

WORKING PAPER · NO. 2020-106

An Instrumental Variable Approach to Dynamic Models

Steven T. Berry and Giovanni Compiani

AUGUST 2020

An Instrumental Variable Approach to Dynamic Models*

Steven T. Berry[†] and Giovanni Compiani[‡]

August 4, 2020

Abstract

We present a new class of methods for identification and inference in dynamic models with serially correlated unobservables, which typically imply that state variables are econometrically endogenous. In the context of Industrial Organization, these state variables often reflect econometrically endogenous market structure. We propose the use of Generalized Instrument Variables methods to identify those dynamic policy functions that are consistent with instrumental variable (IV) restrictions. Extending popular “two-step” methods, these policy functions then identify a set of structural parameters that are consistent with the dynamic model, the IV restrictions and the data. We provide computed illustrations to both single-agent and oligopoly examples. We also present a simple empirical analysis that, among other things, supports the counterfactual study of an environmental policy entailing an increase in sunk costs.

1 Introduction

We propose an instrumental variable (IV) approach to identification and inference in dynamic models in the presence of serially correlated unobservables. Such serial correlation typically leads to dynamic state variables that are econometrically endogenous, which creates problems for identification and inference. As a result, much of the literature to date either assumes away serial correlation in the unobservables or else deals with such correlation in particularly simple fashions.

*We are grateful to Allan Collard-Wexler for generously providing the data, and to Xiaohong Chen, Liran Einav, Phil Haile, Francesca Molinari, Jesse Shapiro, Paulo Somaini for helpful comments and suggestions. We are also grateful to seminar participants at NAMES 2017, Northwestern, Harvard/MIT, Stanford, Duke, 2018 and 2019 Conference on Dynamic Models, Boston University, ASSA 2019.

[†]Yale University, Cowles Foundation and NBER. Email: steven.berry@yale.edu.

[‡]University of Chicago, Booth School of Business. Email: Giovanni.Compiani@chicagobooth.edu.

We mostly focus on applications to dynamic models of Industrial Organization (IO). These models often feature state variables that measure various kinds of “market structure,” such as the number of firms, the number of retail outlets, the vector of current productivity levels of firms and so forth. As an example, consider a simple model of entry and exit where profits depend on the number of firms in the market, N . Standard empirical approaches assume that any variables affecting firms’ profits and not captured by the data—e.g. unobserved demand or supply shocks—are independent over time. This implies that “market structure” N is independent of the contemporaneous unobservables and thus the effect of N on, say, firm entry can be directly observed in the data. In contrast, when the unobservables are persistent over time, markets with large N are likely to be more systematically profitable in terms of unobservables. Thus, the observed entry probabilities reflect the correlation between unobservables and N and we cannot directly observe the “causal effect” of N on entry. This is a classic endogeneity problem.

In the dynamic context, IV methods provide a natural and economically meaningful solution to this familiar IO problem of endogenous market structure. More broadly, this paper is part of the research agenda that relates the formal identification of IO models to classic IV intuition, as in standard equilibrium models of supply and demand. The goal is to address a persistent critique of IO models that claims they are typically not well-identified. Specific examples of this agenda include Cournot-style models, as in Bresnahan (1989), differentiated products demand and supply market equilibrium, as in Berry and Haile (2014), cross-sectional market structure (“static entry” models), as in Tamer (2003), and auction cost heterogeneity, as in Somaini (2015).¹

1.1 Idea of the Paper

Our paper builds on the intuition of classic “two-step” methods, following on Hotz and Miller (1993) (henceforth, HM), that distinguish between the identification of [i] the structural parameters of an underlying dynamic model and [ii] the policy function that results from the solution of that dynamic model evaluated at the true value of the structural parameters. It is this policy function that (according to the model) generates the data.

The task of identification and inference is made much easier by the assumption that unobservable shocks are distributed independently over time. This is made clear in Rust (1987) and exploited in the HM “conditional choice probability” or “CCP” approach. In the related IO literature, the shocks are then typically assumed to also be private information.²

¹In addition, as in much of the auction literature, there are many formal IO identification arguments that do not so clearly involve instrumental variables.

²See Pesendorfer and Schmidt-Dengler (2008), Bajari et al. (2007) and Pakes et al. (2007) for a discussion of HM style methods, with pure i.i.d. private information shocks, extended to a dynamic oligopoly context with possibly multiple equilibria. An early review of this approach is in Akerberg et al. (2007).

Under these assumptions, the dynamic policy function is often point-identified “directly from the data.” For instance, in the case of dynamic discrete choice models, estimating the policy boils down to estimation of conditional probabilities. The structural parameters are then identified as those that are consistent with the observed policy function.

However, the simplicity of these methods depends critically on the econometric exogeneity of dynamic states. Once unobservables are allowed to be serially correlated, the dynamic states become econometrically endogenous. This is because the dynamic states reflect past values of the unobservables which, due to serial correlation, are typically *not* independent of the current unobservable entering the policy function. The econometric endogeneity problem here is classic in its form: the “right hand side” state variables in the dynamic policy function are correlated with the unobservables that enter the same function.

In order to tackle the endogeneity of the dynamic states, we rely on instrumental variables. These instruments have the classic features that they [1] do not directly enter today’s policy decision, [2] are assumed to be exogenous (independent of the unobservables), and yet [3] are correlated with the current state, likely because they effected past policy decisions that are correlated with present states. In a dynamic entry model, an example would be past market size or past regulatory environments that influenced past decisions to enter a market. In the presence of sunk costs, these past decisions will continue to be correlated with current market structure, even if current entry decisions are only driven by current market size and current regulations. We discuss further examples of possible instrumental variables after we have formally defined key features of the model.

Traditional IV and panel data methods face a difficult problem in our context: the policy function is derived from the “structural” dynamic model and this typically implies that the policy function is not additively separable in the serially correlated unobservable(s). The nonseparability of the policy function in unobservables creates difficulties for both identification and inference. Luckily, there is a large recent literature on the nonparametric identification of functions with nonseparable unobservables and econometrically endogenous right-hand side variables, sometimes mixed with a classic panel data structure. In the easiest possible examples for us, the dynamic policy function will be point-identified even in the presence of serial correlation, but more general cases may lead only to set identification. To consider more general cases, we leverage an existing large literature on identification and inference in partially identified models, including Manski and Tamer (2002), Tamer (2003), Manski (2003), Chernozhukov et al. (2007), Berry and Tamer (2007), Ciliberto and Tamer (2009), Beresteanu et al. (2011), Galichon and Henry (2011), Chesher (2010) and Andrews and Shi (2013).

One paper that sums up and extends an IV style literature on this topic is Chesher and Rosen (2017) (henceforth, CR), who discuss a class of “Generalized Instrumental Variable” (henceforth, GIV) methods. In addition to emphasizing an appropriate IV framework for the identification of a very broad class of dynamic policy functions, CR closely build on the

work of Galichon and Henry (2011) and Beresteanu et al. (2011) to characterize the sharp identified set. This characterization will help us build intuition about how instruments serve to (set) identify policy functions.

The identifying power of these instrumental variable methods is increased by the presence of multiple periods of data. In particular, we note that even in the absence of any instrumental variables, nonseparable policy functions can be usefully restricted purely from the presence of multiple periods of data, as in the nonseparable error, nonparametric panel data papers of Altonji and Matzkin (2005) and Athey and Imbens (2006).³

We illustrate our approach in a simple single-agent entry and exit model. This minimal example allows us to build intuition about the sources of identification as well as to explore how the number of time periods, the presence of exogenous covariates and the strength of the instruments affect the identified set for the structural parameters. We then apply the method to data from the US ready-mix concrete industry and consider a counterfactual policy that increases the magnitude of the sunk costs of entry into the market. When we compare our approach to three different methods that assume away serial correlation in the unobservables, we find that the latter results are significantly different than ours. Moreover, the sign of the bias varies across the three methods. Two approaches tend to over-predict the responses to the policy in terms of both the number of firms and the fraction of new entrants. This stems from the fact that, in the counterfactual, the unobservables exhibit too much volatility over time when serial correlation is ruled out. On the other hand, a third approach estimates a very large sunk cost (as a way to match the persistence in the data without appealing to serially correlated unobservables) and thus predicts no response to the policy change. The difficulty in a priori signing the bias from standard, more restrictive methods further motivates our contribution.

1.2 Some Related Papers

The literature on the identification and estimation of dynamic problems is immense and we can only highlight a set of related literatures here.

We are obviously not the first authors to consider the issue of serially correlated unobservables in dynamic models, including dynamic games. Outside of the “two-step” literature following on CCP methods, there is an important set of papers emphasizing computational approaches to estimation that allow for some form of persistent unobservables, sometimes in the form of a limited number of “discrete types” of agents. A classic single-agent example is Keane and Wolpin (1997). A classic oligopoly example is the full-solution approach of Ericson and Pakes (1995) and Pakes and McGuire (2001), who emphasize that serially correlated unobservables are an important feature of realistic dynamic models in IO. These

³These papers do not explicitly consider the fully dynamic problems that we consider here, but instead focus on nonparametric analogues of non-dynamic panel data-style arguments.

computationally oriented papers do not typically discuss formal identification.

Work on the identification of mixture models, as in Kasahara and Shimotsu (2009) and Hu and Shum (2012),⁴ does provide some formal results on identification of discrete dynamic policy functions with persistent unobservables. This work again emphasizes limited forms of discrete heterogeneity. In Berry and Compiani (2020) we show that our framework includes the class of models they consider as a special case. A key restriction in Kasahara and Shimotsu (2009) is that the variation in unobservables is in some well-defined sense lower-dimensional than the variation in the observed data. In particular, the degree of point-identified heterogeneity is limited by the time-series dimension of the data. As a complementary result, our set-identification approach is applicable to settings with as few as two time periods irrespective of the dimension of the unobservable. Of course, if the data exhibits too little variation our identified sets may be so large as to be of little use. In our empirical application, we obtain informative results with fewer than 500 cross-sectional observations.

On the estimation side, Arcidiacono and Miller (2011) provide maximum likelihood computational methods for the structural parameters of dynamic models with discrete persistent heterogeneity. Norets (2009) proposes a Bayesian estimation method for dynamic discrete choice models with serially correlated unobservables. Additional full-solution approaches allowing for serial correlation include Blevins (2016) and Reich (2018). Neither discusses identification formally.

Our paper is also related to the literature on dynamic panel data models that are not derived from explicit dynamic optimization (e.g., Altonji and Matzkin (2005) and Athey and Imbens (2006)) and to the large literature on distinguishing between state dependence and unobserved heterogeneity (see, e.g., Heckman and Singer (1984), Israel (2005), Dubé et al. (2010)). Similar to these papers, we face the challenge of disentangling the roles of past actions and persistent unobservables in driving current outcomes. Our IV intuition is very much in line with this literature on dynamic panels with state dependence. For example, Israel (2005) argues for the usefulness of past exogenous shocks that shift the current state but do not affect today’s decision conditional on the current state. More recently, Heckman et al. (2016) study identification and estimation of dynamic treatment effects allowing for time-invariant unobserved heterogeneity. Honoré and Tamer (2006) discuss set identification in a dynamic panel context.

Even without serially correlated unobservables, there are typically no formal point-identification results for models with continuous actions,⁵ which is why Bajari et al. (2007) (“BBL”) relies on set-identification, as do we. Unlike BBL, we characterize the sharply iden-

⁴See also Hu et al. (2015).

⁵Identification in the case of no serial correlation and discrete actions is considered in Magnac and Thesmar (2002) and related papers.

tified set for the structural parameters in the presence of persistent unobservables using the “generalized instrumental variable (GIV)” framework of CR and related literature.

A recent paper by Kalouptsidei et al. (2020) shares our IV intuition and is in many ways closest to our spirit. They show that in a class of dynamic discrete choice models with serially correlated market-level unobservables, one can obtain Euler equations that point-identify the profit parameters. This approach leads to computationally light linear IV estimators that are robust to endogeneity problems caused by the market-level unobservables. However, it does not address identification of the joint distribution of the unobservables over time and thus cannot be used to perform counterfactuals requiring that distribution as an input. The paper provides some interesting examples of IV potential applications with serially correlated market-level states, including durable goods demand, land use and dynamic labor supply. They discuss possible instruments in these settings. The examples and instruments could be applicable to our methods as well.

The rest of this paper is organized as follows. Section 2 introduces the general model and identification framework for single-agent problems. Section 3 illustrates the approach via a simple entry example. Section 4 extends the analysis to the oligopoly setting. Section 5 contains the empirical application and Section 6 concludes. Appendix A specializes the GIV framework to the case of discrete persistent unobserved heterogeneity that has been used in prior literature. Appendix B provides more details on our leading single-agent example, whereas Appendix C presents an additional computed oligopoly model. Finally, Appendix D discusses the second step of our empirical application in more detail and Appendix E contains a proof omitted from the text.

2 Model and Identification

In this section, we present the formal model. After introducing variables and notation, we focus on the single-agent case and illustrate the GIV approach through a simple monopolist entry example. We then extend the analysis to settings with multiple agents in each market.

2.1 Variables and Notation

We consider a model that generates data on a large set of markets, with one or more agents⁶ per market, and a fixed (perhaps small) number of time periods denoted by $t = 1, \dots, T$. We may additionally have access to some subset of variables for prior periods, $t < 1$. In the general oligopoly model, markets are indexed by i and firms within markets are indexed by j . We do not model cross-market interactions. Our simpler examples will involve a single firm per market.

⁶Since many IO dynamic models involve firms making decisions over time, we use the words “firms” and “agents” interchangeably throughout the paper.

In each market in each time period, each firm takes an action (or actions) denoted a_{ijt} . These actions contribute over time to the firm's observed current state(s), denoted x_{ijt} . The set of feasible actions for a firm with state x_{ijt} is denoted $\mathcal{A}(x_{ijt})$. As one example, in an entry model there might be a scalar action a_{ijt} , equal to one or zero, that indicates the decision to operate in the market in period $t + 1$. A scalar state x_{ijt} might then be whether firm j operates in market i in period t .

There are also observed exogenous states, w_{ijt} , that evolve separately from the firms' actions. Some or all of the exogenous states may be shared across firms. In some cases, we may observe some partial information on exogenous variables from before the beginning of our full panel dataset. We denote these variables, which will later prove useful as instruments, as r_i .

In addition, there are unobserved (to us) state(s) u_{ijt} that also evolve exogenously from the actions of firms. For example, in an entry model, u_{ijt} may represent the component of fixed costs not captured by the data. Within market, the unobservables may be correlated both across time and firms. The u_{ijt} are the only variables that the firms observe but we do not. In the oligopoly context, we treat the serially correlated component of u_{ijt} as commonly observed by all firms. In some cases, it is also useful to model an independent (over time and firms) component that is private information to the firm.⁷

Suppose that there are a maximum of J firms within each market. We define

$$a_{it} \equiv (a_{i1t}, a_{i2t}, \dots, a_{iJt}) \tag{1}$$

and we define the market-time vectors x_{it} , w_{it} and u_{it} in a similar fashion. As further notation, we let across-time, within-market vectors of variables (and their respective supports) be denoted $a_i = (a_{i1}, \dots, a_{iT}) \in \mathbb{A}^T$, $x_i = (x_{i1}, \dots, x_{iT}) \in \mathbb{X}^T$, $w_i = (w_{i1}, \dots, w_{iT}) \in \mathbb{W}^T$, and $u_i = (u_{i1}, \dots, u_{iT}) \in \mathbb{U}^T$.

The probability that the vector u_i of unobservables (across time and firms within market) lies in the set $\mathcal{S} \subset \mathbb{U}$ is denoted by

$$\Phi(\mathcal{S}; \theta_u), \tag{2}$$

where the vector θ_u parameterizes the distribution of the vector of market unobservables across time and firms. The parameter θ_u will often, *inter alia*, control the degree of serial correlation in the unobservables. The single-period profit of firm j in market i in period t is given by the function

$$\pi_j(a_{it}, x_{it}, w_{it}, u_{it}; \theta_\pi). \tag{3}$$

The subscript j on the single-period profit function indicates the natural property that firm

⁷The distinction between the full information serially correlated unobservable and the private independent unobservable is similar to the distinction between the variables ν^1 and ν^2 in Pakes et al. (2015). See also Appendix A for more details on this point.

j 's profits depend differently on its own elements of $(a_{ijt}, x_{ijt}, w_{ijt}, u_{ijt})$ as opposed to its rivals'. The unknown parameters of the single period profit function are θ_π . The full vector of structural parameters, θ , then includes the unknown parameters of the single-period profit function and of the distribution of unobservables: $\theta = (\theta_\pi, \theta_u)$.⁸

2.2 Single Firm per Market

We begin with the single-agent case, returning to dynamic oligopoly in section 4. In this special case, we treat each firm (agent) as operating in its own “market” and so we drop the j firm subscripts in $(a_{ijt}, x_{ijt}, w_{ijt}, u_{ijt})$, leaving (for example) a_{it} as the action of the firm in market i at time t . In the single-firm case, we will shorthand the phrase “firm in market i ” as “firm i .”

