estimator (i.e., the variance of MLE) because the moment conditions in (42) are not the optimal ones.

Hotz et al. present several Mont Carlo experiments that illustrate that this estimator can have large bias in small samples. To the finite sample bias of the Hotz-Miller estimator now we should add the bias due to the simulation error. Despite these problems, this is a very interesting and useful estimator. The estimator can be applied to models with continuous state variables and it can be extended to deal with continuous decision variables as well. Altug and Miller (1998) applied this method to estimate a model of female labor supply where the decision variable, hours of work, is continuous and censored. Other applications include Miller and Sanders (1997) on welfare participation, and Hollifield, Miller and Sandas (2004) on limit orders markets.

3.2 Estimation of Eckstein-Keane-Wolpin models

Under the label 'Eckstein-Keane-Wolpin' we grouped most applications which have not used Rust's DP-conditional logit model. Specifically, we listed the following four departures from that framework: (1) Unobservables which do not satisfy assumption AS; (2) Observable but choice-censored payoff or state variables; (3) Permanent unobserved heterogeneity; (4) Unobservables which are correlated across choice alternatives. The prevalent estimation criterion for Eckstein-Keane-Wolpin models has been FIML because the partial likelihood approach exploited in Rust's framework does not give consistent estimators under (2) and (3). The bulk of this section will address the estimation of models allowing for (3) and (4), i.e., methods for finite mixtures of likelihoods (section 3.2.1) and Keane and Wolpin's simulation and interpolation method (section 3.2.2). In the next two paragraphs we briefly discuss the estimation of models featuring (1) and (2).[46]

[46] Several of the seminal papers in this literature, such as Miller (1984), Wolpin (1984, 1987) and Pakes (1986), precede Rust's contribution and 'deviated' from his framework. All these models considered binary choices and a single unobservable state variable which resulted in a 'threshold' decision rule.

The main consequence of departing from AS is that the econometric model may not be *saturated*. A discrete choice model is *saturated* if for any value of the observable state variables and of the structural parameters the model predicts a strictly positive choice probability for any of the choice alternatives. Non-saturation causes econometric and computational problems in maximum likelihood estimation (see Rust, 1994). A natural way of dealing with this issue is to allow for measurement error in the state and/or choice variables. Wolpin (1987) and Keane and Wolpin (1997, 2001) are examples of this approach to deal with non-saturated models.

The problem with censored payoff or state variables was illustrated in the occupational choice model of Example 2, where wages are observed if individuals work but the distribution of wages across occupations cannot be estimated consistently from wage data alone because observed wages are a choice-censored sample. The full likelihood uses the structural behavioral model to correct for sample selection bias.[47]

3.2.1 Finite mixture models (permanent unobserved heterogeneity)

The finite mixture framework and FIML estimation: Consider a more general version of the finite mixture model introduced in Section 2, with permanent unobserved heterogeneity in the utility function, payoff function, and transition probabilities. Individuals in the population belong to one of L unobserved types indexed by ℓ. The vector ω^ℓ represents unobserved heterogeneity, and $\pi_{\ell|x_{i1}}$ denotes the mass (to be estimated) of type ℓ individuals conditional on the individual's initial value of the state variables. We distinguish three components in ω^ℓ, i.e., $\omega^\ell = (\omega_u^\ell, \omega_Y^\ell, \omega_f^\ell)$, which correspond to heterogeneity in utility, payoff and transition rules, respectively. In example 2 there is heterogeneity in utility and payoffs but not in transition probabilities because the laws of motion of schooling and experience are deterministic. The set of structural parameters consists of: $\pi = \{\pi_{\ell|x} : \ell = 1, 2, ..., L; x \in X\}$; the set of values $\Omega = \{\omega^\ell : \ell = 1, 2, ..., L\}$; and the vector $\theta = (\theta_u, \theta_Y, \theta_f)$, that is invariant across types. As we presented in section 2, the contribution of agent i to the conditional log-likelihood in this mixture model is $l_i(\theta, \Omega, \pi) = \log(\sum_{\ell=1}^{L} \pi_{\ell|x_{i1}} \; L_i(\theta, \omega^\ell))$, where $L_i(\theta, \omega^\ell)$ is $\prod_{t=1}^{T_i} P(a_{it}|x_{it}, \omega^\ell, \theta) \; f_Y(y_{it}|a_{it}, x_{it}, \omega_Y^\ell, \theta) \prod_{t=1}^{T_i-1} f_x(x_{i,t+1}|a_{it}, x_{it}, \omega_f^\ell, \theta_f)$. Conditional on type, the likelihood factors into conditional choice probabilities, payoff probabilities, and state transition probabilities as in Rust's model. However, unlike the log-likelihood in (10), this

[47] There are approaches other than FIML to estimate consistently the wage equation controlling for selection. For instance, the model can provide exclusion restrictions (i.e., variables that affect occupational choice but not wages) which can be used to estimate the wage equation using a Heckman two-step method. The first step of this method is a *reduced form* finite-mixture, multinomial probit model for occupational choice where the vector of explanatory variables consists of the terms of a polynomial in x_{it}.

log-likelihood is not additively separable because the type proportions appear inside the log.[48] Consistent estimates of θ_Y and θ_f cannot be obtained separately because agents' choices, which appear in the payoff and transition probabilities, are not independent of the individual unobserved types. FIML estimation corrects the endogeneity bias. This approach to allow for persistent individual heterogeneity is based on the seminal work by Heckman and Singer (1984).

The computational cost of estimation may be much larger than in a similar model without permanent unobserved heterogeneity because the likelihood of the choice history is maximized in the full parameter vector and furthermore the number of times the DP problem needs to be solved is multiplied by the number of types. As noted before this is the reason why finite mixtures are used and the number of types is kept small. Following Heckman and Singer, if L is unknown the number of types can be estimated by increasing L until the likelihood at the FIML estimates 'does not increase'. However, this search is expensive and typically L is set a priori. One should be careful not to choose a large value for L. If the value of L in the estimated model is larger than its true value in the population, the model is not identified (e.g., there are multiple combinations of the $\pi_{\ell|x_{i1}}$ parameters that can explain the data equally well). This is another reason why most applications have considered a small number of types.

Sequential EM algorithm (ESM): Arcidiacono and Jones (2003) propose a clever algorithm which makes the two-step partial likelihood approach compatible with the finite mixture model. Their insight is that additive separability of the log-likelihood, which is the basis for the two-step partial likelihood strategy, can be recovered in a 'multi-cycle' or 'sequential' version of the well known Expectation-Maximization (EM) algorithm. A key element of Arcidiacono and Jones' ESM algorithm is the 'posterior' probability that individual i belongs to unobserved type ℓ given her observed history of choices and states. For the sake of simplicity, ignore for the moment the initial conditions problem and assume that $\pi_{\ell|x_{i1}} = \pi_\ell$. We also assume that conditional on type, Rust's assumption CI-Y holds, i.e., conditional on type the payoff function y_{it} is independent of the transitory shocks ε_{it}. Let $\widetilde{a}_i$, $\widetilde{x}_i$ and $\widetilde{y}_i$ be individual i's histories of actions, states, and payoff variables, respectively. From Bayes' theorem this is

$$\Pr(\ell|\widetilde{a}_i,\widetilde{x}_i,\widetilde{y}_i;\theta,\Omega,\pi) = \frac{\pi_\ell \ \Pr\left(\widetilde{a}_i,\widetilde{x}_i,\widetilde{y}_i \mid \ell;\theta,\pi\right)}{\Pr\left(\widetilde{a}_i,\widetilde{x}_i,\widetilde{y}_i \mid \theta,\Omega,\pi\right)} = \frac{\pi_\ell \ L_i(\theta,\omega^\ell)}{\exp\left\{l_i(\theta,\Omega,\pi)\right\}} \tag{44}$$

where the functions $L_i(\theta,\omega^\ell)$ and $l_i(\theta,\Omega,\pi)$ are the likelihoods that we have defined in equation (16). It can be shown that the FIML estimator $(\hat{\theta},\hat{\Omega},\hat{\pi})$ that maximizes the likelihood function

[48] If there is not permanent unobserved heterogeneity in the transition probabilities, then we can write the full likelihood as the sum of two partial likelihoods and $\boldsymbol{\theta}_f$ can be still estimated consistently by maximizing the partial likelihood associated with transition data and without solving the DP problem.

(16) satisfies the following conditions:

$$\begin{aligned} \text{(a)} \quad \hat{\pi}_\ell &= \frac{1}{N}\sum\nolimits_{i=1}^{N} \Pr(\ell \mid \widetilde{a}_i, \widetilde{x}_i, \widetilde{y}_i; \hat{\theta}, \hat{\Omega}, \hat{\pi}) \\ \text{(b)} \quad (\hat{\theta}, \hat{\Omega}) &= \arg\max_{\{\theta,\Omega\}} \sum\nolimits_{i=1}^{N} \sum\nolimits_{\ell=1}^{L} \Pr(\ell \mid \widetilde{a}_i, \widetilde{x}_i, \widetilde{y}_i; \hat{\theta}, \hat{\Omega}, \hat{\pi}) \ \log L_i(\theta, \omega^\ell) \end{aligned} \tag{45}$$

Condition (a) is quite intuitive since it states that unconditional type proportions and individual posterior probabilities have to be mutually consistent. Condition (b) states that $(\hat{\theta}, \hat{\Omega})$ also maximizes a mixture of log-likelihoods, weighted by posterior type probabilities. In this maximization the posterior weights are kept fixed and appear outside the logs, so the mixture of log likelihoods is once again separable in choice, payoff and state transition factors.

These properties motivate Arcidiacono and Jones' sequential version of the EM algorithm. The algorithm is initialized with an arbitrary vector $\{\hat{\theta}_0, \hat{\Omega}_0, \hat{\pi}_0\}$ in the space of the structural parameters. Given $\{\hat{\theta}_0, \hat{\Omega}_0, \hat{\pi}_0\}$ we obtain a new vector $\{\hat{\theta}_1, \hat{\Omega}_1, \hat{\pi}_1\}$ as follows:

"E" step: Compute $P_{i\ell 0} \equiv \Pr(\ell | \widetilde{a}_i, \widetilde{x}_i, \widetilde{y}_i; \hat{\theta}_0, \hat{\Omega}_0, \hat{\pi}_0)$ as $\hat{\pi}_{\ell 0} L_i(\hat{\theta}_0, \hat{\omega}_0^\ell) / \exp\{l_i(\hat{\theta}_0, \hat{\Omega}_0, \hat{\pi}_0)\}$.

