

DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

A SIMPLE ESTIMATOR FOR DYNAMIC MODELS WITH SERIALLY
CORRELATED UNOBSERVABLES

Yingyao Hu
Johns Hopkins

Matthew Shum
Caltech

Wei Tan
SUNY-Stony Brook

SOCIAL SCIENCE WORKING PAPER 1324

May 2010

A Simple Estimator for Dynamic Models with Serially Correlated Unobservables

Yingyao Hu

Matthew Shum

Wei Tan

Abstract

We present a method for estimating Markov dynamic models with unobserved state variables which can be serially correlated over time. We focus on the case where all the model variables have discrete support. Our estimator is simple to compute because it is noniterative, and involves only elementary matrix manipulations. Our estimation method is nonparametric, in that no parametric assumptions on the distributions of the unobserved state variables or the laws of motions of the state variables are required. Monte Carlo simulations show that the estimator performs well in practice, and we illustrate its use with a dataset of doctors' prescription of pharmaceutical drugs.

JEL classification numbers: C13, C33, C51

Key words: dynamic models, serially-correlated unobservables, identification

A Simple Estimator for Dynamic Models with Serially Correlated Unobservables*

Yingyao Hu

Matthew Shum

Wei Tan

1 Introduction

In this paper, we consider nonparametric identification in Markovian dynamic models where the agent may have a serially-correlated unobserved state variable. These models have been the basis for much of the recent empirical applications of dynamic models. Throughout, by “unobservable”, we mean variables which are observed by agents, and affect their decisions, but are unobserved by the researcher.

Consider a dynamic optimization model described by the sequence of variables

$$\{(W_{t+1}, X_{t+1}^*), (W_t, X_t^*), \dots, (W_1, X_1^*)\}$$

where W_t denotes the observed variables for the optimizing agent in period t , and X_t^* denotes the unobserved variables, which we allow to vary over time and be serially-correlated.

In empirical dynamic models, the observed variables W_t typically consists of two types of variables:

$$W_t \equiv (Y_t, M_t),$$

where Y_t denotes the choice, or control variable in period t , and M_t denotes the state variables which are observed by both the optimizing agent and the researcher. We assume that the serially-correlated variable X_t^* is observed by the agent prior to making his choice of Y_t in period t , but the researcher never observes X_t^* . For simplicity, we assume that the variables Y_t, M_t, X_t^* are each scalar-valued.

Main Results: This paper focuses on the identification and estimation of the density

$$f(W_t, X_t^* | W_{t-1}, X_{t-1}^*), \tag{1}$$

*The authors can be reached at yhu@jhu.edu, mshum@caltech.edu, and wtan@notes.cc.sunysb.edu. We thank Wei Zhao for extraordinary research assistance.

which corresponds to the law of motion of the choice and state variables along the optimal path of the dynamic optimization problem. In Markovian dynamic settings, the law of motion can be factored into two components of interest:

$$\begin{aligned}
 f(W_t, X_t^* | W_{t-1}, X_{t-1}^*) &= f(Y_t, M_t, X_t^* | Y_{t-1}, M_{t-1}, X_{t-1}^*) \\
 &= \underbrace{f(Y_t | M_t, X_t^*)}_{\text{CCP}} \cdot \underbrace{f(M_t, X_t^* | Y_{t-1}, M_{t-1}, X_{t-1}^*)}_{\text{law of motion for state variables}}.
 \end{aligned} \tag{2}$$

The first term denotes the conditional choice probabilities (CCP) for the agents' actions in period t , conditional on the current state (M_t, X_t^*) . In the Markov setting, agents' optimal strategies typically depends just on the current state variables (M_t, X_t^*) , but not past values. The second term is the Markovian law of motion for the state variables (M_t, X_t^*) , along the dynamically optimal path. As shown in Hotz and Miller (1993) and Magnac and Thesmar (2002), once these two structural components are known, it is possible to recover the “deep” structural elements of the model, including the period utility functions.

In Section 2, we show that, under reasonable assumptions, three observations of $\{W_t, W_{t-1}, W_{t-2}\}_i$ across many agents i suffices to identify the law of motion $f(W_t, X_t^* | W_{t-1}, X_{t-1}^*)$. Moreover, when the model variables (Y_t, M_t, X_t^*) are all discrete, the identification arguments are constructive, and lead naturally to a simple estimation procedure involving only elementary matrix manipulations.

Section 3 contains results from simulation exercise, which highlights the good performance of our estimator even in moderately-sized samples. In section 4, we present an empirical illustration to advertising and doctors' prescription behavior in a pharmaceutical drug market.

Related literature Recently, there has been a growing literature related to identification and estimation of dynamic optimization models. Papers include Hotz and Miller (1993), Rust (1994), Aguirregabiria and Mira (2002), Magnac and Thesmar (2002), Hong and Shum (2009), Kasahara and Shimotsu (2009).¹ Our main contribution relative to this literature is to provide nonparametric identification and estimation results for the case where there are agent-specific unobserved state variables, which are serially correlated over time.²

¹A parallel literature has also developed in dynamic games; see Aguirregabiria and Mira (2007), Pesendorfer and Schmidt-Dengler (2008), Bajari, Benkard, and Levin (2007), Pakes, Ostrovsky, and Berry (2007), and Bajari, Chernozhukov, Hong, and Nekipelov (2007).

²The class of models considered in this paper also resemble models analyzed in the dynamic treatment effects literature in labor economics (eg. Cunha, Heckman, and Schennach (2006), Abbring and Heckman

Furthermore, our identification procedure is novel because it is based on recent developments in measurement error econometrics. Specifically, we show that Hu’s (2008) identification results for misclassification models can be applied to Markovian dynamic optimization models, and we use those results to establish nonparametric identification.

A few recent papers have considered estimation methodologies for dynamic models with serially-correlated unobservables. Imai, Jain, and Ching (2009) and Norets (2009) consider Bayesian estimation of these models, and Arcidiacono and Miller (2006) develop an EM-algorithm for estimating dynamic games where the unobservables are assumed to follow a discrete Markov process. Siebert and Zulehner (2008) estimate a dynamic product choice game for the computer memory industry where each firm experiences a serially-correlated productivity shock, and Gallant, Hong, and Khwaja (2009) and Blevins (2008) develops simulation estimators for dynamic games with serially-correlated unobservables, utilizing state-of-the-art recursive importance sampling (“particle filtering”) techniques. However, all these papers focus on estimation of parametric models in which the parameters are assumed to be identified, whereas this paper concerns estimation based on nonparametric identification results.

Finally, the models we consider in this paper fall under the rubric of “Hidden state” Markov (HSM) models, which have been considered in the computer science and machine learning literature (see Ghahramani (2001) for a survey). The identification results contained in this paper are new relative to this literature; moreover, the estimator we propose here has the virtue of being *non-iterative*, which makes it attractive relative to the EM-algorithm, an iterative procedure which is typically used to estimate HSM models.

