Econometric Theory, **31**, 2015, 521–538.
doi:10.1017/S0266466614000437

IDENTIFICATION IN DISCRETE MARKOV DECISION MODELS

SORAWOOT SRISUMA
University of Surrey

We derive conditions for the identification of the structural parameters in Markov decision model under the assumptions of Rust (1987, *Econometrica* 55, 999–1033) when the payoff function is parametrically specified. Identification in this class of dynamic problems is difficult to establish since the parameters of interest enter the value function nonlinearly, and the value function is only defined implicitly as a fixed point of some functional equation. We show it is sufficient to verify identification in the pseudomodel, which is more tractable as it is originally designed to reduce the computational burden in the estimation problem, for the identification of the data generating parameter of the underling model. Our results extend naturally to a class of dynamic discrete action games commonly used in empirical industrial organizations.

1. INTRODUCTION

The discrete Markov decision model that is surveyed in Rust (1994) provides a useful framework to study dynamic decision problems in many areas of economics. Importantly, it also forms the foundation to model dynamic games in the empirical industrial organization literature (see Aguirregabiria and Mira, 2010 and Bajari, Hong, and Nekipelov, 2013 for recent surveys). Several estimation methodologies have been proposed to estimate a class of decision problems that satisfies additive separability and conditional independence (Rust, 1987), where the main object of interest is the finite dimensional structural parameter that parameterizes the payoff function of the economic agent. The identification conditions for these parameters are difficult to establish due to the nonlinear structure inherent in the dynamic optimization framework. This paper derives more tractable conditions for the identification of the structural parameters in these models.

Existing identification results in this literature are nonparametric in nature. In a single agent setting, a well-known negative result of Rust (1994) says that without any structure a dynamic decision model is nonparametrically not identified. Magnac and Thesmar (2002) show that nonparametric identification of the payoff function is possible under additive separability and conditional

I thank the editor, a co-editor, and anonymous referees for valuable suggestions that have greatly improved the paper. I also thank Xiaohong Chen, Oliver Linton, and Hashem Pesaran for helpful comments. Address correspondence to School of Economics, University of Surrey, Guildford, Surrey, GU2 7XH, United Kingdom; e-mail: s.srisuma@surrey.ac.uk.

© Cambridge University Press 2014

https://doi.org/10.1017/S0266466614000437 Published online by Cambridge University Press

independence assumptions by imposing additional further restrictions on the payoffs. Pesendorfer and Schimdt-Dengler (2008) use a similar approach to provide conditions to identify dynamic discrete games, and Bajari, Chernozhukov, Hong, and Nekipelov (2009) extend their results for games when the observed state variables are continuous. Although nonparametric identification results are fundamental, our work is motivated by the fact that most of the methodology and empirical literature employ a parametric payoff function.[1]

One obstacle in the study of identification can be attributed to the nonlinearity generated by the recursive structure of the dynamic programming problem in the form of a value function. The intractability of value functions is a well-known computational issue for estimation. Indeed, earlier emphasis in the literature has been on developing feasible estimation methods focusing on the value function, which generally has no closed-form and is defined implicitly as a fixed point of some nonlinear function equation (e.g., see Rust, 1994, 1996). An alternative estimation approach, aimed to reduce the computational burden, following Hotz and Miller (1993) instead estimates a "pseudomodel" in two steps, where the value function is replaced by an easy to compute *policy value function* in the first stage. The two-step approach is particularly attractive for the estimation of dynamic games as it has been shown to circumvent incomplete models that may arise due to the presence of multiple equilibria. Examples of two-step estimators include: (single agent problems) Hotz and Miller (1993), Hotz, Miller, Sanders, and Smith (1994), Aguirregabiria and Mira (2002), Srisuma and Linton (2012); (dynamic games) Aguirregabiria and Mira (2007), Bajari, Benkard, and Levin (2007), Bajari et al. (2012), Pakes, Ostrovsky, and Berry (2007), and Pesendorfer and Schmidt-Dengler (2008).

The identification problem we consider in this paper is set in a parametric framework. Although our results are also directly applicable in semiparametric models where the infinite dimensional nuisance parameters are nonparametrically identified (e.g., Bajari et al., 2012; Srisuma and Linton, 2012). General conditions for the identification of nonlinear parametric models have been provided by Wegge (1965) and Rothenberg (1971).[2] However, these results are often difficult to apply to a specific nonlinear model. Our strategy is to study identification in the pseudomodel. We first show that identification in the pseudomodel is sufficient for the identification of the true model. Then we exploit the fact that policy value function can generally be defined as a linear transformation of the payoff function, so that the pseudomodel provides a tractable framework for the parameterization of the structural parameters. We then give conditions on the observables for the model to be identified. We also emphasize that establishing the identification of the pseudomodel is important in itself, since it is necessary for consistent estimation of all two-step estimators cited above.

We focus on the single agent model to develop our intuition and arguments. Our identification strategy for single agent models is directly applicable to a popular of dynamic games where players are known to play pure strategies almost

https://doi.org/10.1017/S0266466614000437 Published online by Cambridge University Press

surely (a.s.). Thus, to avoid additional notations and discussions on the multiplicity problem specific to games, we shall present our main results for a single agent model where the setup is more transparent. The paper is organized as follows. Section 2 introduces the decision problem of interest and defines the model, pseudomodel, and identification concept. Section 3 gives the main results of the paper. Section 4 briefly describes how our results can be extended to dynamic games. Section 5 concludes with a discussion on the identification of related decision problems. All proofs can be found in the Appendix.

2. MARKOV DECISION PROCESSES

Our starting point is the basic framework of a general single agent decision problem described in Rust (1994), which we use to define the statistical model derived from the decision problem and illustrate the intractability of the parameter parameterization. We introduce the pseudomodel through the pseudodecision problem in Section 2.2. We define and discuss identifiability in Section 2.3.

2.1. Basic Setup

We consider the decision process of a forward looking agent who solves an infinite horizon intertemporal decision problem. The random variables in the model are the control (choice) and state variables, denoted by $a_t \in A$ and $s_t \in S$, respectively. The action set A is finite, $A = \{0, 1, \ldots, K\}$. The state $s_t = (x_t, \varepsilon_t) \in X \times \mathcal{E}$ can be partitioned into two components where, from the perspective of the econometrician, x_t is observable and $\varepsilon_t = (\varepsilon_t(0), \ldots, \varepsilon_t(K))$ is unobservable. At each period t, the agent observes s_t and chooses an action $a_t(s_t)$ in order to maximize her discounted expected utility. The payoffs are time separable and are represented by $u_\theta(a_t, s_t)$ in each period, which is known up to $\theta \in \Theta$, where Θ is a subset of $\mathbb{R}^p$. The agent's action in period t affects the uncertain future states according to the first order Markov transition law $F(s_{t+1}|s_t, a_t)$. The next period utility is subjected to discounting at the rate $\beta \in (0, 1)$. Formally, for any time t, the agent is represented by a triple of primitives (θ, β, F), who is assumed to behave according to an optimal decision rule, $\{a_\tau(s_\tau)\}_{\tau=t}^{\infty}$, in solving the following sequential problem

$$W_\theta(s_t) = \max_{\{a_\tau(s_\tau)\}_{\tau=t}^{\infty}} E\left[\sum_{\tau=t}^{\infty} \beta^{\tau-t} u_\theta(a_\tau(s_\tau), s_\tau) \middle| s_t\right] \text{ s.t. } a_\tau(s_\tau) \in A \text{ for all } \tau \geq t, \tag{1}$$

where W_θ denotes the value function. Under some regularity conditions, there exists a stationary Markovian *optimal policy function* $\delta_\theta : S \to A$ so that $\delta_\theta(s_t) = \delta_\theta(s_{t+\tau})$ for any $s_t = s_{t+\tau}$ and any t, τ, where

$$\delta_\theta(s_t) = \arg\max_{a \in A} \left\{u_\theta(a, s_t) + \beta E\left[W_\theta(s_{t+1}) | s_t, a_t = a\right]\right\}, \tag{2}$$

https://doi.org/10.1017/S0266466614000437 Published online by Cambridge University Press

and W_θ is the unique solution to the *Bellman's equation*

$$W_\theta(s_t) = \max_{a \in A} \left\{ u_\theta(a, s_t) + \beta E\left[W_\theta(s_{t+1}) | s_t, a_t = a \right] \right\}. \quad (3)$$

Sufficient conditions for the existence of stationary optimal policies for the decision problem described above can be found in Section 3 of Rust (1994). Note that δ_θ and W_θ also generally depend on (β, F). We suppress the explicit dependence on these primitives for the ease notational since our focus is on the identification of θ, as (β, F) are often either assumed to be known or nonparametrically identified in practice. We shall also impose the following modeling assumptions throughout the paper.

