

CCP and the Estimation of Nonseparable Dynamic Models*

Dennis Kristensen [†] Lars Nesheim [‡]
UCL, CeMMAP, CREATES and IFS UCL, CeMMAP and IFS

Áureo de Paula [§]
UCL, São Paulo School of Economics,
CeMMAP and IFS

Preliminary and Incomplete.

First Version: November 2014. This Version: November 2015

Abstract

In this paper we generalize the so-called CCP estimator of Hotz and Miller (1993) to a broader class of dynamic discrete choice (DDC) models that allow period payoff functions to be non-separable in observable and unobservable (to the econometrician) variables. Such non-separabilities are common in applied microeconomic environments and our generalized CCP estimator allows for computationally simple estimation in this class of DDC models. We first establish invertibility results between conditional choice probabilities (CCPs) and value functions and use this to derive a policy iteration mapping in our more general framework. This is used to develop a pseudo-maximum likelihood estimator of model parameters that side-step the need for solving the model. To make the inversion operational, we use Mathematical Programming with Equilibrium Constraints (MPEC), following Judd and Su (2012) and demonstrate its

*We thank Victor Aguirregabiria, Steve Berry, Andrew Chesher, Phil Haile, Xavier d'Haultfoeuille, Jim Heckman, Bo Honoré, Hiro Kasahara, John Kennan, Pedro Mira, Jean-Marc Robin, Petra Todd and Ken Wolpin together with participants at the Econometric Society Summer Meetings 2014 in Toulouse, and various seminars for helpful comments and suggestions. This research was supported by the Economic and Social Research Council through the ESRC Centre for Microdata Methods and Practice grant RES-589-28-0001. Kristensen acknowledges support from the Danish National Research Foundation (through a grant to CREATES) and the European Research Council (grant no. ERC-2012-StG 312474). de Paula acknowledges support from the European Research Council (grant no. ERC-2013-StG 338187).

[†]University College London, London, UK, CeMMAP, London, UK, and IFS, London, UK. E-mail: d.kristensen@ucl.ac.uk

[‡]University College London, London, UK, CeMMAP, London, UK, and IFS, London, UK. E-mail: l.nesheim@ucl.ac.uk

[§]University College London, London, UK, São Paulo School of Economics, São Paulo, Brazil, CeMMAP, London, UK, and IFS, London, UK. E-mail: apaula@ucl.ac.uk

applicability in a series of Monte Carlo experiments that replicate the model used in Keane and Wolpin (1997).

1 Introduction

Full maximum-likelihood estimation of dynamic discrete choice (DDC) models is a computationally demanding task due to the need for solving the model for any given candidate value of the unknown parameters. This approach leads to the so-called nested-fixed point (NFXP) algorithm as described in Rust (1988). To reduce the computational cost Hotz and Miller (1993) developed the so-called conditional choice probability (CCP) estimator where the usual mapping taking value functions into CCP's used in NFXP is replaced by a policy iteration mapping that maps CCP's into CCP's. This then allows the econometrician to estimate model parameters from an initial nonparametric estimator of the CCP's without the need of solving the model repeatedly. The CCP estimator was further refined and a full asymptotic analysis of it developed by, amongst others, Aguirregabiria and Mira (2002), Aguirregabiria (2004), Kasahara and Shimotsu (2008), and Kasahara and Shimotsu (2012).

The Hotz-Miller CCP estimator was developed under the assumption that the unobserved state variables enter additively in the per-period utility function, which rules out utility specifications employed in, for example, labour. We here extend their estimation method to allow for the unobserved state variable to enter non-additively and to be multivariate. This extension allows econometricians to employ the CCP estimator in more general models so broadens the scope for computationally simple estimation of DDC models.

As a first step, we show that the arguments Hotz and Miller (1993) and Aguirregabiria and Mira (2002) can be extended to the above mentioned broader class of models. This involves deriving a policy iteration mapping in our more general setting, and analyzing its properties. The invertibility result should be of independent interest since it could be used when verifying identification of the model as done in Magnac and Thesmar (2002), who work in the same setting as Hotz and Miller (1993). We then proceed to show that the asymptotic arguments of Aguirregabiria (2004), Kasahara and Shimotsu (2008), and Kasahara and Shimotsu (2012) continue to hold in our setting. In particular, we develop an asymptotic theory for our generalized CCP estimator, showing that the estimator is first-order equivalent to the NFXP estimator and that higher-order improvements are obtained through iterations.

The policy iteration mapping for the generalized CCP estimator is in general not available on closed form since it involves integrals that cannot be computed analytically. However, the integrals are low-dimensional and, using state-of-the-art numerical integration techniques, can be computed fast and accurately. Through various numerical exercises (**TBC**) we (hope to) show that the CCP estimator remains computationally attractive in our more general set-up, while still yielding precise estimates that are comparable to the NFXP estimator.

We also discuss some extensions, including random-coefficient models and measurement errors, that can be accommodated for in our framework. The remainder of the paper is organized as follows: In the next section, we present the class of DDC models for which our generalized CCP estimator can be applied to. In Section 3, we derive a policy iteration mapping in our setup and

analyze its properties. The policy iteration mapping is then used to develop the generalized CCP estimator in Section 5. Section 6 discusses various aspects of the numerical implementation of the estimator, while Section 7 contains the results of our numerical experiments. In Section 8, we conclude by discussing some extensions.

2 Framework

Consider a succession of periods $t = 0, 1, \dots, T$, with $T \leq \infty$ where an individual i chooses among $J + 1$ alternatives $d_t \in \mathcal{D} = \{1, \dots, J + 1\}$. In each period, alternative d gives per-period utility

$$u_d(x_t, \epsilon_t)$$

where ϵ_t is unobserved by the econometrician whereas x_t is observed. We collect the state variables in $z_t = (x_t, \epsilon_t)$ which is a controlled Markov process with transition dynamics described by $F_{z_{t+1}|z_t, d_t}(z_{t+1}|z_t, d_t)$. An Individual chooses d_t that maximises lifetime utility given the current state,

$$d_t = \arg \max_{d \in \mathcal{D}} E \left[\sum_{i=0}^T \beta^i u_d(z_{t+i}) \middle| z_t \right].$$

When $T = \infty$, the corresponding value function, $W(z_t)$, solves the Bellman equation taking the form

$$W(z_t) = \max_{d \in \mathcal{D}} \left\{ u_d(z_t) + \beta \int W(z_{t+1}) dF_{z_{t+1}|z_t, d_t}(z_{t+1}|z_t, d_t) \right\}.$$

Conditions for the existence of a solution to the above problem are given, for example, in Bhattacharya and Majumdar (1989). For $T < \infty$, value functions are accordingly indexed by time period and the problem can be solved backwards (see, e.g., Rust (1994)). For ease of notation, we focus our exposition on the infinite horizon case. Many of the ideas can nonetheless be appropriately transported to the finite horizon case. We follow the literature in assuming Conditional Independence (CI):

Assumption CI. (a) The state variables are Markovian and evolve according to

$$F_{z_{t+1}|z_t, d_t}(z_{t+1}|z_t, d_t) = F_{\epsilon_{t+1}|x_{t+1}}(\epsilon_{t+1}|x_{t+1})F_{x_{t+1}|x_t, d_t}(x_{t+1}|x_t, d_t). \quad (\text{CI})$$

Assumption CI restricts the dependence in the (x_t, ϵ_t) process. First, x_{t+1} is a sufficient statistic for ϵ_{t+1} implying that any serial dependence between ϵ_t , and ϵ_{t+1} is transmitted entirely through x_{t+1} . Second, the distribution of x_{t+1} depends only on x_t and not on ϵ_t . In applied work, it is typical to assume that ϵ_t is iid and independent of everything else. This assumption allows us to express the solution to the DDC model in terms of the so-called integrated (or smoothed) value function $V(x_t)$ defined by

$$V(x_t) = \int_{\mathcal{E}} W(x_t, \epsilon_t) dF_{\epsilon_t|x_t}(\epsilon_t|x_t).$$

More specifically, the Bellman equation can be rewritten as

$$V(x_t) = \Gamma(V)(x_t), \quad (1)$$

where Γ is the smoothed Bellman operator defined as

$$\Gamma(V)(x_t) = \int_{\mathcal{E}} \max_{d_t \in \mathcal{D}} \left\{ u_d(x_t, \epsilon_t) + \beta \int_{\mathcal{X}} V(x_{t+1}) dF_{x_{t+1}|x_t, d}(x_{t+1}|x_t, d_t) \right\} dF_{\epsilon_t|x_t}(\epsilon_t|x_t). \quad (2)$$