2 Nonparametric identification in discrete dynamic models

Here we present our main identification result. For each agent i , $\{W_1, \dots, W_T\}_i$ is observed, for $T \geq 3$. Let $\Omega_{<t} = \{W_{t-1}, \dots, W_1, X_{t-1}^*, \dots, X_1^*\}$ denote the history of the process up to period $t - 1$.

Assumption 1

$$\textit{First-order Markov: } f(W_t, X_t^* | W_{t-1}, X_{t-1}^*, \Omega_{<t-1}) = f(W_t, X_t^* | W_{t-1}, X_{t-1}^*); \quad (3)$$

(2007), Heckman and Navarro (2007)).

Assumption 2 *Limited feedback*

- (i) $f(Y_t | M_t, X_t^*, Y_{t-1}, M_{t-1}, X_{t-1}^*) = f(Y_t | M_t, X_t^*),$
- (ii) $f(X_t^* | M_t, Y_{t-1}, M_{t-1}, X_{t-1}^*) = f(X_t^* | M_t, M_{t-1}, X_{t-1}^*).$

The first-order Markov assumption is standard in most empirical dynamic models, beginning from the the genre-creating papers of Rust (1987), Pakes (1986), Miller (1984), Keane and Wolpin (1994). Assumption (2) limits the feedback patterns in the Markov law of motion $f(W_t, X_t^* | W_{t-1}, X_{t-1}^*)$. Assumption 2(i) is motivated completely by the state-contingent aspect of the optimal policy function in Markov dynamic optimization models. This assumption is actually stronger than necessary for identification, but it allows us to achieve identification only using three periods of data.

Assumption 2(ii) implies that X_t^* is independent of Y_{t-1} conditional on M_t, M_{t-1} and X_{t-1}^* . Hence, it eliminates direct feedback from Y_{t-1} to X_t^* , but it allows for indirect feedback via M_t and M_{t-1} . In practice, this assumption should be verified on a model-by-model basis.³

Next, we restrict attention to stationary Markovian dynamic models. In a stationary setting, the law of motion $f(W_t, X_t^* | W_{t-1}, X_{t-1}^*)$ is time-invariant. In what follows, we will remove time subscripts from time-invariant functions, and just use primes (')s to denote next period values.

Assumption 3 *Stationarity of Markov kernel:*

$$f(W_t, X_t^* | W_{t-1}, X_{t-1}^*) = f(W', X'^* | W, X^*), \quad \forall 0 \leq t \leq T.$$

Finally, this paper focuses on the case where all the model variables are discrete:

Assumption 4 *For all periods t , Y_t and X_t^* are discrete-valued with J points of support. Without loss of generality,⁴ we assume that their common support is $\{1, 2, 3, \dots, J\}$.*

³The limited feedback assumption here is more restrictive than in Hu and Shum (2008), which used the limited feedback assumption

$$f(W_t^* | X_t^*, W_{t-1}, X_{t-1}^*) = f(W_t^* | X_t^*, W_{t-1}).$$

However, there is a tradeoff in that the extra restrictions in this paper allow us to achieve identification given only three time-contiguous observations per agent, whereas four observations are required for identification in Hu and Shum (2008).

⁴This is without loss of generality because our identification is fully nonparametric, and does not rely on particular functional forms. That is, for any arbitrary function $f(Y_{1t})$, where Y_{1t} has discrete support $\{y^1, y^2, \dots, y^J\}$ we could define another function $\tilde{f}(\tilde{y})$ such that $\tilde{f}(\tilde{y}) = f(Y_{1t} = y^{\tilde{y}})$ for all $\tilde{y} = 1, \dots, J$.

2.1 Identification argument: overview

The conditional independence assumptions 1-2 imply that the Markov law of motion (1) can be factored into

$$\begin{aligned} f(W', X^* | W, X^*) &= f(Y', M', X^* | Y, M, X^*) \\ &= \underbrace{f(Y' | M', X^*)}_{\text{CCP}} \cdot \underbrace{f(X^* | M', M, X^*)}_{\text{Law of motion for } X^*} \cdot \underbrace{f(M' | Y, M, X^*)}_{\text{Law of motion for } M}. \end{aligned} \quad (4)$$

We will identify these three components of $f(W_t, X_t^* | W_{t-1}, X_{t-1}^*)$ in turn.

Consider the joint distribution of $\{Y_t, M_t, Y_{t-1}, M_{t-1}, Y_{t-2}\}$, which is observed in the data. As shown in the Appendix, Assumptions 1-2 imply that

$$\begin{aligned} &f(Y_t, M_t, Y_{t-1} | M_{t-1}, Y_{t-2}) \\ &= \sum_{X_{t-1}^*} f(Y_t | M_t, M_{t-1}, X_{t-1}^*) f(M_t, Y_{t-1} | M_{t-1}, X_{t-1}^*) f(X_{t-1}^* | M_{t-1}, Y_{t-2}). \end{aligned} \quad (5)$$

The first two functions on the right-hand side of the equation can be written in terms of the three components of $f(W', X^* | W, X^*)$ from equation (4):

$$\begin{aligned} f(Y_t | M_t, M_{t-1}, X_{t-1}^*) &= \sum_{x_t^*} \underbrace{f(Y_t | M_t, X_t^*)}_{\text{CCP}} \cdot \underbrace{f(X_t^* | M_t, M_{t-1}, X_{t-1}^*)}_{\text{Law of motion for } X^*} \\ f(M_t, Y_{t-1} | M_{t-1}, X_{t-1}^*) &= \underbrace{f(M_t | Y_{t-1}, M_{t-1}, X_{t-1}^*)}_{\text{Law of motion for } M} \cdot \underbrace{f(Y_{t-1} | M_{t-1}, X_{t-1}^*)}_{\text{CCP}}. \end{aligned} \quad (6)$$

Our identification argument proceeds by first showing how to identify the functions on the left-hand side of Equation (6) from the observed distribution of $f(Y_t, M_t, Y_{t-1} | M_{t-1}, Y_{t-2})$. Subsequently, we show that once these LHS elements are identified, then so are the equilibrium CCP's and laws of motions, and hence also Markov equilibrium law of motion $f(W', X^* | W, X^*)$.

2.2 Identification argument: Details

Our identification argument is related to the recent econometric literature on misclassification models. Hu (2008) shows that, in a general nonlinear setting with misclassification error, three “measurements” of a latent variable are enough to achieve identification. For fixed values of (M_t, M_{t-1}) , we see that (Y_t, Y_{t-1}, Y_{t-2}) enter equation (5) separately in the first, second, and third terms. Hence, we use (Y_t, Y_{t-1}, Y_{t-2}) as three “measurements” of the latent variable X_{t-1}^* .