As is classic in much of the literature following on Rust (1987), we assume that the observed endogenous states of the firm evolve according to the transition probability function

$$\Gamma(x_{it+1}|a_{it}, x_{it}, w_{it}), \quad (4)$$

where Γ gives the probability of each possible future state conditional on the firm’s own action and observable states. As a special case, this could describe deterministic state transitions, where some state occurs with a conditional probability of one. For instance, in a dynamic entry model, the current state (whether the firm is in or out of the market) is equal to the action taken last period. Table 1 gives some examples of actions, states and transition processes that might occur in the IO context.

Table 1: Some Single Agent IO Examples

State, x_{it}	Action, a_{it}	$\mathcal{A}(x_{it})$	Transition
Capital	Investment	\mathbb{R}^+	$x_{it+1} = \tilde{\lambda}x_{it} + a_{it}$
Out/In	Entry/Exit	$\{0, 1\}$	$x_{it+1} = a_{it}$
Retail	# of Stores	\mathcal{I}^+	$x_{it+1} = a_{it}$
Quality	R&D	\mathbb{R}^+	$x_{it+1} \sim f(x_{it}, a_{it})$

Similarly, the observable exogenous states are first-order Markov, with conditional prob-

⁸Following standard practice in the dynamics literature, we assume that the discount factor is known throughout.

abilities given by⁹

$$H(w_{it+1}|w_{it}). \quad (5)$$

In our single-firm examples, we will focus on the special case where the firm-level unobservable u_{it} follows a first-order Markov process. Our leading example here will be a model of first-order serial correlation where u_{it} is a scalar that obeys

$$u_{it} = \rho u_{it-1} + \nu_{it} \sqrt{1 - \rho^2}. \quad (6)$$

In the simplest case, the period t vector ν_{it} innovation might be assumed to have a simple parameterized distribution. The parameter θ_u then includes those parameters plus the serial correlation parameter ρ . In Appendix A, we consider more general forms of unobserved heterogeneity, including cases with a persistent component and an idiosyncratic component in the spirit of Heckman and Singer (1984), Keane and Wolpin (1997), Arcidiacono and Miller (2011) and related literature.

In this case, the firm's dynamic problem is given by the classic Bellman equation:

$$V(x_{it}, w_{it}, u_{it}) = \max_{a_{it} \in \mathcal{A}(x_{it})} (\pi(a_{it}, x_{it}, w_{it}, u_{it}; \theta_\pi) + \delta E_{\theta_u} [V(x_{it+1}, w_{it+1}, u_{it+1}) | a_{it}, x_{it}, w_{it}, u_{it}]). \quad (7)$$

where δ denotes the discount factor and V the value function. Note that, since the function V does not depend on time, we assume a stationary environment. The expected value function in this expression is

$$E_{\theta_u} [V(x_{it+1}, w_{it+1}, u_{it+1}) | a_{it}, x_{it}, w_{it}, u_{it}] = \int \int \int V(x_{it+1}, w_{it+1}, u_{it+1}) d\Gamma(x_{it+1} | a_{it}, x_{it}, w_{it}) dH(w_{it+1} | w_{it}) d\tilde{\Phi}(u_{it+1} | u_{it}; \theta_u) \quad (8)$$

where $\tilde{\Phi}(u_{it+1} | u_{it}; \theta_u)$ denotes the conditional distribution of u_{t+1} given u_t . Note that this is similar to Rust (1987), but we do not make Rust's full conditional independence assumption. Specifically, in order to allow for serial correlation in the unobservables, we do not drop the conditioning on the past unobservable in $\tilde{\Phi}(u_{it+1} | u_{it}; \theta_u)$.

In the single-agent case, there is a unique solution for the value function and we assume standard conditions such that there is a unique policy function consistent with that value

⁹Our framework also allows for the case where the transition processes for x_{it} and w_{it} in (4) and (5) depend on u_{it} (or components of it) provided that they can be identified from the data. Since we know of no empirical models featuring this dependence, we focus on the case where the transitions do not depend on u_{it} throughout the paper.

function.¹⁰ We let σ denote this policy function, so that

$$a_{it} = \sigma(x_{it}, w_{it}, u_{it}), \quad \sigma \in \mathcal{F}. \quad (9)$$

In many cases, reasonable assumptions on the single-period return function and the transition processes imply that the policy function must obey certain qualitative restrictions, such as monotonicity. These restrictions can then be imposed on the set of possible policy functions \mathcal{F} .

For the identification argument presented next, it will also be useful to define the different (counterfactual) policy functions that would be generated by any other possible parameter vectors $\theta = (\theta_\pi, \theta_u)$. In the single-firm case, these policy functions, generated by the model and the unique solution to Bellman's equations, will be denoted by

$$a_{it} = \sigma_\theta(x_{it}, w_{it}, u_{it}). \quad (10)$$

2.3 Identification of the Single-Agent Model

We first focus on identification in the single-agent case. For purposes of identification, we assume that we observe the true distribution of the data, which we denote by

$$P(a_i, x_i, w_i, r_i)$$

This is equivalent to seeing a T -period panel on a very large (in fact, infinite) cross-section of firms or agents. We look to identify (possibly set-identify) the parameters θ . Nothing in our general discussion of identification requires these to be finite dimensional, but in practice the models we consider in the simulations and empirical application are all finite-dimensional. Note that in the case where (a_i, x_i, w_i) are discrete, the single-period profit function may be fully flexibly characterized by a finite number of parameters, one for each combination of (a_i, x_i, w_i) values.

The potential instruments in the model consist of the exogenous variables

$$z_i = (r_i, w_i)$$

The critical assumption that allows for our instrumental variable approach is independence of the instrument and the unobservables:¹¹

$$z_i \perp u_i$$

¹⁰For the technical conditions guaranteeing a unique policy function, see Stokey et al. (1989).

¹¹While we focus on this restriction throughout the paper, CR show that the GIV approach may also be applied under weaker assumptions, such as mean or quantile independence.

Note that the assumption that w_i be exogenous is standard in the existing literature. In addition, we require excluded instruments r_i to deal with the endogeneity of the dynamic states x_i . Table 2 gives some ideas of possible instruments r_i in different contexts. As is usual with discussions of potential instruments, the required independence assumption may be more or less appropriate in different real-world cases.

Table 2: Examples of Possible Instruments r_i

State	Example Instruments
Capital	Past investment cost
Out/In of Market	Past market population, past regulation
# of Stores	Distance from headquarters, interacted with time
Quality	Past R&D shocks, age of firm

In studying identification of the model, we follow the classic “two-step” approach. First, we discuss the (set-)identification of the policy function and serial correlation parameters, using GIV techniques. Given the results of the first step, we then discuss the identification of the structural parameters of the profit function using a broad generalization of existing approaches.

2.4 First Step: Identification of the Policy

The broad idea is to (set-)identify the policy function from classic instrumental variables conditions, extended to cases where the policy function is highly nonlinear in the states. The GIV framework achieves this and it allows us to deal with the following complications arising in many dynamic models of interest:

1. the incompleteness of the model, i.e. the fact that the exogenous variables do not uniquely pin down the endogenous variables (Tamer (2003));
2. the fact that, if the dynamic states and actions are discrete—as in entry/exit models—the policy function is known to be generally only partially identified in the absence of a model for the endogenous explanatory variables (Chesher (2010));
3. lack of point-identification of the parameters, even in the absence of problems 1 and 2, e.g. due instruments that are not strong enough.

In applications, we may have all or none of these problems. If the model and data generating process in fact imply point-identification, then the sharply identified set will collapse to the

true parameter value. In the single agent case, an incomplete model can follow from the presence of unknown initial conditions, i.e. the fact that the joint distribution of (x_{i1}, u_{i1}) is not known.¹² Traditionally, solutions to the initial conditions problem include either [i] parameterizing the initial joint distribution of states and unobservables or [ii] specifying some process for the past history of the firm that uses the model parameters to construct that same initial joint distribution.¹³ We argue that if the parameterization in method [i] is so flexible as to not impact the resulting identified set, then we might just as well look for the sharply identified set that does not restrict the initial distribution.

As in Tamer (2003), any given action a_{it} naturally leads to conditions on sets of unobservables. In particular, following CR and using similar notation, if the sequence (a_i, x_i, w_i) occurs, then u_i must be in the inverse image set

$$\mathcal{U}(a_i, x_i, w_i, \sigma) \equiv \{u_i : \sigma(x_{it}, w_{it}, u_{it}) = a_{it}, \forall t\} \quad (11)$$

The condition

$$\{u_i \in \mathcal{U}(a_i, x_i, w_i, \sigma)\} \quad (12)$$

is then a necessary condition for the observed event (a_i, x_i, w_i) . If the model is incomplete, however, that condition is not sufficient for the event: when the exogenous variables (w_i, u_i) do not uniquely pin down the endogenous variables (x_i, a_i) , it can happen that (12) is satisfied but the event (a_i, x_i, w_i) does not occur. Then, Corollary 1 in CR gives the following characterization of the identified set for the policy and parameters for the unobservables, (σ, θ_u) : a pair (σ, θ_u) is in the identified set if and only if for all closed sets $\mathcal{S} \subset \mathbb{U}$ and for all z

$$\Pr(\mathcal{U}(a_i, x_i, w_i, \sigma) \subseteq \mathcal{S} | z) \leq \Phi(\mathcal{S}; \theta_u). \quad (13)$$

In this last equation, the left-hand side is the conditional probability of the outcomes $y_i = (a_i, x_i)$, which, according to σ , have $\{u_i : u_i \in \mathcal{S}\}$ as a necessary condition. For a given σ and z , this probability is observed in the data. The right-hand side is the probability of that necessary condition wrt the distribution of u_i , which by assumption does not depend on z . Further, given θ_u , this term is known and can be computed in closed form or via simulation.¹⁴

In the case with discrete a_{it}, x_{it}, w_{it} , CR shows that to obtain the sharp identified set for

¹²See Anderson and Hsiao (1981), Arellano and Bond (1991) and Blundell and Bond (1998), among others. Honoré and Tamer (2006) emphasize how the initial conditions problem leads to partial identification in nonlinear dynamic panel data models.

¹³Collard-Wexler (2014) employs both solutions. While the results of his counterfactuals are robust, a few parameter estimates vary substantially across the two methods, suggesting that the way in which the initial conditions problem is addressed matters in general.

¹⁴See Berry (1992) and Ciliberto and Tamer (2009) for other uses of simulation in models characterized by moment conditions.

θ , one only needs to check sets \mathcal{S} (labeled “core-determining”) that belong to a collection $Q(\sigma, z_i)$. This collection includes the “elemental” sets, $\mathcal{U}(a_i, x_i, w_i, \sigma)$, associated with individual realizations of the observables, as well as unions of overlapping sets of that form, excluding cases of strict subsets.

The CR approach operates in the space of the unobservables and it builds on earlier results that apply to the space of observables (Beresteanu et al. (2011), Galichon and Henry (2011)). For us, the result is useful because it completely characterizes the inequality restrictions that define the sharply identified set of policy functions. We illustrate this in the simple example of Section 3. However, the number of these restrictions can grow quite large in realistic problems. In these cases, one may not be able to list all the CR inequalities needed to obtain the sharp identified set. However, the CR characterization is still helpful to build intuition for selecting which inequalities to impose. We do this in the empirical application of Section 5 and show that, while we do not get sharp identification, the results are still informative.

When the necessary conditions (13) are actually *necessary and sufficient* for particular (sets of) actions, then the associated inequalities become strict equalities. In a complete model, all of the necessary conditions are equalities. However, as usual, this does not guarantee that the parameters are point-identified (since, e.g., the instruments might not be strong enough) and so, in the absence of a proof of point-identification, we might still want to consider set-identification.

To formalize the argument above, we define the set of policy functions that are identified exclusively by the IV conditions and the data, with no use of the dynamic model. In particular, for a given θ_u and a given data generating process, we define

$$\Sigma^{IV}(\theta_u) \equiv \{\sigma : \text{condition (13) holds } \forall \mathcal{S} \in Q(\sigma, z) \text{ and } \forall z\} \cap \mathcal{F} \quad (14)$$

We emphasize that $\Sigma^{IV}(\theta_u)$ is a subset of the space of admissible policies \mathcal{F} and, as such, incorporates all the natural economic restrictions—e.g., monotonicity—that one may be willing to impose on \mathcal{F} . Definition (14) immediately gives the following characterization of the sharp identified set for (σ, θ_u) .

Theorem 1. *The sharp identified set for (σ, θ_u) is given by $\{(\sigma, \theta_u) : \sigma \in \Sigma^{IV}(\theta_u), \Sigma^{IV}(\theta_u) \neq \emptyset\}$.*

2.5 Second Step: Identification of the Profit Parameters

We now show how one can use the sharp identified set for (σ, θ_u) from the first step to characterize the identified set for the structural dynamic parameters entering the profit function. These results can be viewed as a generalization of the approach in Hotz et al. (1994) (HMSS). As noted, for any $\theta = (\theta_\pi, \theta_u)$, we can use the Bellman equation to compute

the implied policy σ_θ , defined as follows:

$$\sigma_\theta(x_{it}, w_{it}, u_{it}) \equiv \operatorname{argmax}_{a_{it} \in \mathcal{A}(x_{it})} \left(\pi(a_{it}, x_{it}, w_{it}, u_{it}, \theta_\pi) + \delta E_{\theta_u} [V(x_{it+1}, w_{it+1}, u_{it+1}) | a_{it}, x_{it}, w_{it}, u_{it}] \right).$$

The sharply identified set of parameters is then given by all (θ_π, θ_u) pairs whose associated policy function is not rejected by the GIV restrictions. We formalize this in the following result.

Theorem 2. *The identified set for the structural parameters (θ_π, θ_u) is given by*

$$\Theta_{ID} \equiv \{\theta = (\theta_\pi, \theta_u) : \sigma_\theta(x_{it}, w_{it}, u_{it}) \in \Sigma^{IV}(\theta_u)\}. \quad (15)$$

Note that Theorem 2 imposes both the dynamic model and the GIV restrictions. This is the sharp identified set because any θ in this set generates a policy function that cannot be rejected by the data plus the sharp GIV conditions.

A natural question that arises is how to recover the profit parameters θ_π from any given (σ, θ_u) pair. This question has been investigated extensively for models without serial correlation in the unobservables. Here we extend some of those methods to the case with serial correlation. In particular, we show that, under certain conditions, the parameter θ_π can be conveniently recovered by solving a system of linear equations. To this end, we define the value of taking action a as follows:

$$v(a, x, w, u; \sigma, \theta) = \pi(a, x, w, u; \theta_\pi) + E_{\theta_u} \left[\sum_{t=1}^{\infty} \delta^t \pi(\sigma(x_t, w_t, u_t), x_t, w_t, u_t; \theta_\pi) \middle| a, x, w, u \right], \quad (16)$$

Given that the true policy σ is optimal, it must be the case that

$$\sigma(x, w, u) = \operatorname{arg max}_a v(a, x, w, u; \sigma, \theta) \quad (17)$$

for every (x, w, u) . Thus, if a value θ_π is in the identified set, it must be that it solves (17) for some θ_u and some $\sigma \in \Sigma^{IV}(\theta_u)$. Note that, given a pair (σ, θ_u) from the first stage, verifying this condition is a static optimization problem and is therefore much easier than solving the associated Bellman equation.

In order to illustrate how the condition in (17) can be used in practice, we extend an argument made by HMSS for dynamic discrete choice models with i.i.d. unobservables. Our approach applies to a wide class of models with discrete actions and continuous, possibly

serially-correlated unobservables.¹⁵ More specifically, we use indifference conditions implied by (17) to write a system of linear equations in the profit parameters.