Sequential "M" step: For $\{P_{i\ell 0}\}$ fixed, obtain $\{\hat{\theta}_1, \hat{\Omega}_1, \hat{\pi}_1\}$ using:

$$\text{(a)} \quad \hat{\pi}_{\ell 1} = \frac{1}{N}\sum_{i=1}^{N} P_{i\ell 0}$$

$$\text{(b1)} \quad (\hat{\theta}_{f1}, \hat{\omega}_{f1}^\ell) = \arg\max_{\{\theta_f, \omega_f^\ell\}} \sum_{i=1}^{N}\sum_{\ell=1}^{L} P_{i\ell 0} \left[\sum_{t=1}^{T_i - 1} \log f_x(x_{i,t+1} | a_{it}, x_{it}, \omega_f^\ell, \theta_f)\right]$$

$$\text{(b2)} \quad (\hat{\theta}_{Y1}, \hat{\omega}_{Y1}^\ell) = \arg\max_{\{\theta_Y, \omega_Y^\ell\}} \sum_{i=1}^{N}\sum_{\ell=1}^{L} P_{i\ell 0} \left[\sum_{t=1}^{T_i} \log f_Y(y_{it} | a_{it}, x_{it}, \omega_Y^\ell, \theta_Y)\right]$$

$$\text{(b3)} \quad (\hat{\theta}_{u1}, \hat{\omega}_{u1}^\ell) = \arg\max_{\{\theta_u, \omega_u^\ell\}} \sum_{i=1}^{N}\sum_{\ell=1}^{L} P_{i\ell 0} \left[\sum_{t=1}^{T_i} \log P(a_{it} | x_{it}, \omega_u^\ell, \hat{\omega}_{Y1}^\ell, \hat{\omega}_{f1}^\ell, \theta_u, , \hat{\theta}_{Y1}, \hat{\theta}_{f1})\right]$$

Then, use $\{\hat{\theta}_1, \hat{\Omega}_1, \hat{\pi}_1\}$ as the initial value and apply the "E" step and the sequential "M" step again. We proceed until convergence in $\{\hat{\theta}, \hat{\Omega}, \hat{\pi}\}$. Arcidiacono and Jones show that, if the algorithm converges, it obtains consistent and asymptotically normal estimators. These estimators are not asymptotically efficient because the sequential partial likelihood approach is used in every M-step. The algorithm requires multiple maximization steps, but each of them may be much less costly than full information maximization. Arcidiacono and Jones illustrate their method in a Monte Carlo experiment based on a model of schooling choices. In their experiments, the ESM delivers estimators much faster than FIML, with a small loss of precision.

Initial conditions: As illustrated in Example 2, if the model has permanent unobserved heterogeneity then the first observation x_{i1} on which the likelihood is conditioned is potentially an endogenous variable because it is not independent of the individual's type: i.e., $\pi_{\ell|x_{i1}} \equiv \Pr(\ell|x_{i1}) \neq \Pr(\ell)$. Following Heckman (1981), there are two standard solutions to this problem. The first solution, which has been the most common approach in life-cycle models, is to complement the conditional likelihood derived from the structural DP model with an auxiliary model for the distribution of types conditional on the initial value of the state variables. A sufficiently flexible multinomial logit model can approximate arbitrarily well the distribution of type ℓ as a function of x_{i1}. That is, for $\ell \leq L-1$:

$$\pi_{\ell|x_{i1}} = \frac{\exp\{x_{i1}\ \theta_{\pi,\ell}\}}{1 + \sum_{\ell'=1}^{L-1} \exp\left\{x_{i1}\ \theta_{\pi,\ell'}\right\}} \tag{46}$$

where $\theta_\pi \equiv \{\theta_{\pi,\ell} : \ell = 1, 2, ..., L-1\}$ is a vector of parameters to estimate, and $\theta_{\pi,L}$ is normalized to zero. This seems like the most reasonable approach if the researcher thinks that the structural model does not apply to pre-sample periods. For instance, in Keane-Wolpin occupational choice model, the authors did not believe that their structural model explained schooling at age 16, but still treated this variable as endogenous.

In the second approach, it is assumed that the structural behavioral model explains the distribution of the initial values of the state variables. Consider again Keane-Wolpin's occupational choice model. The initial age in their data and model (i.e., $t = 1$) is 16 years old. For the sake of illustration, suppose that at some out-of-sample age lower than 16 (i.e., $t_0 < 1$) the state variables took the same value for all individuals. That is, all the individuals at age t_0 have the same level of formal education and the same (zero) labor market experience. Therefore, for every individual i, $x_{it_0} = x_0$ where x_0 is known to the researcher. Also, assume that individuals' choice probabilities at ages in the sample (i.e., $t \geq 1$) can be extrapolated to ages younger than 16. Then, given the choice probabilities of the structural model, we can obtain the probabilities $\Pr(x_{i1}|x_{it_0} = x_0, \omega^\ell, \theta)$ for every type ℓ and every value x_{i1} in the sample. By Bayes rule, $\pi_{\ell|x_{i1}} = \pi_\ell \Pr(x_{i1}|x_0, \omega^\ell, \theta) / \sum_{\ell'=1}^{L} \pi_{\ell'} \Pr(x_{i1}|x_0, \omega^{\ell'}, \theta)$. Therefore, given the probabilities $\Pr(x_{i1}|x_0, \omega^\ell, \theta)$, we can construct the conditional log-likelihood function $\log(\sum_{\ell=1}^{L} \pi_{\ell|x_{i1}} L_i(\theta, \omega^\ell))$, where now the vector of parameters π contains the unconditional mass probabilities of each type $\{\pi_\ell : \ell = 1, 2, ..., L\}$, which are primitive parameters of the structural model and are estimated together with θ. The key assumption for the validity of this approach is the extrapolation of choice probabilities for periods $t < 1$. If this assumption is correct, this "structural" approach to deal with the initial conditions problem provides more efficient estimates of the structural parameters than the first, "reduced form" approach. However, the approach has two important limitations:

it is computationally much more intensive, and it relies on out-of-sample extrapolations which may be not realistic in some applications.

In some applications some individuals' histories may be left-censored which implies that individuals have different initial periods. For instance, in Keane and Wolpin's occupational choice model suppose we do not observe all the individuals since age 16 but from very different initial ages, e.g., age 24, 28, etc. If we use the "reduced form" approach to deal with the initial conditions problem, we will have to allow the parameters $\theta_{\pi,\ell}$ to vary in a flexible way with the individual's initial age. In this context, the number of θ_π parameters to estimate may be very large and therefore the estimation of all the parameters can be quite inefficient. Keane and Wolpin (2001) propose and implement a simulation estimation method which deals with this problem, and more generally with the problem of missing state or choice data, which is in the same spirit as the "structural approach" that we have described above. They simulate complete outcome histories and match them to incomplete observed histories in order to compute the probabilities of the latter, allowing for measurement error in order to avoid degeneracy. They use it to estimate a dynamic model of schooling choices with savings decisions and borrowing constraints. More recently, an alternative approach to the initial conditions problem has been explored by Aguirregabiria and Mira (2007) (see section 4.2).

3.2.2 Nested Backwards Induction with Simulation and Interpolation

Keane and Wolpin's simulation and interpolation method has been the most widely used for applications with large state spaces and unobservables which are correlated across choices, beginning with the occupational choice model. The estimation criterion is FIML and individual contributions to the likelihood are the finite mixtures shown in (16). Conditional on an individual's unobserved type Assumption CI-X holds so the likelihood factors into conditional choice probabilities and the solution of the DP problem is characterized by the Emax function. However, the choice-specific unobservables do not have extreme value distributions and are correlated so the CCP's and Emax functions do not have closed forms. Computing them involves solving J-dimensional integrals at every point x in the state space. Keane and Wolpin use Monte Carlo integration to simulate these multiple integrals. Furthermore, for every time period the Emax integrals are simulated at a subset of the state space points only, and their value at every other point is interpolated using a regression function which is fit to the points in the subset.

As before, solving the model essentially amounts to obtaining the Emax function. In a finite horizon model, this is done by backwards induction. Let $\bar{V}_{\ell t}(x_t)$ be the integrated value function,

or Emax function, at period t and for type ℓ, as defined in section 2.1. Let $\{\varepsilon_t^{(r)} : r = 1, 2, ..., R\}$ be R random draws of histories of an individual's unobservables. Using these random draws we can construct simulation-based approximations (simulators) for the Emax function. Starting at the last period T, the simulator of $\tilde{V}_{\ell T}(x)$ is:[49]

$$\tilde{V}_{\ell T}(x) = \frac{1}{R}\sum_{r=1}^{R} \max_{a\in A}\left\{U_T(a, x, \varepsilon_T^{(r)}, \omega^\ell, \theta)\right\} \tag{47}$$

At period $t < T$ we already know the simulator of next period's Emax function, $\tilde{V}_{\ell,t+1}(.)$. Then, the simulator of the Emax function at period t is:

$$\tilde{V}_{\ell t}(x) = \frac{1}{R}\sum_{r=1}^{R} \max_{a\in A}\left\{U_t(a, x, \varepsilon_t^{(r)}, \omega^\ell, \theta) + \beta \sum_{x'\in X} \tilde{V}_{\ell,t+1}(x')\ f_x(x'|a, x)\right\} \tag{48}$$

Note that these Emax values should be calculated at every point x in the support X. This can be very costly for DP problems with large state spaces. In order to alleviate this computational burden the method obtains simulated-based approximations to the Emax function only at a (randomly chosen) subset $\bar{X}_t$ of the state points every period. The Emax at other points are obtained as the predicted values from a regression function which is estimated from the points in $\bar{X}_t$. In the model of career decisions, the arguments of the regression function were $\{\overline{v}_t(a, x) - \max[\overline{v}_t(1, x), \ldots, \overline{v}_t(J, x)]\}$, where:

$$\overline{v}_t(a, x) \equiv E_{\varepsilon_t}\left[U_t(a, x, \varepsilon_t, \omega^\ell, \theta)\right] + \beta \sum_{x'\in \bar{X}_{t+1}} \tilde{V}_{\ell,t+1}(x')\ f_x(x'|a, x) \tag{49}$$

That is, the interpolating function depends on the state through choice-specific value functions only. This interpolating function worked very well in this example but the arguments are costly to compute and the full state space has to be spanned in either simulation or interpolation. Using a polynomial in the state variables is much cheaper because in order to approximate the Emax at state x, we don't need to compute $\{\overline{v}_t(a, x)\}$ and because of this we don't need to know Emaxes at all states that may visited in the future from x. Monte Carlo experiments reported in Keane and Wolpin (1994) show that the method performed a bit less well in this case. However, polynomial approximations have been used in most subsequent applications.

Given the Emax functions, the conditional choice probabilities and the conditional density of the payoff variables can be obtained using simulation.[50] The parameters are estimated by FIML.

[49] We describe a version of Keane-Wolpin method that uses a simple frequency simulator. However, more efficient simulators can be used.

[50] Simulation of CCP's is needed only at sample points; kernel smoothing is used in this case, in order to avoid empty cells and to enable the use of gradient methods in the maximization of the likelihood. Note that the same draws of $\varepsilon's$ are used for the simulation of CCPs and the conditional density of wages. See Hajivassiliou and Rudd (1994) for details on simulation of different types of limited dependent variables models.

As Keane and Wolpin noted, the approximation errors in the Emax functions enters non-linearly in the CCP's. Therefore, simulated CCP's are biased and this implies that estimators of structural parameters are not consistent.[51] To put this problem into perspective, note that approximation error in the Emax is not the only source of inconsistency; e.g., discretization of continuous variables, approximate convergence of the Bellman operator in infinite horizon problems, etc. Furthermore, the large computational gain involved has allowed researchers using this method to produce many interesting applications, estimating models with very large state spaces and richer structures than would otherwise be possible.

3.3 Other issues in the estimation of single-agent models

3.3.1 Serially correlated, time-variant unobservables

We start this section reviewing methods and applications for models where the unobservables follow stochastic process with serial correlation. The seminal paper by Pakes (1986) on patent renewal was the first application to estimate this type of model. There are two main issues in the estimation of this class of models. First, the observable state variable x_t is not a sufficient statistic for the current choice and the probability of an individual's choice history cannot be factored into CCP's conditional on x_t alone. Therefore, computing that individual contribution involves solving an integral of dimension T_i. Pakes' paper was one the first econometric applications that used Monte Carlo simulation techniques to approximate high dimensional integrals.[52] Since Pakes's paper, there have been very important contributions in the areas of Monte Carlo integration (see Geweke, 1996) and simulation-based estimation methods (see Stern, 1997). In the context of dynamic discrete choice models, an important contribution was the development of the *Geweke-Hajivassiliou-Keane* (GHK) simulator. This is a very efficient importance-sampling simulator of multinomial probabilities in discrete choice models with normally distributed unobservables. The use of this simulator reduces significantly the approximation error and thus the bias and variance of simulation-based estimators.