Given the discreteness assumption (4), all the functions in equation (5) are probability

mass functions (abbreviated pmf afterwards), which can be represented in the form of matrices. In what follows, we use capital letters to denote random variables, while lowercase letters denote a particular realization. For any given (m_t, y_{t-1}, m_{t-1}) in the support of (M_t, Y_{t-1}, M_{t-1}) and $i, j, k, l \in \mathcal{S} \equiv \{1, 2, \dots, J\}$, we define the following J -dimensional square matrices

$$\begin{aligned}
\mathbf{A} &= \left[f(y_t = i, m_t, y_{t-1} | m_{t-1}, y_{t-2} = j) \right]_{i,j} ; & \mathbf{B} &= \left[f(y_t = i | m_t, m_{t-1}, x_{t-1}^* = k) \right]_{i,k} ; \\
\mathbf{C} &= \left[f(x_{t-1}^* = k | m_{t-1}, y_{t-2} = j) \right]_{k,j} ; & \mathbf{D}_1 &= \text{diag} \left\{ \left[f(y_{t-1} | m_t, m_{t-1}, x_{t-1}^* = k) \right]_k \right\} ; \\
\mathbf{D}_2 &= \text{diag} \left\{ \left[f(m_t | m_{t-1}, x_{t-1}^* = k) \right]_k \right\} ; & \mathbf{E} &= \left[f(y_t = i, m_t | m_{t-1}, y_{t-2} = j) \right]_{i,j} ; \\
\mathbf{F} &= \left[f(x_t^* = l | m_t, m_{t-1}, x_{t-1}^* = k) \right]_{l,k} ; & \mathbf{G} &= \left[f(y_t = i | m_t, x_t^* = l) \right]_{i,l} .
\end{aligned} \tag{7}$$

The \mathbf{D}_1 and \mathbf{D}_2 matrices are diagonal matrices. Among the above matrices, only \mathbf{A} and \mathbf{E} are observed, but the rest are unobserved. Clearly, the identification of a matrix, e.g., \mathbf{B} , is equivalent to that of its corresponding pmf, e.g., $f(Y_t | m_t, m_{t-1}, X_{t-1}^*)$.

Given the matrix definitions above, equation (5) can be written as (for fixed (m_t, y_{t-1}, m_{t-1})):

$$\mathbf{A} = \mathbf{B} \cdot \mathbf{D}_1 \cdot \mathbf{D}_2 \cdot \mathbf{C}. \tag{8}$$

Integrating out y_{t-1} in equation 5 yields

$$f(Y_t, M_t | M_{t-1}, Y_{t-2}) = \sum_{X_{t-1}^*} f(Y_t | M_t, M_{t-1}, X_{t-1}^*) f(M_t | M_{t-1}, X_{t-1}^*) f(X_{t-1}^* | M_{t-1}, Y_{t-2}) \tag{9}$$

which, in matrix notation, is (for any given (m_t, m_{t-1})),

$$\mathbf{E} = \mathbf{B} \cdot \mathbf{D}_2 \cdot \mathbf{C}. \tag{10}$$

If the matrix \mathbf{E} is invertible, then we could postmultiply equation (8) by the inverse of equation (10) to get

$$\mathbf{A} \cdot \mathbf{E}^{-1} = \mathbf{B} \cdot \mathbf{D}_1 \cdot \mathbf{B}^{-1} \tag{11}$$

The right-hand side of the above equation is an eigenvalue-eigenvector decomposition of the observed matrix $\mathbf{A} \cdot \mathbf{E}^{-1}$. In this decomposition, the columns of \mathbf{B} (corresponding to

the pmf's $f(Y_t|m_t, m_{t-1}, X_{t-1}^*)$ are the eigenvectors, and the diagonal elements in \mathbf{D}_1 (corresponding to the functions $f(y_{t-1}|m_t, m_{t-1}, X_{t-1}^*)$) are the eigenvalues.

The next two assumptions ensure the existence and uniqueness of this decomposition. Both assumptions are directly testable from the observed data $\{W_t\}$.

Assumption 5 *For any (m_t, y_{t-1}, m_{t-1}) , \mathbf{A} is invertible.*

Assumption 6 *Diag(\mathbf{D}_1) contains J distinctive values. That is, for any (y_{t-1}, m_t, m_{t-1}) , and $\bar{x}_{t-1}^* \neq \tilde{x}_{t-1}^* \in \mathcal{S}$, $f(y_{t-1}|m_t, m_{t-1}, \bar{x}_{t-1}^*) \neq f(y_{t-1}|m_t, m_{t-1}, \tilde{x}_{t-1}^*)$.*

From equation (8), the invertibility of \mathbf{A} (assumption 5) immediately implies the invertibility of \mathbf{B} , \mathbf{C} , \mathbf{D}_1 , and \mathbf{D}_2 . Hence, from equation (10), \mathbf{E} is invertible, and its inverse is given by $\mathbf{C}^{-1} \cdot \mathbf{D}_2^{-1} \cdot \mathbf{B}^{-1}$. Hence, the eigenvalue-eigenvector decomposition (11) is valid.

Assumption 6 ensures that all the J eigenvalues in the decomposition, corresponding to the elements of $Diag(\mathbf{D}_1)$ are distinct. This is testable from the data because it requires the observed matrix $(\mathbf{A} \cdot \mathbf{E}^{-1})$ to have J^2 distinctive eigenvalues. This amounts to testing whether the characteristic polynomial (the determinant of $\lambda I - \mathbf{A} \cdot \mathbf{E}^{-1}$) has J^2 distinctive roots. If all the J^2 eigenvalues are distinct, their corresponding eigenvectors are linearly independent, and the decomposition (11) is unique up to the ordering of the eigenvalues. But like all such decompositions, it is unique only up to a normalization and an ordering of the eigenvectors. Because each eigenvector in \mathbf{B} is a pmf, we should appropriately normalize each column so that it sums to one.

Determining the right ordering of the eigenvectors in \mathbf{B} is important for the identification, because each column corresponds to particular values for the unobserved X_{t-1}^* . In order to pin down the right ordering, additional assumptions must be made. Typically, these assumptions will depend on the type of model under consideration.

Here we make one which is quite general, and should be satisfied by many models. Since $f(Y_t|m_t, m_{t-1}, X_{t-1}^*)$, is identified from the \mathbf{B} matrix, we assume that this marginal pmf is stochastically increasing in X_t^* :

Assumption 7 *$f(Y_t|m_t, m_{t-1}, x_{t-1}^*)$ is stochastically increasing (in the sense of first-order stochastic dominance) in x_{t-1}^* , for fixed (m_t, m_{t-1}) .*

Given this assumption, we can pin down the values of x_{t-1}^* corresponding to each column of the eigenvector matrix \mathbf{B} as follows. For each column j , with elements $B_{i,j}$, we compute the column mean:

$$\mu_j \equiv \sum_{i=1}^J i \times B_{i,j}.$$

Assumption 7 implies that the columns of \mathbf{B} should be ordered such that the column means μ_j are increasing. Hence, without loss of generality, we can set $x_{t-1}^*(j)$, the value of x_{t-1}^* corresponding to column j , to the rank of its column mean, i.e.