Assumption M1. The discounting factor β is known.

Assumption M2. The transitional distribution of the states has the following factorization: $F(x_{t+1}, \varepsilon_{t+1} | x_t, \varepsilon_t, a_t) = Q(\varepsilon_{t+1} | x_t) G(x_{t+1} | x_t, a_t)$, where Q is the conditional distribution function of ε_t given x_t, and G denotes the transition law of x_{t+1} conditioning on x_t and a_t.

Assumption M3. (i) u_θ is *additive separable,* $u_\theta(a_t, x_t, \varepsilon_t) = \pi_\theta(a_t, x_t) + \sum_{a \in A} \varepsilon_t(a) \mathbf{1}[a_t = a]$; (ii) for each x, the conditional distribution of ε_t given $x_t = x$ is absolutely continuous with respect to (w.r.t.) the Lebesgue measure full support on $\mathbb{R}^{K+1}$ and the first moment of$\varepsilon_t(a)$ exists for all a.

Assumption M4. x_t has finite support, $X = \left\{ x^1, \ldots, x^J \right\}$.

Assumptions M1–M4 are standard in the (identification and estimation) literature for Markov decision problems (e.g., see Rust, 1987; Hotz and Miller, 1993; Aguirregabiria and Mira, 2002; Magnac and Thesmar, 2002). Under M1–M3 the primitives of the model reduce to θ. We define the statistical model for the Markov decision processes as follows.

DEFINITION S1. *The true model is a collection of conditional choice probabilities* $\{\mathcal{P}_\theta\}_{\theta \in \Theta}$ *such that for all a:*

$$\begin{aligned} \mathcal{P}_\theta(a|x_t) &= \Pr[\delta_\theta(x_t, \varepsilon_t) = a | x_t] \ \ a.s., \ \ \textit{where} \\ \delta_\theta(s_t) &= \arg\max_{a \in A} \left\{ u_\theta(a, s_t) + \beta E\left[W_\theta(s_{t+1}) | x_t, a_t = a \right] \right\}, \\ W_\theta(s_t) &= \max_{a \in A} \left\{ u_\theta(a, s_t) + \beta E\left[W_\theta(s_{t+1}) | x_t, a_t = a \right] \right\}. \end{aligned} \quad (4)$$

Note that the conditioning set for the discounted expected future returns at s_{t+1} only depends on x_t, a_t since ε_t has no additional information for s_{t+1} on under M2. The intractability of the parameterization in $\mathcal{P}_\theta$ is well-documented in the methodology literature, as the following simple example from Rust (1987) illustrates.

https://doi.org/10.1017/S0266466614000437 Published online by Cambridge University Press

Example

At each period the manager inspects a bus and decides whether he should replace the existing engine or to service it. The tradeoffs are: there is an immediate fixed-cost of replacement but the maintenance costs and risks of future breakdowns increase as the bus continues to run on an old engine. Here a_t takes value 1 if the decision is to replace the engine and 0 otherwise, x_t denotes the mileage on the bus, and $\varepsilon_t = (\varepsilon_t(0), \varepsilon_t(1))$ is a vector of unobserved costs (of goodwill, etc.) associated with choice 0 and 1, respectively. The payoff function that represents the cost to replace the bus engine is described as follows

$$u_\theta(a_t, s_t) = \begin{cases} -\theta_1 - c(0, \theta_2) + \varepsilon_t(1) & \text{if } a_t = 1 \\ -c(x_t, \theta_2) + \varepsilon_t(0) & \text{if } a_t = 0 \end{cases},$$

where $\theta = (\theta_1, \theta_2)$ with θ_1 denoting the fixed engine replacement cost, and $c(x, \theta_2)$ is a parametric function describing the maintenance costs, parameterize by θ_2. The transition law for the mileage is taken to be known, otherwise G can be estimated (parametrically or nonparametrically) independently of θ; Rust (1987) discretizes X and estimates G using a multinomial distribution. Costs of future periods are then discounted by β. In the special case when ε_t has i.i.d. (independent and identically distributed) extreme value distribution, the choice probability of replacing the engine with mileage x is:

$$\begin{aligned} \mathcal{P}_\theta(1|x) = \frac{\exp\{-\theta_1 - c(0,\theta_2) + \beta E[W_\theta(s_{t+1})|x_t = x, a_t = 1]\}}{\exp\{-\theta_1 - c(0,\theta_2) + \beta E[W_\theta(s_{t+1})|x_t = x, a_t = 1]\}} \\ + \exp\{-c(x,\theta_2) + \beta E[W_\theta(s_{t+1})|x_t = x, a_t = 0]\}, \end{aligned} \quad \textbf{(5)}$$

which takes a familiar logit form with an addition of a dynamic term. However, $E[W_\theta(s_{t+1})|x_t = \cdot, a_t = \cdot]$ generally does not have a closed form, it is defined implicitly as a fixed point of a functional map, $\Upsilon_\theta : \mathcal{W} \to \mathcal{W}$, where $\mathcal{W}$ is a space of functions whose domain is $A \times X$ such that for all $(a, x) \in A \times X$ and $w \in \mathcal{W}$:

$$\Upsilon_\theta(w)(a,x) = \int E[\max_{a' \in A} \{u_\theta(a', s_t) + \beta w(a', x_t)\} | x_t = x'] dG(x'|x, a). \quad \textbf{(6)}$$

By applying Blackwell's theorem and its generalization it is often straightforward to show that Υ_θ is a contraction map (e.g., Theorem 1 in Rust, 1987), in addition, under the conditional logit framework Υ_θ simplifies to

$$\Upsilon_\theta(w)(a,x) = \int \log \begin{pmatrix} \exp\{-\theta_1 - c(0,\theta_2) + \beta w(1, x')\} \\ + \exp\{-c(x', \theta_2) + \beta w(0, x')\} \end{pmatrix} dG(x'|x, a).$$

Thus $E[W_\theta(s_{t+1})|x_t = \cdot, a_t = \cdot]$, hence $\mathcal{P}_\theta$, generally depends on θ in a complicated manner. In fact, the expression in the previous display is unusually transparent since the function $E[\max_{a' \in A}\{u_\theta(a', s_t) + \beta w(a', x_t)\}|x_t = \cdot]$, that enters Υ_θ in equation (6), generally does not have a closed-form when ε_t does not have i.i.d. extreme value distribution.

https://doi.org/10.1017/S0266466614000437 Published online by Cambridge University Press

2.2. Pseudodecision Problem

Estimation procedures that solve the dynamic optimization problem for every θ in order to evaluate $\mathcal{P}_\theta$ is often called a full-solution approach. This is typically characterized by the need to approximate some functional of W_θ by solving a nonlinear functional equation, for instance, as in the Example above, by applying fixed-point iteration on Υ_θ in equation (6). We next introduce the pseudomodel that underlies the popular two-step approach of Hotz and Miller (1993), which is designed to reduce the computational burden of estimating dynamic decision problems.

Suppose we are interested in θ^*, which generates the data according to δ^* that equals δ_θ for $\theta = \theta^*$as defined in (2). The policy value function generated by δ^*-, which we denote by V_θ, is the discounted expected payoff an economic agent would receive from making her decision based on δ^* in every period:

$$V_\theta(s_t) = E\left[\sum_{\tau=t}^{\infty} \beta^{\tau-t} u_\theta\left(\delta^*(s_\tau), s_\tau\right) \middle| s_t\right].$$

We define the pseudomodel as follows.

DEFINITION S2. *The pseudomodel is a collection of conditional choice probabilities* $\{P_\theta\}_{\theta\in\Theta}$ *such that for all a:*

$$\begin{aligned} P_\theta(a|x_t) &= \Pr[\alpha_\theta(s_t) = a|x_t] \quad a.s., \quad \textit{where} \\ \alpha_\theta(s_t) &= \arg\max_{a\in A}\left\{u_\theta(a, s_t) + \beta E\left[V_\theta(s_{t+1})|x_t, a_t = a\right]\right\}, \\ V_\theta(s_t) &= u_\theta\left(\delta^*(s_t), s_t\right) + \beta E\left[V_\theta(s_{t+1})|s_t\right]. \end{aligned} \tag{7}$$

P_θ describes the distribution of the optimal action that solves a pseudodecision problem, which differs from the true decision problem in (4) by replacing W_θ with V_θ. Particularly the pseudoproblem is a static optimization problem in the sense that V_θ does not incorporate Bellman's principal of optimality. Furthermore, by construction $\alpha_{\theta^*} = \delta^*$, hence $P_{\theta^*} = \mathcal{P}_{\theta^*}$. In what follows we shall simply denote P_{θ^*} by $\mathcal{P}^*$.