Fix a pair $(\tilde{\sigma}, \tilde{\theta}_u)$ from the first step and let $\tilde{\theta}_\pi$ denote a value for the profit parameters that is consistent with $(\tilde{\sigma}, \tilde{\theta}_u)$ given the model. For now, we do not assume that such a $\tilde{\theta}_\pi$ is unique; later, we will provide conditions that ensure it is unique. If there is no such $\tilde{\theta}_\pi$, then the model rejects the pair $(\tilde{\sigma}, \tilde{\theta}_u)$. Suppose that, given (x, w) and a pair of actions, a and a' , there exists a value of the unobservable, say $\tilde{u}(a, a', x, w)$, such that

$$v\left(a, x, w, \tilde{u}(a, a', x, w); \tilde{\sigma}, \left(\tilde{\theta}_\pi, \tilde{\theta}_u\right)\right) = v\left(a', x, w, \tilde{u}(a, a', x, w); \tilde{\sigma}, \left(\tilde{\theta}_\pi, \tilde{\theta}_u\right)\right) \quad (18)$$

We do not require that $\tilde{u}(a, a', x, w)$ be uniquely defined by (18), only that it exist.

Assumption 1. *The variables (a, x, w) take discrete values and for each (a, x, w) there is an action $a' \neq a$ such that there is at least one $\tilde{u}(a, a', x, w)$ satisfying the indifference condition in (18).*

Assumption 1 is high-level. The next assumption provides more primitive sufficient conditions.

Assumption 2. *(i) The variables (a, x, w) take discrete values; (ii) the support of U is connected; (iii) the action-specific value function v is continuous in u ; (iv) for each (x, w) , $\tilde{\sigma}(x, w, \cdot)$ takes at least two distinct values.*

Assumption 2(ii) is a standard support restriction; Assumption 2(iii) is also standard and can be verified using results in Stokey et al. (1989); Assumption 2(iv) can be directly verified by inspecting the $\tilde{\sigma}$ from the first step.

Lemma 1. *Assumption 2 implies Assumption 1.*

Proof. Fix any $(x, w) \in \mathbb{X} \times \mathbb{W}$. By Assumption 2(iv) and the definition of v , there are two actions, a', a'' , and two values of the unobservable, u', u'' , such that

$$\begin{aligned} v\left(a', x, w, u'; \tilde{\sigma}, \left(\tilde{\theta}_\pi, \tilde{\theta}_u\right)\right) &\geq v\left(a'', x, w, u'; \tilde{\sigma}, \left(\tilde{\theta}_\pi, \tilde{\theta}_u\right)\right) \\ v\left(a'', x, w, u''; \tilde{\sigma}, \left(\tilde{\theta}_\pi, \tilde{\theta}_u\right)\right) &\geq v\left(a', x, w, u''; \tilde{\sigma}, \left(\tilde{\theta}_\pi, \tilde{\theta}_u\right)\right) \end{aligned}$$

Define

$$d(u) \equiv v\left(a', x, w, u; \tilde{\sigma}, \left(\tilde{\theta}_\pi, \tilde{\theta}_u\right)\right) - v\left(a'', x, w, u; \tilde{\sigma}, \left(\tilde{\theta}_\pi, \tilde{\theta}_u\right)\right)$$

¹⁵See also Pesendorfer and Schmidt-Dengler (2010).

and note that $d(u') \geq 0 \geq d(u'')$. Thus, by the Intermediate Value Theorem for general metric spaces,¹⁶ there exists at least one u''' such that $d(u''') = 0$, which proves the claim. \square

Next, we show that these assumptions ensure that a minimal necessary condition for uniquely recovering θ_π is satisfied, i.e. that there be at least as many equations as unknowns. We consider the case in which profits are parameterized in an entirely flexible way with the elements of θ_π representing the (deterministic) single-period profits for each combination of (a, x, w) . Under more restrictive parameterization of profits, it may be possible to recover θ_π under weaker conditions.

Lemma 2. *Under Assumption 1, the parameters $\tilde{\theta}_\pi$ associated with the pair $(\tilde{\sigma}, \tilde{\theta}_u)$ from the first step satisfy a system of equations with at least as many equations as the cardinality of $\mathbb{A} \times \mathbb{X} \times \mathbb{W}$.*

Proof. The result follows immediately from the fact that, given Assumption 1, one can write at least as many equation of the form (18) as the cardinality of $\mathbb{A} \times \mathbb{X} \times \mathbb{W}$. \square

We now show that, when the single-period profits are linear in θ_π , the above yields a system of linear equations. The coefficients of this system are known given a candidate $(\tilde{\sigma}, \tilde{\theta}_u)$ from the first step and can be computed via forward-simulation as in BBL.

Assumption 3. *The single-period profit function is linear in θ_π .*

Lemma 3. *Fix a pair $(\tilde{\sigma}, \tilde{\theta}_u)$ from the first step. Under Assumptions 1 and 3, the parameters $\tilde{\theta}_\pi$ associated with $(\tilde{\sigma}, \tilde{\theta}_u)$ satisfy a system of linear equations, with at least as many equations as the cardinality of $\mathbb{A} \times \mathbb{X} \times \mathbb{W}$. Further, the coefficients of the system are known given the model, $(\tilde{\sigma}, \tilde{\theta}_u)$, and the transition functions in (4)-(5).*

Proof. HMSS show that when then the single-period profit is linear in θ_π , then the action-specific value function v is also linear in θ_π . This trivially extends to the case of serially correlated unobservables, so that we can write

$$v\left(a, x, w, u; \tilde{\sigma}, \left(\tilde{\theta}_\pi, \tilde{\theta}_u\right)\right) = h_0(a, x, w, u, \tilde{\sigma}, \tilde{\theta}_u) + h_1(a, x, w, u, \tilde{\sigma}, \tilde{\theta}_u)\tilde{\theta}_\pi,$$

for some known functions h_0, h_1 .

For each (a, x, w) , Assumption 1 then guarantees that we can write

$$\left[h_1(a, x, w, \tilde{u}(a, a', x, w), \tilde{\sigma}, \tilde{\theta}_u) - h_1(a', x, w, \tilde{u}(a, a', x, w), \tilde{\sigma}, \tilde{\theta}_u) \right] \tilde{\theta}_\pi = \tag{19}$$

$$h_0(a', x, w, \tilde{u}(a, a', x, w), \tilde{\sigma}, \tilde{\theta}_u) - h_0(a, x, w, \tilde{u}(a, a', x, w), \tilde{\sigma}, \tilde{\theta}_u)$$

for some $a' \neq a$. This establishes the result. \square

¹⁶See, e.g., Theorem 4.22 in Rudin (1976).

Given this linear system, we have two cases depending on whether a standard (and directly verifiable) rank condition is satisfied. If the rank condition holds, we can solve for the unique $\tilde{\theta}_\pi$ in closed form. Otherwise, we obtain multiple values of $\tilde{\theta}_\pi$ associated with the first-step parameters.

One might wonder whether the value(s) of θ_π obtained in the second step are always consistent with the Bellman equation. This, paired with sharpness of the identified set for (σ, θ_u) from the first step, would imply that the identified set for θ_π is also sharp. In the next section, we address this for our illustrative entry/exit example and show that indeed the two-step approach delivers sharp identification of the structural parameters.

The results above focus on the case where the actions and states are discrete. However, a similar logic applies to models with continuous actions as well. We illustrate this via the following example.

Example 1 (Continuous-choice stochastic accumulation). *Consider the “stochastic accumulation problem” from Pakes (1994) and Doraszelski and Pakes (2007). In this model, firm i chooses the level of investment a_{it} , a continuous variable, based on its current efficiency or quality x_{it} (often taken to be discrete). The distribution of efficiency or quality at $t + 1$ is assumed to be stochastically increasing in a_{it} . Let*

$$v(a, x, w, u; \theta) = \pi(a, x, w, u, \theta_\pi) + \delta E_{\theta_u} \left[\sum_k V(x' = k, w', u') \Gamma(x' = k | x, a) \Big| w, u \right], \quad (20)$$

where both the single-period profit and the transition $\Gamma(x' = k | x, a)$ are differentiable in a .

The optimal action then satisfies

$$\frac{\partial \pi(a, x, w, u, \theta_\pi)}{\partial a} + \delta E_{\theta_u} \left[\sum_k V(x' = k, w', u') \frac{\partial \Gamma(x' = k | x, a)}{\partial a} \Big| w, u \right] = 0. \quad (21)$$

Given a candidate $(\tilde{\sigma}, \tilde{\theta}_u)$, then we can reject a candidate $\tilde{\theta}_\pi$ unless

$$\frac{\partial \pi(\tilde{\sigma}(x, w, u), x, w, u, \tilde{\theta}_\pi)}{\partial a} + \delta E_{\tilde{\theta}_u} \left[\sum_k V(x' = k, w', u') \frac{\partial \Gamma(x' = k | x, \tilde{\sigma}(x, w, u))}{\partial a} \Big| w, u \right] = 0. \quad (22)$$

As in the discrete case, linearity of the profit function is inherited by the value function and forward simulation can be used to approximate the coefficients on $\tilde{\theta}_\pi$. Thus, equation (22) gives a continuum of linear equations that can be used to back out the profit parameters.

We conclude this section by noting that our proposed second step does not require the use of any inequality conditions. This is contrast to existing approaches, such as BBL, which

require considering perturbations of the policy function from the first stage and imposing the implied inequalities even in absence of serially correlated unobservables.

2.6 Inference

So far, we have focused on identification of the model. We now briefly discuss how to obtain confidence regions for the structural parameters. In doing so, we again follow the two steps outlined above.

The sharp identified set for the first-step parameters (σ, θ_u) is characterized by the inequalities (13). Therefore, estimates of the identified set and confidence regions for (σ, θ_u) can be obtained by applying methods from the by now large literature on moment inequalities models (see, e.g., Chernozhukov et al. (2007), Andrews and Soares (2010), Beresteanu et al. (2011), Galichon and Henry (2011), Andrews and Shi (2013), and Chernozhukov et al. (2013)). As pointed out by CR, one issue that often arises is that the number of inequalities characterizing the sharp identified set is large relative to the sample size, or even infinite. For example, in our empirical application, we consider the entry and exit patterns from a cross section of markets over twelve years. Since there can be zero, one or two firms in the market at any given point in time, the number of inequalities associated with just the “elemental” sets is $3^{12} = 531,441$. Fortunately, recent results provide some guidance on how to deal with the “many inequalities” case. References include Menzel (2009), Chernozhukov et al. (2018), and Andrews and Shi (2017). In our application, we use one of the bootstrap procedures proposed by Chernozhukov et al. (2018) to obtain valid confidence regions for the parameters. Roughly speaking, their approach provides an econometrically disciplined way of determining the subset of moment inequalities that are most informative about the parameter values.

The second step in our approach maps the first-step parameters into the primitive single-period profit parameters. As shown in Lemma 2, this map only depends on the model and does not involve the data. In particular, all that is needed is knowledge of the action-specific value functions given the first-step parameters. Standard forward-simulation methods can be used for this purpose, as illustrated in Section 3.2. With a large number of simulation draws, the error from the second stage will be negligible relative to sampling error. Alternatively, one could adjust the standard errors to account for any noise from the second step.

The two-step approach is not the only way to conduct inference. Another option is to solve the model for each candidate value of the structural parameter θ and verify whether the implied policy function satisfies the GIV restrictions. This full-solution method may be viable in some cases (e.g., when the states take a small number of values), but typically would require one to parameterize the profit function. A third approach—often used in the empirical literature—parameterizes both the policy function and the profit function. This “double parameterization” has the undesirable feature that the functional form used

for the policy might be inconsistent with the profit parameterization in the sense that no choice of the structural parameters leads—via the model—to the chosen functional form for the policy. Thus, in the numerical illustration of Section 3.3, as well as in the empirical application (Section 5), we focus on the first two approaches: the two-step procedure and the full-solution method.

3 A Simple Example

In this section, we consider a minimal single-agent model that illustrates identification via GIV restrictions. In the example, the state is whether a firm is “In” or “Out” of the market in the prior period, $x_{it} \in \{0, 1\}$, and the action today is whether to be active in the market today, $a_{it} \in \{0, 1\}$. Exit is reversible and there are no exogenous profit shifters w .

A firm that is already in the market ($x_{it} = 1$) and decides to stay in ($a_{it} = 1$) earns a single-period profit equal to $\pi - \epsilon_{it}$, while a firm that is out of the market ($x_{it} = 0$) and decides to enter ($a_{it} = 1$) earns the same single-period profit minus a sunk cost γ . Whenever a firm decides to be inactive ($a_{it} = 0$), it earns zero profits. We interpret ϵ_{it} as a shock reflecting variation in per-period fixed costs and assume that it follows a first-order autocorrelation process,

$$\epsilon_{it} = \rho\epsilon_{i,t-1} + \nu_{it}\sqrt{1 - \rho^2}, \quad (23)$$

where ν_{it} is distributed standard normal. The resulting model then has three structural parameters: π , γ , and ρ . The policy function that generates the data is

$$a_{it} = \sigma(x_{it}, u_{it}), \quad (24)$$


where $u_{it} \sim \text{Unif}(0, 1)$ can be normalized to be the quantile of ϵ_{it} .¹⁷ We assume that the dynamic model generates the natural monotonicity results that σ is weakly increasing in x_{it} and weakly decreasing in u_{it} .

3.1 First step

The first step consists in characterizing the identified set for (σ, ρ) . If we focus on only one period of data, the policy function in (24) is a nonparametric binary choice model with endogeneity and monotonicity restrictions, similar to Chesher (2010). Given monotonicity in u_{it} , the policy function is fully described by two policy cutoffs, $\tau(x)$, for $x \in \{0, 1\}$, as illustrated in Figure 1.

¹⁷To see this, one can write $a_{it} = \bar{\sigma}(x_{it}, F_\epsilon^{-1}(F_\epsilon(\epsilon_{it})))$, where $\bar{\sigma}$ is a nonparametric function and F_ϵ is the cdf of ϵ_{it} . Equation (24) then follows by defining $u_{it} = F_\epsilon(\epsilon_{it})$ and $\sigma(x_{it}, \cdot) = \bar{\sigma}(x_{it}, F_\epsilon^{-1}(\cdot))$.

Figure 1: Policy Cutoffs in the One Period Case



As in Manski (1988) Chesher (2010), even one period of data will generate nontrivial bounds on the policy function. As a simple example of GIV restrictions, Table 3 illustrates the inverse image sets associated with the example (in column 2) as well as the inequalities implied by the GIV restrictions (in the last columns of the table). Note that [1] there are

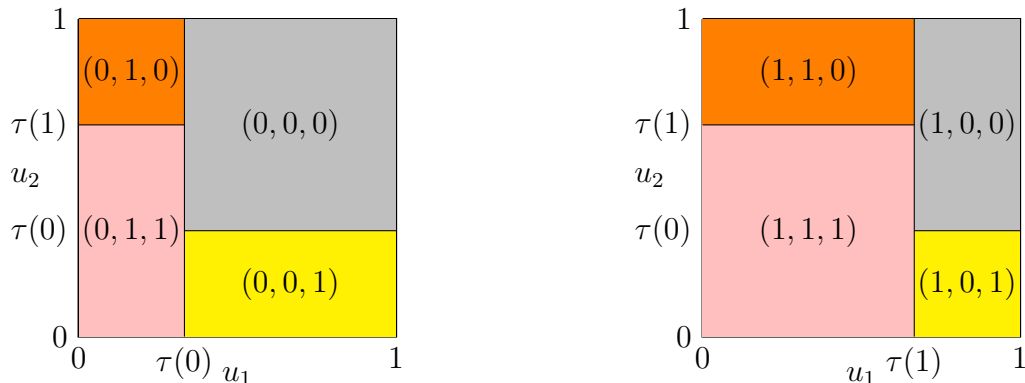
Table 3: Inverse Image Sets & Inequalities for the One-Period Example

a	x	$\mathcal{S} = \mathcal{U}(a, x, \sigma)$	$\Pr(\mathcal{U}(a_i, x_i, \sigma) \subseteq \mathcal{S} z)$	$\leq \Phi(\mathcal{S}; \theta_u)$
1	1	$(0, \tau(1))$	$\Pr((1, 1) z) + \Pr((1, 0) z)$	$\leq \tau(1)$
1	0	$(0, \tau(0))$	$\Pr((1, 0) z)$	$\leq \tau(0)$
0	1	$(\tau(1), 1)$	$\Pr((0, 1) z)$	$\leq 1 - \tau(1)$
0	0	$(\tau(0), 1)$	$\Pr((0, 0) z) + \Pr((0, 1) z)$	$\leq 1 - \tau(0)$

nontrivial bounds even in the absence of IVs, but [2] instrumental variable variation is helpful to tighten those bounds. Note also that, by themselves, the restrictions in Table 3 place no restrictions on θ_u . It is not surprising that it is impossible, in the example, to learn anything about serial correlation from restrictions on the single-period policy function. However, with multiple periods of data, restrictions on the policy function may rule out some values of serial correlation, even without reference to the structural model.

