

Komarova, Tatiana; Sanches, Fabio; Junior, Daniel Silva; Srisuma, Sorawoot

Article

Joint analysis of the discount factor and payoff parameters in dynamic discrete choice models

Quantitative Economics

Provided in Cooperation with:

The Econometric Society

Suggested Citation: Komarova, Tatiana; Sanches, Fabio; Junior, Daniel Silva; Srisuma, Sorawoot (2018) : Joint analysis of the discount factor and payoff parameters in dynamic discrete choice models, Quantitative Economics, ISSN 1759-7331, The Econometric Society, New Haven, CT, Vol. 9, Iss. 3, pp. 1153-1194, <https://doi.org/10.3982/QE675>

This Version is available at:

<https://hdl.handle.net/10419/217125>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



<https://creativecommons.org/licenses/by-nc/4.0/>

Joint analysis of the discount factor and payoff parameters in dynamic discrete choice models

TATIANA KOMAROVA

Department of Economics, London School of Economics

FABIO SANCHES

Department of Economics, PUC-Rio

DANIEL SILVA JUNIOR

Department of Economics, City University of London

SORAWOOT SRISUMA

School of Economics, University of Surrey

Most empirical and theoretical econometric studies of dynamic discrete choice models assume the discount factor to be known. We show the knowledge of the discount factor is not necessary to identify parts, or even all, of the payoff function. We show the discount factor can be generically identified jointly with the payoff parameters. On the other hand, it is known the payoff function cannot be nonparametrically identified without any a priori restrictions. Our identification of the discount factor is robust to any normalization choice on the payoff parameters. In IO applications, normalizations are usually made on switching costs, such as entry costs and scrap values. We also show that switching costs can be nonparametrically identified, in closed-form, independently of the discount factor and other parts of the payoff function. Our identification strategies are constructive. They lead to easy to compute estimands that are global solutions. We illustrate with a Monte Carlo study and the dataset used in [Ryan \(2012\)](#).

KEYWORDS. Discount factor, dynamic discrete choice problem, identification, estimation, switching costs.

JEL CLASSIFICATION. C14, C25, C51.

Tatiana Komarova: t.komarova@lse.ac.uk

Fabio Sanches: fmiessi@gmail.com

Daniel Silva Junior: d.silva-junior@lse.ac.uk

Sorawoot Srisuma: s.srisuma@surrey.ac.uk

We are grateful to the Editor and three anonymous referees for their suggestions that led to many improvements on the paper. We also thank Kirill Evdokimov, Emmanuel Guerre, Koen Jochmans, Arthur Lewbel, Oliver Linton, Aureo de Paula, Martin Pesendorfer, Yuya Sasaki, Richard Spady, Pasquale Schiraldi, and seminar participants at CREATES, GRIPS, Johns Hopkins University, LSE, NUS, Queen Mary University of London, Stony Brook University, University of Cambridge, University of São Paulo, and participants at Bristol Econometrics Study Group (2014), AMES (Taipei), EMES (Toulouse) and IAAE (London) for comments and discussions.

1. INTRODUCTION

The stationary dynamic discrete decision model surveyed in [Rust \(1994\)](#) has been a subject of much research in econometric theory and empirical studies. The primitives of the model consist of the period payoff function, Markov transition law, and discount factor. A well-known characteristic of a dynamic decision model is that it is not identified. For example, [Manski \(1993\)](#) pointed out in general that the discount factor and payoff function cannot be jointly identified nonparametrically. Most positive identification results in the literature until recently focus on identifying payoff parameters while assuming other primitives to be known; for example, see [Magnac and Thesmar \(2002\)](#), and also [Pesendorfer and Schmidt-Dengler \(2008\)](#) and [Bajari, Chernozhukov, Hong, and Nekipelov \(2009\)](#). Meanwhile, empirical studies typically parameterize the payoff function, parameterize at least part of the distribution of the unobserved variables, and assume the discount factor to be known.

In this paper, we are interested in identifying the discount factor jointly with the payoff function under the linear-in-parameter specification. This parametric model is the most commonly used specification in practice. When there are finite states, the linear specification can represent any nonparametric function. Most empirical studies assume the value of the discount factor to be known without any formal justification in this setting. To the best of our knowledge, we are not aware of any prior identification study involving the discount factor in a general parametric model. We provide conditions under which both the discount factor and payoff parameters can be identified, and propose an easy to compute estimator for them. Other positive identification results on the discount factor in the literature use a nonparametric approach. They use exclusion restrictions in the form of variables affecting future utilities but not current utilities to identify the discount factor; for example, see [Dubé, Hitsch, and Jindal \(2014\)](#), [Wang \(2014\)](#), [Fang and Wang \(2015\)](#), and [Ching and Osborne \(2017\)](#). We do not rely on these assumptions.

A nonparametric payoff function without any restriction cannot be identified even if the discount factor is known. The fundamental identification characteristic in a discrete choice model can be traced to the static random utility model of [McFadden \(1974\)](#), where utility is ordinal and its level cannot be identified. Some form of normalization has to be made. [Aguirregabiria and Suzuki \(2014, AS hereafter\)](#) recently highlighted the undesirable effects that an arbitrary normalization can have on unnormalized parameters and counterfactual studies, and emphasize the importance of identifiable objects without any normalization; also see [Kalouptsi, Scott, and Souza-Rodrigues \(2016, 2017\)](#). An important question then is whether our identification result is robust against misspecifying the normalization choice.

We verify that our identification of the discount factor is robust against any normalization choice. On the other hand the payoff parameters are generally not individually robust. But some of their meaningful combinations are. To this end, we also contribute to the literature by providing a nonparametric framework to identify the payoff parameters that arise from changing in the actions of players between time periods. We call

these *switching costs*.¹ For example, in an entry/exit model, they are entry cost and scrap value. Individually, the *entry cost* and *scrap value* cannot be separately identified but their difference, namely the *sunk entry cost*, can be identified. We show that switching costs can be written explicitly in terms of the observed choice probabilities, independently of the discount factor as well as other (nonswitching costs) components of the payoff function. AS has already shown the sunk entry costs in several IO models can be identified in this fashion. We extend these results to sunk investment costs that can arise from firm investing and divesting, as well as individual switching costs themselves under other a priori restrictions.

A general discussion on the nonidentification of the dynamic model we consider can be found in Rust (1994). Positive identification is possible when more structures are imposed on the primitives. Magnac and Thesmar (2002) have shown the problem of identifying the payoff parameters nonparametrically when all other primitives of the model are assumed to be known can be reduced to a study of solutions to a linear system; also see Pesendorfer and Schmidt-Dengler (2008) and Bajari et al. (2009). We are interested in the payoff parameters as well as the discount factor. The discount factor enters the decision problem recursively and thereby introduces nonlinearity in the model.

Magnac and Thesmar (2002, Section 4.2) suggested that exclusion or parametric restrictions can be used to identify the discount factor. For the former, their Proposition 4 illustrates in a simple two-period model the discount factor is in fact typically overidentified. The identifying restriction they use is that: for some states, utilities in the first period are the same but differ in the second period. This idea has been elaborated and applied in different empirical contexts by Dubé, Hitsch, and Jindal (2014), Wang (2014), Fang and Wang (2015), and Ching and Osborne (2017) among others. On the other hand, while it may be plausible to assume identification is possible in a parametric model we are not aware of any theoretical result that has verified this to be true. In particular, establishing parametric identification in a general nonlinear model is a nontrivial task; see Komunjer (2012) for a recent illustration. We prove identification using an empirical model that is linear in the payoff parameters conditioning on the discount factor. We construct a one-dimensional criterion function to be used for identification. It exploits the conditional linear structure to profile out the payoff parameters and reduce the nonlinear nature of the problem to just one dimension. The criterion function we construct to establish identification has a sample counterpart that can be used for estimation.

In many IO applications, switching costs are often the essence of a dynamic decision problem and can even be the central object of the dynamic model itself (e.g., see Slade (1998), and also the general discussions in Akerberg et al. (2007) and Pesendorfer (2013)). Our study on the switching costs takes a nonparametric approach. We identify combinations of the switching costs by exploiting empirically motivated exclusion and testable independence assumptions. A key step involves eliminating common future expected discounted payoffs that arise from different states. Our result does not depend on the discount factor and some other components of the payoff function. The

¹We use the term switching costs that shares the same spirit as generic adjustment costs and other inertia. Examples of usages in various fields of economics and marketing include the cost to change in health insurance plan, changing of credit and other utility providers, and retailer's decisions on sales promotions.

robust identification result of this nature has precedence in the literature but has not been highlighted.² For example, an inspection of Proposition 2 in [Aguirregabiria and Suzuki \(2014\)](#) will reveal that the same implication of our Theorem 2 has already been obtained for a binary action game of entry/exit.³ We provide closed-form expressions for switching costs and their combinations in terms of only the observed choice probabilities. They can therefore be trivially estimated. They also suggest overidentification tests can be constructed by comparing against other estimates of switching costs obtained under additional assumptions on the model primitives.

Throughout the paper, our identification results are obtained using an empirical model under the assumption that the choice and transition probabilities are nonparametrically identified. These same probabilities are used to compute expected payoffs in a *pseudo-decision problem* for all values of the model parameters as opposed to the *actual (or full-solution) model* where equilibrium probabilities are used. The choice probabilities implied by our empirical model can be used to construct pseudo-likelihood functions as done in [Aguirregabiria and Mira \(2002\)](#) and [Kasahara and Shimotsu \(2008\)](#). This empirical model is used because it is tractable. It forms the basis for any *two-step* estimation procedures, following [Hotz and Miller \(1993\)](#), which are preferred on computational grounds over a full-solution approach such as the nested fixed-point algorithm of [Rust \(1987\)](#). The estimator we propose in this paper will be based on the two-step approach of [Sanches, Silva, and Srisuma \(2016\)](#) with computational simplicity in mind. It is worth noting that, although consistent, a simple two-step estimator like ours tend to have larger finite sample bias and is less efficient than estimators that enforce the equilibrium restriction of the model. Equilibrium constraints can be imposed during estimation with additional computational cost, also without the need to solve out a dynamic optimization problem (cf. [Rust \(1987\)](#)). For example, [Aguirregabiria and Mira \(2002\)](#) and [Egedal, Lai, and Su \(2015\)](#) have shown the fully efficient maximum likelihood estimator can be obtained in this way.

When the data come from a single time series, or when they are pooled across short panels of multiple homogeneous markets, the choice and transition probabilities are nonparametrically identified under weak conditions. In practice, many datasets are short panels, where it would be more reasonable to assume some form of unobserved heterogeneity exists across markets. A flexible yet tractable way to model unobserved heterogeneity in this literature is to use a finite mixture model. For example, [Aguirregabiria and Mira \(2007\)](#) suggested economic agents' payoffs have a time-invariant unobserved market specific component that is unobserved to the econometrician, therefore, markets of different types have different equilibrium distributions on the observables. [Kasahara and Shimotsu \(2009\)](#) and [Arcidiacono and Miller \(2011\)](#) have given conditions so that the probabilities for each mixture type can be nonparametrically identified under different frameworks, thereby extending the scope of applying

²In one instance, for a slightly different model with a mixed continuous-discrete decision variable, [Hong and Shum \(2010\)](#) rely on a deterministic state transition rule to define a *pairwise-difference* estimator that matches on (and thereby avoid computing) future expected discounted payoffs from different states.

³We thank an anonymous referee for pointing this out to us.

two-step estimation methods to models with unobserved heterogeneity. All identification results in our paper are valid in such setting as long as we can identify the type specific probabilities to be able to set up the corresponding pseudo-decision problem. Specifically, the degree of overidentification on the model primitives increases proportionally to the number of mixture types.

The class of decision problems we consider is a special case of dynamic games described in Aguirregabiria and Nevo (2013) and Bajari, Hong, and Nekipelov (2013). All of our intuition and results are applicable to these games. The majority of this paper focuses on the single agent model for notational simplicity and clarity of idea, and to abstract ourselves away from game specific issues (such as multiple equilibria). For the same reasoning given for models with unobserved heterogeneity, the portability of our results to dynamic games is immediate as long as the choice and transition probabilities can be consistently estimated nonparametrically. The numerical studies of our proposed estimators are in fact performed in a dynamic game setting. The details on extending our single agent's results to games can be found in the [Appendix](#).

We perform a Monte Carlo study of our proposed estimators using the simulation design in Pesendorfer and Schmidt-Dengler (2008). We then use the same dataset as used in Ryan (2012) to estimate a dynamic game played between firms in the US Portland cement industry. In our version of the game, firms choose whether to enter the market as well as decide on the capacity level of operation (five different levels). We assume firms compete in a capacity constrained Cournot game, so the period variable profit can be estimated directly from the data as done in Ryan. The dynamic parameters we estimate are the discount factor, fixed operating cost, and 25 switching cost parameters. We estimate the model twice. Once using the data from before 1990 and once after 1990. The separation date coincides with implementation of the 1990 Clean Air Act Amendments (1990 CAAA). Our estimates on switching costs generally appear sensible, having correct signs and relative magnitudes. They show that firms entering the market with a higher capacity level incur larger costs, and suggest that increasing capacity level is generally costly while a reduction can return some revenue. We find that operating and entry costs are generally higher after the 1990 CAAA, which supports Ryan's key finding. We are also able to estimate the discount factor to be within the commonly assumed range with a reasonable precision.

The remainder of the paper is organized as follows. Section 2 introduces the theoretical model and the basic modeling assumptions. Section 3 gives a joint identification result on the discount factor and the payoff parameters under the linear-in-parameter specification. Section 4 studies nonparametric identification of the switching costs. Section 5 illustrates the performance and use of our estimator with simulated and real data. Section 6 concludes. The [Appendix](#) contains details for extending our identification results to dynamic games and further results on identifying the discount factor.

NOTATION. We use $\rho(\mathbf{A})$, $CS(\mathbf{A})$, \mathbf{A}^\top , \mathbf{A}^{-1} , and \mathbf{A}^\dagger to respectively denote the rank, column space, transpose, inverse, and Moore–Penrose inverse of matrix \mathbf{A} . For any positive integers p , q , we let \mathbf{I}_p and $\mathbf{0}_{p \times q}$, respectively, denote the identity matrix of size p and a $p \times q$ matrix of zeros.

2. BASIC MODELING FRAMEWORK

We begin by describing an infinite time horizon dynamic discrete choice model as in Rust (1987, 1994).⁴ Given our empirical examples and application below, we shall sometimes refer to our representative economic agent as a *firm* and her payoffs as *profits*. Let $t \in \{1, 2, \dots, \infty\}$ denote time. The random variables in our model are the action and state variables, which we denote by a_t and s_t , respectively. a_t takes values from a finite set of alternatives $A = \{0, 1, \dots, J\}$. s_t contains two components, $s_t \equiv (x_t, \varepsilon_t) \in X \times \mathbb{R}^{J+1}$ with $X \subseteq \mathbb{R}$. x_t is public information to both the firm and the econometrician, while $\varepsilon_t \equiv (\varepsilon_t(0), \dots, \varepsilon_t(J)) \in \mathbb{R}^{J+1}$ is private information only observed by the firm. Future states are uncertain. Today's action and states affect outcomes for states in the future. The evolution of the states is summarized by a Markov transition law $P(s_{t+1}|s_t, a_t)$. The firm's period payoff function is $u(a_t, s_t) \in \mathbb{R}$. Future period's payoffs are discounted at the rate $\beta \in [0, 1)$. At time t , the firm observes s_t and chooses an action optimally. Specifically, we assume $a_t = \alpha(s_t)$ so that:

$$\begin{aligned} \alpha(s) &= \arg \max_{a \in A} \{u(a, s) + \beta E[V(s_{t+1})|s_t = s, a_t = a]\}, \quad \text{where} \\ V(s) &= \max_{a \in A} \{u(a, s) + \beta E[V(s_{t+1})|s_t = s, a_t = a]\}. \end{aligned} \quad (1)$$

Using the optimal decision rule, we can remove the max operator and write the value function as

$$V(s) = E \left[\sum_{t=0}^{\infty} \beta^t u(a_t, s_t) \middle| s_0 = s \right]. \quad (2)$$

The expectation operators in the displays above integrate out variables with respect to the probability distribution induced by the equilibrium choice probabilities and Markov transition law. As standard in the literature, we assume the following assumptions.

ASSUMPTION M.

- (i) (*Additive Separability*) For all a, x, ε :

$$u(a, x, \varepsilon) = \pi(a, x) + \varepsilon(a).$$

- (ii) (*Conditional Independence*) The transition distribution of the states has the following factorization for all $x', \varepsilon', x, \varepsilon, a$:

$$P(x', \varepsilon'|x, \varepsilon, a) = Q(\varepsilon')G(x'|x, a),$$

where Q is the cumulative distribution function of ε_t and G denotes the transition law of x_{t+1} conditioning on x_t, a_t . Furthermore, ε_t has finite first moments, and a positive, continuous, and bounded density on \mathbb{R}^{J+1} .

- (iii) (*Finite Observed State*) $X = \{1, \dots, K\}$.

⁴The notation for an infinite time stationary model is much simpler relative to a finite time horizon one. Our identification strategy is valid for finite time horizon models, and with or without absorbing states.

The primitives of the model under this setting consist of (π, β, Q, G) . Throughout the paper, we shall assume (G, Q) to be known. G can be identified from the data when (a_t, x_t, x_{t+1}) are observed. Consistent estimation of the joint distribution of (a_t, x_t, x_{t+1}) holds under weak conditions with a single time series, as well as repeated observations from short panels when there is no other unobserved heterogeneity. Q is typically assumed known in most empirical applications. Conditions for the identification of Q exist when x_t is a continuous variable using a large support type argument; for example, see Aguirregabiria and Suzuki (2014, Proposition 1), Buchholz, Shum, and Xu (2016, Lemma 4), and Chen (2014, Theorem 4). Our results do not depend on any continuity assumption to achieve identification as we take x_t to be a discrete random variable.

Our subsequent analysis use the fact that we can identify the choice probability from data as the starting point, which in turn is informative about (π, β) . More specifically, for any $a > 0$, let $\Delta v(a, x) \equiv v(a, x) - v(0, x)$, where $v(a, x)$ denotes the choice-specific value function that serves as the mean utility in a discrete choice modeling:

$$\begin{aligned} v(a, x) &= \pi(a, x) + \beta E[V(s_{t+1}) | x_t = x, a_t = a], \\ \Pr[a_t = a | x_t = x] &= \Pr[\Delta v(a, x) - \Delta v(a', x) > \varepsilon_t(a') - \varepsilon_t(a) \text{ for all } a' \neq a]. \end{aligned} \quad (3)$$

By inverting the choice probabilities (Hotz and Miller (1993)), we can recover $\Delta v(a, x)$ for all $a > 0, x$.

