

DISCUSSION PAPER SERIES

DP13240

LINEAR IV REGRESSION ESTIMATORS FOR STRUCTURAL DYNAMIC DISCRETE CHOICE MODELS

Myrto Kalouptsi, Paul T. Scott and Eduardo Souza-
Rodrigues

**INDUSTRIAL ORGANIZATION, LABOUR
ECONOMICS AND PUBLIC ECONOMICS**



LINEAR IV REGRESSION ESTIMATORS FOR STRUCTURAL DYNAMIC DISCRETE CHOICE MODELS

Myrto Kalouptsi, Paul T. Scott and Eduardo Souza-Rodrigues

Discussion Paper DP13240
Published 11 October 2018
Submitted 08 October 2018

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **INDUSTRIAL ORGANIZATION, LABOUR ECONOMICS AND PUBLIC ECONOMICS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Myrto Kalouptsi, Paul T. Scott and Eduardo Souza-Rodrigues

LINEAR IV REGRESSION ESTIMATORS FOR STRUCTURAL DYNAMIC DISCRETE CHOICE MODELS

Abstract

In structural dynamic discrete choice models, the presence of serially correlated unobserved states and state variables that are measured with error may lead to biased parameter estimates and misleading inference. In this paper, we show that instrumental variables can address these issues, as long as measurement problems involve state variables that evolve exogenously from the perspective of individual agents (i.e., market-level states). We define a class of linear instrumental variables estimators that rely on Euler equations expressed in terms of conditional choice probabilities (ECCP estimators). These estimators do not require observing or modeling the agent's entire information set, nor solving or simulating a dynamic program. As such, they are simple to implement and computationally light. We provide constructive identification arguments to identify the model primitives, and establish the consistency and asymptotic normality of the estimator. A Monte Carlo study demonstrates the good finite-sample performance of the ECCP estimator in the context of a dynamic demand model for durable goods.

JEL Classification: N/A

Keywords: N/A

Myrto Kalouptsi - myrto@fas.harvard.edu
Harvard University

Paul T. Scott - ptscott@gmail.com
NYU

Eduardo Souza-Rodrigues - e.souzarodrigues@utoronto.ca
University of Toronto

Linear IV Regression Estimators for Structural Dynamic Discrete Choice Models *

Myrto Kalouptsi[†], Paul T. Scott[‡], Eduardo Souza-Rodrigues[§]

Harvard University, CEPR and NBER, NYU Stern, University of Toronto

September 2018

Abstract

In structural dynamic discrete choice models, the presence of serially correlated unobserved states and state variables that are measured with error may lead to biased parameter estimates and misleading inference. In this paper, we show that instrumental variables can address these issues, as long as measurement problems involve state variables that evolve exogenously from the perspective of individual agents (i.e., market-level states). We define a class of linear instrumental variables estimators that rely on Euler equations expressed in terms of conditional choice probabilities (ECCP estimators). These estimators do not require observing or modeling the agent’s entire information set, nor solving or simulating a dynamic program. As such, they are simple to implement and computationally light. We provide constructive identification arguments to identify the model primitives, and establish the consistency and asymptotic normality of the estimator. A Monte Carlo study demonstrates the good finite-sample performance of the ECCP estimator in the context of a dynamic demand model for durable goods.

Keywords: dynamic discrete choice, unobserved states, instrumental variables, identification, Euler equations.

JEL: C13; C35; C36; C51; C61

*We are grateful to Victor Aguirregabiria, Chris Conlon, Mike Dickstein, Bo Honoré, Rob McMillan, Yuanyuan Wan, and Stan Zin for many helpful comments. Jan Dee provided outstanding research assistance. Much of the material of this paper was formerly part of our paper “Identification of Counterfactuals in Dynamic Discrete Choice Models”, NBER Working Paper 21527.

[†]Harvard University, Department of Economics, Littauer Center 124, Cambridge, MA 02138. email: myrto@fas.harvard.edu

[‡]Stern School of Business, New York University, Kaufman Management Center, 44 W. 4th St., New York, NY 10012. email: pscott@stern.nyu.edu.

[§]Department of Economics, University of Toronto, Max Gluskin House, 150 St. George St., Toronto, Ontario M5S 3G7, Canada. email: e.souzarodrigues@utoronto.ca

1 Introduction

Instrumental variables methods are widely used to address omitted variables and measurement error problems in reduced-form models. In this paper, we show that instrumental variables can play the same role in structural dynamic discrete choice (DDC) models, as long as the measurement problems involve market-level state variables (also known as forcing variables), i.e., variables that evolve exogenously from the perspective of individual agents. To that end, we define a class of linear IV regression estimators for structural dynamic discrete choice models, which we call Euler Equations in Conditional Choice Probabilities (ECCP) estimators. ECCP estimators are a discrete-choice analog to the Euler equation approach for models with continuous choice variables developed by Hall (1978), Hansen and Sargent (1980), Hansen and Sargent (1982), and Hansen and Singleton (1982).¹

Structural DDC models have proven useful in a variety of important fields – for example, in the study of labor markets, firm dynamics, consumer demand, and environmental problems. Standard methods for estimating DDC models require the computation of continuation value functions, by either solving the full dynamic problem (Rust, 1987) or measuring the continuation values by forward-solving (Hotz and Miller, 1993; Pesendorfer and Schmidt-Dengler, 2008) or forward-simulating (Hotz et al., 1994; Bajari et al., 2007). In this context, Arcidiacono and Miller (2011) and Aguirregabiria and Magesan (2013) represent two major departures from prior work, providing strategies for estimating model parameters without calculating continuation values, thereby reducing the computational burden of estimating DDC models substantially. Arcidiacono and Miller (2011) define the *finite dependence* property and show how it allows the econometrician to construct conditions from an agent’s optimization problem in which the continuation values cancel out.² Aguirregabiria and Magesan (2013) propose a representation of the discrete choice problem as a continuous decision problem in which the decision variables are choice probabilities. Based on this representation, they derive first-order conditions for optimization that are expressed in terms of choice probabilities, and that are similar to Euler equations for continuous decision problems.

We build on these methodological contributions to allow for measurement problems in market-level state variables, including serially correlated unobserved states, endogeneity problems, and measurement error. It is widely acknowledged that the presence of such unobservables is an important concern in empirical applications, and that ignoring them may lead to biased estimates and misleading inference. Empirical work exploring finite dependence in structural DDC models that allows for such measurement problems has appeared in the recent applied literature (Scott,

¹Euler equation estimators rely on agents’ behavior in adjacent time periods, drawing on restrictions that are necessary for dynamic optimization.

²Finite dependence requires that, starting from any two different states, there exist two finite sequence of actions that will lead to the same distribution over states. See Altug and Miller (1998) for an early application exploiting finite dependence, and Arcidiacono and Ellickson (2011) for further discussion of the role of finite dependence in DDC models.

2013; Traiberman, 2018; De Groote and Verboven, 2018; Diamond et al., 2018).³ However, the underlying class of models has yet to receive a comprehensive econometric treatment. This paper aims to fill that gap.

To that end, we provide sufficient conditions for both parametric and nonparametric identification of DDC model primitives. The identification arguments are constructive and lead naturally to an estimator, the ECCP estimator. Like other conditional-choice-probability-based estimators (Hotz and Miller, 1993; Aguirregabiria and Mira, 2002; Pesendorfer and Schmidt-Dengler, 2008), the ECCP estimator involves two steps: first estimating conditional choice probabilities, and then estimating parameters of the model. Unlike other CCP estimators, the second step amounts to estimating a linear regression equation, making it easy to implement and computationally light.⁴ Both steps can be implemented using standard tools available in econometrics software packages such as STATA and R. We establish the consistency and asymptotic normality of the estimator, and we illustrate the finite-sample performance of the ECCP estimator in a Monte Carlo study of dynamic demand for durable goods. We find that it performs well in finite samples while estimation techniques ignoring measurement problems in state variables can be substantially biased.

To identify and estimate structural parameters, the ECCP approach exploits moment restrictions implied by dynamic optimization. These moment restrictions can be constructed as long as (a) the state variables can be decomposed into *agent-specific* state variables and *market-level* states, and (b) there is finite dependence in the agent-specific state variables. Under these conditions, ECCP equations (like Euler equations in general) involve relationships between current behavior and expected future behavior. Assuming that agents have rational expectations, the ECCP equation with *observed* (rather than expected) behavior in successive time periods can serve as a valid estimating equation. Furthermore, unobserved variables and measurement error can be represented in ECCP equations in the same way that they are represented in standard regression equations, and handled similarly through the use of instrumental variables.⁵

The ECCP approach allows researchers to deal with endogeneity problems using standard linear instrumental variables techniques. No assumptions are needed regarding the evolution of the unobservable shocks, except that they satisfy exclusion restrictions (i.e., they are uncorrelated with instrumental variables). For example, cost shifters may be used as instruments for unobserved demand shocks (De Groote and Verboven, 2018); lagged observed states form another possible source of instruments available to researchers.

Unlike other approaches to estimating DDC models, the ECCP approach does not require a

³Scott (2013) studies how US agricultural policies affect farmers' land use choices; De Groote and Verboven (2018) investigate the adoption of solar photovoltaic systems for electricity production; Traiberman (2018) focuses on workers' occupational reallocations from trade liberalization; Diamond et al. (2018) study the welfare impacts of rent control on tenants and landlords.

⁴As long as agents' payoffs are linear-in-parameters, the resulting ECCP equation will also be linear-in-parameters.

⁵The estimating equations involve CCPs in successive time periods, and so require sufficiently rich panel data so that CCPs can be estimated separately for each cross section.

full model of the agent’s information set. This is a significant advantage, as modeling market-levels state variables can be both conceptually and computationally demanding in many applications. For instance, it may be difficult for the econometrician to model the evolution of certain state variables (e.g., when government policy is in flux as in De Groot and Verboven (2018)); observed market states may be insufficient to capture the true extent of market heterogeneity (e.g., unobserved local price variation as in Scott (2013)); or it may be difficult to reliably estimate how market-level state variables evolve, as when the dimensionality of the state space is large relative to the length of the panel data (e.g., when the state-space involves industry-level productivity shocks for a variety of industries as in Traiberman (2018)).

While ECCP estimators avoid the need for a full model of how state variables evolve, counterfactual analysis typically requires a full model to solve and simulate a counterfactual dynamic problem. Thus, while the ECCP approach has an advantage in requiring weaker modeling assumptions for estimation, it is not immediately clear whether that advantage carries through to counterfactual analysis. In our Monte Carlo exercise, we consider a dynamic demand model with correlation between the observed price and an unobserved demand shock. In computing long-run demand elasticities from this model, we find that using the ECCP approach for estimation and then imposing restrictive assumptions on unobservable shocks only for counterfactuals outperforms an approach that imposes the same restrictive assumptions in both estimation and counterfactual simulation.⁶

Related Literature. This paper relates to several important prior studies examining the identification and estimation of structural DDC models. In addition to the previously mentioned contributions (which focused mostly on estimation), there exists a growing literature on identification that builds on the seminal work of Rust (1994) and Magnac and Thesmar (2002). Arcidiacono and Miller (2017) investigate the nonparametric identification of DDC models for both stationary and nonstationary environments in the presence of long and short panel data, relying on *single-action* finite dependence to eliminate continuation values.⁷ Our results build on those of Arcidiacono and Miller (2017) by (a) allowing for endogeneity problems in market-level state variables, and (b) showing that parametric restrictions commonly imposed in applied work allow researchers to relax the single-action requirement and identify model parameters under general patterns of finite dependence. Blevins (2014) shows how models with discrete and continuous choice variables can be identified in the presence of (observed) continuous states under the conditional independence assumption on the unobservables. Our results can be combined with those of Blevins (2014) to gen-

⁶Specifically, we set the unobserved demand shock at its long-run mean when calculating the long-run demand elasticity. The unobserved demand shock has zero mean with no loss of generality, given that a non-zero mean term is absorbed by the constant term in the linear regression.

⁷Single-action finite dependence requires that the same action must be chosen repeatedly finitely many times in order to reset the distribution of the state variables. The empirical papers cited above (Scott, 2013; Traiberman, 2018; De Groot and Verboven, 2018; Diamond et al., 2018) make use of *one-period* finite dependence (specifically, terminal or renewal actions), which is a special case of single-action finite dependence.

erate an augmented set of moment restrictions that incorporates both the ECCP and the standard Euler equations for continuous variables.

Serially correlated unobserved state variables in structural dynamic models are a widespread problem without a standard econometric solution. Arcidiacono and Miller (2011), Hu and Shum (2012), Blevins et al. (2017), and Berry and Compiani (2017) represent four other approaches to estimating DDC models with unobserved state variables. Arcidiacono and Miller (2011) restrict individual unobserved heterogeneity to have a discrete distribution that is invariant over time (see also Kasahara and Shimotsu (2009)). Hu and Shum (2012) allow the unobserved state to follow a Markov process, but require it to be a scalar, to have the same cardinality as the action space, and to be realized before the realization of the observed states. Blevins et al. (2017) use particle filtering methods to allow for an unobservable state that follows a first-order autoregressive process. Berry and Compiani (2017) propose the use of lagged exogenous state variables as instrumental variables and obtain partial identification of payoff function parameters in a discrete choice setting. Whereas the ECCP approach allows for market-level unobserved heterogeneity very flexibly, these papers allow for individual-level unobserved heterogeneity with stronger restrictions on the nature of that heterogeneity. As such, the ECCP approach complements, and can be combined with, these other contributions.

Models of dynamic demand is one area where dynamic discrete choice models have been estimated while allowing for endogeneity concerns with market-level variables (i.e., prices being correlated with market demand shocks). Existing studies impose strong functional form restrictions on how observed and unobserved state variables evolve (Hendel and Nevo, 2006; Gowrisankaran and Rysman, 2012; Melnikov, 2013). Such restrictions – in particular, inclusive value sufficiency – limit the dimensionality of the state space to render the problem tractable, facilitating an estimation approach that relies on solving agents’ dynamic problem. However, such restrictions effectively impose that in two different states (e.g., high price with low unobserved product quality vs. low price with high unobserved product quality), consumers must have the same expectations. The ECCP approach avoids both the need to solve the dynamic problem and the need to specify a restrictive process for how state variables evolve in the estimation procedure.⁸

Morales et al. (2015) and Dickstein and Morales (2018) pioneered the use of Euler-equation-like estimators for DDC models using moment inequalities. The ECCP approach, like most estimation approaches for DDC models since Rust (1987), relies on the existence of conditionally independent individual payoff shocks with a distribution that is known *ex-ante* (e.g., logit). In contrast, the moment inequalities approach allows researchers to impose less structure on payoff shocks, requiring minimal distributional assumptions on the error term. However, the moment inequalities approach yields only partially identified parameter estimates, and it has only been shown to be robust to

⁸Although the ECCP approach can accommodate finitely many unobserved types, as in Scott’s (2013) use of Arcidiacono and Miller’s (2011) strategy, it does not, on its own, accommodate richer unobserved consumer heterogeneity (e.g., continuously distributed random coefficients on flow utilities). We regard extensions of the ECCP approach that incorporate richer forms of consumer heterogeneity as an important avenue for future work.

a limited set of endogeneity concerns (e.g., when the error terms are fixed effects that can be differenced out; see Pakes (2010)).

The remainder of the paper is organized as follows. Section 2 describes the framework, and introduces three applied examples (demand for durable goods, dynamic land use choice, and adoption of technology). Section 3 derives the ECCP equations, and illustrates them in the context of the applied examples. Section 4 discusses the identification results. Section 5 presents the ECCP estimator and establishes its asymptotic properties. Section 6 presents the Monte Carlo evidence, and Section 7 concludes. (All proofs can be found in the Appendix.)

2 Model

Time is discrete and the horizon is infinite. There are N agents operating independently in M independent markets, such as geographical locations. Every period t , agent i in market m chooses an action $a_{imt} \in \mathcal{A} = \{0, \dots, A\}$, $A < \infty$, with the goal of maximizing her expected discounted sum of payoffs

$$E \left[\sum_{\tau=0}^{\infty} \beta^{\tau} \Pi_{im,t+\tau} | \mathcal{I}_{imt} \right],$$

where Π_{imt} denotes per-period payoffs, $\beta \in (0, 1)$ is the discount factor, and $E[\cdot | \mathcal{I}_{imt}]$ denotes the expectation operator conditioned on the information set \mathcal{I}_{imt} available to agent i in market m at time t .