A second important issue is that a DP problem with continuous state variables - observable or unobservable - cannot be solved exactly and needs to be approximated using interpolation methods or polynomial approximations.[53] To illustrate this issue consider the occupational choice

[51] A bias remains as long as interpolation is used, even if the number of simulation draws goes to infinity.

[52] See also the seminal work by Lerman and Manski (1981).

[53] This second issue is not a problem in Pakes' patent renewal model. A nice feature of that application is that the particular structure of the model (i.e., optimal stopping problem and the specification of the stochastic process of ε_t) is such that it is possible to obtain an exact recursive solution of the threshold values that characterize the optimal decision rule.

model in example 2 where, for every choice alternative, we omit the random effect $\omega_i(a)$ but relax the IID assumption in $\varepsilon_{it}(a)$. For instance, suppose that $\varepsilon_{it}(a)$ follows an AR(1) process, $\varepsilon_{it}(a) = \rho_a \varepsilon_{i,t-1}(a) + \xi_{it}(a)$. The value function of this DP problem depends on ε_{it}, which is a vector of continuous variables, and cannot be solved exactly. Note that the problem cannot be solved by considering the integrated value function, as in models where ε_{it} is iid, because it is no longer the case that $\bar{V}(x_{it}) \equiv \int V(x_{it}, \varepsilon_{it}) dG_\varepsilon(\varepsilon_{it})$ fully characterizes the solution of the DP problem.[54] There are two classes of approximation methods to solve this type of DP problems: interpolation methods, and polynomial approximations. Stinebrickner (2000) discusses these methods in the context of a dynamic discrete structural models with serial correlation and he presents some examples to illustrate the relative strengths of the various approximation approaches. His experiments suggest that, at least for models with normally distributed variables, interpolating methods based on Hermite and Gauss-Legendre quadrature perform very well even when the degree of serial correlation is high. Bound, Stinebrickner and Waidman (2005) have recently applied Gauss-Legendre quadrature interpolation to solve and estimate a structural model of retirement where a component of an individual's health is unobserved to the researcher and it is a continuous and serially correlated random variable. The study in Benitez-Silva et al (2005) presents a very extensive comparison of different strategies for solving dynamic programming problems with continuous, serially correlated state variables. One of the methods considered in that paper is the parameterized policy iteration method. This solution method has been used by Hall and Rust (2003) in the estimation of a model of inventory investment and price speculation by a durable commodity intermediary. The simulation-interpolation method in Keane and Wolpin (1994) could also be used for this class of models.

A recent paper by Hendel and Nevo (2006) presents an interesting and useful approximation method for the estimation of dynamic demand models with large state spaces.[55] Hendel and Nevo estimate a dynamic model for the demand of a differentiated storable good (laundry detergent) using consumer scanner data. They use the estimated model to study long-run and short-run demand responses to temporary price reductions (i.e., sales promotions).[56] The state space in this model includes prices and advertising expenditures for all brands in all sizes of the product: more than one

[54] Note that the integrated value function $\bar{V}(x_{it}, \varepsilon_{i,t-1}) \equiv \int V(x_{it}, \rho\varepsilon_{it} + \xi_{it}) dG_\xi(\xi_{it})$ fully characterizes the solution of the DP problem. However, $\bar{V}(x_{it}, \varepsilon_{i,t-1})$ has the same dimension as $V(x_{it}, \varepsilon_{it})$.

[55] Hendel and Nevo's approach has clear similarities with the use of Gittins's indexes in the estimation of an occupational choice model in Miller (1984).

[56] The empirical results are really striking: static demand estimates overestimate own-price elasticities by 30% and underestimate cross-price elasticities by 500%. Therefore, the biases of ignoring consumer intertemporal substitution are very large and they have important economic implications in the standard applications of demand systems such as measure of market power, antri-trust and merger analysis or the evaluation of the effects of a new good.

hundred continuous state variables. Hendel and Nevo show that in their model the probability of choosing a brand conditional on quantity does not depend on dynamic considerations, and therefore many of the demand parameters can be estimated from a static brand-choice model without solving the DP problem. Once these parameters are estimated, it is possible to construct a single index (or inclusive value) for each quantity choice (four quantity choice alternatives). These four indexes summarize all the relevant information in prices and advertising expenditures for the current utility of an individual. Then, Hendel and Nevo assume that these inclusive values follow a first order Markov process: i.e., all the information in current prices and advertising that is relevant to predict next period inclusive values can be summarized in today's inclusive values. Under this assumption, more than hundred state variables can be summarized in just four state variables. This is a very convenient approach in the estimation of dynamic demand models of differentiated products.[57] Even under these assumptions the solution of the DP problem in this model is cumbersome: the vector of state variables consists of the inclusive value, consumer inventory, and an unobserved shock in the marginal utility of consumption, which are all continuous variables. Thus, Hendel and Nevo should use interpolation techniques to approximate the solution of the DP problem. They use the *parameterized policy iteration* method in Benitez-Silva et al (2005) as the 'inner algorithm' in a Rust's nested fixed point algorithm.

3.3.2 Approximation error and inference

Numerical methods provide only approximations to the solution of dynamic decision models with continuous state variables. Therefore, the researcher cannot calculate the exact likelihood function (or other sample criterion function) of the model but only an approximated likelihood based on his approximated solution to the model. An important question is what are the implications for statistical inference of using an approximated likelihood function. The literature on simulation-based estimation has dealt with the implications of simulation errors on the asymptotic properties of estimators. However, much less is known when the approximation error does not come from using Monte Carlo simulators but from other numerical methods such as interpolation techniques. The standard practice in applications that use interpolation techniques has been to conduct inference as if the exact solution of the model were used and to ignore the effects of approximation errors. In this context, the recent paper by Fernandez-Villaverde, Rubio-Ramirez and Santos (2006) contains some

[57] Nevertheless, under realistic specifications for the stochastic process of prices and brand-choice parameters, the assumption of a Markov process for inclusive value can be clearly rejected. On the positive side, this is an assumption that can be tested empirically and that can be relaxed to a certain extent, e.g., a higher order Markov process. See Erdem, Imai and Keane (2003) for a similar application with different assumptions and estimation method.

important contributions to this difficult topic. They show that convergence of the approximated policy function to the exact policy function does not necessarily implies that the approximated likelihood function also converges to the exact likelihood. Some additional conditions are needed for convergence of the likelihood function. In particular, in addition to regularity conditions to have a well defined likelihood function, the optimal decision rule and the transition rule of the state variables, as functions of the vector of structural parameters, should be continuously differentiable and have bounded partial derivatives. They also propose a likelihood ratio test to check for the importance of errors in the approximated likelihood. Suppose that a researcher is using interpolation methods to approximate the solution of a DP model and that he solves and estimates the model under two different levels of approximation error, e.g., two different grids of points in the state space. We can interpret the two different approximations as two competing models. We want to test if the data significantly support one approximation over the other one. Let $\hat{l}^{(1)}$ and $\hat{l}^{(2)}$ be the maximum values of the two likelihood functions under the two levels of approximation error. The likelihood ratio statistic is $LR = \hat{l}^{(1)} - \hat{l}^{(2)}$. Vuong (1989) develops the asymptotic behavior of this statistic for both nested and non-nested models and his results are general enough to include the case we consider here. Based on this test, Fernandez-Villaverde et al. suggest to use an increasing approach to choose the degree of accuracy in the numerical solution of the model. That is, to increase the accuracy of the numerical solution until the likelihood ratio test cannot reject that the less accurate solution is statistically equivalent to the more accurate solution.

3.3.3 Bayesian methods

The computational cost of evaluating the likelihood in structural dynamic discrete choice models has sofar made Bayesian inference in these models intractable.[58] If a Markov Chain Montecarlo algorithm (MCMC) is used, the number of likelihood evaluations that is required is typically much larger than in the case of algorithms tailored to classical likelihood-based inference. In a recent paper, Imai , Jain and Ching (2007) propose an estimation method which promises to alleviate this problem significantly. At each iteration, a modified Metropolis-Hastings step is combined with a single iteration in the Bellman equation which updates the value function and replaces the full solution of the DP problem. Therefore, estimation and solution proceed simultaneously. This is in the same spirit as Aguirregabiria and Mira's NPL algorithm, but gradual updating of the DP solution is carried out in 'value function space' rather than in 'conditional choice probability space'.

[58] An exception is Lancaster (1997) who demonstrates the feasibility of Bayesian inference in a search model with closed form solutions. Geweke and Keane (1995) show how to do Bayesian inference when the future component of the value function is replaced by a flexible approximation in order to avoid the burden of full solution.

Imai et al show that the approximate solution converges to the full solution and the approximate posterior converges to the true posterior. They illustrate their method in Montecarlo experiments with uninformative priors for models of firm entry and exit with random effects and continuous state variables.

3.3.4 Validation of dynamic structural models

One of the most attractive features of dynamic structural models is that they can be used to predict the effects of counterfactual policy changes (or *ex ante policy evaluation*). Clearly, a prerequisite for such an exercise is that the researcher has enough confidence in the estimated structural model, i.e., the model needs to be validated. Some of the most commonly used criteria for validation of dynamic structural models have been: (a) conformity with the researcher's priors on "admissible" ranges of parameter values, based on economic theory and/or previous empirical work; (b) informal assessment of goodness of fit; (c) formal specification tests such as goodness of fit and over-identifying restrictions tests (see Andrews, 1988, and Rust, 1994). However, these criteria may seem insufficient for the purpose of credible policy evaluation. First, many researchers are concerned about identification. Second, goodness of fit and over-identifying restrictions tests are of limited usefulness when the estimated model is the result of pretesting (or "structural data mining", as referred by Wolpin, 2007). Alternative specifications with very similar *in-sample* performance may provide very different *out-of-sample* predictions of the effects of counterfactual policies. In fact, as pointed out by Wolpin (2007), it can be the case that some model features that contribute to improve in-sample goodness-of-fit may have a negative effect on the performance of the model for the prediction of counterfactual experiments. Then, how should we choose among these models? First, if the structural model will be use to predict a counterfactual policy, it seems reasonable that the model should be judged in terms of its ability to predict that particular policy. In this sense, the "best" model depends on the type of counterfactual policy one wants to predict. Social experiments and exogenous regime swifts provide very useful information for the credible validation of a structural model. Subject to identification issues, one can estimate the structural model using the control-group subsample and then use the experimental group to evaluate how well the model predicts the effects of the policy intervention which was the object of the social experiment. Once a model has been selected in this way, it can be used to predict extrapolations of the policy in the social experiment. This idea has been used before in dynamic structural models by Lumsdaine, Stock and Wise (1992), to predict retirement behavior under alternative pension plans, and by Todd and Wolpin (2006), to predict the effect of subsidies on children school attendance in Mexico.

However, for good or for bad, social experiments and regime swifts are a rarity. In a recent paper, Keane and Wolpin (2007) propose an approach for model validation that is in this same spirit but can be used when non-experimental data is available. The approach consists in holding out a part of the sample that has the characteristics of a policy change similar (in some way) to the counterfactual policy that we want to evaluate. The hold out sample plays the same role as the treatment group sample when a social experiment is available: it is not used for estimation but it is used to validate the predictive ability of the model. For instance, a researcher who wants to predict the effect of a child care subsidy program on a female labor supply model may estimate model using the subsample of females with two or more children and holdout the sample of females with less than two children. This idea seems very interesting and useful, though the fact the subsample selection is not random introduces nontrivial econometric issues.