$$\forall j : \quad x_{t-1}^*(j) = \text{rank}_{\mu_1, \dots, \mu_J} \mu_j.$$

From the eigenvalue-eigenvector decomposition in Eq. (11), we can identify the \mathbf{B} and \mathbf{D}_1 matrices. Since \mathbf{B} and \mathbf{D}_1 are invertible, the product $\mathbf{H} \equiv \mathbf{D}_2 \cdot \mathbf{C}$ is also identified as $\mathbf{D}_1^{-1} \cdot \mathbf{B}^{-1} \cdot \mathbf{A}$ for any (m_t, y_{t-1}, m_{t-1}) . Because D_2 is a diagonal matrix, the matrix \mathbf{H} corresponds to the product of two probability mass functions, for any (m_t, m_{t-1}) :

$$\mathbf{H} = \left[f(m_t | m_{t-1}, x_{t-1}^* = k) \cdot f(x_{t-1}^* = k | m_{t-1}, y_{t-2} = j) \right]_{k,j}. \quad (12)$$

The following claim (proved in the appendix) shows that identification of \mathbf{H} implies identification of \mathbf{D}_1 and \mathbf{C} :

Claim (*) *Identification of \mathbf{H} implies identification of \mathbf{D}_2 and \mathbf{C} .*

Consequently, we can identify $f(M_t, Y_{t-1} | M_{t-1}, X_{t-1}^*)$ as

$$f(M_t, Y_{t-1} | M_{t-1}, X_{t-1}^*) = f(Y_{t-1} | M_t, M_{t-1}, X_{t-1}^*) f(M_t | M_{t-1}, X_{t-1}^*)$$

where the two functions on the right-hand side correspond to the matrices \mathbf{D}_1 and \mathbf{D}_2 , respectively. At this point, we have identified $f(M_t, Y_{t-1} | M_{t-1}, X_{t-1}^*)$ and $f(Y_t | M_t, M_{t-1}, X_{t-1}^*)$, which are the two functions on the left-hand side of equations (6), thus completing the first part of our identification argument.

Next, using the second equation in (6), we can factor $f(M_t, Y_{t-1} | M_{t-1}, X_{t-1}^*)$ to recover the CCP $f(Y' | M', X^*)$, and the law of motion for M , $f(M' | Y, M, X^*)$. These are two of the three components which constitute the Markov equilibrium law of motion in Eq. (4).

For the third component $f(X^{*'} | M', M, X^*)$, the law of motion for the unobserved state variables X^* , we use the first equation in (6). This equation can be written in matrix notation as

$$\mathbf{B} = \mathbf{G} \cdot \mathbf{F}$$

for a given (m_t, m_{t-1}) because assumption 3 implies $f_{Y' | M', X^{*'}} = f_{Y_t | M_t, X_t^*}$. Because \mathbf{B} is invertible (see the earlier discussion following Assumption 5), then so are \mathbf{F} and \mathbf{G} .

Hence, the law of motion for X^* , corresponding to the matrix \mathbf{F} , can be recovered as:

$$\mathbf{F} = \mathbf{G}^{-1} \cdot \mathbf{B}. \quad (13)$$

Hence, our identification argument is complete:

Theorem 8 *Under the assumptions 1, 2, 3, 4, 5, 6, and 7, the density $f(W_t, W_{t-1}, W_{t-2})$, for any $t \in \{3, \dots, T\}$, uniquely determines the Markov equilibrium law of motion $f(W', X^* | W, X^*)$.*

3 A Monte Carlo simulation example

Based on the nonparametric identification results in the previous section, we present some simulation results which utilize the constructive identification proof for nonparametric estimation of a binary-choice dynamic optimization model.

3.1 Details of test model

We consider a stationary binary choice dynamic optimization model. This model consists of three variables: Y, M, X^* . Each variable is binary, and takes values in $\{0, 1\}$. Because the model is stationary, we use primes $'$ to denote next-period values

Following the restriction in Assumption 2(ii), we parametrized the law of motion for X^* as:

$$Pr(X^{*'} = 1 | X^*, M, M') = \frac{\exp(\phi(X^*, M, M'))}{1 + \exp(\phi(X^*, M, M'))}$$

where

$$\phi(X^*, M, M') = \alpha_1 \cdot I(X^* = 1) + \alpha_2 \cdot I(M = 1) + \alpha_3 \cdot I(M' = 1) + \alpha_4 \cdot I(M = 1) \cdot I(X^* = 1)$$

Similarly, the law of motion for the observed state variable M is parametrized as:

$$Pr(M' = 1 | M, Y, X^*) = \frac{\exp(\psi(X^*, M, Y))}{1 + \exp(\psi(X^*, M, Y))}$$

with

$$\psi(X^*, M, Y) = \gamma_1 \cdot I(X^* = 1) + \gamma_2 \cdot I(M = 1) + \gamma_3 \cdot I(Y = 1)$$

For what follows, it is useful to define $S \equiv (M, X^*)$ with

$$\begin{aligned} S = 1 &\Leftrightarrow M = 0, X^* = 0 \\ S = 2 &\Leftrightarrow M = 0, X^* = 1 \\ S = 3 &\Leftrightarrow M = 1, X^* = 0 \\ S = 4 &\Leftrightarrow M = 1, X^* = 1 \end{aligned}$$

Also, note that

$$Pr(S'|S, Y) = Pr(X^*|M, M', X^*) \cdot Pr(M'|M, Y, X^*)$$

The per-period utility functions are given by

$$u(Y, M, X^*, \epsilon_Y) = \underbrace{\lambda_1 \cdot I(Y = 1) + \lambda_2 \cdot I(M = 1) + \lambda_3 \cdot (I(X^* = 1) - 0.5) \cdot I(Y = 1)}_{\equiv \bar{u}(Y, S)} + \epsilon_Y$$

where ϵ_0 , and ϵ_1 are i.i.d. extreme value distributed with unit variance.

Using Eq. (15) from Aguirregabiria and Mira (2007), the optimal choice probabilities are implicit defined by the following functional equation

$$\begin{aligned} Pr(Y|S) &= \frac{\exp[\bar{u}(Y, S) + \beta \sum_{S'} V(S'; P) \cdot Pr(S'|S, Y)]}{\sum_{\tilde{Y} \in (0,1)} \exp[\bar{u}(\tilde{Y}, S) + \beta \sum_{S'} V(S'; P) \cdot Pr(S'|S, \tilde{Y})]} \\ &\equiv \Psi(\Theta; P) \end{aligned} \quad (14)$$

where P denotes the $|Y| \times |S|$ vector of choice probabilities $Pr(Y|S)$.