We illustrate this by considering two periods of data. Now there are eight elemental inverse image sets $\mathcal{U}(a_i, x_i, \sigma)$, in the space of (u_{i1}, u_{i2}) , that depend on $(x_{i1}, a_{i1} = x_{i2}, a_{i2})$. These are illustrated in Figure 2. The left panel gives the four elemental sets associated with the initial condition $x_{i1} = 0$, while the right panel gives the sets associated with $x_{i1} = 1$. For a given initial condition, the model is complete (the sets do not overlap), but across initial conditions the sets do overlap, reflecting incompleteness. For example, there are values of (u_{i1}, u_{i2}) that are consistent with both the sequence $(1, 1, 1)$ and the sequence $(0, 0, 0)$. If the initial x_{i1} was exogenous, the model would be complete.

Figure 2: Elemental Inverse Image Sets Labeled as (x_{i1}, a_{i1}, a_{i2})



The sets in Figure 2 allow us to build some intuition about identification in this class of models.¹⁸ Recall that the probability of each of the eight events associated with different (x_{i1}, a_{i1}, a_{i2}) must be less than the probability weight placed by the distribution of (u_{i1}, u_{i2}) over the regions of the elemental sets. In Figure 2, the joint density of (u_{i1}, u_{i2}) , which varies with the serial correlation parameter ρ , places the relevant probability weight over the various regions. Note that in this example, with two time periods, we can rule out some values of ρ without any use of the dynamic model. For example, perfect correlation, $\rho = 1$, collapses the joint density down to a straight line across the diagonal of each box. If in large samples we observe the events $(0, 1, 0)$ or $(1, 0, 1)$ conditional on any value of the instrument, then we can reject $\rho = 1$ since the associated GIV inequalities of the form (13) have a positive left-hand side and the right-hand side equal to zero.

As another piece of intuition, consider an instrument associated with a probability equal to one for initial condition $x_{i1} = 1$. The event probabilities associated with the right-hand side panel of Figure 2 then sum to one and all of the associated inequality restrictions hold with equality. These equalities are exactly the same as those that would be implied by maximum likelihood applied to the model with an exogenous initial condition $x_{i1} = 1$. Thus, if MLE point-identifies the parameters $(\tau(0), \tau(1), \rho)$, then GIV identifies the same parameter values in this special case.

3.2 Second Step

Next, we apply Lemma 3 to recover the structural parameters (π, γ) given a pair (σ, ρ) from the first step. For each value of $x \in \{0, 1\}$, firms are indifferent between being in and out of

¹⁸The core-determining collection of sets also includes unions of partially overlapping elemental sets, which we show in Appendix B.

the market when $u_{it} = \tau(x)$. This gives two indifference conditions involving action-specific value functions of the form (18).¹⁹ Since the action-specific value functions do not have a closed form expression, we show how to approximate them via forward-simulation. For any a and x ,

$$v^s(a, x, \tau(x), \sigma, \theta) = \pi \frac{1}{S} \sum_{s=1}^S \sum_{t=0}^{\bar{T}} \delta^t a_t^s - \gamma \frac{1}{S} \sum_{s=1}^S \sum_{t=0}^{\bar{T}} \delta^t a_t^s 1\{x_t^s = 0\} - \sum_{s=1}^S \sum_{t=0}^{\bar{T}} \delta^t \epsilon_t^s a_t^s,$$

where: (i) ϵ_t^s is set to the $\tau(x)$ -th quantile of ϵ_{it} for $t = 0$ and for $t \geq 1$ is drawn using (23) and the correlation ρ from step 1, (ii) x_0^s and a_0^s are set to x and a , respectively, (iii) a_t^s for $t \geq 1$ is determined by the policy σ from step 1, and (iv) S, \bar{T} are large numbers. The two equalities $v^s(0, x, \tau(x), \sigma, \theta) = v^s(1, x, \tau(x), \sigma, \theta)$ for $x \in \{0, 1\}$ then give a system of linear equations in (π, γ) .

As mentioned above, one might wonder whether the sharpness of the identified set for (σ, ρ) from the first step is inherited by the identified for θ_π from the second step. To address this, we show that any value π obtained in the second step is consistent with the model's Bellman equation. Thus, if the identified set for (σ, ρ) from the first step is sharp, the resulting identified set for π from the second step is also sharp.

Lemma 4. *Let $(\tilde{\sigma}, \tilde{\rho})$ be any pair from the first step with $\tilde{\sigma}$ weakly decreasing in u and $\tilde{\rho} \geq 0$, and let $\tilde{\pi}$ be any value of π returned by the second step. Then,*

$$\tilde{\sigma}(x, u) = \arg \max_a v(a, x, u; \tilde{\sigma}, \tilde{\pi}, \tilde{\rho}) \tag{25}$$

for all x, u , i.e. the policy $\tilde{\sigma}$ from the first stage solves the Bellman equation associated with $\tilde{\pi}$ and $\tilde{\rho}$.

Proof. See Appendix E. □

3.3 Numerical Illustration

We now compute the identified set for the structural parameters in an instance of the entry model described above. We will pay special attention to how the identified set changes with the number of time periods, the strength of the instruments, and the presence of exogenous profit shifters. In order to abstract from sampling error and focus on the shape of the identified set, we draw a large number of cross-sectional markets (50,000).

We set the deterministic profit parameter π to 0.5, the sunk cost γ to 1.5, and the correlation parameter ρ to 0.75, so that there is persistent unobserved heterogeneity. We generate

¹⁹In order to fully map our entry example into Lemma 3, note that we implicitly have used the two additional restrictions that single-period profits are zero when firms decide to be out. Thus, we have in total four equations corresponding to the cardinality of $\mathbb{A} \times \mathbb{X}$.

time-invariant excluded instruments z taking the values $\{0, 1\}$ with equal probabilities and set $x_{i1} = z_i$ for a fraction of markets in the data equal to 0.50 or 0.75. We call this fraction “IV strength” and note that it is equal to the square root of the R^2 coefficient in the regression of the endogenous state x_{i1} on the IV (plus a constant).²⁰ We compute the three-dimensional identified set for (π, γ, ρ) and we plot its projections onto the space of profits $(\pi, \pi - \gamma)$ and the space of sunk cost and correlation parameters. Note that, since we are treating both the policy and the deterministic profit function fully flexibly, the full-solution method and the two-step approach give the same identified set.

First, we consider how the identified set varies with the IV strength as well as the number of time periods ($T = 2$ and $T = 10$). In the case with $T = 2$, we are able to list all of the GIV restrictions implied by the model and thus obtain sharpness of the identified set, as outlined in Section 3.1. On the other hand, with $T = 10$, the number of inequalities in the GIV core-determining class becomes very large. So, instead of listing all the inequalities, we use those corresponding to the sharp two-period GIV identified set as well as those associated with several observable events over the ten time periods.²¹ We pick events that intuitively should help us shrink the identified set. Specifically, we use the events “the firm enters at least once,” “the firm exits at least once,” “the firm enters at least once and exits at least once,” and “the number of firms in the market, x_{it} , does not change for at least six consecutive periods.” To build some intuition, consider the latter event. We would expect this to help rule out values of ρ close to zero, since it yields inequalities where the sample probabilities on the left-hand side are large (the data exhibits a lot of persistence given that we set $\rho = 0.75$) and the model necessary conditions on the right-hand side are relatively small (when ρ is close to zero the model predicts little persistence in the observables). A similar argument applies to the other events we include. Figure 3 show that, when IV strength is low and $T = 2$, the identified set is quite large. On the other hand, as expected, the set shrinks considerably as the number of time period grows or the IV becomes stronger (Figure 4).

As a comparison, we report estimates obtained via two standard methods (MLE and GMM) that assume away serial correlation in the unobservables, consistent with most of the existing literature. In the MLE approach, we pool all observations along the cross-section and time-series dimensions and maximize the resulting one-period likelihood. In contrast, in the GMM approach, we use moments based on the two-period and three-period transitions as well, while still restricting ρ to be zero. Both methods—and particularly GMM—tend to overestimate the sunk cost. Intuitively, a model that assumes no serial correlation in the unobservables will load all the persistence in the data onto the sunk cost, thus overestimating its magnitude. In addition, Figures 7 and 8 in Appendix B.2 show that MLE and GMM

²⁰Specifically, IV strength = 0.50 corresponds to $R^2 = 0.25$ and IV strength = 0.75 corresponds to $R^2 = 0.56$.

²¹Since it is hard to obtain a closed form for the model probabilities associated with these events, we approximate the probabilities via simulation.

with $\rho = 0$ tend to underestimate the profit parameters.

Figure 3: Low IV strength

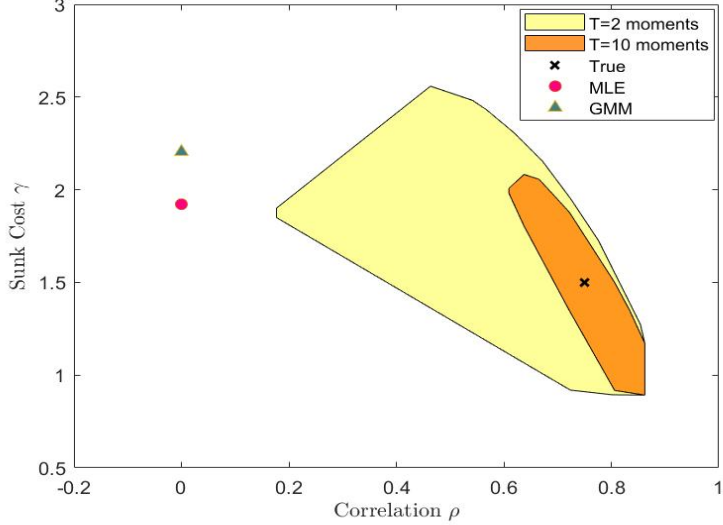



Figure 4: High IV strength



Next, we explore the impact of adding an exogenous covariate w_{it} to the model. We specify the firm per-period profit as follows:

$$\pi_{it} = \begin{cases} \alpha w_{it} - \beta - \epsilon_{it} & \text{if } a_{it} = 1, x_{it} = 1 \\ \alpha w_{it} - \beta - \gamma - \epsilon_{it} & \text{if } a_{it} = 1, x_{it} = 0 \\ 0 & \text{if } a_{it} = 0 \end{cases}$$

so that w_{it} could be interpreted as a measure of market size such as population, αw_{it} represents variable profits, β is fixed costs, and γ is again the sunk cost of entry. Further, we let ϵ_{it} follow the AR(1) process in (23), with ν_{it} distributed $\mathcal{N}(0, \sigma_\nu^2)$. Following Pakes et al. (2007), we assume that the term $\alpha w_{it} - \beta$ has already been estimated outside the dynamic model and we focus instead on the parameters γ, ρ and σ .²² When generating the data, we set $\alpha = 1.5, \beta = 1, \gamma = 1.5, \rho = 0.75, \sigma = 1$, and we focus on the case where IV strength is low. Regarding the distribution of the covariate, we let the initial w_{i0} for each market take the values $\{0.15, 1.00, 1.65\}$ with equal probabilities and evolve according to the transition matrix

$$\begin{bmatrix} 0.6 & 0.2 & 0.2 \\ 0.2 & 0.6 & 0.2 \\ 0.2 & 0.2 & 0.6 \end{bmatrix}$$


Note that the parameter values are chosen in such a way that the single-period profit function when $w_{it} = 1$ is the same as in the model without w_{it} .

Because w_{it} is exogenous, it can be used as an additional conditioning variable in the GIV inequalities along with the excluded IV. Thus, adding w_{it} to the model increases the number of inequalities that each candidate parameter value must satisfy in order to be included in the identified set. We would then expect the identified set to be smaller in the model with the exogenous covariate.²³ As shown in Figure 5, this is indeed the case.

²²Note that, unlike the case without w_{it} , here we are not allowed to normalize σ since the non-stochastic part of the period-profits is modeled parametrically. Accordingly, for this design, we report the standardized sunk cost $\frac{\gamma}{\sigma}$ to ensure comparability with the plots for the case without w .

²³However, note that, because w_{it} varies over time in the data, the policy function also changes relative to the “no w_{it} ” case (even for the value of w_{it} that makes the per-period profit identical to that in the “no w_{it} ” model). Thus, it need not be the case that the identified set in this section is a subset of that in the “no w_{it} ” model.

Figure 5: Effect of exogenous covariate



Finally, we illustrate how the identified sets translate into bounds on counterfactual quantities. For brevity, we focus on the case with $T = 2$, strong IV and no exogenous covariate w_{it} . We consider three counterfactual scenarios: (i) an increase in the sunk cost of entry by 4.00—corresponding to 267% of its true value—which we call “the sunk cost counterfactual” for brevity; (ii) an increase in the fixed cost by 0.50—corresponding to 50% of its true value—and a simultaneous decrease in the sunk cost by 0.125—or 8% of its true value—which we call “the fixed cost counterfactual”; (iii) a 1.25 subsidy to entry—corresponding to 83% of the true value of the sunk cost—which we call “the subsidy counterfactual.” The sunk cost counterfactual is meant to simulate a policy, such as environmental regulation, that only constrains new entrants. On the other hand, the fixed cost counterfactual corresponds to a policy restricting both incumbents and entrants. We simultaneously decrease the sunk cost in this counterfactual to reflect the fact that complying with the regulation may be easier for new entrants than for incumbents (e.g., due to retrofitting costs). Finally, the subsidy counterfactual mimics a policy encouraging entry of new firms that might be using cleaner or otherwise better technology.

The procedure we employ to assess the impact of these shocks is as follows. For each market in the data, we draw many time series for the unobservables and look at how the number of firms as well as the fraction of new firms—defined as firms that enter after the

policy change—evolve 10 years after each of the three counterfactual changes.²⁴ We then average across markets as well as realizations of the unobservables.

Tables 4 and 5 show the results for the GIV approach as well as the MLE and GMM models with $\rho = 0$. One can see that assuming away serial correlation in the unobservables leads to bias in the estimated reaction to the policy changes, in terms of both the number of firms and the fraction of new firms. The effects of the policies tend to be overstated, in particular for the fixed cost and the subsidy counterfactuals. Intuitively, when $\rho = 0$ the unobservables drawn in the counterfactuals will exhibit much less persistence relative to their true distribution. This, in turn, leads to an excessive amount of predicted entry and exit.

Table 4: Number of firms

	True	GIV	MLE $\rho = 0$	GMM $\rho = 0$
Baseline	0.659	(0.612, 0.701)	0.817	0.763
Sunk cost (Δ)	0.059	(0.017, 0.123)	0.012	-0.154
Fixed cost (Δ)	-0.167	(-0.289, -0.139)	-0.769	-0.700
Subsidy (Δ)	-0.023	(-0.052, -0.011)	-0.167	-0.140

Note: The first row shows the average number of firms in the market in the absence of policy changes. Rows 2-4 show the change—relative to the first row—in the average number of firms in the market 10 years after the policy change (an increase in the sunk cost, a simultaneous increase in the fixed cost and decrease in the sunk cost, and a subsidy to entry, respectively). The dgp has $T = 2$, strong IV and no w_{it} covariate.

²⁴This requires choosing an initial value for the unobservable. For each market, we compute the bounds on u_0 implied by (x_0, a_0) and the model, and repeat the exercise twice setting u_0 equal to the upper and lower bounds, respectively. We then take the convex hull of the two resulting outcomes.

