

Machine Learning for Dynamic Discrete Choice

Vira Semenova *

October 25, 2018

Abstract

Dynamic discrete choice models often discretize the state vector and restrict its dimension in order to achieve valid inference. I propose a novel two-stage estimator for the set-identified structural parameter that incorporates a high-dimensional state space into the dynamic model of imperfect competition. In the first stage, I estimate the state variable's law of motion and the equilibrium policy function using machine learning tools. In the second stage, I plug the first-stage estimates into a moment inequality and solve for the structural parameter. The moment function is presented as the sum of two components, where the first one expresses the equilibrium assumption and the second one is a bias correction term that makes the sum insensitive (i.e., Neyman-orthogonal) to first-stage bias. The proposed estimator uniformly converges at the root- N rate and I use it to construct confidence regions. The results developed here can be used to incorporate high-dimensional state space into classic dynamic discrete choice models, for example, those considered in Rust (1987), Bajari et al. (2007), and Scott (2013).

1 Introduction

In empirical work on dynamic models, economists often make specification choices - for example, discretize state space or select covariates for flow utility - in order to achieve computational tractability and precise estimates (e.g. Pakes and Schankerman (1984), Pakes (1986), Rust (1987), Ryan (2012)). Typically, economists have little intuition about which covariates to select or how to discretize a continuous state variable (Lanjouw et al. (1998)). Unfortunately, counterfactual predictions of dynamic models are sensitive to specification choices and are difficult to interpret when

*I am deeply grateful to my advisors Victor Chernozhukov, Whitney Newey, and Anna Mikusheva for their guidance and encouragement. I am thankful to Alberto Abadie, Chris Ackerman, Sydnee Caldwell, Denis Chetverikov, Ben Deaner, Mert Demirer, Jerry Hausman, Peter Hull, Tetsuya Kaji, Kevin Li, Elena Manresa, Rachael Meager, Denis Nekipelov, Cory Smith, Sophie Sun, Roman Zarate for helpful comments.

a model is misspecified. As discussed in Athey (2017), there has been recent interest in data-driven model selection based on modern machine learning tools. Moreover, (Belloni et al. (2016), Chernozhukov et al. (2017a), Chernozhukov et al. (2017b), Wager and Athey (2016)) have shown how to leverage these tools into high-quality estimates of causal parameters. Because dynamic models are more challenging to analyze, the model selection in dynamic models remains an open question.

This paper estimates a dynamic model of imperfect competition with a high-dimensional state space. Consider the bus engine replacement model from Rust (1987) as a special case. An agent decides whether to replace the bus engine in each period. The agent incurs a fixed cost in the case of replacement and a cost proportional to the current mileage in the case of maintenance. One can imagine that the future bus mileage depends on a vector of observed exogenous characteristics, such as current traffic and weather conditions, encoded in a high-dimensional vector. While these characteristics do not enter into the per-period utility of the owner, they affect the mileage's law of motion in the case of maintenance decision, and hence are taken into account in the agent's optimal renewal policy. Therefore, the expected value of bus ownership depends on a high-dimensional state vector consisting of the current mileage and exogenous characteristics. I am interested in the identified set of the possible values of the cost parameters that are rationalized by the agent's optimal behavior. In addition to the model in Rust (1987), the methods developed here apply to a broad variety of dynamic models: for example, those considered in Bajari et al. (2007) or Scott (2013), etc. that point- or partially identify their parameters under various assumptions about the agent's optimal behavior.

The main difficulty of this approach is estimating the value function. The plug-in (naive) approach of Bajari et al. (2007) consists of two steps. In the first step, one estimates the state variable's law of motion and the equilibrium policy function, effectively recovering agents' equilibrium beliefs. In the second step, one estimates the value function by drawing a sequence of states and actions from the conditional distributions estimated in the first step and averaging over multiple simulation draws. If the state variable is high-dimensional, we must regularize the estimate of the first-stage parameter in order to achieve consistency in high dimension. An inherent cost of these methods is bias that converges slower than parametric rate. As a result, first-stage bias carries over into the second stage, resulting in a low-quality estimate of the identified set.

The major challenge of this paper is to overcome transmitting the bias from the first to the second stage. A basic idea, proposed in a point-identified case, is to make the moment equation insensitive, or, formally, Neyman-orthogonal, to the biased estimation of the first-stage parameter

(Neyman (1959), Chernozhukov et al. (2017a)). The second idea consists in the use of different samples to estimate the equilibrium beliefs at the first stage and to compute the sample average of the bias correction term at the second stage. Using different samples in the first and the second stages allows me to employ modern machine learning methods to estimate the first stage parameters.