We can then write the decision problem as

$$d_t = \arg \max_{d \in \mathcal{D}} [u_d(x_t, \epsilon_t) - u_{J+1}(x_t, \epsilon_t) + \beta \Delta_d(x_t)],$$

where

$$\Delta_d(x_t) = \int_{\mathcal{X}} V(x_{t+1}) dF_{x_{t+1}|x_t, d}(x_{t+1}|x_t, d) - \int_{\mathcal{X}} V(x_{t+1}) dF_{x_{t+1}|x_t, d}(x_{t+1}|x_t, J+1) \quad (3)$$

is the choice-specific relative expected value function. Notice here that $\Delta_{J+1}(x_t) = 0$. The optimal decision rule d_t induces CCP's defined as

$$\mathbb{P}_d(x) := P(d_t = d|x_t = x) = Q_d(\Delta|x),$$

where $Q(\Delta|x) = [Q_d(\Delta|x)]_{d=1}^J$ is the CCP mapping defined as

$$Q_d(\Delta|x) \equiv P(u_d(x, \epsilon) + \beta \Delta_d(x) \geq u_j(x, \epsilon) + \beta \Delta_j(x), j \in \mathcal{D}). \quad (4)$$

Consider now a parametric specification of the utility function and the state dynamics, $u_d(z_t; \theta)$ and $F_{z_{t+1}|z_t, d_t}(z_{t+1}|z_t, d_t; \theta)$, where θ contains the unknown parameters. Given observations (d_{it}, x_{it}) , $i = 1, \dots, n$ and $t = 1, \dots, T$, of n individuals across T periods, the NFXP estimator of θ is then defined as the maximizer of the likelihood function that takes the form

$$\mathcal{L}_{nT}(\theta) = \prod_{i=1}^n \prod_{t=1}^T Q_{d_{it}}(\Delta_{\theta}|x_{it}; \theta) f_{x_t|x_{t-1}, d_{t-1}}(x_{it}|x_{it-1}, d_{it-1}; \theta).$$

Here, $f_{x_t|x_{t-1}, d_{t-1}}$ is the density of $F_{x_t|x_{t-1}, d_{t-1}}$ w.r.t. some appropriate measure, while Δ_{θ} is given by eq. (3) with $V = V_{\theta}$ being the fixed point of $V_{\theta} = \Gamma_{\theta}(V_{\theta})$. As we search for the maximizer of the likelihood, the MLE, we need to solve the programme in terms of V_{θ} repeatedly in order to compute $Q_d(\Delta_{\theta}|x_{it}; \theta)$, leading to a so-called nested fixed algorithm as outlined in Rust (1988). This can be computationally very costly, and often only an approximate version of V_{θ} is available. Furthermore, for a given (approximate) solution, one in general also needs to approximate $Q_d(\Delta_{\theta}|x_{it}; \theta)$ by simulation (e.g., Keane and Wolpin (1997)). Using these estimates we can obtain an estimate for the likelihood function which, once maximised, will give a Simulated Maximum Likelihood estimator (or SMM as in Pakes (1986)). In this case, the simulated version of $Q_d(\Delta_{\theta}|x_{it}; \theta)$ will typically be non-smooth so a non-gradient based optimisation method (Nelder-Meade, simulated annealing, genetic algorithms, etc.) needs to be employed or the objective function needs to be smoothed out (somehow) (Kristensen and Schjerning (2014)).

It is nevertheless possible to reduce the computational burden above using an insight from Hotz and Miller (1993) who considered models with additively separable utility and choice-specific scalar unobservables,

$$u_d(z_t) = u_d(x_t) + \epsilon_{dt} \text{ with } \epsilon_t = (\epsilon_{1t}, \dots, \epsilon_{J+1t})'. \quad (\text{AS})$$

This assumption (together with Extreme Value Type 1 residuals) is also employed by Rust (1987), to obtain the likelihood function numerically. Hotz and Miller (1993) show that, with two alternatives ($J = 1$) for simplicity¹, that the CCP mapping can be expressed as

$$Q_1(\Delta_\theta | x_{it}) = F_{\epsilon_{2t} - \epsilon_{1t} | x_t}(u_1(x_t) - u_2(x_t) - \beta \Delta_1(x_t) | x_t),$$

which is invertible such that

$$\Delta_1(x_t) = \left(F_{\epsilon_{2t} - \epsilon_{1t} | x_t}^{-1}(\mathbb{P}_1(x_t) + u_2(x_t) - u_1(x_t)) \right) / \beta.$$

Under Extreme Value Type 1 distributed residuals, for example,

$$\Delta_1(x_t) = \ln \frac{\mathbb{P}_1(x_t) + u_2(x_t) - u_1(x_t)}{1 - \mathbb{P}_1(x_t) - u_2(x_t) + u_1(x_t)} / \beta$$

The existence and computation of the inversion is not so transparent with more than two alternatives and general distributions for the residuals but can be done. The invertibility is established in Proposition 1 of Hotz and Miller (1993). This result was then used to derive an estimator of θ that bypassed the need for solving the model for each putative value of θ required for the implementation of the NFXP estimator.

Assumption AS, though adequate in a few settings and convenient for some of the estimation strategies, is not always desirable or employed (e.g., Todd and Wolpin (2006)). In this case, the usual strategy so far has been to implement the NFXP estimator described above. The goal of this paper is to generalize the results of Hotz and Miller (1993) to the above more general framework where errors are allowed to enter the utility function non-additively and the choice-specific errors may not be scalar: In the next two sections, we first establish that under certain conditions on the utility function, $Q(\Delta | x)$ defines an invertible mapping between the CCP's and the relative value functions, and that V_θ can be expressed as a functional of the CCP's. We then show how these two mappings allow us to bypass solving the model in order to learn about the parameters from data, and so lead to a computationally simpler and faster estimator of θ .

3 Policy Iteration Mapping

In this section, we show that there exists a mapping $\Psi(\mathbb{P})$ for which the CCP's is a unique fixed point, $\mathbb{P} = \Psi(\mathbb{P})$. As a first step in deriving this mapping and establishing the fixed-point result, we analyze the properties of Q . For notational simplicity, we suppress any dependence on the model parameters θ in the following.

¹With more alternatives, this expression will involve multiple integrals over the space of unobservables.

3.1 Invertibility of CCP Mapping

We here show that the CCP mapping $Q(\Delta|d)$ is invertible w.r.t. the vector of relative expected value function, $\Delta(\mathbf{x}_t) = (\Delta_1(\mathbf{x}_t), \dots, \Delta_J(\mathbf{x}_t))$, defined in eq. (3). To establish the invertibility result, we adopt the following Random Utility (RU) Assumption:

Assumption RU. The conditional distribution of the random utility functions

$$U_t := (u_1(x_t, \epsilon_t), \dots, u_{J+1}(x_t, \epsilon_t))$$

given $x_t = x$, has a continuous distribution $F_{U_t|x_t}(u|x)$ with rectangular support on \mathbb{R}^{J+1} .

Assumption RU is implied by the following more primitive conditions that covers the case of a scalar alternative-specific unobservable:

Lemma 1 *Assumption RU holds under the following conditions: (i) $\epsilon_t = (\epsilon_{1t}, \dots, \epsilon_{J+1t})$ and $u_d(x_t, \epsilon_t) = u_d(x_t, \epsilon_{dt})$; (ii) $u_d(x, \cdot)$ is strictly increasing for every x ; (iii) $F_{\epsilon_t|x_t}(\epsilon_t|x_t)$ is absolutely continuous with rectangular support.*

Proof. Under (i)-(ii), the mapping $\epsilon_t \mapsto u(x_t, \epsilon_t) = (u_1(x_t, \epsilon_{1t}), \dots, u_{J+1}(x_t, \epsilon_{J+1t}))$ is one-to-one. Furthermore, (ii) implies that $u_d(x, \cdot)$ is differentiable (Lebesgue-)almost everywhere (though not necessarily everywhere). We can therefore employ the theorem regarding differentiable transformations of continuous random variables to obtain that $U_t|x_t$ has a density given by

$$f_{U_t|x_t}(u|x) = f_{\epsilon_t|x_t}(u^{-1}(x, u)|x) \left| \frac{\partial u^{-1}(x, u)}{\partial u} \right|,$$

where $u^{-1}(x, u) = (u_1^{-1}(x, u_1), \dots, u_{J+1}^{-1}(x, u_{J+1}))$ is the inverse of $u(x, \epsilon)$ w.r.t. ϵ . Finally, given that $\epsilon_t|x_t$ has rectangular support and $\epsilon_{dt} \mapsto u_d(x_t, \epsilon_{dt})$ is strictly increasing, then $U_t|x_t$ must also have rectangular support. ■

Another specification that falls within the framework of Assumption RU is random coefficient-type models. Consider, for example, the following utility specification:

$$u_d(x_t, \epsilon_t) = \theta_t' x_t + \eta_{dt}, \tag{5}$$

where $\epsilon_t = (\theta_t, \eta_t)$ with $\eta_t = (\eta_{1t}, \dots, \eta_{J+1t})$ and θ_t being random coefficients that vary across individuals.