Table 5: Fraction of new firms

	True	GIV	MLE $\rho = 0$	GMM $\rho = 0$
Baseline	0.591	(0.460, 0.781)	0.574	0.516
Sunk cost (Δ)	-0.193	(-0.349, -0.098)	-0.291	-0.270
Fixed cost (Δ)	0.186	(0.102, 0.288)	0.429	0.426
Subsidy (Δ)	0.124	(0.066, 0.209)	0.366	0.379

Note: The first row shows the average fraction of new firms in the market in the absence of policy changes. Rows 2-4 show the change—relative to the first row—in the average fraction of new firms in the market 10 years after the policy change (an increase in the sunk cost, a simultaneous increase in the fixed cost and decrease in the sunk cost, and a subsidy to entry, respectively). The dgp has $T = 2$, strong IV and no w_{it} covariate.

4 Dynamic Oligopoly

We now extend the analysis to the setting where multiple firms interact in each market. In the oligopoly case, each firm’s equilibrium policy is its single-agent best reply to its rivals’ equilibrium strategies. The firm still solves a value function problem similar to (7), but its expectations of the future evolution of endogenous market states depend on its action as well as the equilibrium actions of its rivals.

Throughout this section, we assume complete information for the serially correlated components of unobservables. If the serially correlated unobservables are not common knowledge across all players, standard equilibrium concepts such as perfect Bayesian equilibrium often become intractable in that they imply that the entire history of play enters the current state. Fershtman and Pakes (2012) propose a tractable framework to deal with persistent sources of asymmetric information and our methods might be usefully merged with theirs. In Appendix C, we illustrate a computed oligopoly problem that features both [i] serially correlated errors that are observed by all the firms and [ii] private information shocks that are independent over time. The combination of observed (by rivals) errors and private information errors is reminiscent of the discussion in Pakes, Porter, Ho, and Ishii (2015).

Since there are multiple firms per market, here we require our original notation of i for the market and j for the firm. If the equilibrium policies of firm j ’s rivals are given by the function σ_{-j} , then the firm’s expected equilibrium state transition probabilities are given by

$$\tilde{\Gamma}_j(x_{it+1}|a_{ijt}, x_{it}, w_{it}, \sigma_{-j}(x_{it}, w_{it}, u_{it})). \quad (26)$$

This notation allows for a rich set of possible state transitions models, including oligopoly variations on our earlier single-firm examples.

Firm j ’s equilibrium Bellman equation then depends on the equilibrium strategies of its

rivals:

$$V_j(x_{it}, w_{it}, u_{it}, \sigma_{-j}) = \quad (27)$$

$$\max_{a_{ijt} \in \mathcal{A}(x_{ijt})} (\pi_j(a_{it}, x_{it}, w_{it}, u_{it}; \theta_\pi) + \delta E_{\theta_u} [V_j(x_{it+1}, w_{it+1}, u_{it+1}, \sigma_{-j}) | a_{it}, x_{it}, w_{it}, u_{it}]).$$

The expected Bellman's equation is

$$E[V_j(x', w', u') | a, x, w, u, \sigma_{-j}] = \int \int \int V_j(x', w', u', \sigma_{-j}) d\tilde{\Gamma}(x' | a_j, x, w, \sigma_{-j}(x, w, u)) dQ(w' | w) d\tilde{\Phi}(u' | u; \theta_u).$$

Associated with this dynamic program is a best response strategy for firm j , which we assume is unique,²⁵ denoted by $\bar{\sigma}_j(\sigma_{-j}, \theta)$. The vector of best response strategies is then the J -vector

$$\bar{\sigma}(\sigma, \theta) = (\bar{\sigma}_1(\sigma_{-1}, \theta), \dots, \bar{\sigma}_J(\sigma_{-J}, \theta)).$$

Any vector of equilibrium strategies, σ^* , must satisfy the fixed point

$$\sigma^* = \bar{\sigma}(\sigma^*, \theta). \quad (28)$$

We can then define the set of possible equilibrium policy function vectors as

$$\Sigma^{EQ}(\theta) = \{\sigma^* : \sigma^* = \bar{\sigma}(\sigma^*, \theta)\}.$$

We adopt the same approach as in earlier papers and assume that, even if the underlying model admits multiple equilibria, the firms themselves always play the same policy function when at the same state vectors.²⁶ The true policy function that generates the data is then an element of the set $\Sigma^{EQ}(\theta)$, where θ is the true parameter that generates our data.

The sharply identified set of parameters in the oligopoly case is the same as in the single agent case, except with the further restriction that the policies associated with θ are a vector of equilibrium policies:

$$\Theta_{ID} \equiv \{\theta = (\theta_\pi, \theta_u) : \text{there exists } \sigma^* \in \Sigma^{EQ}(\theta) \text{ such that } \sigma^* \in \Sigma^{IV}(\theta_u)\}, \quad (29)$$

That is, a parameter vector θ is in the identified set if there is a policy vector that both [i] is not rejected by the IV restrictions and the data (given θ_u) and [ii] is a vector of equilibrium strategies given θ .

In practice, we may recover Θ_{ID} via a two-step procedure, just like in the single-agent

²⁵Again, Stokey et al. (1989) provide standard conditions guaranteeing uniqueness of the best response strategy. Note that we assume uniqueness of the best reply but not of the overall equilibrium strategy profile.

²⁶This approach was adopted approximately simultaneously in Bajari, Benkard, and Levin (2007), Pakes, Ostrovsky, and Berry (2007) and Pesendorfer and Schmidt-Dengler (2008).

case. First, note that the argument for the first step from Section 2.4 immediately extends to the oligopoly setting, since displays (11) to (14) continue to hold when the variables $a_{it}, x_{it}, w_{it}, u_{it}$ and the function σ are vector-valued. Thus, we can characterize the identified set for the policy vector $\Sigma^{IV}(\theta_u)$ using the GIV restrictions. Second, one can easily extend the dynamic best reply condition (17) to the oligopoly setting as follows:

$$\sigma_j(x, w, u) = \arg \max_{a_j} v_j(a_j, \sigma_{-j}(x, w, u), x, w, u; \sigma, \theta) \quad (30)$$

where v_j is the oligopoly analogue of (16) for firm j . Note that, in defining v_j for the oligopoly case, we treat the actions of rival firms as being generated by σ_{-j} (in addition to generating the future actions of firm j based on σ_j). As in the single-agent case, the best reply in (30) is therefore a static optimization problem and does not require solving any Bellman equations. Furthermore, it is a general condition that can be used to characterize the identified set in any problem. Specifically, the identified set for θ_π is characterized as the values of the profit parameters that solve (30) for some θ_u and some $\sigma \in \Sigma^{IV}(\theta_u)$.

As in the single-agent case, in many oligopoly cases we will be able to find simple indifference conditions that are necessary for the best reply equation in (30) to hold. In that case, the second-step search for the identified set of single-period profit function parameters may once again be characterized as the set of solutions to a system of linear equations. In the next section, we illustrate this via an oligopoly empirical application. The exact form of the indifference conditions yielding the system of linear equations for our application is given in Appendix D.

It may be useful to compare our oligopoly procedure to other two-step procedures in the literature. Our two-step procedure of the last paragraph is quite similar to Bajari, Benkard, and Levin (2007), with two differences. First, the policy functions are identified via GIV conditions. Second, we obtain the sharply identified set via the static best-reply condition (perhaps through the implied linear indifference conditions), whereas BBL suggest the use of various inequality conditions that are motivated by the same best reply condition.

We can also compare our approach to “full-solution” approaches that search across the set of “structural” parameters (θ_π, θ_u) , at each point solving for a (one hopes) unique equilibrium. We believe that our set-identification approach clarifies identification issues in a way that is often hard to do with full-solution approaches. There also a number of more practical differences. Full-solution methods become particularly hard in the case of possible multiple equilibria.²⁷ Once an equilibrium is calculated for a given parameter value, a full computational method must then compare the model’s predictions to data. If the model features a natural solution to the initial conditions problem, then the fit to data might be done via

²⁷See the online appendix of Doraszelski and Satterthwaite (2010) for examples of multiplicity. Further examples are in Pesendorfer and Schmidt-Dengler (2010). Borkovsky et al. (2010) provides a homotopy method for exploring a range of possible multiple equilibria.

maximum likelihood (as in, *e.g.*, Rust (1987) and Igami (2018)) or else via a fit of data moments to moments predicted from the model (Pakes and McGuire (2001)). In contrast, our method naturally accounts for unrestricted initial conditions and never requires a computed solution to the equilibrium fixed-point problem, or even a solution to the single-agent contraction mapping. In our method, the “fit to data” is provided by the GIV method. In the absence of an initial conditions problem, the GIV approach often collapses into an MLE or method of moments approach, reducing the difference between the methods.

The computational tradeoff is that our method requires us to find a confidence region for the identified set of policy functions. The nature of the tradeoff here may vary with the fine details of the problem and (in the absence of multiple equilibria issues) might favor either method. Note that we could, if computationally advantageous, also employ a full solution method. In this case, the “fit to data” for a computed equilibrium would involve testing the GIV conditions for the policy function implied by the equilibrium behavior. In our empirical example, we employ a full-solution GIV approach (assuming a unique equilibrium and allowing for an initial conditions problem) as well as the two-step approach that we have discussed in this section (which does not require a unique equilibrium).

5 Empirical Application to Environmental Policy Style Counterfactuals

In order to illustrate the approach, we apply it to the ready-mix concrete industry studied by Collard-Wexler (2014) (henceforth, CW). CW quantifies the magnitude of the sunk cost of entry in each of many isolated markets in the US and uses these estimates to assess how persistent the effects of a horizontal merger are in this industry. More specifically, CW first estimates the firms’ policy functions based on data on the number of ready-mix concrete plants and demand shifters. Given the policies, the paper then simulates the evolution of a market following a merger to monopoly and evaluates how long it takes for a second firm to enter.²⁸ CW imposes an intuitively appealing parametric form for the policy functions, rather than deriving them from an underlying dynamic model. One of our approaches below will roughly mimic this approach.

We use the same data and modeling framework as CW, but estimate all of the structural parameters as opposed to just the policy functions and the serial correlation parameter.²⁹ This allows us to address the counterfactual effects of policies that affect the “structural”

²⁸Lazarev et al. (2018) also consider counterfactual questions that only require knowledge of the dynamic policy functions, but not the structural single-period profit parameters.

²⁹Collard-Wexler (2013) estimates a full model of industry dynamics using firm-level data under the assumption that the dynamic states are econometrically exogenous. We use coarser market-level data, but address the endogeneity of market structure.

profit function. In particular, we consider the effects of policies—similar to the environmental policies in the cement industry study of Ryan (2012)—that alter sunk costs. In addition, since the GIV approach accommodates incomplete models, we are able to tackle the initial conditions problem in a flexible way.³⁰

Note that our application is intentionally simplified to serve as an example within a longer methodological paper. In particular, for this worked empirical example, we want to avoid the large, growing and important literature on methods that solve the computational challenge of methods that involve both (1) many parameters and (2) many inequality restrictions. Examples of this literature include Chen et al. (2018) and Kaido et al. (2019). To keep the number of parameters small, we use an intentionally simplified state space and, even then, we place further parametric restrictions. This gives us models with as few as five parameters so that we can easily apply a multidimensional grid search to compute the required confidence regions. Our restriction to problems amenable to a grid search is obviously strong given the state of the literature, but it removes complex computational choices from the list of problems we need to tackle.

Table 6 summarizes the variables we use. The data on number of plants and construction employment is the same as in CW and we refer the reader to that paper for more details. Briefly, the number of plants variable measures how many firms are active in each isolated town, while construction employment captures demand for ready-mix concrete. We follow CW and treat construction employment as exogenous, while the number of concrete plants is endogenous. In addition, we obtain data on past household income growth at the county level from the US Census website. We assume that past income growth is excluded from the current profits of concrete firms and that, conditional on current construction employment, past income growth is independent of within sample unobserved shocks to profitability. Past income growth therefore serves as an excluded instrument in our model.

Table 7 shows results for a “quasi first-stage regression.” This table presents an ordered probit model with the number of firms as the dependent variable. We see that the coefficients on both exogenous variables are positive and precisely estimated. The result suggests that our excluded instrument is “relevant,” even when conditioning on current demand. Of course, the true reduced form of the model is not an ordered probit and we present this merely as a descriptive result.

We conclude our brief descriptive analysis by reporting in Table 8 the transition probabilities for the number of plants, which shows that the variable exhibits substantial persistence

³⁰CW addresses the initial conditions problem by simulating the probabilities of the initial states via a modification of the GHK algorithm. This requires assuming that the industry has been following the same set of policies for a long time.

over time.³¹

Table 6: Summary statistics

Variable	Mean	St.Dev.	Min	Max
Number of plants	0.97	0.93	0.00	6.00
Construction Employment	519	819	3	17,772
% Household Income Growth 1969-1989	0.15	0.11	-0.16	0.69

Note: Fully balanced panel of 428 markets between 1994 and 2005.

Table 7: Ordered probit results

Log Construction Employment	0.14**
Income Growth 1969-1989	0.22**
Likelihood-Ratio Test p-value	0.00

Note: Dependent variable is number of plants. ** denotes significance at the 95% level.

³¹The results in Tables 7 and 8, as well as those from structural estimation, are based on a discretized version of the original data. Specifically, since more than 90% of the original observations at the market-year level have two or fewer plants, we censor the number of plants at two, which reduces the number of parameters to estimate in the structural model. Similarly, we discretize construction employment and household income growth. For each variable, we define a high and a low value depending on whether a given observation is above or below the median. Because construction employment essentially has no over-time variation after discretization, we take its value in the first year of the sample for each market and assume it is constant over time (and that firms know it). This means that, when we condition on construction employment to compute the moments implied by the model, the conditioning variable is scalar—as opposed to having dimension equal to the number of time periods—which leads to much more precisely estimated moments.

Table 8: Transition probabilities conditional on instrument

x_{it}	x_{it+1}		
	0	1	2
0	0.92	0.07	0.01
1	0.05	0.89	0.06
2	0.01	0.10	0.89

Note: Each row displays the probabilities of different values of x_{it+1} conditional on a value of x_{it} .

We now turn to the structural analysis. As in CW, we estimate a version of the Last-In First-Out model developed by Abbring and Campbell (2010). Again, we refer the reader to CW for more details on what assumptions are imposed and why this is a suitable oligopoly model in the context of the ready-mix concrete industry. Importantly, in our simple model, the Abbring-Campbell assumptions imply uniqueness of the equilibrium (see footnote 32) within the class of Last-In-First-Out equilibria. However, we also present results from a two-step approach that allows for multiple equilibria.

Some—but not all—of the results presented below rely on the following specification for the single-period profit function:

$$\pi_{it} = \begin{cases} \alpha_{x_{it}} w_i - \beta + \epsilon_{it} & \text{if was in at } t-1, \text{ stays in at } t \\ \alpha_{x_{it}} w_i - \beta - \gamma + \epsilon_{it} & \text{if was out at } t-1, \text{ enters at } t \\ 0 & \text{if is out at } t \end{cases} \quad (31)$$

where w_i denotes construction employment (in thousands) in market i , $\alpha_{x_{it}}$ is a coefficient that depends on the number x_{it} of active firms in market i at time t , β represents the intercept of the variable profit function as well as any fixed costs, γ is the sunk cost of entry, and ϵ_{it} denotes a potentially serially correlated unobservable shock to profitability. We assume that, conditional on ϵ_{it-1} , ϵ_{it} is equal to ϵ_{it-1} with probability ρ and is drawn uniformly from the $[-1, 1]$ interval with probability $1 - \rho$.³² Further, we impose the natural restriction $\alpha_2 \leq \frac{\alpha_1}{2}$, i.e. that per-firm variable profits (weakly) decrease with the number of competitors.

As mentioned above, because we have twelve years of data, the number of moment conditions defining the sharp identified set is extremely large and it is not practical to use

³²This parametric specification for the joint distribution of the unobservables is used by Abbring and Campbell (2010), who show that it satisfies their Assumption 3. This, along with other mild assumptions, ensures existence and uniqueness of a Markov-perfect equilibrium in Last-In First Out strategies. Note that we are able to map our model to that of Abbring and Campbell (2010), which features a scalar exogenous state variable, since the covariate w_i is time-invariant in our setting after the discretization described in footnote 31. If w_i also varied over time, then we would have two exogenous state variables— w_{it} and ϵ_{it} —and uniqueness of the equilibrium might no longer hold. However, note that even in this case our two-step approach would be able to (partially) identify the model parameters under the standard assumption that the same equilibrium is played at a given state in the data.

all of them. However, the goal of obtaining a small confidence region for the identified set suggests that we want to use a large number of them. This creates a potential problem, as many methods for obtaining confidence regions will not perform well when the number of moment inequalities is very large. A tension thus arises between the desire for a small confidence region (on one hand) and the desire for accurate inference (on the other hand). To resolve this tension, we report confidence sets for the structural parameters obtained via the two-step multiplier bootstrap approach proposed by Chernozhukov et al. (2018) (henceforth, CCK).³³ The method is intended for cases where the number of moment inequalities is very large, even larger than the dimension of the data.