The payoff function Π_{imt} depends on the state variables $\mathbf{s}_{imt} = (k_{imt}, w_{mt}, \eta_{mt}, \varepsilon_{imt})$, where (k_{imt}, w_{mt}) are observed by the econometrician, while $(\eta_{mt}, \varepsilon_{imt})$ are not. The observed states $k_{imt} \in \mathcal{K}$ are “controlled”: their evolution can be affected by the agent’s actions; such states may include a firm’s capital stock, size or type of product. The market-level observed as well as unobserved states, w_{mt} and η_{mt} , cannot be affected by the agent’s actions; such states may include market demand variables, aggregate input prices, or the government policy environment. We collect these market-level states into the vector $\omega_{mt} = (w_{mt}, \eta_{mt}) \in \Omega$. We assume \mathcal{K} is finite, as usually done in the literature, but we allow ω_{mt} to be continuous. Neither assumption is important and our results apply to both discrete and continuous states. Our results do not require us to specify the dimension of η_{mt} , so there may be many unobserved market-level state variables. Finally, the unobservable state $\varepsilon_{imt} = (\varepsilon_{0imt}, \dots, \varepsilon_{Aimt})$ is i.i.d. across agents and time with a distribution function that is absolutely continuous with respect to Lebesgue measure in \mathbb{R}^{A+1} .

As usual, in each period t agents observe the state variables \mathbf{s}_{imt} , make choices a_{imt} , flow payoffs are then realized, and states evolve. Agent’s information set \mathcal{I}_{imt} therefore includes all current and past state variables \mathbf{s}_{imt} , as well as all past actions. We assume \mathbf{s}_{imt} follows a controlled first-order

Markov process with transition distribution function that factors as,

$$F(\mathbf{s}_{imt+1}|a, \mathbf{s}_{imt}) = F^k(k_{imt+1}|a, k_{imt}, w_{mt}) F^\omega(\omega_{mt+1}|\omega_{mt}) F^\varepsilon(\varepsilon_{imt+1}). \quad (1)$$

Equation (1) limits our focus to settings with small decision makers, as opposed to dynamic games. It says that market-level state variables are perceived as exogenous by individual agents, ruling out settings where an individual agent's decision can have aggregate impacts. It also rules out settings where there are externalities acting through the agent-level state variables k , such as when one agent's state depends on a neighbor's decision.

The per period payoff is given by,

$$\Pi(a, \mathbf{s}_{imt}) = \bar{\pi}(a, k_{imt}, w_{mt}) + \xi(a, k_{imt}, \omega_{mt}) + \varepsilon_{aimt}. \quad (2)$$

While it is standard practice to include additively separable idiosyncratic shocks ε_{imt} in dynamic discrete choice models, the other unobserved term of the payoff function, ξ , deserves some discussion. The $\bar{\pi}(a, k, w)$ term depends on observed market-level state-variables, w , and so the $\xi(a, k, \omega)$ term captures the impact of *unobserved* market-level state variables on the flow payoff. Thus, ξ may reflect mis-measured profits or unobservable costs. It is important to stress that ξ is a function of state variables, and it need not be a state variable itself. That is, the value of ξ at time t may not be a sufficient statistic for the distribution of future values of ξ at time $t + 1$. In the background, there are some (potentially) more informative state variables ω that agents use to form their expectations about the future. Consequently, ξ need not evolve according to a first-order Markov process, while ω does.

In addition, $\bar{\pi}$ and ξ may be correlated because they may depend on the same state variables w , or because w and η may not be independent to each other. These are reasons to consider the use of instrumental variables to identify and estimate the model. Finally, note that ξ may accommodate measurement errors in $\bar{\pi}$. Without loss of generality, we assume that ξ is mean zero, and, to simplify notation, define

$$\pi \equiv \bar{\pi} + \xi.$$

Next, we introduce three applied examples to illustrate what the ξ term can capture in practice, and we preview the advantages of the ECCP approach in the context of each example.

Example 1: Durable Demand. In our Monte Carlo study, we consider a model of dynamic demand for a durable good. Each period, a consumer chooses whether or not to purchase the good, in turn discarding the old version of the good if she already owns a unit. The state variable that a consumer controls is a dummy variable indicating whether or not she owns the good. When the consumer owns the good, there is a chance of product failure, leaving the consumer without the product in the subsequent period.

We consider two exogenous state variables: an observable price, and an unobservable quality

shock captured by ξ . Other approaches to estimating dynamic demand systems (e.g. Hendel and Nevo (2006); Gowrisankaran and Rysman (2012)) rely on solving a consumer’s dynamic problem and imposing strong functional form restrictions on how observed and unobserved state variables evolve. Such restrictions – the so-called “inclusive value sufficiency” – limit the dimensionality of the state space to render the problem tractable. However, they effectively impose that in two different states (e.g., high price with low quality vs. low price with high quality), consumers must have the same expectations. The ECCP approach avoids both the need to solve the dynamic problem and the need to specify a restrictive process for how state variables evolve.

Example 2: Land Use Change. Scott (2013) adopts this framework to model farmers’ land use choices. In each period, each farmer chooses whether to plant crops or not. The farmer’s controlled state variable is the number of years since the field was last in crops (reflecting the condition of the plot). The farmer’s payoff consists of returns from the chosen land use and involve both observed (e.g. crop prices and yields) and unobserved revenues and costs – for example, expected returns may be calculated based on input and output prices, but price data may only be available at a regional level, leading to measurement error in returns coming from the unobserved local price variation. The ξ term captures this measurement error in expected returns.

Measurement error can plausibly render ξ serially correlated. High-priced localities in one year are likely to be high-priced localities the next year. The ECCP approach can be implemented in the presence of such serially correlated unobserved local price variation.

Furthermore, in the context of agricultural markets, it is difficult to model the evolution of *observed* market-level state variables, given the large set of variables that can influence farmers’ expected returns (e.g., technological conditions, uncertain government policies, crop stocks, etc.).⁹ As previously mentioned, the ECCP approach does not require specifying what all the relevant market-level state variables are in the decision making process, nor specifying (and estimating) how all such variables evolve.

Example 3: Technology Adoption. De Groote and Verboven (2018) study the adoption of renewable energy technologies for electricity production: the solar photovoltaic (PV) systems. In every period, a household may either choose to not adopt a PV, or it may choose to adopt one of several available PV alternatives. The adoption decision is a terminating state; not adopting provides the option of waiting for decreased prices or increased adoption subsidies. In this context, the ξ term captures unobserved quality shocks (or it could be interpreted as unobserved adoption cost shocks). These quality shocks could be correlated with price, and the ECCP makes it straightforward for De Groote and Verboven to instrument for price using cost shifters.

Given that the adoption subsidies are large (and that their levels change substantially over time), government policy is an important observable state variable. The ECCP approach allows for

⁹For example, see Wright (2014) for a discussion of how traditional models of competitive storage fail to explain recent grain price movements. Wright argues that changes in government policy are crucial in explaining recent grain market behavior.

the estimation of the dynamic model without requiring an explicit specification of the law of motion for government policy. Traditional approaches would either assume that changes in government policy are either fully anticipated or complete surprises (Kalouptsi, 2018); the ECCP approach only requires an assumption that agents have rational expectations about such changes.

2.1 Value Functions and Choice Probabilities

Let $V(\mathbf{s}_{imt})$ be the value function of the dynamic programming problem, i.e., the expected discounted stream of payoffs under optimal behavior. By Bellman's principle of optimality,

$$V(\mathbf{s}_{imt}) = \max_{a \in \mathcal{A}} \{ \Pi(a, \mathbf{s}_{imt}) + \beta E[V(\mathbf{s}_{imt+1}) | a, \mathbf{s}_{imt}] \}.$$

Following the literature, we define the *ex ante value function*:

$$V(k_{imt}, \omega_{mt}) \equiv \int V(k_{imt}, \omega_{mt}, \varepsilon_{imt}) dF^\varepsilon(\varepsilon_{imt}),$$

and the *conditional value function*:

$$v_a(k_{imt}, \omega_{mt}) = \pi(a, k_{imt}, \omega_{mt}) + \beta E[V(k_{imt+1}, \omega_{mt+1}) | a, k_{imt}, \omega_{mt}]. \quad (3)$$

The agent's optimal policy is given by the conditional choice probabilities (CCPs):

$$p_a(k, \omega) = \int 1 \{ v_a(k, \omega) + \varepsilon_a \geq v_j(k, \omega) + \varepsilon_j, \text{ for all } j \in \mathcal{A} \} dF^\varepsilon(\varepsilon), \quad (4)$$

where $1\{\cdot\}$ is the indicator function. Define the $(A+1) \times 1$ vector of conditional choice probabilities $p(k, \omega) = \{p_a(k, \omega) : a \in \mathcal{A}\}$.¹⁰

Finally, it is worth noting that for any (a, k, ω) , there exists a real-valued function $\psi_a(\cdot)$ derived only from F^ε that satisfies the following equality (Arcidiacono and Miller (2011, Lemma 1)):

$$V(k, \omega) - v_a(k, \omega) = \psi_a(p(k, \omega)). \quad (5)$$

When ε_{imt} follows the type 1 extreme value distribution, then $\psi_a(p(k, \omega)) = \gamma - \ln p_a(k, \omega)$, where γ is the Euler constant.¹¹

¹⁰Choice probabilities are invariant to scale normalizations; here we normalize the scale parameter to one. I.e., we take $\Pi_a = \bar{\pi}_a + \xi_a + \sigma \varepsilon_a$, with $\sigma = 1$.

¹¹The Arcidiacono-Miller Lemma can be derived from the Hotz-Miller inversion (Proposition 1 of Hotz and Miller (1993)). Chiong et al. (2016) propose a novel approach that can calculate ψ_a for a broad set of distributions F^ε .

3 ECCP Equations

Assume the available data set is $\{y_{imt} = (a_{imt}, k_{imt}, w_{mt}, z_{mt}) : i = 1, \dots, N; m = 1, \dots, M; t = 1, \dots, T\}$, where the vector z_{mt} consists of instrumental variables, as explained below. For our identification arguments (presented in Section 4), we assume the joint distribution of y_{imt} , $\Pr(y)$, is known.

Even though the market state ω_{mt} is not fully observed, the conditional choice probabilities $p_a(k, \omega_{mt})$ for a particular market m and time period t can be estimated given a sufficiently rich cross section of agents within market m . That is, if we observe many agents with each value of the individual state within each market and time period, then $p_a(k, \omega_{mt})$ can be estimated using a simple frequency estimator. The law of motion for the agent-controlled states, $F^k(k_{imt+1}|a, k_{imt}, w_{mt})$, can also be estimated using a simple frequency estimator (when it is not known in advance). For our identification results, we treat $p_a(\cdot)$ and $F^k(\cdot)$ as known objects.¹² In contrast, the law of motion for market-level states $F^\omega(\omega_{mt+1}|\omega_{mt})$ cannot be estimated without strong assumptions. (As state ω_{mt} is not fully observed, F^ω cannot be estimated directly from the data.) Following the literature, we assume that the discount factor β and the distribution of idiosyncratic shocks F^ε are known (which implies $\psi_a(\cdot)$, for all a , are known as well).

The objective is to identify the per-period payoff function $\bar{\pi}$. It is well-known that the standard dynamic discrete choice model, in which $\omega = w$ and $\xi = 0$, so that there are no unobserved states, is nonparametrically not identified (Rust, 1994; Magnac and Thesmar, 2002). Indeed, intuitively, $\bar{\pi}$ has $(A + 1)S$ parameters (where S is the size of the state space), while there are only AS independent CCPs (given that the CCP's have to add up to unity conditional on each state). Thus, there are S free payoff parameters and S restrictions will need to be imposed. Although there are several ways to do so, restrictions that suffice to identify the model parameters are all equivalent to pre-specifying the payoff for some reference action J in all states (for a detailed discussion see Kalouptsi et al. (2017)).

Since the model here is more general than the standard model without unobserved states, we need to impose such restrictions for identification. In addition, we assume rational expectations, finite dependence, and valid instruments, which we discuss in detail next.

Assumption 1. (*Rational expectations*) *Agent's expectations conditional on the information set \mathcal{I}_{imt} correspond to the conditional expectations of the true data generating process given \mathcal{I}_{imt} .*

Rational expectations is a common assumption in the literature. However, in contrast to the standard procedure (Rust, 1987; Aguirregabiria and Mira, 2002), we make use of *realized* values of agents' future expected payoffs as a noisy measure of agents' expected future payoffs. The use of realized values allows us to relax typical assumptions about how agents form beliefs about

¹²See Appendix A.3 for a formal justification. In practice, some smoothing across markets or individual states may be needed.

the evolution of the state variables. To that end, we rely on expectational errors (also known as forecast errors), defined as follows:

Definition 1. (*Expectational error*) For any function $h(k, \omega)$ and particular realization $\omega^* \in \Omega$,

$$e^h(k', \omega, \omega^*) \equiv E_{\omega'|\omega} [h(k', \omega') | \omega] - h(k', \omega^*),$$

$$e^h(a, k, \omega, \omega^*) \equiv \sum_{k'} e^h(k', \omega, \omega^*) F^k(k' | a, k, \omega),$$

where k' and ω' denote next period values for k and ω .

The expectational error $e^h(k', \omega, \omega^*)$ is the prediction error of $h(k', \omega')$ for a particular realized value of the individual state k' ; and $e^h(a, k, \omega, \omega^*)$ is the corresponding prediction error conditioned on k, ω , and a , integrating over the realizations of the individual state k' . An important property of the expectational error (and that we make use of later on) is that it is mean independent of variables that belong to the agent's information set.¹³

Lemma 1. *Suppose Assumption 1 holds. Then,*

- (i) *For any action a and individual state k , the expectational error term $e^h(a, k, \omega_{mt}, \omega_{mt+1}^*)$ is mean zero given the information set available to the agent: $E[e^h(a, k, \omega_{mt}, \omega_{mt+1}^*) | \mathcal{I}_{imt}] = 0$.*
- (ii) *For $z_{mt} \in \mathcal{I}_{imt}$, $E[e^h(a, k, \omega_{mt}, \omega_{mt+1}^*) | z_{mt}] = 0$, for all a and k .*
- (iii) *Expectational errors are serially uncorrelated.*

Expectational errors are useful as they allow us to “dispose” of the actual expectations. Indeed, combining the conditional value function (3) with the Arcidiacono-Miller Lemma (5), we obtain:

$$\pi(a, k_{imt}, \omega_{mt}) = V(k_{imt}, \omega_{mt}) - \beta E[V(k_{imt+1}, \omega_{mt+1}) | a, k_{imt}, \omega_{mt}] - \psi_a(k_{imt}, \omega_{mt}), \quad (6)$$

where $\psi_a(k, \omega)$ is a short-cut notation for $\psi_a(p(k, \omega))$. Next, note that the expectation $E[V(k_{imt+1}, \omega_{mt+1}) | a, k_{imt}, \omega_{mt}]$ is given by:

$$\begin{aligned} & \sum_{k'} \int_{\omega'} V(k', \omega') dF^{\omega}(\omega' | \omega_{mt}) F^k(k' | a, k_{imt}, \omega_{mt}) \\ &= \sum_{k'} (E_{\omega'|\omega_{mt}} [V(k', \omega') | \omega_{mt}]) F^k(k' | a, k_{imt}, \omega_{mt}) \\ &= \sum_{k'} V(k', \omega_{mt+1}) F^k(k' | a, k_{imt}, \omega_{mt}) + e^V(a, k_{imt}, \omega_{mt}, \omega_{mt+1}), \end{aligned} \quad (7)$$

where $e^V(\cdot)$ is the expectational error of the value function $V(\cdot)$. Substituting (7) in (6), we obtain

¹³In models in which agents have perfect foresight, there is no prediction error and so $e^h(\cdot) = 0$. In contrast, when agents have biased beliefs, their conditional expectations do not necessarily coincide with the true (data generating) conditional expectations, which implies $E[e^h(a, k_{imt}, \omega_{mt}, \omega_{mt+1}^*) | \mathcal{I}_{imt}] \neq 0$.

the following equation, which is the basis for our identification results:

$$\begin{aligned} & \pi(a, k_{imt}, \omega_{mt}) + \beta e^V(a, k_{imt}, \omega_{mt}, \omega_{mt+1}) \\ = & V(k_{imt}, \omega_{mt}) - \beta \sum_{k'} V(k', \omega_{mt+1}) F^k(k'|a, k_{imt}, \omega_{mt}) - \psi_a(k_{imt}, \omega_{mt}). \end{aligned} \quad (8)$$

Next, we simplify the notation and use (m, t) subscripts to denote functions that depend on ω_{mt} . We rewrite payoffs as $\pi_{mt}(a, k_{imt}) \equiv \pi(a, k_{imt}, \omega_{mt})$, while $\bar{\pi}_{amt}$, ξ_{amt} , $V_{mt}(k_{imt})$, $p_{amt}(k_{imt})$ and $\psi_{amt}(k_{imt})$ are similarly defined. We also make use of matrix notation, so that π_{amt} is a $K \times 1$ vector that stacks $\pi_{mt}(a, k_{imt})$ for all $k \in \mathcal{K}$ (and similarly for the vectors V_{mt} , p_{amt} and ψ_{amt}). Therefore, (8) in matrix form becomes:

$$\pi_{amt} + \beta e_{am,t,t+1}^V = V_{mt} - \beta F_{amt}^k V_{mt+1} - \psi_{amt} \quad (9)$$

for all a , where $e_{am,t,t+1}^V$ stacks $e_{mt,t+1}^V(a, k) \equiv e^V(a, k_{imt}, \omega_{mt}, \omega_{mt+1})$ for all $k \in \mathcal{K}$, and F_{amt}^k is the $K \times K$ transition matrix for market m at time period t with (n, l) element equal to $\Pr(k_{imt+1} = k_l | a, k_{imt} = k_n, \omega_{mt})$.