For estimation purposes, the main insight of Hotz and Miller (1993) centers on applying an inversion result (see Lemma HM) so that $E[V_\theta(s_{t+1})|x_t = \cdot, a_t = \cdot]$ can be estimated using a nonparametric estimator of $\mathcal{P}^*$ for each θ. Since V_θ is a solution to a linear equation, in contrast W_θ is defined through the max operator, evaluating P_θ numerically is significantly easier than $\mathcal{P}_\theta$ as the former does not require one to solve a dynamic optimization problem. The inherent linear structure that defines V_θ is crucial in our identification analysis since we claim the parameterization of θ in $E[V_\theta(s_{t+1})|x_t = \cdot, a_t = \cdot]$ is much more tractable relative to $E[W_\theta(s_{t+1})|x_t = \cdot, a_t = \cdot]$. Importantly, by the law of iterated expectations, under M1: $E[V_\theta(s_{t+1})|x_t, a_t] = E[E[.V_\theta(s_{t+1})|x_{t+1}]|x_t, a_t]$, where $E[.V_\theta(s_t)|x_t]$ is the solution to a linear equation

$$E[V_\theta(s_t)|x_t] = E[u_\theta(\delta^*(s_t), s_t)|x_t] + \beta E[E[V_\theta(s_{t+1})|x_{t+1}]|x_t],$$

https://doi.org/10.1017/S0266466614000437 Published online by Cambridge University Press

which is obtained by taking conditional expectation of $V_\theta(s_t)$ w.r.t. x_t. Let $m_\theta = E[V_\theta(s_t)|x_t = \cdot]$ and $g_\theta = E[V_\theta(s_{t+1})|x_t = \cdot, a_t = \cdot]$, then we can succinctly summarize their relationships with the per period payoff function π_θ in a linear functional notation as follows:

$$m_\theta = r_\theta + \mathcal{L} m_\theta, \tag{8}$$

$$g_\theta = \mathcal{H} m_\theta, \tag{9}$$

where r_θ is the ex-ante expected payoff given state x_t using δ^*, namely $r_\theta(x) = E\left[u_\theta(\delta^*(s_t), s_t)|x_t = x\right]$ for any $x \in X$; and for any $m : X \to \mathbb{R}^J$ the linear operator $\mathcal{L}$ generates discounted expected next-period value of its operands $\mathcal{L}m(x) = \beta E\left[m(x_{t+1})|x_t = x\right]$ for any $x \in X$; $\mathcal{H}$ is another conditional expectation operator that generates the choice specific expected next period values of its operands operator $\mathcal{H}m(x,a) = E\left[m(x_{t+1})|x_t = x, a_t = a\right]$ for any $a, x \in A \times X$. In contrast to a general solution of a fixed point, the structure of g_θ is more tractable in terms of π_θ. In particular $\mathcal{L}, \mathcal{H}$ are linear operators that do not depend on θ, and $(I - \mathcal{L})^{-1}$ exists (see below), so that $g_\theta = \mathcal{H}(I - \mathcal{L})^{-1} r_\theta$. We now revisit the engine replacement problem introduced earlier and consider the choice probabilities for the pseudomodel.

Example (continued)
The choice probability of replacing the engine with mileage x implied by the pseudodecision problem is:

$$P_\theta(1|x) = \frac{\exp\{-\theta_1 - c(0,\theta_2) + \beta g_\theta(1,x)\}}{\exp\{-\theta_1 - c(0,\theta_2) + \beta g_\theta(1,x)\} + \exp\{-c(x,\theta_2) + \beta g_\theta(0,x)\}},$$

which differs from $\mathcal{P}_\theta(1|x)$ (defined in (5)) by replacing $E[W_\theta(s_{t+1})|x_t = \cdot, a_t = \cdot]$ with g_θ. In order to decompose g_θ, one starts with r_θ and solve for m_θ. Since X is finite, equation (8) can then be represented by the following matrix equation

$$\begin{pmatrix} m_\theta(x^1) \\ \vdots \\ m_\theta(x^J) \end{pmatrix} = \begin{pmatrix} r_\theta(x^1) \\ \vdots \\ r_\theta(x^J) \end{pmatrix} + \begin{pmatrix} l_{11} & \cdots & l_{1,J} \\ \vdots & \ddots & \vdots \\ l_{J,1} & \cdots & l_{J,J} \end{pmatrix} \begin{pmatrix} m_\theta(x^1) \\ \vdots \\ m_\theta(x^J) \end{pmatrix},$$

where $L = \left(l_{ij}\right)$ is a $J \times J$-matrix with $l_{ij} = \beta \Pr\left[x_{t+1} = x^j | x_t = x^i\right]$. Using Hotz and Miller (1993)'s inversion theorem (see Lemma HM), it can be shown that $r_\theta\left(x^j\right) = \gamma^* + \mathcal{P}^*\left(1|x^j\right)\left(-\theta_1 - c(0,\theta_2) - \mathcal{P}^*\left(1|x^j\right)\right) + \mathcal{P}^*\left(0|x^j\right)\left(-c\left(x^j,\theta_2\right) - \log(\mathcal{P}^*\left(0|x^j\right)\right)$, where γ^* is the Euler's constant and $\mathcal{P}^*$ is the choice probabilities induced by δ^*. Since L is a stochastic matrix multiplied by $\beta \in (0,1)$, by the dominant diagonal theorem, the matrix $I - \mathcal{L}$ is invertible and m_θ can be uniquely recovered from $(I - L)^{-1} r_\theta$. Since A is finite, equation (9) can also be represented by a matrix equation. For each a:

$$\begin{pmatrix} g_\theta(a,x^1) \\ \vdots \\ g_\theta(a,x^J) \end{pmatrix} = \begin{pmatrix} h^a_{1,1} & \cdots & h^a_{1,J} \\ \vdots & \ddots & \vdots \\ h^a_{J,1} & \cdots & h^a_{J,J} \end{pmatrix} \begin{pmatrix} m_\theta(x^1) \\ \vdots \\ m_\theta(x^J) \end{pmatrix},$$

https://doi.org/10.1017/S0266466614000437 Published online by Cambridge University Press

where $H^a = (h^a_{ij})$ is a $J \times J$-matrix with $h^a_{ij} = \Pr\left[x_{t+1} = x^j | x_t = x^i, a_t = a\right]$. By letting $H = [H^{1\top}, H^{2\top}]^\top$, then g_θ can be represented by a $2J$-dimensional vector $((H^1 m_\theta)^\top, (H^2 m_\theta)^\top)^\top$.

The example above illustrates that the parameterization in $E[V_\theta(s_{t+1})|x_t = \cdot, a_t = \cdot]$ is generally more tractable than that of $E[W_\theta(s_{t+1})|x_t = \cdot, a_t = \cdot]$, in particular the former can generally be written as some linear combination of elements in π_θ. In the estimation literature, r_θ has been shown to be identified under various modeling assumptions (see the list of two-step estimator papers cited in Section 1), and $\mathcal{L}$ is nonparametrically identified under weak conditions. Even if we relax M4 and allow x_t to be a continuous variable equations (8) and (9) still hold. In that more general case m_θ is the solution to a general type II integral equation, which solves a well-posed inverse problem since the operator $\mathcal{L}$ is a contraction map in commonly used metric spaces.[3] $\mathcal{H}$ is also identified as a conditional expectation of the observables, therefore g_θ is identified for every θ. Numerous estimators for $(r_\theta, m_\theta, g_\theta, \mathcal{L}, \mathcal{H})$ have been proposed in the literature and have also been shown to be uniformly consistent for any compact subset of Θ and $A \times X$. Since $g_\theta = \mathcal{H}(I - \mathcal{L})^{-1} r_\theta$, $\mathcal{H}(I - \mathcal{L})^{-1}$ is a nonparametric object that can typically be represented or approximated by a matrix, objective functions that involve $E[V_\theta(s_{t+1})|x_t = \cdot, a_t = \cdot]$, through (7), are generally much easier to compute than those with $E[W_\theta(s_{t+1})|x_t = \cdot, a_t = \cdot]$ (see (4)), which is the main motivation for using the pseudomodel.

2.3. Scope of Identification

We define identification using the notion of observational equivalence between two parameters in terms of their implied distribution functions. Let $\mathrm{P}_\theta(\cdot|x_t)$ denote any conditional probability function defined on A that is known up to the unknown parameter $\theta \in \Theta$.

DEFINITION I1. *For a collection of conditional choice probabilities* $\{\mathrm{P}_\theta\}_{\theta\in\Theta}$, *two parameter points* θ *and* θ' *are said to be* observationally equivalent *in* $\{\mathrm{P}_\theta\}_{\theta\in\Theta}$ *if* $\mathrm{P}_\theta(a|x_t) = \mathrm{P}_{\theta'}(a|x_t)$ *almost surely for all* a *in* A.