In the above, $V(S; P)$ denotes the value function, given some set of choice probabilities P . Using Eq. (14) in Aguirregabiria and Mira (2007), $V(S; P)$ is the corresponding element of the $|S|$ -dimensional vector defined as:

$$(I - \beta F)^{-1} \left\{ \sum_{\tilde{Y} \in (0,1)} P(\tilde{Y}) * [\bar{u}(\tilde{Y}) + \epsilon(\tilde{Y})] \right\} \quad (15)$$

where $*$ denotes elementwise multiplication, and F is the $|S|$ -dimensional square matrix with (i, j) -element equal to

$$Pr(S' = j|S = i) \equiv \sum_{Y=(0,1)} P(Y|S = i) \cdot Pr(S' = j|S = i, Y). \quad (16)$$

Moreover, $P(\tilde{Y})$ is the $|S|$ -vector consisting of elements $Pr(\tilde{Y}|S)$, $\bar{u}(\tilde{Y})$ is the $|S|$ -vector

of per-period utilities $\bar{u}(\tilde{Y}; S)$, and $\epsilon(\tilde{Y})$ is an $|S|$ -vector where each element is $E[\epsilon_{\tilde{Y}}|\tilde{Y}, S]$. For the logit assumptions, the closed-form is

$$E[\epsilon_{\tilde{Y}}|\tilde{Y}, S] = \text{Euler's constant}(0.57721) - \log(P(\tilde{Y}|S)).$$

Given the derivations above, it is possible to simulate, for fixed values of the parameters $\Theta \equiv \{\alpha_1, \alpha_2, \alpha_3, \gamma_1, \gamma_2, \gamma_3, \lambda_1, \lambda_2\}$ as well as the discount rate β , sequences of the variables $\{Y_t, X_t^*, M_t\}_i$ for agents $i = 1, \dots, N$ and $t = 1, 2, 3$. In the simulations reported here, we consider the number of agents $N \in \{800, 1000, 5000\}$.

3.2 Results

After simulating a dataset in the manner above, we mimic our identification argument from the previous section to recover the structural components from equation (2), using only the variables (Y, M) . That is, the matrices \mathbf{A} and \mathbf{E} (as defined previously), which contain the pmf's for $f(y_t, m_t, y_{t-1}|m_{t-1}, y_{t-2})$ and $f(y_t = i, m_t|m_{t-1}, y_{t-2} = j)$, respectively, are estimated directly from the data, and the matrix manipulations in Equations (8)-(13) are performed on these matrices to obtain estimates of the CCP's $f(Y|M, X^*)$ and the Markov laws of motions for the state variables (M, X^*) .

Table 1 contains the simulation results for the nonparametric estimates of the conditional choice probabilities of the model. Tables 2 and 3 contain the estimates for the Markov laws of motion for, respectively, the observed state variable M and the unobserved state variable X^* .

Table 1: Simulation results for CCP's $Pr(Y = 0|M, X^*)$

Nonparametric results from Monte Carlo simulation experiments.
Each experiment replicated 100 times.

				N = 800^a		N = 3000		N = 5000	
<i>M</i>	<i>X*</i>	True	Avg. ^b	Stdev. ^c	Avg.	Stdev.	Avg.	Stdev.	
0	0	0.9102	0.8611	0.1363	0.8958	0.0817	0.9046	0.0707	
0	1	0.0661	0.0798	0.0391	0.0694	0.0220	0.0676	0.0204	
1	0	0.9064	0.8524	0.1108	0.8948	0.0618	0.9024	0.0466	
1	1	0.0654	0.0692	0.0380	0.0641	0.0237	0.0656	0.0178	

^a N denotes number of agents in simulated dataset

^bAveraged across all replications

^cStandard deviation across all replications.

Overall, the results are quite encouraging. In Table 1, which contains the results for the conditional choice probabilities, we see that even when $N = 800$, the average

estimates are close to the truth, with the exception of the $(M = 0, X^* = 1)$ case. For this case, however, the results improve as the number of agents N in the simulated datasets increases.

The results are qualitatively similar for the results in Table 2, which are for $Pr(M' = 0|M, Y, X^*)$, the law of motion for the observed state variable M along the dynamically optimal path. For these results, the two cases with the less accurate results are $(M = 1, Y = 0, X^* = 0)$ and $(M = 1, Y = 1, X^* = 0)$.

Table 2: Simulation results for Law of motion $Pr(M' = 0|M, Y, X^*)$

Nonparametric results from Monte Carlo simulation experiments.

Each experiment replicated 100 times.

			N = 800^a			N = 3000		N = 5000	
<i>M</i>	<i>Y</i>	<i>X*</i>	True	Avg. ^b	Stdev. ^c	Avg.	Stdev.	Avg.	Stdev.
0	0	0	0.5000	0.5143	0.1089	0.4977	0.0510	0.4953	0.0338
0	0	1	0.7109	0.6904	0.2442	0.7310	0.1564	0.7305	0.1109
0	1	0	0.4134	0.4016	0.0869	0.4194	0.0427	0.4112	0.0368
0	1	1	0.6341	0.6090	0.3097	0.6525	0.1973	0.6479	0.1325
1	0	0	0.6457	0.5641	0.4302	0.5138	0.3582	0.5860	0.3487
1	0	1	0.8176	0.8206	0.0284	0.8195	0.0138	0.8174	0.0100
1	1	0	0.5622	0.6314	0.3636	0.5299	0.3330	0.5517	0.3066
1	1	1	0.7595	0.7457	0.1370	0.7603	0.0190	0.7573	0.0162

^a N denotes number of agents in simulated dataset

^bAveraged across all replications

^cStandard deviation across all replications.

The estimation results for $Pr(X^{*'} = 0|M, M', X^*)$, the law of motion for the unobserved state variable X^* , reported in Table 3, are more erratic. This may not be surprising, as these law of motions are affected the most by X^* , the unobserved state variable. But even here, for about half of the eight cases, the estimation performs remarkably well, even with the modest dataset size ($N = 800$). However, in the cases $(M = 0, M' = 1, X^* = 1)$, $(M = 1, M' = 0, X^* = 0)$, $(M = 1, M' = 0, X^* = 1)$, and $(M = 1, M' = 1, X^* = 1)$, the estimation performs less accurately. However, even in these cases, the magnitudes of the estimates are correct, but the actual estimates are not too close to the truth on average.

Table 3: Simulation results for Law of motion $Pr(X^{*t} = 0|M, M', X^*)$
 Nonparametric results from Monte Carlo simulation experiments.
 Each experiment replicated 100 times.

			N = 800 ^a		N = 3000		N = 5000		
M	M'	X^*	True	Avg. ^b	Stdev. ^c	Avg.	Stdev.	Avg.	Stdev.
0	0	0	0.5000	0.5898	0.1815	0.5159	0.0994	0.5073	0.0627
0	0	1	0.1192	0.1102	0.0500	0.1180	0.0318	0.1187	0.0242
0	1	0	0.7109	0.7732	0.1941	0.7198	0.1242	0.7174	0.1002
0	1	1	0.0431	0.0420	0.0524	0.0425	0.0348	0.0414	0.0313
1	0	0	0.4013	0.4674	0.1626	0.4215	0.0948	0.4184	0.0655
1	0	1	0.0832	0.1330	0.1387	0.0872	0.0504	0.0796	0.0310
1	1	0	0.6225	0.6668	0.1505	0.6324	0.1870	0.6429	0.0675
1	1	1	0.0293	0.1124	0.1444	0.0407	0.0471	0.0352	0.0310

^a N denotes number of agents in simulated dataset

^bAveraged across all replications

^cStandard deviation across all replications.