3.3.5 Relaxing rational expectations using subjective expectations data

The assumption of rational expectations has been ubiquitous in this literature. Nevertheless, this is a very strong assumption in many applications. Information is costly and individuals typically make predictions using different sets of partial information. Though the rational expectations assumptions is made for identification reasons (i.e., observed choices may be consistent with many alternative specifications of preferences and expectations), it can induce serious biases in our predictions of counterfactual experiments. Data on expectations can be used to relax or to validate assumptions about expectations, and it can make the predictions of dynamic structural models more credible. Manski (2004) presents an excellent review on the use of data on subjective expectations in microeconometric decision models. In the context of dynamic discrete choice structural models, an example is the study by van der Klaauw and Wolpin (2005) on Social Security and savings. The study uses data on subjective expectations from the Health and Retirement Survey (expectations over the own retirement age and longevity) and from the Survey of Economic Expectations (expectations over future changes in Social Security policy).

4 Estimation methods for dynamic discrete games

As we did for single agent models, we distinguish between the structural parameters in the utility function, θ_u, in the payoff function, θ_Y, and in the transition rules of the state variables, θ_f. Under assumptions CI-X and CI-Y, for the estimation of θ_Y and θ_f we do not have to deal with the problem of calculating a MPE of the game. We can estimate θ_Y and θ_f as the

vectors that maximize the partial likelihood functions $\sum_{m=1}^{M}\sum_{t=1}^{T_m-1}\log f_Y(y_{imt}|a_{imt},x_{mt};\theta_Y)$ and $\sum_{m=1}^{M}\sum_{t=1}^{T_m-1}\log f_x(x_{m,t+1}|a_{mt},x_{mt};\theta_f)$, respectively. We assume that β is known and that θ_Y and θ_f have been estimated in a first step and we focus on the estimation of the parameters in θ_u, which requires that we exploit the equilibrium structure of the game. Following most of the papers in this literature, we start considering that the observations in the data have been generated by only one Markov Perfect equilibrium.

ASSUMPTION One-MPE-Data: Define the distribution of a_{mt} conditional on x_{mt} in market m at period t as $\mathbf{P}_{mt}^0 \equiv \{\Pr(a_{mt}=a|x_{mt}=x):(a,x)\in A^N\times X\}$. (A) For every observation (m,t), $\mathbf{P}_{mt}^0=\mathbf{P}^0$. (B) Players expect $\mathbf{P}^0$ to be played in future (out of sample) periods. (C) The observations $\{a_{mt},x_{mt}\}$ are independent across markets and $\Pr(x_{mt}=x)>0$ for all x in X.

Assumption (A) establishes that the data has been generated by only one Markov Perfect equilibrium. Thus even if the model has multiple equilibria, the researcher does not need to specify an equilibrium selection mechanism because the equilibrium that has been selected will be identified from the conditional choice probabilities in the data. We will discuss in section 4.2 how this assumption can be relaxed. Assumption (B) is necessary in order to accommodate dynamic models. Without it, we cannot compute the expected future payoffs of within-sample actions unless we specify the beliefs of players regarding the probability of switching equilibria in the future.

Following the notation in section 2.3, let $\Lambda(v^{\mathbf{P}}(\theta))\equiv\{\Lambda(a_i|v_i^{\mathbf{P}}(.,x,\theta)):(i,a_i,x)\in I\times A\times X\}$ be the equilibrium mapping of the dynamic game. Furthermore, let $v_i^{\mathbf{P}}(a_i,x,\theta)$ hereafter denote Hotz-Miller's representation of choice-specific value functions, adapted to this context (see details below). Aguirregabiria and Mira (2007) show that, for given θ, a vector of conditional choice probabilities $\mathbf{P}$ is a MPE of the game if and only if satisfies the fixed point condition, $\mathbf{P}=\Lambda(v^{\mathbf{P}}(\theta))$. For the description of several estimators in this literature, it is convenient to define the following *pseudo likelihood function*:

$$Q(\theta,\mathbf{P})=\sum_{m=1}^{M}\sum_{t=1}^{T_m}\sum_{i=1}^{N}\ln\Lambda(a_{imt}|v_i^{\mathbf{P}}(x_{mt},\theta)) \tag{50}$$

where $\mathbf{P}$ is an arbitrary vector of players' choice probabilities. This is a "pseudo" likelihood because the choice probabilities $\Lambda(a_{imt}|v_i^{\mathbf{P}}(x_{mt},\theta))$ are not necessarily equilibrium probabilities associated with θ, but just best responses to arbitrary beliefs $\mathbf{P}$ about other players' behavior. The MLE can be defined as:

$$\hat{\theta}_{MLE}=\arg\max_{\theta\in\Theta}\left\{\sup_{\mathbf{P}\in(0,1)^{N|X|}}Q(\theta,\mathbf{P})\quad subject\ to:\mathbf{P}=\Lambda(v^{\mathbf{P}}(\theta))\right\} \tag{51}$$

For given θ, the expression $\mathbf{P} = \Lambda(v^{\mathbf{P}}(\theta))$ defines the set of vectors $\mathbf{P}$ which are equilibria associated with that value of the structural parameters. For some values of θ that set may contain more than one $\mathbf{P}$ and therefore to obtain the MLE one should maximize over the set of equilibria. Under standard regularity conditions, multiple equilibria does not affect the standard properties of the MLE which in this model is root-M consistent, asymptotically normal and efficient. However, in practice, this estimator can be very difficult to implement. This is particularly the case if we use an algorithm that for each trial value of θ computes all the vectors $\mathbf{P}$ which are an equilibrium associated with θ and then selects the one with maximum value for $Q(\theta, \mathbf{P})$. Finding all the Markov Perfect equilibria of a dynamic game can be very difficult even for relatively simple models.

4.1 Two-Step methods

Several recent papers (Aguirregabiria and Mira (2007), Bajari, Benkard and Levin (2007), Jofre-Bonet and Pesendorfer (2003), Pakes, Ostrovsky and Berry (2007), and Pesendorfer and Schmidt-Dengler (2007)) have proposed different versions and extensions of Hotz-Miller CCP estimator to the estimation of dynamic games. An interesting aspect of the application of the CCP estimator to dynamic games is that this method deals with the problem of multiple equilibria, i.e., it avoids the optimization of the (pseudo) likelihood with respect to $\mathbf{P}$. Under assumption '*One-MPE-Data*', players' choice probabilities can be interpreted as players' beliefs about the behavior of their opponents. Given these beliefs, one can interpret each player's problem as a game against nature with a unique optimal decision rule in probability space, which is the player's best response. While equilibrium probabilities are not unique functions of structural parameters, the best response mapping is a unique function of structural parameters and players' beliefs about the behavior of other players. These methods use best response functions evaluated at consistent nonparametric estimates of players' beliefs.

We now describe different variants of this estimator in the context of dynamic games. As in the case of single agent models, the CCP method is particularly useful in models where the utility function is linear in parameters. Therefore, we assume that $u_i(a_t, x_t, \theta_u) = z_i(a_t, x_t)'\theta_u$, where $z_i(a_t, x_t)$ is a vector of known functions. Let $\mathbf{P}$ be a vector of conditional choice probabilities, for every player, state and action. Following the same approach as in single agent models, the alternative-specific value functions can be written as follows:

$$v_i^{\mathbf{P}}(a_i, x_t) = \tilde{z}_i^{\mathbf{P}}(a_i, x_t)\ \theta_u + \tilde{e}_i^{\mathbf{P}}(a_i, x_t) \tag{52}$$

where $\tilde{z}_i^{\mathbf{P}}(a_i, x_t)$ is the expected and discounted sum of current and future $z_i's$ that originate from

(a_i, x_t) given that all players behave now and in the future according to the probabilities in $\mathbf{P}$; and $\tilde{e}_i^{\mathbf{P}}(a_i, x_t)$ is the expected and discounted sum of the stream $\{\varepsilon_{i,t+j}(a_{i,t+j}) : j = 1, 2, ...\}$ given that the sequence of players' actions is generated by the choice probabilities in $\mathbf{P}$. More formally, we have that:

$$\begin{aligned} \tilde{z}_i^{\mathbf{P}}(a_i, x_t) &\equiv \sum_{a_{-i}} \left(\prod_{j \neq i} P_j(a_j|x_t) \right) \left[z_i(a_i, a_{-i}, x_t) + \beta \sum_{x'} f_x(x'|a_i, a_{-i}, x_t)\ W_{zi}^{\mathbf{P}}(x') \right] \\ \tilde{e}_i^{\mathbf{P}}(a_i, x_t) &\equiv \beta \sum_{a_{-i}} \left(\prod_{j \neq i} P_j(a_j|x_t) \right) \left[\sum_{x'} f_x(x'|a_i, a_{-i}, x_t)\ W_{ei}^{\mathbf{P}}(x') \right] \end{aligned} \tag{53}$$

where a_{-i} is the vector with the actions of all players other than i. $W_{zi}^{\mathbf{P}}(x)$ and $W_{ei}^{\mathbf{P}}(x)$ are valuation operators. Let $\mathbf{W}_i^{\mathbf{P}}$ be the matrix $\{(W_{zi}^{\mathbf{P}}(x), W_{ei}^{\mathbf{P}}(x)) : x \in X\}$. Then, this valuation operator is defined as the unique solution in $\mathbf{W}$ to the fixed point problem: $\mathbf{W} = \sum_{a \in A^N} [\prod_{j=1}^N \mathbf{P}_j(a_j)] * \{[\mathbf{z}_i(a), \mathbf{e}_i(a)] + \beta\ \mathbf{F}_x(a)\mathbf{W}\}$, and $\mathbf{e}_i(a)$ has the same definition as in the single agent model.

Let $\hat{\mathbf{P}}$, $\hat{\theta}_Y$ and $\hat{\theta}_f$ be consistent estimators of $\mathbf{P}^0$, θ_Y^0 and θ_f^0, respectively. Based on these initial estimates, we can obtain a two-step estimator of θ_u as a GMM estimator that solves the sample moment conditions:

$$\sum_{m=1}^{M} \sum_{i=1}^{N} \sum_{t=1}^{T_m - 1} H_i(x_{mt}) \begin{bmatrix} I\{a_{imt} = 1\} - \Lambda\left(0 \mid \tilde{z}_i^{\hat{\mathbf{P}}}(., x_{mt})\theta_u + \tilde{e}_i^{\hat{\mathbf{P}}}(., x_{mt})\right) \\ \vdots \\ I\{a_{imt} = J\} - \Lambda\left(J \mid \tilde{z}_i^{\hat{\mathbf{P}}}(., x_{mt})\theta_u + \tilde{e}_i^{\hat{\mathbf{P}}}(., x_{mt})\right) \end{bmatrix} = 0 \tag{54}$$

where $H(x_{mt})$ is a matrix with dimension $\dim(\theta_u) \times J$ with functions of x_{mt} which are used as instruments. The estimator is root-M consistent and asymptotically normal. This estimator is used by Jofre-Bonet and Pesendorfer (2003) and Pakes, Ostrovsky and Berry (2007). An attractive feature of this method of moments estimator, emphasized by Pakes, Ostrovsky and Berry (2007), is that when the matrix of instruments $H_i(x_{mt})$ does not depend on the nonparametric estimator $\hat{\mathbf{P}}$, this estimator can have lower finite sample bias than the pseudo maximum likelihood and the minimum distance estimators that we describe in the following paragraphs. We return to this issue at the end of this section.