3. IDENTIFYING THE DISCOUNT FACTOR WITH LINEAR-IN-PARAMETER PAYOFFS

In this section, we assume the payoff function takes on a linear-in-parameter specification. Section 3.1 defines the identification concept for the discount factor and payoff parameters. Section 3.2 provides some representation lemmas that will be useful for defining a criterion function to study identification. Section 3.3 gives the identification result.

3.1 Definition of parametric identification

We will assume Assumption M and the following assumption throughout this section.

ASSUMPTION P (Linear-in-Parameter). *For all a, x :*

$$\pi(a, x; \theta) = \pi_0(a, x) + \theta^\top \pi_1(a, x),$$

where π_0 is a known real value function, π_1 is a known p -dimensional vector value function and θ belongs to \mathbb{R}^p .

Assumption P can be interpreted as nonparametric. For example, it can represent an unrestricted nonparametric function of π by assigning a parameter for each possible pair of a and x . However, such function is too rich and cannot be identified. We will maintain the parametric appearance for π as we will not be exploiting any nonparametric restriction in our identification study of the discount factor.

The role of π_0 is to represent the payoff components that are identifiable without the knowledge of the discount factor or other model primitives. In practice, π_0 and possibly parts of π_1 may have to be estimated (e.g., see Section 5.2). For the purpose of identification, they can be treated as known. The primitives in this setting are (β, θ) . They belong to $\mathcal{B} \times \Theta$ where $\mathcal{B} = [0, 1)$ and $\Theta = \mathbb{R}^p$. We are interested in the data generating discount factor and payoff parameters, which we denote by β_0 and θ_0 , respectively.

We begin by defining the parametric choice-specific value function (cf. equation (3)):

$$v(a, x; \beta, \theta) \equiv \sum_{t=0}^{\infty} \beta^t E[\pi(a_t, x_t; \theta) + \varepsilon_t(a_t) | a_0 = a, x_0 = x]. \quad (4)$$

Then we denote the differences in these value functions when action a is chosen relative to action 0 by $\Delta v(a, x; \beta, \theta) \equiv v(a, x; \beta, \theta) - v(0, x; \beta, \theta)$. It is important to emphasize that the stochastic process $\{a_t, x_t, \varepsilon_t\}_{t=0}^{\infty}$ that defines the right-hand side of equation (4) follows an optimal controlled process consistent with (β_0, θ_0) , whose distribution is identified by the observed probabilities from the data. Therefore, $\Delta v(a, x; \beta, \theta)$ is identified for all $(a, x) \in A \times X$ and $(\beta, \theta) \in \mathcal{B} \times \Theta$. Furthermore, $\Delta v(a, x; \beta_0, \theta_0)$ is also identified by Hotz–Miller’s inversion. We shall use the mapping $(\beta, \theta) \mapsto \{\Delta v(a, x; \beta, \theta)\}_{(a,x) \in A \times X}$ as a basis of our identification study.

More formally, we take each pair (β, θ) to be a *structure* of our empirical model and its implied choice-specific values, denoted by $\mathcal{V}_{\beta, \theta} \equiv \{\Delta v(a, x; \beta, \theta)\}_{(a,x) \in A \times X}$, to be its corresponding *reduced form*. We then define identification using the notion of *observational equivalence* in terms of the differences in expected payoffs.

DEFINITION I1 (Observational Equivalence). Any distinct (β, θ) and (β', θ') in $\mathcal{B} \times \Theta$ are *observationally equivalent* if and only if $\mathcal{V}_{\beta, \theta} = \mathcal{V}_{\beta', \theta'}$.

DEFINITION I2 (Point Identification). An element in $\mathcal{B} \times \Theta$, say (β, θ) , is *point identified* if and only if (β', θ') and (β, θ) are not observationally equivalent for all $(\beta', \theta') \neq (\beta, \theta)$ in $\mathcal{B} \times \Theta$.

For our identification study, we define our statistical model to be the collection of all reduced forms, namely: $\{\mathcal{V}_{\beta, \theta}\}_{(\beta, \theta) \in \mathcal{B} \times \Theta}$. All statements made on identification in Section 3 are in the context of this statistical model unless explicitly stated otherwise. Alternatively, we can also define a statistical model based on probability distributions as in the traditional econometrics studies on identification. Specifically, the model implied choice probabilities for each (β, θ) are

$$\mathcal{P}_{\beta, \theta} \equiv \{\Pr[\Delta v(a, x; \beta, \theta) - \Delta v(a', x; \beta, \theta) > \varepsilon(a') - \varepsilon(a) \text{ for all } a' \neq a]\}_{(a,x) \in A \times X}.$$

It is known there is a one-to-one relation between $\{\mathcal{V}_{\beta, \theta}\}_{(\beta, \theta) \in \mathcal{B} \times \Theta}$ and $\{\mathcal{P}_{\beta, \theta}\}_{(\beta, \theta) \in \mathcal{B} \times \Theta}$; see Matzkin (1991), Hotz and Miller (1993), and Norets and Takahashi (2013). Therefore, identification for our decision problem can be equivalently established with either $\{\mathcal{V}_{\beta, \theta}\}_{(\beta, \theta) \in \mathcal{B} \times \Theta}$ or $\{\mathcal{P}_{\beta, \theta}\}_{(\beta, \theta) \in \mathcal{B} \times \Theta}$. Note that one can interpret elements in $\mathcal{P}_{\beta, \theta}$ as the implied choice probabilities for an economic agent who solves a *pseudo-decision problem* where the expected payoff for taking each action is given by equation (4).

3.2 Some representation lemmas

Under Assumptions **M** and **P**, it shall be useful to separate out the contributions of the expected discounted payoffs in (4) as follows:

$$\begin{aligned} v(a, x; \beta, \theta) &= \pi_0(a, x) + \beta \sum_{t=0}^{\infty} \beta^t E[\pi_0(a_t, x_t) | a_0 = a, x_0 = x] \\ &\quad + \beta \sum_{t=0}^{\infty} \beta^t E[\varepsilon_t(a_t) | a_0 = a, x_0 = x] \\ &\quad + \theta^\top \left(\pi_1(a, x) + \beta \sum_{t=0}^{\infty} \beta^t E[\pi_1(a_t, x_t) | a_0 = a, x_0 = x] \right). \end{aligned}$$

Subsequently, by defining $\Delta\pi_l(a, x) \equiv \pi_l(a, x) - \pi_l(0, x)$ for $l = 0, 1$, we have:

$$\begin{aligned} \Delta v(a, x; \beta, \theta) &= \Delta\pi_0(a, x) + \beta \sum_{t=0}^{\infty} \beta^t (E[\pi_0(a_t, x_t) | a_0 = a, x_0 = x] \\ &\quad - E[\pi_0(a_t, x_t) | a_0 = 0, x_0 = x]) \\ &\quad + \beta \sum_{t=0}^{\infty} \beta^t (E[\varepsilon_t(a_t) | a_0 = a, x_0 = x] - E[\varepsilon_t(a_t) | a_0 = 0, x_0 = x]) \\ &\quad + \theta^\top \left(\Delta\pi_1(a, x) + \beta \sum_{t=0}^{\infty} \beta^t (E[\pi_1(a_t, x_t) | a_0 = a, x_0 = x] \right. \\ &\quad \left. - E[\pi_1(a_t, x_t) | a_0 = 0, x_0 = x]) \right). \end{aligned}$$

The decomposition of Δv helps us distinguish how β and/or θ affect different parts of the per-period payoffs. Lemma 1 summarizes this in a matrix form.

LEMMA 1. *Under Assumptions **M** and **P**, for all $a > 0$, $\Delta v(a, x; \beta, \theta)$ can be collected in the following vector form for all $(\beta, \theta) \in \mathcal{B} \times \Theta$:*

$$\begin{aligned} \Delta \mathbf{v}^a(\beta, \theta) &= \Delta \mathbf{R}_0^a + \beta \Delta \mathbf{H}^a (\mathbf{I}_K - \beta \mathbf{L})^{-1} \mathbf{R}_0 \\ &\quad + \beta \Delta \mathbf{H}^a (\mathbf{I}_K - \beta \mathbf{L})^{-1} \boldsymbol{\epsilon} \\ &\quad + (\Delta \mathbf{R}_1^a + \beta \Delta \mathbf{H}^a (\mathbf{I}_K - \beta \mathbf{L})^{-1} \mathbf{R}_1) \theta, \end{aligned} \tag{5}$$

where the elements in the above display are collected and explained in Tables **A** and **B**.

PROOF. This is a special case of Lemma R in Sanches, Silva, and Srisuma (2016). \square

All vectors and matrices in Tables **A** and **B** are either known or estimable from the choice and transitional probabilities. The tables will serve as a useful reference for con-

TABLE A. The matrices consist of (differences in) expected payoffs and probabilities. The latter represent conditional expectations for any function ψ of x_{t+1} .

Matrix	Dimension	Representing
$\Delta \mathbf{R}_1^a$	$K \times p$	$\Delta \pi_1(a, \cdot)$
\mathbf{R}_1	$K \times p$	$\pi_1(a, \cdot)$
\mathbf{L}	$K \times K$	$E[\psi(x_{t+1}) x_t = \cdot]$
\mathbf{H}^a	$K \times K$	$E[\psi(x_{t+1}) x_t = \cdot, a_t = a]$
$\Delta \mathbf{H}^a$	$K \times K$	$E[\psi(x_{t+1}) x_t = \cdot, a_t = a] - E[\psi(x_{t+1}) x_t = \cdot, a_t = 0]$

TABLE B. The $K \times 1$ vectors represent (differences in) expected payoffs.

Vector	Representing
ϵ	$E[\varepsilon_t(a_t) x_t = \cdot]$
$\Delta \mathbf{R}_0^a$	$\Delta \pi_0(a, \cdot)$
\mathbf{R}_0	$E[\pi_0(a_t, x_t) x_t = \cdot]$
$\Delta \mathbf{H}^a(\mathbf{I}_K - \beta \mathbf{L})^{-1} \mathbf{R}_0$	$\sum_{t=0}^{\infty} \beta^t (E[\pi_0(a_t, x_t) a_0 = a, x_0 = x] - E[\pi_0(a_t, x_t) a_0 = 0, x_0 = x])$
$\Delta \mathbf{H}^a(\mathbf{I}_K - \beta \mathbf{L})^{-1} \mathbf{R}_1$	$\sum_{t=0}^{\infty} \beta^t (E[\pi_1(a_t, x_t) a_0 = a, x_0 = x] - E[\pi_1(a_t, x_t) a_0 = 0, x_0 = x])$
$\Delta \mathbf{H}^a(\mathbf{I}_K - \beta \mathbf{L})^{-1} \epsilon$	$\sum_{t=0}^{\infty} \beta^t (E[\varepsilon_t(a_t) a_0 = a, x_0 = x] - E[\varepsilon_t(a_t) a_0 = 0, x_0 = x])$

structuring the necessary components we use for defining the criterion function in Section 3.3.

Given that we can identify $\Delta \mathbf{v}^a(\beta_0, \theta_0)$ for all $a > 0$, to identify (β_0, θ_0) , it is sufficient to show that for all $(\beta, \theta) \neq (\beta_0, \theta_0)$, $\Delta \mathbf{v}^a(\beta, \theta) \neq \Delta \mathbf{v}^a(\beta_0, \theta_0)$ for some a . Our next lemma provides a characterization as to how changing β and θ can affect $\Delta \mathbf{v}^a$.

LEMMA 2. Under Assumptions M and P, for any $a > 0$ and $(\beta, \theta), (\beta', \theta') \in \mathcal{B} \times \Theta$:

$$\Delta \mathbf{v}^a(\beta, \theta) - \Delta \mathbf{v}^a(\beta, \theta') = (\Delta \mathbf{R}_1^a + \beta \Delta \mathbf{H}^a(\mathbf{I}_K - \beta \mathbf{L})^{-1} \mathbf{R}_1)(\theta - \theta'), \quad (6)$$

$$\Delta \mathbf{v}^a(\beta', \theta') - \Delta \mathbf{v}^a(\beta, \theta') = (\beta - \beta') \Delta \mathbf{H}^a(\mathbf{I}_K - \beta' \mathbf{L})^{-1} (\mathbf{I}_K - \beta \mathbf{L})^{-1} (\mathbf{R}_0 + \mathbf{R}_1 \theta' + \epsilon). \quad (7)$$

And (β, θ) is identifiable if and only if there is no other (β', θ') such that for all $a > 0$:

$$\Delta \mathbf{v}^a(\beta', \theta') - \Delta \mathbf{v}^a(\beta, \theta') = \Delta \mathbf{v}^a(\beta, \theta) - \Delta \mathbf{v}^a(\beta, \theta').$$

PROOF. It follows from some algebra based on equation (5). \square

Lemma 2 illustrates the nature of the identification problem we have at hand. We highlight the following particulars:

(i) If the discount rate is assumed to be known, from (6), a sufficient condition for $\Delta \mathbf{v}^a(\beta_0, \theta) \neq \Delta \mathbf{v}^a(\beta_0, \theta')$ when $\theta \neq \theta'$ is that $\Delta \mathbf{R}_1^a + \beta \Delta \mathbf{H}^a(\mathbf{I}_K - \beta \mathbf{L})^{-1} \mathbf{R}_1$ has full column rank for some $a > 0$. Also see Theorem 3 in Srisuma (2015).

(ii) If the payoff function is assumed to be known, from (7), a sufficient condition for $\Delta \mathbf{v}^a(\beta', \theta_0) \neq \Delta \mathbf{v}^a(\beta, \theta_0)$ when $\beta \neq \beta'$ is that $(\mathbf{R}_0 + \mathbf{R}_1 \theta' + \epsilon) \neq 0$ and $\Delta \mathbf{H}^a$ is invertible some $a > 0$.

(iii) Suppose p is large relative to K . Then for any $a > 0$ such that $\Delta\mathbf{R}_1^a + \beta\Delta\mathbf{H}^a(\mathbf{I}_K - \beta\mathbf{L})^{-1}\mathbf{R}_1$ has rank K , and for any $\theta', \beta \neq \beta'$ that $\Delta\mathbf{v}^a(\beta', \theta') \neq \Delta\mathbf{v}^a(\beta, \theta')$, by equating (6) and (7), we can always find θ such that $\Delta\mathbf{v}^a(\beta', \theta') = \Delta\mathbf{v}^a(\beta, \theta)$.

Point (i) shows that sufficient conditions for identification of the payoff parameters when the discount rate is assumed known can be easily stated and verified. More generally, the sufficient condition for the identification of the payoff parameter can be stated in terms of the full column rank of the matrix that stacks together $\Delta\mathbf{R}_1^a + \beta\Delta\mathbf{H}^a(\mathbf{I}_K - \beta\mathbf{L})^{-1}\mathbf{R}_1$ over a . In the case we are able to identify the payoff function outside of the dynamic model, (ii) shows that the discount factor can also be identified and provide one type of sufficient conditions that can be readily checked. Point (iii) shares the intuition along the line of Manski (1993) that when the parameterization on the payoff function is too rich, (β, θ) may not be identifiable in $\mathcal{B} \times \Theta$.

From Lemma 2, it is also apparent that we should be able to identify (β_0, θ_0) jointly when the change in the vector of expected payoffs from altering the discount factor moves in a different direction to the change caused by altering the payoff parameters.

3.3 Sum of squares criterion function

The study of identification involving the discount factor is complicated due to the fact that $\mathcal{V}_{\beta, \theta}$ is nonlinear in (β, θ) . However, for a given β , we can see from (5) that $\mathcal{V}_{\beta, \theta}$ is linear in θ . We use profiling to exploit the conditional linearity to simplify the identification problem for a nonlinear model with $p + 1$ parameters to a one-dimensional problem.

Let $\mathbf{m}^a(\beta, \theta) \equiv \Delta\mathbf{v}^a(\beta_0, \theta_0) - \Delta\mathbf{v}^a(\beta, \theta)$. Then we can write, using (5):

$$\begin{aligned}\mathbf{m}^a(\beta, \theta) &= \mathbf{a}^a(\beta) - \mathbf{B}^a(\beta)\theta, \\ \mathbf{a}^a(\beta) &\equiv \Delta\mathbf{v}^a(\beta_0, \theta_0) - \Delta\mathbf{R}_0^a - \beta\Delta\mathbf{H}^a(\mathbf{I}_K - \beta\mathbf{L})^{-1}(\mathbf{R}_0 + \boldsymbol{\epsilon}), \\ \mathbf{B}^a(\beta) &\equiv \Delta\mathbf{R}_1^a + \beta\Delta\mathbf{H}^a(\mathbf{I}_K - \beta\mathbf{L})^{-1}\mathbf{R}_1.\end{aligned}$$

It is clear that $\mathbf{m}^a(\beta, \theta)$ is linear in θ for any given β . We can stack together the system of equations above across a . In doing so, we obtain the following vector value function, $\mathbf{m}: \mathcal{B} \times \Theta \rightarrow \mathbb{R}^{KJ}$:

$$\mathbf{m}(\beta, \theta) = \mathbf{a}(\beta) - \mathbf{B}(\beta)\theta, \quad (8)$$

where $\mathbf{a}(\beta)$ is a $KJ \times 1$ vector and $\mathbf{B}(\beta)$ is a $KJ \times p$ matrix.

Let $\mathcal{M}(\beta, \theta) \equiv \|\mathbf{m}(\beta, \theta)\|$, that is, $\mathcal{M}(\beta, \theta)$ is the Euclidean norm of $\mathbf{m}(\beta, \theta)$. Then by construction,

$$\mathcal{M}(\beta, \theta) = 0 \quad \text{if } (\beta, \theta) = (\beta_0, \theta_0),$$

and any other (β, θ) such that $\mathcal{M}(\beta, \theta) = 0$ is observationally equivalent to (β_0, θ_0) by the property of the norm. Therefore, \mathcal{M} has the necessary property to serve as a criterion for identification.