Note that the time- t value function term, V_{mt} , can be removed by simply differencing equation (9) across two different actions, a and j :

$$\psi_{jmt} - \psi_{amt} = \pi_{amt} - \pi_{jmt} + \beta (e_{am,t,t+1}^V - e_{jm,t,t+1}^V) - \beta (F_{jmt}^k - F_{amt}^k) V_{mt+1}. \quad (10)$$

Finally, separating payoffs into the observable and unobservable components:

$$\begin{aligned} \psi_{jmt} - \psi_{amt} = & \bar{\pi}_{amt} - \bar{\pi}_{jmt} + \xi_{amt} - \xi_{jmt} + \beta (e_{am,t,t+1}^V - e_{jm,t,t+1}^V) \\ & - \beta (F_{jmt}^k - F_{amt}^k) V_{mt+1}. \end{aligned} \quad (11)$$

Equation (11) is one step away from being a regression equation. Indeed, in a static model (i.e., $\beta = 0$), equation (11) simplifies further to $\psi_{jmt} - \psi_{amt} = \bar{\pi}_{amt} - \bar{\pi}_{jmt} + \xi_{amt} - \xi_{jmt}$. On the left hand side we have an observed dependent variable; in the logit model, $\psi_{jmt} - \psi_{amt} = \ln\left(\frac{p_{amt}}{p_{jmt}}\right)$. The right-hand side includes the payoff functions $\bar{\pi}$, and the unobservables ξ . If the payoffs for one reference action were known or pre-specified (e.g., set $\bar{\pi}_{jmt} = 0$, which is common practice), then $\bar{\pi}_{amt}$ can be estimated using standard methods for regression models. When the ξ_{mt} terms are correlated with w_{mt} (which is an argument of $\bar{\pi}_{amt}$), then we need instruments for w_{mt} . In this case, we can estimate the model parameters using standard instrumental variables estimators.

Thus, the remaining term to be dealt with before (11) can be used for identification and estimation of a dynamic model are the time- $(t + 1)$ value function terms. As we show below, these terms can be eliminated using applications of the Arcidiacono-Miller Lemma, as long as the

transition process F^k satisfies the finite dependence property.¹⁴

Definition 2. (*Finite Dependence*) A pair of choices a and j satisfies τ -period finite dependence if there exists two sequences of actions (a, a_1, \dots, a_τ) and (j, j_1, \dots, j_τ) such that, for all t ,

$$F_{amt}^k F_{a_1 mt+1}^k \cdots F_{a_\tau mt+\tau}^k = F_{jmt}^k F_{j_1 mt+1}^k \cdots F_{j_\tau t+\tau}^k. \quad (12)$$

We say that τ -period finite dependence holds for the model if all pairs of actions satisfy τ -period finite dependence.

Specifically, the τ -period finite dependence property holds if, starting from any two distributions of individual states at the beginning of time period t , there are sequences of actions (not necessarily optimal) that result in the same distribution of state variables at $t + \tau$.¹⁵ *Single-action* τ -period finite dependence is a special case that requires the sequences of actions to be of the type (a, J, \dots, J) and (j, J, \dots, J) , for any actions a and j , and some action J .

Common special cases of one-period dependence are renewal and terminal actions. Action J is a renewal action if, taking action J in period $t + 1$ leads to the same distribution of states at the beginning of time period $t + 2$, regardless of which state the agent was in during period t . Examples of renewal actions are replacing the bus engine (Rust, 1987), planting crops (Scott, 2013), choosing occupations (Traiberman, 2018), and migration decisions (Diamond et al., 2018). A terminal action ends the decision making process. Examples of terminal actions include a worker retiring (Rust and Phelan, 1997), a mortgage owner defaulting (Bajari et al., 2016), and a household adopting a PV system (De Groote and Verboven, 2018). For both renewal and terminal actions, (12) simplifies to

$$F_{amt}^k F_{Jmt+1}^k = F_{jmt}^k F_{Jmt+1}^k. \quad (13)$$

for all t and all a, j .

Other models require more than one period to eliminate dependence in state variables. Altug and Miller (1998) consider female labor supply with human capital appreciation and depreciation (in which full depreciation of human capital occurs if a woman stays out of the workforce for τ periods). Bishop (2008) considers a model of migration with relocation costs that depend on whether the household previously lived in the same region they move to.

ECCP estimation with one-period finite dependence. For the sake of exposition, we first focus on deriving a ECCP regression equation with one-period finite dependence (i.e., for

¹⁴See Arcidiacono and Ellickson (2011) for further discussion of finite dependence. Note that we make use of deterministic sequences of actions in Definition 2, while Arcidiacono and Miller (2011) allows for stochastic sequences. Our parametric identification result can be extended to incorporate stochastic sequences.

¹⁵The terms in the sequences (a, a_1, \dots, a_τ) and (j, j_1, \dots, j_τ) depend on the particular initial pair of actions (a, j) chosen; for ease of exposition, we do not incorporate this dependence on the initial pairs into our notation.

models with renewal and terminal actions). In Section 4, we show that the ECCP approach extends to models with τ -period finite dependence.

Let J be a renewal or terminal action, and take two sequences (a, J) and (j, J) , for any actions a and j . If we recursively substitute V_{mt+1} , in equation (11), we obtain

$$\begin{aligned} \psi_{jmt} - \psi_{amt} = & \bar{\pi}_{amt} - \bar{\pi}_{jmt} + \xi_{amt} - \xi_{jmt} + \beta (e_{am,t,t+1}^V - e_{jm,t,t+1}^V) \\ & - \beta (F_{jmt}^k - F_{amt}^k) (\bar{\pi}_{Jmt+1} + \xi_{Jmt+1} + \psi_{Jmt+1}). \end{aligned} \quad (14)$$

because the V_{t+2} portions of the value function cancel conditional on action J being chosen in period $t + 1$.¹⁶

One can view (14) as an Euler equation. Traditionally, Euler equations express intertemporal first-order conditions implied by optimal dynamic behavior. While equation (14) is derived directly from optimality conditions, and while the condition relates choice probabilities in different time periods, it is not immediately obvious that it can be viewed as a first-order condition. However, following Aguirregabiria and Magesan (2013), if we treat the choice probabilities as choice variables themselves, we could derive equation (14) as a first-order condition, formalizing the analogy between (14) and traditional Euler equations for dynamic problems with continuous choice variables.¹⁷

Rearranging (14) we obtain

$$\psi_{jmt} - \psi_{amt} + \beta (F_{jmt}^k - F_{amt}^k) \psi_{Jmt+1} = \bar{\pi}_{amt} - \bar{\pi}_{jmt} - \beta (F_{jmt}^k - F_{amt}^k) \bar{\pi}_{Jmt+1} + u_{ajmt}, \quad (15)$$

where the unobservable term is $u_{ajmt} = \tilde{\xi}_{ajmt} + \tilde{e}_{ajmt}^V$, with

$$\tilde{\xi}_{ajmt} = (\xi_{amt} - \xi_{jmt}) + \beta (F_{amt}^k - F_{jmt}^k) \xi_{Jmt+1}, \quad (16)$$

$$\tilde{e}_{ajmt}^V = \beta (e_{am,t,t+1}^V - e_{jm,t,t+1}^V). \quad (17)$$

The Euler equation (15) can be used to construct moment restrictions that are well-suited for identification and estimation of the payoff function $\bar{\pi}$. More simply, we can treat it as a (nonparametric) regression model, where the right-hand-side contains the parameters of interest $\bar{\pi}$ and the unobservables u , while the left-hand-side contains the endogenous observed variables.

Next, we return to each of our applied examples, illustrating the regression equations from the ECCP equation (15) in each context. For the first example, we also outline the steps involved in

¹⁶Formally, use (9) for J in $t + 1$ to solve for V_{mt+1} and replace the latter in (9) for any a in t :

$$\begin{aligned} \pi_{amt} + \beta e_{am,t,t+1}^V = & V_{mt} - \psi_{amt} \\ & - \beta F_{amt}^k [\pi_{Jmt+1} + \beta F_{Jmt+1}^k E_{t+1} [V_{mt+2}] + \psi_{Jmt+1}]. \end{aligned}$$

Next, evaluate the above at a and j and subtract to obtain (14) using property (13).

¹⁷Aguirregabiria and Magesan (2013) do not allow for unobserved state variables as we do, but their approach to deriving Euler equations by treating choice probabilities as choice variables could be applied in our setting.

the derivation of the ECCP equation, for the derivation for the special case is considerably simpler than the general derivation above.

Example 1 (continued): Dynamic Demand. Each period t , consumer i in market m decides whether or not to purchase a unit of a durable good at price w_{mt} . The choice set is $\mathcal{A} = \{b, nb\}$, where b means buying the good, and nb means not buying the good.

Consumer i controls state $k_{imt} \in \{0, 1\}$ where $k_{imt} = 0$ if the consumer does not have a unit of the good at the beginning of time period t , and $k_{imt} = 1$ when she already owns it. If consumer i chooses not to buy a new unit of the good ($a_{imt} = nb$) when she already owns it ($k_{imt} = 1$), then there is probability ϕ of product failure, resulting in $k_{imt+1} = 0$. Formally, state k evolves as follows:

$$Pr(k_{imt+1} = 1 | k_{imt}, a_{imt}, w_{mt}) = \begin{cases} 1 & \text{if } a_{imt} = b \\ 0 & \text{if } a_{imt} = nb, k_{imt} = 0 \\ 1 - \phi & \text{if } a_{imt} = nb, k_{imt} = 1. \end{cases}$$

Note that action $a_{imt} = b$ is a renewal action.

There is one observed exogenous state variable, the price w_{mt} , and one unobserved exogenous state, a quality shock ξ_{mt} .¹⁸ The consumer enjoys the following flow utility if purchasing the product:

$$\pi(b, k_{imt}, \omega_{mt}) = \theta_0 + \theta_1 w_{mt} + \xi_{mt},$$

where $\theta_1 < 0$ so that demand slopes down. When not purchasing the product, the consumer enjoys utility

$$\pi(nb, k_{imt}, \omega_{mt}) = \begin{cases} \theta_0 & \text{if } k_{imt} = 1 \\ 0 & \text{if } k_{imt} = 0. \end{cases}$$

Notice that the θ_0 term appears in the utility function when the good is being consumed, either through purchase or because the consumer already owns the good. θ_0 can be interpreted as the flow value of consumption. Since the demand shock ξ enters into consumption only conditional on purchase, we can interpret it as a quality shock that the consumer only cares about when the product is newly purchased, such as the quality of the in-store experience or some non-durable services associated with the durable good.

As an illustration, we derive the Euler equation (14) and corresponding regression in the context of this model. Assume logit errors. Starting with the Arcidiacono-Miller Lemma (5), we have the following relationship between choice probabilities and conditional value functions:

$$\ln \left(\frac{p_b(k_{imt}, \omega_{mt})}{p_{nb}(k_{imt}, \omega_{mt})} \right) = v_b(k_{imt}, \omega_{mt}) - v_{nb}(k_{imt}, \omega_{mt}).$$

¹⁸Formally, we can define η_{mt} as the unobserved quality of the product (the state variable), while the function $\xi(a, k, \omega_{mt})$ equals η_{mt} when $a = b$, and equals zero when $a = nb$.

In estimating this model, it suffices to focus on $k_{imt} = 0$. Notice that when $k_{imt} = 0$ (when the agent does not own a unit of a good), the state variable k_{imt} is a deterministic function of a_{imt} . (This simplifies the derivation of the Euler equation substantially.) We expand the conditional value function given $k_{imt} = 0$, and introduce expectational errors, as follows:

$$\begin{aligned} \ln \left(\frac{p_b(0, \omega_{mt})}{p_{nb}(0, \omega_{mt})} \right) &= \theta_0 + \theta_1 w_{mt} + \xi_{mt} + \beta (V(1, \omega_{mt+1}) - V(0, \omega_{mt+1})) \\ &\quad + \beta (e^V(1, \omega_{mt}, \omega_{mt+1}) - e^V(0, \omega_{mt}, \omega_{mt+1})). \end{aligned} \quad (18)$$

We now exploit the fact that purchasing the good is a renewal action, resulting in the state $k' = 1$ regardless of what the initial state k is. As a result, when we substitute for $V(1, \omega_{mt+1})$ and $V(0, \omega_{mt+1})$ in equation (18) using the Arcidacono-Miller Lemma (5), the time- $(t+2)$ value functions cancel, leaving

$$\begin{aligned} \ln \left(\frac{p_b(0, \omega_{mt})}{p_{nb}(0, \omega_{mt})} \right) &= \theta_0 + \theta_1 w_{mt} + \xi_{mt} + \beta [-\ln p_b(1, \omega_{mt+1}) + \ln p_b(0, \omega_{mt+1})] \\ &\quad + \beta (e^V(1, \omega_{mt}, \omega_{mt+1}) - e^V(0, \omega_{mt}, \omega_{mt+1})). \end{aligned} \quad (19)$$

Equation (19) is the Euler equation for this model, and we can construct a regression equation by rearranging it to have all the choice probabilities on the left-hand side:

$$Y_{mt} = \theta_0 + \theta_1 w_{mt} + u_{mt}, \quad (20)$$

where

$$Y_{mt} = \ln \left(\frac{p_b(0, \omega_{mt})}{p_{nb}(0, \omega_{mt})} \right) + \beta \ln \left(\frac{p_b(1, \omega_{mt+1})}{p_b(0, \omega_{mt+1})} \right), \quad (21)$$

and

$$u_{mt} = \xi_{mt} + \beta (e^V(1, \omega_{mt}, \omega_{mt+1}) - e^V(0, \omega_{mt}, \omega_{mt+1})). \quad (22)$$

Given an instrumental variable that is correlated with the price w_{mt} but not with the error term u_{mt} , we can estimate equation (20) using a linear instrumental variables regression.

Example 2 (continued): Land Use Change. Scott (2013) assumes that in every period t , farmer i in county m chooses whether to plant crops or not, $\mathcal{A} = \{crops, other\}$. The state $k_{imt} \in \{0, 1, \dots, K\}$ equals the number of years since the field was last in crops, up to some limit K . This field state reflects vegetal cover and the state of the terrain. Formally, $k'(a, k) = 0$ if $a = crops$, and $k'(a, k) = \min\{k+1, K\}$ if $a = other$. Note that *crops* is a renewal action. The farmer's payoff consists of expected returns from the chosen land use and involve both observed (e.g. crop prices and yields), represented by R , and unobserved returns and costs, represented by

ξ . The payoff function can be written

$$\pi(a, k, \omega_{mt}) = \theta_0(a, k) + \theta_1 R(a, w_{mt}) + \xi_{mt}(a, k),$$

where θ_0 are switching costs parameters. Assuming logit errors, equation (15) can be rewritten for each k as

$$Y_{mt}(k) = \tilde{\theta}_0(k) + \theta_1 [R(crops, w_{mt}) - R(other, w_{mt})] + u_{mt},$$

where

$$Y_{mt}(k) = \ln\left(\frac{p_{crops,mt}(k)}{p_{other,mt}(k)}\right) + \beta \ln\left(\frac{p_{crops,mt+1}(0)}{p_{crops,mt+1}(k'(other,k))}\right),$$

$$\begin{aligned} \tilde{\theta}_0(k) &= \theta_0(crops, k) - \theta_0(other, k) \\ &\quad + \beta(\theta_0(crops, 0) - \theta_0(crops, k'(other, k))), \end{aligned}$$

$$\begin{aligned} u_{mt} &= \xi_{mt}(crops, k) - \xi_{mt}(other, k) \\ &\quad + \beta(\xi_{mt+1}(crops, 0) - \xi_{mt+1}(crops, k'(other, k))) \\ &\quad + \beta(e_{m,t,t+1}^V(0) - e_{m,t,t+1}^V(k'(other, k))). \end{aligned}$$

Again, given estimates of CCPs, one can construct $Y_{mt}(k)$ and estimate $\tilde{\theta}_0(k)$ and θ_1 using linear IV regressions. Under the assumption that $\theta_0(other, k) = 0$ for all k , we can recover the switching costs parameters $\theta_0(crops, k)$ from the regression equation intercepts $\tilde{\theta}_0(k)$.