The first contribution of this paper is to derive a Neyman-orthogonal moment for the value function. This equation is presented as the sum of two terms, where the first term is the value function itself and the second term is a bias correction term that makes the sum insensitive (i.e., Neyman-orthogonal) to the first-stage bias. To derive the bias correction term, I characterize the value function as a solution to the recursive (Bellman) equation, which equates the expected value at the current state to the sum of expected immediate payoff and the expected discounted future payoff. This equation can be viewed as a semiparametric moment equation, where the parametric component gives the value function's moment (i.e., weighted average) and the nuisance parameter consists of the first-stage parameters appearing in the value function (i.e., conditional choice probability and state variable's law of motion). The first-stage parameter appears both inside and outside the value function. For example, the conditional probability of a given choice appears as a weight on immediate and discounted future payoffs corresponding to that choice. Applying the implicit function theorem to this equation, I derive the bias correction term and show how to approximate it by simulation. This derivation is novel: this is the first result in the literature that derives the bias correction term for a moment function that is not available in the closed form.

The second contribution of this paper is to extend the general theory of moment inequalities proposed by Chernozhukov et al. (2007) to allow for moment functions that depend on a first-stage nuisance parameter that can be high-dimensional (e.g., conditional choice probability). I characterize the identified set as the minimizer of the criterion function that penalizes the incorrect sign of the moment inequality (e.g., the sign that contradicts the optimality assumption). I show that, if the moment function is insensitive with respect to the biased estimation of its nuisance parameter at each point of the space of the structural parameter, plugging in the first-stage estimate of the nuisance parameter into the moment function delivers a high-quality sample criterion function. In particular, the estimator of the identified set obtained by inverting the sample criterion function converges at the same rate as if the true value of the first-stage nuisance parameter were known. Furthermore, inferential statistics based on the estimated moment function have a non-degenerate large sample distribution and are used to construct confidence regions for the identified set by subsampling.

This paper leaves a number of open questions. First, this article assumes that the first-stage nuisance parameter is identified. This assumption does not hold for every application (e.g., Ciliberto and Tamer (2009)). Second, the high-dimensional state space introduces a wealth of feasible suboptimal Markov policies to choose from for the construction of moment inequalities. In this paper I take the set of chosen alternatives as given, leaving the optimal choice of these alternatives for future research.

The structure of the paper is as follows. Section 2 gives a brief overview of the results. Section 3 gives the low-level sufficient conditions for the dynamic discrete choice model in Bajari et al. (2007). Section 4 presents an asymptotic theory for the identified sets defined by the semiparametric moment inequalities.

1.1 Literature review

This paper is built on three lines of work: estimation and inference in partially identified models and Neyman-orthogonal semiparametric estimation.

The first line of research, see e.g. Rosen (2006), Chernozhukov et al. (2007), Romano and Shaikh (2010) develops a framework for the estimation and inference of identified sets that are partially identified by moment inequalities. Extending this framework, Kaido and White (2014) allow the moment function to depend on an identified low-dimensional first-stage parameter in addition to the target. Furthermore, I allow the first-stage parameter to be a high-dimensional vector or a highly complex function and estimate it by modern machine learning methods.

Within the first line of my research, my application is most connected to the estimation of dynamic models of imperfect competition. Specifically, I build on Bajari et al. (2007) who identifies the structural parameter as a solution to a set of moment inequalities that embody the assumptions about the agent's optimal behavior. This paper proposed a two-stage algorithm to estimate the parameter, where in the first-stage one estimates the state variables's law of motion and equilibrium policy function, and then plugs them into a moment equation derived from the equilibrium assumption. However, to achieve valid inference, Bajari et al. (2007) imposed parametric restrictions in the first stage. Extending this algorithm, I drop these restrictions and estimate first-stage parameters by machine learning methods.

The second line of research (Ibbragimov and Hasminskii (1981), Andrews (1994), Newey (1994), van der Vaart (1998)) is concerned with obtaining a root- N consistent and asymptotically normal estimate for a low-dimensional target parameter in the presence of a nuisance parameter. In

this literature, a two-stage statistical procedure is insensitive, or, formally, Neyman-orthogonal, to the estimation error of the first-stage parameter (Neyman (1959)). Extending the orthogonality idea from a parametric to semiparametric setup was done in Newey (1994), Robins et al. (1994), Robins and Rotnitzky (1995). Combining Neyman-orthogonality and sample splitting, Chernozhukov et al. (2017a) and Chernozhukov et al. (2017b) incorporated modern machine learning methods to estimate low-dimensional target parameters defined by semiparametric moment equations. Subsequently, the idea has been extended to the case of high-dimensional target parameter in Chernozhukov et al. (2017c) and Chernozhukov et al. (2018). This paper translates the idea of Neyman-orthogonality from point- to set-identified case.