Next, we profile out θ in order to reduce the dimensionality on \mathcal{M} by exploiting its least squares structure. For each β , run a regression of $\mathbf{a}(\beta)$ on $\mathbf{B}(\beta)$, we can define:

$$\theta^*(\beta) \equiv (\mathbf{B}(\beta)^\top \mathbf{B}(\beta))^\dagger \mathbf{B}(\beta)^\top \mathbf{a}(\beta). \quad (9)$$

So that $\theta^*(\beta)$ is a least squares solution to $\min_{\theta \in \Theta} \mathcal{M}(\beta, \theta)$. Then we define:

$$\mathcal{M}^*(\beta) \equiv \mathcal{M}(\beta, \theta^*(\beta)). \quad (10)$$

By construction, it also holds that

$$\mathcal{M}^*(\beta) = 0 \quad \text{if } \beta = \beta_0.$$

In this way, we have reduced the parameter space in the identification problem to a one-dimensional one. Furthermore, the domain of the parameter space is on a small interval: $[0, 1)$. The reasoning is analogous to profiling in an estimation routine. Particularly, we can ignore any θ that does not lie in $\arg \min_{\theta \in \Theta} \mathcal{M}(\beta, \theta)$ since necessarily,

$$\mathcal{M}(\beta, \theta) > \mathcal{M}(\beta, \theta^*(\beta)) \geq 0.$$

Therefore, (β_0, θ_0) is identified when $\mathcal{M}^*(\beta)$ has a unique minimum *and* $\min_{\theta \in \Theta} \mathcal{M}(\beta_0, \theta)$ has a unique solution.

THEOREM 1. *Under Assumptions M and P, (β_0, θ_0) is identifiable in $\{\mathcal{V}_{\beta, \theta}\}_{(\beta, \theta) \in \mathcal{B} \times \Theta}$ if*

$$\mathcal{M}^*(\beta) = 0 \quad \text{if and only if } \beta = \beta_0,$$

and $\mathbf{B}(\beta_0)$ has full column rank.

PROOF. Suppose (β_0, θ_0) is identifiable. If there is $\beta' \neq \beta_0$ such that $\mathcal{M}^*(\beta') = 0$, then $\Delta \mathbf{v}^a(\beta_0, \theta_0) = \Delta \mathbf{v}^a(\beta', \theta^*(\beta'))$ for all a by the property of the norm. Since Θ is a closed set, by the projection theorem, $\theta^*(\beta')$ exists and is the unique element in Θ . This leads to a contradiction since (β_0, θ_0) and $(\beta', \theta^*(\beta'))$ are observationally equivalent. Next, suppose that $\mathbf{B}(\beta_0)$ does not have full column rank. Let θ' be another element in $\arg \min_{\theta \in \Theta} \mathcal{M}(\beta_0, \theta)$ that differs from θ_0 . Since $\mathcal{M}(\beta_0, \theta) \geq 0$ for all $\theta \in \Theta$ and $\mathcal{M}(\beta_0, \theta_0) = 0$, $\mathcal{M}(\beta_0, \theta') = 0$. Thus (β_0, θ_0) and (β_0, θ') are observationally equivalent, also a contradiction. \square

COMMENTS ON THEOREM 1.

(i) *High Level Assumptions.* Conditions in Theorem 1 are high level as we do not relate them to the underlying primitives of the model. However, they are statements made on objects that are observed or can be consistently estimated nonparametrically. In the [Appendix](#), we give more detailed conditions for \mathcal{M}^* to have a unique minimum; see Theorem 4.

(ii) *Feasible Check and Estimation.* Since we have reduced the identification problem to a single-parameter that can reside only in a narrow range, there is no need to refer to complicated results for the identification of a general nonlinear model. We can use the sample parts of components in Tables A and B to consistently estimate $\mathcal{M}^*(\beta)$ for all β . So one can plot the sample counterpart of \mathcal{M}^* over \mathcal{B} for an exhaustive analysis of the problem. Once the minimum of \mathcal{M}^* is found, the corresponding rank matrix can then be

checked. This suggests one natural way to estimate the discount factor, namely by grid search. In practice, we can detect an identification problem if the sample counterpart of \mathcal{M}^* contains a flat region at the minimum, or when the sample counterpart of $\mathbf{B}(\beta_0)$ does not have full column rank.

(iii) *Identification in the empirical model.* It is clear that positive identification of (β_0, θ_0) in our empirical model is sufficient for identifying (β_0, θ_0) in the full-solution model. Therefore, our identification results in this paper can be used to establish identification in the full-solution model. However, the implication may not be necessary, and we do not make any other claim on the identification of the full-solution model. The identification study in the full-solution model is much more complicated since it is less tractable analytically; for a further discussion, we refer the reader to [Srisuma \(2015\)](#).

By inspecting the proof of [Theorem 1](#), it is clear there are some separation between the identifiability of β_0 and θ_0 . In particular, we have defined $\theta^*(\beta)$ using a generalized inverse of the matrix $\mathbf{B}(\beta)^\top \mathbf{B}(\beta)$. Therefore, β_0 can be identified even if θ_0 is not.

The full column rank condition on $\mathbf{B}(\beta_0)$, however, is not an innocuous assumption when we view [Assumption P](#) as a representation of a nonparametric function. In practice, this is often delivered by exclusion assumptions or more generally by *normalization* of payoff parameters. In the next section, we will focus on payoff parameters that we call *switching costs*. We will revisit the question of identifiability of the discount factor under different normalization choice in [Section 4.3](#).

4. NONPARAMETRIC IDENTIFICATION OF SWITCHING COSTS

In this section, we consider payoff functions under nonparametric restrictions that allow us to obtain closed-form expressions for the switching costs parameters. In [Section 4.1](#), we define a switching cost function and explain the assumptions required for our identification result. [Section 4.2](#) gives the identification result. [Section 4.3](#) relates the identification of the discount factor under [Assumption P](#) to models with switching costs.

4.1 *Switching costs*

The payoff function cannot be nonparametrically identified without any restrictions. Economic theory can help guide how to impose structures on the payoff function. A main consideration in making a dynamic discrete decision is how a change in one's action from the previous period immediately affect today's payoffs. Actions from the past are therefore often important components of the state variables. We will consider restrictions focusing on *switching costs*.

In order to highlight the role of switching costs, we distinguish past actions from other state variables. At time t , we denote actions from the previous period by w_t , so that $w_t \equiv a_{t-1}$. We denote the switching cost from changing action from w to a by $SC^{w \rightarrow a}$. Subsequently, in this section we shall maintain an updated version of [Assumption M](#) where x_t is replaced with (w_t, x_t) everywhere. In addition, we impose the following assumptions.

ASSUMPTION N.

(i) (*Decomposition of Profits*): For all a, w, x :

$$\pi(a, w, x) = \mu(a, x) + \phi(a - w, w, x),$$

such that $\phi(0, w, x) = 0$.

(ii) (*Conditional Independence*): The distribution of x_{t+1} conditional on a_t and x_t is independent of w_t .

The decomposition of π in N(i) may appear peculiar at first, but it is typical in many empirical IO applications. We will give an interpretation of its components within the context of an IO application. The defining feature of μ is that it excludes past actions. μ can represent the firm's operational profit in the current period, such as variable profits and operational costs, which does not depend on actions from the past. ϕ is the *switching cost function* that takes nonzero values only when a change of action occurs. Note that, by construction, we have

$$\phi(a - w, w, x) = SC^{w \rightarrow a}(x) \cdot \mathbf{1}[w \neq a], \quad (11)$$

where $\mathbf{1}[\cdot]$ denotes the indicator function.

Assumption N(ii) imposes that knowing actions from the past does not help predict future state variables when the present action and other observable state variables are known. Note that N(ii) is not implied by M(ii). In many applications, $\{x_t\}$ is simply assumed to be a strictly exogenous first order Markov process. Specifically, this implies x_{t+1} is independent of a_t conditional on x_t in addition to N(ii). In any case, unlike M(ii), N(ii) is a restriction made on the observables so it can be tested directly from the data. Later on we shall show how x_t can be modified to contain past actions so N(ii) can be weakened to allow for dependence of other state variables with past actions.

Even under Assumption N(i) identification issue persists (e.g., see the discussion in Aguirregabiria and Suzuki (2014)). $SC^{w \rightarrow a}$ cannot be identified for all $w \neq a$ without any further restrictions. Some of their differences, however, can be identified. For example, identification is possible if we *normalize* some baseline switching costs to be known. We will look at different restrictions that can be used to identify individual or combination of the switching costs. Before giving the formal result, we provide an intuition as to why Assumption N is helpful for identifying the switching costs. It will also illustrate the key steps of our identification strategy.

Exclusion and independence restrictions Consider a two-period entry/exit decision problem. Let $A = \{0, 1\}$, where 0 denotes exit and 1 denotes entry. Then $SC^{0 \rightarrow 1}$ and $SC^{1 \rightarrow 0}$, respectively, have interpretations of entry cost and scrap value. In this case, we can write

$$\phi(a - w, w, x) = SC^{0 \rightarrow 1}(x) \cdot a(1 - w) + SC^{1 \rightarrow 0}(x) \cdot (1 - a)w. \quad (12)$$

The choice-specific value function (cf. (3)) in this model is

$$v(a, w, x) = \pi(a, w, x) + \beta E[\pi(a_{t+1}, w_{t+1}, x_{t+1}) | a_t = a, w_t = w, x_t = x].$$

Let $\Delta v(w, x) \equiv v(1, w, x) - v(0, w, x)$. At time t , a firm will enter if and only if $\Delta v(w, x) > \varepsilon_t(0) - \varepsilon_t(1)$. We can identify Δv from the observed choice probabilities.

The role of our assumptions is to isolate today's switching costs from the remaining components in the choice-specific value function. Specifically, we apply N(i) to decompose the profit function in the current period and use N(ii) to simplify the expected future profits. We can then rewrite the equation above as

$$\begin{aligned} v(a, w, x) &= \lambda(a, x) + \phi(a - w, w, x), \quad \text{where} \\ \lambda(a, x) &= \mu(a, x) + \beta E[\pi(a_{t+1}, a, x_{t+1}) | a_t = a, x_t = x]. \end{aligned}$$

Crucially note that the conditional expectation on future profits in λ no longer depends on w_t under N(ii) due to the law of iterated expectation. We treat λ as a nuisance parameter. It is a nonparametric object that depends on all primitives in the model. Let $\Delta\lambda(x) \equiv \lambda(1, x) - \lambda(0, x)$. Using equation (12), we have

$$\Delta v(w, x) = \Delta\lambda(x) + SC^{0 \rightarrow 1}(x) \cdot (1 - w) - SC^{1 \rightarrow 0}(x) \cdot w. \quad (13)$$

It is now clear we can identify a combination of the switching costs by differencing out $\Delta\lambda$ in the equation above:

$$\Delta v(1, x) - \Delta v(0, x) = -SC^{0 \rightarrow 1}(x) - SC^{1 \rightarrow 0}(x). \quad (14)$$

In an entry/exit game, the quantity $-SC^{0 \rightarrow 1} - SC^{1 \rightarrow 0}$ represents the *sunk entry cost* that a firm cannot recover back once it decides to leave the market after entering. Equation (14) shows the sunk entry cost can be identified independently of β and μ . On the other hand, it is well known that entry cost and scrap value cannot be nonparametrically identified separately in this particular model. In an empirical work, an unidentified object gets normalized. It is clear from equation (14) that either the entry cost or scrap value can be identified if one of them is assumed to be known. For example, a common assumption is to normalize the scrap value to be zero, the entry cost can be estimated conditionally on this value along with the other parameters.

The identification strategy above can be generalized substantially. Results for a more general single agent decision model under M and N can be obtained with little modification. But extending our single agent's results to dynamic games is more complex. It requires additional notation and a more general notion of a difference, characterized by a projection matrix, is used. We defer the details for dynamic games to the [Appendix](#).

4.2 Closed-form identification

We start by providing an expression for the differences in choice-specific valuations that generalizes equation (13). For any $a > 0$, let $\Delta v(a, w, x) \equiv v(a, w, x) - v(0, w, x)$, $\Delta\lambda(a, x) \equiv \lambda(a, x) - \lambda(0, x)$, and $\Delta\phi(a, w, x) \equiv \phi(a - w, w, x) - \phi(-w, w, x)$. Lemma 3 generalizes equation (13).

LEMMA 3. *Under Assumptions M and N, we have for all i , $a > 0$ and w, x :*

$$\Delta v(a, w, x) = \Delta\lambda(a, x) + \Delta\phi(a, w, x), \quad (15)$$

where

$$\begin{aligned}\Delta\lambda(a, x) &\equiv \mu(a, x) - \mu(0, x) + \beta(\tilde{m}(a, x) - \tilde{m}(0, x)), \\ \tilde{m}(a, x) &\equiv E[m(a, x_{t+1})|a_t = a, x_t = x], \\ m(w, x) &\equiv E[V(s_t)|w_t = w, x_t = x].\end{aligned}$$

PROOF. Using the law of iterated expectation, the value function as defined in equation (2) satisfies: $E[V(s_{t+1})|a_t, w_t, x_t] = E[m(w_{t+1}, x_{t+1})|a_t, w_t, x_t]$ under **M(ii)**. $E[m(w_{t+1}, x_{t+1})|a_t, w_t, x_t]$ can be simplified further to $E[\tilde{m}(a_t, x_t)|a_t, x_t]$ after another application of the law of iterated expectation and imposing **N(ii)**. The remainder of the proof then follows from the definitions of the terms defined within the main text. \square

The components of Δv consist of $\Delta\lambda$ and $\Delta\phi$. We treat $\Delta\lambda$ as a nuisance parameter. $\Delta\phi$ contains the switching costs of interest, for any a, w, x :

$$\Delta\phi(a, w, x) = SC^{w \rightarrow a}(x) \cdot \mathbf{1}[w \neq a] - SC^{w \rightarrow 0}(x) \cdot \mathbf{1}[w \neq 0]. \quad (16)$$

As seen previously, we can identify the differences in $\Delta\phi$ by eliminating $\Delta\lambda$. This can be done by looking at the differences of $\Delta v(a, w, x)$ across different w while holding (a, x) fixed.

THEOREM 2. *Under Assumptions **M** and **N**, we have for all $a > 0$ and x, w, w' :*

$$\Delta\phi(a, w, x) - \Delta\phi(a, w', x) = \Delta v(a, w, x) - \Delta v(a, w', x). \quad (17)$$

Theorem 2 follows immediately from Lemma 3. Equation (17) tells us that we can always identify *some* combinations of the switching costs nonparametrically. Importantly, the identified objects do not depend on β or μ .

COMMENTS ON THEOREM 2.

(i) Certain differences in $\Delta\phi$ in equation (17) are economically meaningful. We have already introduced the sunk entry cost in the entry/exit model as an example. The notion of sunk costs naturally generalizes to other irreversible investment costs with a varying degree of commitment. More specifically consider an investment or capacity game where it costs a firm to choose $a_t > a_{t-1}$ and, conversely, a firm can divest to recover some of these costs by choosing $a_t < a_{t-1}$. In this case, $-SC^{a' \rightarrow a} - SC^{a \rightarrow a'}$ with $a > a'$ represents a sunk investment cost for a firm that increases its investment level from a' to a then divests back to a' . Using equations (16) and (17), Corollaries 1 and 2 give closed-form expressions for identifying the sunk investment costs.

COROLLARY 1. *For all $a > 0, x$:*

$$-SC^{0 \rightarrow a}(x) - SC^{a \rightarrow 0}(x) = \Delta v(a, a, x) - \Delta v(a, 0, x).$$

COROLLARY 2. For all $a, a' > 0, x$:

$$-SC^{a' \rightarrow a}(x) - SC^{a \rightarrow a'}(x) = \Delta v(a, a, x) + \Delta v(a', a', x) - \Delta v(a, a', x) - \Delta v(a', a, x).$$

(ii) We would prefer to identify the switching costs individually. However, without further information, they are not identified nonparametrically for this type of models; for example, see [Aguirregabiria and Suzuki \(2014\)](#) for a thorough discussion. But identification can be achieved if we are willing to impose some constraints on the switching costs. One example is by assuming symmetry of switching costs between any two actions, which would be reasonable in applications with logistical or physical adjustment costs such as the traditional menu costs (e.g., see [Slade \(1998\)](#)). Corollary 3 shows that individual switching costs under symmetry are identified. Its proof follows immediately from Corollaries 1 and 2.

COROLLARY 3. For all a, a', x , suppose that $SC^{a' \rightarrow a}(x) = SC^{a \rightarrow a'}(x)$, then for any $a, a' > 0$:

$$SC^{0 \rightarrow a}(x) = -(\Delta v(a, a, x) - \Delta v(a, 0, x))/2,$$

$$SC^{a \rightarrow a'}(x) = -(\Delta v(a, a, x) + \Delta v(a', a', x) - \Delta v(a, a', x) - \Delta v(a', a, x))/2.$$

(iii) It is frequent in many applications that some components of the switching costs are taken to be known. Typically, this is done by way of a normalization assumption. The most commonly used normalization assumes that taking action 0 yields zero payoff. For example, for an entry or investment game with entry, such assumption means a firm has no recovery value of assets upon leaving the market. In other cases, some institutional or other external knowledge outside of the dynamic model are used. For example, [Kalouptside \(2014\)](#) uses data on resale value of second hand ships to identify the scrap values and entry costs directly. In another example, in a study of promotion pricing decisions, [Myśliwski, Sanches, Silva, and Srisuma \(2017\)](#) rely on anecdotal evidence to assume a cost is incurred to producers when a sale promotion is on while there is no costs for switching back to the regular price. In these cases, we can identify individual switching costs directly as Corollary 4 shows.

COROLLARY 4. For all a' , suppose $SC^{a' \rightarrow 0}(x) = \phi_0(w, x)$ then for any a, a', x :

$$SC^{a' \rightarrow a}(x) = \Delta v(a, a', x) - \Delta v(a, a, x) + \phi_0(a', x) - \phi_0(a, x). \quad (18)$$

It is important to highlight that assigning incorrect values to ϕ_0 generally leads to incorrect values of $SC^{w \rightarrow a}$. On the other hand, it is easy to verify that certain combinations of switching costs, including those in Corollaries 1 and 2, are robust against any choice of ϕ_0 .

(iv) Generally Corollaries 1 and 2 can be informative on the validity of a particular normalization choice since they have been derived without any normalization. For example, let us go back to the discussion on investment game at the end of our first comment where there is a divestment opportunity. In this context, it would be natural to

assume that $-SC^{a' \rightarrow a} - SC^{a \rightarrow a'} = c_0$ for some positive c_0 when $a > a'$. Then, given both $-SC^{a' \rightarrow a}$ and $SC^{a \rightarrow a'}$ are positive, it must be the case that $-SC^{a' \rightarrow a}$ is bounded below by c_0 .

(v) When $\mathcal{A} = \{0, 1\}$ our Theorem 1 implies the sunk entry cost can be identified without any normalization. Proposition 2 in Aguirregabiria and Suzuki (2014) has established the same result using a different argument.