As in the Monte Carlo simulations from Section 3, we impose two sets of GIV moments: (i) the full list of moments characterizing the sharp identified based on the first two time periods;³⁴ and (ii) several inequalities corresponding to observable events over the entire twelve years in the data.³⁵ This second set is chosen to include events that are intuitively likely to distinguish differing levels of serial correlation. For example, if the event “no entry or exit occurs in the twelve-year period” is very common, this might indicate a high degree of serial correlation. In addition to our list of intuitively information aggregated events, we include the inequalities associated with the complement of each of these events and we make sure that the left-hand side of each GIV inequality incorporates not only the observable event associated with that inequality, but also all observable events whose necessary conditions (in the space of unobservables) are subsets of the set whose probability is on the right hand side of the inequality.

Computationally, we follow CCK in generating a large number of draws to approximate the integrals corresponding to these events and ignoring the corresponding simulation error. Simple diagnostics suggest that the simulation variance is indeed negligible relative to the sample variance. Since all our conditioning variables are discrete, we can easily turn the conditional moment inequalities into unconditional ones. We obtain a total of with 614

³³For each candidate parameter value, we use as test statistic the maximum of the sample moments (at that parameter value) divided by their standard deviations and scaled by the square root of the sample size. We then compare the test statistic to the CCK critical value obtained with 250 bootstrap draws.

³⁴We compute these moments by only using the first two years in the panel. In principle, one could obtain more moments by taking all subsets of the data with two consecutive years. Because we already obtain informative results when only using the first two years, we do not pursue this extension in the paper.

³⁵These events are: “some entry and some exit occur,” “there are always zero firms,” “there is always one firm,” “there are always two firms,” “there is at least one period with zero firms,” “there is at least one period with one firm,” “there is at least one period with two firms,” “the number of firms goes from zero to one,” “the number of firms goes from one to two,” “the number of firms goes from zero to two,” “the number of firms goes from two to zero,” “the number of firms changes at least once,” “the number of firms changes exactly once,” “the number of firms changes exactly twice,” “the number of firms goes from one to zero,” “the number of firms goes from two to one or zero,” “the number of firms in the market is unchanged for at least five consecutive periods,” “the number of firms in the market is unchanged for at least nine consecutive periods.”

inequalities.

We estimate our GIV model in two ways. First, we use a traditional full-solution approach in which we solve the model for each candidate θ and verify whether the implied policy functions satisfy the GIV restrictions. This method requires a model with a unique equilibrium. Our model satisfies the unique equilibrium conditions of Abbring and Campbell if the exogenous shifters w_{it} are fixed over time (see footnote 32). This is a strong restriction that we maintain in our empirical specification. We also estimate a version of our two-step method, which does not require uniqueness of the equilibrium. As in many more elaborate two-step empirical exercises, to hold down the computation burden we impose a functional form on the policy function. In particular, we use a five-parameter specification consisting of four entry thresholds (corresponding to two values of x_{it} and two values of w_i) plus a shifter that, when added to the entry thresholds, gives the corresponding exit thresholds. This choice of parameterization fairly closely matches the policies that are obtained via the full-solution method. In addition, we impose the specification of single-period profits given in (31). This “double parameterization”—on both the policies and the single-period profits—mimics the double parameterized specifications that are common in the CCP literature. We also estimated a more flexible two-step model in which the single-period profits are unrestricted, i.e. each of the eight possible discrete values of $\pi(x, w, a)$ is treated as a separate parameter. Unsurprisingly given the limited sample size and variation in the data, this more flexible model gives wide confidence intervals for the counterfactuals. However, this might be a viable option when larger samples with more variation in the endogenous states are available.

In addition, we estimate three models that set $\rho = 0$ and thus assume away serial correlation. Two are full-solution methods corresponding to the MLE and GMM approaches used as benchmarks in the numerical illustration of Section 3.3. A third method mimics the standard two-step approaches with exogenous states. In particular, we use the same double parameterization as in the GIV two-step procedure described above, but we set $\rho = 0$ and estimate the policy in the first step by MLE.

Tables 9 and 10 display confidence intervals for the structural parameters in specification (31) based on the full-solution and two-step approaches, respectively. First, both GIV approaches give estimates of ρ that are positive and significantly different than zero, indicating substantial persistence in the unobservables over time. Second, the GMM approach with $\rho = 0$ considerably over-estimates the sunk cost γ . Intuitively, like in the simulations of Section 3.3, the GMM approach requires a very large sunk cost in order to rationalize the observed persistence in the number of firms over time since, with $\rho = 0$, sunk costs are the only source of stickiness in the model. On the other hand, the sunk cost is not over-estimated by the MLE model with $\rho = 0$. To understand this, recall that the MLE model pools all observations together and maximizes the single-period likelihood, thus effectively ignoring the time dimension of the data, whereas the GMM model is designed to match the time

transitions of the states.

Table 9: Full-solution parameter estimates

	Full-solution GIV	MLE $\rho = 0$	GMM $\rho = 0$
α_1	(0.17,0.18)	(0.042,0.045)	(0.10,0.20)
α_2	(0.010,0.011)	(0.01,0.02)	(-0.013,0.015)
β	(0.260,0.269)	(0.038,0.043)	(-0.04,-0.02)
γ	(2.582,2.585)	(1.44, 1.46)	(6.40,10.56)
ρ	(0.65,0.76)	—	—

Note: The results impose the specification for single-period profits in (31), but do not restrict the policy functions. The intervals have 95% confidence level.

Table 10: Two-step parameter estimates

	Two-step GIV	Two-step $\rho = 0$
α_1	(0.10,0.40)	(0.006,0.010)
α_2	(-0.0004,0.218)	(-0.005,-0.002)
β	(0.14,0.54)	(-0.036,-0.030)
γ	(-0.0001,1.766)	(1.62,1.68)
ρ	(0.74,0.87)	—

Note: The results impose the parameterization of the single-period profits in (31) as well as the five-parameter specification for the policy functions described in the text. The intervals have 95% confidence level.

In order to investigate whether allowing for endogeneity of market structure makes a difference for policy-relevant questions, we turn to counterfactual analysis. We consider the impact of an increase in the sunk cost on the number and composition of firms in the market.³⁶ This can be thought of as arising from environmental regulation, such as a mandate for new firms to invest in technology to reduce polluted water from running off concrete operations.

³⁶Specifically, we take the midpoint of the (projection of the) full-solution identified set for γ and increase the sunk cost by 10% of that value.

Tables 11 and 12 shows the counterfactual results. One can see that the approaches assuming away serial correlation in the unobservables yield significantly different predictions relative to the GIV methods. In particular, maximum likelihood with $\rho = 0$ tends to over-estimate the response to the counterfactual policy in terms of both the change in the number of firms and the decrease in the percentage of new firms (defined as firms that enter the market after the policy change). Intuitively, the i.i.d. assumption forces the unobservables to vary too much from one period to the next, which translates into excessive variation in the implied market outcomes relative to the model with serial correlation. Table 12 shows a similar pattern for the two-step methods. In contrast, the GMM approach with $\rho = 0$ predicts no change at all in the number or composition of firms. This is because GMM estimates a very large sunk cost, relative to which the policy change is second-order. It is also interesting to compare the two GIV methods. While the two-step approach delivers wider confidence intervals than the full-solution method, the latter confidence intervals are not contained in the former. This is because the two models are not nested: the two-step procedure is more restrictive in that it places a parametric structure on the policy functions, but on the other hand it is more flexible in that it allows for multiple equilibria.

In sum, these policy counterfactual results show the large bias that may result from models that artificially set the serial correlation parameter to zero, as is common in much of the literature. Further, the relatively precise estimates from our full-solution procedure show that dropping restrictive initial conditions assumptions need not lead to uninformative results. In part, this stems from our use of new econometric techniques for moment inequalities that allow us to employ very large number of moment inequalities in an econometrically disciplined and correct way.

Table 11: Counterfactual Outcomes: Full-solution approaches

	Full-solution GIV	MLE $\rho = 0$	GMM $\rho = 0$
# of Firms			
Baseline	(0.47,0.50)	(0.60,0.66)	(0.30,0.64)
Increase in Sunk Cost	(-0.032,-0.031)	(-0.15,-0.11)	0
% of New Firms			
Baseline	(7.7,9.7)	(27.6,32.0)	(0.7,0.9)
Increase in Sunk Cost	(-2.3,-2.2)	(-15.4,-13.7)	0

Note: For each outcome of interest (number of firms and percentage of new firms), the “baseline” numbers refer to the average outcomes in the absence of policy changes, whereas the “increase in sunk cost” numbers represent the change—relative to the baseline—in the counterfactual scenario where the sunk cost is higher. All intervals have 95% confidence level.

Table 12: Counterfactual Outcomes: Two-step approaches

	Two-step GIV	Two-step $\rho = 0$
# of Firms		
Baseline	(0.74,1.314)	(0.89,1.02)
Increase in Sunk Cost	(-0.18,-0.01)	(-0.22,-0.19)
% of New Firms		
Baseline	(15.2,38.9)	(22.1,27.9)
Increase in Sunk Cost	(-0.9,-0.1)	(-8.6,-6.8)

Note: For each outcome of interest (number of firms and percentage of new firms), the “baseline” numbers refer to the average outcomes in the absence of policy changes, whereas the “increase in sunk cost” numbers represent the change—relative to the baseline—in the counterfactual scenario where the sunk cost is higher. All intervals have 95% confidence level.

6 Conclusion

In this paper, we have proposed an approach to identification and inference in dynamic models with serially correlated unobservables. We tackle the resulting endogeneity of dynamic states by relying on the type of instrumental variables intuition that is commonly used in static models. In order to characterize the identified sets for quantities of interest and obtain confidence regions, we leverage recent results in the econometrics literature on partially identified models and the associated inference literature. Our empirical application extends work by Collard-Wexler on dynamic entry models with serially correlated unobservables to consider policy counterfactuals that are motivated by classic questions in environmental economics. We find that approaches ignoring serial correlation can significantly misstate the effects of policies that affect the underlying profitability of an industry, such as an environmental regulation that affects the sunk cost of entry.

This paper opens several avenues for future research. Most importantly, it would be interesting to apply the proposed approach to a wider class of empirical settings and see how accounting for the endogeneity of market structure affects additional counterfactual policy results.

Appendix A: Discrete Unobserved Heterogeneity

A.1 Set-up

In this section, we specialize our method to a single-agent problem with persistent unobserved discrete types. As a particularly simple example, we again consider the monopoly entry problem but now focus on the case in which firm heterogeneity is discrete, in the spirit of Heckman and Singer (1984), Keane and Wolpin (1997), Arcidiacono and Miller (2011) and related literature. In this case, the unobservable that enters the single-period profit function is sometimes modeled as

$$\epsilon_{it} = \lambda_{it} + \eta_{it}, \quad (32)$$

where η_{it} is continuously distributed and independent over t and i while λ_{it} is persistent and takes on a discrete set of values:

$$\lambda_{it} \in \{\bar{\lambda}_1, \bar{\lambda}_2, \dots, \bar{\lambda}_K\}.$$

For example, the single-period return to operating in market t in a single-agent dynamic entry model could be specified as

$$\pi(a_{it}, x_{it}, w_{it}; \theta_\pi) - (\lambda_{it} + \eta_{it}),$$

so that $(\lambda_{it} + \eta_{it})$ represents the fixed cost of operation, including a persistent and a transitory component.

A particularly simple example would be a pure “discrete firm-effects” model where λ_{it} does not vary over time, so that $\lambda_{it} \equiv \lambda_i$. More generally, λ_{it} might evolve across values according to a parameterized distribution

$$M(\lambda_{it} | \lambda_{it-1}; \theta_u). \quad (33)$$

We can think of this transition as being generated by a function

$$\lambda_{it} = h(\lambda_{it-1}, \nu_{it}, \theta_u) \quad (34)$$

where ν_{it} is a scalar or vector that is independent over time and markets, with distribution $\Phi(\nu_{it}; \theta_u)$.

This suggests a generalization to a broader set of models where we treat λ_{it} as the “persistent component” of the unobservables, ν_{it} as the innovation in that component and η_{it} as additional transient (independent over time) shocks to single-period profits. We then define u_{it} as the (possibly vector-valued) unobservable that enters the policy function. Note that in the discrete heterogeneity case $u_{it} = (\lambda_{it}, \eta_{it})$. As in (32), it might sometimes be helpful to distinguish ϵ_{it} as the particular combination of period t unobservables that affect single-period returns. However, as a more general case we can just let single-period profits depend on u_{it} .

The serial correlation case in the body of the text is particularly simple because the scalar u_{it} is simultaneously the persistent component (λ_{it}) that effects future expectations and is also the component (ϵ_{it}) that effects single period returns. That is, in the single-agent AR(1) serial correlation case, $u_{it} \equiv \lambda_{it}$ and $u_{it} \equiv \epsilon_{it}$. In the discrete heterogeneity case outlined above, those simple relationships do not hold. The term that enters single period profits, $\epsilon_{it} = \lambda_{it} + \eta_{it}$, is different from the term, λ_{it} , that conditions future expectations.

Having defined the firm-policy relevant vector of unobservables u_{it} in this way, the general discussion in the main body of identification via GIV in Section 2.3 goes through in its entirety. In the next subsection, we further illustrate the details of the discrete heterogeneity case and tie our GIV treatment to the treatment of discrete heterogeneity in Kasahara and Shimotsu (2009) and related literature.

A.2 Details on the Discrete Heterogeneity Case

In the single-agent case with discrete types of the last section, the single period return is still $\pi(a_{it}, x_{it}, w_{it}, u_{it}, \theta_\pi)$, but now

$$u_{it} = (\lambda_{it}, \eta_{it}). \quad (35)$$

The Bellman equation is

$$V(x_{it}, w_{it}, u_{it}) = \max_{a_{it} \in \mathcal{A}(x_{it})} (\pi(a_{it}, x_{it}, w_{it}, u_{it}; \theta_\pi) + \delta E_{\theta_u} [V(x_{it+1}, w_{it+1}, u_{it+1}) | a_{it}, x_{it}, w_{it}, \lambda_{it}]). \quad (36)$$

This takes the same form as (7), except that the future expectation is conditioned on λ_{it} , rather than u_{it} and we again recall the new definition of u_{it} in (35).

The associated policy function is still $a_{it} = \sigma_{it}(x_{it}, w_{it}, u_{it})$. Consider the specialization to the case of Section 3 with binary actions and states, and with no role for w_{it} . That is, a_{it} and x_{it} both take on the values 0 or 1. Further, assume only two possible values for λ , fixed over time. If the policy function is monotonic in the independent shock η , then under appropriate conditions the dynamic firm policy function is described by four cutoff rules in η , one for each combination of x_{it} and λ_{it} . This is twice as many cutoff rules as existed in the serial correlation example, where there was a cutoff in u_{it} for each value of x_{it} .

To be more formal, consider an AR(1) example following on the simplest example in the main text, where the inverse image set of unobservables in market i is

$$\mathcal{U}(a_i, x_i, w_i, \sigma) = \{u_i : (-1 + 2a_{it})u_{it} < (-1 + 2a_{it})\tau(x_{it}, w_{it}), \forall t\}. \quad (37)$$

Note that the term $(-1 + 2a_{it})$ is just a device to switch the sign of the inequality depending on whether or not $a_{it} = 1$. In the discrete heterogeneity case, the corresponding inverse image set is

$$\mathcal{U}(a_i, x_i, w_i, \sigma) = \{(\eta_i, \lambda_i) : (-1 + 2a_{it})\eta_{it} < (-1 + 2a_{it})\tau(x_{it}, w_{it}, \lambda_{it}), \forall t\}. \quad (38)$$

It is clear that (37) involves fewer unknown policy function parameters τ than (38). This is one clear advantage of the AR(1) case.