DEFINITION I2. *A collection of conditional choice probabilities* $\{\mathrm{P}_\theta\}_{\theta\in\Theta}$ *is said to be* identified *at* θ_0 *if there is no other* θ *in* Θ *which is observationally equivalent.*

We say that a collection of conditional functions $\{\mathrm{P}_\theta\}_{\theta\in\Theta}$ is identified if it is identified at every point θ in Θ. The definition of identification employed in this paper is chosen to reflect the estimation problem, where most known objective functions in the literature are explicit functionals of the conditional choice probabilities. Compared to related work in the literature, similar to Magnac and Thesmar (2002), our definition of identification takes the primitive (θ) as the *structure* of the model and uses its implied distribution function ($\mathcal{P}_\theta$) as the *reduced form*, whilst Rust (1994, Sect. 3.5) defines the decision rule (δ_θ) as the

https://doi.org/10.1017/S0266466614000437 Published online by Cambridge University Press

reduced form. The two definitions are equivalent under M1–M4 in the sense that for any $\theta, \theta' \in \Theta$:

$$\delta_\theta(s_t) = \delta_{\theta'}(s_t) \text{ a.s.} \Leftrightarrow \mathcal{P}_\theta(s_t) = \mathcal{P}_{\theta'}(s_t) \text{ a.s.}$$

In Section 3 we shall focus on the identification of the pseudomodel. Our first result shows that the identification of $\{P_\theta\}_{\theta\in\Theta}$ has an immediate implication for the identification of $\{\mathcal{P}_\theta\}_{\theta\in\Theta}$.

PROPOSITION 1. *Suppose the pseudomodel $\{P_\theta\}_{\theta\in\Theta}$ is generated from the decision rule δ^* that satisfies (2) with primitives θ^*. If $\{P_\theta\}_{\theta\in\Theta}$ is identified at θ^*, then there can be no other decision rule δ' that satisfies (2) with primitives $\theta' \neq \theta^*$ such that $\delta'(s_t) = \delta^*(s_t)$ almost surely.*

3. IDENTIFICATION

3.1. Equivalent Condition for Identification

Let $v_\theta = \pi_\theta + \beta g_\theta$, following the notations used in (8) and (9), so that for all a, x: $v_\theta(a,x) = \pi_\theta(a,x) + \beta g_\theta(a,x)$. Then we can write choice probabilities of the pseudomodel as follows:

$$P_\theta(a|x) = \Pr\left[v_\theta(a,x_t) + \varepsilon_t(a) \geq v_\theta\left(a', x_t\right) + \varepsilon_t\left(a'\right) \text{ for all } a' \neq a \,\middle|\, x_t = x \right].$$

Define $\Delta v_\theta(a,x)$ to be $v_\theta(a,x) - v_\theta(0,x)$ for all (a,x). Under M1–M4, Hotz and Miller (1993)'s inversion theorem (also see Lemma 8 of Matzkin, 1991) allows $\{\Delta v_\theta(a,x_t)\}_{a>0}$ to be uniquely recovered from $\{P_\theta(a|x_t)\}_{a>0}$. Therefore the identifiability condition reduces to whether the parameterization of θ in Δv_θ is unique on $A \times X$. A version of the recoverability result is stated below for the ease of reference.

For all $a > 0$ we define a map $\phi_{a,x} : \mathbb{R}^K \to [0,1]$, such that for any real K-vector $\widetilde{v} = (\widetilde{v}_1, \ldots, \widetilde{v}_K)$,

$$\phi_{a,x}(\widetilde{v}) = \int_{\varepsilon_a=-\infty}^{\infty} \int_{\varepsilon_0=-\infty}^{\widetilde{v}_a+\varepsilon_a} \int_{\varepsilon_1=-\infty}^{\widetilde{v}_a-\widetilde{v}_1+\varepsilon_a} \cdots \int_{\varepsilon_{a-1}=-\infty}^{\widetilde{v}_a-\widetilde{v}_{a-1}+\varepsilon_a} \int_{\varepsilon_{a+1}=-\infty}^{\widetilde{v}_a-\widetilde{v}_{a+1}+\varepsilon_a}$$

$$\cdots \int_{\varepsilon_K=-\infty}^{\widetilde{v}_a-\widetilde{v}_K+\varepsilon_a} dQ(\varepsilon_0, \ldots, \varepsilon_K|x).$$

By construction $P_\theta(a|x) = \phi_{a,x}(\Delta v_\theta(1,x), \ldots, \Delta v_\theta(K,x))$ for all (a,x). Then define a map $\Phi_x = \left(\phi_{1,x}, \ldots, \phi_{K,x}\right)^\top$ so that $\Phi_x : \mathbb{R}^K \to \Delta^K$, where Δ^K denotes the K-simplex. The following Lemma follows from Proposition 1 of Hotz and Miller (1993).

LEMMA HM. *Suppose Assumptions M1–M4 hold. Then Φ_x is invertible for all $x \in X$.*

https://doi.org/10.1017/S0266466614000437 Published online by Cambridge University Press

The invertibility of Φ_x can be used to recover the indifferent-value vector for any given choice probabilities (and vice-versa). The identification problem of $\{P_\theta\}_{\theta\in\Theta}$ is then equivalent to the question whether Δv_θ is uniquely parameterized on $A \times X$. Since A and X are finite we can represent Δv_θ in a matrix form. The following notations shall be useful for Theorems 1 and 2.

Preliminary Notations for Theorems 1 and 2

For all θ, let $v_\theta^a = (v_\theta(a, x^1), \ldots, v_\theta(a, x^J))$ for all a, and $\mathbf{v}_\theta = (v_\theta^0, \ldots, v_\theta^K)^\top$, define π_θ^a and $\boldsymbol{\pi}_\theta$ analogously. For any k let: I_k denote an identity matrix of size k; $\mathbf{H}$ denote a block-diagonal matrix $\mathbf{diag}(H^0, H^1, \ldots, H^K)$, where H^a denotes a $J \times J$ matrix whose ij-th element is $\Pr[x_{t+1} = x^j | x_t = x^i, a_t = a]$; $\Delta\mathbf{H}$ denote a block-diagonal matrix $\mathbf{diag}(H^1 - H^0, \ldots, H^K - H^0)$; $\mathbf{M}_k = (I_{k+1} \otimes M)$ where $M = (I - L)^{-1}$ and L denotes a $J \times J$ matrix whose ij-th element is $\beta \Pr[x_{t+1} = x^j | x_t = x^i]$; $\mathbf{P}_0 = \begin{bmatrix} P^0 & P^1 & \ldots & P^K \\ \vdots & \vdots & \ldots & \vdots \\ P^0 & P^1 & \ldots & P^K \end{bmatrix}$ where $P^k = diag(\mathcal{P}^*(k|x^1), \ldots, \mathcal{P}^*(k|x^J))$ and $\mathbf{P}_1 = \begin{bmatrix} P^1 & P^2 & \ldots & P^K \\ \vdots & \vdots & \ldots & \vdots \\ P^1 & P^2 & \ldots & P^K \end{bmatrix}$. Define $\Delta v_\theta^a = (v_\theta(a, x^1) - v_\theta(0, x^1), \ldots, v_\theta(a, x^J) - v_\theta(0, x^J))$ for all $a > 0$, and $\Delta\mathbf{v}_\theta = (\Delta v_\theta^1, \ldots, \Delta v_\theta^K)^\top$, define $\Delta\pi_\theta^a$ and $\Delta\boldsymbol{\pi}_\theta$ analogously. Let $\mathbf{D}$ denote the $JK \times J(K+1)$ matrix that performs the transformation $\mathbf{D}\mathbf{v}_\theta = \Delta\mathbf{v}_\theta$. Finally, for any θ, θ', let $\Delta\mathbf{v}_{\theta,\theta'} = \Delta\mathbf{v}_\theta - \Delta\mathbf{v}_{\theta'}$, $\boldsymbol{\pi}_{\theta,\theta'} = \boldsymbol{\pi}_\theta - \boldsymbol{\pi}_{\theta'}$, and $\Delta\boldsymbol{\pi}_{\theta,\theta'} = \Delta\boldsymbol{\pi}_\theta - \Delta\boldsymbol{\pi}_{\theta'}$. Then the following representation Lemmas hold.