The results of Theorem 2 and Corollaries 1 to 4 are constructive. We can replace the unknown Δv using the empirical choice probabilities. The sample analog estimators can be computed without any optimization. Given the empirical literature is concerned with the computational cost our closed-form identification result can substantially reduce the number of parameters to be estimated in a model. Such estimators will be consistent and asymptotically normal as long as the initial choice probabilities have these properties.

4.3 Identification and normalization

We have emphasized that normalizations of switching costs are necessary in many situations. The validity of the identification of payoff parameters is not robust against incorrect normalization choice. We now ask: to what extent the identification of the discount factor depends on the specific normalization choice on the payoff parameters?

In the empirical literature, the discount factor is customarily assumed to be known while the focus on identification falls on which payoff parameters can (or cannot) be identified. A particular normalization choice is made, for example, by assigning a value to an unknown parameter as previously explained. Such normalization assumption is always made independent to the choice of the discount factor. The identification problem on the payoff parameters considered in practice therefore mathematically translates to the matrix $\mathbf{B}(\beta)$ in equation (8) being rank deficient for all β . In particular, it is also implicitly assumed that the linear dependence relation between the column vectors of $\mathbf{B}(\beta)$ are the same for all β .

Recall that $\mathbf{B}(\beta)$ is a $KJ \times p$ matrix. For the remainder of this subsection, we shall assume $\rho(\mathbf{B}(\beta)) = r < p$ for all β , such that:

$$\mathbf{B}(\beta) = [\mathbf{B}_1(\beta) : \mathbf{B}_2(\beta)],$$

where $\mathbf{B}_1(\beta)$ is a matrix consisting of the first r columns of $\mathbf{B}(\beta)$ with $CS(\mathbf{B}_1(\beta)) = CS(\mathbf{B}(\beta))$, and $\mathbf{B}_2(\beta)$ is a matrix containing the last $(p - r)$ columns of $\mathbf{B}(\beta)$. It will now be convenient to reintroduce here $\mathcal{M}(\beta, \theta) = \|\mathbf{a}(\beta) - \mathbf{B}(\beta)\theta\|$ from Section 3.3, along with equations (9) and (10), respectively:

$$\begin{aligned} \theta^*(\beta) &\equiv (\mathbf{B}(\beta)^\top \mathbf{B}(\beta))^\dagger \mathbf{B}(\beta)^\top \mathbf{a}(\beta), \\ \mathcal{M}^*(\beta) &\equiv \mathcal{M}(\beta, \theta^*(\beta)). \end{aligned}$$

When we present our Theorem 1, we stated that “ (β_0, θ_0) is identified when $\mathcal{M}^*(\beta)$ has a unique minimum and $\min_{\theta \in \Theta} \mathcal{M}(\beta_0, \theta)$ has a unique solution.” The issue associated

with normalizing payoff parameters only concerns the latter, as we know $\mathcal{M}(\beta_0, \theta)$ has a unique minimum at θ_0 if and only if $\mathbf{B}(\beta)$ has full column rank. Since $\mathbf{B}(\beta)$ is rank deficient, $\mathcal{M}(\beta_0, \theta)$ has a linear subspace of minimizers. Normalization is a way to select an element from this subspace. This is a separate issue to whether $\mathcal{M}^*(\beta)$ has a unique minimum or not. One way to clearly illustrate this is the following.

Since $CS(\mathbf{B}_2(\beta)) \subset CS(\mathbf{B}_1(\beta))$, there exists an $r \times (p - r)$ matrix $\mathbf{\Gamma}$ such that $\mathbf{B}_2(\beta) = \mathbf{B}_1(\beta)\mathbf{\Gamma}$.⁵ Making a normalization on the payoff parameters corresponds to fixing a value of θ_2 . For any (β, θ_2) , we can define $\theta_1^*(\beta, \theta_2)$ to be the minimizer of $\|\mathbf{a}(\beta) - \mathbf{B}_1(\beta)\theta_1 - \mathbf{B}_1(\beta)\mathbf{\Gamma}\theta_2\|$, so that

$$\theta_1^*(\beta, \theta_2) = (\mathbf{B}_1(\beta)^\top \mathbf{B}_1(\beta))^{-1} \mathbf{B}_1(\beta)^\top \mathbf{a}(\beta) - \mathbf{\Gamma}\theta_2.$$

We can then profile out θ_1 , and define:

$$\mathcal{M}^*(\beta, \theta_2) \equiv \|\mathbf{a}(\beta) - \mathbf{B}_1(\beta)\theta_1^*(\beta, \theta_2) - \mathbf{B}_1(\beta)\mathbf{\Gamma}\theta_2\|.$$

Substituting $\theta_1^*(\beta, \theta_2)$ into the right-hand side of the display above, we get

$$\mathcal{M}^*(\beta, \theta_2) \equiv \|\mathbf{a}(\beta) - \mathbf{B}_1(\beta)(\mathbf{B}_1(\beta)^\top \mathbf{B}_1(\beta))^{-1} \mathbf{B}_1(\beta)^\top \mathbf{a}(\beta)\|.$$

We see that $\mathcal{M}^*(\beta, \theta_2)$ is simply the norm of the residual one gets from an orthogonal projection of $\mathbf{a}(\beta)$ onto $CS(\mathbf{B}_1(\beta))$. Importantly, $\mathcal{M}^*(\beta, \theta_2)$ does not depend on θ_2 . From the projection theory in linear algebra, $\mathcal{M}^*(\beta)$ and $\mathcal{M}^*(\beta, \theta_2)$ are necessarily equal. This residual will also be identical if we project $\mathbf{a}(\beta)$ on the linear span of any other r linear combinations of the columns in $\mathbf{B}(\beta)$ as long as it equals $CS(\mathbf{B}(\beta))$. Therefore, our argument holds without any loss of generality on how we select $\mathbf{B}_1(\beta)$. In practice, a researcher has to perform this selection when she decides upon her normalization choice. Subsequently, the discount factor can be identified regardless of how we normalize the payoff parameters. We state this result as a proposition.

PROPOSITION 1. *If the discount factor can be identified, it can be identified for all normalization choices on the payoff parameters.*

Our discussion here also leads to another empirical fact that may not be obvious a priori. Suppose a researcher specifies a payoff function in practice that satisfies both **P** and **N**. Then there are two different ways to estimate the switching costs based on our parametric and nonparametric identification approaches. We have shown in Section 4 that some combinations of the switching costs can be identified without any normalization using the nonparametric approach. We are interested to know whether the parametric approach taken in Section 3, which relies on a possibly incorrect normalization choice, can consistently estimate these combinations.

The answer is positive. Consider any combination of the switching costs, which can be written explicitly in terms of the differences in choice-specific valuations (e.g., sunk

⁵For instance, this is a consequence of Theorem 6.2.4 in Mirsky (1955).

costs, and more generally Corollaries 1 and 2). A vector of such combinations can be represented by $\Sigma \mathbf{a}_0$ for some matrix Σ . Then for any $\tilde{\theta}$ such that $(\beta_0, \tilde{\theta})$ is observationally equivalent to (β_0, θ_0) we also have $\Sigma \mathbf{a}_0 = \Sigma \mathbf{B}(\beta_0) \theta_0 = \Sigma \mathbf{B}(\beta_0) \tilde{\theta}$, that is, the combinations of switching costs described by $\Sigma \mathbf{B}(\beta_0)$ identify the same objects.

5. NUMERICAL ILLUSTRATION

We now illustrate the use of our identification strategies and implement the suggested estimators in the previous sections. Section 5.1 gives results from a Monte Carlo study taken from Pesendorfer and Schmidt-Dengler (2008). Section 5.2 estimates a discrete investment game using the data from Ryan (2012).

5.1 Monte Carlo study

The simulation design is the two-firm dynamic entry game taken from Section 7 in Pesendorfer and Schmidt-Dengler (2008). In period t , each firm i has two possible choices, $a_{it} \in \{0, 1\}$; with $a_{it} = 1$ denoting entry. The only observed state variables are previous period's actions, $w_t = (a_{1t-1}, a_{2t-1})$. Using their notation, firm 1's period payoffs are described as follows:

$$\pi_1(a_{1t}, a_{2t}, x_t; \theta) = a_{1t}(\mu_1 + \mu_2 a_{2t}) + a_{1t}(1 - a_{1t-1})F + (1 - a_{1t})a_{1t-1}W, \quad (19)$$

where μ_1, μ_2, F , and W are respectively the monopoly profit, duopoly profit, entry cost, and scrap value. The latter two components are switching costs. Each firm also receives additive private shocks that are i.i.d. $\mathcal{N}(0, 1)$. The game is symmetric and Firm's 2 payoffs are defined analogously. The data generating parameters are set as: $(\mu_{10}, \mu_{20}, F_0, W_0) = (1.2, -1.2, -0.2, 0.1)$ and $\beta_0 = 0.9$. Pesendorfer and Schmidt-Dengler (2008) show there are three distinct equilibria for this game.

It is easy to verify the model satisfies both Assumptions MN and MP in the Appendix, which are the dynamic game's generalization of Assumptions N and P. Therefore, we can estimate the model in at least two different ways. We consider the following two estimation methods. Method A profiles out all the payoff parameters using the OLS expression and use grid search to estimate the discount factor. Method B first estimates the entry cost in closed-form independently before profiling out the other payoff parameters and use grid search to estimate the discount factor. We will also be interested to see how sensitive our estimates are with respect to the normalization choice.

For each equilibrium, we perform 10,000 simulations with sample sizes $N = 100, 1000, 10,000$. Since the entry cost and scrap value cannot be jointly identified, we estimate the model under different normalized values for W . We report the bias and standard deviation (in *italics*) for $(\hat{\beta}, \hat{\mu}_1, \hat{\mu}_2, \hat{F})$ and the sunk entry cost (\widehat{SUNK}). We use the **bold** font to highlight the statistics that correspond to the correctly assumed choice of W . We estimate the sunk entry cost for Methods A and B by first estimating the entry cost and combine it with the assumed scrap value. In addition, we also estimate the sunk entry cost without normalizing the scrap value according to Example 1 in the Appendix (also see Corollary 1). We label the columns of statistics for the sunk entry estimator

TABLE 1. Data generated from equilibrium 1 in Pesendorfer and Schmidt-Dengler (2008)

N	W	Method A			Method B			N-N
		0	0.1	0.2	0	0.1	0.2	–
100	$\hat{\beta}$	–0.0809	–0.0806	–0.0799	–0.0752	–0.0768	–0.0738	–
		<i>0.2697</i>	<i>0.2691</i>	<i>0.2686</i>	<i>0.2619</i>	<i>0.2640</i>	<i>0.2596</i>	–
	$\hat{\mu}_1$	–0.0418	–0.0253	–0.0071	–0.0631	–0.0450	–0.0291	–
		<i>0.2974</i>	<i>0.3050</i>	<i>0.3150</i>	<i>0.3693</i>	<i>0.3774</i>	<i>0.3858</i>	–
	$\hat{\mu}_2$	0.0627	0.0815	0.0988	0.0963	0.1141	0.1313	–
		<i>0.2970</i>	<i>0.2991</i>	<i>0.3029</i>	<i>0.4779</i>	<i>0.4801</i>	<i>0.4831</i>	–
	\hat{F}	0.0446	–0.0554	–0.1552	–0.0019	–0.1017	–0.2021	–
		<i>0.2836</i>	<i>0.2835</i>	<i>0.2839</i>	<i>0.5692</i>	<i>0.5699</i>	<i>0.5702</i>	–
	\widehat{SUNK}	0.0554	0.0554	0.0552	0.1019	0.1017	0.1021	0.0477
		<i>0.2836</i>	<i>0.2835</i>	<i>0.2839</i>	<i>0.5692</i>	<i>0.5699</i>	<i>0.5702</i>	<i>0.5935</i>
1000	$\hat{\beta}$	–0.0356	–0.0372	–0.0380	–0.0328	–0.0339	–0.0343	–
		<i>0.1741</i>	<i>0.1790</i>	<i>0.1801</i>	<i>0.1677</i>	<i>0.1695</i>	<i>0.1715</i>	–
	$\hat{\mu}_1$	–0.0051	0.0090	0.0229	–0.0028	0.0110	0.0244	–
		<i>0.1032</i>	<i>0.1129</i>	<i>0.1251</i>	<i>0.1066</i>	<i>0.1152</i>	<i>0.1265</i>	–
	$\hat{\mu}_2$	–0.0046	0.0091	0.0231	–0.0084	0.0050	0.0185	–
		<i>0.0934</i>	<i>0.0946</i>	<i>0.0992</i>	<i>0.1190</i>	<i>0.1204</i>	<i>0.1246</i>	–
	\hat{F}	0.0958	–0.0042	–0.1042	0.1000	0.0000	–0.1000	–
		<i>0.0901</i>	<i>0.0901</i>	<i>0.0902</i>	<i>0.1480</i>	<i>0.1480</i>	<i>0.1480</i>	–
	\widehat{SUNK}	0.0042	0.0042	0.0042	0.0001	0.0001	0.0001	–0.0132
		<i>0.0901</i>	<i>0.0901</i>	<i>0.0902</i>	<i>0.1480</i>	<i>0.1480</i>	<i>0.1480</i>	<i>0.1573</i>
10,000	$\hat{\beta}$	–0.0005	–0.0003	–0.0005	–0.0005	–0.0007	–0.0005	–
		<i>0.0204</i>	<i>0.0158</i>	<i>0.0204</i>	<i>0.0204</i>	<i>0.0238</i>	<i>0.0205</i>	–
	$\hat{\mu}_1$	–0.0104	–0.0004	0.0097	–0.0101	0.0000	0.0100	–
		<i>0.0298</i>	<i>0.0299</i>	<i>0.0309</i>	<i>0.0302</i>	<i>0.0310</i>	<i>0.0312</i>	–
	$\hat{\mu}_2$	–0.0093	0.0007	0.0108	–0.0098	0.0003	0.0103	–
		<i>0.0297</i>	<i>0.0298</i>	<i>0.0300</i>	<i>0.0355</i>	<i>0.0356</i>	<i>0.0358</i>	–
	\hat{F}	0.0992	–0.0008	–0.1008	0.0998	–0.0002	–0.1002	–
		<i>0.0282</i>	<i>0.0282</i>	<i>0.0282</i>	<i>0.0437</i>	<i>0.0437</i>	<i>0.0437</i>	–
	\widehat{SUNK}	0.0008	0.0008	0.0008	0.0002	0.0002	0.0002	–0.0011
		<i>0.0282</i>	<i>0.0282</i>	<i>0.0282</i>	<i>0.0437</i>	<i>0.0437</i>	<i>0.0437</i>	<i>0.0454</i>

with no normalization by N-N. Tables 1–3 provides results that correspond to the data generated according to the three equilibria as enumerated in Pesendorfer and Schmidt-Dengler (2008), respectively.

The findings are in line with the theory part of the paper. First, it shows the discount factor can be consistently estimated. The consistency property is robust against the normalization choice of the scrap value. The sunk entry cost can also be consistently estimated independently of the scrap value used. When the model is correctly specified in the sense, we correctly assume $W = W_0$ all estimators are consistent, while misspecifying the scrap value cause biases to all estimators of the individual payoff parameters. The estimation results from Methods A and B, as well as N-N for the sunk entry, are qualitatively the same across all equilibria. The performances between estimation methods seem to depend on the equilibrium and sample size. Method A performs better in Equi-

TABLE 2. Data generated from equilibrium 2 in Pesendorfer and Schmidt-Dengler (2008)

N	W	Method A			Method B			N-N
		0	0.1	0.2	0	0.1	0.2	
100	$\hat{\beta}$	-0.0675	-0.0691	-0.0704	-0.0667	-0.0660	-0.0684	-
		<i>0.2501</i>	<i>0.2523</i>	<i>0.2542</i>	<i>0.2493</i>	<i>0.2477</i>	<i>0.2513</i>	-
	$\hat{\mu}_1$	-0.2087	-0.1899	-0.1726	-0.1185	-0.1027	-0.0835	-
		<i>0.3978</i>	<i>0.4135</i>	<i>0.4286</i>	<i>0.4495</i>	<i>0.4572</i>	<i>0.4718</i>	-
	$\hat{\mu}_2$	0.3264	0.3447	0.3623	0.1847	0.2025	0.2196	-
		<i>0.5430</i>	<i>0.5454</i>	<i>0.5500</i>	<i>0.6563</i>	<i>0.6605</i>	<i>0.6641</i>	-
	\hat{F}	-0.0630	-0.1632	-0.2632	0.0942	-0.0058	-0.1058	-
		<i>0.4166</i>	<i>0.4161</i>	<i>0.4159</i>	<i>0.5515</i>	<i>0.5515</i>	<i>0.5515</i>	-
\widehat{SUNK}	0.1630	0.1632	0.1632	0.0058	0.0058	0.0058	-0.0455	
	<i>0.4166</i>	<i>0.4161</i>	<i>0.4159</i>	<i>0.5515</i>	<i>0.5515</i>	<i>0.5515</i>	<i>0.5991</i>	
1000	$\hat{\beta}$	-0.0296	-0.0302	-0.0314	-0.0318	-0.0306	-0.0304	-
		<i>0.1584</i>	<i>0.1600</i>	<i>0.1625</i>	<i>0.1637</i>	<i>0.1603</i>	<i>0.1594</i>	-
	$\hat{\mu}_1$	-0.0275	-0.0139	0.0003	-0.0096	0.0028	0.0158	-
		<i>0.1631</i>	<i>0.1739</i>	<i>0.1872</i>	<i>0.1596</i>	<i>0.1691</i>	<i>0.1807</i>	-
	$\hat{\mu}_2$	0.0494	0.0626	0.0763	0.0267	0.0394	0.0523	-
		<i>0.2108</i>	<i>0.2159</i>	<i>0.2234</i>	<i>0.2047</i>	<i>0.2097</i>	<i>0.2162</i>	-
	\hat{F}	0.0767	-0.0233	-0.1233	0.1006	0.0006	-0.0994	-
		<i>0.1526</i>	<i>0.1526</i>	<i>0.1526</i>	<i>0.1495</i>	<i>0.1495</i>	<i>0.1495</i>	-
\widehat{SUNK}	0.0233	0.0233	0.0233	-0.0006	-0.0006	-0.0006	-0.0052	
	<i>0.1526</i>	<i>0.1526</i>	<i>0.1526</i>	<i>0.1495</i>	<i>0.1495</i>	<i>0.1495</i>	<i>0.1638</i>	
10,000	$\hat{\beta}$	-0.0001	-0.0002	-0.0004	-0.0002	-0.0004	-0.0002	-
		<i>0.0093</i>	<i>0.0127</i>	<i>0.0183</i>	<i>0.0130</i>	<i>0.0183</i>	<i>0.0128</i>	-
	$\hat{\mu}_1$	-0.0147	-0.0046	0.0056	-0.0127	-0.0025	0.0073	-
		<i>0.0399</i>	<i>0.0405</i>	<i>0.0425</i>	<i>0.0381</i>	<i>0.0398</i>	<i>0.0387</i>	-
	$\hat{\mu}_2$	-0.0036	0.0064	0.0166	-0.0063	0.0039	0.0138	-
		<i>0.0639</i>	<i>0.0642</i>	<i>0.0649</i>	<i>0.0608</i>	<i>0.0613</i>	<i>0.0610</i>	-
	\hat{F}	0.0968	-0.0032	-0.1032	0.0995	-0.0005	-0.1005	-
		<i>0.0487</i>	<i>0.0487</i>	<i>0.0487</i>	<i>0.0464</i>	<i>0.0464</i>	<i>0.0464</i>	-
\widehat{SUNK}	0.0032	0.0032	0.0032	0.0005	0.0005	0.0005	-0.0002	
	<i>0.0487</i>	<i>0.0487</i>	<i>0.0487</i>	<i>0.0464</i>	<i>0.0464</i>	<i>0.0464</i>	<i>0.0508</i>	

librium 1, and generally in smaller samples. We may be able to attribute the difference in smaller samples performance to the fact that Method A fully exploits the correctly specified parametric form of the payoff function while the others use nonparametric estimators. At larger sample sizes, there appear to be no dominating estimation methods for Equilibria 2 and 3.