The two-step method in Aguirregabiria and Mira (2007) is a pseudo maximum likelihood (PML) estimator that maximizes in θ_u the criterion function $Q(\theta_u, \hat{\theta}_Y, \hat{\theta}_f, \hat{\mathbf{P}})$. The values $\tilde{z}_i^{\hat{\mathbf{P}}}$ and $\tilde{e}_i^{\hat{\mathbf{P}}}$ are calculated as described above: i.e., solving for the matrices $\mathbf{W}_i^{\hat{\mathbf{P}}}$ and then applying the expressions in (53). This method is a particular case of the class of GMM estimators defined in equation (54). The (pseudo) likelihood equations that define this estimator can be expressed as in (54) with a matrix $H_i(x_{mt})$ equal to $diag\{\partial \log \Lambda_{imt}(j)/\partial \theta_u : j = 0, 1, ..., J\}$. This PML method is asymptotically

efficient within the class of GMM estimators described by conditions (54). However, in contrast to the case of single-agent models, this estimator is less efficient asymptotically than the partial MLE. This is because the initial nonparametric estimator of $\mathbf{P}^0$ and the PML estimator of θ_u^0 are not asymptotically independent and therefore there is an efficiency loss from using an inefficient initial estimator of $\mathbf{P}^0$.

Pesendorfer and Schmidt-Dengler (2007) consider the following class of minimum distance estimators:

$$\hat{\theta}_u = \arg\min_{\theta_u} \left[\hat{\mathbf{P}} - \mathbf{\Lambda}\left(v^{\hat{\mathbf{P}}}(\theta)\right)\right]' \mathbf{A}_M \left[\hat{\mathbf{P}} - \mathbf{\Lambda}\left(v^{\hat{\mathbf{P}}}(\theta)\right)\right] \tag{55}$$

where $\mathbf{A}_M$ is a weighting matrix that converges in probability to a non-stochastic positive definite matrix $\mathbf{A}_0$ as M goes to infinity. Different choices of weighting matrices give rise to distinct estimators within this class. Under standard regularity conditions, all the estimators in this class are consistent and asymptotically normal. Minimum distance estimation theory establishes that the efficient estimator in this class is the one where the weighting matrix $\mathbf{A}_0$ is:

$$\left(\left[\left(\mathbf{I} \vdots \mathbf{0}\right) - \nabla_{\mathbf{P},\theta_Y,\theta_f}\mathbf{\Lambda}\left(v^{\mathbf{P}_0}(\theta^0)\right)\right]' \Sigma \left[\left(\mathbf{I} \vdots \mathbf{0}\right) - \nabla_{\mathbf{P},\theta_Y,\theta_f}\mathbf{\Lambda}\left(v^{\mathbf{P}_0}(\theta^0)\right)\right]\right)^{-1} \tag{56}$$

where $(\mathbf{I}\vdots\mathbf{0})$ is the identity matrix vertically stacked with a matrix of zeros; $\nabla_{\mathbf{P},\theta_Y,\theta_f}\mathbf{\Lambda}$ is the Jacobian matrix of $\mathbf{\Lambda}$ with respect to $\mathbf{P}$, θ_Y and θ_f; and Σ is the variance matrix of the initial estimators $\hat{\mathbf{P}}$, $\hat{\theta}_Y$ and $\hat{\theta}_f$. Note that this optimal weighting matrix depends on θ_u^0. Therefore, the efficient estimator is obtained in three steps: estimate $\hat{\mathbf{P}}$, $\hat{\theta}_Y$ and $\hat{\theta}_{f0}$ and their variance; obtain an inefficient minimum distance estimator of θ_u^0; finally, construct a consistent estimator of the optimal weighting matrix and obtain the efficient estimator. Pesendorfer and Schmidt-Dengler show that this efficient estimator is asymptotically equivalent to MLE.

In models with continuous state variables or with large state spaces, the computation of continuation values $\tilde{z}_i^{\hat{\mathbf{P}}}$ and $\tilde{c}_i^{\hat{\mathbf{P}}}$ can be infeasible or extremely burdensome. Bajari, Benkard and Levin (2007) propose a method that builds on and extends the simulation-based CCP estimator that we have described in section 3.1.4. Their method has two important features that distinguish it from the other methods that we review here: it can be applied to models with continuous decision variables (as long as the utility function satisfies $\partial^2 u(a_i, a_{-i}, x, \varepsilon_i)/\partial a_i \partial \varepsilon_i \geq 0$), and to models where the parameters are not point identified (i.e., set identification). In fact, these two features of their method also contribute to the literature on estimation of single-agent dynamic structural models. Bajari, Benkard and Levin (BBL) propose an estimator that minimizes a set of moment inequalities and that can be applied to a general class of dynamic structural models under assumptions

AS, IID and CI-X, including dynamic games with either discrete or continuous decision and state variables. Define $W_i^{\mathbf{P}}(x) \equiv (W_{zi}^{\mathbf{P}}(x), W_{ei}^{\mathbf{P}}(x))$, and split the vector of choice probabilities $\mathbf{P}$ into the sub-vectors $\mathbf{P}_i$ and $\mathbf{P}_{-i}$, where $\mathbf{P}_i$ are the probabilities associated to player i and $\mathbf{P}_{-i}$ contains the probabilities of the other players. The model implies that for any state $x \in X$ and any $\mathbf{P}_i \neq \mathbf{P}_i^0$ the following inequality should hold:

$$W_i^{\mathbf{P}_i^0, \mathbf{P}_{-i}^0}(x) \begin{pmatrix} \theta_u \\ \mathbf{1} \end{pmatrix} \geq W_i^{\mathbf{P}_i, \mathbf{P}_{-i}^0}(x) \begin{pmatrix} \theta_u \\ \mathbf{1} \end{pmatrix} \tag{57}$$

Let H be a set of values for $(i, x, \mathbf{P})$. If the set H is large enough and θ_u^0 is identified, then θ_u^0 uniquely minimizes the population criterion function:

$$\sum_{\{i,x,\mathbf{P}\} \in H}^{2} \min \left\{ 0 \; ; \; \left(W_i^{\mathbf{P}_i^0, \mathbf{P}_{-i}^0}(x) - W_i^{\mathbf{P}_i, \mathbf{P}_{-i}^0}(x) \right) \begin{pmatrix} \theta_u \\ \mathbf{1} \end{pmatrix} \right\}^2 \tag{58}$$

This criterion function penalizes departures from the inequalities in (57). The Bajari-Benkard-Levin (BBL) estimator of θ_u^0 minimizes a simulation-based sample counterpart of this criterion function. More precisely,

$$\hat{\theta}_u = \arg\min_{\theta_u \in \Theta} \sum_{\{i,x,\mathbf{P}\} \in H} \left\{ 0 \; ; \; \left(\tilde{W}_i^{\hat{\mathbf{P}}_i, \hat{\mathbf{P}}_{-i}}(x) - \tilde{W}_i^{\mathbf{P}_i, \hat{\mathbf{P}}_{-i}}(x) \right) \begin{pmatrix} \theta_u \\ \mathbf{1} \end{pmatrix} \right\}^2 \tag{59}$$

where $\hat{\mathbf{P}}$ is a nonparametric estimator of $\mathbf{P}^0$ (there is also initial estimator of θ_Y^0 and θ_f^0 but we have omitted them for the sake of notational simplicity), and $\tilde{W}_i$ is a simulator of W_i which is obtained as described in section 3.1.4. The estimator is root-M consistent and asymptotically normal. The asymptotic variance of the estimator depends not only on the variance of $\hat{\mathbf{P}}$ but also on the number of simulations and, very importantly, on the choice of the set of "deviations" with respect to the optimal policy contained in H. Bajari et al. describe a bootstrap procedure to calculate standard errors (see also Chernozhukov, Hong and Tamer, 2007). Ryan (2006) uses this method to estimate a dynamic oligopoly model of the US cement industry. In his model, firms compete in quantities in a static equilibrium, but they are subject to capacity constraints. Firms invest in future capacity and this decision is partly irreversible (and therefore dynamic). Note that the decision variable in this model, investment, is a censored continuous variable. Ryan estimates the parameters in demand and marginal costs using data on prices, quantities and the static equilibrium conditions. In a second step he estimates investment costs as well as entry and exit costs using BBL's inequality estimator. Ryan allows entry costs to vary before and after year 1990 when several amendments were introduced in the Clean Air Act. He finds that these amendments raised significantly the sunk costs of entry in the cement industry.

The main advantage of these two-step estimators is their computational simplicity. However, they have two important limitations. The first problem is finite sample bias. The initial nonparametric estimator can be very imprecise in the small samples available in actual applications, and this can generate serious finite sample biases in the two-step estimator of structural parameters. In dynamic games with heterogeneous players the number of observable state variables is proportional to the number of players and therefore the so called *curse of dimensionality in nonparametric estimation* (and the associated bias of the two-step estimator) can be particularly serious. The sources of this finite sample bias can be illustrated using the moment conditions in (54): (1) if the matrix of instruments $H_i(x_{mt})$ depends on the nonparametric estimator $\hat{\mathbf{P}}_0$, then there is a finite sample correlation between these instruments and the "errors" $I\{a_{imt} = j\} - \Lambda(j|\ \tilde{z}_{imt}\theta_u + \tilde{e}_{imt})$; and (2) the choice probabilities $\Lambda(j|\ \tilde{z}_{imt}\theta_u + \tilde{e}_{imt})$ are complicated nonlinear functions of the nonparametric estimator $\hat{\mathbf{P}}_0$, and the expected value of a nonlinear function is not equal to the function evaluated at the expected value. As argued by Pakes, Ostrovsky and Berry (2007), the first source of bias is present in the pseudo maximum likelihood estimator and the minimum distance estimator but not in a simple method of moments estimator. However, the second source of bias appears in all these two steps estimators and it can be very important as illustrated in the Monte Carlo experiments of several papers (see the Monte Carlo experiments in Hotz et al., 1994, Aguirregabiria and Mira, 2002 and 2007, or Pesendorfer and Schmidt-Dengler, 2007).

A second important limitation of these two-step methods is the restrictions imposed by the IID assumption. Ignoring persistent unobservables, if present, can generate important biases in the estimation of structural parameters.

4.2 Sequential estimation

We have described in section 3.1.3 how a recursive or sequential CCP procedure is a bias reduction method that deals with the problem of finite sample bias of the two-step CCP estimator. This procedure can be particularly useful in the context of dynamic games with heterogeneous players because it is in this context where the finite sample bias of the two-step estimator can be very serious. The sequential CCP method or NPL algorithm also deals with the issue of permanent unobserved heterogeneity. Here we follow Aguirregabiria and Mira (2007) and describe the NPL method for a model with permanent unobserved market heterogeneity. Consider the entry-exit model in Example 4 but extended to include unobserved market heterogeneity. The profit of an

active firm is:

$$U_{imt}(1) = \theta_{RS} \log(S_{mt}) - \theta_{RN} \log\left(1 + \sum_{j \neq i} a_{jmt}\right) - \theta_{FC,i} - \theta_{EC,i}(1 - a_{im,t-1}) + \omega_m + \varepsilon_{imt} \quad (60)$$

where ω_m is a random effect interpreted as a time-invariant market characteristic affecting firms' profits, which is common knowledge to the players but unobservable to the econometrician. We assume that ω_m has a discrete and finite support $\Omega = \{\omega^1, \omega^2, \ldots, \omega^L\}$, and it is independently and identically distributed across markets with probability mass function $\pi_\ell \equiv \Pr(\omega_m = \omega^\ell)$. Furthermore, ω_m does not enter into the conditional transition probability of x_{mt}, i.e., $\Pr(x_{m,t+1}|a_{mt}, x_{mt}, \omega_m) = f_x(x_{m,t+1}|a_{mt}, x_{mt})$. This assumption implies that the transition probability function f_x can still be estimated from transition data without solving the model.