Example 3 (continued): Technology Adoption. In every period t , a household i in region m may either choose to not adopt a photovoltaic system, $a = 0$, or it may choose to adopt one of the available PV alternatives, so that $\mathcal{A} = \{0, 1, \dots, A\}$. The adoption decision ($a > 0$) leads to a terminal state; not adopting provides the option of waiting for when the prices may have decreased or when the subsidies for adoption (or the quality) may have increased. Each alternative $a \neq 0$ is characterized by observable attributes w_{amt} (including capacity sizes, upfront investment prices, electricity cost savings, benefits from subsidies for adoption), and by the unobserved quality (which can be captured by ξ). The vector of observed state variables is given by $w_{mt} = (w_{1mt}, \dots, w_{Amt})$.

De Groot and Verboven (2018) assume the state k_{imt} is a dummy variable equal to zero if no solar panel has been installed, and equal to the type of solar panel installed (i.e., the corresponding action in \mathcal{A}) if one is already installed. The household no longer makes a PV adoption decision when $k_{imt} \geq 1$.

De Groot and Verboven assume logit errors and specify a linear-in-parameters flow payoff. Specifically, if no PV has been adopted at t (i.e., $k_{imt} = 0$), the payoff from adopting option $a \neq 0$ is

$$\pi(a, 0, \omega_{mt}) = w_{amt}\theta + \xi_{amt},$$

and the payoff of the outside option is set to zero: $\pi(0, k, \omega_{mt}) = 0$ for all k and ω_{mt} .

Given that all actions $a > 0$ are terminal, (15) can be established for any a given $k = 0$. Here we take $j = 0$ and use $J = 1$ as the terminal action in period $t + 1$ to obtain for any $a \geq 2$,

$$Y_{amt} = (w_{amt} - \beta w_{1mt+1}) \theta + u_{a0mt},$$

where

$$Y_{amt} = \ln \left(\frac{p_{amt}(0)}{p_{0mt}(0)} \right) - \beta \ln p_{1mt+1}(0),$$

$$u_{a0mt} = \xi_{amt} - \beta \xi_{1mt+1} + \beta (e_{m,t,t+1}^V(a) - e_{m,t,t+1}^V(0)),$$

noting that there is no time- $t + 1$ choice probability term corresponding to $k_{imt+1} = a$ because no decision is made once a PV system has been installed; i.e., $p_{1mt+1}(k) = 0$ for $k \geq 1$.¹⁹

As usual, prices may correlate with unobserved quality ξ_{amt} (which are part of the error term u_{a0mt}); one needs therefore to instrument for prices to estimate the model parameters. De Groot and Verboven use the prices of Chinese modules as instruments, arguing that these prices are plausibly exogenous to demand shocks in the Belgian market (and therefore a valid instrument) and that the modules are an important component of solar PV installation costs (therefore, making module prices a strong instrument).

4 Identification

We now discuss the identification of the payoff function $\bar{\pi}$ under more general forms of finite dependence. As the above examples demonstrate finite dependence allows us to cancel continuation values and construct valid estimating equations while allowing for unobservable market-level states (and measurement errors) without having to specify how those states evolve exactly.

Take two sequences of actions (a, a_1, \dots, a_τ) and (j, j_1, \dots, j_τ) satisfying τ -period finite dependence (see Definition 2).²⁰ Extending equation (15) to τ -period dependence, we obtain the

¹⁹In deriving the regression equation with terminal actions, it is convenient to assume that the idiosyncratic errors have a mean-zero distribution, e.g., the standard type-1 extreme value distribution demeaned by Euler's constant. As a result, the expected flow payoffs in the terminal state do not depend on whether or not the agent still receives an idiosyncratic shock in that state.

²⁰Recall that the sequences depend on the particular initial pair of actions (a, j) chosen. That is, we use a_d to denote the d -th action following the initial action a when the alternative initial action is j . Note that we must assume access to a panel data with $T > \tau$.

following equation (see Appendix A.1 for a detailed derivation):

$$\begin{aligned}
& \psi_{jmt} - \psi_{amt} + F_{jmt}^k \sum_{d=1}^{\tau} \beta^d \Lambda_{jmt d} \psi_{j_d mt+d} - F_{amt}^k \sum_{d=1}^{\tau} \beta^d \Lambda_{amt d} \psi_{a_d mt+d} \\
&= \bar{\pi}_{amt} - \bar{\pi}_{jmt} + F_{amt}^k \sum_{d=1}^{\tau} \beta^d \Lambda_{amt d} \pi_{a_d mt+d} - F_{jmt}^k \sum_{d=1}^{\tau} \beta^d \Lambda_{jmt d} \pi_{j_d mt+d} \\
&+ u_{ajmt}
\end{aligned} \tag{23}$$

where the $K \times K$ matrix $\Lambda_{amt d}$ is observed (estimable), and is defined recursively

$$\begin{aligned}
\Lambda_{amt d} &= I, & \text{for } d = 1, \\
\Lambda_{amt d} &= \Lambda_{amt, d-1} F_{a_d mt+d}^k, & \text{for } d \geq 2,
\end{aligned}$$

and the unobservable term is now $u_{ajmt} = \tilde{\xi}_{ajmt} + \tilde{e}_{ajmt}^V$, with

$$\begin{aligned}
\tilde{\xi}_{ajmt} &= \xi_{amt} + F_{amt}^k \sum_{d=1}^{\tau} \beta^d \Lambda_{amt d} \xi_{a_d mt+d} \\
&\quad - \xi_{jmt} - F_{jmt}^k \sum_{d=1}^{\tau} \beta^d \Lambda_{jmt d} \xi_{j_d mt+d},
\end{aligned} \tag{24}$$

$$\begin{aligned}
\tilde{e}_{ajmt}^V &= \beta e_{am, t, t+1}^V + F_{amt}^k \sum_{d=1}^{\tau} \beta^{d+1} \Lambda_{amt d} e_{a_d, m, t+d, t+d+1}^V \\
&\quad - \beta e_{jm, t, t+1}^V - F_{jmt}^k \sum_{d=1}^{\tau} \beta^{d+1} \Lambda_{jmt d} e_{j_d, m, t+d, t+d+1}^V.
\end{aligned} \tag{25}$$

As discussed in the context of the applied examples, to identify and estimate the model parameters based on the regression equation (23), it is key to access valid and relevant instrumental variables. We therefore assume the researcher has access to such instruments.

Assumption 2. (*Instrumental Variables*) *There exist instruments z_{mt} such that:*

(i) *For all functions $q(w_{mt})$ with finite expectation, if $E[q(w_{mt}) | z_{mt}] = 0$ almost surely, then $q(w_{mt}) = 0$ almost surely,*

(ii) *$E[\tilde{\xi}_{ajmt} | z_{mt}] = 0$, for all a and j , and*

(iii) *$E[\tilde{e}_{ajmt}^V | z_{mt}] = 0$, for all a and j .*

Assumption 2.(i) is the well-known ‘‘completeness condition,’’ which is the nonparametric analog of the standard rank condition for linear models (Newey and Powell, 2003). Assumptions 2.(ii) and 2.(iii) are usual exclusion restrictions, requiring mean independence between the instruments and both the structural errors ξ and the expectational errors e^V . Recall that Lemma 1 shows that,

when agents have rational expectations, and when instruments belong to agents' information set, expectational errors e^V are mean independent of z_{mt} , which implies Assumption 2.(iii).²¹

We now turn to our main identification results.

Proposition 1. *Suppose (β, F^ϵ) are known and Assumption 2 holds. Assume that, for all pair of actions a and j , the single-action τ -period finite dependence property holds for action J . Assume also that the payoff $\bar{\pi}(J, k, w)$ is known for all (k, w) . Then, given the joint distribution of observables $\Pr(y)$, where $y_{imt} = (a_{imt}, k_{imt}, w_{mt}, z_{mt})$, the flow payoff $\bar{\pi}(a, k, w)$ is identified for all (a, k, w) .*

The primitive of interest $\bar{\pi}$ is nonparametrically identified provided that the flow payoff of action J is known and that single-action finite dependence holds for that same action J . Arcidiacono and Miller (2017) have obtained identification results under similar conditions for nonstationary models with short panels (but with no endogeneity problems).

In general, DDC models require restrictions on the payoff function for identification (Rust, 1994; Magnac and Thesmar, 2002). In the literature, identification typically relies on a restriction of the form $\pi(J, \mathbf{s}) = 0$ for all states \mathbf{s} and an *arbitrary* action J .²² Proposition 1, however, requires such a restriction on a *specific* action J ; this is an unusually strong requirement in the dynamic discrete choice literature. For instance, in a setting with a renewal or terminal action, Proposition 1 requires that the renewal or terminal action's payoffs to be known (i.e., restricted *ex-ante*). The applied examples presented above do not impose this restriction on payoffs. This means that, while sufficient for identification, assuming $\bar{\pi}(J, k, w)$ is known for the *specific* action J is not necessary for identification.

Similarly, while single-action finite dependence is part of the sufficient conditions for identification, it is not necessary. The next proposition shows that, at the cost of imposing parametric restrictions on payoffs, identification can be obtained under more general notion of finite dependence, and it does not require the payoff of a specific action J to be known. There exists therefore a clear trade-off between the identification of parametric and nonparametric models in the present context.

Proposition 2. *Suppose (β, F^ϵ) are known and that τ -period finite dependence holds. Assume a linear-in-parameters flow payoff: $\bar{\pi}(a, k, w) = x(a, k, w)\theta$, where $\theta \in \mathbb{R}^P$, and $x(a, k, w)$ is a known function. Let X_{amt} be a $K \times P$ matrix with elements given by $x(a, k, w_{mt})$, so that*

²¹Strictly speaking, Lemma 1 implies that $E[e_{ajmt}^V|z_{mt}] = 0$ while condition (iii) of Assumption 2 requires $E[\tilde{e}_{ajmt}^V|z_{mt}] = 0$, where \tilde{e}_{ajmt}^V includes e_{ajmt}^V as well as potentially $e_{ajm,t+d}^V$ terms with $d \geq 1$. As discussed in the proof of Lemma 1, an instrument z_{mt} in the time- t information set will also be uncorrelated with future expectational error terms.

²²Such identifying restrictions matter for some, but not all counterfactuals; see Aguirregabiria (2010), Norets and Tang (2014), Arcidiacono and Miller (2017), and Kalouptsi et al. (2017) for more details.

$\bar{\pi}_{amt} = X_{amt}\theta$, and define

$$\begin{aligned} \tilde{X}_{ajmt} \equiv & X_{amt} + F_{amt}^k \sum_{d=1}^{\tau} \beta^d \Lambda_{amtd} X_{a_dmt+d} \\ & - X_{jmt} - F_{jmt}^k \sum_{d=1}^{\tau} \beta^d \Lambda_{jmt d} X_{j_dmt+d}. \end{aligned} \quad (26)$$

Denote the $K \times 1$ vector on the left hand side of (23) by Y_{ajmt} . Stack equation (23) for all feasible combinations of actions $(a, j) \in \mathcal{A}$ to obtain the following equation

$$Y_{mt} = \tilde{X}_{mt}\theta + u_{mt}, \quad (27)$$

where Y_{mt} , \tilde{X}_{mt} , and u_{mt} stack the vectors Y_{ajmt} , \tilde{X}_{ajmt} and u_{ajmt} , respectively. Let Z_{mt} be an $L \times K$ matrix of instrumental variables with $L \geq P$. The parameter θ is identified provided $E[Z_{mt}u_{mt}] = 0$ and $\text{rank}(E[Z_{mt}\tilde{X}_{mt}]) = P$.

While Proposition 2 does not require the payoff parameters of a particular action to be known, the rank condition $\text{rank}(E[Z_{mt}\tilde{X}_{mt}]) = P$ still implicitly limits the number of parameters than can be identified (see, e.g., the dynamic land use change model discussed in Example 2). However, the condition is straightforward to check for a given data set and parametric specification.

We now discuss the two components of the residual, $u_{mt} = \tilde{\xi}_{mt} + \tilde{e}_{mt}^V$, in turn. We start with the expectational error term, \tilde{e}_{mt}^V . As previously mentioned, when the instruments belong to the agent's contemporaneous information set, and when agents have rational expectations, the instruments are uncorrelated with \tilde{e}_{mt}^V . Yet, it is important to note that, while the expectational error terms are mean-zero given \mathcal{I}_{imt} , there is a distinction between the way e^V averages out in cross-sectional dimension M and the way it averages out in the time-series dimension T . For instance, if in a given time period t , a macro shock affects all agents in all markets, then the prediction errors e^V do not average out to zero asymptotically when $M \rightarrow \infty$. For this averaging out to happen as the cross-section becomes large, we do indeed need to make a substantive assumption about the correlation across markets. Such type of assumption may be appropriate for some applications, but it may not be appropriate when, for example, markets are geographically defined and state variables include prices whose movements over time are strongly correlated around the world (e.g., transportable commodities).

In contrast, in a large- T setting, asymptotic convergence of e^V terms does not require a substantive assumption about how markets are correlated. Instead, rational expectations guarantees that expectational error terms are serially uncorrelated (see Lemma 1). Intuitively, when $T \rightarrow \infty$, macro shocks wash out in the limit. We investigate this possibility in our Monte Carlo exercise by exploring the finite sample performance of the ECCP estimator with different sample sizes M and T , both in the presence and absence of macro shocks.

Another concern is that, in some applications, the right-hand side variable \tilde{X}_{mt} from equation (26) may include covariates observed at $(t + 1)$ or subsequent periods. Such variables will generally be correlated with the expectational error component of the residual \tilde{e}_{mt}^V . Recall that the expectational error term at t captures the difference between realized value functions at $(t + 1)$ and their time- t expectations, and note that the realization of the value function at $(t + 1)$ will be correlated with realized covariates at $(t + 1)$. Thus, even if the unobserved shock component of the residual $\tilde{\xi}_{mt}$ is exogenous, expectational error terms can create a mechanical endogeneity problem if future values of covariates are included in the regressors \tilde{X}_{mt} . This problem can be avoided when future values of covariates difference out of the regressors \tilde{X}_{mt} (as is the case in Examples 1 and 2 above), or when appropriate instruments in the contemporaneous information set exist for those future values of covariates. While this endogeneity problem does not arise in some applied examples we consider, it is a central concern in Dickstein and Morales (2018).

We now turn the discussion to the unobservable term $\tilde{\xi}_{mt}$. Recall that ξ_{mt} is a function of observed and unobserved market-level states $\omega_{mt} = (w_{mt}, \eta_{mt})$. If (a) ξ_{mt} depends on η_{mt} , but not on w_{mt} , and (b) the observed and unobserved market-level states, w_{mt} and η_{mt} , evolve independently, then, for any (a, k) , w_{mt} and ξ_{mt} are independent and w_{mt} can be taken as exogenous in the regression (23). In contrast, if ξ_{mt} depends on w_{mt} , or if w_{mt} and η_{mt} do not evolve independently, then w_{mt} and ξ_{mt} are correlated, in which case it may be reasonable to use lagged w_{mt} as instruments, as done by Scott (2013).²³

Another possibility is identification of demand by using supply shifters as instruments (or vice versa). To instrument for PV installation costs and electricity prices, De Groote and Verboven (2018) use supply cost shifters such as the Chinese price index of PV modules on the European market (which are the most important cost component of PV installations), and oil prices.

The linear regression approach can be combined with randomized control trials or quasi-natural experiments to identify the model parameters. For instance, Diamond et al. (2018) extend the linear IV approach to a set of differences-in-differences regressions that explore variation in the assignment of rent control due to a 1994 ballot initiative in San Francisco. Combining quasi-experiments with structural dynamic models in the present context seems to be a promising venue for future applied research.