5.2 Empirical illustration

We next estimate a simplified version of an entry-investment game based on the model studied in Ryan (2012); using the same dataset as him. In what follows, we provide a brief description of the data, and highlight the main differences between our empirical

TABLE 3. Data generated from equilibrium 3 in Pesendorfer and Schmidt-Dengler (2008)

N	W	Method A			Method B			N-N
		0	0.1	0.2	0	0.1	0.2	
100	$\hat{\beta}$	-0.0649	-0.0641	-0.0658	-0.0695	-0.0649	-0.0663	-
		<i>0.2459</i>	0.2427	<i>0.2472</i>	<i>0.2526</i>	0.2450	<i>0.2471</i>	-
	$\hat{\mu}_1$	-0.2070	-0.1907	-0.1725	-0.1116	-0.0986	-0.0807	-
		<i>0.3991</i>	0.4108	<i>0.4261</i>	<i>0.4724</i>	0.4804	<i>0.4920</i>	-
	$\hat{\mu}_2$	0.3263	0.3420	0.3588	0.1801	0.1961	0.2130	-
		<i>0.5460</i>	0.5484	<i>0.5551</i>	<i>0.7092</i>	0.7109	<i>0.7158</i>	-
	\hat{F}	-0.0677	-0.1676	-0.2672	0.0897	-0.0103	-0.1103	-
		<i>0.4224</i>	0.4227	<i>0.4230</i>	<i>0.5987</i>	0.5988	<i>0.5988</i>	-
\widehat{SUNK}	0.1677	0.1676	0.1672	0.0103	0.0103	0.0103	-0.0370	
	<i>0.4224</i>	0.4227	<i>0.4230</i>	<i>0.5987</i>	0.5988	<i>0.5988</i>	<i>0.6455</i>	
1000	$\hat{\beta}$	-0.0320	-0.0322	-0.0333	-0.0326	-0.0324	-0.0319	-
		<i>0.1634</i>	0.1643	<i>0.1666</i>	<i>0.1647</i>	0.1648	<i>0.1638</i>	-
	$\hat{\mu}_1$	-0.0237	-0.0104	0.0041	-0.0060	0.0071	0.0199	-
		<i>0.1677</i>	0.1796	<i>0.1932</i>	<i>0.1678</i>	0.1790	<i>0.1900</i>	-
	$\hat{\mu}_2$	0.0500	0.0633	0.0771	0.0251	0.0383	0.0511	-
		<i>0.2130</i>	0.2188	<i>0.2264</i>	<i>0.2174</i>	0.2235	<i>0.2305</i>	-
	\hat{F}	0.0766	-0.0234	-0.1234	0.1014	0.0014	-0.0986	-
		<i>0.1549</i>	0.1550	<i>0.1550</i>	<i>0.1604</i>	0.1604	<i>0.1604</i>	-
\widehat{SUNK}	0.0234	0.0234	0.0234	-0.0014	-0.0014	-0.0014	-0.0061	
	<i>0.1549</i>	0.1550	<i>0.1550</i>	<i>0.1604</i>	0.1604	<i>0.1604</i>	<i>0.1785</i>	
10,000	$\hat{\beta}$	-0.0003	-0.0003	-0.0003	-0.0001	-0.0004	-0.0002	-
		<i>0.0159</i>	0.0158	<i>0.0156</i>	<i>0.0093</i>	0.0163	<i>0.0128</i>	-
	$\hat{\mu}_1$	-0.0146	-0.0046	0.0054	-0.0128	-0.0026	0.0073	-
		<i>0.0410</i>	0.0414	<i>0.0420</i>	<i>0.0399</i>	0.0415	<i>0.0410</i>	-
	$\hat{\mu}_2$	-0.0033	0.0067	0.0167	-0.0062	0.0039	0.0138	-
		<i>0.0648</i>	0.0649	<i>0.0650</i>	<i>0.0646</i>	0.0650	<i>0.0650</i>	-
	\hat{F}	0.0965	-0.0035	-0.1035	0.0992	-0.0008	-0.1008	-
		<i>0.0496</i>	0.0496	<i>0.0496</i>	<i>0.0497</i>	0.0497	<i>0.0497</i>	-
\widehat{SUNK}	0.0035	0.0035	0.0035	0.0008	0.0008	0.0008	0.0002	
	<i>0.0496</i>	0.0496	<i>0.0496</i>	<i>0.0497</i>	0.0497	<i>0.0497</i>	<i>0.0553</i>	

model and that of Ryan (2012). Then we present and discuss our estimates of the model primitives.

Data We download Ryan's data from the Econometrica webpage.⁶ There are two sets of data. One contains aggregate prices and quantities for all the US regional markets from the US Geological Survey's Mineral Yearbook. The other contains the capacities of plants and plant-level information that Ryan has collected for the Portland cement industry in the United States from 1980 to 1998. Data on plants includes the name of the firm that owns the plant, the location of the plant, the number of kilns in the plant, and kiln characteristics. Following Ryan, we assume that the plant capacity equals the sum of the

⁶<https://www.econometricsociety.org/content/supplement-costs-environmental-regulation-concentrated-industry-0>.

capacity of all kilns in the plant and that different plants are owned by different firms. We observe that plants' names and ownerships change frequently. This can be due to either mergers and acquisitions or to simple changes in the company name. We do not treat these changes as entry/exit movements. We check each observation in the sample using the kiln information (fuel type, process type, year of installation, and plant location) installed in the plant. If a plant changes its name but keeps the same kiln characteristics, we assume that the name change is not associated to any entry/exit movement. This way of preparing the data enables us to match most of the summary statistics of plant-level data in Table 2 of Ryan. Any discrepancies most likely can be attributed to the way we treat the change in plants' names, which may differ to Ryan in a very small number of cases.

Dynamic game Ryan models a dynamic game played between firms that own cement plants in order to measure the welfare costs of the 1990 Clean Air Act Amendments (1990 CAAA) on the US Portland cement industry. The decision for each firm is first whether to enter (or remain in) the market or exit, and if it is active in the market then how much to invest or divest. Firm's investment decisions is governed by its capacity level. The firm's profit is determined by variable payoffs from the competition in the product market with other firms, as well as switching costs from the entry and investment/divestment decisions. There are two action variables in Ryan's model. One is a binary choice used to model entry. The other is a continuous variable used to model the level of investment. Past actions are the only observed endogenous state variables in the game. The aggregate data that are used to construct variable profits, through a static Cournot game with capacity constraints between firms, are treated as exogenous.

We consider a discrete game that extends the single agent model in the paper as described in the [Appendix](#). The main departure from [Ryan \(2012\)](#) is that we combine the entry decision along with the capacity level into a single discrete variable. We set the action space to be an ordinal set $\{0, 1, 2, 3, 4, 5\}$, where 0 represents exit/inactive, and the positive integers are ordered to denote entry/active with different capacity levels. The payoff for each firm has two additive separable components. One depends on the observables while the other is an unobserved shock. The observable component can be broken down into variable profits, operating cost, and switching costs. We assume the variable profit is determined by the players competing in a capacity constrained Cournot game. The operating cost is a fixed profit that incurs whenever $a_{it} > 0$. The switching costs capture the essence of firms' entry and investment decisions. Lastly, each firm receives unobserved profit shocks for each action with a standard i.i.d. type-1 extreme value distribution.

Estimation The period expected payoff for each firm as a function of the observables consists of variable profits, operating costs and switching costs. The variable profit is derived from a capacity constrained Cournot game constructed from the same demand and cost functions estimated as in Ryan's paper. The operating and switching costs parameters enter the payoff function additively and are parameters to be estimated using the dynamic model. These operating cost is nonzero whenever $a_{it} > 0$. For the switching costs, we normalize the payoff for choosing action 0 to be zero. There are a total of 25

switching cost parameters to be estimated.⁷ The payoff function used in our empirical model satisfies Assumptions MN and MP in the Appendix. So we estimate the model using Methods A and B as described in Section 5.1. We also test if the two estimates of the switching costs statistically differ. Instead of using nonparametric estimator, similar to Ryan, we use a multinomial logit to estimate the choice and transition probabilities in the first stage. More specifically, method A profiles out the 26 linear coefficients and uses grid search to estimate the discount factor. Method B first estimates the 25 switching cost parameters in closed-form using the closed-form expression in Section 4, treat them as known, before profiling and performing the grid search. We also estimate the sunk entry and investment values based on the estimates from Methods A and B, as well as nonparametrically without normalization (cf. Corollaries 1 and 2, and see the discussion in the Appendix).

We estimate the standard errors, as well as computing the p -value of the Wald statistics to test if the switching costs estimators from methods A and B differ by bootstrapping. Our bootstrap sample is generated using the multinomial logit choice and transition probabilities for each player in each market in the same manner as a parametric bootstrap; cf. Kasahara and Shimotsu (2008) and Pakes, Ostrovsky, and Berry (2007). We use 500 bootstrap samples and report the standard errors in *italics*.

Results We estimate the model twice. Once using the data from before the implementation of the 1990 CAAA and another after. We allow the equilibria over the two time periods to differ. But, for illustrational purposes, we assume the data are generated from the same equilibrium in all markets within each time period and there is no other source of unobserved heterogeneity.⁸

Tables 4 and 5 compile the results from estimating switching costs using the data from the years 1980 to 1990 and 1991 to 1998, respectively. Tables 6 and 7 give the estimates for the discount factor and fixed operating cost using the data from the corresponding periods. Table 8 compares the estimates of the sunk entry costs and sunk investment costs.

The signs and relative magnitudes of individually estimated switching costs almost uniformly make sensible economic sense. For example, by reading down the columns in Tables 4 and 5, we see that entering at higher capacity level generally implies higher cost (negative payoff), and increasing the capacity level should be costly while divestment can return revenue for firms. This is quite an impressive finding in particular for

⁷Ryan (2012) models the switching costs differently. The fixed operating cost is normalized to be zero. Nonzero investment and divestment costs are drawn from two distinct independent normal distributions, whose means and variances are estimated using the methodology described in Bajari, Benkard, and Levin (2007).

⁸Recently Otsu, Pesendorfer, and Takahashi (2016) propose several tests to detect differences in the probability distribution of data across markets. If a test rejects then there is evidence data across markets should not be pooled together, which can point to possible violation of single equilibrium assumption and/or misspecification in terms of omitting other unobserved heterogeneity. They actually suggest Ryan's data in general should not be pooled together across markets. In particular, there is a strong evidence against pooling data between 1980 and 1990, while the data from 1991 to 1998 did not get rejected by some of their poolability tests.

TABLE 4. Results from estimating switching costs using data from the years 1980 to 1990.

Method A						
	$a_{it-1}=0$	$a_{it-1}=1$	$a_{it-1}=2$	$a_{it-1}=3$	$a_{it-1}=4$	$a_{it-1}=5$
$a_{it}=1$	-3.300 <i>0.985</i>	- <i>0.985</i>	2.265 <i>0.680</i>	5.080 <i>0.707</i>	7.956 <i>0.766</i>	10.770 <i>0.929</i>
$a_{it}=2$	-10.502 <i>0.937</i>	-5.243 <i>0.719</i>	- <i>0.719</i>	5.528 <i>0.887</i>	10.609 <i>0.998</i>	15.810 <i>1.117</i>
$a_{it}=3$	-23.266 <i>1.405</i>	-15.439 <i>1.010</i>	-7.624 <i>0.683</i>	- <i>0.683</i>	7.996 <i>0.923</i>	16.050 <i>1.237</i>
$a_{it}=4$	-41.023 <i>2.003</i>	-30.620 <i>1.850</i>	-20.196 <i>1.430</i>	-9.808 <i>1.094</i>	- <i>1.094</i>	11.648 <i>1.442</i>
$a_{it}=5$	-52.879 <i>2.281</i>	-50.648 <i>2.585</i>	-39.027 <i>2.041</i>	-25.756 <i>1.395</i>	-11.949 <i>1.537</i>	- <i>1.537</i>

Method B						
	$a_{it-1}=0$	$a_{it-1}=1$	$a_{it-1}=2$	$a_{it-1}=3$	$a_{it-1}=4$	$a_{it-1}=5$
$a_{it}=1$	-2.776 <i>0.269</i>	- <i>0.269</i>	2.540 <i>0.333</i>	5.333 <i>0.567</i>	8.014 <i>0.967</i>	11.696 <i>1.113</i>
$a_{it}=2$	-10.483 <i>0.689</i>	-5.197 <i>0.365</i>	- <i>0.365</i>	5.243 <i>0.368</i>	10.466 <i>0.718</i>	15.893 <i>1.110</i>
$a_{it}=3$	-23.279 <i>1.339</i>	-15.427 <i>0.920</i>	-7.769 <i>0.474</i>	- <i>0.474</i>	7.732 <i>0.640</i>	16.134 <i>1.006</i>
$a_{it}=4$	-41.422 <i>1.808</i>	-31.007 <i>1.594</i>	-20.797 <i>1.078</i>	-10.416 <i>0.682</i>	- <i>0.682</i>	10.852 <i>0.864</i>
$a_{it}=5$	-54.378 <i>1.911</i>	-52.892 <i>2.232</i>	-41.874 <i>1.844</i>	-28.792 <i>1.659</i>	-16.091 <i>1.835</i>	- <i>1.835</i>

Specification	Test
Statistic	14.069
<i>p</i> -value	<i>0.961</i>

Method B, which shows that the observed probabilities alone can generate switching costs estimates that capture well some key features of a complicated structural model. The switching cost estimates from both Methods A and B are similar. The Wald statistics do not find the two switching costs estimators to be statistically different.⁹ Therefore, we do not reject the capacity constrained Cournot game specification based on comparing the switching costs estimates. Comparing Tables 4 and 5 shows the entry and switching costs increase after the implementation of 1990 CAAA. Higher entry costs is a key finding in Ryan's paper as new entrants face more stringent regulations than incumbents. An increase in switching costs can be partly attributed to the new plants using newer

⁹Our test statistic takes a standard quadratic form of the difference between the switching costs estimates from methods A and B. Its asymptotic distribution under the null hypothesis (of no difference) is a Chi-squared random variable with 25 degree of freedoms.

TABLE 5. Results from estimating switching costs using data from the years 1991 to 1998.

Method A						
	$a_{it-1}=0$	$a_{it-1}=1$	$a_{it-1}=2$	$a_{it-1}=3$	$a_{it-1}=4$	$a_{it-1}=5$
$a_{it}=1$	-6.962 <i>1.530</i>	- -	4.449 <i>1.514</i>	9.881 <i>1.501</i>	15.125 <i>1.689</i>	20.264 <i>1.634</i>
$a_{it}=2$	-17.038 <i>1.723</i>	-8.291 <i>1.364</i>	- -	9.872 <i>1.714</i>	18.531 <i>1.860</i>	26.722 <i>1.527</i>
$a_{it}=3$	-35.489 <i>2.444</i>	-23.412 <i>1.866</i>	-11.411 <i>1.371</i>	- -	12.961 <i>1.955</i>	24.283 <i>1.614</i>
$a_{it}=4$	-51.544 <i>3.061</i>	-50.043 <i>3.419</i>	-33.220 <i>3.278</i>	-16.363 <i>2.825</i>	- -	16.524 <i>3.561</i>
$a_{it}=5$	-64.018 <i>4.514</i>	-63.994 <i>4.524</i>	-61.481 <i>4.502</i>	-48.514 <i>3.683</i>	-24.374 <i>2.056</i>	

Method B						
	$a_{it-1}=0$	$a_{it-1}=1$	$a_{it-1}=2$	$a_{it-1}=3$	$a_{it-1}=4$	$a_{it-1}=5$
$a_{it}=1$	-5.653 <i>0.726</i>	- -	5.294 <i>0.704</i>	10.730 <i>1.109</i>	16.264 <i>1.703</i>	21.567 <i>1.378</i>
$a_{it}=2$	-17.746 <i>1.379</i>	-9.278 <i>0.780</i>	- -	8.774 <i>0.857</i>	17.461 <i>1.364</i>	25.754 <i>1.218</i>
$a_{it}=3$	-36.098 <i>2.282</i>	-24.537 <i>1.767</i>	-11.950 <i>1.128</i>	- -	11.862 <i>1.221</i>	23.489 <i>1.401</i>
$a_{it}=4$	-51.840 <i>2.202</i>	-50.425 <i>2.649</i>	-33.468 <i>2.397</i>	-16.760 <i>1.904</i>	- -	16.753 <i>2.025</i>
$a_{it}=5$	-64.236 <i>6.712</i>	-64.355 <i>6.771</i>	-61.706 <i>6.713</i>	-48.272 <i>5.695</i>	-24.093 <i>3.389</i>	

Specification	Test
Statistic	13.196
<i>p</i> -value	<i>0.975</i>

TABLE 6. Results from estimating the discount factor and fixed operating cost using data from the years 1980 to 1990.

Method A	
Discount factor	Operating cost
0.956	-1.679
<i>0.084</i>	<i>0.489</i>

Method B	
Discount factor	Operating cost
0.999	-1.523
<i>0.075</i>	<i>0.649</i>

TABLE 7. Results from estimating the discount factor and fixed operating cost using data from the years 1991 to 1998.