Lemma 2 *Assumption RU holds for the random coefficient specification in eq. (5) if $\{\theta_t\} \perp \eta_t|x_t$ with $\theta_t|x_t$ and $\eta_t|x_t$ being continuously distributed with rectangular supports.*

Proof. Under (i)-(ii), $U_t|x_t$ has a continuous distribution with density

$$f_{U_t|x_t}(u|x) = \int_{\Theta} f_{\eta_t|x_t}(u - \theta'x|x) f_{\theta_t|x}(\theta|x) d\theta.$$

Since $\theta_t|x_t$ and $\eta_t|x_t$ has rectangular support and vary independently of each other, $U_{d,t} = \theta'x_t + \eta_t$ also has rectangular support. ■

We now derive an expression for the CCP mapping $Q(\Delta|x)$ under Assumption RU. To this end, observe that alternative d is chosen whenever

$$u_d(x_t, \epsilon_t) - u_j(x_t, \epsilon_t) + \beta(\Delta_d(x_t) - \Delta_j(x_t)) \geq 0 \Leftrightarrow u_d(x_t, \epsilon_t) + \beta(\Delta_d(x_t) - \Delta_j(x_t)) \geq u_j(x_t, \epsilon_t)$$

for all $j \in \mathcal{D}$. Thus, $Q_d(\Delta|x)$, $\Delta \in \mathbb{R}^J$, as defined in eq. (4), can be expressed as

$$\begin{aligned} Q_d(\Delta|x) &= \int_{u_d=-\infty}^{\infty} \int_{u_1=-\infty}^{u_d+\beta(\Delta_d-\Delta_1)} \dots \int_{u_{J+1}=-\infty}^{u_d+\beta\Delta_d} f_{U_t|x_t}(u|x) du_1 \dots du_{J+1} \\ &= \int_{-\infty}^{\infty} \partial_d F_{U_t|x_t}(u_d + \beta(\Delta_d - \Delta_1), \dots, u_d, \dots, u_d + \beta\Delta_d|x) du_d, \end{aligned} \quad (6)$$

with ∂_d denoting the partial derivative with respect to the d -th argument, u_d ; this derivative is well-defined due to the assumption of $F_{U_t|x_t}$ being absolutely continuous. It is noteworthy that Q_d depends on β, u and x .

To obtain the inversion result that follows, we also require that model be saturated (in the sense of Rust (1994), Section 3.2). This essentially requires CCPs not to be zero and is encoded in the following Saturation (S) Assumption:

Assumption S. $Q_d(\Delta|x) \in (0, 1)$ for any d and x .

This assumption rules out choices that are always dominated regardless of the realization of payoffs and it is important for CCPs to reveal information about relative value functions Δ . Under Additive Separability (AS), Norets and Takahashi (2013) demonstrate that this holds whenever the distribution of unobservables is absolutely continuous with respect to the Lebesgue measure on \mathbb{R}^{J+1} and has full support. They also show that Assumption S fails when the support is bounded.² In general, since choices are determined by comparing alternatives, the relevant support is that of the differences in unobservables. In our more general case, this boils down to requiring that

$$\beta\Delta \in \Delta\mathcal{U}(x) := \text{supp} \{(u_1(x_t, \epsilon_t) - u_{J+1}(x_t, \epsilon_t), \dots, u_J(x_t, \epsilon_t) - u_{J+1}(x_t, \epsilon_t)) | x_t = x\}.$$

This holds, for example, when $\text{supp} \{(u_1(x_t, \epsilon_t), \dots, u_{J+1}(x_t, \epsilon_t)) | x_t = x\} = [0, \infty)^{J+1}$ as in a dynamic Roy model where payoffs are given by wages in each of the $J + 1$ sectors. We are now in shape to establish the following result:

Theorem 3 *Under Assumptions CI, RU and S, for any $x \in \mathcal{X}$, $\Delta \mapsto Q(\Delta|x) \equiv [Q_j(\Delta|x)]_{j=1}^J$ is continuously differentiable and invertible.*

²One can try to emulate their results in our context by identifying the unobservables there with our random period payoffs and their utility indices with our relative value functions. The application of their results is nevertheless not immediate to us, as our random utilities and relative value functions are structurally related.

Proof. For any $j \in \mathcal{D}$, $Q_j(\Delta|x)$ is the integral of the function $\partial_j F_{U_t|x_t}(u_j + \beta(\Delta_j - \Delta_1), \dots, u_j, \dots, u_j + \beta\Delta_j|x)$ with respect to u_j . Since this function is positive and integrable, the Dominated Convergence Theorem allows us to express the derivative of Q_j with respect to Δ_k as the integral of the derivative of that function with respect to Δ_k . Moreover, because $Q_j(\Delta|x) \in [0, 1]$, its derivative with respect to Δ_k is finite almost-everywhere (or else, Q_j would be infinite), and the Jacobian of $Q(\Delta|x)$ (with respect to Δ) is given by $D(\Delta|x) = [D_{jk}(\Delta|x)]_{j,k=1}^J$ where

$$D_{jj}(\Delta|x) = \sum_{k \neq j} h_{jk}(\Delta|x), \quad D_{jk}(\Delta|x) = -h_{jk}(\Delta|x) \text{ for } j \neq k,$$

with

$$h_{jk}(\Delta|x) := \beta \int_{u_j = -\infty}^{\infty} \partial_{jk}^2 F_{U_t|x_t}(u_j + \beta(\Delta_j - \Delta_1), \dots, u_j, \dots, u_j + \beta\Delta_j|x) du_j$$

for $k \neq j$, and

$$h_{jj}(\Delta|x) := \int_{u_j = -\infty}^{\infty} \partial_{jj}^2 F_{U_t|x_t}(u_j + \beta(\Delta_j - \Delta_1), \dots, u_j, \dots, u_j + \beta\Delta_j|x) du_j.$$

Notice that since $Q_j(\Delta|x)$ is finite, then $h_{j1}(\Delta|x), \dots, h_{jJ}(\Delta|x)$ must be finite almost-everywhere. Furthermore, Assumption S implies that $h_{j1}(\Delta|x), \dots, h_{jJ}(\Delta|x)$ are also strictly positive. To see this, note that for $j \in \mathcal{D}$,

$$\begin{aligned} & F_{U_t|x_t}(u_j + \beta(\Delta_j - \Delta_1), \dots, u_j + \beta(\Delta_j - \Delta_J), u_j + \beta\Delta_j|x) \\ &= \int_{-\infty}^{u_j + \beta(\Delta_j - \Delta_1)} \dots \int_{-\infty}^{u_j + \beta(\Delta_j - \Delta_J)} \int_{-\infty}^{u_j + \beta\Delta_j} f_{U_t|x_t}(v_1, \dots, v_{J+1}) dv, \end{aligned}$$

which in turn gives:

$$\begin{aligned} & \partial_{jk}^2 F_{U_t|x_t}(u_j + \beta(\Delta_j - \Delta_1), \dots, u_j + \beta(\Delta_j - \Delta_J), u_j + \beta\Delta_j|x) \\ &= \int_{U_{jk}(\Delta)} f_{U_t|x_t}(v_1, \dots, v_{j-1}, u_j, v_{j+1}, \dots, v_{k-1}, u_j + \beta(\Delta_j - \Delta_k), v_{k+1}, \dots, v_{J+1}) d\tilde{v}, \end{aligned}$$

where $\tilde{v} := (v_1, \dots, v_{j-1}, v_{j+1}, \dots, v_{k-1}, v_{k+1}, \dots, v_{J+1})$ and

$$U_{jk}(\Delta) = \prod_{d \neq \{j,k\}} [-\infty, u_d + \beta(\Delta_j - \Delta_d)].$$

Recall that the support of $U_t|x_t = x$ is rectangular and takes the form $\mathcal{U}(x) = \prod_{i=1}^{J+1} [\underline{u}_i, \bar{u}_i]$ for some $-\infty \leq \underline{u}_j < \bar{u}_j \leq +\infty$. In particular, from the expression above, $\partial_{jk}^2 F_{U_t|x_t} > 0$ if and only if $\underline{u}_j \leq u_j \leq \bar{u}_j$, $\underline{u}_k \leq u_j + \beta(\Delta_j - \Delta_k) \leq \bar{u}_k$, $\underline{u}_d \leq u_j + \beta(\Delta_j - \Delta_d)$ for $j \in \mathcal{J} \setminus \{j, k\}$ and $\underline{u}_{J+1} \leq u_j + \beta\Delta_j$. These restrictions in turn implies that $h_{jk}(\Delta|x) > 0$ if and only if $\beta\Delta \in \Delta\mathcal{U}(x) := \text{supp}\{(U_{1,t} - U_{J+1,t}, \dots, U_{J,t} - U_{J+1,t})|x_t = x\}$. This in turn is equivalent to Assumption S.