Within the second line research, my application is most connected to the estimation of dynamic discrete choice models under point-identification (Bajari et al. (2010), Bajari et al. (2015), Arcidiacono et al. (2013), Chernozhukov et al. (2017b)). Specifically, Chernozhukov et al. (2017b) introduces high-dimensional state space into a dynamic discrete choice model whose choice set contains a renewal choice and derives a Neyman-orthogonal moment equation for the structural parameter in that model. In this paper, I address the cases that do not have renewal choice property and derive the Neyman-orthogonal moment equation for the value function directly. This result is applicable to both point- and set-identified cases with discrete and continuous choice sets.

2 Set-Up and Motivation

I am interested in an identified set defined by moment inequalities

$$\Theta_I := \{\theta : \mathbb{E}m(D, \theta, \eta_0) \leq 0\}, \tag{1}$$

where D is the data vector distributed as P_D , θ is the parameter of interest, and η is an identified yet unknown parameter of the distribution P_D whose true value is η_0 . For example, in the bus engine replacement model of (Rust (1987)) the parameter of interest, θ , is a vector of operational and replacement costs, the data vector D contains bus mileage and other observed bus characteristics, η contains mileage's law of motion and the conditional probabilities of bus replacement, and the inequality restriction comes from the assumption that agent behaves optimally. I allow the state variable, w , to be high dimensional and estimate the parameter η by modern machine learning methods.

The examples below demonstrate how a high-dimensional state w may appear in the dynamic

discrete choice model.

Example 1. Engine Replacement Model from Rust (1987) with a High-Dimensional State Variable A single agent makes a binary decision $a \in \mathcal{A} = \{0, 1\}$ whether to replace a bus engine in each period $t \in \{1, 2, \dots, \infty\}$. His per-period utility function is

$$\pi(a, s, \epsilon) = \begin{cases} -R + \epsilon(0), & a = 0, \\ -\mu \cdot s + \epsilon(1), & a = 1, \end{cases} \quad (2)$$

where $s \in \mathcal{R}$ is the bus mileage, R is replacement costs, μs is operational cost, a is the decision of the agent, and $\epsilon = (\epsilon(0), \epsilon(1))$ is a vector of private shocks associated with each decision. The state variable w consists of the mileage s and additional high-dimensional vector of exogenous variables x (e.g., engine manufacturer characteristics) that I assume do not change with time. After the replacement decision ($a = 0$) the mileage s_{next} resets to 1 with probability one. After the maintenance decision ($a = 1$), the mileage s_{next} follows a first-order Markov process

$$s_{next} = \rho_0(w) + e, \quad e \sim N(0, 1), \quad (3)$$

where $\rho_0(w)$ is the conditional expectation function of the mileage tomorrow s_{next} given the state today w and $e \sim N(0, 1)$ is an independent $N(0, 1)$ shock. The target parameter $\theta = (R, \mu)$ consists of the replacement and operational cost parameters. The observed data vector consists of the current state, action, and future state (i.e, $D = (w, a, w_{next})$). Under the assumptions discussed below, the unknown yet identified high-dimensional parameters consist of the conditional choice probability $\gamma(w) = P(a = 1|w)$ and the transition function $\rho(w)$.

Example 2. Entry Game with a Long-Lived and a Short-Lived Player

In each period $t \in \{1, 2, \dots, \infty\}$ Apple decides whether to issue a new model of a phone ($a_{Apple} = 0$) or keep the existing one ($a_{Apple} = 1$). In addition to Apple, a short-lived potential entrant (Player 2) decides whether to issue a fake ($a_P = 1$) or not ($a_P = 0$). Each period Apple faces a new copy of player 2. Player 2's actions do not influence the motion of the state, and he has no dynamic incentives. In each period both players observe a state vector $w = (s, x)$ that consists of the age of the current make s and the vector x of short-lived player's characteristics (e.g., country, information about intellectual property rights protection) that does not change with time (i.e, $x_t = x_0$).