LEMMA 1. *Suppose Assumptions M1–M4 hold. Then for all θ, θ':*

$$\Delta\mathbf{v}_{\theta,\theta'} = \mathbf{D}\left(I_{J(K+1)} + \beta\mathbf{H}\mathbf{M}_{K+1}\mathbf{P}_0\right)\boldsymbol{\pi}_{\theta,\theta'}.$$

LEMMA 2. *Suppose Assumptions M1–M4 and R1 (defined below) hold. Then for all θ, θ',*

$$\Delta\mathbf{v}_{\theta,\theta'} = \left(I_{JK} + \beta\,\Delta\mathbf{H}\mathbf{M}_K\mathbf{P}_1\right)\Delta\boldsymbol{\pi}_{\theta,\theta'}.$$

There are two notable features of the representation Lemmas. First, for any θ, θ', the RHS of $\Delta\mathbf{v}_{\theta,\theta'}$ can always be written as a product of some matrix that can be identified from the data and a vector that depends on the structural parameters. This follows directly from the fact that all future discounted expected returns can be written as linear transforms of $\boldsymbol{\pi}_\theta$. The other is that $\Delta\mathbf{v}_{\theta,\theta'}$ only depends on the distribution of ε_t through the equilibrium induced transition and choice probabilities. Therefore conditions for identification consist only of two separate sources: one on the parameterization in $\boldsymbol{\pi}_\theta$ and the other on objects matrices that can be identified from the data.

https://doi.org/10.1017/S0266466614000437 Published online by Cambridge University Press

3.2. Main Results

THEOREM 1. *Suppose Assumptions M1–M4 hold. Then* $\{P_\theta\}_{\theta\in\Theta}$ *is identified at* θ_0 *if and only if the intersection between* $\left\{\boldsymbol{\pi}_{\theta_0,\theta} : \theta \in \Theta\backslash\{\theta_0\}\right\}$ *and the null space of* $\mathbf{D}\left(I_{J(K+1)} + \beta\mathbf{H}\mathbf{M}_{K+1}\mathbf{P}_0\right)$ *is empty.*

It immediately follows from Lemma 1 that the intersection between the null space of $\mathbf{D}(I_{J(K+1)} + \beta\mathbf{H}\mathbf{M}_{K+1}\mathbf{P}_0)$ and $\left\{\boldsymbol{\pi}_{\theta_0,\theta} : \theta \in \Theta\backslash\{\theta_0\}\right\}$ completely characterizes the set of $P_\theta's$ that are observationally equivalent to P_{θ_0}. However, the implication from the characterization may not be immediately obvious since it can be difficult to visualize geometrically. We also note that the possibility for nonidentification is not necessarily trivial if the parameterization of $\boldsymbol{\pi}_\theta$ is too rich since $\mathbf{D}\left(I_{J(K+1)} + \beta\mathbf{H}\mathbf{M}_{K+1}\mathbf{P}_0\right)$ is always noninjective (as it has more columns than rows).[4] Our next two sets of results are easier to verify. Each makes use of an additional restriction on the payoff function commonly found in the literature. The first assumes the knowledge of the payoff from an outside option. The second relies on the linear-in-parameter specification.

Assumption R1. For all θ, x, $\pi_\theta(0, x) = \underline{\pi}(x)$.

THEOREM 2. *Suppose Assumptions M1–M4 and R1 hold, and whenever* $\theta \neq \theta'$, $\Delta\boldsymbol{\pi}_{\theta,\theta'}$ *is not a null vector. Then* $\{P_\theta\}_{\theta\in\Theta}$ *is identified at* θ_0 *if and only if the intersection between* $\left\{\Delta\boldsymbol{\pi}_{\theta_0,\theta} : \theta \in \Theta\backslash\{\theta_0\}\right\}$ *and the null space of* $I_{KJ} + \beta\Delta\mathbf{H}\mathbf{M}_K\mathbf{P}_1$ *is empty.*

The proof follows immediately from Lemma 2. Assumption R1 is commonly used in the nonparametric identification literature, which can also be motivated in particular empirical applications (e.g., in entry games $a_t = 0$ often denotes firms not participating hence getting zero payoff). Although note that, in contrast to static discrete choice models, R1 is not a normalization condition without loss of generality; cf. Theorem 1 (also see Theorems 3 and 4). Whilst the requirement that π_θ is well-parameterized, in the sense that there are no redundant parameters, is obviously a necessary condition for identification. Different specification of π_θ can be checked on a case-by-case basis. Since $I_{KJ} + \beta\Delta\mathbf{H}\mathbf{M}_K\mathbf{P}_1$ is a square matrix, then we have a simple sufficient condition for identification independent of θ.

COROLLARY 1. *Suppose Assumptions M1–M4 and R1 hold, and whenever* $\theta \neq \theta'$ *there exists* a, x *so that* $\Delta\pi_\theta(a, x) - \Delta\pi_{\theta'}(a, x) \neq 0$. *If* $I_{KJ} + \beta\Delta\mathbf{H}\mathbf{M}_K\mathbf{P}_1$ *is nonsingular then* $\{P_\theta\}_{\theta\in\Theta}$ *is identified.*

It is worth noting a feature of our results thus far that the conditions for identification consist of two separate sources: one on the parameterization in $\boldsymbol{\pi}_\theta$ and the other on matrices that can be identified from the data. This follows directly from the fact that all future discounted expected returns can be written as linear transforms of $\boldsymbol{\pi}_\theta$. Next we consider the case when π_θ is linear-in-parameters, where we can revert back to the linear functional notations instead of the cumbersome vectorized notations used for Theorems 1 and 2.

https://doi.org/10.1017/S0266466614000437 Published online by Cambridge University Press

Assumption R2. The payoff function is linear-in-parameters, i.e., $\pi_\theta(a,x) = \theta^\top \pi_0(a,x)$ for all a, x, θ, where π_0 is a p-dimensional vector of real functions.

For Theorems 3 and 4 it shall be useful to represent $\{\Delta v_0(a,x)\}_{a>0}$ using a matrix $\Delta V(x)$, where

$$\Delta V(x) = \begin{bmatrix} \Delta v_0(1,x)^\top \\ \vdots \\ \Delta v_0(K,x)^\top \end{bmatrix} \quad \text{for all } x, \tag{10}$$

where Δv_0 is defined within the following Lemma.

LEMMA 3. *Under Assumptions M1–M4 and R2, for all* θ, θ', $\Delta v_\theta - \Delta v_{\theta'} = (\theta - \theta')^\top \Delta v_0$, *where* $v_0 = \pi_0 + \beta \mathcal{H}(I - \mathcal{L})^{-1} r_0$ *with* $r_0(x) = E[\pi_0(\delta^*(s_t), x_t) | x_t = x]$ *for all* x *and* $\Delta v_0(a,x) = v_0(a,x) - v_0(0,x)$ *for all* $x, a > 0$.

Under the linear-in-parameters specification, Lemma 3 shows that the identification of $\{P_\theta\}_{\theta\in\Theta}$ is analogous to a familiar problem of no multicollinearity condition in a linear model. We give conditions for identification of this special case in the next Theorem.

THEOREM 3. *Under Assumptions M1–M4 and R2, and* $E[\Delta V(x_t)^\top \Delta V(x_t)]$ *exists. Then* $\{P_\theta\}_{\theta\in\Theta}$ *is identified if and only if* $E\left[\Delta V(x_t)^\top \Delta V(x_t)\right]$ *has rank* p.

Given an explicit specification of the parameterization under R2 we can adapt Theorem 3 to incorporate a priori relationships between the elements of θ. Our next theorem gives a weaker set of conditions for identification when the rank condition in Theorem 3 may fail to hold but θ is known to satisfy a set of linear restrictions.

THEOREM 4. *Suppose Assumptions M1–M5 hold, and* $E\left[\Delta V(x_t)^\top \Delta V(x_t)\right]$ *exists. Let* $\Theta_0 = \{\theta \in \Theta : R^\top \theta = \gamma_0\}$ *where, for some* $d < p$, R *is a* $p \times d$*-matrix with full column rank, and* γ *be a* d*-dimensional vector. Then* $\{P_\theta\}_{\theta\in\Theta_0}$ *is identified if and only if there exists a* $p \times (p-d)$*-matrix* D *whose column space equals the orthogonal complement of the column space of* R *such that* $D^\top E\left[\Delta V(x_t)^\top \Delta V(x_t)\right] D$ *has rank* $p-d$.

The requirements on the column space of D in Theorem 4 do not impose any restriction since such matrix always exists (although not unique) as $\mathbb{R}^p$ can always be written as a direct sum between the column space of (any matrix) R and its orthogonal complement. The essence of the proof is the reparameterization of θ to incorporate linear constraints that reduces the dimension of the free parameters; such method is familiar from the classical study of constrained least square problems (for instance, see Amemiya, 1985, Sect. 1.4.2).