However, Kasahara and Shimotsu (2009) exploit a clear advantage of the discrete heterogeneity case in handling initial conditions. If we are willing to assume that persistent heterogeneity takes only a small number of discrete values, then the initial conditions are similarly described by a small set of additional unknown parameters. In particular, the initial conditions of the model are completely characterized by the unknown terms

$$\Pr(x_{i1} | \lambda_{i1} = \bar{\lambda}_k, w_{i1}),$$

which is a relatively small number of unknowns when λ_{i1} takes on a small number of values. Further, the number of unknown initial condition parameters does not increase in T . This simple specification of the initial conditions then allows Kasahara and Shimotsu (2009) to complete the dynamic model, which in turn allows them to specify a set of strict equalities that are necessary to rationalize the observed data. In the fully discrete case, Kasahara and Shimotsu (2009) then discuss (point) identification of transition probabilities conditional on each possible value of λ . They provide sufficient conditions involving combinations of [a] a minimum number of time periods, [b] sufficient variation in observed (unconditional on λ) multiperiod transition probabilities in the data, [c] limits on the number of discrete points of support for λ and [d] limits on the variation in λ over time. These results have the further advantage of possibly identifying joint transitions for (x_{it}, w_{it}) that depend on λ , which is more general than the base specification in the main body of our text.

In the single-agent case, then, the chief advantage of the discrete heterogeneity model lies in a naturally low-dimensional parameterization of the initial conditions, following directly from the assumption that the persistent heterogeneity is, in a formal sense, of low dimension relative to the overall variation in the data over time. It is possible that the advantages of the discrete heterogeneity assumption are highest when the time dimension is of intermediate length. When the time dimension is too short, the large number of combined parameters of the policy functions and the initial conditions will overwhelm the data even given the advantages of equality restrictions. At long time periods, GIV inequalities applied to the AR(1) serial correlation case may do reasonably well since the issue of incomplete initial conditions may, in any case, be small in long panels, as highlighted by Honoré and Tamer (2006).

A.3 Oligopoly

Here we extend the discussion of the last subsection to oligopoly. As an example, again consider a simple “entry” model where the return to firm j of operating in market i at time t is

$$\bar{\pi}_j(x_{it}, w_{it}, \theta_\pi) - (\lambda_{ijt} + \eta_{ijt}).$$

A convenient set of assumptions is that λ_{ijt} is discrete, persistent and commonly known by all firms at time t while η_{ijt} is continuously distributed, i.i.d. over firms and time and is also private information to firm j . Note that the private continuously distributed η_{ijt} can help to ensure existence of equilibrium in an otherwise discrete dynamic game, as in Doraszelski and Satterthwaite (2010). The assumption that λ is commonly observed prevents issues of informational updating and strategic information revelation with respect to persistent firm types. This often then preserves the first-order Markov structure of the oligopoly strategies.

Given the model of the last paragraph, the unobservables that enter firm j ’s single-period profit function are then

$$u_{ijt} = (\lambda_{ijt}, \lambda_{i,-j,t}, \eta_{ijt}). \quad (39)$$

Assuming (as usual in the empirical literature) that firms all play the same strategy whenever they reach the same state (even when the underlying game admits multiple equilibria), the equilibrium policy functions once again take the form,

$$a_{ijt} = \sigma_j(x_{it}, w_{it}, u_{ijt}),$$

with u_{ijt} now defined by (39). Once again, in a discrete game the results by Kasahara and Shimotsu (2009) may or may not provide point identification of cutoff decision rules in the quantiles of η_{ijt} (as well as transition matrices for x_{it} and w_{it} that potentially depend on λ_{it} .) These transitions and cutoffs will in some commonly used cases point identify the structural parameters and in other cases they may not. Either way, the GIV framework generalizes the result.

In particular, the inverse image set in the simple oligopoly entry example is very close to the single agent version in (38)

$$\mathcal{U}(a_i, x_i, w_i, \sigma) = \{(\eta_{ijt}, \lambda_{it}) : (-1 + 2a_{ijt})\eta_{ijt} < (-1 + 2a_{ijt})\tau(x_{it}, w_{it}, \lambda_{it}), \forall(j, t)\}. \quad (40)$$

Note that the cutoffs $\tau(x_{it}, w_{it}, \lambda_{it})$ take on a finite number of values as long as x_{it} , w_{it} , and λ_{it} all take on a finite number of values. One could allow for continuous λ , but then the function τ might easily take on a continuum of values.

GIV inequality constraints associated with (40) can once again be transformed into equalities by introducing additional parameters for initial conditions. However, the oligopoly case naturally multiplies the

number of these initial condition parameters. If each firm has a potentially different λ and if the number of potential values for λ is K , then we have K^J initial condition parameters. It is not immediately clear whether it is easier to include these parameters in the identification and/or estimation problem and take advantage of the resulting equality restriction, or whether it is better to use the GIV inequalities without the initial condition parameters. The answer is likely to be problem specific.


Another alternative to the discrete persistent heterogeneity model just described is a model that extends the AR(1) case from the main text to the multiple agents settings. In such a model, each agent has one continuously distributed persistent unobservable that is common knowledge. Continuity of the unobservable implies that we can establish existence of an equilibrium without the need for additional private information i.i.d. shocks. However, in general, the image sets $\mathcal{U}(a_i, x_i, w_i, \sigma)$ needed as an input to characterize the CR identified set for the policy functions will not be determined by a finite number of thresholds. Instead, they will typically be defined by curves in the space of unobservables, which makes implementation hard or outright infeasible. Therefore, we view the model described in this section as preferable for our purposes and we estimate a variant of it in Appendix C.

One could also consider examples of discrete/continuous games of the sort envisioned by Bajari et al. (2007)'s discussion of possible set identification in a world without persistent heterogeneity. Perhaps some combination of discrete and continuous persistent λ variables will be appropriate in that case.


Appendix B: Additional Details on Single-Agent Example

B.1 Non-Elemental Sets


First, we present the non-elemental sets belonging to the collection of core determining sets for the example considered in Section 3 with $T = 2$ time periods. These are the sets given by the overlapping unions of elemental sets, excluding cases of strict subsets. Below each set we specify the events associated with the elemental sets that we are taking the union of (top three sets) or the non-elemental sets that we taking the union of (bottom two sets).




(a) $(0,0,0)+(1,1,0)$




(b) $(0,0,0)+(1,1,1)$



(c) $(1,1,1)+(0,0,1)$



(d) $(a)+(b)$



(e) $(b)+(c)$

B.2 Projection of Identified Sets on the Space of Profit Parameters

Next, we complement the results in Section 3.3 by showing the projections of the identified sets onto the space of the profit parameters. We focus on the model without exogenous profit shifters w_{it} . One can see that the GMM and MLE models that assume away serial correlation tend to underestimate both profit parameters.

Figure 7: Low IV strength





Figure 8: High IV strength



Appendix C: Computed Example of Dynamic Oligopoly

In this section, we illustrate how our approach applies to the dynamic oligopoly setting by considering a model with discrete persistent common knowledge unobservables and private information i.i.d. shocks, along the lines of the example in Appendix A.3.

Consider the problem faced by two firms choosing how many stores to open in each of several markets over time. We assume that the firms choose between having one or two stores open in each market at any given point in time. Given the number of stores in a market, the two firms engage in Bertrand competition and each of them charges the same price across its own stores. We let j index stores and $k(j)$ be the firm owning store j . Consumers view the products as horizontally differentiated across stores. Specifically, let

$$u_{ijt} = \bar{\delta} - \alpha p_{k(j)t} + \epsilon_{ijt}$$

be the utility that consumer i gets from buying from store j at time t . Assuming that ϵ_{ijt} is i.i.d. extreme value across consumers and stores, the market share of firm k at time t is given by

$$s_{kt}(p_t) = \frac{x_{kt} \exp^{\bar{\delta} - \alpha p_{kt}}}{1 + \sum_{r=1}^2 x_{rt} \exp^{\bar{\delta} - \alpha p_{rt}}},$$

where $p_t = (p_{1t}, p_{2t})$ and x_{kt} denotes the number of stores that firm k has open at time t . The first order condition for firm k 's static problem at time t is then

$$p_{kt} = mc_{kt} - \frac{s_{kt}(p_t)}{\frac{\partial}{\partial p_{kt}} s_{kt}(p)}$$

where mc_{kt} denotes the marginal cost of firm k at time t , which is assumed to be constant across firm k 's stores. The system of first-order conditions implicitly determines the equilibrium prices in each market at any point in time and this in turn determines each firm's variable profits. Let $\pi_k^*(x_t, w_t)$ be firm k 's variable profit at time t as a function of the endogenous states $x_t = (x_{1t}, x_{2t})$ and of market size w_t , which we assume to be exogenous. Further, each firm incurs a fixed cost for each open store at time t and a sunk cost if it decides to open a new store in the next period. The timing is such that at time t a firm chooses a_{kt} , i.e. the number of stores at time $t + 1$, and incurs the associated sunk cost, if any, at time t . The final specification for flow profits is as follows:

$$\pi_{kt}(a_{kt}, x_t, w_t, \lambda_{kt}, \nu_{kt}) = \pi_k^*(x_t, w_t) - \beta x_{kt} \lambda_{kt} - \gamma (a_{kt} - x_{kt}) \mathbb{I}\{a_{kt} \geq x_{kt}\} \cdot \nu_{kt},$$

where β is the fixed cost, γ is the sunk cost and λ_{kt} and ν_{kt} are shocks that are unobserved to the econometrician. We assume that λ_{kt} is common knowledge and possibly correlated over time and across players. This is the structural shock that may be correlated with the state x_t , thus leading to an endogeneity problem. On the other hand, ν_{kt} is assumed to be i.i.d. over time and across players and to be observed by player k but not by the other firm. Adding this shock to the profit specification helps show existence of the dynamic equilibrium and compute said equilibrium in practice.

In the simulation, we set $\alpha = 0.5$, $\bar{\delta} = 2$ on the demand side and $mc_{kt} = 0.5$, $\beta = 0.7$, $\gamma = 1$ on the cost side. Moreover, we assume λ_{kt} can only take the two values 0.25 and 0.75, so that at any time t there are four possible combinations of $\lambda_t = (\lambda_{1t}, \lambda_{2t})$. The joint probability distribution for $(\lambda_t, \lambda_{t+1})$ is given by the

following matrix

$$\begin{bmatrix} 0.175 & 0.025 & 0.025 & 0.025 \\ 0.025 & 0.175 & 0.025 & 0.025 \\ 0.025 & 0.025 & 0.175 & 0.025 \\ 0.025 & 0.025 & 0.025 & 0.175 \end{bmatrix}$$

In other words, the λ_t process is fairly persistent over time. Further note that the matrix above is characterized by just one parameter, which is convenient computationally. In particular, this correlation structure implies that, given any value of λ_t , λ_{t+1} will be equal to λ_t with probability $\rho = 0.7$ and will switch to any of the other three values with 0.1 probability each. Moreover, the shock ν_{kt} takes one of the two values 0.5 and 1 with equal probability and is i.i.d. over time and across players. The exogenous state w is either 0.8 or 2 and, for simplicity, is taken to be constant over the time span from $t = 1$ to $t = T$, where T is the number of observed time periods. In the simulation, we let $T = 2$.³⁷

Finally, we use past values of w as instruments. For each firm, the instrument has a correlation of 0.81 and 0.56 with x from period 1 and 2, respectively.

The structural parameters are the fixed cost β , the sunk cost γ and the correlation parameter ρ for the λ_t process. To characterize the identified for $\theta = (\beta, \gamma, \rho)$, we proceed in two steps. First, we find all the policies that satisfy the CR conditions. Note that, in this simple example, there are only 64 possible states and therefore a policy is a vector of length 64 with elements taking two possible values (corresponding to choosing one or two stores). This makes it possible to enumerate all possible (monotonic) policies and check the GIV conditions for each of them. Only 5,335 policies survive at the end. This is the sharp GIV identified set for the policy.

In the second step, we go from the identified set for the policy to the identified set for the structural parameters. Specifically, for each candidate value θ in a grid, we check whether there exists a policy in the GIV identified set such that when the opponent plays according to that policy, it is optimal for a firm to play the same policy. Here we exploit the fact that the two firms face the same problem and thus there has to be a symmetric equilibrium. Further, note that this procedure does not rely on uniqueness of the equilibrium and thus can be applied without any changes even if a given value of θ is associated with multiple equilibria.

Figures 9 and 10 show projections of the identified set for the structural parameters.

³⁷In order to focus on identification of the parameters and abstract from inference issues, we draw a very large sample of markets (50,000).

Figure 9: Projection of identified set for β, γ





Figure 10: Projection of identified set for ρ



Appendix D: The Second Step in the Empirical Application

In this appendix, we write out the system of equations that we use to perform the second step of our procedure in the empirical application. In our sample, each market can have 0, 1 or 2 firms. We denote by 1 the first firm to enter the market and by 2 the second firm to enter. Abbring and Campbell (2010) show that in equilibrium firm 1 will also be the last to exit. Each market is characterized by the level of construction employment (W), which we discretize into “high” and “low,” respectively labeled \bar{w} , \underline{w} . Denote by $\tau_{j,In}(w)$ the entry threshold for firm j when construction employment is set at w . This is the value of the unobservable that makes firm j indifferent between staying out and entering the market. Similarly, let $\tau_{j,Out}(w)$ be firm j ’s exit threshold. By the properties of the Last-In-First-Out equilibrium and since profits are increasing in both W and the unobservable, we have (i) $\tau_{j,In}(w) \geq \tau_{j,Out}(w)$, for all j and w ; (ii) $\tau_{1,In}(w) \leq \tau_{2,In}(w), \tau_{1,Out}(w) \leq \tau_{2,Out}(w)$ for all w ; and (iii) $\tau_{j,In}(\bar{w}) \leq \tau_{j,In}(\underline{w}), \tau_{j,Out}(\bar{w}) \leq \tau_{j,Out}(\underline{w})$ for all j .

The single-period profit parameters to be recovered are $\pi(x, w)$ —the deterministic part of variable profits as a function of x , the number of firms in the market and w —and $\gamma(x, w)$ —the sunk cost of entry as a function of x and w . In total, there are four π and four γ parameters. Since we will have eight equations, we will be able to recover all of the eight parameters. However, it seems reasonable to assume that sunk costs are invariant to x and w and so the results in the paper impose that restriction.

We let $\theta \equiv (\pi, \gamma)$. Given a candidate profile of strategies $\tilde{\sigma}$ and a candidate serial correlation parameter $\tilde{\rho}$ from the first stage, the goal is to recover θ . Let $v_j(a_j, w, \epsilon, \tilde{\sigma}, \theta)$ be firm j ’s value function associated with

the action a_j at states (w, ϵ) when all firms follow the strategy profile $\tilde{\sigma}$ and the profit are parameterized by θ . Now we evaluate v_j at $a_j = 1$ (firm j chooses to be in the market) and $\epsilon = \tau_{j,In}(w)$ for each firm j and each value of w . At that value of the unobservable, j is indifferent between staying out of the market and entering. Thus, the value function is equal to zero, the value of staying out. The value of entering can be expressed as a linear function of the profit parameters via simulation as follows

$$v^s(1, w, \tau_{j,In}(w), \tilde{\sigma}, \theta) = \sum_{\tilde{x} \in \{1,2\}} \sum_{\tilde{w} \in \{\underline{w}, \bar{w}\}} \pi(\tilde{x}, \tilde{w}) \frac{1}{S} \sum_{s=1}^S \sum_{t=0}^{\bar{T}} \delta^t 1 \{x_t^s = \tilde{x}, w_t^s = \tilde{w}, a_{jr}^s = 1 \forall r \leq t\} - \gamma(j, w) + \sum_{s=1}^S \sum_{t=0}^{\bar{T}} \delta^t \epsilon_t^s 1 \{a_{jr}^s = 1 \forall r \leq t\}, \quad (41)$$

where (i) a_{jt}^s —and thus x_t^s —are determined by the candidate profile of strategies $\tilde{\sigma}$; (ii) $x_0^s = j$ for all s , since if firm 1 is indifferent between entering and staying out then firm 2 will stay out and, conversely, if firm 2 is indifferent between entering and staying out then firm 1 will be in; (iii) w_t^s is set to w for $t = 0$ and after that is drawn based on the law of motion for the exogenous covariates recovered outside the dynamic model; and (iv) ϵ_t^s is set to $\tau_{j,In}(w)$ for $t = 0$ and after that is drawn from the distribution with candidate serial correlation parameter $\tilde{\rho}$ from the first stage. Note that, since exit is irreversible, the sunk cost γ is paid only at $t = 0$. Equating (41) to zero gives one linear equation in θ for each firm and each value of w , i.e. four equations in total.