4 Empirical illustration: advertising and doctors' prescriptions of pharmaceutical drugs

Next, we consider a small scale empirical illustration, where we estimate a simple model relating doctors' prescriptions of the cholesterol-lowering drug *Crestor* to drug advertising or *detailing*, which are the visits of the pharmaceutical representatives ("pharmreps" for short) to the doctor. In this model, agents i are doctors, and t denotes weeks. The variables are:

- Y_{it} : share of *Crestor* prescribed by doctor i in week t to his patients
- M_{it} (the observed state variable): a binary variable indicating whether doctor i was visited by a *Crestor* pharmrep in the two weeks prior to week t .
- X_{it}^* (unobserved state variable): doctor i 's unobserved preference shock related to *Crestor*. This could be due to word of mouth, or the knowledge that doctor i has from prescribing *Crestor* to his patients.

Based on the framework in this paper, the structural components we want to recover are: (i) choice probability $f(Y_t|M_t, X_t^*)$; (ii) law of motion for advertising $f(M_t|M_{t-1}, X_{t-1}^*, Y_{t-1})$; and (iii) law of motion for unobserved state variable $f(X_t^*|M_t, M_{t-1}, X_{t-1}^*)$.

Market background *Crestor* (active ingredient resuvastatin) is one of the so-called "statin" cholesterol-lowering drugs which, as a group, constitutes the largest drug market

worldwide, in terms of both sales and prescriptions. Because *Crestor* was the most powerful statin drug when it entered the market in the Fall of 2003, some extra precautions were taken regarding it. In particular, the FDA label for Crestor, which appeared in its first form on August 12, 2003, contained a warning regarding patients of Asian descent, who appeared in studies to retain much higher levels of drug concentration in their blood, relative to Caucasian users:⁵

Pharmacokinetic studies show an approximate 2-fold elevation in median exposure in Japanese subjects residing in Japan and in Chinese subjects residing in Singapore when compared with Caucasians residing in North America and Europe. No studies directly examining Asian ethnic population residing in the U.S. are available, so the contribution of environmental and genetic factors to the observed increase in *rosuvastatin* drug levels have not been determined.

In the March 2, 2005 version of the FDA, this precaution was strengthened, on the basis of studies on US subjects:

Pharmacokinetic studies, including one conducted in the US, have demonstrated an approximate 2-fold elevation in median exposure in Asian subjects when compared with Caucasian control group.

Exploiting this contraindication, we focus here on estimating the differential effects of advertising on doctors' prescriptions of *Crestor* to Asians vs. non-Asian patients. This provides a unique test for the informational content of advertising: if advertising is informative, it should have a more negative effect on prescription probabilities to Asian patients, relative to non-Asian patients. Here, we will test whether this is true, within the modeling framework of this paper, which allows advertising to be "endogenous", in the sense that it is influenced by lagged values of the doctor-specific shock X_{t-1}^* as well as the doctor's previous choice variable Y_{t-1} .⁶

Data description The data used for the empirical analysis is from a category of prescription drug used to treat high levels of blood cholesterol; such a drug is commonly referred to as a Statin drug. The data includes a panel of representative physicians from the United States. These data are obtained from a pharmaceutical consulting firm *ImpactRX*.⁷ The dataset is unique in that, for each physician, we observe a sample of prescriptions between January 1st 2004 and December 31st 2004. In addition, we also have a record of all the detailing visits made by pharmaceutical sales representatives

⁵See Food and Drug Administration (2003, 2005).

⁶For more information on this market, and also additional empirical evidence, see Shum and Tan (2007).

⁷See Manchanda and Narayanan (2009) for another paper which utilizes a similar dataset.

during the same period. We construct our data by combining the prescription data and detailing data.

For the purpose of the empirical application here, we aggregated all data to the (doctor-week) level, as described above. In addition, while the identification argument above was “cross-sectional” in nature, being based upon observations of three observations of $\{Y_t, M_t\}$ per individual, in the estimation we exploited the long time series data we have for each subject, and pooled every “three time-contiguous observations” $\{Y_{i,\tau}, M_{i,\tau}\}_{\tau=t-2}^{\tau=t}$ across all doctors i , and all weeks $\tau = 3, \dots, T$. Formally, this is justified under the assumption that the process $\{Y_t, M_t\}$ is stationary and ergodic for each subject and each round. Under these assumptions, the ergodic theorem ensures that the (across time and subjects) sample frequencies used to construct the matrices \mathbf{A} and \mathbf{E} converge towards population counterparts.

Asians make up a small percentage of the patients in our sample. Only 260 doctors have a three-week sequence where Asian patients were seen in each week. Hence, in order to estimate the model, we considered different definitions of the binary advertising variable M_t , so that the resulting matrix \mathbf{A} would be invertible, for all eight combinations of (M_t, Y_t, M_{t-1}) (in order to satisfy Assumption 5). Moreover, because weeks in the dataset are defined according to calendar weeks (ie. Monday to Sunday), but pharmreps can visit on any weekday, not all of a doctor’s patients in a given (calendar) week may have been affected by the detailing that occurred that week. Accordingly, after some trial and error, we settled on the following definitions of the binary variables Y_t and M_t :

- $Y_{it} \in \{0, 1\}$: $Y_t = 1$, if for doctor i in week t , there are observations of *Crestor* prescribed. Otherwise, $Y_{it} = 0$.
- $M_{it} \in \{0, 1\}$: $M_t = 1$, if more than 35% of the patients who visited doctor i in week t visited within two weeks of the most recent *Crestor* pharmrep visit. Otherwise, $M_{it} = 0$.

Table 4 contains summary statistics from the raw data. We see that, moving from $M = 0$ to $M = 1$, the probability of prescribing *Crestor* to non-Asian patients increases by 10% (from 24.5% to 34.9%, which represents roughly a 40% increase. For Asian patients, this probability increases by 9% from 10.9% to 20%, which is almost a doubling of the prescription probability for these patients. Hence, in the raw data, advertising does not appear informative. The right-hand side of this table presents the transition probabilities for advertising.

Estimation results Table 5 contains the estimates of the *Crestor* prescription probabilities, conditional on both M and X^* . These were estimated by pooling together all

Table 4: Summary statistics from the data

$M =$	$Y =$	$P(Y = 1 M)$		$P(M' = 1 Y, M)$	
		Asians	non-Asians	Asians	non-Asians
0		0.1086	0.2452		
0	0			0.1731	0.1816
0	1			0.5147	0.5780
1		0.2000	0.3489		
1	0			0.3158	0.2396
1	1			0.5882	0.6505
# obs ^a		260	22928		
$P(Y = 1)$		0.1308	0.2811		
$P(M = 1)$		0.3000	0.3323		

^aEach observation is a (doctor-week).