In consequence, we have that:

$$|D_{jj}(\Delta|x)| = \sum_{i=1, \dots, J, i \neq j} |D_{ji}(\Delta|x)| + h_{jJ+1}(\Delta|x) > \sum_{i=1, \dots, J, i \neq j} |D_{ji}(\Delta|x)|.$$

This implies that $D(\Delta|x)$ has a dominant diagonal (see McKenzie (1959), p. 47) for $\Delta \in \Delta\mathcal{U}(x)$. Because it is also equal to its diagonal form (see McKenzie (1959), p. 60), all its principal minors are positive (see McKenzie (1959), Theorem 4' on p. 60). Then, $D(\Delta|x)$ is a P-matrix (i.e., all its principal minors are positive) (see Gale and Nikaidô (1965), p. 84). We can then apply Theorem 4 in Gale and Nikaidô (1965) which states that $Q(\Delta|x)$ is univalent (i.e., invertible) from $\Delta\mathcal{U}(x)$ onto $Q(\Delta\mathcal{U}(x)|x)$. ■

The result above establishes that one can travel from CCP's to (relative) value functions. This generalises the results of Hotz and Miller (1993) (Proposition 1) to the non separable case: Hotz and Miller (1993) imposes the assumption of Additive Separability (AS). Berry (1994), studying the estimation of static differentiated product demand models, also demonstrates the same invertibility result (again under AS) using contraction mapping arguments. Our proof uses (global) univalence results typically used in the general equilibrium literature. Alternatively, one could also prove the invertibility using the results in Pallas (1959) (after mapping the conditional choice probabilities into the whole Euclidean space as in Chiappori, Komunjer, and Kristensen (2013)), which provides conditions for a given mapping of \mathbb{R}^J onto \mathbb{R}^J to be a diffeomorphism or those in Berry, Gandhi, and Haile (2013).³

We end this section by observing that $Q(\Delta|x)$ also induces a functional $\Lambda(V)$ that maps the integrated value function into CCP's. With $M(V) = [M_d(V)]_{d=1}^J$ defined as

$$M_d(V)(x_t) := \int_{\mathcal{X}} V(x_{t+1}) dF_{x_{t+1}|x_t, d}(x_{t+1}|x_t, d) - \int_{\mathcal{X}} V(x_{t+1}) dF_{x_{t+1}|x_t, d}(x_{t+1}|x_t, J+1), \quad (7)$$

$d = 1, \dots, J$, we obtain from eqs. (3) and (4) that

$$\mathbb{P} = \Lambda(V) := Q(M(V)). \quad (8)$$

3.2 Policy Valuation Operator

The next step in order to extend the results of Aguirregabiria and Mira (2002) to our more general set-up is to show that there also exists an invertible mapping $\varphi(\mathbb{P})$ taking the first J CCP's, $\mathbb{P}(x)$ into the the intergrated value function V , $V = \varphi(\mathbb{P})$. As a first step towards establishing the

³Arguments similar to those provided in Proposition 1 from Hotz and Miller (1993) could also be adapted to the nonseparable case. Nevertheless, as pointed out in Berry, Gandhi, and Haile (2013), that proof establishes only local invertibility. In addition, under certain restrictions on the distribution of unobservables and preferences, the contraction mapping used in Berry (1994) or Berry, Levinsohn, and Pakes (1995) could also be used to demonstrate injectivity. (One can use the contraction mapping there by identifying the unobservables there with our random per-period payoffs and their utility indices with our relative value functions. The application of those results is not immediate in our case, as our random utilities and relative value functions are structurally related.)

existence of φ , we first note that the smoothed value function $V(\mathbf{x})$ solves eq. (1). The smoothed Bellman operator $\Gamma(V)$ is a contraction mapping under the following condition, c.f. Theorem 1 in Rust, Traub, and Wozniakowski (2002):

Assumption CO. The vector of state variables, x_t , has compact support.

More general conditions for the contraction property to hold, allowing for unbounded state space, can be found in Bhattacharya and Majumdar (1989); see also Norets (2010). We here maintain Assumption CO since this facilitates some of the more technical arguments in the following.

The potentially costly step in solving the Bellman equation in (1) is solving a large number of maximization problems. As noted by Hotz and Miller (1993) and Aguirregabiria and Mira (2002), an alternative Bellman operator can be expressed in terms of the choice probabilities \mathbb{P} . This alternative Bellman operator, which we denote as $\Gamma^*(V, \mathbb{P})$ maps integrated value functions into integrated value functions. Note that we explicitly state Γ^* 's dependence on \mathbb{P} . As in eq. (5) of Aguirregabiria and Mira (2002), using that the data-generating CCP's $\mathbb{P}_d(x_t)$ are induced by the optimal decision rule, Γ^* can be written as a functional of (V, \mathbb{P}) ,

$$\Gamma^*(V, \mathbb{P})(x_t) = \sum_{d \in \mathcal{D}} \mathbb{P}_d(x_t) \left\{ \bar{U}_d(\mathbb{P})(x_t) + \beta \int_{\mathcal{X}} V(x_{t+1}) dF_{x_{t+1}|x_t, d}(x_{t+1}|x_t, d) \right\}. \quad (9)$$

where $\bar{U}_d(\mathbb{P})(x_t) := E[u_d(x_t, \epsilon_t)|x_t, d]$ denotes expectations w.r.t. ϵ_t conditional on the optimal decision rule. Given that \mathbb{P} is optimal from the discussion above, we have that $\mathbb{P} = Q(\Delta|x)$. As a result, in our setup \bar{U}_d satisfies

$$\bar{U}_d(\mathbb{P})(x_t) = \int_{\mathcal{E}} u_d(x_t, \epsilon_t) \mathbb{I}_d(\mathbb{P})(x_t, \epsilon_t) dF_{\epsilon_t|x_t}(\epsilon_t|x_t) / \mathbb{P}_d(x_t) \quad (10)$$

where the function \mathbb{I}_d is an indicator that equals one if choice d is optimal,

$$\mathbb{I}_d(\mathbb{P})(x_t, \epsilon_t) = \mathbb{I}\{u_d(x_t, \epsilon_t) - u_j(x_t, \epsilon_t) + \beta(\Delta_d(\mathbb{P})(x_t) - \Delta_j(\mathbb{P})(x_t)) \geq 0 : \forall j\}. \quad (11)$$

Here, $\Delta(\mathbb{P})(x_t) \equiv Q^{-1}(\mathbb{P})(x_t)$ is well-defined due to Theorem 3.

The operator Γ^* is a contraction w.r.t. V for any given \mathbb{P} . As a result, it has a unique fixed point and the optimal value function is the unique solution to the above integral equation. Since Γ^* has been expressed as an operator w.r.t. (V, \mathbb{P}) , its fixed point V can be expressed as a functional of \mathbb{P} , $V = \varphi(\mathbb{P})$. The functional $\varphi(\mathbb{P})$ is implicitly defined by

$$\varphi(\mathbb{P}) = \Gamma^*(\varphi(\mathbb{P}), \mathbb{P}). \quad (12)$$

Note that the above definition of $\varphi(\mathbb{P})$ implicitly makes use of Theorem 3 since we have employed that $\Delta = Q^{-1}(\mathbb{P})$ in its derivation.

The solution of (12) cannot be written explicitly since the operator Γ^* is not available in closed form. Thus, in general, numerical approximations must be employed, see Srisuma and Linton

(2012). However, when $\mathbf{x}_t \in \{\mathbf{x}^1, \dots, \mathbf{x}^M\}$ has discrete support, the solution is easily computed. In that case the operator equation is a matrix equation and can be easily inverted as discussed in Section 5.