Method A	
Discount factor	Operating cost
0.938	-2.079
<i>0.162</i>	<i>1.10</i>
Method B	
Discount factor	Operating cost
0.946	-1.893
<i>0.160</i>	<i>0.948</i>

(or better maintained) equipment that require more certification and testing than previously.

We find the discount factor estimates to be around the range that are usually assumed in empirical work (between 0.9 and 0.95) apart from the estimate using Method B before the 1990 CAAA that appears close to the boundary.¹⁰ Although our estimates suggest firms face a lower borrowing rate than in Ryan, we do not reject the hypothesis that $\beta = 0.9$ as assumed in his paper. We also find a small increase in the fixed operating costs after the implementation of 1990 CAAA.

Finally Table 8 reports sunk costs using different estimation methods. The estimates from Methods A and B can be found by computing $-SC^{a' \rightarrow a} - SC^{a \rightarrow a'}$ using individual switching costs in Tables 4 and 5. The N-N approach estimates the same object without the assumption that the payoff is zero upon choosing action 0. The signs and magnitudes of the sunk cost estimates are plausible. We find the sunk investment costs between any two capacity levels increase as the gap between levels grow, while we find the costs to be of similar magnitude when compared within the same capacity difference bands. We also find the sunk costs to have increased after the implementation of 1990 CAAA.

6. CONCLUDING REMARKS

We show the discount factor can be identified jointly with the payoff function under the linear-in-parameter specification. The key property we exploit is the conditional linearity of the choice-specific value functions for a given value of the discount factor. The discount factor can in fact be identified even if the payoff parameters cannot be identified. This has an important implication since many empirical problems have to normalize

¹⁰The infinite time expected discounted payoffs with respect to each action is unbounded with $\beta = 1$. However, the differences between diverge very slowly when we approximate them with a Neumann sum, and the objective function appears to be well-defined numerically even as β is very close to 1.

TABLE 8. Results from estimating the sunk entry and investment costs.

a_{it}	a_{it-1}	Before			After		
		Method A	Method B	N-N	Method A	Method B	N-N
1	0	3.30	2.78	2.78	6.96	5.65	5.66
		<i>0.36</i>	<i>0.27</i>	<i>0.27</i>	<i>1.53</i>	<i>0.73</i>	<i>0.70</i>
2	0	10.50	10.48	10.48	17.04	17.75	17.74
		<i>0.94</i>	<i>0.69</i>	<i>0.69</i>	<i>1.72</i>	<i>1.38</i>	<i>1.49</i>
3	0	23.27	23.28	23.28	35.49	36.10	36.10
		<i>1.41</i>	<i>1.34</i>	<i>1.34</i>	<i>2.44</i>	<i>2.28</i>	<i>2.18</i>
4	0	41.02	41.42	41.42	51.54	51.84	51.83
		<i>2.00</i>	<i>1.81</i>	<i>1.80</i>	<i>3.06</i>	<i>2.20</i>	<i>1.61</i>
5	0	52.88	54.38	54.25	64.02	64.24	64.22
		<i>2.28</i>	<i>1.91</i>	<i>2.00</i>	<i>4.51</i>	<i>6.71</i>	<i>6.34</i>
2	1	2.98	2.66	2.44	3.84	3.98	3.30
		<i>1.22</i>	<i>2.54</i>	<i>0.25</i>	<i>0.31</i>	<i>0.61</i>	<i>0.36</i>
3	2	2.10	2.53	2.56	1.54	3.18	3.22
		<i>1.18</i>	<i>2.30</i>	<i>0.26</i>	<i>0.30</i>	<i>0.73</i>	<i>0.33</i>
4	3	1.81	2.68	2.58	3.40	4.90	4.81
		<i>1.52</i>	<i>4.33</i>	<i>0.28</i>	<i>0.42</i>	<i>2.45</i>	<i>0.50</i>
5	4	0.30	5.24	2.87	7.85	7.34	7.30
		<i>2.50</i>	<i>4.75</i>	<i>0.33</i>	<i>1.74</i>	<i>4.58</i>	<i>2.14</i>
3	1	10.36	10.09	10.01	13.53	13.81	13.05
		<i>1.22</i>	<i>2.12</i>	<i>0.75</i>	<i>0.79</i>	<i>1.24</i>	<i>0.98</i>
4	2	9.59	10.33	10.29	14.69	16.01	16.07
		<i>1.54</i>	<i>3.31</i>	<i>0.77</i>	<i>0.81</i>	<i>2.13</i>	<i>1.25</i>
5	3	9.71	12.66	10.91	24.23	24.78	24.21
		<i>1.45</i>	<i>4.83</i>	<i>0.91</i>	<i>1.37</i>	<i>6.09</i>	<i>5.22</i>
4	1	22.66	22.99	22.76	34.92	34.16	34.02
		<i>1.78</i>	<i>3.29</i>	<i>1.37</i>	<i>1.45</i>	<i>1.93</i>	<i>1.42</i>
5	2	23.22	25.98	24.05	34.76	35.95	34.79
		<i>1.83</i>	<i>4.64</i>	<i>1.79</i>	<i>1.59</i>	<i>6.89</i>	<i>6.34</i>
5	1	39.88	41.20	40.21	43.73	42.79	41.67
		<i>2.40</i>	<i>4.68</i>	<i>2.60</i>	<i>2.08</i>	<i>6.82</i>	<i>6.40</i>

parts of the payoff parameters. Our result shows the discount factor can be identified independently of these normalization choices.

We also contribute to a recent interest in the robust identification of combinations of switching costs without any normalization as studied in [Aguirregabiria and Suzuki \(2014\)](#); also see [Kalouptsi, Scott, and Souza-Rodrigues \(2016, 2017\)](#). We provide closed-form identification results on switching costs that do not depend on the knowledge of the discount factor and other parts of the payoff function. We show some costs, such as sunk entry and investment costs, can be identified in this way. We show the same combinations of switching costs can be identified for linear models in two steps. In the first step, some normalization is made in order to identify each switching cost individually. Even when an incorrect normalization is used, thus the implied switching costs are in-

correct individually, we show certain combinations of these costs can still be correctly identified.

Our parametric and nonparametric identification approaches deliver quite different flavors of results. But there are overlapping implications when the payoff function satisfies both Assumptions **N** and **P**, as we then have two different ways to identify the switching costs. However, there are notable distinctions where our nonparametric results remain valid but the analysis under Assumption **P** is no longer appropriate. First, a researcher may want to use a nonlinear parametric specification on parts of the payoffs outside of the switching costs. One example of this is to impose positivity on the variable profits. Our nonparametric identification results do not depend on the specification of the variable profit function. Second, our nonparametric identification strategy holds pointwise for each observed state. Therefore, it is immediately applicable for models with continuous states; for example, see [Srisuma and Linton \(2012\)](#).

Finally, our main message is that one should generally attempt to identify and estimate the discount factor in dynamic decision problems and games. Clearly, we do not expect the linear specification to be necessary for identification. But analyzing models with nonlinear parametric payoff functions will be substantially more difficult. Similarly, outside of discrete choice models, for example, for games with supermodular payoff functions (see [Bajari, Benkard, and Levin \(2007\)](#) and [Srisuma \(2013\)](#)), joint identification and estimation of the discount factor and payoff parameters should also be possible. However, in this case even the practical implementation can be burdensome when the payoff functions take a linear-in-parameter structure.

APPENDIX

The Appendix contains two parts: [A.1](#) extends the results on identification of switching costs to dynamic games. [A.2](#) provides a sufficient condition for the identification of the discount factor. Since the single agent decision problem is a special case of a game, we also present the results in [A.2](#) in the context of a game.

A.1 Identification of the switching costs in dynamic games

We shall keep our description of the basic elements of the game very brief. The notation we use directly extends what we describe in Sections [2](#) and [3](#). Consider a game with I players, indexed by $i \in \mathcal{I} = \{1, \dots, I\}$. The random variables in the game are the actions: $a_t \equiv (a_{it}, a_{-it}) \in A^I$, $A = \{0, 1, \dots, J\}$; past actions $w_t \equiv (w_{it}, w_{-it}) \in A^I$; $s_{it} \equiv (w_t, x_t, \varepsilon_{it}) \in A^I \times X \times \mathbb{R}^{J+1}$, where $X = \{1, \dots, K\}$, and $\varepsilon_{it} \equiv (\varepsilon_{it}(0), \dots, \varepsilon_{it}(J)) \in \mathbb{R}^{J+1}$; and we let $s_t \equiv (w_t, x_t, \varepsilon_{1t}, \dots, \varepsilon_{It})$.

In an equilibrium $a_{it} = \alpha_i(s_{it})$ for all i , such that

$$\begin{aligned} \alpha_i(s_i) = \max_{a_i \in A} \{ & E[u_i(a_{it}, a_{-it}, s_i) | s_{it} = s_i, a_{it} = a_i] \\ & + \beta E[V_i(s_{it+1}) | s_{it} = s_i, a_{it} = a_i] \}, \end{aligned} \quad (20)$$

where u_i and V_i are player i 's payoff and value function respectively; in particular,

$$V_i(s_i) = \sum_{t=0}^{\infty} \beta^t E[u_i(a_{it}, a_{-it}, s_{it}) | s_{i0} = s_i].$$

Assumption MN updates Assumptions M and N for games.

ASSUMPTION MN.

(i) (*Additive Separability*) For all $a_i, a_{-i}, w, x, \varepsilon_i$:

$$u_i(a_i, a_{-i}, w, x, \varepsilon_i) = \pi_i(a_i, a_{-i}, w, x) + \varepsilon_i(a_i).$$

(ii) (*Conditional Independence I*) The transition distribution of the states has the following factorization for all $x', \varepsilon', x, \varepsilon, a$:

$$P(x', \varepsilon' | x, \varepsilon, w, a) = \prod_{i=1}^I Q_i(\varepsilon'_i) G(x' | x, w, a),$$

where Q_i is the cumulative distribution function of ε_{it} and G denotes the transition law of x_{t+1} conditioning on x_t, a_t . Furthermore, ε_{it} has finite first moments, and a positive, continuous, and bounded density on \mathbb{R}^{J+1} .

(iii) (*Finite Observed State*) $X = \{1, \dots, K\}$.

(iv) (*Decomposition of Profits*): For all a, w, x :

$$\pi_i(a_i, a_{-i}, w, x, \varepsilon) = \mu_i(a_i, a_{-i}, x) + \phi_i(a_i - w_i, w_{-i}, x),$$

such that $\phi_i(0, w_{-i}, x) = 0$.

(v) (*Conditional Independence II*): The distribution of x_{t+1} conditional on a_t and x_t is independent of w_t .

Beside from explicitly separating out past actions from other observed state variables, MN(i) to MN(iii) are standard in the dynamic discrete choice game literature; for example, see Aguirregabiria and Mira (2007), Bajari, Benkard, and Levin (2007), Pakes, Ostrovsky, and Berry (2007), and Pesendorfer and Schmidt-Dengler (2008). MN(iv) extends N(i). It assumes that strategic interactions can affect payoffs in μ_i directly but not ϕ_i , while past actions enter ϕ_i but not μ_i . The exclusion restrictions we impose are quite natural for components of μ_i such as per-period variable profits and operation costs, while switching costs that occur for each player are determined by her own actions. It will be useful to sometimes represent the switching cost using a more intuitive notation (cf. equation (11)):

$$\phi_i(a_i - w_i, w_{-i}, x) = SC_i^{w_i \rightarrow a_i}(w_{-i}, x).$$

MN(v) is a direct extension of N(ii).

As with the single agent case, our identification study will be based on the choice-specific value function:

$$v_i(a_i, w, x) = E[\pi_i(a_i, a_{-it}, w_t, x)|w_t = w, x_t = x] + \beta E[V_i(s_{t+1})|w_t = w, x_t = x, a_t = a],$$

which can be recovered from

$$\begin{aligned} & \Pr[a_{it} = a_i | w_t = w, x_t = x] \\ &= \Pr[\Delta v_i(a_i, w, x) - \Delta v_i(a'_i, w, x) > \varepsilon_{it}(a'_i) - \varepsilon_{it}(a_i) \text{ for all } a'_i \neq a_i], \end{aligned}$$

where $\Delta v_i(a_i, w, x) \equiv v_i(a_i, w, x) - v_i(0, w, x)$. Let also, $\Delta \lambda_i(a_i, a_{-i}, x) \equiv \lambda_i(a_i, a_{-i}, x) - \lambda_i(0, a_{-i}, x)$ and $\Delta \phi_i(a_i, w, x) \equiv \phi_i(a_i - w_i, w_{-i}, x) - \phi_i(-w_i, w_{-i}, x)$. Lemma 4 is a generalization of Lemma 1.

LEMMA 4. *Under Assumption MN, we have for all $i, a_i > 0$ and w, x :*

$$\Delta v_i(a_i, w, x) = E[\Delta \lambda_i(a_i, a_{-it}, x)|w_t = w, x_t = x] + \Delta \phi_i(a_i, w, x),$$

where

$$\begin{aligned} \Delta \lambda_i(a_i, a_{-i}, x) &\equiv \pi_i(a_i, a_{-i}, x) - \pi_i(0, a_{-i}, x) + \beta(\tilde{m}_i(a_i, a_{-i}, x) - \tilde{m}_i(0, a_{-i}, x)), \\ \tilde{m}_i(a_i, a_{-i}, x) &\equiv E[m_i(w_{t+1}, x_{t+1})|a_{it} = a_i, a_{-it} = a_{-i}, x_t = x], \\ m_i(w, x) &\equiv E[V_i(s_{it})|w_t = w, x_t = x]. \end{aligned}$$

PROOF. It follows immediately from applying the law of iterated expectations (cf. the proof of Lemma 1). \square

Since we have finite actions and states, we can collect $\Delta v_i(a_i, w, x)$ across w for each (i, a_i, x) into a vector of size $(J+1)^I$. Using a matrix form, we have

$$\Delta \mathbf{v}_i(a_i, x) = \mathbf{Z}_i(x) \Delta \boldsymbol{\lambda}_i(a_i, x) + \mathbf{Q}_i(a_i, x) \boldsymbol{\phi}_i(a_i, x), \quad (21)$$

where $\Delta \mathbf{v}_i(a_i, x) = (\Delta v_i(a_i, w, x))_{w \in A^I}$, $\Delta \boldsymbol{\lambda}_i(a_i, x) = (\Delta \lambda_i(a_i, a_{-i}, x))_{a_{-i} \in A^{I-1}}$, $\mathbf{Z}_i(x)$ represents the matrix of conditional probabilities for computing a conditional expectation of a_{-it} given $(w_t = w, x_t = x)$, $\mathbf{Q}_i(a_i, x) \boldsymbol{\phi}_i(a_i, x)$ represents $(\Delta \phi_i(a_i, w, x))_{a \in A^I}$ with $\boldsymbol{\phi}_i(a_i, x) = (\phi_i(a_i - w_i, w_{-i}, x))_{w_i \in A, w_{-i} \in A^{I-1}}$ and $\mathbf{Q}_i(a_i, x)$ is a matrix of indicators (consisting of 0's and 1's) that pick up switching costs as appropriate.

Theorem 3 generalizes the closed-form identification of switching costs in Theorem 1 for dynamic games.

THEOREM 3. *Assume that Assumption MN holds. Let \mathbf{D} be an $\ell_1 \times (J+1)^I$ matrix with $\rho(\mathbf{D}) = \ell_1$ such that $(J+1)^{I-1} < \ell_1 \leq (J+1)^I$. Denote $\mathbf{D}\mathbf{Z}_i(x)$ by $\tilde{\mathbf{Z}}$ and $\rho(\tilde{\mathbf{Z}})$ by ℓ_2 . Suppose also $\mathbf{D}\mathbf{Q}_i(a_i, x) \boldsymbol{\phi}_i = \tilde{\mathbf{Q}} \tilde{\boldsymbol{\phi}} + \boldsymbol{\phi}_0$ for some ℓ_3 -dimensional vectors $\tilde{\boldsymbol{\phi}}$ and $\boldsymbol{\phi}_0$ that consist of elements, possibly combinations, of $\boldsymbol{\phi}_i$ such that $\ell_3 \leq \ell_1 - \ell_2$, and $\tilde{\mathbf{Q}}$ is an $\ell_1 \times \ell_3$ matrix with $\rho(\tilde{\mathbf{Q}}) = \ell_3$. If $\rho([\tilde{\mathbf{Z}} : \tilde{\mathbf{Q}}]) = \ell_2 + \ell_3$, then*

$$\tilde{\boldsymbol{\phi}} = (\tilde{\mathbf{Q}}^\top \tilde{\mathbf{P}} \tilde{\mathbf{Q}})^{-1} \tilde{\mathbf{Q}}^\top \tilde{\mathbf{P}} (\mathbf{D} \Delta \mathbf{v}_i(a_i, x) - \boldsymbol{\phi}_0), \quad (22)$$

where $\tilde{\mathbf{P}} = I_{\ell_1} - \tilde{\mathbf{Z}}(\tilde{\mathbf{Z}}^\top \tilde{\mathbf{Z}})^\dagger \tilde{\mathbf{Z}}^\top$.

Before presenting the proof to Theorem 3 some explanations on the notation will be useful. The crucial interpretation of our result rests on the relation: $\mathbf{D}\mathbf{Q}_i(a_i, x)\boldsymbol{\phi}_i = \tilde{\mathbf{Q}}\tilde{\boldsymbol{\phi}} + \boldsymbol{\phi}_0$. The goal of Theorem 3 is to identify components, or combinations, of $(\phi_i(a_i, w, x))_{w \in A^I}$ using choice-specific value functions in equation (21) for a given (i, a_i, x) . We denote the object of interest by $\tilde{\boldsymbol{\phi}}$. We use $\boldsymbol{\phi}_0$ to account for components of switching costs that can be identified outside the dynamic model from the data or by normalization. Therefore, $(\mathbf{D}, \tilde{\mathbf{Q}})$ are user-chosen matrices and are completely known. For identification, we can also treat $\tilde{\mathbf{Z}}_i$ as known since $\mathbf{Z}_i(x)$ is a matrix of observed choice probabilities.