The introduction of unobserved market heterogeneity also implies that we can relax the assumption of only '*One MPE in the data*' to allow for different market types to have different equilibria. Let $\mathbf{P}^0_{mt} \equiv \{\Pr(a_{mt} = a|x_{mt} = x, m, t) : (a, x) \in A^N \times X\}$ be the distributions of a_{mt} conditional on x_{mt} in market m at period t. We assume that $\mathbf{P}^0_{mt} = \mathbf{P}^0_\ell$, where ℓ is the type of market m. Each market type has its own MPE. Though we still assume that only one equilibrium is played in the data *conditional on market type*, the data generating process may correspond to multiple equilibria. Markets which are observationally equivalent to the econometrician may have different probabilities of entry and exit because the random effect component of profits ω is different. Furthermore, though, in our example, market heterogeneity ω_m is payoff-relevant, this variable may also play (in part) the role of a *sunspot.*

The vector of structural parameters now includes the distribution of firm types: $\pi \equiv \{\pi_\ell : \ell = 1, 2, ..., L\}$ and $\Omega = \{\omega^1, \omega^2, \ldots, \omega^L\}$. The (conditional) pseudo likelihood function has the following finite mixture form:

$$Q(\theta, \pi, \Omega, \{\mathbf{P}_\ell\}) = \sum_{m=1}^{M} \ln\left(\sum_{\ell=1}^{L} \pi_{\ell|x_{m1}} \left[\prod_{t=1}^{T} \prod_{i=1}^{N} \Lambda\left(a_{imt} \mid \tilde{z}_i^{\mathbf{P}_\ell}(x_{mt})\theta_u + \tilde{e}_i^{\mathbf{P}_\ell}(x_{mt}) + \omega^\ell\right)\right]\right) \quad (61)$$

where $\pi_{\ell|x}$ is the conditional probability $\Pr(\omega_m = \omega^\ell|x_{m1} = x)$. It is clear that firms' incumbent statuses at period 1, which are components of the vector x_{m1}, are not independent of market type, i.e., more profitable markets tend to have more incumbent firms. Therefore, $\pi_{\ell|x_{m1}}$ is not equal to the unconditional probability π_ℓ. Under the assumption that x_{m1} is drawn from the stationary distribution induced by the MPE, we can obtain the form of $\pi_{\ell|x_{m1}}$. Let $\mathbf{p}^*(\mathbf{P}_\ell) \equiv \{p^*(x|\mathbf{P}_\ell) : x \in X\}$ be the stationary distribution of x induced by the equilibrium $\mathbf{P}_\ell$ and the transition $f_x(.|., \theta_f)$. This stationary distribution can be very simply obtained as the solution to the system of linear

equations:

$$\begin{aligned} p^*(x|\mathbf{P}_\ell) &= \sum_{x_0 \in X} p^*(x_0|\mathbf{P}_\ell)\ \Pr(x|x_0, \mathbf{P}_\ell) \\ &= \sum_{x_0 \in X} p^*(x_0|\mathbf{P}_\ell)\ \left(\sum_{a \in A^N} \left[\prod_{j=1}^{N} P_{\ell j}(a_j|x_0) \right] f_x(x|a, x_0) \right) \end{aligned} \tag{62}$$

Then, by Bayes' rule, we have that:

$$\pi_{\ell|x_{m1}} = \frac{\pi_\ell\ p^*(x_{m1}|\mathbf{P}_\ell)}{\sum_{\ell'=1}^{L} \pi_{\ell'}\ p^*(x_{m1}|\mathbf{P}_{\ell'})} \tag{63}$$

The NPL estimator is obtained using an iterative procedure similar to the one we have described in section 3.1.3 for a model without heterogeneity. The main difference is that now we have to calculate the steady-state distributions $\mathbf{p}^*(\mathbf{P}_\ell)$ to deal with the initial conditions problem. However, the pseudo likelihood approach also reduces very significantly the cost of dealing with the initial conditions problem. The reason is that given $\mathbf{P}_\ell$ the steady-state distributions do not depend on the structural parameters in θ_u. Therefore, the distributions $\mathbf{p}^*(\mathbf{P}_\ell)$ remain constant during any pseudo maximum likelihood estimation and they are updated only between two pseudo maximum likelihood estimations when new choice probabilities are obtained. This implies a very significant reduction in the computational cost associated with the initial conditions problem. Aguirregabiria and Mira (2007) also consider a distribution of ω_m that simplifies the computation of the NPL: $\omega_m = \sigma_\omega \omega_m^*$ where ω_m^* is a discretized standard normal. Therefore, the support of ω_m^* and the probabilities π_ℓ are known and the only parameter to be estimated is σ_ω. Given this distribution of ω_m, the probabilities $\pi_{\ell|x_{m1}}$ also remain constant during any pseudo maximum likelihood estimation. The algorithm proceeds as follows. We start with L arbitrary vectors of players' choice probabilities, one for each market type: $\{\hat{\mathbf{P}}_{\ell 0} : \ell = 1, 2, ..., L\}$. Then, we perform the following steps. Step 1: For every market type we obtain the steady-state distribution of x_{m1} and the probabilities $\{\pi_{\ell|x_{m1}}\}$. Step 2: We obtain the pseudo maximum likelihood estimator of θ_u and σ_ω as: $(\hat{\theta}_{u1}, \hat{\sigma}_{\omega 1}) = \arg\max_{(\theta_u, \sigma_\omega)} Q(\theta_u, \sigma_\omega, \hat{\theta}_Y, \hat{\theta}_f, \{\hat{\mathbf{P}}_{\ell 0}\})$. Step 3: Update the vector of players' choice probabilities using the best response probability mapping. That is, for market type ℓ, $\hat{\mathbf{P}}_{\ell 1} = \mathbf{\Lambda}(v^{\hat{\mathbf{P}}_{\ell 0}}(\hat{\theta}_{u1}, \hat{\sigma}_{\omega 1}, \omega_\ell^*, \hat{\theta}_Y, \hat{\theta}_f))$. If, for every type ℓ, $||\hat{\mathbf{P}}_{\ell 1} - \hat{\mathbf{P}}_{\ell 0}||$ is smaller than a predetermined small constant, then stop the iterative procedure and choose $(\hat{\theta}_{u1}, \hat{\sigma}_{\omega 1})$ as the NPL estimator. Otherwise, repeat steps 1 to 4 using $\{\hat{\mathbf{P}}_{\ell 1}\}$.

Collard-Wexler (2006) uses this method to estimate a dynamic oligopoly model of entry and exit in the US ready-mix concrete industry. He finds that including unobserved market heterogeneity increases very significantly the estimate of the effect of competitors on profits (i.e., the parameter

θ_{RN} in equation (60)). While the effect of a second competitor on profits is positive in the model without market effects, it is negative and significant when market heterogeneity is included.

5 General equilibrium models

In this section we describe the method proposed by Lee and Wolpin (LW) to estimate competitive equilibrium models. The model is that of Example 3, which is a simplified version of the model in LW with only one sector and two occupations as in Lee (2005). Total factor productivity z_t is assumed to follow an AR(1) process. This is the only source of exogenous aggregate uncertainty which is explicitly modelled and it implies that individuals solving the occupational choice model face uncertainty about future skill prices. Future skill prices depend on future TFP and on future cross-sectional distributions of schooling and occupation-specific experience. However, including these distributions in the vector of state variables $\widetilde{X}_t$ would make the dimension of the state space so large as to make solution and estimation infeasible. Lee and Wolpin assume that current and lagged values of skill prices provide a good approximation to the information contained in these distributions that is relevant to predict future skill prices. More specifically, LW assume that the evolution of skill prices is described by the following system of difference equations:

$$\ln r_{a,t+1} - \ln r_{a,t} = \eta_{a0} + \sum_{k=1}^{2} \eta_{ak}(\ln r_{k,t} - \ln r_{k,t-1}) + \eta_{a3}(\ln z_{t+1} - \ln z_t) \tag{64}$$

where $\eta \equiv \{\eta_{a0}, \eta_{a1}, \eta_{a2}, \eta_{a3} : a = 1, 2\}$ is a vector of parameters. Under this assumption the vector of aggregate state variables that individuals use to predict future prices is $\widetilde{X}_t = (z_t, r_{1t}, r_{2t}, r_{1,t-1}, r_{2,t-1})$. Therefore, the equilibria they consider are approximations to the full rational expectations equilibria. This approach is in the spirit of Krusell and Smith (1998). It is important to note that the vector η is determined in equilibrium as a function of the structural parameters, but is not itself one of the structural parameters or primitives of the model.

The vector of structural parameters of the model is $\theta = (\theta_u, \theta_y, \theta_\pi, \theta_Y, \theta_z, \theta_n)$ where: θ_u and θ_y represents the parameters in utility function and wage equations, respectively; θ_π is the distribution of types (by cohort, schooling at age 16 and gender); θ_Y contains the parameters in the aggregate production function; θ_z has the coefficients in the stochastic process of total factor productivity; and θ_n represents the parameters in the stochastic process followed by the number of pre-school children.

A rational expectations equilibrium in this model can be described as a value of the vector η, say $\eta^*(\theta)$, that solves a fixed point problem. The following description of the equilibrium mapping

also provides an algorithm to compute the fixed point. Consider an arbitrary value of η, say η_0.

Step 1 (Optimal individual behavior): Given η_0 individuals use equation (64) to form expectations about future prices, and to solve their occupational choice problems.
Step 2 (Solve for market clearing skill prices): Given initial conditions for TFP, skill prices and the distribution of state variables for all individuals alive at $t = 1$:

a. Simulate a sequence of values of TFP for $t = 1, 2, ..., T$, drawing from the AR(1) process defined by θ_z.

b. Guess skill prices $\{r_{at}\}$ at $t = 1$ using the TFP draw and equation (64) Draw idiosyncratic shocks for all individuals alive at $t = 1$ and simulate their choices using the solutions from step 1. Obtain aggregate skill supplies S_{1t} and S_{1t} for $t = 1$.

c. Given these skill supplies, use the market clearing conditions to obtain a new value of skill prices, that is:

$$r'_{at} = \frac{\alpha_{at}}{S_{at}} \left(z_t \ {S_{1t}}^{\alpha_{1t}} \ {S_{2t}}^{\alpha_{2t}} \ K_t^{1-\alpha_{1t}-\alpha_{2t}} \right) \quad \text{for } a = 1, 2 \text{ for } t = 1$$

d. In general, the new skill prices $\{r'_{at}\}$ will not be the same as the original guess $\{r_{at}\}$. Replace $\{r_{at}\}$ by $\{r'_{at}\}$ and repeat steps b-c until convergence.

e. Repeat steps b-c-d for $t = 2, .., T$. Let $\{r_{at}(\eta_0, \theta) : t = 1, 2, ..., T\}$ be the sequence of skill prices that we obtain upon convergence.
Step 3 (Update Beliefs): Use this new sequence of skill prices and the sequence of TFP to obtain a new value of η, say η_1, as the vector of OLS coefficients for the 'regression equation' in (64).
Step 4 (Impose Self-Fulfilling Beliefs): If $\eta_1 = \eta_0$, then η_0 is a rational expectations equilibrium associated with θ, i.e., individuals' beliefs are self-fulfilling. Otherwise, we start again in Step 1 using η_1 instead of η_0.