In each period both players receive a privately observed shock. Apple's utility function is given

by

$$\pi(a, w, \epsilon) = \begin{cases} -R + \delta_1 a_P + \epsilon(0), & a_{Apple} = 0, \\ -\mu \cdot s + \delta_2 a_P + \epsilon(1), & a_{Apple} = 1, \end{cases} \quad (4)$$

where R is the fixed cost of replacing the current model with a new one, $-\mu \cdot s$ is the profit from the current make that decays with age, and $\epsilon = (\epsilon(0), \epsilon(1))$ is a vector of Apple's shocks associated with each decision. After the replacement decision ($a_{Apple} = 1$) the age s_{next} resets to 1 with probability one. After the maintenance decision ($a_{Apple} = 2$) the age s_{next} increases by 1 with probability 1

$$s_{next} = s + 1.$$

The target parameter $\theta = (R, \mu, \delta_1 - \delta_0)$ consists of the cost parameters and the difference between the interaction parameters $\delta_1 - \delta_0$. The observed data vector $D = (w, a, w_{next})$ consists of the current state, w , action profile $a = (a_{Apple}, a_P)$, and the future state, w_{next} . Under the assumptions discussed below, the unknown yet identified high-dimensional parameters consist of the conditional choice probability of both players: $\gamma_A(w) = P(a_{Apple} = 1|w)$ and $\gamma_P(w) := P(a_P = 1|w)$.

Consider the setting of Example 1. I assume that the agent follows a Markov policy $\sigma(w, \epsilon)$ that maps the current state $w \in \mathcal{W} \subset \mathcal{R}^{d_w}$ and the shock vector $\epsilon \in \mathcal{R}^2$ into the action space $\mathcal{A} = \{0, 1\}$. The value function $V(w; \theta; \sigma)$ of a Markov policy σ is given by

$$V(w; \theta; \sigma) = \mathbb{E}\left[\sum_{t=0}^{\infty} \beta^t \pi(\sigma(w_t, \epsilon_t), w_t, \epsilon_t) | w\right],$$

where $\beta < 1$ is a discount factor. This function can also be written recursively:

$$V(w; \theta; \sigma) = \mathbb{E}_\epsilon[\pi(\sigma(w, \epsilon), w, \epsilon(\sigma(w, \epsilon))) + \beta \mathbb{E}[V(w_{next}; \theta; \sigma) | w, \sigma(w, \epsilon)]], \quad (5)$$

where the first and second summands show the expected current profit and the future expected discounted value, respectively. Define the choice-specific value function as

$$\begin{aligned} v(a, w) &:= -R_0(1 - a) + (-\mu_0 s)a + \beta \mathbb{E}[V(w_{next}; \theta_0; \sigma^*) | w, a = 0](1 - a) \\ &+ \beta \mathbb{E}[V(w_{next}; \theta_0; \sigma^*) | w, a = 1]a, \quad a \in \mathcal{A}, \end{aligned}$$

where the current deterministic utility is evaluated at the action a and the future discounted value is evaluated at the optimal strategy σ^* , conditional on the current action a . The symbols (R_0, μ_0) stand for the true values of the cost parameters. The choice a is optimal if and only if its total utility is greater than or equal to the utility of any other choice $a' \in \mathcal{A}$:

$$v(a, w) + \epsilon(a) \geq v(a', w) + \epsilon(a'), \quad a' \in \mathcal{A}.$$

Then the optimal Markov policy $\sigma^*(w, \epsilon)$ has a cutoff form:

$$\sigma^*(w, \epsilon) = \begin{cases} 1, & \epsilon(1) - \epsilon(0) \geq -v(1, w) + v(0, w) \\ 0, & \text{otherwise.} \end{cases} \quad (6)$$

To identify the difference $(v(1, w) - v(0, w))$ I make the following standard assumption (e.g., Hotz and Miller (1993), Scott (2013)) that I will maintain throughout the paper. In general case, the model will involve several players K . Denote their choice sets by $\mathcal{A}_k, k \in \{1, 2, \dots, K\}$.

ASSUMPTION 1 (Independent logit errors). *The components of the private shock vector $\epsilon = \{(\epsilon_k(j))_{j \in \mathcal{A}_k}\}_{k=1}^K \in \mathcal{E}$ are identically and independently distributed with a type 1 extreme value distribution whose distribution function is equal to $F(t) = \exp(-\exp(-t))$.*

As discussed in Hotz and Miller (1993), the vector of differences of the choice-specific value functions can be expressed as

$$v(a_1, w) - v(1, w) = \log \frac{P(a_1|w)}{P(1|w)} = \log \frac{\gamma_0(w)}{1 - \gamma_0(w)}, \quad (7)$$

where $\gamma_0(w) = P(a = 1|w)$ is the probability of the decision to maintain the engine conditional on the state w which is identified. For expositional purpose, I consider a simple suboptimal Markov policy: the choice of the decision based on the coin toss:

$$\mathbb{E}(\sigma(w, \epsilon) = 1|w) = \frac{1}{2}. \quad (8)$$

Combining 6 and 7, I recognize that the optimal strategy can be viewed as a function of γ :

$$\sigma^*(w, \epsilon) = \sigma^*(w, \epsilon, \gamma)$$

Therefore, value function $V(w; \theta; \sigma^*)$ can be viewed as the function of γ .