Under R2 the identification problem is determined by a system of linear equations. If the rank conditions in Theorems 3 and 4 are satisfied, then the likelihood criterion and any other objective functions, which depends on the choice

https://doi.org/10.1017/S0266466614000437 Published online by Cambridge University Press

probabilities from the pseudomodel, will be able to deliver a consistent estimator of θ^* under standard regularity conditions. Our identification conditions also have finite sample implications. Since the estimator of ΔV is easy to compute, the empirical analogs of $E\left[\Delta V(x_t)^\top \Delta V(x_t)\right]$ and $D^\top E\left[\Delta V(x_t)^\top \Delta V(x_t)\right] D$ can (and should) be analyzed prior the estimation stage. For instance, suppose $\left\{a_n, x_n, x_n'\right\}_{n=1}^N$ is a random sample from $(\delta^*(s_t), x_t, x_{t+1})$ and $\Delta\widehat{V}$ is the estimator of ΔV. If $\frac{1}{N}\sum_{n=1}^N \Delta\widehat{V}(x_n)^\top \Delta\widehat{V}(x_n)$ does not have full rank, then the intersection of the null spaces of $\{\Delta\widehat{V}(x_n)\}_{n=1}^N$ is nonempty. Thus, by applying Lemma HM, for each θ_0 there are uncountably many other θ's that imply the same choice probabilities so that any objective functions constructed from the choice probabilities cannot have a unique optimizer; analogous conclusions can also be drawn from studying the rank of $D^\top\left(\frac{1}{N}\sum_{n=1}^N \Delta\widehat{V}(x_n)^\top \Delta\widehat{V}(x_n)\right) D$.

COMMENTS ON RELAXING M4: Conceptually, we can naturally extend our identification results so that x_t also contains continuously distributed state variable (see Srisuma and Linton, 2012). In particular, for Theorems 1 and 2 and Corollary 1, $\Delta\mathbf{v}_{\theta,\theta'}$ remains a linear transform of $\boldsymbol{\pi}_{\theta,\theta'}$ and $\Delta\boldsymbol{\pi}_{\theta,\theta'}$ for Lemmas 1 and 2, respectively. However, the sets $\left\{\boldsymbol{\pi}_{\theta_0,\theta} : \theta \in \Theta\backslash\{\theta_0\}\right\}$ and $\left\{\Delta\boldsymbol{\pi}_{\theta_0,\theta} : \theta \in \Theta\backslash\{\theta_0\}\right\}$ become subsets of functions in some Banach spaces, and the matrices of conditional choice and transition probabilities are to be replaced by linear integral operators. Although general conditions for injectivity of linear operators are more difficult to be characterized and verified. On the other hand, Theorems 3 and 4 immediately accommodate continuously distributed x_t without any modification.

We end this section by illustrating how our results can be used, based on observables objects, to show the identification of specific parametric models that are otherwise generally intractable.

Example (continued)
We consider parametric models similar to those studied in Rust (1987), generated from the following functional forms: (a) $c(x,\theta_2) = \exp(\theta_2 x)$, and (b) $c(x,\theta_2) = \theta_2 c_0(x)$ for some known $c_0(\cdot)$.[5] We assume that $x^j < x^{j+1}$ and $\theta \in \Theta = \mathbb{R}^{2++}$.

In order to study identification we construct $\Delta v_\theta(a,x)$. Since $K = 1$, the preliminary notations introduced for Lemmas 1 and 2 can be simplified further. In what follows we let $\imath$ denote $J \times 1$ vector of ones, $\mathbf{P} = diag\left(\mathcal{P}^*\left(1|x^1\right), \ldots, \mathcal{P}^*\left(1|x^J\right)\right)$, and $\mathbf{c}(\theta_2) = \left(c\left(x^1,\theta_2\right), \ldots, c\left(x^J,\theta_2\right)\right)^\top$. Note that we can write

$$\pi_\theta(a,x) = [-\theta_1 - c(0,\theta_2)]\cdot a + [-c(x,\theta_2)]\cdot(1-a).$$

Since $\pi_\theta(1,x) - \pi_\theta(0,x) = -\theta_1 - c(0,\theta_2) + c(x,\theta_2)$, it follows that

$$\Delta\boldsymbol{\pi}_\theta = -\theta_1\cdot\imath - c(0,\theta_2)\cdot\imath + \mathbf{c}(\theta_2).$$

https://doi.org/10.1017/S0266466614000437 Published online by Cambridge University Press

The vector of differences in future expected payoffs, $\{g_\theta(1,x^j)-g_\theta(0,x^j)\}_{j=1}^J$, is $\Delta\mathbf{H}\mathbf{M}_1\mathbf{P}_0\boldsymbol{\pi}_\theta$, where $\Delta\mathbf{H},\mathbf{M}_1$ are now square matrices of size J and

$$\mathbf{P}_0\boldsymbol{\pi}_\theta=-\theta_1\cdot\mathbf{P}\boldsymbol{\iota}-c(0,\theta_2)\cdot\mathbf{P}\boldsymbol{\iota}-(I_J-\mathbf{P})\,\mathbf{c}(\theta_2).$$

Then we can write $\Delta\mathbf{v}_\theta=\Delta\boldsymbol{\pi}_\theta+\beta\Delta\mathbf{H}\mathbf{M}_1\mathbf{P}_0\boldsymbol{\pi}_\theta$ as

$$\begin{aligned}\Delta\mathbf{v}_\theta=&-\theta_1\cdot(I_J+\beta\Delta\mathbf{H}\mathbf{M}_1\mathbf{P})\,\iota-c(0,\theta_2)\cdot(I_J+\beta\Delta\mathbf{H}\mathbf{M}_1\mathbf{P})\,\iota\\&+(I_J-\beta\Delta\mathbf{H}\mathbf{M}_1(I_J-\mathbf{P}))\,\mathbf{c}(\theta_2).\end{aligned}$$

Next we impose specific functional form for $c(x,\theta_2)$.

Case (a): $c(x,\theta_2)=\exp(\theta_2 x)$. For any $\theta,\theta'\in\Theta$, we have

$$\Delta\mathbf{v}_{\theta,\theta'}=-(\theta_1-\theta_1')\cdot(I_J+\beta\Delta\mathbf{H}\mathbf{M}_1\mathbf{P})\iota+(I_J-\beta\Delta\mathbf{H}\mathbf{M}_1(I_J-\mathbf{P}))(\mathbf{c}(\theta_2)-\mathbf{c}(\theta_2')).$$

Note that (a) does not satisfy either R1 or R2. Therefore the condition for identification of $\{P_\theta\}_{\theta\in\Theta}$ at any point θ_0, as characterized by Theorem 1, is equivalent to whether there exists any other $\theta\in\Theta$ that $\Delta\mathbf{v}_{\theta_0,\theta}=\mathbf{0}$, which has the geometric interpretation that θ_0 is identifiable if the one-dimensional curve $\theta_2\mapsto(I_J-\beta\Delta\mathbf{H}\mathbf{M}_1(I_J-\mathbf{P}))(\mathbf{c}(\theta_{02})-\mathbf{c}(\theta_2))$ does not intersect the line spanned by $(I_J+\beta\Delta\mathbf{H}\mathbf{M}_1\mathbf{P})\iota$ in $\mathbb{R}^J$. If we impose R1, for instance by taking θ_{01} to be known, then $\Delta\mathbf{v}_{\theta,\theta'}=(I_J-\beta\Delta\mathbf{H}\mathbf{M}_1(I_J-\mathbf{P}))\left(\mathbf{c}(\theta_2)-\mathbf{c}\left(\theta_2'\right)\right)$. Then, a sufficient condition for identification, following from Corollary 1, is for $I_J-\beta\Delta\mathbf{H}\mathbf{M}_1(I_J-\mathbf{P})$ to be nonsingular since $\mathbf{c}(\theta_2)\neq\mathbf{c}(\theta_{02})$ for any $\theta_2\neq\theta_{02}$.[6] Although such condition is not necessary. Even if $I_J-\beta\Delta\mathbf{H}\mathbf{M}_1(I_J-\mathbf{P})$ is rank deficient, $\{P_\theta\}_{\theta\in\Theta}$ may still be identified. In particular we can relax the requirement on the null space of $I_J-\beta\Delta\mathbf{H}\mathbf{M}_1(I_J-\mathbf{P})$ by using information on the structure of $c(x,\theta_2)$. For example, given that $\exp\left(\theta_2x^{j+1}\right)-\exp\left(\theta_{02}x^{j+1}\right)>\exp\left(\theta_2x^j\right)-\exp\left(\theta_{02}x^j\right)$ for any $\theta_2>\theta_{02}$ and $j=1,2$,[7] $\{P_\theta\}_{\theta\in\Theta}$ will also be identified even if the null space of $I_J-\beta\Delta\mathbf{H}\mathbf{M}_1(I_J-\mathbf{P})$ is nonempty but also does not contain vectors of the form $\mathbf{d}=(d_1,d_2,d_3)^\top$ where $d_3>d_2>d_1>0$.