PROOF OF THEOREM 3. Note that $\ell_3 \geq 1$ since $\ell_2 \leq \min\{\ell_1, \rho(\mathbf{Z}_i(x))\}$ and $\rho(\mathbf{Z}_i(x)) \leq (J + 1)^{J-1}$. Multiply equation (21) by \mathbf{D} yields

$$\mathbf{D}\Delta\mathbf{v}_i(a_i, x) = \tilde{\mathbf{Z}}\Delta\boldsymbol{\lambda}_i(a_i, x) + \tilde{\mathbf{Q}}\tilde{\boldsymbol{\phi}} + \boldsymbol{\phi}_0.$$

By assumption, $\tilde{\mathbf{P}}\tilde{\mathbf{Q}}$ has full column rank. The result then follows from projecting $\Delta\mathbf{v}_i(a_i, x)$ orthogonally onto the null space of $\tilde{\mathbf{Z}}$ and solve out for $\tilde{\boldsymbol{\phi}}_i$. \square

One systematic approach to apply Theorem 3 in practice is to first write out the matrix equation (21). Then choose \mathbf{D} so that $\mathbf{D}\mathbf{Q}_i(a_i, x)\boldsymbol{\phi}_i$ contains the switching costs of interest, and define $\tilde{\mathbf{Q}}\tilde{\boldsymbol{\phi}} + \boldsymbol{\phi}_0$ appropriately. We now illustrate this identifying strategy with a two-player binary choice game for different types of switching costs.

For notational compactness, we will suppress x_t and assume that $SC_i^{w \rightarrow a}(w_{-i})$ is the same for all w_{-i} . We use $\Delta v_i(w_i, w_{-i}) \equiv v_i(1, w_i, w_{-i}) - v_i(0, w_i, w_{-i})$, $p_{-i}(w) \equiv \Pr[a_{-it} = 1 | w_t = w]$, and $\Delta\lambda_i(a_{-i}) \equiv \Delta\lambda_i(1, a_{-i})$. Then equation (21) represents:

$$\begin{aligned} \begin{bmatrix} \Delta v_i(0, 0) \\ \Delta v_i(0, 1) \\ \Delta v_i(1, 0) \\ \Delta v_i(1, 1) \end{bmatrix} &= \begin{bmatrix} 1 - p_{-i}(0, 0) & p_{-i}(0, 0) \\ 1 - p_{-i}(0, 1) & p_{-i}(0, 1) \\ 1 - p_{-i}(1, 0) & p_{-i}(1, 0) \\ 1 - p_{-i}(1, 1) & p_{-i}(1, 1) \end{bmatrix} \begin{bmatrix} \Delta\lambda_i(0) \\ \Delta\lambda_i(1) \end{bmatrix} \\ &+ \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} SC_i^{0 \rightarrow 1} \\ -SC_i^{1 \rightarrow 0} \end{bmatrix}. \end{aligned} \tag{23}$$

In particular, we have

$$\mathbf{Q}_i(a_i, x)\boldsymbol{\phi}_i = \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & 1 \end{bmatrix} \begin{bmatrix} SC_i^{0 \rightarrow 1} \\ -SC_i^{1 \rightarrow 0} \end{bmatrix}.$$

We consider three examples of potential objects of interest.

EXAMPLE 1 (Sunk Entry Cost). Suppose we want to identify $-SC_i^{0 \rightarrow 1} - SC_i^{1 \rightarrow 0}$ that represents the sunk entry cost in the context of an entry game. We can subtract $\Delta v_i(0, 0)$ from

the first equation in (23) off the remaining three equations. This yields

$$\begin{aligned} \begin{bmatrix} \Delta v_i(0, 1) \\ \Delta v_i(1, 0) \\ \Delta v_i(1, 1) \end{bmatrix} &= \begin{bmatrix} p_{-i}(0, 0) - p_{-i}(0, 1) & p_{-i}(0, 1) - p_{-i}(0, 0) \\ p_{-i}(0, 0) - p_{-i}(1, 0) & p_{-i}(1, 0) - p_{-i}(0, 0) \\ p_{-i}(0, 0) - p_{-i}(1, 1) & p_{-i}(1, 1) - p_{-i}(0, 0) \end{bmatrix} \begin{bmatrix} \Delta \lambda_i(0) \\ \Delta \lambda_i(1) \end{bmatrix} \\ &+ \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix} [-SC_i^{0 \rightarrow 1} - SC_i^{1 \rightarrow 0}]. \end{aligned}$$

In particular, in this case,

$$\begin{aligned} \tilde{\mathbf{Z}} &= \begin{bmatrix} p_{-i}(0, 0) - p_{-i}(0, 1) & p_{-i}(0, 1) - p_{-i}(0, 0) \\ p_{-i}(0, 0) - p_{-i}(1, 0) & p_{-i}(1, 0) - p_{-i}(0, 0) \\ p_{-i}(0, 0) - p_{-i}(1, 1) & p_{-i}(1, 1) - p_{-i}(0, 0) \end{bmatrix}, \\ \mathbf{D} &= \begin{bmatrix} -1 & 1 & 0 & 0 \\ -1 & 0 & 1 & 0 \\ -1 & 0 & 0 & 1 \end{bmatrix}, \quad \tilde{\mathbf{Q}} = \begin{bmatrix} 0 \\ 1 \\ 1 \end{bmatrix}, \quad \tilde{\boldsymbol{\phi}} = -SC_i^{0 \rightarrow 1} - SC_i^{1 \rightarrow 0}, \quad \text{and} \quad \boldsymbol{\phi}_0 = 0. \end{aligned}$$

The sunk entry cost can then be identified by the expression in equation (22).

EXAMPLE 2 (Menu Cost Under Symmetry). Suppose we want to identify $SC_i^{0 \rightarrow 1}$ under the assumption that $SC_i^{0 \rightarrow 1} = SC_i^{1 \rightarrow 0}$. Then equation (23) becomes

$$\begin{aligned} \begin{bmatrix} \Delta v_i(0, 0) \\ \Delta v_i(0, 1) \\ \Delta v_i(1, 0) \\ \Delta v_i(1, 1) \end{bmatrix} &= \begin{bmatrix} 1 - p_{-i}(0, 0) & p_{-i}(0, 0) \\ 1 - p_{-i}(0, 1) & p_{-i}(0, 1) \\ 1 - p_{-i}(1, 0) & p_{-i}(1, 0) \\ 1 - p_{-i}(1, 1) & p_{-i}(1, 1) \end{bmatrix} \begin{bmatrix} \Delta \lambda_i(0) \\ \Delta \lambda_i(1) \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \\ -1 \\ -1 \end{bmatrix} [SC_i^{0 \rightarrow 1}]. \end{aligned}$$

In this case,

$$\begin{aligned} \tilde{\mathbf{Z}} &= \begin{bmatrix} 1 - p_{-i}(0, 0) & p_{-i}(0, 0) \\ 1 - p_{-i}(0, 1) & p_{-i}(0, 1) \\ 1 - p_{-i}(1, 0) & p_{-i}(1, 0) \\ 1 - p_{-i}(1, 1) & p_{-i}(1, 1) \end{bmatrix}, \quad \mathbf{D} = I_4, \\ \tilde{\mathbf{Q}} &= \begin{bmatrix} 1 \\ 1 \\ -1 \\ -1 \end{bmatrix}, \quad \tilde{\boldsymbol{\phi}} = SC_i^{0 \rightarrow 1}, \quad \text{and} \quad \boldsymbol{\phi}_0 = 0. \end{aligned}$$

EXAMPLE 3 (Switching Costs With Normalizations). Suppose we want to identify $SC_i^{0 \rightarrow 1}$ under the assumption that $SC_i^{0 \rightarrow 1} = c_0$. For example, we may be interested in identify-

ing the entry cost under the assumption that the scrap value is c_0 . Then equation (23) becomes

$$\begin{bmatrix} \Delta v_i(0, 0) \\ \Delta v_i(0, 1) \\ \Delta v_i(1, 0) \\ \Delta v_i(1, 1) \end{bmatrix} = \begin{bmatrix} 1 - p_{-i}(0, 0) & p_{-i}(0, 0) \\ 1 - p_{-i}(0, 1) & p_{-i}(0, 1) \\ 1 - p_{-i}(1, 0) & p_{-i}(1, 0) \\ 1 - p_{-i}(1, 1) & p_{-i}(1, 1) \end{bmatrix} \begin{bmatrix} \Delta \lambda_i(0) \\ \Delta \lambda_i(1) \end{bmatrix} + \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \end{bmatrix} [SC_i^{0 \rightarrow 1}] + \begin{bmatrix} 0 \\ 0 \\ -c_0 \\ -c_0 \end{bmatrix}.$$

In this case,

$$\tilde{\mathbf{Z}} = \begin{bmatrix} 1 - p_{-i}(0, 0) & p_{-i}(0, 0) \\ 1 - p_{-i}(0, 1) & p_{-i}(0, 1) \\ 1 - p_{-i}(1, 0) & p_{-i}(1, 0) \\ 1 - p_{-i}(1, 1) & p_{-i}(1, 1) \end{bmatrix}, \quad \mathbf{D} = I_4,$$

$$\tilde{\mathbf{Q}} = \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \end{bmatrix}, \quad \tilde{\boldsymbol{\phi}} = SC_i^{0 \rightarrow 1}, \quad \text{and} \quad \boldsymbol{\phi}_0 = \begin{bmatrix} 0 \\ 0 \\ -c_0 \\ -c_0 \end{bmatrix}.$$

In order to obtain the sunk costs when the number of actions is larger than two, one has to combine identifiable objects across actions; for example, see Corollary 2. Identification of objects for each action can be obtained as the examples above have shown. We use Theorem 3 to estimate the games such as those in our simulation study and the empirical model of capacity game in Section 5 of our paper.

A.2 A sufficient condition for identification of the discount factor

In this part of the Appendix, we give a more analytical approach that ensures identification of the discount factor and payoff parameters in a dynamic game context. We first introduce some additional notation.

For any $x = (x_1, \dots, x_p)^\top \in \mathbb{R}^p$ and $y = (y_1, \dots, y_{p+1})^\top \in \mathbb{R}^{p+1}$, let $\|x\|_{\alpha_1} = \max_{i=1, \dots, p} |x_i|$ and $\|y\|_{\alpha_2} = \max_{i=1, \dots, p} |y_i| + |y_{p+1}|$. Then for a class of $p + 1$ by p real matrices, we denote the matrix norms induced by $(\|\cdot\|_{\alpha_1}, \|\cdot\|_{\alpha_2})$ by $\|\cdot\|_{\alpha_1, \alpha_2}$. We comment that these are not standard induced matrix norms; however, they have simple explicit bounds. In particular, it is easy to verify that, for any matrix $(p + 1) \times p$, $C = (c_{ij})$,

$$\|C\|_{\alpha_1, \alpha_2} \leq \max_{i=1, \dots, p} \sum_{j=1}^p |c_{ij}| + \sum_{j=1}^p |c_{p+1, j}|.$$

We also need the parameter space to be compact. Let $\bar{\boldsymbol{\theta}} \equiv \{\theta \in \boldsymbol{\theta} : \max_{i=1, \dots, p} |\theta_i| \leq \bar{k}\}$ and $\bar{\mathcal{B}} \equiv [0, \bar{b}]$ for some positive \bar{k} and $\bar{b} \in (0, 1)$.

Next, we generalize the setup of Section 4 to dynamic games. The following is a straightforward extension of Assumptions M and P.

ASSUMPTION MP.

(i) (*Additive Separability*) For all $a_i, a_{-i}, x, \varepsilon_i$:

$$u_i(a_i, a_{-i}, x, \varepsilon_i; \theta) = \pi_i(a_i, a_{-i}, x; \theta) + \varepsilon_i(a_i).$$

(ii) (*Conditional Independence I*) The transition distribution of the states has the following factorization for all $x', \varepsilon', x, \varepsilon, a$:

$$P(x', \varepsilon' | x, \varepsilon, w, a) = \prod_{i=1}^I Q_i(\varepsilon'_i) G(x' | x, w, a),$$

where Q_i is the cumulative distribution function of ε_{it} and G denotes the transition law of x_{t+1} conditioning on x_t, a_t . Furthermore, ε_{it} has finite first moments, and a positive, continuous, and bounded density on \mathbb{R}^{J+1} .

(iii) (*Finite Observed State*) $X = \{1, \dots, K\}$.

(iv) (*Linear-in-Parameters*): For all $a_i, a_{-i}, x, \varepsilon_i$:

$$\pi_i(a_i, a_{-i}, x; \theta) = \pi_{i0}(a_i, a_{-i}, x) + \theta^\top \pi_{i1}(a_i, a_{-i}, x),$$

where π_{i0} is a known real value function, π_{i1} is a known p -dimensional vector value function and θ belongs to \mathbb{R}^p .

Our analysis will be based on the parameterized choice-specific value function:

$$v_i(a_i, x; \beta, \theta) = E[\pi_i(a_i, a_{-it}, x; \theta) | x_t = x] + \beta E[V_i(s_{t+1}; \beta, \theta) | x_t = x, a_{it} = a_i], \quad \text{where}$$

$$V_i(s_i; \beta, \theta) = \sum_{t=0}^{\infty} \beta^t E[u_i(a_{it}, a_{-it}, s_{it}; \theta) | s_{i0} = s_i].$$

Let $\Delta v_i(a_i, x; \beta, \theta) \equiv v_i(a_i, x; \beta, \theta) - v_i(0, x; \beta, \theta)$. We can use Δv_i from all players to define an empirical model and the corresponding notion of identification, and observationally equivalence, as in Section 4. We will omit this discussion to avoid repetition.

Our starting point will be the following lemma that generalizes Lemma 2.

LEMMA 5. Under Assumption MP, we have for all $i, a_i > 0$, $\Delta \mathbf{v}_i^{a_i}(\beta, \theta) \equiv (\Delta v_i(a_i, x; \beta, \theta))_{x \in X}$ can be collected in the following vector form for all $(\beta, \theta) \in \mathcal{B} \times \Theta$:

$$\begin{aligned} \Delta \mathbf{v}_i^{a_i}(\beta, \theta) &= \Delta \mathbf{R}_{i0}^{a_i} + \beta \Delta \mathbf{H}_i^{a_i} (\mathbf{I}_K - \beta \mathbf{L})^{-1} \mathbf{R}_{i0} \\ &\quad + (\Delta \mathbf{R}_{i1}^{a_i} + \beta \Delta \mathbf{H}_i^{a_i} (\mathbf{I}_K - \beta \mathbf{L})^{-1} \mathbf{R}_{i1}) \theta \\ &\quad + \beta \Delta \mathbf{H}_i^{a_i} (\mathbf{I}_K - \beta \mathbf{L})^{-1} \boldsymbol{\varepsilon}_i, \end{aligned} \tag{24}$$

where the elements in the above display are collected and explained in Tables C and D.

Our strategy to show identification is to rewrite Lemma 5 in order to set up a mapping that has the data generating parameters its fixed point. One desired relation is the following.

TABLE C. The matrices consist of (differences in) expected payoffs and probabilities. The latter represent conditional expectations for any function ψ of x_{t+1} .

Matrix	Dimension	Representing
$\Delta \mathbf{R}_{i1}^{a_i}$	$K \times p$	$E[\pi_{i1}(a_i, a_{-it}, x_t) - \pi_{i1}(0, a_{-it}, x_t) x_t = \cdot]$
\mathbf{R}_1	$K \times p$	$E[\pi_{i1}(a_t, x_t) x_t = \cdot]$
\mathbf{L}	$K \times K$	$E[\psi(x_{t+1}) x_t = \cdot]$
$\mathbf{H}_i^{a_i}$	$K \times K$	$E[\psi(x_{t+1}) x_t = \cdot, a_{it} = a_i]$
$\Delta \mathbf{H}_i^{a_i}$	$K \times K$	$E[\psi(x_{t+1}) x_t = \cdot, a_{it} = a_i] - E[\psi(x_{t+1}) x_t = \cdot, a_{it} = 0]$

TABLE D. The $K \times 1$ vectors represent (differences in) expected payoffs.

Vector	Representing
$\boldsymbol{\epsilon}_i$	$E[\varepsilon_{it}(a_{it}) x_t = \cdot]$
$\Delta \mathbf{R}_{i0}^{a_i}$	$E[\pi_{i0}(a_i, a_{-it}, x_t) - \pi_{i0}(0, a_{-it}, x_t) x_t = \cdot]$
\mathbf{R}_{i0}	$E[\pi_{i0}(a_t, x_t) x_t = \cdot]$
$(\mathbf{I}_K - \beta_i \mathbf{L})^{-1} \mathbf{R} \mathbf{H}_{ij}$	$\sum_{t=0}^{\infty} \beta^t E[\pi_{ij}(a_t, x_t) x_0 = \cdot]$
$\Delta \mathbf{H}_i^{a_i} (\mathbf{I}_K - \beta_i \mathbf{L})^{-1} \mathbf{R} \mathbf{H}_{ij}$	$\sum_{t=0}^{\infty} \beta^t (E[\pi_{ij}(a_t, x_t) a_{i0} = a_i, x_0 = \cdot] - E[\pi_{ij}(a_t, x_t) a_{i0} = 0, x_0 = \cdot])$
$\Delta \mathbf{H}_i^{a_i} (\mathbf{I}_K - \beta_i \mathbf{L})^{-1} \boldsymbol{\epsilon}_i$	$\sum_{t=0}^{\infty} \beta^t (E[\varepsilon_t(a_t) a_{i0} = a_i, x_0 = \cdot] - E[\varepsilon_t(a_t) a_{i0} = 0, x_0 = \cdot])$

LEMMA 6. Under Assumption MP, (β, θ) is observationally equivalent to (β_0, θ_0) if and only if (β, θ) satisfies

$$\mathbf{c}_i^{a_i} - \mathbf{D}_i^{a_i}(\beta)\theta - \mathbf{E}_i(\beta) = \mathbf{F}_i^{a_i} \begin{pmatrix} \theta \\ \beta \end{pmatrix} \quad (25)$$

for all $i, a_i > 0$, where

$$\begin{aligned} \mathbf{c}_i^{a_i} &= \Delta \mathbf{v}_i^{a_i}(\beta_0, \theta_0) - \Delta \mathbf{R}_{i0}^{a_i}, \\ \mathbf{D}_i^{a_i}(\beta) &= \beta \Delta \mathbf{H}_i^{a_i} (\mathbf{I}_K - \beta \mathbf{L})^{-1} \mathbf{R}_{i1}, \\ \mathbf{E}_i(\beta) &= \beta^2 \Delta \mathbf{H}_i^{a_i} \mathbf{L} (\mathbf{I}_K - \beta \mathbf{L})^{-1} (\mathbf{R}_{i0} + \boldsymbol{\epsilon}_i), \\ \mathbf{F}_i^{a_i} &= [\Delta \mathbf{R}_{i1}^{a_i} : \Delta \mathbf{H}_i^{a_i} (\mathbf{R}_{i0} + \boldsymbol{\epsilon}_i)]. \end{aligned}$$

PROOF. Equation (25) is obtained by re-arranging equation (24), after applying the identity that $(\mathbf{I}_K - \beta \mathbf{L})^{-1} = \mathbf{I}_K + \beta \mathbf{L} (\mathbf{I}_K - \beta \mathbf{L})^{-1}$ and replace $\Delta \mathbf{v}_i^{a_i}(\beta, \theta)$ by $\Delta \mathbf{v}_i^{a_i}(\beta_0, \theta_0)$. Therefore, by construction, (β, θ) satisfies (24) if and only if it is observationally equivalent to (β_0, θ_0) . \square

The following result provides one condition that is sufficient for the identification of (β_0, θ_0) .