Define the moment function $m(w, \theta, \gamma)$ as the difference of the value function evaluated for the suboptimal and the optimal strategies

$$m(w, \theta, \gamma) := V(w; \theta; \sigma; \gamma) - V(w; \theta; \sigma^*; \gamma). \quad (9)$$

Define the identified set Θ_I as

$$\Theta_I := \mathbb{E}[m(w, \theta, \gamma_0)] \leq 0.$$

Because σ^* is an optimal strategy, the inequality above holds for the true parameter θ_0 , and Θ_I is a valid identified set.

2.1 Naive Approach to the Estimation of the Identified Set

The function $m(w, \theta, \eta)$ given in (9) presents two complications. First, the value function $V(w; \theta; \sigma; \eta)$ depends on the unknown nuisance parameter η . Second, even if the value of η is given, the value function is not readily available in the closed form and must be approximated by simulation. Here I focus on the first complication as if the moment function were readily available, leaving the description of the simulation algorithm for Section 3.

A possible, though naive estimator of the value function can be constructed as follows. Consider an ideal scenario where the researcher knows mileage's law of motion. Then the choice probability γ is the only unknown nuisance parameter that appears in the value function. Given an i.i.d sample $(D_i)_{i=1}^N$ from the law P_D , it is split into a main sample J_2 and an auxiliary sample J_1 of equal size $n = \lfloor N/2 \rfloor$ such that $J_1 \sqcup J_2 = \{1, 2, \dots, N\}$. After that, the estimator of the value function $V(w; \theta; \sigma^*; \gamma)$ is constructed as the sample average:

$$\hat{V}(w; \theta; \sigma^*; \hat{\gamma}) := \frac{1}{|J_2|} \sum_{i \in J_2} V(w_i; \theta; \sigma^*; \hat{\gamma}(w_i)),$$

where $\hat{\gamma}$ is estimated on the auxiliary sample J_1 . Unfortunately, for some parameter values $\theta \in \Theta$ this estimator has slower than \sqrt{N} convergence:

$$\sqrt{N} |\hat{V}(w; \theta; \sigma^*; \hat{\gamma}) - \mathbb{E}V(w; \theta; \sigma^*; \gamma_0)| \rightarrow \infty. \quad (10)$$

Therefore, the estimator $\hat{\Theta}_I$ of the identified set Θ_I based on the moment function (9) has suboptimal convergence rates.

The source of the slow convergence (10) can be understood through the following decomposition:

$$\begin{aligned} \sqrt{N}(\hat{V}(w; \theta; \sigma^*; \hat{\gamma}) - \hat{V}(w; \theta; \sigma^*; \gamma_0)) &= \underbrace{\sqrt{N}(\mathbb{E}[V(w_i; \theta; \sigma^*; \gamma_0)] - \frac{1}{|J_2|} \sum_{i \in J_2} V(w_i; \theta; \sigma^*; \gamma_0))}_{\mathbf{a}} \\ &+ \underbrace{\sqrt{N}\mathbb{E}[V(w_i; \theta; \sigma^*; \hat{\gamma}) - V(w_i; \theta; \sigma^*; \gamma_0)]}_{\mathbf{b}} \\ &+ \underbrace{\sqrt{N}(\frac{1}{|J_2|} \sum_{i \in J_2} V(w_i; \theta; \sigma^*; \hat{\gamma}) - \mathbb{E}V(w_i; \theta; \sigma^*; \hat{\gamma}))}_{\mathbf{c}}. \end{aligned}$$

The term **a** is the centered sample average of the value function $V(w; \theta; \sigma^*; \gamma_0)$ evaluated at the true value γ_0 of the choice probability. Due to the sample splitting, the term **c** is a centered sample average of conditional on J_1 i.i.d random variables and is well-behaved. The term **b** stands for

the bias of the value function $V(w; \theta; \sigma^*; \hat{\gamma})$ coming from the biased estimation of the conditional choice probability γ . This term is responsible for the slow convergence (10).