Similarly, we can evaluate v_j at $a_j = 1$ and $\epsilon = \tau_{j,Out}(w)$ for each firm j and each value of w . At that value of the unobservable, j is indifferent between staying in the market and exiting. The simulated version of the value function is

$$v^s(1, w, \tau_{j,Out}(w), \tilde{\sigma}, \theta) = \sum_{\tilde{x} \in \{1,2\}} \sum_{\tilde{w} \in \{\underline{w}, \bar{w}\}} \pi(\tilde{x}, \tilde{w}) \left[\frac{1}{S} \sum_{s=1}^S \sum_{t=0}^{\bar{T}} \delta^t 1 \{x_t^s = \tilde{x}, w_t^s = \tilde{w}, a_{jr}^s = 1 \forall r \leq t\} \right] + \sum_{s=1}^S \sum_{t=0}^{\bar{T}} \delta^t \epsilon_t^s 1 \{a_{jr}^s = 1 \forall r \leq t\}, \quad (42)$$

Relative to (41), now ϵ_0^s is set to $\tau_{j,Out}(w)$. Further, the γ parameters do not show up in (42) since firm j is already in at the initial time. Since in the model exit is irreversible, the value of exiting the market is zero. Thus, equating (42) to zero gives another set of four equations. In total, we have a system of eight linear equations which characterize the set of values of θ consistent with the first-stage candidates $(\tilde{\sigma}, \tilde{\rho})$ and the model.

Appendix E: Proofs Omitted from the Text

Proof of Lemma 4

First, note that the restrictions imposed in the second step imply that (25) is satisfied at $x = \bar{x}, u = \tau(\bar{x})$ for $\bar{x} \in \{0,1\}$. Thus, if we could show that the policy solving (25) is weakly decreasing in u , then we could conclude that that policy is the same as $\tilde{\sigma}$ at all values of x and u , completing the proof. The policy solving (25) is weakly decreasing in u if and only if $\pi(1, x, u) + \delta E[V(1, u') - V(0, u')|u]$ is weakly decreasing in u , where $\pi(1, x, u)$ denotes the single-period profit from choosing $a = 1$. The specification of the single-period profit function implies that $\pi(1, x, u)$ is (strictly) decreasing in u , so it suffices to show

that $E[V(1, u') - V(0, u')|u]$ weakly decreases in u . To this end, we use Corollary 1 in Stokey et al. (1989). Based on the corollary, if we can show that $E[\tilde{V}(1, u') - \tilde{V}(0, u')|u]$ weakly decreasing in u implies that the value function returned by the Bellman equation, $B(\tilde{V})$, satisfies the same property, then we have shown that the fixed point of the Bellman equation, V , satisfies the property, yielding the desired result. Note that

$$B(\tilde{V})(1, u) - B(\tilde{V})(0, u) = \max\{\pi(1, 1, u) + \delta E[V(1, u')|u], \delta E[V(0, u')|u]\} - \max\{\pi(1, 0, u) + \delta E[V(1, u')|u], \delta E[V(0, u')|u]\}$$

Since $\pi(1, 1, u) - \pi(1, 0, u) = \gamma \geq 0$, we have

$$B(\tilde{V})(1, u) - B(\tilde{V})(0, u) = \begin{cases} \gamma & \text{if } \pi(1, 0, u) + \delta E[\tilde{V}(1, u') - \tilde{V}(0, u')|u] \geq 0 \\ \pi(1, 1, u) + \delta E[\tilde{V}(1, u') - \tilde{V}(0, u')|u] & \text{if } 0 \leq \pi(1, 1, u) + \delta E[\tilde{V}(1, u') - \tilde{V}(0, u')|u] \leq \gamma \\ 0 & \text{if } \pi(1, 1, u) + \delta E[\tilde{V}(1, u') - \tilde{V}(0, u')|u] \leq 0 \end{cases}$$

By inspecting this function, one can see that it is weakly decreasing in u if $E[\tilde{V}(1, u') - \tilde{V}(0, u')|u]$ weakly decreases in u . Since $\tilde{\rho} \geq 0$, the distribution of u' is stochastically increasing in u . It then follows that $E[B(\tilde{V})(1, u') - B(\tilde{V})(0, u')|u]$ weakly decreases in u .

References

- ABBRING, J. H. AND J. R. CAMPBELL (2010): “Last-In First-Out Oligopoly Dynamics,” *Econometrica*, 78, 1491–1527.
- ACKERBERG, D., L. BENKARD, S. BERRY, AND A. PAKES (2007): “Econometric Tools for Analyzing Market Outcomes,” in *Handbook of Econometrics*, ed. by J. J. Heckman and E. Leamer, North-Holland, vol. 6A, chap. 63.
- ALTONJI, J. AND R. L. MATZKIN (2005): “Cross-Section and Panel Data Estimators for Nonseparable Models with Endogenous Regressors,” *Econometrica*, 73, 1053–1102.
- ANDERSON, T. AND C. HSIAO (1981): “Estimation of Dynamic Models with Error Components,” *Journal of the American Statistical Association*, 76, 598–606.
- ANDREWS, D. W. AND X. SHI (2013): “Inference based on conditional moment inequalities,” *Econometrica*, 81, 609–666.
- (2017): “Inference Based on Many Conditional Moment Inequalities,” *Journal of Econometrics*, 196, 275–287.
- ANDREWS, D. W. AND G. SOARES (2010): “Inference for Parameters Defined by Moment Inequalities Using Generalized Moment Selection,” *Econometrica*, 78, 119–157.
- ARCIDIACONO, P. AND R. MILLER (2011): “Conditional Choice Probability Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity,” *Econometrica*, 79, 1823–1868.
- ARELLANO, M. AND S. BOND (1991): “Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations,” *Review of Economic Studies*, 58, 277–297.
- ATHEY, S. AND G. IMBENS (2006): “Identification and inference in nonlinear difference-in-difference models,” *Econometrica*, 74, 431–497.
- BAJARI, P., C. L. BENKARD, AND J. LEVIN (2007): “Estimating Dynamic Models of Imperfect Competition,” *Econometrica*, 75, 1331–1370.
- BERESTEANU, A., F. MOLINARI, AND I. MOLCHANOV (2011): “Sharp Identification Regions in Models with Convex Moment Predictions,” *Econometrica*, 79, 1785–1821.
- BERRY, S. AND G. COMPIANI (2020): “Empirical Models of Industry Dynamics with Endogenous Market Structure,” Tech. rep., Yale, in preparation for the Annual Review of Economics.

- BERRY, S. AND E. TAMER (2007): “Identification in Models of Oligopoly Entry,” in *Advances in Economics and Econometrics: Theory and Applications, Ninth World Congress*, ed. by W. N. R. Blundell and T. Persson, Cambridge University Press, vol. 2.
- BERRY, S. T. (1992): “Estimation of a Model of Entry in the Airline Industry,” *Econometrica*, 60, 889–917.
- BERRY, S. T. AND P. A. HAILE (2014): “Identification in Differentiated Products Markets Using Market Level Data,” *Econometrica*, 82, 1749–1797.
- BLEVINS, J. (2016): “Sequential Monte Carlo Methods for Estimating Dynamic Microeconomic Models,” *Journal of Applied Econometrics*, 31, 773–804.
- BLUNDELL, R. AND S. BOND (1998): “Initial conditions and moment restrictions in dynamic panel data models,” *Journal of Econometrics*, 87, 115–143.
- BORKOVSKY, R. N., U. DORASZELSKI, AND Y. KRYUKOV (2010): “A User’s Guide to Solving Dynamic Stochastic Games Using the Homotopy Method,” *Operations Research*, 58, 1116–1132.
- BRESNAHAN, T. (1989): “Empirical Studies of Industries with Market Power,” in *The Handbook of Industrial Organization*, ed. by R. Schmalensee and R. Willig, North-Holland, no. 10 in Handbooks in Economics.
- CHEN, X., T. M. CHRISTENSEN, AND E. TAMER (2018): “Monte Carlo Confidence Sets for Identified Sets,” *Econometrica*, 86, 1965–2018.
- CHERNOZHUKOV, V., D. CHETVERIKOV, AND K. KATO (2018): “Inference on Causal and Structural Parameters using Many Moment Inequalities,” *The Review of Economic Studies*, 86, 1867–1900.
- CHERNOZHUKOV, V., H. HONG, AND E. TAMER (2007): “Estimation and Confidence Regions for Parameter Sets in Econometric Models¹,” *Econometrica*, 75, 1243–1284.
- CHERNOZHUKOV, V., S. LEE, AND A. ROSEN (2013): “Intersection Bounds: Estimation and Inference,” *Econometrica*, 81, 667–737.
- CHERNOZHUKOV, V., D. CHETVERIKOV, AND K. KATO (2018): “Inference on Causal and Structural Parameters Using Many Moment Inequalities,” *Review of Economic Studies*, forthcoming.
- CHESHER, A. (2010): “Instrumental Variables Models for Discrete Outcomes,” *Econometrica*, 78, 575–601.

- CHESHER, A. AND A. ROSEN (2017): “Generalized Instrumental Variable Models,” *Econometrica*, 83, 959–989.
- CILIBERTO, F. AND E. TAMER (2009): “Market Structure and Multiple Equilibria in Airline Markets,” *Econometrica*, 77, 1791–1828.
- COLLARD-WEXLER, A. (2013): “Demand Fluctuations in the Ready-Mix Concrete Industry,” *Econometrica*, 81, 1003–1037.
- (2014): “Mergers and Sunk Costs: An Application to the Ready-Mix Concrete Industry,” *American Economic Journal: Microeconomics*, 6, 407–447.
- DORASZELSKI, U. AND A. PAKES (2007): “A Framework for Applied Dynamic Analysis in I.O.” in *Handbook of Industrial Organization*.
- DORASZELSKI, U. AND M. SATTERTHWAITE (2010): “Computable Markov-perfect industry dynamics,” *The RAND Journal of Economics*, 41, 215–243.
- DUBÉ, J., G. HITSCH, AND P. ROSSI (2010): “State dependence and alternative explanations for consumer inertia,” *RAND Journal of Economics*, 41, 417–445.
- ERICSON, R. AND A. PAKES (1995): “Markov Perfect Industry Dynamics: A Framework for Empirical Work,” *Review of Economic Studies*, 62, 53–82.
- FERSHTMAN, C. AND A. PAKES (2012): “Dynamic Games with Asymmetric Information: A Framework for Empirical Work,” *The Quarterly Journal of Economics*, 127, 1611–1661.
- GALICHON, A. AND M. HENRY (2011): “Set identification in models with multiple equilibria,” *Review of Economic Studies*, 78, 1264–1298.
- HECKMAN, J. AND B. SINGER (1984): “A Method for Minimizing the Impact of Distributional Assumptions in Econometric Models for Duration Data,” *Econometrica*, 52, 271–320.
- HECKMAN, J. J., J. E. HUMPHRIES, AND G. VERAMENDI (2016): “Dynamic treatment effects,” *Journal of Econometrics*, 191, 276–292.
- HONORÉ, B. AND E. TAMER (2006): “Bounds on parameters in panel dynamic discrete choice models,” *Econometrica*, 3, 611–629.
- HOTZ, J. AND R. A. MILLER (1993): “Conditional Choice Probabilities and the Estimation of Dynamic Models,” *Review of Economic Studies*, 60, 497–529.

- HOTZ, V. J., R. A. MILLER, S. SANDERS, AND J. SMITH (1994): “A Simulation Estimator for Dynamic Models of Discrete Choice,” *The Review of Economic Studies*, 61, pp. 265–289.
- HU, Y. AND M. SHUM (2012): “Nonparametric identification of dynamic models with unobserved state variables,” *Journal of Econometrics*, 171, 32–44.
- HU, Y., M. SHUM, W. TAN, AND R. XIAO (2015): “A Simple Estimator for Dynamic Models with Serially Correlated Unobservables,” Tech. rep., Cal Tech, working Paper.
- IGAMI, M. (2018): “Industry Dynamics of Offshoring: The Case of Hard Disk Drives,” *American Economic Journal: Microeconomics*, 10, 67–101.
- ISRAEL, M. (2005): “Tenure dependence in consumer-firm relationships: an empirical analysis of consumer departures from automobile insurance firms,” *RAND Journal of Economics*, 36, 165–192.
- KAIDO, H., F. MOLINARI, AND J. STOYE (2019): “Confidence Intervals for Projections of Partially Identified Parameters,” *Econometrica*, 87, 1397–1432.
- KALOUPTSIDI, M., P. SCOTT, AND E. SOUZA-RODRIGUES (2020): “Linear IV Regression Estimators for Structural Dynamic Discrete Choice Models,” *Journal of Econometrics*, forthcoming.
- KASAHARA, H. AND K. SHIMOTSU (2009): “Nonparametric Identification of Finite Mixture Models of Dynamic Discrete Choices,” *Econometrica*, 77, pp. 135–175.
- KEANE, M. P. AND K. I. WOLPIN (1997): “The Career Decisions of Young Men,” *Journal of Political Economy*, 105, 473–522.
- LAZAREV, J., L. BENKARD, AND A. BODOH-CREED (2018): “Simulating the Dynamic Effects of Horizontal Mergers: U.S. Airlines,” Working paper, NYU.
- MAGNAC, T. AND D. THESMAR (2002): “Identifying Dynamic Discrete Decision Processes,” *Econometrica*, 70, 801–816.
- MANSKI, C. F. (1988): “Identification of Binary Response Models,” *Journal of the American Statistical Association*, 83, 729–738.
- (2003): *Partial Identification of Probability Distributions*, New York: Springer.
- MANSKI, C. F. AND E. TAMER (2002): “Inference on Regressions with Interval Data on a Regressor or Outcome,” *Econometrica*, 70, 519–546.

- MENZEL, K. (2009): “Consistent Estimation With Many Moment Inequalities,” *Journal of Econometrics*, 182, 329–350.
- NORETS, A. (2009): “Inference in Dynamic Discrete Choice Models with Serially Correlated Unobserved State Variables,” *Econometrica*, 77, 1665–1682.
- PAKES, A. (1994): “Dynamic Structural Models, Problems and Prospects: Mixed Continuous Discrete Controls and Market Interactions,” in *Advances in Econometrics*, ed. by C. Sims, Cambridge: Cambridge University, chap. 5, 171–260.
- PAKES, A. AND P. MCGUIRE (2001): “Stochastic Approximation for Dynamic Analysis: Markov Perfect Equilibrium and the ‘Curse’ of Dimensionality,” *Econometrica*.
- PAKES, A., M. OSTROVSKY, AND S. BERRY (2007): “Simple Estimators for the Parameters of Dynamic Games, with Entry/Exit Examples,” *RAND Journal of Economics*, 38, 373–399.
- PAKES, A., J. PORTER, K. HO, AND J. ISHII (2015): “Moment Inequalities and Their Application,” *Econometrica*, 83, 315–334.
- PESENDORFER, M. AND P. SCHMIDT-DENGLER (2008): “Asymptotic Least Squares Estimators for Dynamic Games,” *Review of Economic Studies*, 75, 901–928.
- (2010): “Sequential Estimation of Dynamic Discrete Games: A Comment,” *Econometrica*, 78, 833–842.
- REICH, G. (2018): “Divide and Conquer: Recursive Likelihood Function Integration for Hidden Markov Models with Continuous Latent Variables,” *Operations Research*, 66, 1457–1470.
- RUDIN, W. (1976): *Principles of Mathematical Analysis*, McGraw-Hill Education, 3rd ed.
- RUST, J. (1987): “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55, 999–1033.
- RYAN, S. P. (2012): “The Costs of Environmental Regulation in a Concentrated Industry,” *Econometrica*, 80, 1019–1061.
- SOMAINI, P. (2015): “Identification in Auction Models with Interdependent Costs,” Tech. rep., Stanford.
- STOKEY, N. L., R. E. LUCAS, AND E. C. PRESCOTT (1989): *Recursive Methods in Economic Dynamics*, Cambridge MA: Harvard Univ.
- TAMER, E. (2003): “Incomplete Simultaneous Discrete Response Model with Multiple Equilibria,” *The Review of Economic Studies*, 70, pp. 147–165.