THEOREM 4. Assume that $K \geq p + 1$ and Assumption MP holds. Suppose there exists i, a_i such that: (i) the rank of $\mathbf{F}_i^{a_i}$ is $p + 1$; (ii) there exists a $p + 1$ by K matrix \mathbf{A}_0 such that $\mathbf{A}_0 \mathbf{F}_i^{a_i}$

is nonsingular; and (iii) $\max\{\mathbf{g}_1, \mathbf{g}_2\} < 1$, where

$$\begin{aligned}\mathbf{g}_1 &= \max_{\beta \in \overline{\mathcal{B}}} \left\| (\mathbf{A}_0 \mathbf{F}_i^{a_i})^{-1} \mathbf{A}_0 \Delta \mathbf{H}_i^{a_i} \beta (\mathbf{I}_K - \beta \mathbf{L})^{-1} \mathbf{R}_{1i} \right\|_{\alpha_1, \alpha_2}, \\ \mathbf{g}_2 &= \max_{\beta, \beta' \in \overline{\mathcal{B}}, \theta \in \overline{\Theta}} \left\| (\mathbf{A}_0 \mathbf{F}_i^{a_i})^{-1} \mathbf{A}_0 \Delta \mathbf{H}_i^{a_i} ((\mathbf{I}_K - \beta \mathbf{L})^{-1} (\mathbf{I}_K - \beta' \mathbf{L})^{-1} \mathbf{R}_{1i} \theta \right. \\ &\quad \left. + \mathbf{L} (\mathbf{I}_K - \beta \mathbf{L})^{-1} ((\beta + \beta') \mathbf{I}_K - \beta \beta' \mathbf{L}) (\mathbf{I}_K - \beta' \mathbf{L})^{-1} (\mathbf{R}_{0i} + \boldsymbol{\epsilon}_i)) \right\|_{\alpha_1, \alpha_2}.\end{aligned}$$

Then (β_0, θ_0) is identifiable.

PROOF. First, define $\mathcal{Q}_i^{a_i} : [0, 1] \times \Theta_k \rightarrow \mathbb{R}^{p+1}$ as follows:

$$\mathcal{Q}_i^{a_i}(\beta, \theta) = (\mathbf{A}_0 \mathbf{F}_i^{a_i})^{-1} \mathbf{A}_0 \mathbf{c}_i^{a_i} - (\mathbf{A}_0 \mathbf{F}_i^{a_i})^{-1} \mathbf{A}_0 \mathbf{D}_i^{a_i}(\beta) \theta - (\mathbf{A}_0 \mathbf{F}_i^{a_i})^{-1} \mathbf{A}_0 \mathbf{E}_i(\beta).$$

By construction, from (25), it is easy to see that (β_0, θ_0) is a fixed-point of \mathcal{Q} . Take any $(\beta, \theta), (\beta', \theta') \in \overline{\mathcal{B}} \times \overline{\Theta}$, then

$$\mathcal{Q}_i^{a_i}(\beta, \theta) - \mathcal{Q}_i^{a_i}(\beta', \theta') = -(\mathbf{A}_0 \mathbf{F}_i^{a_i})^{-1} \mathbf{A}_0 (\mathbf{D}_i^{a_i}(\beta) \theta - \mathbf{D}_i^{a_i}(\beta') \theta' + \mathbf{E}_i(\beta) - \mathbf{E}_i(\beta')),$$

where

$$\begin{aligned}\mathbf{D}_i^{a_i}(\beta) \theta - \mathbf{D}_i^{a_i}(\beta') \theta' &= \Delta \mathbf{H}_i^{a_i} (\beta (\mathbf{I}_K - \beta \mathbf{L})^{-1} \mathbf{R}_{1i} \theta - \beta' (\mathbf{I}_K - \beta' \mathbf{L})^{-1} \mathbf{R}_{1i} \theta') \\ &= \Delta \mathbf{H}_i^{a_i} ((\beta - \beta') (\mathbf{I}_K - \beta \mathbf{L})^{-1} (\mathbf{I}_K - \beta' \mathbf{L})^{-1} \mathbf{R}_{1i} \theta \\ &\quad + \beta' (\mathbf{I}_K - \beta' \mathbf{L})^{-1} \mathbf{R}_{1i} (\theta - \theta')), \end{aligned}$$

and

$$\begin{aligned}\mathbf{E}_i(\beta) - \mathbf{E}_i(\beta') &= \Delta \mathbf{H}_i^{a_i} \mathbf{L} (\beta^2 (\mathbf{I}_K - \beta \mathbf{L})^{-1} - \beta'^2 (\mathbf{I}_K - \beta' \mathbf{L})^{-1}) (\mathbf{R}_{i0} + \boldsymbol{\epsilon}_i) \\ &= \Delta \mathbf{H}_i^{a_i} \mathbf{L} ((\beta - \beta') (\mathbf{I}_K - \beta \mathbf{L})^{-1} ((\beta + \beta') \mathbf{I}_K - \beta \beta' \mathbf{L}) (\mathbf{I}_K - \beta' \mathbf{L})^{-1}) (\mathbf{R}_{i0} + \boldsymbol{\epsilon}_i),\end{aligned}$$

which can be shown by making use of the following identities:

$$\begin{aligned}\beta (\mathbf{I}_K - \beta \mathbf{L})^{-1} - \beta' (\mathbf{I}_K - \beta' \mathbf{L})^{-1} &= (\beta - \beta') (\mathbf{I}_K - \beta \mathbf{L})^{-1} (\mathbf{I}_K - \beta' \mathbf{L})^{-1}, \\ \beta^2 (\mathbf{I}_K - \beta \mathbf{L})^{-1} - \beta'^2 (\mathbf{I}_K - \beta' \mathbf{L})^{-1} &= (\beta - \beta') (\mathbf{I}_K - \beta \mathbf{L})^{-1} ((\beta + \beta') \mathbf{I}_K - \beta \beta' \mathbf{L}) (\mathbf{I}_K - \beta' \mathbf{L})^{-1}.\end{aligned}$$

It then follows that

$$\begin{aligned}|\mathcal{Q}_i^{a_i}(\beta, \theta) - \mathcal{Q}_i^{a_i}(\beta', \theta')| &\leq \mathbf{g}_1 \|\theta - \theta'\|_{\alpha_1} + \mathbf{g}_2 |\beta - \beta'| \\ &\leq \max\{\mathbf{g}_1, \mathbf{g}_2\} \left\| \begin{pmatrix} \theta \\ \beta \end{pmatrix} - \begin{pmatrix} \theta' \\ \beta' \end{pmatrix} \right\|_{\alpha_2}.\end{aligned}$$

That is, $\mathcal{Q}_i^{a_i}$ is a contraction, hence it has a unique fixed point. Now suppose (β_0, θ_0) is not identifiable. Then there exists some $(\beta, \theta) \neq (\beta_0, \theta_0)$ that is observationally equivalent to (β_0, θ_0) . By an implication of Lemma 6, (β, θ) must also be a fixed point of $\mathcal{Q}_i^{a_i}$, which is a contradiction. Thus (β_0, θ_0) is identifiable. \square

COMMENTS ON THEOREM 4.

(i) *Compact Domain.* \mathcal{B} cannot include 1 as the expected discounted returns would then be unbounded. Compactness is useful for showing existence of a fixed point. There is also a trade-off in the choice of \bar{b} and \bar{k} in the definitions of $\bar{\mathcal{B}}$ and $\bar{\Theta}$, respectively. For example, smaller \bar{b} and \bar{k} means smaller $\max\{\mathbf{g}_1, \mathbf{g}_2\}$ but this is a restriction on the parameter space.

(ii) *Choice of \mathbf{A}_0 .* The need to select \mathbf{A}_0 can be eliminated altogether by removing some rows in (25) so that we have exactly $p + 1$ equations. In fact, it is not necessary to take equations that only correspond to the states from a particular player i and a_i . Since the parametric structure in (25) is the same for all states, we can select any $p + 1$ equations from any i and a_i and compute the corresponding matrix norms for \mathbf{g}_1 and \mathbf{g}_2 . This gives us different combinations of equations we can use, and we only need the analog of $\max\{\mathbf{g}_1, \mathbf{g}_2\}$ to be less than 1 for one of them to ensure (β_0, θ_0) is identifiable.

(iii) *Rank Deficiency.* We have emphasized in Section 4 that sometimes not all components of the payoff functions can be identified and normalizations are necessary. For example, in the entry/exit game generally the entry cost and scrap value cannot be jointly identified. Then one may consider normalizing, say, the scrap value in order to estimate all the other parameters in the model. Furthermore, we discussed in Section 4.3 that the discount factor can be identified even if an incorrect normalization is used. Relatedly, we can also relax condition (i) in Theorem 4 in this direction and allow $\mathbf{F}_i^{a_i}$ to be rank deficient. In particular, recall from (25) that $\mathbf{F}_i^{a_i} = [\Delta\mathbf{R}_{i1}^{a_i} : \Delta\mathbf{H}_i^{a_i}(\mathbf{R}_{i0} + \boldsymbol{\epsilon}_i)]$, we can allow $\Delta\mathbf{R}_{i1}^{a_i}$ to be rank deficient. In such case, there exists a full rank matrix \mathbf{W} such that $\Delta\mathbf{R}_{i1}^{a_i}\mathbf{W} = [\tilde{\Delta}\mathbf{R}_{i1}^{a_i} : 0]$ where $\tilde{\Delta}\mathbf{R}_{i1}^{a_i}$ has full column rank. Then $\mathbf{F}_i^{a_i}(\frac{\theta}{\beta})$ in (25) becomes $[\tilde{\Delta}\mathbf{R}_{i1}^{a_i} : 0 : \Delta\mathbf{H}_i^{a_i}(\mathbf{R}_{i0} + \boldsymbol{\epsilon}_i)](\frac{\mathbf{W}^{-1}\theta}{\beta})$. Therefore, by inspection, the proof of Theorem 4 can be readily adapted by reparameterizing θ to show the identification of the discount factor is possible.

REFERENCES

- Ackerberg, D., L. Benkard, S. Berry, and A. Pakes (2007), “Econometric tools for analyzing market outcome.” In *Handbook of Econometrics*, Vol. 6 (J. Heckman and E. Leamer, eds.). North-Holland, Amsterdam. [1155]
- Aguirregabiria, V. and P. Mira (2002), “Swapping the nested fixed point algorithm: A class of estimators for discrete Markov decision models.” *Econometrica*, 70, 1519–1543. [1156]
- Aguirregabiria, V. and P. Mira (2007), “Sequential estimation of dynamic discrete games.” *Econometrica*, 75, 1–53. [1156, 1183]

Aguirregabiria, V. and A. Nevo (2013), "Recent developments in empirical IO: Dynamic demand and dynamic games." In *Advances in Economics and Econometrics, Tenth World Congress, Vol. 3, Econometrics*, Cambridge University Press, Cambridge. [1157]

Aguirregabiria, V. and J. Suzuki (2014), "Identification and counterfactuals in dynamic models of market entry and exit." *Quantitative Marketing and Economics*, 12, 267–304. [1154, 1156, 1159, 1166, 1169, 1170, 1181]

Arcidiacono, P. and R. A. Miller (2011), "Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity." *Econometrica*, 79, 1823–1867. [1156]

Bajari, P., C. L. Benkard, and J. Levin (2007), "Estimating dynamic models of imperfect competition." *Econometrica*, 75, 1331–1370. [1177, 1182, 1183]

Bajari, P., V. Chernozhukov, H. Hong, and D. Nekipelov (2009), "Identification and efficient estimation of a dynamic discrete game." Working Paper, University of Minnesota. [1154, 1155]

Bajari, P., H. Hong, and D. Nekipelov (2013), "Game theory and econometrics: A survey of some recent research." In *Advances in Economics and Econometrics, Tenth World Congress, Vol. 3, Econometrics*, Cambridge University Press, Cambridge. [1157]

Buchholz, N., M. Shum, and H. Xu (2016), "Semiparametric estimation of a dynamic discrete choice models." Working Paper, California Institute of Technology. [1159]

Chen, L. (2014), "Identification of discrete choice dynamic programming models with nonparametric distribution of unobservables." Working Paper, Institute of Economics, Academia Sinica. [1159]

Ching, A. and M. Osborne (2017), "Identification and estimation of forward-looking behavior: The case of consumer stockpiling." Working Paper, University of Toronto. [1154, 1155]

Dubé, J. P., G. Hitsch, and P. Jindal (2014), "The joint identification of utility and discount functions from stated choice data: An application to durable goods adoption." *Quantitative Marketing and Economics*, 12, 331–377. [1154, 1155]

Egedal, M., Z. Lai, and C. Su (2015), "Estimating dynamic discrete-choice games of incomplete information." *Quantitative Economics*, 6, 567–597. [1156]

Fang, H. and Y. Wang (2015), "Estimating dynamic discrete choice models with hyperbolic discounting, with an application to mammography decisions." *International Economic Review*, 56, 565–596. [1154, 1155]

Hong, H. and M. Shum (2010), "Pairwise-difference estimation of a dynamic optimization model." *Review of Economic Studies*, 77, 273–304. [1156]

Hotz, V. and R. A. Miller (1993), "Conditional choice probabilities and the estimation of dynamic models." *Review of Economic Studies*, 60, 497–531. [1156, 1159, 1160]

- Kalouptsi, M. (2014), "Time to build and fluctuations in bulk shipping." *American Economic Review*, 104, 564–608. [1169]
- Kalouptsi, M., P. Scott, and E. Souza-Rodrigues (2016), "Identification of counterfactuals in dynamic discrete choice models." NBER Working Paper 21527. [1154, 1181]
- Kalouptsi, M., P. Scott, and E. Souza-Rodrigues (2017), "On the non-identification of counterfactuals in dynamic discrete games." *International Journal of Industrial Organization*, 50, 362–371. [1154, 1181]
- Kasahara, H. and K. Shimotsu (2008), "Pseudo-likelihood estimation and bootstrap inference for structural discrete Markov decision models." *Journal of Econometrics*, 146, 92–106. [1156, 1177]
- Kasahara, H. and K. Shimotsu (2009), "Nonparametric identification of finite mixture models of dynamic discrete choices." *Econometrica*, 77, 135–175. [1156]
- Komunjer, I. (2012), "Global identification in nonlinear models with moment restrictions." *Econometric Theory*, 28, 719–729. [1155]
- Magnac, M. and D. Thesmar (2002), "Identifying dynamic discrete decision processes." *Econometrica*, 70, 801–816. [1154, 1155]
- Manski, C. (1993), "Dynamic choice in social settings." *Journal of Econometrics*, 58, 121–136. [1154, 1163]
- Matzkin, R. L. (1991), "Semiparametric estimation of monotone and concave utility functions for polychotomous choice models." *Econometrica*, 59, 1315–1327. [1160]
- McFadden, D. L. (1974), "Conditional logit analysis of qualitative choice behavior." In *Frontier in Econometrics*, Academic Press, New York. [1154]
- Mirsky, L. (1955), *An Introduction to Linear Algebra*. Oxford University Press, London. [1171]
- Myśliwski, M., F. Sanches, D. Silva, Jr., and S. Srisuma (2017), "Dynamic pricing and consumer loyalty in a multiproduct oligopoly." Working Paper, University College London. [1169]
- Norets, A. and S. Takahashi (2013), "On the surjectivity of the mapping between utilities and choice probabilities." *Quantitative Economics*, 4, 149–155. [1160]
- Otsu, T., M. Pesendorfer, and Y. Takahashi (2016), "Pooling data across markets in dynamic Markov games." *Quantitative Economics*, 7, 523–559. [1177]
- Pakes, A., M. Ostrovsky, and S. Berry (2007), "Simple estimators for the parameters of discrete dynamic games (with entry/exit example)." *RAND Journal of Economics*, 38, 373–399. [1177, 1183]
- Pesendorfer, M. (2013), "Estimation of (dynamic) games: A discussion." In *Advances in Economics and Econometrics, Tenth World Congress, Vol. 3, Econometrics*, Cambridge University Press, Cambridge. [1155]

Pesendorfer, M. and P. Schmidt-Dengler (2008), "Asymptotic least squares estimator for dynamic games." *Review of Economics Studies*, 75, 901–928. [1154, 1155, 1157, 1172, 1173, 1174, 1175, 1183]

Rust, J. (1987), "Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher." *Econometrica*, 55, 999–1033. [1156, 1158]

Rust, J. (1994), "Estimation of dynamic structural models: Problems and prospects part I: Discrete decision processes." In *Proceedings of the 6th World Congress of the Econometric Society* (C. A. Sims, ed.), 119–171, Cambridge University Press, Cambridge. [1154, 1155, 1158]

Ryan, S. P. (2012), "The costs of environmental regulation in a concentrated industry." *Econometrica*, 80, 1019–1061. [1153, 1157, 1172, 1174, 1175, 1176, 1177]

Sanches, F., D. Silva, Jr., and S. Srisuma (2016), "Ordinary least squares estimation of a dynamic game model." *International Economic Review*, 57, 623–634. [1156, 1161]

Slade, M. E. (1998), "Optimal pricing with costly adjustment: Evidence from retail-grocery prices." *Review of Economic Studies*, 65, 87–107. [1155, 1169]

Srisuma, S. (2013), "Minimum distance estimators for dynamic games." *Quantitative Economics*, 4, 549–583. [1182]

Srisuma, S. (2015), "Identification in Markov decision models." *Econometric Theory*, 31, 521–538. [1162, 1165]

Srisuma, S. and O. B. Linton (2012), "Semiparametric estimation of Markov decision processes with continuous state space." *Journal of Econometrics*, 166, 320–341. [1182]

Wang, Y. (2014), "Dynamic implications of subjective expectations: Evidence from adult smokers." *American Economic Journal: Applied Economics*, 6, 1–37. [1154, 1155]

Co-editor Rosa L. Matzkin handled this manuscript.

Manuscript received 4 February, 2016; final version accepted 13 December, 2017; available online 30 January, 2018.