The divergence of \mathbf{b} comes from the combination of two facts: biased estimation of the first stage and the transmission of the bias from the first to the second stage. Since the conditional choice probability $\gamma(w)$ is a function of a high-dimensional state vector w , bias of its machine learning estimate (e.g., ℓ_1 -regularized logistic regression) converges slower than root- N . Because the value function $V(w; \theta; \sigma^*; \gamma)$ is sensitive to this bias, it carries over into the second stage.

Overcoming the Regularization Bias using Orthogonalization. To overcome the translation of the first stage bias into the second stage I add the bias correction term to the value function $V(w; \theta; \sigma^*; \gamma)$ to make it insensitive with respect to the biased estimation of γ . The new moment function for the value function $V(w; \theta; \sigma^*; \gamma)$ takes the form

$$g(D; \theta; \sigma^*; \gamma) := V(w; \theta; \sigma^*; \gamma) + \frac{1}{1-\beta} \Gamma(w; \theta) (1_{a=1} - \gamma(w)),$$

where the function $\Gamma(w; \theta)$ for $\theta = (R, \mu)$ is

$$\begin{aligned} \Gamma(w; \theta) := & -\mu s + R + \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma_0) | w, a = 1] \\ & - \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma_0) | w, a = 0] - \frac{2}{\gamma(w)} - \log \frac{\gamma(w)}{1 - \gamma(w)}. \end{aligned} \quad (11)$$

Because the bias correction term has zero mean

$$\mathbb{E}\left[\frac{1}{1-\beta} \Gamma(w; \theta) (1_{a=1} - \gamma_0(w))\right] = \frac{1}{1-\beta} \mathbb{E}_w \Gamma(w; \theta) \mathbb{E}[(1_{a=1} - \gamma_0(w)) | w] = 0, \quad \theta \in \Theta,$$

the new moment function can replace the old one in the definition of the identified set. Moreover, the new function is insensitive to the bias estimation of γ . As a result, under additional mild regularity conditions, the regularization bias of the choice probability $\hat{\gamma}(w)$ does not translate into the bias of the moment function:

$$\sqrt{N} |\mathbb{E}g(D, \theta, \hat{\gamma}) - g(D, \theta, \gamma_0)| \rightarrow 0.$$

The Role of Sample Splitting in Preventing Overfitting. Another key aspect of the proposed analysis is using different samples J_1 and J_2 for different stages of estimating the value function. Had I used the whole sample to estimate the conditional choice probability $\hat{\gamma}(w)$, the sample average

$$\frac{1}{N} \sum_{i=1}^N g(D_i, \theta, \hat{\gamma}) - \mathbb{E}[g(D_i, \theta, \hat{\gamma})]$$

would not be a sample average of the i.i.d (or weakly dependent) random observations. The relation between the first stage error $\hat{\gamma}(w_i) - \gamma_0(w_i)$ and the value of the moment function $g(D_i, \theta, \gamma_0)$ creates bias, referred to as overfitting bias. To control the overfitting bias in the worst-case scenario:

$$\mathbb{E} \sup_{\gamma \in \mathcal{G}} \left| \sum_{i=1}^N (g(D_i, \theta, \gamma) - g(D_i, \theta, \gamma_0)) - \mathbb{E}[g(D_i, \theta, \gamma) - g(D_i, \theta, \gamma_0)] \right|,$$

one must impose complexity constraints on the class of functions \mathcal{G} that are used to estimate the conditional choice probability. While some nonparametric estimators designed for low-dimensional state spaces obey these constraints, some modern machine learning estimators designed for high-dimensional state variables do not. To accommodate machine learning estimators at the first stage, I use different samples.

Bias Correction Terms for Examples 1 and 2 . Suppose a nuisance parameter can be presented as a conditional expectation $\mathbb{E}[U|w]$. Then bias correction term for conditional choice probability takes the form

$$\alpha(D; \theta; \xi) = \Pi(w; \theta)(U - \mathbb{E}[U|w]),$$

where ξ is an unknown vector-valued function of the state variable w . The true value $\xi_0 = \xi_0(\theta)$ consists of the original conditional expectation function $\mathbb{E}[U|w]$ and the function $\Pi(w; \theta)$

$$\xi_0(\theta) := \{\mathbb{E}[U|w], \Pi(\cdot, \theta)\}. \quad (12)$$

Remark 1 (Example 1, continued). *Consider the setup in Example 1. The nuisance parameter $\eta = (\gamma, \rho)$ consists of the conditional choice probability γ and the transition function ρ defined in (3). The transition function ρ is present in the value function evaluated for the optimal σ^* and the suboptimal σ strategies. The conditional choice probability enters $V(w; \theta; \sigma^*; \eta)$ only through the optimal strategy σ^* , described in (6)-(7). To sum up, the bias correction term for Example 1 is*