Case (b): $c(x,\theta_2)=\theta_2c_0(x)$. For any $\theta,\theta'\in\Theta$, we have

$$\Delta\mathbf{v}_{\theta,\theta'}=-\left(\theta_1-\theta_1'\right)\cdot(I_J+\beta\Delta\mathbf{H}\mathbf{M}_1\mathbf{P})\,\iota+\left(\theta_2-\theta_2'\right)\cdot(I_J-\beta\Delta\mathbf{H}\mathbf{M}_1(I_J-\mathbf{P}))\,\mathbf{c}_0.$$

Let $\Delta\mathbf{V}=[(I_J+\beta\Delta\mathbf{H}\mathbf{M}_1\mathbf{P})\,\iota:(I_J-\beta\Delta\mathbf{H}\mathbf{M}_1(I_J-\mathbf{P}))\,\mathbf{c}_0]$, then it is clear that the sufficient and necessary condition for identification is the full column rank of $\Delta\mathbf{V}$, which is equivalent to the condition in Theorem 3. In particular, $\Delta v_0(1,x)=\left(\Delta v_0^1(1,x),\Delta v_0^2(1,x)\right)^\top$ (cf. Lemma 3), so that $(I_J+\beta\Delta\mathbf{H}\mathbf{M}_1\mathbf{P})\,\iota$ and $(I_J-\beta\Delta\mathbf{H}\mathbf{M}_1(I_J-\mathbf{P}))\,\mathbf{c}_0$ are $\left(\Delta v_0^1\left(1,x^1\right),\ldots,\Delta v_0^1\left(1,x^J\right)\right)^\top$ and $\left(\Delta v_0^2\left(1,x^1\right),\ldots,\Delta v_0^2\left(1,x^J\right)\right)^\top$, respectively.

https://doi.org/10.1017/S0266466614000437 Published online by Cambridge University Press

4. DYNAMIC GAMES

Recently there has been a growing interest in extending econometric methods developed for single agent models to estimate empirical games, especially in the empirical industrial organization literature. However, an additional difficulty that arises in models with strategic interactions is the problem of multiple equilibria. Specifically, for every θ, the game may have more than one (if any) equilibrium that predicts different distributions of optimal actions. But the pseudomodel generally does not suffer from problems of multiple equilibria when players are known to play pure strategies almost surely. Our results in the previous section can be directly applied to the dynamic games studied in Aguirregabiria and Mira (2007), Bajari et al. (2007), Pakes et al. (2007), and Pesendorfer and Schmidt-Dengler (2008). To save space and avoid repetition we do not list out various conditions and theorems for the game counterparts to those of single agent problems.

5. CONCLUDING REMARKS

Existing identification results for nonlinear parametric models can be hard to verify the identification of Markov decision models and dynamic games. We show the identification of the pseudomodel is generally simpler and can also be used to show identification of the data generating parameter from the underlying model. Therefore our results are useful for two-step estimators, and other estimation methods where the objective functions are analytically less tractable but need to assume the model is identified for consistent estimation (such as the nested-fixed point algorithm of Rust, 1987, the iterative versions of estimators in Aguirregabiria and Mira, 2002, 2007, and the constrained optimization method proposed by Egesdal, Lai, and Su, 2013). We derive our identification results directly by studying the implied expected payoffs rather than choice probabilities, since the latter is unnecessarily complicated by a further one-to-one, but nonlinear, transformation (Hotz and Miller, 1993). Relatedly the structural parameters can also be estimated directly based on the implied expected payoffs instead of the traditional choice probabilities; Sanches, Silva, and Srisuma (2013) recently show there are computational advantages in doing so without any theoretical cost (or gain).

The linear structure that defines the policy value functions is a general property that is also present in other Markov decision problems and dynamic games outside the discrete choice framework. For instance, a closely related class of decision problems imposes increasing differences conditions on the payoff functions instead of additive separability, see Hong and Shum (2010) in a single agent setting, and Bajari et al. (2007) and Srisuma (2013) for games. However, the lack of separability means we can no longer apply Hotz and Miller's inversion lemma, the parameterization of for the identification of the pseudomodel for these problems are therefore less tractable than the ones considered in this paper. Nevertheless the general local identification results in Wegge (1965) and Rothenberg (1971) can be applied for any pseudomodels, which are also sufficient for the local identification of the actual model for the same reasons as implied by Proposition 1.[8]

https://doi.org/10.1017/S0266466614000437 Published online by Cambridge University Press

NOTES

1. There are some notable exceptions when an outcome variable is also available, which we do not assume, e.g., see Heckman and Navarro (2007).

2. Other important results on the identification of parameters that may enter a system of equations nonlinearly without specifying parametric distributional assumptions can be found in Fisher (1961, 1965, 1966), Brown (1983), Roehrig (1988), Benkard and Berry (2006), and Komunjer (2012).

3. Suppose X is some compact subset of $\mathbb{R}^L$ and B be a space of bounded real-valued functions defined on X. Consider a Banach space $(B, \|\cdot\|)$ equipped with the sup-norm, i.e., $\|\phi\| = \sup_{x\in X} |\phi(x)|$ for any $\phi \in B$. For any $x \in X$, $\mathcal{L}\phi(x) = \beta E[\phi(x_{t+1})|x_t = x]$, then it follows that $|\mathcal{L}\phi(x)| \leq \beta \sup_{x\in X} |\phi(x)|$. In other words $\|\mathcal{L}\phi\| \leq \beta\|\phi\|$, hence the operator norm $\|\mathcal{L}\|$ is bounded above by β. Since $\beta \in (0,1)$, $\mathcal{L}$ is a contraction. Therefore the inverse of $I - \mathcal{L}$ exists. Furthermore, it is a linear bounded operator and admits a Neumann series representation, $\sum_{\tau=0}^{\infty} \mathcal{L}^{\tau}$ (see Kreyszig, 1989).

4. This is familiar from the literature on nonparametric identification, in which case $\{\boldsymbol{\pi}_{\theta_0,\theta} : \theta \in \Theta \backslash \{\theta_0\}\}$ can be thought of as an arbitrary set of points around (but excluding) the origin in $R^{J(K+1)}$; see the nonidentification result in Magnac and Thesmar (2002, Prop. 2) in a two-period model for a similar intuition.

5. We assume the specifications in (a) and (b) for notational simplicity. Rust (1987, p. 1015) considers four functional forms, (i) to (iv), for $c(x,\theta)$. Rust's specification (ii) only differs from our (a) by an additional multiplicative (scaling) parameter. Rust's specifications (i), (iii), and (iv) also satisfy R2, where the latter two are special cases of our (b).

6. This condition is equivalent to $\Delta\mathbf{HM}_1$ ($I_J - \mathbf{P}$) does not contain β^{-1} as an eigenvalue.

7. $\exp(\theta_2 x)$ is a strictly supermodular function.

8. Establishing local identification can be informative in this literature since several estimators have been proposed to estimate the same (pseudo)model, and it is sometimes unclear whether the failure to identify the parameter of interests is a property of the model or a consequence of a particular methodology (see Srisuma, 2013).

REFERENCES

Aguirregabiria, V. & P. Mira (2002) Swapping nested fixed point algorithm: A class of estimators for discrete Markov decision models. *Econometrica* 70, 1519–1543.

Aguirregabiria, V. & P. Mira (2007) Sequential estimation of dynamic discrete games. *Econometrica* 75, 1–53.

Aguirregabiria, V. & P. Mira (2010) Dynamic discrete choice structural models: A survey. *Journal of Econometrics* 156, 38–67.

Amemiya, T. (1985) *Advanced Econometrics*. Harvard University Press.

Bajari, P., C.L. Benkard, & J. Levin (2007) Estimating dynamic models of imperfect competition. *Econometrica* 75, 1331–1370.

Bajari, P., V. Chernozhukov, H. Hong, & D. Nekipelov (2012) Semiparametric Estimation of a Dynamic Game of Incomplete Information. Working paper, University of Minnesota.

Bajari, P., H. Hong, & D. Nekipelov (2013) Econometrics for Game Theory. *Advances in Economics and Econometrics: Theory and Applications, 10th World Congress*.

Benkard, C.L. & S. Berry (2006) On the nonparametric identification of nonlinear simultaneous equations models: Comment on Brown (1983) and Roehrig (1988). *Econometrica* 74, 1429–1440.

Brown, B.W. (1983) The identification problem in systems nonlinear in the variables. *Econometrica* 51, 175–196.