We have assumed rational expectations throughout the paper so far; the following remarks discuss the possibility of implementing the ECCP approach with other notions of agents' beliefs.

²³While lagged values of w_{mt} are in the contemporaneous information set, and therefore uncorrelated with \tilde{e}_{mt}^V given rational expectations, it places strong restrictions on the evolution of ξ to assume that lagged values of w_{mt} are uncorrelated with $\tilde{\xi}_{mt}$. For instance, if ξ_{mt} is serially correlated, and the reason for considering instruments is to deal with potential correlation between ξ_{mt} and w_{mt} , then $\tilde{\xi}_{mt}$ and w_{mt-1} will typically be correlated (which means further lags of w_{mt} may be needed to obtain valid instruments). However, if w_{mt} is serially correlated but $\tilde{\xi}_{mt}$ is not, as when $\tilde{\xi}_{mt} = \xi_{mt}$ captures i.i.d. measurement error in w , then lagged values of w may serve as valid instruments. Finally, note that if $\tilde{\xi}_{mt}$ involves an additive market-level fixed-effect, it is possible to differenced-out the fixed-effect as usually done in panel data settings.

Remark 1. (*Biased Beliefs*). It is clear that Propositions 1 and 2 do not require rational expectations. Identification of model primitives can therefore be obtained under biased beliefs. While under the rational expectation assumption, e^V is mean zero given current and all past states (Lemma 1), under biased beliefs the true conditional expectation of e^V may not be zero. In such contexts, one may still find plausible instruments that do not correlate with systematic errors e^V .²⁴ An alternative is to construct other moment restrictions based on (23) or (26) that do not require rational expectations. Indeed, Diamond et al. (2018) allow for biased beliefs: they only need to assume that difference in expectations between treatment and control households in the same year are zero on average.

Remark 2. (*Perfect Foresight*). We derive our estimators based on the assumption that agents have rational expectations, but the same estimators can also be derived from the assumption that agents have perfect foresight about the market state ω_{mt} . Recall that our regression equations have residuals for the form $u_{mt} = \tilde{\xi}_{mt} + \tilde{e}_{mt}^V$, where $\tilde{\xi}_{mt}$ involves unobserved shocks to payoffs, and the expectational error term \tilde{e}_{mt}^V comes from uncertainty in the market state ω_{mt} . When agents have perfect foresight about ω_{mt} , the unobservable component of the payoff becomes the entire residual, i.e., $u_{mt} = \tilde{\xi}_{mt}$. This reinterpretation of the residual is of no consequence for the method-of-moments estimators we propose, for the moments depend only on the value of the entire residual u_{mt} .

Note however that in a setting where the potential instruments in Z_{mt} include future values of observables, there will generally be a mechanical endogeneity problem coming from the expectational error terms (as discussed above, expectational error terms are generally correlated with realized values of future variables), but a misspecified perfect foresight model would ignore this endogeneity problem and result in biased parameter estimates.

5 Two-stage Estimation

We follow the tradition of Hotz and Miller (1993) and estimate conditional choice probabilities and transition probabilities in the first step, and estimate the model parameters in the second step. To simplify exposition, we assume a large number of individuals N and markets M , but we hold the number of time periods T fixed.²⁵

Although it is possible to nonparametrically estimate payoff functions $\bar{\pi}$ (following Proposition 1), here we consider a parametric model $\bar{\pi}(a, k, w; \theta_0)$, where $\theta_0 \in \Theta \subset \mathbb{R}^P$ is the parameter of interest. Parametric models estimated using panel data involving large number of cross-sectional observations and small number of time periods is common in applied work.²⁶

²⁴For instance, assumptions on the learning process about the evolution of states may suggest possible sources of instruments. We leave this possibility for future research.

²⁵Asymptotic results can be extended to large T by imposing stationarity and ergodicity assumptions.

²⁶Nonparametric payoff functions $\bar{\pi}$ can be estimated in the second step using estimators proposed by Newey and

5.1 First Stage

In the first stage, we estimate $p_{amt}(\cdot)$ and $F_{amt}^k(\cdot)$ in all markets m and all available time periods t . Denote the estimators for the CCPs and transition probabilities by \widehat{p}_{amt} and \widehat{F}_{amt}^k . Because K is finite, we consider the frequency estimators:²⁷

$$\begin{aligned}\widehat{p}_{amt}(k) &= \frac{\sum_{i=1}^N 1\{a_{imt} = a, k_{imt} = k\}}{\sum_{i=1}^N 1\{k_{imt} = k\}}, \\ \widehat{F}_{amt}^k(k'|k) &= \frac{\sum_{i=1}^N 1\{a_{imt} = a, k_{imt} = k, k_{imt+1} = k'\}}{\sum_{i=1}^N 1\{a_{imt} = a, k_{imt} = k\}}.\end{aligned}$$

We impose the following condition for each market m and each time period t .

Condition 1. The observations $\{a_{imt}, k_{imt} : i = 1, \dots, N\}$ are i.i.d. conditional on the market level state ω_{mt} .

Condition 1 formalizes the idea that ω_{mt} is a common shock affecting all agents i in market m at time period t . As shown in Andrews (2005), this assumption is valid when the sample of individuals are drawn randomly from the population. (Sources of spatial dependence among agents within markets other than the common shock ω can be accommodated in our framework, but is beyond the scope of the paper.)

The probability limits of \widehat{p}_{amt} and \widehat{F}_{amt}^k can be determined following the Law of Large Numbers for exchangeable random variables (see Hall and Heyde (1980)). In addition, the result can be strengthened to a law of iterated logarithm, which is an important input to derive the asymptotic results for the second step. To simplify notation, stack the vectors $(\widehat{p}_{amt}, \widehat{F}_{amt}^k)$ and (p_{amt}, F_{amt}^k) for all actions and states, and denote them, respectively, by $\widehat{\delta}_{mt}$ and δ_{mt} . Note that δ_{mt} is itself random because (p_{amt}, F_{amt}^k) depends on the realization of ω_{mt} . In what follows, we use the Euclidean norm $\|\cdot\|$.²⁸

Lemma 2. *Suppose Condition 1 holds. Then,*

$$\widehat{\delta}_{mt} \rightarrow \delta_{mt} \text{ a.s.}, \tag{28}$$

Powell (2003) or by Escanciano et al. (2016). While the former involves ill-posed inverse issues, the latter combines standard kernel estimation with the computation of a matrix eigenvector problem in a way that avoids the ill-posed inverse problem. Aradillas-Lopez (2015) investigates the semiparametric efficiency properties of estimators of parametric models with rational expectations in which agents' beliefs are treated nonparametrically as "generated regressors." This approach may be useful in our context when k is a continuous variable.

²⁷Results can be adapted to more complicated estimators of CCP and transitions. For instance, we can incorporate time-invariant observables that affect individuals' choices, as well as unobservable heterogeneity as in Kasahara and Shimotsu (2009), and Arcidiacono and Miller (2011).

²⁸Although we do not exploit the asymptotic distribution of $\widehat{\delta}_{mt}$ in this paper, note that it can be obtained following the arguments in Andrews (2005). Specifically, under the regularity conditions stated in Andrews (2005), $\sqrt{N}(\widehat{\delta}_{mt} - \delta_{mt})$ converges in distribution to a mixture of normal distributions that depend on the common shock ω_{mt} .

as $N \rightarrow \infty$. Moreover,

$$\|\widehat{\delta}_{mt} - \delta_{mt}\| = O_{a.s.} \left(\sqrt{\frac{\log \log N}{N}} \right). \quad (29)$$

5.2 Second Stage

Recall that, for each combination of a and j , the unobservable u_{ajmt} is the $K \times 1$ vector satisfying (23). Define the vector $u_{mt}(\theta, \delta_{mt})$ that stacks u_{ajmt} for all feasible combinations of (a, j) , and define the function $g_{mt}(\theta) = h(z_{mt}) u_{mt}(\theta, \delta_{mt})$, where $h(z_{mt})$ is a conformable function of the instrumental variables. The unconditional moment restriction (implied by Assumption 2) is then

$$E[h(z_{mt}) u_{mt}(\theta_0, \delta_{mt})] = 0. \quad (30)$$

Define the function $g(\theta) \equiv E[g_{mt}(\theta)]$; the GMM population criterion function is given by

$$Q(\theta) = g(\theta)' \mathbf{W} g(\theta), \quad (31)$$

where \mathbf{W} is a (non-stochastic) positive-definite weighting matrix. By the identification results, θ_0 is the unique minimizer of $Q(\theta)$.

Next, consider the sample analogue of (31). Define the functions

$$\widehat{g}_{mt}(\theta) = h(z_{mt}) u_{mt}(\theta, \widehat{\delta}_{mt}), \text{ and} \quad (32)$$

$$\widehat{g}_M(\theta) = \frac{1}{M(T-\tau)} \sum_{m=1, t=1}^{M, (T-\tau)} \widehat{g}_{mt}(\theta), \quad (33)$$

where τ reflects the τ -finite dependence assumption (see Definition 2). The GMM criterion function is then given by

$$\widehat{Q}_M(\theta) = \widehat{g}_M(\theta)' \mathbf{W}_M \widehat{g}_M(\theta), \quad (34)$$

where \mathbf{W}_M is a positive-definite weighting matrix that may depend on data. The estimator $\widehat{\theta}_M$ minimizes $\widehat{Q}_M(\theta)$ over Θ .

The next set of conditions suffices for consistency of $\widehat{\theta}_M$.

Condition 2. Conditions for consistency:

- (i) The vector (w_{mt}, z_{mt}) is i.i.d. across markets m .
- (ii) $\mathbf{W}_M \xrightarrow{p} \mathbf{W}$ as $M \rightarrow \infty$.
- (iii) Θ is compact.
- (iv) θ_0 uniquely minimizes $Q(\theta)$ over Θ .
- (v) $\bar{\pi}_{amt}(\theta)$ is continuous at each $\theta \in \Theta$ for all $a \in \mathcal{A}$.
- (vi) $E[\sup_{\theta \in \Theta} \|h(z_{mt}) u_{mt}(\theta, \delta_{mt})\|] < \infty$.
- (vii) $E[\sup_{\theta \in \Theta} \|h(z_{mt}) \nabla_{\delta} u_{mt}(\theta, \delta_{mt})\|^2] \leq B < \infty$, where B is a finite constant.

Condition 2 establishes standard regularity conditions that guarantee the problem is well-behaved. The assumption that market-level variables are independent across markets simplifies the derivation of the asymptotic results (Condition 2(i)), but results can be extended to allow for spatial dependence across markets (Conley, 1999; Andrews, 2005; Kuersteiner and Prucha, 2013). Condition 2(ii)–(v) are standard. Condition 2(vii) is used in conjunction with equation (28) in Lemma 2 of Section 5.1, to guarantee uniform convergence in probability of the criterion function to its population version. The term $\nabla_{\delta} u_{mt}$ is the derivative of the vector u_{mt} with respect to δ , which in turn, depends on the derivative of ψ_a with respect to the CCPs (see equation (23)). As shown in Kalouptsi et al. (2017), ψ_a is indeed a differentiable function of $p(k, \omega)$, provided $p(k, \omega)$ lies strictly between zero and one (which is satisfied when payoffs are bounded and ε_{imt} has full support on \mathbb{R}^{A+1}).

Proposition 3. *Under Conditions 1 and 2, $\widehat{\theta}_M \xrightarrow{p} \theta_0$ as $(M, N) \rightarrow \infty$.*

To obtain the asymptotic distribution of $\widehat{\theta}_M$, we impose the following:

Condition 3. Conditions for asymptotic distribution:

- (i) $\theta_0 \in \text{interior}(\Theta)$.
- (ii) $\bar{\pi}_{amt}(\theta)$ is continuously differentiable in a neighborhood \mathcal{N} of θ_0 with probability approaching one for all $a \in \mathcal{A}$.
- (iii) $E[\|h(z_{mt}) u_{mt}(\theta_0, \delta_{mt})\|^2] < \infty$, and $E[\sup_{\theta \in \mathcal{N}} \|h(z_{mt}) \nabla_{\theta} u_{mt}(\theta, \delta_{mt})\|] < \infty$.
- (iv) $\mathbf{G}'\mathbf{W}\mathbf{G}$ is nonsingular for $\mathbf{G} = E\left[\frac{1}{T-\tau} \sum_{t=1}^{T-\tau} h(z_{mt}) \nabla_{\theta} u_{mt}(\theta_0, \delta_{mt})\right]$.

Similar to Condition 2, Condition 3 imposes standard regularity conditions to make the problem well-behaved. The next proposition follows.

Proposition 4. *Suppose Conditions 1, 2 and 3 hold. Assume $(M \log \log N)/N \rightarrow 0$, as $(M, N) \rightarrow \infty$. Then,*

$$\sqrt{M} \left(\widehat{\theta}_M - \theta_0 \right) \xrightarrow{p} N(0, \mathbf{V})$$

where

$$\mathbf{V} = (\mathbf{G}'\mathbf{W}\mathbf{G})^{-1} \mathbf{G}'\mathbf{W}\Sigma\mathbf{W}\mathbf{G} (\mathbf{G}'\mathbf{W}\mathbf{G})^{-1},$$

and

$$\Sigma = E \left[\left(\frac{1}{T-\tau} \sum_{t=1}^{T-\tau} g_{mt}(\theta_0) \right) \left(\frac{1}{T-\tau} \sum_{t=1}^{T-\tau} g_{mt}(\theta_0) \right)' \right].$$

The asymptotic distribution of $\widehat{\theta}_M$ in Proposition 4 is the same as the distribution of an unfeasible estimator in which δ_{mt} is observed instead of estimated in the first step. The first step estimators $\widehat{\delta}_{mt}$, for all markets and time periods, do not affect the asymptotic variance of $\widehat{\theta}_M$ when the *within* markets observations N is sufficiently large compared to the number of markets M . The rate at which N must increase to eliminate the influence of the first step on the variance of

the second step is $\frac{M \log \log N}{N} \rightarrow 0$ as $(M, N) \rightarrow \infty$. This rate is a direct result of the Law of Iterated Logarithm obtained in Lemma 2.

Consistent estimators for \mathbf{V} can be obtained using standard arguments and are omitted here (see, e.g., Theorem 4.5 in Newey and McFadden (1994)).

6 Monte Carlo

In this section, we present a Monte Carlo experiment to illustrate the performance of the ECCP estimator. We consider the dynamic demand model for a durable good discussed in details in Example 1.

As previously mentioned, the choice set is $\mathcal{A} = \{b, nb\}$, where $a = b$ if the consumer buys the good, and $a = nb$ if she does not buy it. The individual-level state k_{imt} reflects whether she already owns the product or not at the beginning of time period t . We consider two market-level state variables: observed price w_{mt} and unobservable quality ξ_{mt} .²⁹ We also consider an observed cost shifter z_{mt} (that will play the role of an instrumental variable). The price is a function of z_{mt} and ξ_{mt} , determined as follows:

$$w_{mt} = \gamma_0 + \gamma_1 z_{mt} + \gamma_2 \xi_{mt} + \varepsilon_{mt}^w, \quad (35)$$

where ε_{mt}^w is a mean-zero normally distributed i.i.d. price shock with variance σ_w^2 . Note that γ_1 represents how variation in the observed cost shifter z_{mt} passes through to prices. ξ_{mt} is included in the price equation to capture the idea that demand shocks may influence the price.

The supply shifter z_{mt} and unobserved quality ξ_{mt} follow independent AR(1) processes:

$$\xi_{mt+1} = \rho_1 + \rho_2 \xi_{mt} + \varepsilon_{mt}^\xi,$$

$$z_{mt+1} = \rho_3 + \rho_4 z_{mt} + \varepsilon_{mt}^z,$$

where ε_{mt}^ξ and ε_{mt}^z are normally distributed zero-mean i.i.d. shocks with variances σ_ξ^2 and σ_z^2 , respectively. Recall that, because ξ_{mt} has mean zero by assumption, we take $\rho_1 = 0$.

We consider two settings in our simulations, depending on whether the unobservable state ξ is present in the data generating process or not. When there is no unobserved states, we set $\xi_{mt} = 0$ for all m and t (or, equivalently, we set $\sigma_\xi^2 = 0$).