$$\alpha(D; \theta; \xi) = \alpha_{\sigma}^{TRANS}(D; \theta; \xi) - \alpha_{\sigma^*}^{TRANS}(D; \theta; \xi) - \alpha^{CCP}(D; \theta; \gamma), \quad (13)$$

where e.g., $\alpha_{\sigma}^{TRANS}(D; \theta; \xi)$ is the individual bias correction term for ρ for the case of suboptimal strategy σ .

As discussed above, the bias correction term $\alpha^{CCP}(D; \theta; \gamma)$ that corrects the bias of the conditional choice probability γ_0 is

$$\alpha^{CCP}(D; \theta; \gamma) = \frac{1}{1 - \beta} \Gamma(w; \theta)(1_{\{a=1\}} - \gamma(w)), \quad (14)$$

where the function $\Gamma(w, \theta)$ is given in (11). The form for the bias correction term for the transition function ρ is the same regardless whether σ is optimal or not. Let $\tilde{\sigma} \in \{\sigma, \sigma^*\}$ be a Markov policy. Suppose the state vector w has a stationary distribution. Then the bias correction term of the value function $V(w; \theta; \sigma; \eta)$ for the transition function $\rho(\cdot)$ is equal to:

$$\alpha_{\tilde{\sigma}}^{TRANS}(D; \theta; \xi) := \frac{\beta}{1 - \beta} \mathbb{E}\left[\frac{dV(x; \theta; \tilde{\sigma}; \eta_0)}{dx}\Big|_{x=w_{next}}\Big|w, a = 1\right](s_{next} - \rho(w)), \quad (15)$$

where ξ is an unknown vector-valued function of the state variable w . Its true value $\xi_0 = \xi_0(\theta)$ consists of the original nuisance parameter η_0 and the function $\Pi_0(w, \theta)$:

$$\Pi_0(w, \theta) := \mathbb{E}\left[\frac{dV(x; \theta; \sigma; \eta_0)}{dx}\Big|_{x=w_{next}}\Big|w\right],$$

which is equal to the expectation of the derivative of the value function with respect to the state variable evaluated for the future state conditional on the current state w .

The bias correction term for other suboptimal Markov policies is more complicated, but can be derived using the argument of Appendix 5.

Remark 2 (Example 2, continued). Consider the setup in Example 2. The nuisance parameter $\eta = (\gamma_A, \gamma_P)$ consists of the conditional choice probabilities of Apple and Player 2. The choice probability γ_A enters the value function $V(w; \theta; \sigma^*; \eta)$ only through the optimal strategy σ^* , described in (6)-(7). The choice probability γ_P enters the value function $V(w; \theta; \tilde{\sigma}; \eta)$, $\tilde{\sigma} \in \{\sigma^*, \sigma\}$. To sum up, the bias correction term for Example 1 is equal to:

$$\alpha(D; \theta; \eta) := -\alpha_A^{CCP}(D; \theta; \gamma) + \alpha_P^\sigma(D; \theta; \gamma_P) - \alpha_P^{\sigma^*}(D; \theta; \gamma_P),$$

where e.g. $\alpha_A^{CCP}(D; \theta; \gamma)$ is the individual bias correction term for γ_A .

Let $\tilde{\sigma} \in \{\sigma, \sigma^*\}$ be a Markov policy. Suppose the state vector w has a stationary distribution. Then the bias correction term of the value function $V(w; \theta; \sigma; \eta)$ for the γ_P is equal to:

$$\alpha_P^\sigma(D; \theta; \gamma_P) := \Gamma_{\tilde{\sigma}}(w; \theta)(1_{\{a_P=1\}} - \gamma_P(w)),$$

where

$$\Gamma_{\tilde{\sigma}}(w; \theta) = \frac{1}{1 - \beta}(\delta_0 + (\delta_1 - \delta_0)\gamma^{\tilde{\sigma}}(w)),$$

where $\gamma^{\tilde{\sigma}}(w) = \mathbb{P}(\tilde{\sigma}(w, \epsilon) = 1|w)$ is the conditional probability of Apple's decision under the policy $\tilde{\sigma}$. The bias correction term for Apple's conditional choice probability γ_A is

$$\begin{aligned} \alpha_A^{CCP}(D; \theta; \gamma) &= -\mu s + R + (\delta_1 - \delta_0)\gamma_A(w) + \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \eta_0)|w, a = 1] \\ &\quad - \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \eta_0)|w, a = 0]. \end{aligned}$$