Egesdal, M., Z. Lai, & C. Su (2013) Estimating Dynamic Discrete-Choice Games of Incomplete Information. Working paper, University of Chicago Booth School of Business.

Fisher, F.M. (1961) Identifiability criteria in nonlinear systems. *Econometrica* 29, 574–590.

Fisher, F.M. (1965) Identifiability criteria in nonlinear systems: A further note. *Econometrica* 33, 197–205.

https://doi.org/10.1017/S0266466614000437 Published online by Cambridge University Press

Fisher, F.M. (1966) *The Identification Problem in Economics*. McGraw-Hill.
Heckman, J. & S. Navarro (2007) Dynamic discrete choice and dynamic treatment effects. *Journal of Econometrics* 136, 341–396.
Hong, H. & M. Shum (2010) Pairwise-difference estimation of a dynamic optimization model. *Review of Economic Studies* 77, 273–304.
Hotz, V. & R.A. Miller (1993) Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies* 60, 497–531.
Hotz, V., R.A. Miller, S. Sanders, & J. Smith (1994) A simulation estimator for dynamic models of discrete choice. *Review of Economic Studies* 61, 265–289.
Komunjer, I. (2012) Global identification in nonlinear models with moment restrictions. *Econometric Theory* 28, 719–729.
Kreyszig, E. (1989) *Introduction to Functional Analysis with Applications*. Wiley Classics Library.
Magnac, M. & D. Thesmar (2002) Identifying dynamic discrete decision processes. *Econometrica* 70, 801–816.
Maskin, E. & J. Tirole (2001) Markov perfect equilibrium I: Observable actions. *Journal of Economic Theory* 100, 191–219.
Matzkin, R.L. (1991) Semiparametric estimation of monotone and concave utility functions for polychotomous choice models. *Econometrica* 59, 1315–1327.
Pakes, A., M. Ostrovsky, & S. Berry (2007) Simple estimators for the parameters of discrete dynamic games (with entry/exit example). *RAND Journal of Economics* 38, 373–399.
Pesendorfer, M. & P. Schmidt-Dengler (2008) Asymptotic least squares estimator for dynamic games. *Reviews of Economic Studies* 75, 901–928.
Roehrig, C.S. (1988) Conditions for identification in nonparametric and parametric models. *Econometrica* 56, 433–447.
Rothenberg, T.J. (1971) Identification in parametric models. *Econometrica* 39, 577–591.
Rust, J. (1987) Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica* 55, 999–1033.
Rust, J. (1994) Structural estimation of Markov decision processes. In R.F. Engle & D. McFadden (eds.), *Handbook of Econometrics*, Vol. 4. North-Holland.
Rust, J. (1996) Numerical dynamic programming in economics. In H.M. Aumann, D.A. Kendrick, & J. Rust (eds.), *Handbook of Computational Economics*, Vol. 1. Elsevier.
Sanches F., D. Silva, & S. Srisuma (2013) An Alternative Asymptotic Least Squares Estimator for Dynamic Games.
Srisuma, S. (2013) Minimum distance estimators for dynamic games. *Quantitative Economics* 4, 549–583.
Srisuma, S. & O.B. Linton (2012) Semiparametric estimation of Markov decision processes with continuous state space. *Journal of Econometrics* 166, 320–341.
Wegge, L. (1965) Identifiability criteria for a system of equations as a whole. *Australian Journal of Statistics* 7, 67–77.

APPENDIX

Proof of Proposition

Proof of Proposition 1. Suppose there exist $\theta' \neq \theta^*$ such that $\delta'(s_t) = \delta^*(s_t)$ almost surely. Then $V_{\theta'}$ must equal $W_{\theta'}$ by the definitions of the value functions and policy value functions. However, this implies $\alpha_{\theta'}(s_t) = \alpha_{\theta^*}(s_t)$ almost surely, contradicting the assumption that the pseudomodel is identified. ■

https://doi.org/10.1017/S0266466614000437 Published online by Cambridge University Press

Proofs of Lemmas

Proof of Lemma 1. Immediate. ■

Proof of Lemma 2. Immediate. ■

Proof of Lemma 3. Under the linear-in-parameters specification, $r_\theta = \theta^\top r_0 + \underline{r}$, where $r_0 = E[\pi_0(\delta^*(s_t), x_t)|x_t = \cdot]$ and $\underline{r} = E\left[\sum_{a\in A}\varepsilon_t(a)\mathbf{1}[\delta^*(s_t) = a]|x_t = \cdot\right]$. Since $(I-\mathcal{L})^{-1}$ and $\mathcal{H}$ are linear operators, v_θ can be written as $\theta^\top v_0 + v_1$ such that $v_0 = \pi_0 + \beta g_0$, where $g_0 = \mathcal{H}(I-\mathcal{L})^{-1} r_0$, and $v_1 = \beta\mathcal{H}(I-\mathcal{L})^{-1} r_1$. For all $x, a > 0$ and $j = 0, 1$, let $\Delta v_\theta(a,x)$ denote $v_\theta(a,x) - v_\theta(0,x)$ and $\Delta v_j(a,x)$ denote $v_j(a,x) - v_j(0,x)$ so that $\Delta v_\theta(a,x) = \theta^\top \Delta v_0(a,x) + \Delta v_1(a,x)$. ■

Proofs of Theorems

Proof of Theorem 1. Immediate from Lemma 1. ■

Proof of Theorem 2. Immediate from Lemma 2. ■

Proof of Theorem 3. Suppose $\{P_\theta\}_{\theta\in\Theta}$ is identified. If $E[\Delta V(x_t)^\top \Delta V(x_t)]$ has rank less than p, then there exists $\theta \neq \theta'$ such that $\Delta V(x_t)\theta = \Delta V(x_t)\theta'$ almost surely, which is equivalent to $\Delta v_\theta(a, x_t) = \Delta v_{\theta'}(a, x_t)$ almost surely for all $a > 0$ since $v_1^a(x_t)$ does not depend on θ. Lemma HM then implies $P_\theta(a|x_t) = P_{\theta'}(a|x_t)$ almost surely for all $a \in A$, which leads to a contradiction. For the converse, suppose $\{P_\theta\}_{\theta\in\Theta}$ is not identified. Then, using Lemma HM, there exists $\theta \neq \theta'$ such that $\Delta v_\theta(a, x_t) = \Delta v_{\theta'}(a, x_t)$ almost surely for all $a > 0$, which is equivalent to $\Delta V(x_t)\theta = \Delta V(x_t)\theta'$ almost surely. However, this contradicts the full rank condition of $E[\Delta V(x_t)^\top \Delta V(x_t)]$. ■

Proof of Theorem 4. Let $B = [R, D]^\top$. By construction B is nonsingular since $D^\top R = 0$. For any θ,

$$\begin{aligned}\Delta V(x_t)\theta &= \Delta V(x_t) B^{-1} B\theta \\ &= \Delta V(x_t)\left[R(R^\top R)^{-1}, D(D^\top D)^{-1}\right]\begin{bmatrix} R^\top \\ D^\top \end{bmatrix}\theta \\ &= \Delta V(x_t) R(R^\top R)^{-1}\gamma_1 + \Delta V(x_t) D(D^\top D)^{-1}\gamma_2,\end{aligned}$$

where $\gamma_1 = R^\top\theta$ and $\gamma_2 = D^\top\theta$. Suppose $\{P_\theta\}_{\theta\in\Theta_0}$ is identified. If $D^\top E[\Delta V(x_t)^\top \Delta V(x_t)]D$ has rank less than $p - d$, then there exists $\gamma_2 \neq \gamma_2'$ such that $\Delta V(x_t) D(D^\top D)^{-1}\gamma_2 = \Delta V(x_t) D(D^\top D)^{-1}\gamma_2'$ almost surely. Since B is nonsingular there exists $\theta \neq \theta'$ in Θ_0 such that $D^\top\theta = \gamma_2$ and $D^\top\theta' = \gamma_2'$, so that $\Delta V(x_t)\theta = \Delta V(x_t)\theta'$ almost surely. However, by Lemma HM, this leads to a contradiction. For the converse, using Lemma HM, suppose there exists $\theta \neq \theta'$ in Θ_0 such that $\Delta V(x_t)\theta = \Delta V(x_t)\theta'$ almost surely. Then $\Delta V(x_t)\theta - \Delta V(x_t)\theta' = \Delta V(x_t) D(D^\top D)^{-1}\eta$ almost surely, where $\eta = D^\top\theta - D^\top\theta'$ is a null-vector (since B is nonsingular). However, this contradicts the full rank condition of $D^\top E[\Delta V(x_t)^\top \Delta V(x_t)]D$. ■

https://doi.org/10.1017/S0266466614000437 Published online by Cambridge University Press