The data used in LW have annual frequency and consist of the following items from different sources: (1) occupational choice and wage data (from micro surveys); (2) aggregate output and capital stock; (3) number of pre-school children, by cohort age and gender; (4) cohort sizes, by gender; (5) distribution of schooling at age 16, by cohort and gender; and (6) the initial conditions, i.e., the distribution of state variables x_{i1} for all cohorts alive at calendar time $t = 1$. To make the inference problem more tractable, LW treat data in items (2) to (6) as population parameters which are known to the researcher. The parameter θ_n is obtained directly from (3) and is not estimated with the rest of the parameters. Note that data in items (2), (4), (5) and (6) is used in

the solution algorithm we have just described. For the occupational choice and wage data in item (1), the authors combine two different micro surveys: the Current Population Survey (CPS) and the National Longitudinal Survey of the Youth (NLSY). The two datasets are complementary: on the one hand, the CPS covers a much longer period and thus provides a much wider coverage in terms of calendar time, cohorts and ages; on the other hand, the NLSY has full histories as of age 16, so it adds a true panel dimension and, furthermore, experience capital can be constructed for all sampled individuals.

The estimation method that LW use is a Simulated Method of Moments (SMM). The estimation criterion is a weighted average distance between sample and simulated moments, where the weights are the inverses of the estimated variances of the moments. Moments are selected from CPS and NLSY micro data and a very large number of moments is considered (see pages 23-24 in their paper). The estimation procedure is a nested solution-estimation algorithm. The 'outer algorithm' searches for the value of θ that minimizes the sample criterion function. An iteration of this 'outer algorithm' is a Newton iteration. For each value of θ in this gradient search, the 'inside algorithm' solves for an equilibrium of the model using the procedure that we have described in steps 1 to 4 above. Given that equilibrium, the inside algorithm simulates data and calculates the simulated moments associated with a given θ.[59]

It is helpful to compare this estimation procedure with the one for the single-agent occupational choice model in Example 2. In the single-agent model there was not aggregate uncertainty, structural parameters consisted of $(\theta_u, \theta_\pi, \theta_n)$ and of constant skill prices $\bar{r}_1$ and $\bar{r}_2$, and these parameters were estimated using the likelihood of the micro data in item (1). In LW's approximation to a stochastic rational expectations equilibrium, the state vector of an individual agent's problem is augmented with the 5 continuous aggregate variables $\widetilde{X}_t = (z_t, r_{1t}, r_{2t}, r_{1,t-1}, r_{2,t-1})$ and the parameter vector is augmented with parameters (θ_z, η). However, η is not a free parameter but a function of the structural parameters θ implicitly defined as a fixed point of the equilibrium mapping described above. Therefore, the estimation problem is an order of magnitude more complex than in Example 2.

There are several details of LW's method which are worth mentioning. First, note that at step 2 of the solution algorithm (i.e., imposing market clearing condition) we need initial conditions for TFP and skill prices, $\{r_{a1}, z_1 : a = 1, 2\}$, which are unobservable variables for the researcher. The

[59] Lee and Wolpin do not estimate all the model parameters using this nested procedure. The parameters in the stochastic process of the number of pre-school children, θ_n, are estimated separately in a first step using data item (3).

way that LW deal with this issue is by choosing an initial period $t = 1$ which is many periods before the sample period that is used to construct moments conditions from the micro data. In this way, the choice of the initial conditions has negligible influence on the simulated data for the sample period. This amounts to assuming that aggregate initial conditions are consistent with 'steady state' implications of the model, modified by the limited information we have about aggregate trends in the decades that precede the sample period.

A second detail deals with the solution of the individuals' occupational choice problem. The state space of the DP problem is very large and the estimation procedure requires that the DP problem be solved many times, more than once for each cohort, type and candidate parameter value. LW use Keane-Wolpin's simulation-interpolation method to approximate the solution to individuals' occupational choice problem.

Third, a relevant question is why SMM is used instead of Simulated Maximum Likelihood. The likelihood function of this model is particularly complex and costly to evaluate. For instance, full histories of choices and states are not available for all sampled individuals and, in particular, the value of occupation-specific experience capital is not observable. Missing state variables have to be integrated out for each individual contribution to the likelihood. Computing the likelihood using simulation methods is therefore more costly than computing simulated moments.

Fourth, in order to highlight the equilibrium mapping and the solution of the model we have presented a simplified version of LW's solution/estimation algorithm in which no use is made of the time series of output in data item (2). In LW's solution algorithm the sequence of TFP values in step 2 is not simulated. Instead, they impose that the value of z_t which is conjectured to derive aggregate skill supplies in step 2b should be consistent with aggregate production technology, that is $Y_t = z_t \; S_{1t}^{\alpha_{1t}} \; S_{2t}^{\alpha_{2t}} \; K_t^{1-\alpha_{1t}-\alpha_{2t}}$ where S_{1t} and S_{2t} are the derived supplies and Y_t and K_t are the actual time series data. The iterative procedure within step 2 obtains both the market clearing skill rental prices and the TFP values consistent with the data. In step 3 both η and θ_z are updated based on the sequences derived in step 2. In this solution algorithm θ_z is not a free parameter. An 'estimate' of it is obtained as a by-product, conditional on the rest of the parameter vector and data item (2).

Finally, given that LW's method is computationally very intensive, an important question is whether there is a simpler estimation strategy which does not require one to solve for the equilibrium of the model for each trial value of θ. Consider the same model and assume that the data has been generated by a rational expectations equilibrium of this model. It is possible to obtain

an estimator of an augmented parameter vector (θ, η) that does not fully impose the equilibrium restrictions. The main advantage of this alternative approach, in the spirit of Heckman et al, is that it is computationally simpler: its computational burden is of the same order of magnitude as the one in the estimation of the single-agent occupational choice model. However, it has some potential limitations. First, it may be difficult to identify (θ, η) jointly without fully imposing the equilibrium restrictions. Relatedly, there is a loss of efficiency. Even if (θ, η) is identified, imposing the equilibrium restrictions can improve significantly the precision of our estimates. A third limitation is that the estimated model, though statistically consistent, is not internally consistent and one might argue that this detracts from the credibility of counterfactual exercises.[60] There is a trade-off between the increased efficiency and internal consistency of the estimates obtained using Lee and Wolpin's method and the extra computational burden which is involved.[61]

[60] That is, the estimate vector η can be very different to the value that we get using information from skill prices generated from the market clearing conditions. A possible way of dealing with this internal inconsistency is to include this discrepancy as one of the moment conditions in the sample criterion function to minimize.

[61] A solution algorithm like Lee and Wolpin's is still needed to carry out their empirical analysis of the growth of the service sector which is based on counterfactuals.

Appendix: A guide to the use of the programming codes in the companion web page

1. *Rust's nested fixed point algorithm (NFXP).*

2. *Hotz-Miller's CCP.*

3. *Hotz-Miller's CCP with simulation*

4. *Nested pseudo maximum likelihood (NPL).*

5. *Keane-Wolpin simulation-interpolation method.*

6. *Bajari-Benkard-Levin method (BBL).*

7. *Arcidiacono-Jones EM algorithm*

8. *Benitez-Silva, Hall, Hitsch, and Rust's NFXP with Parameterized Policy Iterations.*

References

[1] Aguirregabiria, V. (1999): "The Dynamics of Markups and Inventories in Retailing Firms," *The Review of Economic Studies*, 66, 275-308.

[2] Aguirregabiria, V. (2007): "Another Look at the Identification of Dynamic Discrete Decision Processes: With an Application to Retirement Behavior," manuscript. University of Toronto.

[3] 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.

[4] Aguirregabiria, V. and P. Mira (2007): "Sequential Estimation of Dynamic Discrete Games," *Econometrica*, 75 (1), 1–54.

[5] Aguirregabiria, V., P. Mira, and H. Roman (2007): "An Estimable Dynamic Model of Entry, Exit and Growth in Oligopoly Retail Markets", *American Economic Review,* 97 (2), 449-454.

[6] Altug, S. and R. Miller (1998): "The Effect of Work Experience on Female Wages and Labour Supply," *Review of Economic Studies*, 65, 45-85.

[7] Albrecht, J. and B. Axell (1984): "An Equilibrium Model of Search Unemployment," *Journal of Political Economy*, 92 (5), 824-840.

[8] Andrews, D. (1988): "Chi-Square Diagnostic Tests for Econometric Models" *Journal of Econometrics*, 37 135–156.

[9] Arcidiacono, P. and J. Jones (2003): "Finite Mixture Distributions, Sequential Likelihood, and the EM Algorithm," *Econometrica*, 71, 933-946.

[10] Arcidiacono, P. and R. Miller (2006): "CCP Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity," manuscript. Carnegie Mellon University.

[11] 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), 37-65.

[12] Bajari, P., L. Benkard and J. Levin (2006): "Estimating dynamic models of imperfect competition," *Econometrica*, forthcoming.

[13] Bajari, P. and H. Hong (2005): "Semiparametric Estimation of a Dynamic Game of Incomplete Information," manuscript. Duke University.

[14] Benitez-Silva, H., G. Hall, G. Hitsch, and J. Rust (2005): "A Comparison of Discrete and Parametric Approximation Methods for Solving Dynamic Programming Problems in Economics," manuscript. Department of Economics. University of Maryland.

[15] Bound, J., T. Stinebrickner and T. Waidman (2005): "Using a Structural Retirement Model to Simulate the Effect of Changes to the Oasdi and Medicare Programs," Economic Research Initiative on the Uninsured Working Paper Series. Number 51.

[16] Chernozhukov, V., Hong, H., and E. Tamer (2007): "Parameter Set Inference in a Class of Econometric Models," *Econometrica*, forthcoming.

[17] Collard-Wexler, A. (2006): "Demand Fluctuations and Plant Turnover in the Ready-to-Mix Concrete Industry," Mimeo. New York University.

[18] Das, M. (1992): "A Micro-econometric Model of Capital Utilization and Retirement: The Case of the Cement Industry," *Review of Economic Studies*, 59, 277-297.

[19] Delavande, A. (2006): "Pill, Patch or Shot? Subjective Expectations and Birth Control Choice," manuscript. Universidade Nova de Lisboa.

[20] Doraszelski, U. & Satterthwaite, M. (2003): "Foundations of Markov-perfect industry dynamics: Existence, purification, and multiplicity," Working paper, Hoover Institution, Stanford.

[21] Eckstein, Z., and G. van den Berg (2007): "Empirical labor search: a survey", *Journal of Econometrics*, 136(2), 531-564.

[22] Eckstein, Z., and K. Wolpin (1989): "The specification and estimation of dynamic stochastic discrete choice models," *Journal of Human Resources,* 24, 562-598.

[23] Eckstein, Z. and K. Wolpin (1990): "Estimating a Market Equilibrium Search Model from Panel Data on Individuals," *Econometrica*, 58(4), 783-808.

[24] Eckstein, Z. and K. Wolpin (1999): "Why Youth Drop out of High School: The Impact of Preferences, Opportunities and Abilities," *Econometrica,* 67, 1295-1339.

[25] Erdem, T., S. Imai and M. P. Keane (2003): "Brand and Quantity Choice Dynamics under Price Uncertainty," *Quantitative Marketing and Economics*, 1, 5-64.

[26] Ericson, R., and A. Pakes (1995): "Markov-Perfect industry dynamics: A framework for empirical work", *Review of Economic Studies,* 62, 53-82.

[27] Fernández-Villaverde, J., J. Rubio-Ramírez, and M. Santos (2006): "Convergence Properties of the Likelihood of Computed Dynamic Models," *Econometrica*, 74 (1), 93-119.

[28] Ferrall, C. (2005): "Solving Finite Mixture Models: Efficient Computation in Economics under Serial and Parallel Execution," *Computational Economics*, 25(4), 343-379.