As many applications may feature aggregate shocks, we also consider two scenarios: with and without macro shocks. In the scenario with aggregated shocks, we incorporate them into the term ε_{mt}^z . Specifically, we simulate $\varepsilon_{mt}^z = \varepsilon_{m,t,1}^z + \varepsilon_{t,2}^z$, where the (mean-zero) macro shock $\varepsilon_{t,2}^z$ accounts for a fraction λ_z of the variance in ε_{mt}^z . (Recall that macro shocks wash out in the limit when

²⁹Recall that, formally, we can define η_{mt} as the state variable, while the function $\xi(a, k, \omega_{mt})$ equals η_{mt} when $a = b$, and equals zero when $a = nb$. I.e., the unobserved quality enters into the utility of purchasing the good and not into the utility of not purchasing (whether the good is already owned or not). To simplify notation, we take ξ_{mt} as the state variable.

the number of time periods goes to infinity, but they may not wash out in the limit when the asymptotics is in the number of markets.)

We consider several possible sample sizes with different combinations of M and T . For each sample structure, we generate 5,000 Monte Carlo replications. For the first sample period, we generate the value of the state variables from their steady-state distributions. The supports of the market-level state variables are discretized to take integer values. We solve the individual dynamic optimization problem by value function iteration on the discretized state space.³⁰ The parameters of the data generating process in our Monte Carlo study is summarized in Table 1.

Table 1: Parameters of the Monte Carlo Data Generating Process

Payoff Parameters:	θ_0	1	$\xi \sim \text{Normal AR}(1)$	ρ_1	0
	θ_1	-.1		ρ_2	.2
				σ_ξ^2	0 or 16
Prob. of Product Failure:	ϕ	.1	$z \sim \text{Normal AR}(1)$	ρ_3	0
Discount Factor:	β	.95		ρ_4	.7
				σ_z^2	25
Process for price w_{mt} :	γ_0	40	Aggregate Shocks	λ_z	0 or .7
	γ_1	1			
	γ_2	1			
	σ_w^2	4			

The main parameters of interest are the payoff parameters $\theta = (\theta_0, \theta_1)$, which are estimated based on the regression equation (20). We estimate θ using two ECCP estimators: the first estimator is based on the Ordinary Least Squares (OLS) estimator, while the second one is based on instrumental variables – specifically, the Two-Stage Least Squares (2SLS) estimator.

We also estimate θ using a standard CCP estimator similar to Hotz and Miller (1993). This procedure relies on a full specification of what the state variables are and how they evolve. To implement this procedure, we assume that the price w_{mt} is the only market-level state variable and model its evolution as a first-order Markov process. So, when the unobserved quality is a relevant state (i.e., when $\sigma_\xi^2 > 0$), this strategy is based on a mis-specified model. See Appendix A.4 for details.

Table 2 presents results for a model with no unobservable demand shocks (i.e., $\sigma_\xi^2 = 0$). For each model parameter, each sample structure, and each estimator, we report the average estimate, the relative mean bias (as a percentage of the true parameter), the standard deviation, and the root-mean squared error (RMSE) of the estimator. On the left panel, we present results for the scenario with no macro shocks (i.e., $\lambda_z = 0$), and on the right panel, we present the estimated

³⁰We simulate conditional choice probabilities for each market-period and assume these CCPs are observed by the econometrician. Thus, we abstract away from any first-stage sampling uncertainty in the estimation of choice probabilities and effectively assume a (sufficiently) large number of agents N within each market-year.

results with aggregated shocks ($\lambda_z \neq 0$).

Table 2: Monte Carlo Experiments without Unobserved Demand Shock

			$\sigma_\xi^2 = 0$ True Parameters: $\theta_0 = 1, \theta_1 = -.1$						
			$\lambda_z = 0$			$\lambda_z = .7$			
Estimator	T		40	160	10	40	160	10	160
	M		40	10	160	40	10	160	160
ECCP: OLS	θ_0	Mean Est.	1.01	1.01	1.01	0.92	0.99	0.75	0.99
		Rel. Bias	0.64%	0.74%	0.81%	-8.05%	-1.49%	-24.6%	-1.38%
		SD	0.04	0.04	0.04	0.17	0.09	0.29	0.09
		RMSE	0.04	0.04	0.04	0.19	0.10	0.38	0.09
	θ_1	Mean Est.	-0.10	-0.10	-0.10	-0.10	-0.10	-0.09	-0.10
		Rel. Bias	0.16%	0.18%	0.20%	-1.98%	-0.38%	-6.20%	-0.33%
		SD	1.01e-3	1.04e-3	9.97e-4	4.05e-3	2.22e-3	6.78e-3	2.05e-3
		RMSE	1.02e-3	1.06e-3	1.02e-3	4.50e-3	2.26e-3	9.19e-3	2.07e-3
ECCP: IV (2SLS)	θ_0	Mean Est.	1.01	1.01	1.01	0.91	0.98	0.72	0.98
		Rel. Bias	0.70%	0.79%	0.87%	-9.0%	-1.66%	-28.2%	-1.57%
		SD	0.04	0.04	0.04	0.19	0.10	0.33	0.09
		RMSE	0.04	0.05	0.04	0.21	0.10	0.43	0.10
	θ_1	Mean Est.	-0.10	-0.10	-0.10	-0.10	-0.10	-0.09	-0.10
		Rel. Bias	0.18%	0.19%	0.22%	-2.22%	-0.42%	-7.08%	-0.38%
		SD	1.06e-3	1.08e-3	1.03e-3	4.46e-3	2.40e-3	7.67e-3	2.23e-3
		RMSE	1.07e-3	1.10e-3	1.06e-3	4.98e-3	2.44e-3	0.01	2.26e-3
Standard CCP	θ_0	Mean Est.	1.00	1.00	1.00	1.00	1.00	1.00	1.00
		Rel. Bias	-0.29%	-0.27%	-0.11%	-0.49%	-0.31%	-0.02%	-0.19%
		SD	9.22e-3	8.45e-3	0.01	0.04	0.02	0.07	0.02
		RMSE	9.66e-3	8.87e-3	0.01	0.04	0.02	0.07	0.02
	θ_1	Mean Est.	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10
		Rel. Bias	-0.48%	-0.51%	-0.30%	-0.70%	-0.60%	-0.64%	-0.25%
		SD	4.40e-4	3.88e-4	6.41e-4	7.96e-4	5.27e-4	1.33e-3	5.18e-4
		RMSE	6.48e-4	6.42e-4	7.08e-4	1.06e-3	7.95e-4	1.48e-3	5.73e-4

Notes: 5000 Monte Carlo replications for each sample structure. SD is the standard deviation of the estimators across replications. RMSE is root-mean squared error. Relative Bias is bias as percentage of the true parameter.

In the absence of macro shocks, all estimation strategies appear to be consistent, as expected. Not surprisingly, the standard CCP approach exhibits smaller standard deviations and RMSE when compared to the ECCP OLS and IV estimators.

Table 2 also illustrates that there can be a difference between the number of markets and time periods for the asymptotic properties of ECCP estimators in the scenario with macro shocks ($\lambda_z = .7$). We have several columns with the same sample size $M \times T = 1600$, and within these columns, there is non-trivial bias in the OLS and IV estimates for short panels when the aggregate shocks are present. However, the bias is reduced when the time dimension increases, and for long panels, the ECCP estimators have little or no bias.³¹

Table 3 presents the results for the main specification, where the unobserved demand shock ξ_{mt} is present. Because price w_{mt} and the unobserved quality ξ_{mt} are correlated by construction (since $\gamma_2 \neq 0$), we expect OLS to be biased. Indeed, the OLS estimator is now highly biased for both the intercept θ_0 and price coefficient θ_1 .

The supply shifter z_{mt} provides a valid instrument given that it correlates with prices w_{mt} and is independent of ξ_{mt} . When using the cost shifter as an instrument for price, we see that the ECCP IV estimator has little-to-no bias, either in the absence of aggregate shocks or when T is large (aggregate shocks still pose problems for short panels, as expected).

The standard CCP estimator is now severely biased, for it treats the market state space as including only the observable price w_{mt} while the unobservable demand shock ξ_{mt} also plays a role. The mis-specification is important regardless of the presence or absence of macro shocks. In particular, the relative bias of the standard CCP estimator is larger than the bias of the ECCP IV estimator in the presence of aggregated shocks in short panels.

6.1 Counterfactuals

So far, we have considered only the estimation of the parameters of agents' utility function. Typically, applied researchers are also interested in the outcomes of policy simulations or counterfactuals. In this section, we will consider how the biases in the parameter estimates pass through to biases in counterfactuals, but first, we must consider the question of how to do counterfactuals within the ECCP framework.

Much of the ECCP approach's appeal comes from the fact that it takes seriously the possibility that the econometrician might be facing important measurement issues; e.g., some market-level state variables might not be observed, and/or it might be difficult to specify how they evolve. However, when doing counterfactuals, researchers typically solve for a new equilibrium of the model, which normally involves fully specifying all the relevant state variables and how they evolve. Thus, *prima facie*, ECCP estimation seems to be at odds with doing counterfactual simulations.

A counterfactual is a function of the model parameters, and sometimes that function does not depend (or depends only minimally) on the presence of unobservable variables or on the precise

³¹In our simulations, aggregate shocks only affect the evolution of the observable variables (z_{mt} directly, and w_{mt} indirectly). However, these aggregate shocks in the observables lead to aggregate shocks in the expectational error terms. In turn, aggregate shocks in the expectational error terms implies that our regression equation's residual features aggregate shocks.

Table 3: Monte Carlo Experiments with Unobserved Demand Shock

		$\sigma_{\xi}^2 = 16$	True Parameters: $\theta_0 = 1, \theta_1 = -.1$						
		$\lambda_z = 0$			$\lambda_z = .7$				
Estimator	T	40	160	10	40	160	10	160	
	M	40	10	160	40	10	160	160	
ECCP: OLS	θ_0	Mean Est.	-8.62	-8.63	-8.61	-9.54	-8.81	-11.60	-8.85
		Rel. Bias	-962%	-963%	-961%	-1050%	-981%	-1260%	-985%
		SD	0.59	0.58	0.57	1.79	1.03	2.77	0.89
		RMSE	9.64	9.64	9.62	10.70	9.86	12.90	9.89
	θ_1	Mean Est.	0.14	0.14	0.14	0.16	0.15	0.22	0.15
		Rel. Bias	-240%	-241%	-240%	-264%	-245%	-315%	-246%
		SD	0.01	0.01	0.01	0.04	0.02	0.06	0.02
		RMSE	0.24	0.24	0.24	0.27	0.25	0.32	0.25
ECCP: IV (2SLS)	θ_0	Mean Est.	1.02	1.00	1.02	0.97	1.02	0.87	0.99
		Rel. Bias	1.62%	-0.20%	1.82%	-3.49%	1.97%	-12.8%	-0.76%
		SD	0.77	0.77	0.75	0.83	0.78	0.90	0.20
		RMSE	0.77	0.77	0.75	0.83	0.78	0.91	0.20
	θ_1	Mean Est.	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10	-0.10
		Rel. Bias	0.50%	-0.05%	0.42%	-1.00%	0.44%	-3.25%	-0.19%
		SD	0.02	0.02	0.02	0.02	0.02	0.02	4.96e-3
		RMSE	0.02	0.02	0.02	0.02	0.02	0.02	4.96e-3
Standard CCP	θ_0	Mean Est.	0.26	0.26	0.25	0.26	0.26	0.26	0.17
		Rel. Bias	-74.4%	-74.4%	-74.7%	-74.3%	-74.3%	-73.5%	-83.3%
		SD	0.03	0.03	0.03	0.04	0.03	0.05	0.02
		RMSE	0.75	0.75	0.75	0.74	0.74	0.74	0.83
	θ_1	Mean Est.	-0.01	-0.01	-0.01	-0.01	-0.01	-0.01	2.12e-3
		Rel. Bias	-88.1%	-87.8%	-89.2%	-88.4%	-87.8%	-89.9%	-102%
		SD	5.38e-3	5.12e-3	5.74e-3	6.33e-3	5.41e-3	9.24e-3	3.66e-3
		RMSE	0.09	0.09	0.09	0.09	0.09	0.09	0.10

Notes: 5000 Monte Carlo replications for each sample structure. SD is the standard deviation of the estimators across replications. RMSE is root-mean squared error. Relative Bias is bias as percentage of the true parameter.

specification of how state variables evolve. Therefore, the modeling issues that motivate the ECCP approach need not undermine the use of parameter estimates for counterfactual analysis. De Groote and Verboven (2018) provide a clear example. They use an ECCP estimator to estimate the rate of time discounting of Belgian households in deciding whether to install solar photovoltaic

systems (the ECCP estimator allows them to flexibly include demand shocks and avoid specifying a process for how government policy evolved). They find that households' estimated discount rate is considerably lower than the interest rate that the Belgian government can borrow at. As they argue, this disparity means that it would be more cost effective for the government to support solar PV installations with up-front payments, rather than the ongoing payments that the government actually used. This conclusion follows intuitively from the disparity in discount rates and plausibly is not affected in an important way by how government policy and unobservable states evolve. The conclusion, however, may be highly sensitive to biases in the *estimation* of the discount factor. In other words, the estimation of a mis-specified model may crucially affect policy recommendations.

In what follows, we show that some counterfactuals – specifically, long-run demand elasticities from our durable good demand model – are robust to the omission of unobservable state variables that are present in the data generating process. Or more specifically, long-run demand elasticities are well approximated by a model in which we set ξ at its long-run mean (i.e., $\xi = 0$). Furthermore, we show that the biases that result from leaving the unobservable shocks out of the counterfactual simulations can be smaller than the biases that result from using an estimation approach that is not robust to their presence.

We perform both a *real* and *feasible* counterfactuals for the durable good demand model. Our *real* counterfactuals take the parameter estimates from various estimations above and plug them into a counterfactual that uses the true data generating process (notably including the true law of motion for the unobservable demand shock ξ_{mt}). That is, the real counterfactuals rely on our understanding of an unobservable that an econometrician who was not simulating the data would not have access to. Our *feasible* counterfactuals, in contrast, simulate a simple model that an econometrician could easily implement: a model that sets $\xi = 0$.³²

The counterfactual we consider is an increase in the mean price level (formally, we increase γ_0 by .01), and we calculate the long-run change in the demand level. That is, we calculate the unconditional probability of purchase $Pr(a = b)$ in the steady state after solving the consumer's dynamic problem. We present this counterfactual in the form of a *long-run demand elasticity*, i.e., the ratio of the percentage change in the probability of purchase to the percentage change in the long-run price.

Table 4 shows the counterfactuals from the ECCP (OLS and IV) and standard CCP estimators based on the parameter estimates from the above simulations with $M = 160$ and $T = 160$. A first observation is that the real and feasible counterfactuals at the true parameters differ by a factor of about 10%. Second, consistent with the biases in the underlying parameter estimates, we find that the ECCP IV estimates yield very little bias in the counterfactuals relative to the true values while the other estimators result in substantially biased counterfactuals. Furthermore, the biases

³²To solve the feasible counterfactual, we need to specify how the exogenous state variable w_{mt} evolves. We consider the residual from the pricing equation (35) as the econometrician can measure it. I.e., $w = \gamma_0 + \gamma_1 z + \nu$, where $\nu = \gamma_2 \xi + \varepsilon^w$. So, we calculate the true evolution of ν given the underlying processes and assume the econometrician is able to estimate it.

Table 4: Sample size, structure and bias

True value			ECCP		Standard CCP
			OLS	IV	
Real LRE	-1.106	Mean Estimate	60.15	-1.104	0.01471
		Relative Bias	-5540%	-0.1561%	-101.3%
		SD	16.62	0.04227	0.02545
		RMSE	63.48	0.04231	1.121
Feasible LRE	-1.022	Mean Estimate	-1.187e4	-1.064	0.03888
		Relative Bias	1.162e6%	4.114%	-103.8%
		SD	1.382e6	0.1184	0.06774
		RMSE	1.382e6	0.1256	1.063

Notes: 5000 replications with sample structure $M = T = 160$. SD is the standard deviation across replications. RMSE is root-mean squared error.

Relative Bias is bias as percentage of the true parameter.

in the long-run elasticities from the OLS and standard CCP estimators (whether we consider the real or feasible versions) are larger than the gap between the real and feasible estimators.

Evidently, counterfactuals are not always robust to setting ξ at its unconditional mean. The broader point we make in this section is that robustness to the presence of unobserved shocks can be assessed through a procedure similar to what we do here. That is, when researchers are concerned about the presence of unobservables, they might adopt a robust estimation approach that delivers consistent estimates of important parameters despite the unobservables. Then, when it comes to counterfactual simulations, they can perform the simulation in several ways to assess whether and how the results of interest might be sensitive to the presence of unobservables and how they evolve.