Table 5: Estimates for CCP $P(Y = 1|X^*, M)$

$X^* =$	$M =$	Asians	non-Asians
0	0	0.0913 (0.1104)	0.1114 (0.0058)
0	1	0.1420 (0.0983)	0.1474 (0.0109)
1	0	0.6806 (0.2302)	0.6318 (0.0174)
1	1	0.9208 (0.2496)	0.6568 (0.0171)

doctors, and all three-week sequences in which at least one patient of a specified ethnicity was seen each week. The estimates show that, for both $X^* = 0$ and $X^* = 1$, advertising raises the probability of prescribing *Crestor* for both Asians and non-Asians. However, the magnitudes of these changes are quite different than in the raw data, and show that the “causal effects” of advertising (once X^* is controlled for) are quite distinct than the raw values presented in Table 4. For instance, when $X^* = 0$ (which we interpret to be the case when doctors are pessimistic about *Crestor*’s effectiveness), we see that advertising raises the prescription probability by only 3% for non-Asians, but by 5% for Asians. Thus, these results echo the trend in the raw data, that advertising raises the prescription probabilities disproportionately more for Asians than for non-Asians.

Table 6: Estimates for law of motion $P(M' = 1|M, Y, X^*)$

$X^* =$	$M =$	$Y =$	Asians	non-Asians
0	0	0	0.1727 (0.0884)	0.1754 (0.0052)
0	0	1	0.2481 (0.2272)	0.2022 (0.0220)
0	1	0	0.5062 (0.1388)	0.5669 (0.0107)
0	1	1	0.1629 (0.2586)	0.6527 (0.0324)
1	0	0	0.9168 (0.3143)	0.2538 (0.0204)
1	0	1	0.5037 (0.3135)	0.2676 (0.0119)
1	1	0	0.4198 (0.2935)	0.6319 (0.0276)
1	1	1	0.0000 (0.3935)	0.6532 (0.0141)

Tables 6 and 7 present, respectively, the estimates for the law of motion for advertising, M , and for the unobserved state variable X^* . The estimates for the law of motion for M show that X^* has an effect on next period’s advertising. This indicates that advertising is endogenous in the sense that it is related to serially correlated shocks X^* which also affect doctors’ prescription behavior. At the same time, the estimates for the law of motion of X^* show that current and past values of M also feedback to the realization of the shock X^* .⁸

⁸In both tables 6 and 7, we set some transition probabilities to 0 (resp, 1) when the estimated probabilities became < 0 (resp, > 1).

Table 7: Estimates for law of motion $P(X^{*'} = 1|M', M, X^*)$

$M' =$	$M =$	$X^* =$	Asians	non-Asians
0	0	0	0.0322 (0.1868)	0.0010 (0.0072)
0	0	1	0.9933 (0.4641)	0.9546 (0.0357)
0	1	0	0.0115 (0.1954)	0.0140 (0.0215)
0	1	1	0.9933 (0.2047)	0.9580 (0.0498)
1	0	0	0.0804 (0.2360)	0.0095 (0.0213)
1	0	1	0.7223 (0.2610)	1.0000 (0.0124)
1	1	0	0.0067 (0.2236)	0.0097 (0.0200)
1	1	1	0.3849 (0.3191)	1.0000 (0.0090)

Therefore, the results from this small empirical application show that advertising appears to cause larger increases in the *Crestor* prescription probability to Asian patients relative to non-Asian patients. Because Asian patients were contraindicated for this drug, our findings appear to refute the hypothesis that advertising is informative. While this finding is striking, we reiterate the caveat that there are not many Asian patients in the dataset and, as we remarked above, we chose the particular definition of M in order for the estimation to proceed. For other definitions of M , the \mathbf{A} matrices were not invertible for all combinations of (y_t, m_t, m_{t-1}) , and we were not able to obtain results. Hence, we view this exercise as more an illustration of our identification results rather than a full-blown empirical application.

5 Conclusion

In this paper, we present a methodology for the estimation of dynamic models, in the case when all the variables of the model are discrete. Monte Carlo simulations showed that the estimator performs quite well in practice, and a short empirical application was provided for estimating the effect of advertising on pharmaceutical prescription probabilities, while allowing advertising to be affected by serially correlated preference shocks which also affect doctors' prescription behavior.

In ongoing work, we are also considering the extension of the methods presented here

to the case of multi-agent dynamic games, in which there are agent-specific unobserved state variables which are serially correlated.

References

- ABBRING, J., AND J. HECKMAN (2007): “Econometric Evaluation of Social Programs, Part III: Distributional Treatment Effects, Dynamic Treatment Effects, Dynamic Discrete Choice, and General Equilibrium Policy Evaluation,” in *Handbook of Econometrics, Vol. 6B*, ed. by J. Heckman, and E. Leamer, chap. 72. North-Holland.
- 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.
- (2007): “Sequential Estimation of Dynamic Discrete Games,” *Econometrica*, 75, 1–53.
- ARCIDIACONO, P., AND R. MILLER (2006): “CCP Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity,” Manuscript, Duke University.
- BAJARI, P., L. BENKARD, AND J. LEVIN (2007): “Estimating Dynamic Models of Imperfect Competition,” *Econometrica*, 75, 1331–1370.
- BAJARI, P., V. CHERNOZHUKOV, H. HONG, AND D. NEKIPELOV (2007): “Nonparametric and Semiparametric Analysis of a Dynamic Game Model,” Manuscript, University of Minnesota.
- BLEVINS, J. (2008): “Sequential MC Methods for Estimating Dynamic Microeconomic Models,” Duke University, working paper.
- CUNHA, F., J. HECKMAN, AND S. SCHENNACH (2006): “Estimating the Technology of Cognitive and Noncognitive Skill Formation,” manuscript, University of Chicago.
- FOOD AND DRUG ADMINISTRATION (2003, 2005): “Labels for *Crestor*,” Available at http://www.fda.gov/cdei/foi/label/2003/21366_crestor_lbl.pdf.
- GALLANT, R., H. HONG, AND A. KHWAJA (2009): “Estimating a Dynamic Oligopolistic Game with Serially Correlated Unobserved Production Costs,” manuscript, Duke University.
- GHAHRAMANI, Z. (2001): “An Introduction to Hidden Markov Models and Bayesian Networks,” *International Journal of Pattern Recognition and Artificial Intelligence*, 15, 9–42.
- HECKMAN, J., AND S. NAVARRO (2007): “Dynamic discrete choice and dynamic treatment effects,” *Journal of Econometrics*, 136, 341–396.
- HONG, H., AND M. SHUM (2009): “Pairwise-Difference Estimation of a Dynamic Optimization Model,” *Review of Economic Studies*, 77, 273–304.
- HOTZ, J., AND R. MILLER (1993): “Conditional Choice Probabilities and the Estimation of Dynamic Models,” *Review of Economic Studies*, 60, 497–529.
- HU, Y. (2008): “Identification and Estimation of Nonlinear Models with Misclassification Error Using Instrumental Variables: a General Solution,” *Journal of Econometrics*, 144, 27–61.
- HU, Y., AND M. SHUM (2008): “Nonparametric Identification of Dynamic Models with Unobserved State Variables,” Johns Hopkins University, Dept. of Economics working paper #543.
- IMAI, S., N. JAIN, AND A. CHING (2009): “Bayesian Estimation of Dynamic Discrete Choice Models,” *Econometrica*, 77, 1865–1899.