[29] Geweke, J. (1996): "Monte Carlo simulation and numerical integration," in H. Amman, D. Kendrick, and J. Rust (eds.) *Handbook of Computational Economics*, chapter 15, pages 731-800. North-Holland. Amsterdam.

[30] Geweke, J. and M. Keane (2000). "Bayesian Inference for Dynamic Discrete Choice Models without the Need for Dynamic Programming." in book *Simulation Based Inference and Econometrics: Methods and Applications*, Mariano, Schuermann and Weeks (eds.), Cambridge University Press, 100-131.

[31] Gilleskie, D. (1998): "A Dynamic Stochastic Model of Medical Care Use and Work Absence," *Econometrica,* 66, 1-45.

[32] Hajivassiliou, V. and P. Rudd (1994): "Classical Estimation Methods for LDV Models Using Simulation," in *The Handbook of Econometrics*, Volume 4, D. McFadden and R. Engle (eds.). North-Holland: Amsterdam.

[33] Hall, G. and J. Rust (2005): "Simulated Minimum Distance Estimation of a Model of Optimal Commodity Price Speculation with Endogenously Sampled Prices," manuscript. Department of Economics. University of Maryland.

[34] Heckman, J. (1981): "The incidental parameters problem and the problem of initial conditions in estimating a discrete time - discrete data stochastic process," in C. Manski and D. McFadden (eds.), *Structural Analysis of Discrete Data with Econometric Applications*. MIT Press.

[35] Heckman, J., L. Lochner, and C. Taber (1998): "Explaining Rising Wage Inequality: Explanations With A Dynamic General Equilibrium Model of Labor Earnings With Heterogeneous Agents," *Review of Economic Dynamics*, 1(1), 1-58.

[36] Heckman, J., and B. Singer (1984): "A method for minimizing the impact of distributional assumptions in economic models for duration data," *Econometrica,* 52, 271-320.

[37] Heckman, J. and S. Navarro (2007): "Dynamic Discrete Choice and Dynamic Treatment Effects," *Journal of Econometrics,* 136(2), 341-396.

[38] Hendel, I., and A. Nevo (2006): "Measuring the Implications of Sales and Consumer Inventory Behavior," *Econometrica,* 74(6), 1637-1674.

[39] Hollifield, B. R. Miller, and P. Sandas (2004): "Empirical Analysis of Limit Order Markets," *Review of Economic Studies,* 71, 1027-1063.

[40] Hotz, J., and R.A. Miller (1993): "Conditional choice probabilities and the estimation of dynamic models," *Review of Economic Studies,* 60, 497-529.

[41] Hotz, J., R.A. Miller, S. Sanders, and J. Smith (1994): "A simulation estimator for dynamic models of discrete choice," *Review of Economic Studies,* 61, 265-89.

[42] Imai, S., N. Jain, and A. Ching (2007): "Bayesian Estimation of Dynamic Discrete Choice Models," Mimeo. University of Toronto.

[43] Jofre-Bonet, M. and M. Pesendorfer (2003): "Estimation of a Dynamic Auction Game," *Econometrica*, 71, 1443–1489.

[44] Kahneman, D. and A. Tversky (1979): "Prospect Theory: An Analysis of Decision under Risk," *Econometrica*, 47, 263-292.

[45] Kano, K. (2006): "Menu Costs, Strategic Interactions and Retail Price Movements," manuscript. Queen's University.

[46] Karlstrom, A., M. Palme and I. Svensson (2004): "A dynamic programming approach to model the retirement behaviour of blue-collar workers in Sweden," *Journal of Applied Econometrics*, 19, 795-807.

[47] Kasahara, H. and K. Shimotsu (2006a): "Nested Pseudo-likelihood Estimation and Bootstrap-based Inference for Structural Discrete Markov Decision Models," manuscript. Department of Economics. The University of Western Ontario.

[48] Kasahara, H. and K. Shimotsu (2006b): "Nonparametric Identification and Estimation of Finite Mixture Models of Dynamic Discrete Choices," manuscript. Department of Economics. The University of Western Ontario.

[49] Keane, Michael P and Wolpin, Kenneth I (1994): "The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence," The Review of Economics and Statistics, 76(4), 648-72.

[50] Keane, M. and K. Wolpin (1997): "The career decisions of young men," Journal of Political Economy, 105, 473-522.

[51] Keane, M. and K. Wolpin (2000): "Estimating Welfare Effects Consistent with Forward-Looking Behavior. Part I: Lessons from a Simulation Exercise," PIER Working Paper 01-019. University of Pennsylvania.

[52] Keane, M. and K. Wolpin (2000): "Estimating Welfare Effects Consistent with Forward-Looking Behavior. Part II: Empirical Results," PIER Working Paper 01-020. University of Pennsylvania.

[53] Keane, M. and K. Wolpin (2001): "The Effect of Parental Transfers and Borrowing Constraints on Educational Attainment," *International Economic Review*, 42, 1051-1103.

[54] Keane, M. and K. Wolpin (2007): "Exploring the Usefulness of a Non-Random Holdout Sample for Model Validation: Welfare Effects on Female Behavior," *International Economic Review*, forthcoming.

[55] Kennet, M. (1993) "Did Deregulation Affect Aircraft Engine Maintenance? An Empirical Policy Analysis," *The RAND Journal of Economics,* 24, 542-558.

[56] Kennet, M. (1994): "A Structural Model of Aircraft Engine Maintenance," *Journal of Applied Econometrics*, 9, 351-368.

[57] Krusell, P. and A. Smith (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106(5), 867-896.

[58] Lancaster, A. (1997): "Exact Structural Inference in Optimal Job-Search Models," *Journal of Business and Economic Statistics*, 15(2), 165-179.

[59] Lee, Donghoon (2005) "An estimable dynamic general equilibrium model of work, schooling and occupational choice". *International Economic Review*, forthcoming.

[60] Lee, D. and K. Wolpin (2006): "Intersectoral labor mobility and the growth of the service sector". *Econometrica*, 74 (1), 1-46.

[61] Lerman, S. and C. Manski (1981): "On the Use of Simulated Frequencies to Approximate Choice Probabilities," in C. Manski and D. McFadden (eds.), *Structural Analysis of Discrete Data with Econometric Applications*, MIT Press, Cambridge, MA.

[62] Lorincz, S. (2005): "Persistence Effects in a Dynamic Discrete Choice Model: Application to Low-End Computer Servers." Discussion Papers 2005/10. Institute of Economics Hungarian Academy of Sciences

[63] Lumsdaine, R., J. Stock, and D. Wise (1992): "Three Models of Retirement: Computational Complexity versus Predictive Validity," in David A. Wise (ed.), *Topics in the economics of aging*, pp. 16-60. University of Chicago Press. Chicago.

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

[65] Manski, C. (1993): "Dynamic choice in social settings: Learning from the experience of others", *Journal of Econometrics* 58, pp. 121-36.

[66] Manski, C. (2004): "Measuring Expectations," *Econometrica*, 72 (5), 1329-1376.

[67] Miller, R. (1984): "Job Matching and Occupational Choice", Journal of Political Economy, 92, 1086-1120.

[68] Miller, R. (1997): "Estimating Models of Dynamic Optimization with Microeconomic Data," in H. Pesaran and P. Smidth (eds.) *Handbook of Applied Econometrics: Microeconomics (Vol 2)*. Blackwell.

[69] Miller, R., and S. Sanders (1997): "Human Capital Development and Welfare Participation", Carnegie-Rochester Conference Series on Public Policy, 46, 1-44.

[70] Mortensen, D. and C. Pissarides: "New developments in models of search in the labor market", Handbook of Labor Economics vol. 2, Chap. 39, 2567-2627.

[71] Pakes, A. (1986): "Patents as Options: Some Estimates of the Value of Holding European Patent Stocks," *Econometrica,* 54, 755-84.

[72] Pakes, A. (1994): "Dynamic structural models, problems and prospects", in C. Sims (ed.) *Advances in Econometrics. Sixth World Congress*, Cambridge University Press.

[73] Pakes, A., M. Ostrovsky, and S. Berry (2004): "Simple estimators for the parameters of discrete dynamic games (with entry/exit examples)," manuscript, Harvard University.

[74] Pesendorfer, M. and Schmidt-Dengler (2007): "Asymptotic Least Squares Estimators for Dynamic Games," forthcoming in *The Review of Economic Studies.*

[75] Rota, P. (2004): "Estimating Labor Demand with Fixed Costs," *International Economic Review*, 45, 25-48.

[76] Rios-Rull, J. V. (1999): "Computation of Equilibria in Heterogeneous Agent Models," in *Computational Methods for the Study of Dynamic Economies: An Introduction,* R. Marimon and S.Scott (eds.). Oxford University Press. Oxford.

[77] Rust, J. (1987): "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher," *Econometrica,* 55, 999-1033.

[78] Rust, J. (1994): "Estimation of dynamic structural models, problems and prospects: discrete decision processes", in C. Sims (ed.) *Advances in Econometrics. Sixth World Congress*, Cambridge University Press.

[79] Rust, J. (1994): "Structural estimation of Markov decision processes," in R. E. Engle and McFadden (eds.) *Handbook of Econometrics Volume 4*, North-Holland. Amsterdam.

[80] Rust, J. and G. Rothwell (1995): "Optimal Response to a Shift in Regulatory Regime: The Case of the US Nuclear Power Industry," *Journal of Applied Econometrics*, 10, S75-S118.

[81] Rust, J. and C. Phelan (1997): "How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets," *Econometrica*, 65, 781-832.

[82] Ryan, S. (2004): "The Costs of Environmental Regulation in a Concentrated Industry," manuscript, MIT Department of Economics.

[83] Sanchez-Mangas, R. (2002): "Pseudo Maximum Likelihood Estimation of a Dynamic Structural Investment Model," Working Paper 02-62, Statistics and Econometrics Series. Departamento de Estadística y Econometría. Universidad Carlos III de Madrid.

[84] Slade, M. (1998): "Optimal Pricing with Costly Adjustment: Evidence from Retail Grocery Stores", *Review of Economic Studies*, 65, 87-108.

[85] Stern, S. (1997): "Simulation-based estimation," *Journal of Economic Literature,* 35, 2006–2039.

[86] Stinebrickner, T. (2000): "Serially Correlated Variables in Dynamic Discrete Choice Models," *Journal of Applied Econometrics*, 15 (6), 595-624.

[87] Sturm, R. (1991) "A Structural Economic Model of Operating Cycle Management in European Nuclear Power Plants," manuscript, RAND Corporation.

[88] Todd, P. and K. Wolpin (2006): "Assessing the Impact of a School Subsidy Program in Mexico." *American Economic Review*, 96 (5), 1384–1417.

[89] Van der Klaauw, W. and K. Wolpin (2005): "Social Security and the Retirement and Savings Behavior of Low Income Households," PIER Working Paper 05-020. Penn Institute for Economic Research. Department of Economics. University of Pennsylvania.

[90] Wolpin, K. (1984) "An Estimable Dynamic Stochastic Model of Fertility and Child Mortality," *Journal of Political Economy*, 92, 852-874.

[91] Wolpin, K. (1987) "Estimating a Structural Search Model: The Transition from Schooling to Work," *Econometrica*, 55, 801-818.

[92] Wolpin, K. (1996): "Public-Policy Uses of Discrete-Choice Dynamic Programming Models," *American Economic Review*, 86, 427-32.

[93] Wolpin, K. (2007): "Model Validation and Model Comparison Ex Ante Policy Evaluation, Structural Estimation, and Model Selection," *American Economic Review,* 97 (2), 48–52.