7 Conclusion

In this paper we propose (and provide a comprehensive econometric treatment of) a class of linear instrumental variables estimators for structural dynamic discrete choice models: the ECCP estimators. This class of estimators shares many of the advantages of the continuous-choice Euler equation approach originally developed by Hall (1978), Hansen and Sargent (1980), Hansen and Sargent (1982), and Hansen and Singleton (1982).

We provide constructive identification results that lead naturally to estimators, we establish the consistency and asymptotic normality of the estimators, and we provide evidence that they perform well in finite-samples based on a Monte Carlo study of a dynamic demand for durable goods.

References

- Aguirregabiria, V. (2010). Another look at the identification of dynamic discrete decision processes: An application to retirement behavior. *Journal of Business & Economic Statistics* 28(2), 201–218.
- Aguirregabiria, V. and A. Magesan (2013). Euler equations for the estimation of dynamic discrete choice structural models. In *Structural Econometric Models*, pp. 3–44. Emerald Group Publishing Limited.
- Aguirregabiria, V. and P. Mira (2002). Swapping the nested fixed point algorithm: A class of estimators for discrete markov decision models. *Econometrica* 70(4), 1519–1543.
- Altug, S. and R. A. Miller (1998). The effect of work experience on female wages and labour supply. *The Review of Economic Studies* 65(1), 45–85.
- Andrews, D. W. (2005). Cross-section regression with common shocks. *Econometrica* 73(5), 1551–1585.
- Aradillas-Lopez, A. (2015). Semiparametric efficiency in models with rational expectations.
- Arcidiacono, P. and P. B. Ellickson (2011). Practical methods for estimation of dynamic discrete choice models. *Annual Review of Economics* 3(1), 363–394.
- Arcidiacono, P. and R. A. Miller (2011). Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica* 79(6), 1823–1867.
- Arcidiacono, P. and R. A. Miller (2017). Identifying dynamic discrete choice models off short panels. Working paper.
- Bajari, P., C. L. Benkard, and J. Levin (2007). Estimating dynamic models of imperfect competition. *Econometrica* 75(5), 1331–1370.
- Bajari, P., C. S. Chu, D. Nekipelov, and M. Park (2016). Identification and semiparametric estimation of a finite horizon dynamic discrete choice model with a terminating action. *Quantitative Marketing and Economics* 14(4), 271–323.
- Berry, S. and G. Compiani (2017). An instrumental variables approach to dynamic models.
- Bishop, K. C. (2008). A dynamic model of location choice and hedonic valuation. *Unpublished, Washington University in St. Louis*.
- Blevins, J. R. (2014). Nonparametric identification of dynamic decision processes with discrete and continuous choices. 5(3), 531–554.

- Blevins, J. R., A. Khwaja, and N. Yang (2017). Firm expansion, size spillovers, and market dominance in retail chain dynamics. *Management Science*.
- Chiong, K. X., A. Galichon, and M. Shum (2016). Duality in dynamic discrete-choice models. *Quantitative Economics* 7(1), 83–115.
- Conley, T. G. (1999). GMM estimation with cross sectional dependence. *Journal of Econometrics* 92(1), 1 – 45.
- De Groote, O. and F. Verboven (2018). Subsidies and myopia in technology adoption: Evidence from solar photovoltaic systems.
- Diamond, R., T. McQuade, and F. Qian (2018). The effects of rent control expansion on tenants, landlords, and inequality: Evidence from san francisco. Technical report, National Bureau of Economic Research.
- Dickstein, M. J. and E. Morales (2018). What do exporters know? *The Quarterly Journal of Economics* (forthcoming).
- Escanciano, J. C., S. Hoderlein, A. Lewbel, O. Linton, and S. Srisuma (2016). Nonparametric euler equation identification and estimation.
- Gowrisankaran, G. and M. Rysman (2012). Dynamics of consumer demand for new durable goods. *Journal of political Economy* 120(6), 1173–1219.
- Hall, P. and C. C. Heyde (1980). *Martingale limit theory and its application*. Academic press.
- Hall, R. E. (1978, December). Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. *Journal of Political Economy* 86(6), 971–87.
- Hansen, L. P. and T. J. Sargent (1980). Formulating and estimating dynamic linear rational expectations models. *Journal of Economic Dynamics and control* 2, 7–46.
- Hansen, L. P. and T. J. Sargent (1982). Instrumental variables procedures for estimating linear rational expectations models. *Journal of Monetary Economics* 9(3), 263–296.
- Hansen, L. P. and K. J. Singleton (1982). Generalized instrumental variables estimation of nonlinear rational expectations models. *Econometrica: Journal of the Econometric Society*, 1269–1286.
- Hendel, I. and A. Nevo (2006, November). Measuring the implications of sales and consumer inventory behavior. *Econometrica* 74(6), 1637–1673.
- Horvath, L. and B. S. Yandell (1988). Asymptotics of conditional empirical processes. *Journal of Multivariate Analysis* 26(2), 184–206.

- Hotz, V. J. and R. A. Miller (1993, July). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies* 60(3), 497–529.
- Hotz, V. J., R. A. Miller, S. Sanders, and J. Smith (1994, April). A simulation estimator for dynamic models of discrete choice. *Review of Economic Studies* 61(2), 265–89.
- Hu, Y. and M. Shum (2012). Nonparametric identification of dynamic models with unobserved state variables. *Journal of Econometrics* 171(1), 32–44.
- Kalouptsi, M. (2018). Detection and impact of industrial subsidies: The case of chinese ship-building. *Review of Economic Studies* 85(2), 1111–1158.
- Kalouptsi, M., P. T. Scott, and E. A. Souza-Rodrigues (2017). Identification of counterfactuals in dynamic discrete choice models. Working paper.
- Kasahara, H. and K. Shimotsu (2009). Nonparametric identification of finite mixture models of dynamic discrete choices. *Econometrica* 77(1), pp. 135–175.
- Kuersteiner, G. M. and I. R. Prucha (2013). Limit theory for panel data models with cross sectional dependence and sequential exogeneity. *Journal of Econometrics* 174(2), 107 – 126.
- Magnac, T. and D. Thesmar (2002). Identifying dynamic discrete decision processes. *Econometrica* 70(2), 801–816.
- Melnikov, O. (2013). Demand for differentiated durable products: The case of the us computer printer market. *Economic Inquiry* 51(2), 1277–1298.
- Morales, E., G. Sheu, and A. Zahler (2015). Extended gravity. *NBER Working Paper* 21351.
- Newey, W. K. and D. McFadden (1994). Large sample estimation and hypothesis testing. *Handbook of Econometrics* 4, 2111–2245.
- Newey, W. K. and J. L. Powell (2003). Instrumental variable estimation of nonparametric models. *Econometrica* 71(5), 1565–1578.
- Norets, A. and X. Tang (2014). Semiparametric inference in dynamic binary choice models. *The Review of Economic Studies* 81(3), 1229–1262.
- Pakes, A. (2010). Alternative models for moment inequalities. *Econometrica* 78(6), 1783–1822.
- Pesendorfer, M. and P. Schmidt-Dengler (2008). Asymptotic least squares estimators for dynamic games. *The Review of Economic Studies* 75(3), 901–928.
- Rust, J. (1987). Optimal replacement of gmc bus engines: an empirical model of harold zurcher. *Econometrica* 55(5), 999–1033.

- Rust, J. (1994). Structural estimation of markov decision processes. *Handbook of econometrics 4*(4), 3081–3143.
- Rust, J. and C. Phelan (1997). How social security and medicare affect retirement behavior in a world of incomplete markets. *Econometrica: Journal of the Econometric Society*, 781–831.
- Scott, P. T. (2013). Dynamic discrete choice estimation of agricultural land use. *Working Paper*.
- Souza-Rodrigues, E. A. (2016). Nonparametric regression with common shocks. *Econometrics 4*(3), 36.
- Traiberman, S. (2018). Occupations and import competition evidence from danish matched employee-employer data.
- Wright, B. (2014). Global biofuels: key to the puzzle of grain market behavior. *Journal of Economic Perspectives 28*(1), 73–98.

A Appendix

A.1 General ECCP Equation Derivation

In this section, we offer a general derivation of Euler equations in conditional choice probabilities relying on finite dependence as defined in Section 3.³³

Arcidiacono and Miller (2011) show that the conditional value functions, v_a , can be represented by functions of flow payoffs and conditional choice probabilities for *any* sequence of future choices, optimal or not. To derive such a representation, begin with an arbitrary initial state ω_{mt} . Consider a sequence of actions from t to $t+\tau$ (where $\tau \geq 1$). Suppose the initial choice is a , and let j denote another arbitrary element of the choice set \mathcal{A} . Let $a_d \in \mathcal{A}$ denote the d -th choice in the sequence following a .

Recall equation (9), which is a vector across rows of the individual state k and absorbs the aggregate state ω_{mt} into mt subscripts:

$$\pi_{amt} + \beta \varepsilon_{am,t,t+1}^V = V_{mt} - \beta F_{amt}^k V_{mt+1} - \psi_{amt}.$$

We then substitute for V_{mt+1} using equation (9) again, using a_1 as the action instead of a :

$$\begin{aligned} \pi_{amt} + \beta \varepsilon_{am,t,t+1}^V &= V_{mt} - \psi_{amt} - \beta F_{amt}^k (\pi_{a_1 mt+1} + \beta \varepsilon_{a_1,m,t+1,t+2}^V + \psi_{a_1 mt+1}) \\ &\quad - \beta^2 F_{amt}^k F_{a_1 mt+1}^k V_{mt+2}. \end{aligned}$$

Repeated substitution of V_{mt+d} above leads to:

$$\begin{aligned} \pi_{amt} + \beta \varepsilon_{am,t,t+1}^V &= V_{mt} - \psi_{amt} \\ &\quad - F_{amt}^k \left[\sum_{d=1}^{\tau} \beta^d \Lambda_{amt d} (\pi_{a_d mt+d} + \beta \varepsilon_{a_d,m,t+d,t+d+1}^V + \psi_{a_d mt+d}) \right] \\ &\quad - \beta^{\tau+1} F_{amt}^k \Lambda_{amt \tau} V_{mt+\tau+1} \end{aligned} \tag{36}$$

where the matrices $\Lambda_{amt d}$ are defined recursively:

$$\begin{aligned} \Lambda_{amt d} &= I, & \text{for } d = 1, \\ \Lambda_{amt d} &= \Lambda_{amt,d-1} F_{a_d mt+d}^k, & \text{for } d \geq 2, \end{aligned}$$

Next, finite dependence allows us to eliminate the $V_{mt+\tau+1}$, resulting in an ECCP equation that forms the basis of our identification arguments. Given τ -period finite dependence, for a pair

³³This derivation (and Proposition 2) can be extended to a notion of finite dependence in which the sequence of actions leading to convergence of the individual state variables may be mixed actions.

of actions (a, j) , we can construct sequences (a, a_1, \dots, a_τ) and (j, j_1, \dots, j_τ) such that³⁴

$$F_{amt}^k F_{a_1 mt+1}^k \cdots F_{a_\tau mt+\tau}^k = F_{jmt}^k F_{j_1 mt+1}^k \cdots F_{j_\tau mt+\tau}^k,$$

i.e.,

$$F_{amt}^k \Lambda_{amt\tau} = F_{jmt}^k \Lambda_{jmt\tau}. \quad (37)$$

We then difference equation (36) across the two sequences of actions. Because of (37), the last term cancels, and the result is equation (23).

A.2 Identification

A.2.1 Proof of Lemma 1

We omit the subscripts i and m to simplify notation. Suppose Assumption 1 holds.

(i) From the definition of $e^h(a, k, \omega_t, \omega_{t+1}^*)$,

$$\begin{aligned} E[e^h(a, k, \omega_t, \omega_{t+1}^*) | \mathcal{I}_t] &= E\left[\sum_{k'} e^h(k', \omega_t, \omega_{t+1}^*) F^k(k'|a, k, \omega_t) | \mathcal{I}_t\right] \\ &= E\left[\sum_{k'} \left(\int_{\omega'} h(k', \omega') dF^\omega(\omega'|\omega_t) - h(k', \omega_{t+1}^*)\right) F^k(k'|a, k, \omega_t) | \mathcal{I}_t\right] \\ &= \sum_{k'} \int_{\omega'} h(k', \omega') dF^\omega(\omega'|\omega_t) F^k(k'|a, k, \omega_t) \\ &\quad - \sum_{k'} \int_{\omega_{t+1}^*} h(k', \omega_{t+1}^*) dF^\omega(\omega_{t+1}^*|\omega_t) F^k(k'|a, k, \omega_t) \\ &= 0. \end{aligned}$$

(ii) By the law of iterated expectations,

$$E[e^h(a, k, \omega_t, \omega_{t+1}^*) | z_t] = E[E[e^h(a, k, \omega_t, \omega_{t+1}^*) | \mathcal{I}_t] | z_t] = 0,$$

where the second equality follows from (i).

Note also that, given that the time- t information set \mathcal{I}_t includes current and past variables, Lemma 1 also implies that $E[e^h(a, k, \omega_t, \omega_{t+1}^*) | z_{t-d}]$, for all a, k and any $d \geq 1$. In particular, $E[e^h(a, k, \omega_{t+d}, \omega_{t+d+1}^*) | z_t]$.

(iii) Next, fix a and k , and simplify notation further by defining $e^h(a, k_t, \omega_t, \omega_{t+1}^*) = e_{t+1}^h$. Note that not only current and past states (k, ω) belong to the information set available to agents \mathcal{I}_t , but also past prediction errors. I.e., $\{e_t^h, e_{t-1}^h, \dots, e_1^h\} \in \mathcal{I}_t$. We can then let $z_t = e_{t-d}^h$ for $d \geq 1$ and use result (ii) above to establish that $E[e_{t-d}^h e_t^h] = 0$. Thus, expectational errors are serially

³⁴Recall that the terms in the sequences depend on the particular initial pair of actions (a, j) chosen.

uncorrelated.

A.2.2 Proof of Proposition 1

Assume single-action τ -period dependence holds for action J . Then, equation (23) simplifies to

$$\begin{aligned} & (\psi_{jmt} - \psi_{amt}) + (F_{jmt}^k - F_{amt}^k) \sum_{d=1}^{\tau} \beta^d \Lambda_{Jmtd} \psi_{Jmt+d} \\ &= \bar{\pi}_{amt} - \bar{\pi}_{jmt} - (F_{jmt}^k - F_{amt}^k) \sum_{d=1}^{\tau} \beta^d \Lambda_{Jmtd} \bar{\pi}_{Jmt+d} + u_{ajmt}. \end{aligned} \quad (38)$$

where the matrix Λ_{Jmtd} is defined recursively

$$\begin{aligned} \Lambda_{Jmtd} &= I, & \text{for } d = 1 \\ \Lambda_{Jmtd} &= \Lambda_{Jmtd-1} F_{Jmt+d}^k, & \text{for } d \geq 2, \end{aligned}$$

and the unobservable term is $u_{ajmt} = \tilde{\xi}_{ajmt} + \tilde{e}_{ajmt}^V$, with

$$\tilde{\xi}_{ajmt} = (\xi_{amt} - \xi_{jmt}) - (F_{jmt}^k - F_{amt}^k) \sum_{d=1}^{\tau} \beta^d \Lambda_{Jmtd} \xi_{Jmt+d}, \quad (39)$$

$$\tilde{e}_{ajmt}^V = \beta (e_{am,t,t+1}^V - e_{jm,t,t+1}^V) - (F_{jmt}^k - F_{amt}^k) \sum_{d=1}^{\tau} \beta^d \Lambda_{Jmtd} e_{Jm,t+d,t+d+1}^V. \quad (40)$$

For any known (and conformable) function $h(z_{mt})$, multiply both sides of (38) and take the expectation. We eliminate the error terms $\tilde{\xi}_{ajmt}$ and \tilde{e}_{ajmt}^V by Assumption 2.(ii)–(iii). Then,

$$\begin{aligned} & E \left[h(z_{mt}) \left((\psi_{jmt} - \psi_{amt}) + (F_{jmt}^k - F_{amt}^k) \sum_{d=1}^{\tau} \beta^d \Lambda_{Jmtd} \psi_{Jmt+d} \right) \right] \\ &= E \left[h(z_{mt}) \left(\bar{\pi}_{amt} - \bar{\pi}_{jmt} - (F_{jmt}^k - F_{amt}^k) \sum_{d=1}^{\tau} \beta^d \Lambda_{Jmtd} \bar{\pi}_{Jmt+d} \right) \right], \end{aligned} \quad (41)$$

where the expectations are taken over $(z_{mt}, w_{mt}, \dots, w_{mt+\tau})$.