The benefit of this representation is that solution of the Bellman equation using maximization over a high dimensional state space is replaced with inversion of a linear operator equation. In effect, data on \mathbb{P} are used to compute optimal behaviour obviating the need to solve the maximization problem repeatedly.

3.3 Policy Iteration Mapping

In previous subsections we showed that the choice probabilities can be written as a functional Λ of the integrated value function V , and that V can also be written as a functional φ of the choice probabilities, where Λ and φ are defined in eq. (8) and (12), respectively. We now combine the two mappings Λ and φ to obtain that \mathbb{P} must satisfy

$$\mathbb{P} = \Psi(\mathbb{P}), \quad \Psi(\mathbb{P}) := \Lambda(\varphi(\mathbb{P})). \quad (13)$$

The mapping Ψ defines our policy iteration mapping. The following theorem is a generalization of Propositions 1-2 in Aguirregabiria and Mira (2002) (see also Proposition 1 in Kasahara and Shimotsu (2008)) stating some important properties of the policy iteration mapping:

Theorem 4 *Let Ψ and φ be defined in eqs. (13) and (12), respectively. Under Assumptions CI, RU and CO:*

- (a) $\mathbb{P} \mapsto \Psi(\mathbb{P})$ has a unique fixed point at \mathbb{P} .
- (b) The Frechet derivatives of $\varphi(\mathbb{P})$ and $\Psi(\mathbb{P})$ w.r.t. \mathbb{P} are zero at the fixed point \mathbb{P} .

Proof. Let V_0 and \mathbb{P}_0 denote the value function and choice probabilities induced by the optimal decision rule.

The proof of (a) follows along the same lines as the proof of Proposition 1(a) in Aguirregabiria and Mira (2002): First note that the integrated Bellman operator $V \mapsto \Gamma(V)$ is a contraction mapping. Suppose now that \mathbb{P} is some fixed point of $\Psi(\mathbb{P}) = \Lambda(\varphi(\mathbb{P}))$. That is, $\mathbb{P} = \Lambda(\varphi(\mathbb{P}))$ and so, using eq. (9) and the definition of φ ,

$$\begin{aligned} \Gamma(\varphi(\mathbb{P})) &= \sum_{d \in \mathcal{D}} \Lambda_d(\varphi(\mathbb{P})) \{ \bar{U}_d(\Lambda(\varphi(\mathbb{P}))) + \beta F_d \varphi(\mathbb{P}) \} \\ &= \sum_{d \in \mathcal{D}} \mathbb{P}_d \{ \bar{U}_d(\mathbb{P}) + \beta F_d \varphi(\mathbb{P}) \} \\ &= \varphi(\mathbb{P}). \end{aligned}$$

We conclude that if \mathbb{P} is a fixed point of Ψ , then $\varphi(\mathbb{P})$ is a fixed point of Γ . Now, suppose that \mathbb{P}_1 and \mathbb{P}_2 are both fixed points of Ψ . Then $\varphi(\mathbb{P}_1)$ and $\varphi(\mathbb{P}_2)$ are both fixed points of Γ and

so $\varphi(\mathbb{P}_1) = \varphi(\mathbb{P}_2)$ since Γ is a contraction. But then $\mathbb{P}_1 = \Lambda(\varphi(\mathbb{P}_1)) = \Lambda(\varphi(\mathbb{P}_2)) = \mathbb{P}_2$. Thus, Ψ has a unique fixed point which is \mathbb{P}_0 .

Part (b) follows by the same arguments as in the proof of Proposition 1 in Kasahara and Shimotsu (2008): They consider the case of additively separable utility on the form $u_d(x_t) + e_{dt}$ where $u_d(x_t)$ and e_t denote the observed and unobserved component of the utility. Now, set $u_d(x_t) := 0$ and $e_{dt} := U_{dt} = u_d(x_t, \epsilon_t)$ and recall that the random utilities $U_t = (U_{1t}, \dots, U_{Jt}) \in \mathcal{U}$ has conditional distribution $F_{U_t|x_t}(u|x)$ satisfying Assumption RU. Thus, e_{dt} , as defined above, satisfies the conditions of Kasahara and Shimotsu (2008) and their arguments apply. For example,

$$\begin{aligned} \bar{U}_d(\mathbb{P})(x) &= E[U_{dt}|x_t, d_t] \\ &= \int_{\mathcal{U}} U_{dt} \mathbb{I}\{U_{dt} - U_{jt} + \beta [Q_d^{-1}(\mathbb{P})(\mathbf{x}_t) - Q_d^{-1}(\mathbb{P})(\mathbf{x}_t)] \geq 0 : j \in \mathcal{D}\} dF_{U_t|x_t}(U_t|x_t), \end{aligned}$$

and we can recycle the arguments in the proof of Lemma 1 of Aguirregabiria and Mira (2002) to obtain that

$$\frac{\partial \left[\sum_{j \in \mathcal{D}} \mathbb{P}_j \bar{U}_j(\mathbb{P}) \right]}{\partial \mathbb{P}} = Q^{-1}(\mathbb{P}).$$

It now easily follows by the same arguments as in Kasahara and Shimotsu (2008) that the Frechet derivatives of $\varphi(\mathbb{P})$ and $\Psi(\mathbb{P}, Q^{-1}(\mathbb{P}))$ are zero at the fixed point. ■

The above properties will prove important when analyzing the CCP estimator in the next section. In particular, they ensure that the first-step estimation of \mathbb{P} , will not affect the first-order properties of the resulting estimator of θ .

4 Generalized CCP Estimator

Having derived the policy iteration mapping Ψ , we are now in a position to extend the CCP-based estimator of Hotz and Miller (1993) and Aguirregabiria and Mira (2002) to our more general framework. Let $f_{x_{t+1}|x_t, d_t}$, $f_{\epsilon_{t+1}|x_{t+1}}$ and $u_d(\cdot)$ be indexed by parameter vectors θ_1 , θ_2 and θ_3 respectively. We collect these in $\theta = (\theta_1, \theta_2, \theta_3)$. Given the policy iteration mapping Ψ that we derived in the previous section, we can now rewrite the likelihood function in terms of the CCP's,

$$\mathcal{L}_{n,T}(\theta) = \prod_{i=1}^n \prod_{t=1}^T \mathbb{P}_{d_{it}}(x_{it}; \theta) f_{x_{t+1}|x_t, d_t}(x_{it}|x_{it-1}, d_{it-1}, \theta_1)$$

subject to $\mathbb{P}(x; \theta)$ solving the fixed point problem

$$\mathbb{P}(x; \theta) = \Psi(\mathbb{P}(\cdot; \theta))(x; \theta). \quad (14)$$

Maximizing $\mathcal{L}_{n,T}(\theta)$ with respect to θ corresponds to the NFXP algorithm, except that we now express the fixed-point problem in terms of CCP's instead of value functions. As before, computing $\mathcal{L}_{n,T}(\theta)$ for a putative value of θ requires solving eq. (14) w.r.t. $\mathbb{P}(x; \theta)$; this involves repeated computation of $\Psi(\mathbb{P})$ as we search for the solution in the space of CCP's.

We can sidestep the calculation of $\mathbb{P}(x; \theta)$ for each putative parameter value θ as we search for a MLE by utilizing that the choice probabilities should satisfy eq. (13): First obtain an initial estimator of the CCP's, say, $\hat{\mathbb{P}}^{(0)}$. This could, for example, be the nonparametric estimator given by

$$\hat{\mathbb{P}}_d^{(0)}(x) = \sum_{i=1}^n \sum_{t=1}^T \mathbb{I}\{d_{it} = d, x_{it} = x\} / (nT),$$

when x_t has discrete support. We will, however, not restrict our attention to this particular estimator and will, for example, allow for kernel smoothed CCP estimators for the case when x_t has continuous support as in Linton and Srisuma (2012). We then use $\hat{\mathbb{P}}^{(0)}$ to define the following pseudo-likelihood function,

$$\hat{\mathcal{L}}_{n,T}(\theta) = \prod_{i=1}^n \prod_{t=1}^T \Psi_{d_{it}}(\hat{\mathbb{P}}^{(0)})(x_{it}; \theta) f_{x_{t+1}|x_t, d_t}(x_{it}|x_{it-1}, d_{it-1}, \theta_1). \quad (15)$$

Maximizing $\hat{\mathcal{L}}_{n,T}(\theta)$, instead of $\mathcal{L}_{n,T}(\theta)$, entails computational gains since it does not involve solving eq. (14) w.r.t. $\mathbb{P}(x; \theta)$ for a putative parameter value θ .