- KASAHARA, H., AND K. SHIMOTSU (2009): “Nonparametric Identification of Finite Mixture Models of Dynamic Discrete Choice,” *Econometrica*, 77, 135–175.
- KEANE, M., AND K. WOLPIN (1994): “The Solution and Estimation of Discrete Choice Dynamic Programming Models by Simulation and Interpolation: Monte Carlo Evidence,” *Review of Economics and Statistics*, 76, 648–672.
- MAGNAC, T., AND D. THESMAR (2002): “Identifying Dynamic Discrete Decision Processes,” *Econometrica*, 70, 801–816.
- MANCHANDA, P., AND S. NARAYANAN (2009): “Heterogeneous Learning and the Targeting of Marketing Communication for New Products,” *Marketing Science*, 28, 424–441.
- MILLER, R. (1984): “Job Matching and Occupational Choice,” *Journal of Political Economy*, 92, 1086–1120.
- NORETS, A. (2009): “Inference in dynamic discrete choice models with serially correlated unobserved state variables,” *Econometrica*, 77, 1665–1682.
- PAKES, A. (1986): “Patents as Options: Some Estimates of the Value of Holding European Patent Stocks,” *Econometrica*, 54(4), 755–84.
- PAKES, A., M. OSTROVSKY, AND S. BERRY (2007): “Simple Estimators for the Parameters of Discrete Dynamic Games (with Entry Exit Examples),” *RAND Journal of Economics*, 38, 373–399.
- PESENDORFER, M., AND P. SCHMIDT-DENGLER (2008): “Asymptotic Least Squares Estimators for Dynamic Games,” *Review of Economic Studies*, 75, 901–928.
- RUST, J. (1987): “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55, 999–1033.
- (1994): “Structural Estimation of Markov Decision Processes,” in *Handbook of Econometrics*, Vol. 4, ed. by R. Engle, and D. McFadden, pp. 3082–146. North Holland.
- SHUM, M., AND W. TAN (2007): “Is Advertising Informative? Evidence from Contraindicated Drug Prescriptions,” work in progress.
- SIEBERT, R., AND C. ZULEHNER (2008): “The Impact of Market Demand and Innovation on Market Structure,” Purdue University, working paper.

1 Derivation of auxiliary results

1.1 Derivation of Equation (5)

Consider the observed density $f(W_t, W_{t-1}, W_{t-2})$. Assumptions 1 and 2(i) imply

$$\begin{aligned}
& f(W_t, W_{t-1}, W_{t-2}) \\
&= \sum_{X_t^*, X_{t-1}^*} f(W_t, X_t^* | W_{t-1}, W_{t-2}, X_{t-1}^*) f(W_{t-1}, W_{t-2}, X_{t-1}^*) \\
&= \sum_{X_t^*, X_{t-1}^*} f(Y_t | M_t, X_t^*) f(X_t^* | M_t, Y_{t-1}, M_{t-1}, X_{t-1}^*) f(M_t | Y_{t-1}, M_{t-1}, X_{t-1}^*) f(Y_{t-1} | M_{t-1}, X_{t-1}^*) f(X_{t-1}^*, M_{t-1}, Y_{t-2}, M_{t-2}) \\
&= \sum_{X_t^*, X_{t-1}^*} f(Y_t | M_t, X_t^*) f(X_t^* | M_t, Y_{t-1}, M_{t-1}, X_{t-1}^*) f(M_t, Y_{t-1} | M_{t-1}, X_{t-1}^*) f(X_{t-1}^*, M_{t-1}, Y_{t-2}, M_{t-2}).
\end{aligned}$$

After integrating out M_{t-2} , assumption 2(ii) then implies

$$\begin{aligned}
& f(Y_t, M_t, Y_{t-1}, M_{t-1}, Y_{t-2}) \\
&= \sum_{X_{t-1}^*} \left(\sum_{X_t^*} f(Y_t | M_t, X_t^*) f(X_t^* | M_t, M_{t-1}, X_{t-1}^*) \right) f(M_t, Y_{t-1} | M_{t-1}, X_{t-1}^*) f(X_{t-1}^*, M_{t-1}, Y_{t-2})
\end{aligned}$$

The expression in the parenthesis can be simplified as $f(Y_t | M_t, M_{t-1}, X_{t-1}^*)$. We then have

$$\begin{aligned}
& f_{Y_t, M_t, Y_{t-1} | M_{t-1}, Y_{t-2}} \\
&= \sum_{X_{t-1}^*} f(Y_t | M_t, M_{t-1}, X_{t-1}^*) f(M_t, Y_{t-1} | M_{t-1}, X_{t-1}^*) f(X_{t-1}^*, M_{t-1}, Y_{t-2})
\end{aligned} \tag{17}$$

as claimed in Equation (5). ■

1.2 Proof of Claim (*)

Define

$$h(j, k; m_t, m_{t-1}) \equiv f(m_t | m_{t-1}, x_{t-1}^* = k) \cdot f(x_{t-1}^* = k | m_{t-1}, y_{t-2} = j)$$

Identification of \mathbf{H} is equivalent to identification of the $h(\dots)$ function.

By integrating $h(k, j; m_t, m_{t-1})$ over m_t , we can identify the $f(x_{t-1}^* = k | m_{t-1}, y_{t-2} = j)$ function:

$$\begin{aligned}
\int h(k, j; m_t, m_{t-1}) dm_t &= \int f(m_t | m_{t-1}, x_{t-1}^* = k) \cdot f(x_{t-1}^* = k | m_{t-1}, y_{t-2} = j) dm_t \\
&= f(x_{t-1}^* = k | m_{t-1}, y_{t-2} = j) \left[\int f(m_t | m_{t-1}, x_{t-1}^* = k) dm_t \right] \\
&= f(x_{t-1}^* = k | m_{t-1}, y_{t-2} = j)
\end{aligned}$$

because $f(m_t | m_{t-1}, x_{t-1}^*)$ is a probability density function. Consequently, $f(m_t | m_{t-1}, x_{t-1}^*)$ is

also identified as

$$f(m_t|m_{t-1}, x_{t-1}^*) = \frac{h(x_{t-1}^*, y_{t-2}; m_t, m_{t-1})}{f(x_{t-1}^*|m_{t-1}, y_{t-2})}.$$

Hence, from knowledge of \mathbf{H} , we are able to identify the function corresponding to the matrices \mathbf{D}_2 and \mathbf{C} also. ■