The LHS of (41) can be recovered from the data (using the results of Lemma 3, in Appendix A.3.3). Then, for any two primitives $\bar{\pi}$ and $\bar{\pi}'$,

$$\begin{aligned} & E \left[h(z_{mt}) \left(\bar{\pi}_{amt} - \bar{\pi}_{jmt} - (F_{jmt}^k - F_{amt}^k) \sum_{d=1}^{\tau} \beta^d \Lambda_{Jmtd} \bar{\pi}_{Jmt+d} \right) \right] \\ &= E \left[h(z_{mt}) \left(\bar{\pi}'_{amt} - \bar{\pi}'_{jmt} - (F_{jmt}^k - F_{amt}^k) \sum_{d=1}^{\tau} \beta^d \Lambda_{Jmtd} \bar{\pi}'_{Jmt+d} \right) \right]. \end{aligned}$$

By the completeness condition (Assumption 2.(i)),

$$\begin{aligned} & \bar{\pi}_{amt} - \bar{\pi}_{jmt} - (F_{jmt}^k - F_{amt}^k) \sum_{d=1}^{\tau} \beta^d \Lambda_{Jmtd} \bar{\pi}_{Jmt+d} \\ &= \bar{\pi}'_{amt} - \bar{\pi}'_{jmt} - (F_{jmt}^k - F_{amt}^k) \sum_{d=1}^{\tau} \beta^d \Lambda_{Jmtd} \bar{\pi}'_{Jmt+d}, \end{aligned} \quad (42)$$

for almost all $(w_{mt}, \dots, w_{mt+\tau})$. Consider (42) for $j = J$. Because $\bar{\pi}_J(k, w)$ is known for all observed states (k, w) , we conclude that $\bar{\pi}_{amt} = \bar{\pi}'_{amt}$ almost surely.

A.2.3 Proof of Proposition 2

Equation (27) is a linear regression equation, and $E[Z_{mt}u_{ajmt}] = 0$ and $\text{rank}(E[Z_{mt}\tilde{X}_{ajmt}]) = P$ are the standard orthogonality and rank conditions, respectively, for parameter identification.

A.3 Estimation

A.3.1 First Step

Proof of Lemma 2. Given that $\{a_{imt}, k_{imt} : i = 1, \dots, N\}$ are i.i.d. conditional on ω_{mt} , the first part of the Lemma (the almost sure convergence) follows by an immediate application of the Law of Large Numbers for exchangeable random variables (see Hall and Heyde (1980), p. 202, (7.1)).

The second part is obtained in three steps. First, Horvath and Yandell (1988) presents a Law of Iterated Logarithm (LIL) applied to both kernel and nearest neighbor estimators for conditional probabilities (see their Corollary 5.1). The i.i.d. sample in Horvath and Yandell (1988) can be replaced by the assumption that the sample is i.i.d. conditional on the common shocks following the arguments in Souza-Rodrigues (2016).³⁵ The LIL then holds for almost all ω_{mt} . Finally, it is straightforward to adapt the kernel regression results to simple frequency estimators (i.e., use simple indicator functions as kernel functions).

A.3.2 Second Step

Recall that $g_{mt}(\theta) = h(z_{mt})u_{mt}(\theta, \delta_{mt})$. Define the following functions:

$$\tilde{g}_M(\theta) = \frac{1}{M(T-\tau)} \sum_{m=1, t=1}^{M, (T-\tau)} g_{mt}(\theta). \quad (43)$$

and

$$\tilde{Q}_M(\theta) = \tilde{g}_M(\theta)' \mathbf{W}_M \tilde{g}_M(\theta). \quad (44)$$

³⁵Souza-Rodrigues (2016) establishes the asymptotic properties of the kernel regression estimator for cross-sectional data in the presence of common shocks.

The criterion function $\tilde{Q}_M(\theta)$ is similar to $\hat{Q}_M(\theta)$ but makes use of δ_{mt} instead of the estimator $\hat{\delta}_{mt}$. I.e., $\tilde{Q}_M(\theta)$ is an unfeasible GMM criterion function, while $\hat{Q}_M(\theta)$ is feasible. The unfeasible estimator $\tilde{\theta}_M$ (approximately) minimizes $\tilde{Q}_M(\theta)$ over Θ .

Proof of Proposition 3. A straightforward application of Theorem 2.6 in Newey and McFadden (1994) proves that the unfeasible estimator $\tilde{\theta}_M$ is a consistent estimator of θ_0 . To show that the feasible estimator $\hat{\theta}_M$ is consistent as well, it suffices to show that $\hat{Q}_M(\theta)$ converges in probability to $\tilde{Q}_M(\theta)$ uniformly over Θ . To do so, define the difference $v_{mt} = \hat{g}_{mt}(\theta) - g_{mt}(\theta)$, and

$$v_M(\theta) = \frac{1}{M(T-\tau)} \sum_{m=1, t=1}^{M, (T-\tau)} v_{mt}(\theta).$$

Then,

$$\begin{aligned} \hat{Q}_M(\theta) &= [\tilde{g}_M(\theta) + v_M(\theta)]' \mathbf{W}_M [\tilde{g}_M(\theta) + v_M(\theta)] \\ &= \tilde{Q}_M(\theta) + v_M(\theta)' \mathbf{W}_M v_M(\theta) + 2\tilde{g}_M(\theta)' \mathbf{W}_M v_M(\theta). \end{aligned}$$

Given Condition 2(ii), it suffices to show that both $\tilde{g}_M(\theta)$ and $v_M(\theta)$ converge to zero in probability uniformly over Θ .

By Conditions 2(i),(iii),(v), and (vi), $\tilde{g}_M(\theta)$ satisfies the uniform Weak Law of Large Numbers, and therefore converges in probability to zero uniformly over Θ as $M \rightarrow \infty$. Now consider $v_M(\theta)$. Note that

$$v_{mt} = h(z_{mt}) \left(u_{mt}(\theta, \hat{\delta}_{mt}) - u_{mt}(\theta, \delta_{mt}) \right),$$

and take a mean-value expansion of $u_{mt}(\theta, \hat{\delta}_{mt})$ about δ_{mt} :

$$u_{mt}(\theta, \hat{\delta}_{mt}) - u_{mt}(\theta, \delta_{mt}) = \nabla_{\delta} u_{mt}(\theta, \delta_{mt}^*) \left(\hat{\delta}_{mt} - \delta_{mt} \right),$$

where δ_{mt}^* lies between $\hat{\delta}_{mt}$ and δ_{mt} . Next, note that

$$\begin{aligned} E \left[\sup_{\theta \in \Theta} \|v_M(\theta)\| \right] &\leq \frac{1}{M(T-\tau)} \sum_{m=1, t=1}^{M, (T-\tau)} E \left[\sup_{\theta \in \Theta} \|h(z_{mt}) \nabla_{\delta} u_{mt}(\theta, \delta_{mt}^*)\| \left\| \hat{\delta}_{mt} - \delta_{mt} \right\| \right] \\ &\leq \frac{B}{M(T-\tau)} \sum_{m=1, t=1}^{M, (T-\tau)} E \left[\left\| \hat{\delta}_{mt} - \delta_{mt} \right\|^2 \right] \end{aligned} \quad (45)$$

where the second inequality follows from the Cauchy–Schwarz inequality and Condition 2(vii). Because $\left\| \hat{\delta}_{mt} - \delta_{mt} \right\| \xrightarrow{p} 0$, as $N \rightarrow \infty$, by Lemma 2, we have that $E \left[\left\| \hat{\delta}_{mt} - \delta_{mt} \right\|^2 \right] = o(1)$, and,

so, the right-hand-side of (45) converges to zero as $N \rightarrow \infty$ for all M and T . We conclude that

$$\sup_{\theta \in \Theta} \left\| \widehat{Q}_M(\theta) - \widetilde{Q}_M(\theta) \right\| \xrightarrow{p} 0, \text{ as } (M, N) \rightarrow \infty.$$

Proof of Proposition 4. By standard arguments (see Theorem 3.2 in Newey and McFadden (1994), the unfeasible estimator $\widetilde{\theta}_M$ satisfies

$$\widetilde{\theta}_M - \theta_0 = -[\mathbf{G}'\mathbf{W}\mathbf{G}]^{-1} \mathbf{G}'\mathbf{W}g(\theta_0) + o_p\left(1/\sqrt{M}\right), \quad (46)$$

and is asymptotically normal,

$$\sqrt{M} \left(\widetilde{\theta}_M - \theta_0 \right) \xrightarrow{p} N(0, \mathbf{V}),$$

under Conditions 3(i)-(iv). The asymptotic distribution of the feasible estimator $\widehat{\theta}_M$ is the same as the asymptotic distribution of the unfeasible $\widetilde{\theta}_M$ provided

$$\left\| \widehat{\theta}_M - \widetilde{\theta}_M \right\| = o_p\left(\frac{1}{\sqrt{M}}\right).$$

From (46), it is clear that

$$\widetilde{\theta}_M - \widehat{\theta}_M = [\mathbf{G}'\mathbf{W}\mathbf{G}]^{-1} \mathbf{G}'\mathbf{W}v_M(\theta_0) + o_p\left(1/\sqrt{M}\right).$$

So,

$$\left\| \widehat{\theta}_M - \widetilde{\theta}_M \right\| \leq \left\| [\mathbf{G}'\mathbf{W}\mathbf{G}]^{-1} \right\| \|\mathbf{G}\| \|\mathbf{W}\| \|v_M(\theta_0)\| + o_p\left(1/\sqrt{M}\right).$$

Note that

$$E[\|v_M(\theta_0)\|] \leq \frac{B}{M(T-\tau)} \sum_{m=1, t=1}^{M, (T-\tau)} \left(E \left[\left\| \widehat{\delta}_{mt} - \delta_{mt} \right\|^2 \right] \right)^{1/2}$$

by Condition 2(vii). Because $E \left[\left\| \widehat{\delta}_{mt} - \delta_{mt} \right\|^2 \right] = O\left(\frac{\log \log N}{N}\right)$, by Lemma 2, we have that

$\|v_M(\theta_0)\| = O_p\left(\sqrt{\frac{\log \log N}{N}}\right)$, which implies

$$\sqrt{M} \left\| \widehat{\theta}_M - \widetilde{\theta}_M \right\| = O_p\left(\sqrt{\frac{M \log \log N}{N}}\right) = o_p(1)$$

provided $\frac{M \log \log N}{N} \rightarrow 0$.

A.3.3 Additional Result

The next lemma provides a result that is used in Proposition 1. Proposition 1 claims that, for a known function f of $\delta_{mt}^\tau = (\delta_{mt}, \dots, \delta_{mt+\tau})$, quantities of the type $E[h(z_{mt}) f(\delta_{mt}^\tau)]$ can be recovered

from the data. (More specifically, $f(\delta_{mt}^\tau)$ in the proof of Proposition 1 corresponds to the term in parenthesis on the LHS of equation (41).)

Lemma 3. *Suppose the vector (w_{mt}, z_{mt}) is i.i.d. across markets m . Assume*

$$E \left[\|h(z_{mt}) \nabla_{\delta} f(\delta_{mt}^\tau)\|^2 \right] \leq C < \infty.$$

Then

$$\frac{1}{M} \sum_{m=1}^M h(z_{mt}) f(\widehat{\delta}_{mt}^\tau) \xrightarrow{P} E[h(z_{mt}) f(\delta_{mt}^\tau)],$$

as $(M, N) \rightarrow \infty$.³⁶

Proof. First, note that

$$\frac{1}{M} \sum_{m=1}^M h(z_{mt}) f(\widehat{\delta}_{mt}^\tau) = \frac{1}{M} \sum_{m=1}^M h(z_{mt}) f(\delta_{mt}^\tau) + \frac{1}{M} \sum_{m=1}^M h(z_{mt}) \left[f(\widehat{\delta}_{mt}^\tau) - f(\delta_{mt}^\tau) \right].$$

The first term on the right-hand-side converges in probability to $E[h(z_{mt}) f(\delta_{mt}^\tau)]$ as $M \rightarrow \infty$ by the Weak Law of Large Numbers. Applying a mean-value expansion on the second term, we get

$$\frac{1}{M} \sum_{m=1}^M h(z_{mt}) \left[f(\widehat{\delta}_{mt}^\tau) - f(\delta_{mt}^\tau) \right] = \frac{1}{M} \sum_{m=1}^M h(z_{mt}) \nabla_{\delta} f(\delta_{mt}^{\tau*}) \left[\widehat{\delta}_{mt}^\tau - \delta_{mt}^\tau \right]$$

where $\delta_{mt}^{\tau*}$ lies between $\widehat{\delta}_{mt}^\tau$ and δ_{mt}^τ . Next, note that

$$\begin{aligned} E \left[\left\| h(z_{mt}) \nabla_{\delta} f(\delta_{mt}^{\tau*}) \left[\widehat{\delta}_{mt}^\tau - \delta_{mt}^\tau \right] \right\|^2 \right] &\leq \left(E \left[\|h(z_{mt}) \nabla_{\delta} f(\delta_{mt}^{\tau*})\|^2 \right] E \left[\left\| \widehat{\delta}_{mt}^\tau - \delta_{mt}^\tau \right\|^2 \right] \right)^{1/2} \\ &\leq C \left(E \left[\left\| \widehat{\delta}_{mt}^\tau - \delta_{mt}^\tau \right\|^2 \right] \right)^{1/2}, \end{aligned}$$

where the first inequality follows from the Cauchy–Schwarz inequality, and the second inequality from the regularity condition $E \left[\|h(z_{mt}) \nabla_{\delta} f(\delta_{mt}^\tau)\|^2 \right] \leq C < \infty$. By Lemma 2, $E \left[\left\| \widehat{\delta}_{mt}^\tau - \delta_{mt}^\tau \right\|^2 \right]$ converges to zero as $N \rightarrow \infty$, which implies

$$\frac{1}{M} \sum_{m=1}^M h(z_{mt}) \left[f(\widehat{\delta}_{mt}^\tau) - f(\delta_{mt}^\tau) \right] \xrightarrow{P} 0, \text{ as } N \rightarrow \infty, \text{ for all } M.$$

■

³⁶The same result applies if (w_{mt}, z_{mt}) is stationary and ergodic, if we average the term $[h(z_{mt}) f(\widehat{\delta}_{mt}^\tau)]$ over $T - \tau$ time periods, and if we take $(T, N) \rightarrow \infty$.

A.4 The Standard CCP Estimator

Here we explain the standard CCP approach implemented in the Monte Carlo experiment to estimate the model parameters. By “standard,” we mean involving a full specification of how all state variables evolve, and not relying on Euler equations. Following Hotz and Miller (1993), this CCP approach avoids the computational burden of solving the dynamic problem within the estimation algorithm associated with Rust’s (1987) nested fixed point approach.

The estimation here follows section 2.1 of Kalouptsidei et al. (2017) and we refer readers to it for details. Estimation begins by estimating choice probabilities conditional on individual states and the modeled exogenous state variable, i.e., $p(k, w)$. Let F_b represent the stochastic matrix for observable state variables (k, w) conditional on buying the product, and let F_{nb} represent the stochastic matrix when the action is not buying the product. Kalouptsidei et al. (2017) shows that

$$\pi_b = A\pi_{nb} + \mathbf{b},$$

where $A = (I - \beta F_b)(I - \beta F_{nb})^{-1}$ and $\mathbf{b} = A\psi_{nb} - \psi_b$, where ψ_a stacks $\psi_a(p(k, w))$ across all values of (k, w) .

We estimate the payoff parameters θ using a Minimum Distance estimator, i.e., by minimizing the L2 norm of

$$\pi_b(\theta) - A\pi_{nb}(\theta) - \mathbf{b}.$$

Given the parameterization, this is achieved by a linear regression of the vector \mathbf{b} on the matrix

$$\left[(\mathbf{1} - A\mathbf{k}) \quad , \mathbf{w} \right],$$

where $\mathbf{1}$ is a vector of ones, \mathbf{k} is a dummy vector equal to one in states where the good is owned, and \mathbf{w} is the vector of prices.³⁷

³⁷Note that one can estimate the model parameters either by minimizing the distance between \mathbf{b} and $\pi_b(\theta) - A\pi_{nb}(\theta)$, or by minimizing the distance between the (nonparametrically) estimated CCP, p , and the CCP generated by the model, $p(\theta)$. See Pesendorfer and Schmidt-Dengler (2008).