We discuss in more detail how the PMLE that maximizes $\hat{\mathcal{L}}_{n,T}(\theta)$ can be implemented in the next section. We would however already here like to point out that the evaluation of $\Psi(\mathbb{P})(x; \theta)$ is complicated by the fact that it involves computing the inverse $Q^{-1}(\mathbb{P}; \theta)$, which enters $\Psi(\mathbb{P})(x; \theta)$ through $\mathbb{I}_d(\mathbb{P})(x_t, \epsilon_t)$ as defined in eq. (11). The numerical inversion of Q can be time consuming, and so one may wish to avoid this issue. This can be done by redefining $\mathbb{I}_d(\mathbb{P})(x_t, \epsilon_t)$ as an operator w.r.t. both \mathbb{P} and Δ ,

$$\mathbb{I}_d(\mathbb{P}, \Delta)(x_t, \epsilon_t) = \mathbb{I}\{u_d(x_t, \epsilon_t) - u_j(x_t, \epsilon_t) + \beta(\Delta_d(x_t) - \Delta_j(x_t)) \geq 0 : \forall j\}, \quad (16)$$

and the letting $\Psi_{d_{it}}(\mathbb{P}, \Delta)$ be the corresponding policy iteration mapping. Instead of maximizing $\hat{\mathcal{L}}_{n,T}(\theta)$, we can then estimate θ by maximizing

$$\tilde{\mathcal{L}}_{n,T}(\theta, \Delta) = \prod_{i=1}^n \prod_{t=1}^T \Psi_{d_{it}}(\hat{\mathbb{P}}^{(0)}, \Delta)(x_{it}; \theta) f_{x_{t+1}|x_t, d_t}(x_{it}|x_{it-1}, d_{it-1}, \theta_1) \quad (17)$$

w.r.t. (θ, Δ) subject to Δ satisfying

$$\hat{\mathbb{P}}^{(0)}(x) = Q(\Delta)(x; \theta). \quad (18)$$

By formulating the problem in this way, we sidestep the need for computing Q^{-1} . Moreover, Proposition 1 in Judd and Su (2012) gives that the two PML estimators are equivalent if we ignore the numerical issues involved in maximizing either of the two likelihoods and inverting Q . We will therefore in the statistical analysis not differentiate between the two estimators. It should be noted though that the above constrained optimization problem involves maximizing over the space for Δ which in general is an infinite-dimensional function space when x_t is continuous. However, if x_t is discrete then Δ lies in a Euclidean space which facilitates solving the above maximization problem.

If one does not have many observations, one drawback of the PMLE is its reliance on non-parametric estimation of the CCP's. But we can use iterations to improve its efficiency properties

as in Aguirregabiria and Mira (2002). Starting with the initial estimate $\hat{\mathbb{P}}^{(0)}$, we take as input in iteration $K \geq 1$ an estimator $\hat{\mathbb{P}}^{(K)}$ and then let $\hat{\theta}^{(K+1)}$ maximize

$$\hat{\mathcal{L}}_{n,T}(\theta) = \prod_{i=1}^n \prod_{t=1}^T \Psi_{d_{it}}(\hat{\mathbb{P}}^{(K)})(x_{it}; \theta) f_{x_{t+1}|x_t, d_t}(x_{it}|x_{it-1}, d_{it-1}, \theta_1). \quad (19)$$

We use these estimates to update the CCP's that are then used in the next step,

$$\hat{\mathbb{P}}^{(K+1)} = \Psi\left(\hat{\mathbb{P}}^{(K)}; \hat{\theta}^{(K+1)}\right).$$

This iterative process follows the suggestion in Aguirregabiria and Mira (2002). One could alternatively employ one of the modified versions proposed in Kasahara and Shimotsu (2008). Similarly, one can develop an iterative version of the alternative maximization problem defined in eqs. (17)-(18).

Finally, to further facilitate implementation, one can use the following two-step procedure in each iteration when the parameter vector θ can be decomposed into θ_1 , which only enters the dynamics of x_t , and (θ_2, θ_3) , which enters the specification of the the distribution of $\epsilon_t|x_t$ and the per-period utility: First compute

$$\tilde{\theta}_1 = \arg \max_{\theta_1 \in \Theta_1} \prod_{i=1}^n \prod_{t=1}^T f_{x_{t+1}|x_t, d_t}(x_{it}|x_{it-1}, d_{it-1}, \theta_1),$$

and then iterate over

$$(\tilde{\theta}_2^{(K+1)}, \hat{\theta}_3^{(K+1)}) = \arg \max_{\theta_k \in \Theta_k, k=2,3} \prod_{i=1}^n \prod_{t=1}^T \Psi_{d_{it}}(\hat{\mathbb{P}}^{(K)})(x_{it}; \hat{\theta}_1, \theta_2, \theta_3).$$

Again, there will in general be a loss of statistical efficiency from doing so, but this may be off-set by computational gains.

To analyze the properties of the (iterative) PMLE, we will assume that the initial CCP estimator satisfies the following condition with $\|\mathbb{P}\|_\infty = \sup_{x \in \mathcal{X}} |\mathbb{P}(x)|$ denoting the uniform norm:

Assumption P. (a) $\|\hat{\mathbb{P}}^{(0)} - \mathbb{P}\|_\infty = o_P(1)$; (b) for some $\gamma > 0$, $\|\hat{\mathbb{P}}^{(0)} - \mathbb{P}(x)\|_\infty = O_P(n^{-\gamma})$.

Assumption P allows for a broad class of initial estimators, including nonparametric kernel smoothers when x_t is continuously distributed. Under Assumption P(a), we will establish convergence rate of the estimator, while P(b) will be used to show that the PMLE is first order equivalent (that is, up to an error term of order $o_P(1/\sqrt{n})$) to the MLE.

In addition to Assumption P, we also need to impose some regularity conditions on the model to ensure identification and smoothness:

Assumption PML. $\{(d_{it}, x_{it}) : t = 1, \dots, T\}$, $i = 1, \dots, n$, are i.i.d.; the true parameter value $\theta_0 = (\theta_{0,1}, \theta_{0,2}, \theta_{0,3})$ lies in the interior of a compact parameter space Θ such that $f_{x_t|x_{t-1}}(x'|x, d; \theta_1) = f_{x_t|x_{t-1}}(x'|x, d; \theta_{0,1})$, $\mathbb{P}_d(x, \theta) = \mathbb{P}_d(x)$ and $\Psi_d(\mathbb{P})(x; \theta) = \mathbb{P}_d(x)$ implies $\theta = \theta_0$; $\hat{\theta}_1 = \theta_{0,1} + O_P(1/\sqrt{n})$; the MLE satisfies $\sqrt{n}(\hat{\theta}_{\text{MLE}} - \theta_0) \rightarrow^d N(0, \mathcal{I}^{-1})$; $\Psi_d(\mathbb{P})(x; \theta)$ is three

times Frechet-differentiable w.r.t. \mathbb{P} and is strictly positive for all $(d, x, \mathbb{P}; \theta)$; the s th Frechet derivative $D^s \Psi_d(\mathbb{P})(x; \theta)$ w.r.t. \mathbb{P} satisfies $E \left[\sup_{\theta \in \Theta} \|D^s \Psi_d(\mathbb{P})(x_t; \theta)\|^2 \right] < \infty$, $s = 1, 2, 3$; $\Psi_d(\mathbb{P})(x, ; \theta)$ is three times continuously differentiable w.r.t. θ , and its derivatives are uniformly bounded and Lipschitz continuous.

The above conditions are identical to the ones in Assumption 4 of Kasahara and Shimotsu (2008). These are high-level conditions, and could be replaced by more primitive conditions involving the underlying utility function and transition densities for (x_t, ϵ_t) as in Rust (1988). We will not pursue this here though.

We first analyze the effects from replacing the true set of choice probabilities, \mathbb{P} , with an estimator, $\hat{\mathbb{P}}^{(K)}$, and how this effect evolves as the number of iterations K increases. Following the arguments of Kasahara and Shimotsu (2008), we obtain the following higher-order expansion showing how the first-step estimation of the CCPs and the number of iterations affect the PMLE:

Theorem 5 *Under Assumptions CI, RU, CO, P(a) and PML, as $n \rightarrow \infty$,*

$$\begin{aligned} \|\hat{\theta}^{(K)} - \hat{\theta}\| &= O_P \left(n^{-1/2} \|\hat{\mathbb{P}}^{(K-1)} - \mathbb{P}\|_\infty \right) + O_P \left(\|\hat{\mathbb{P}}^{(K-1)} - \mathbb{P}\|_\infty^2 \right), \\ \|\hat{\mathbb{P}}^{(K)} - \hat{\mathbb{P}}\|_\infty &= O_P \left(\|\hat{\theta}^{(K)} - \hat{\theta}\| \right), \end{aligned}$$

where $\hat{\theta}$ is the MLE. and $\hat{\mathbb{P}}$ is the CCP it implies.

Proof. The proof is identical to the one of Proposition 2 of Kasahara and Shimotsu (2008) except that their Proposition 1 is replaced by our Theorem 4. The proof in Kasahara and Shimotsu (2008) also relies on their Lemmas 7-8, but by inspection it is easily seen that these lemmas remain valid in our setting by the same arguments as employed in the proof of Theorem 4: They consider the case where the per-period utility is additively separable. That is, it takes the form $u_d(x_t) + e_{dt}$, where we use $u_d(x_t)$ and e_{dt} to denote the observed and unobserved component of the utility as imposed their work. Setting $u_d(x_t) := 0$ and $e_{dt} := u_d(x_t, \epsilon_t)$, our assumptions CI and RU now imply Assumptions 1-3 of Kasahara and Shimotsu (2008). Furthermore, our Assumption PML implies their Assumption 4. ■

One particular consequence of the above result is that if the convergence rate in Assumption P(b) is $\gamma > 1/4$, then the PMLE is first-order equivalent to the exact MLE:

Corollary 6 *Under Assumptions CI, RU, CO, P(b) and PML, $\|\hat{\theta}^{(K)} - \hat{\theta}\| = O_P(n^{-(K-1)/2-2\gamma})$ for $K = 1, 2, \dots$. In particular, if $\gamma > 1/4$ in Assumption P(b) then for any fixed $K \geq 1$,*

$$\sqrt{n}(\hat{\theta}^{(K)} - \theta_0) \rightarrow^d N(0, \mathcal{I}_T^{-1}(\theta_0)),$$

where $\mathcal{I}_T(\theta)$ is the T -period asymptotic information matrix,

$$\mathcal{I}_T(\theta) = - \sum_{t=1}^T E \left[\frac{\partial^2 \log P(d_t | x_t, \theta)}{\partial \theta \partial \theta'} + \frac{\partial^2 \log f_{\mathbf{x}_{t+1} | \mathbf{x}_t, d_t}(\mathbf{x}_{it} | \mathbf{x}_{it-1}, d_{it-1}, \theta)}{\partial \theta \partial \theta'} \right].$$

The above result shows that asymptotically, the PMLE is first-order equivalent to the MLE when the initial CCP estimator converges with rate $o_P(n^{-1/4})$ in which case the first-step estimation of \mathbb{P} has no impact on the estimation of θ . However, as Theorem 5 shows, the higher-order properties of the PMLE and the MLE will in general not be identical and so we recommend to that the bootstrap procedure as proposed in Kasahara and Shimotsu (2008) is implemented in order to get a better approximation of the finite-sample distribution of the PMLE.

5 Numerical implementation

Here we will discuss the numerical implementation of the CCP estimator developed in the previous section.

5.1 Computation of Q

Recall that in general

$$Q_d(\Delta|x) = \int_{u_d=-\infty}^{\infty} \partial_d F_{U_t|x_t}(u_d + \beta(\Delta_d - \Delta_1), \dots, u_d, \dots, u_d + \beta\Delta_d|x) du_d,$$

which requires knowledge of $\partial_d F_{U_t|x_t}$. Under the conditions of Lemma 1, where $u_d(x, \epsilon) = u_d(x, \epsilon_d)$, we obtain

$$Q_d(\Delta|x) = \int_{\epsilon_d=-\infty}^{\infty} \partial_d F_{\epsilon_t|x_t}(u_1^{-1}(x, u_d(x, \epsilon_d) + \beta(\Delta_d - \Delta_1)), \dots, \epsilon_d, \dots, u_{J+1}^{-1}(x, u_d(x, \epsilon_d) + \beta\Delta_d)|x) d\epsilon_d.$$

Assuming $\partial_d F_{\epsilon_t|x_t}$ is available on analytical form, this can be computed using numerical integration.

5.2 Computation of PMLE

This section will describe our numerical implementation of the estimator.

1. Estimate \mathbb{P}^0 .
2. Estimate θ_1 of data on (x_t, d_t) .
3. Guess (θ, V, Δ) .
4. Compute $\Lambda(d_{it}|x_{it}, V, \theta)$ and the gradient with respect to (θ, V) . This requires us to compute a J dimensional numerical integral using methods based on Genz and Bretz (2009).
5. Compute constraint (??) and its gradient: $V - \Gamma^*(V, \mathbb{P}^0, \Delta)$. This step merely requires computation of several sums and simple integrals.
6. Compute inversion constraint $\mathbb{P}^0 - Q(\Delta, x_{it})$. This step requires computation of a J dimensional numerical integral as above.

It should be noted that when x_t is discrete valued then there are a finite number of constraints. If x_t is continuous, we will approximate $V = \sum_k \alpha_{vk} B_k(x)$. In this latter case the maximization is over (θ, α_v) . Recall also that Γ^* depends implicitly on the inverse $\Delta = Q^{-1}(\mathbb{P})$. We first discuss the case where the support of x_t is discrete and then the case where it is continuous.

5.2.1 Discrete x_t

Let $\mathcal{X} = \{x_1, \dots, x_M\}$ be the discrete support of x_t . Then \mathbb{P} can be expressed as a $(J \times M)$ -matrix,

$$\mathbb{P} := (\mathbb{P}_1, \dots, \mathbb{P}_J)^\top \in \mathbb{R}^{J \times M}.$$

5.2.2 Continuous x_t

As in Linton and Srisuma (2012).

6 Numerical Results

TBA

7 Extensions

7.1 Wages

Consider an occupational choice model where the econometrician observes the wages received at the chosen occupation. In this case, suppose that the utility in occupation d is given by $u(w_d; \gamma)$, where w_d is the wage in that sector and $u(\cdot; \gamma)$ is increasing (e.g., $u(w_d; \gamma) = w_d^{1-\gamma}/(1-\gamma)$ or simply w_d as in Keane and Wolpin (1997)). For simplicity, we omit the parameter γ in what follows. Assume in addition that

$$\ln w_d = x^\top \theta_d + \epsilon_d$$

as is common. Notice that the probability density function for $\ln w_d$ given that sector d is chosen and given x is

$$\frac{\partial_d F_{\epsilon_t | \mathbf{x}_t}(u_1^{-1}(x, u_d(x, \epsilon_d) + \beta(\Delta_d - \Delta_1)), \dots, \epsilon_d, \dots, u_{J+1}^{-1}(x, u_d(x, \epsilon_d) + \beta \Delta_d) | x)}{\mathbb{P}_d(d | x)} \Bigg|_{\epsilon_d = \ln w_d - \mathbf{x}^\top \theta_d}.$$

Then, the analysis proceeds as before with the estimator (in iteration K) obtained by maximizing

$$\hat{\mathcal{L}}_{n, T}(\theta) = \prod_{i=1}^n \prod_{t=1}^T \Psi_{d_{it}}(\hat{\mathbb{P}}^{(K)})(x_{it}; \theta) f_{\ln w_d | \mathbf{x}, d}(\ln w_{id_{it}} | x_{it}, d_{it}) \times f_{\mathbf{x}_{t+1} | \mathbf{x}_t, d_t}(x_{it} | x_{it-1}, d_{it-1}, \theta_1)$$

7.2 Types

If there are L time-invariant types indexed by $l \in \{1, \dots, L\}$ (which do not affect the transition law for the state variables), the pseudo-likelihood now becomes:

$$\mathcal{L}_{n,T}(\theta, \pi) = \prod_{i=1}^n \sum_{l=1}^L \pi_l \prod_{t=1}^T \Psi_{d_{it}}(\hat{\mathbb{P}}^{(K)}(l))(x_{it}; \theta, l) f_{x_{t+1}|x_t, d_t}(x_{it}|x_{it-1}, d_{it-1}, \theta_1),$$

that we maximize w.r.t. θ . Notice that now all the appropriate functions are conditional on the type l . Once initial values for the (type specific) conditional choice probabilities, we can proceed as before (see Aguirregabiria and Mira (2007)). Alternatively, in each iteration we can follow Arcidiacono and Miller (2011) and, letting $\hat{\mathbb{P}}^{(K)} \equiv (\hat{\mathbb{P}}^{(K)}(1), \dots, \hat{\mathbb{P}}^{(K)}(L))$ be the k -th iteration type-specific CCPs, use a sequential EM algorithm that:

1. (E-Step) Given current estimates $(\hat{\theta}^K, \hat{\pi}_l^K)$, calculate

$$\hat{\mathbb{P}}_{d_{nt}}^{K+1}(l, x_{nt}^-) = \frac{\hat{\pi}_l^K \prod_{t=1}^T \Psi_{d_{nt}}(\hat{\mathbb{P}}^{(K)}(l))(x_{nt}^-; \hat{\theta}^K, l)}{\sum_{s'} \hat{\pi}_{l'}^K \prod_{t=1}^T \Psi_{d_{nt}}(\hat{\mathbb{P}}^{(K)}(l'))(x_{nt}^-; \hat{\theta}^K, l')}$$

which follows from Bayes' rule;

2. (M-Step) Update the estimate to $\hat{\theta}^{K+1}$ by maximizing the likelihood function above substituting $\hat{\mathbb{P}}_{d_{nt}}^{K+1}(l, x_{nt}^-)$ for π_l . This would be done using MPEC.
3. (M-Step) Update $\hat{\pi}_s^{K+1}$ as

$$\frac{1}{N} \sum_{n=1}^N \hat{\mathbb{P}}_{d_{nt}}^{K+1}(l, x_{nt}^-)$$

4. Update the type-specific CCPs using:

$$\hat{\mathbb{P}}_1^{K+1}(l, x_{nt}^-) = \frac{\sum_n \sum_t d_{nt} \mathbb{P}^{K+1}(l_n = l | d_{nt}, x_{nt}^-; \hat{\theta}^K, \hat{\pi}_l^K, \hat{\mathbb{P}}^K) \mathbf{1}(x_{nt} = x)}{\sum_n \sum_t \mathbb{P}^{K+1}(l_n = l | d_{nt}, x_{nt}^-; \hat{\theta}^K, \hat{\pi}_l^K, \hat{\mathbb{P}}^K) \mathbf{1}(x_{nt} = x)}$$

The value at which one initializes $\hat{\mathbb{P}}^{(0)}$ should be immaterial (in the limit) in either case. Arcidiacono and Miller (2011) mention that “[n]atural candidates for these initial values come from estimating a model without any unobserved heterogeneity and perturbing the estimates” (p.1848). Aguirregabiria and Mira (2002) show that their estimator (without types) converges to a “root of the likelihood function” if initialized at arbitrary CCPs (and is consistent if initialized at a consistent estimator of the CCPs). Arcidiacono and Miller (2011) claim that (if identified) their estimator is consistent (but the proof is not in the main text). (Arcidiacono and Miller (2011) allow the type to be time-varying.) An alternative is to obtain estimates for the CCPs from the suggestions in Kasahara and Shimotsu (2009) or Bonhomme, Jochmans, and Robin (to appear). One will also need to say something about the *initial conditions* problem.

References

- AGUIRREGABIRIA, V. (2004): “Pseudo maximum likelihood estimation of structural models involving fixed-point problems,” *Economics Letters*, 84, 335–340.
- AGUIRREGABIRIA, V., AND P. MIRA (2002): “Swapping the Nested Fixed Point Algorithm: A Class of Estimators for Discrete Markov Decision Models,” *Econometrica*, 70(4), 1519–1543.
- (2007): “Sequential Estimation of Dynamic Discrete Games,” *Econometrica*, 75(1), 1–53.
- ARCIDIACONO, P., AND R. MILLER (2011): “Conditional Choice Probability Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity,” *Econometrica*, 79(6), 1823–1867.
- BERRY, S. (1994): “Estimating Discrete-Choice Models of Product Differentiation,” *The Rand Journal of Economics*, 25(2), 242–262.
- BERRY, S., A. GANDHI, AND P. HAILE (2013): “Connected substitutes and invertibility of demand,” *Econometrica*, 81(5), 2087–2111.
- BERRY, S., LEVINSOHN, AND A. PAKES (1995): “Automobile Prices in Market Equilibrium,” *Econometrica*, 63(4), 841–890.
- BHATTACHARYA, R., AND M. MAJUMDAR (1989): “Controlled Semi-Markov Models: The Discounted Case,” *Journal of Statistical Planning and Inference*, 21, 365–381.
- BONHOMME, S., K. JOCHMANS, AND J.-M. ROBIN (to appear): “Estimating Multivariate Latent-Structure Models,” *The Annals of Statistics*.
- CHIAPPORI, P.-A., I. KOMUNJER, AND D. KRISTENSEN (2013): “On the Nonparametric Identification of Multiple Choice Models,” Columbia University Working Paper.
- GALE, D., AND H. NIKAIIDÔ (1965): “The Jacobian matrix and global univalence of mappings,” *Mathematische Annalen*, 159, 81–93.
- GENZ, A., AND F. BRETZ (2009): *Computation of multivariate normal and t probabilities*, vol. 45. Springer Heidelberg.
- HOTZ, V. J., AND R. A. MILLER (1993): “Conditional Choice Probabilities and the Estimation of Dynamic Models,” *Review of Economic Studies*, 60, 497–529.
- JUDD, K., AND C.-L. SU (2012): “Constrained Optimization Approaches to Estimation of Structural Models,” *Econometrica*.
- KASAHARA, H., AND K. SHIMOTSU (2008): “Pseudo-likelihood estimation and bootstrap inference for structural discrete Markov decision models,” *Journal of Econometrics*, 146, 92–106.

- (2009): “Nonparametric Identification of Finite Mixture Models of Dynamic Discrete Choices,” *Econometrica*, 77(1), 135–175.
- (2012): “Sequential Estimation of Structural Models with a Fixed Point Constraint,” *Econometrica*, 80(5), 2303–2319.
- KEANE, M., AND K. WOLPIN (1997): “The Career Decisions of Young Men,” *Journal of Political Economy*, 105(3).
- KRISTENSEN, D., AND B. SCHJERNING (2014): “Implementation and Estimation of Discrete Markov Decision Models by Sieve Approximations,” UCL and University of Copenhagen Working Paper.
- LINTON, O., AND S. SRISUMA (2012): “Semiparametric Estimation of Markov Decision Processes with Continuous State Space,” *Journal of Econometrics*, 166, 320–341.
- MAGNAC, T., AND D. THESMAR (2002): “Identifying Dynamic Discrete Decision Processes,” *Econometrica*, 70, 801–816.
- MCKENZIE, L. (1959): “Matrices with Dominant Diagonals and Economic Theory,” in *Mathematical Methods in the Social Sciences*, ed. by K. Arrow, S. Karlin, and P. Suppes. Stanford University Press, Stanford.
- NORETS, A. (2010): “Continuity and differentiability of expected value functions in dynamic discrete choice models,” *Quantitative Economics*, 1, 305–322.
- NORETS, A., AND S. TAKAHASHI (2013): “On the Surjectivity of the Mapping Between Utilities and Choice Probabilities,” *Quantitative Economics*, 4, 1490155.
- PAKES, A. (1986): “Patents as Options: Some Estimates of the Value of Holding European Patent Stocks,” *Econometrica*, 54(4), 755–784.
- PALLAIS, R. (1959): “Natural Operations on Differential Forms,” *Transactions of the American Mathematical Society*, 92, 125–141.
- RUST, J. (1987): “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zucher,” *Econometrica*, 55(5), 999–1033.
- RUST, J. (1988): “Maximum Likelihood Estimation of Discrete Control Processes,” *SIAM Journal of Control and Optimization*, 26, 1006–1024.
- RUST, J. (1994): “Structural Estimation of Markov Decision Processes,” in *Handbook of Econometrics*, ed. by R. Engle, and D. McFadden, vol. 4, chap. 51. Elsevier.
- RUST, J., F. TRAUB, AND H. WOZNIKOWSKI (2002): “Is There a Curse of Dimensionality for Contraction Fixed Points in the Worst Case?,” *Econometrica*, 70, 285–329.

SRISUMA, S., AND O. LINTON (2012): “Semiparametric estimation of Markov decision processes with continuous state space,” *Journal of Econometrics*, 166, 320–341.

TODD, P. E., AND K. I. WOLPIN (2006): “Assessing the Impact of a School Subsidy Program in Mexico: Using a Social Experiment to Validate a Dynamic Behavioral Model of Child Schooling and Fertility,” *American Economic Review*, 96(5), 1384–1417.