To sum up, the bias correction term $\alpha(D; \theta; \eta)$ is equal to:

$$\alpha(D; \theta; \eta) := -\alpha_A^{CCP}(D; \theta; \gamma) + \frac{1}{1-\beta}(\delta_1 - \delta_0)(\gamma_0^\sigma(w) - \gamma_{A,0}(w))(1_{\{a_P=1\}} - \gamma_P(w)).$$

In some point-identified problems the value function $V(w; \theta; \sigma; \eta)$ is only evaluated at the optimal strategy σ^* and the true parameter value θ_0 . Then the bias correction term is evaluated only at the true value θ_0 . In particular, in the Examples 1 and 2 the function $\Gamma(w, \theta_0)$ can be further simplified as follows:

$$\begin{aligned} \Gamma(w, \theta_0) &= \underbrace{-\mu_0 s + R_0 + \beta \mathbb{E}[V(w_{next}; \theta_0; \sigma^*; \gamma_0)|w, a = 1] - \beta \mathbb{E}[V(w_{next}; \theta_0; \sigma^*; \gamma_0)|w, a = 0]}_{v(1,w) - v(0,w)} \\ &= -\frac{2}{\gamma(w)} - \log \frac{\gamma(w)}{1 - \gamma(w)} \\ &= \log \frac{\gamma(w)}{1 - \gamma(w)} + \left(-\frac{2}{\gamma(w)} - \log \frac{\gamma(w)}{1 - \gamma(w)}\right) \\ &= -\frac{2}{\gamma(w)}, \end{aligned}$$

where we have used (7) in the second line.

2.2 Overview of the Asymptotic Results

I will now introduce the estimator of the identified set Θ_I , leaving the formal definition to (37). Suppose there exists a function $g(D, \theta, \xi)$ that preserves the expectation of the moment function $m(w, \theta, \eta)$

$$\mathbb{E}g(D, \theta, \xi) = \mathbb{E}m(w, \theta, \eta)$$

and is insensitive to the biased estimation of its own nuisance parameter ξ around its true value $\xi_0 = \xi_0(\theta)$, where $\xi_0(\theta)$ is an identified vector-valued parameter of the distribution P_D for each $\theta \in \Theta$. In many relevant cases such as Example 1, ξ_0 contains the original nuisance parameter η , but may contain more unknown parameters of the distribution P_D .

To estimate the identified set Θ_I I represent it as the minimizer of the criterion function $Q(\theta, \xi_0)$

$$\Theta_I := \arg \min_{\theta \in \Theta} Q(\theta, \xi_0), \tag{16}$$

where $Q(\theta, \xi_0)$ is

$$Q(\theta, \xi_0) = \|\mathbb{E}g(D, \theta, \xi_0)\|_+^2,$$

and its sample analog is

$$Q_N(\theta, \hat{\xi}) := \left\| \frac{1}{N} \sum_{i=1}^N g(D_i, \theta, \hat{\xi}_i) \right\|_+^2. \quad (17)$$

My goal is to use different samples in the first and the second stages in order to avoid overfitting. Yet, simple sample splitting has a drawback that only one half of the sample is used for second-stage estimation, which can lead to loss of efficiency in small samples. In order to use the whole sample for the second stage yet keep the sample splitting idea, I use cross-fitting procedure described below.

Definition 1 (Cross-fitting). *1. For a random sample of size N , denote a K -fold random partition of the sample indices $[N] = \{1, 2, \dots, N\}$ by $(J_k)_{k=1}^K$, where K is the number of partitions and the sample size of each fold is $n = N/K$. Also for each $k \in [K] = \{1, 2, \dots, K\}$ define $J_k^c = \{1, 2, \dots, N\} \setminus J_k$.*

2. For each $k \in [K]$, construct an estimator $\hat{\xi}(V_{i \in J_k^c})$ of the nuisance parameter value ξ_0 using only the data from J_k^c . For any observation $i \in J_k$, define an estimated signal $\hat{\xi}_i := \hat{\xi}(V_{i \in J_k^c})$.

Definition 2 (Definition of the Set Estimator). *Let $\hat{\xi} = \hat{\xi}(\theta)$ be the first-stage estimator of the nuisance parameter constructed in Definition 1. Let the criterion function $Q_N(\theta, \xi)$ be as in (17) and \hat{c} be a positive number. The estimator $\hat{\Theta}_I$ of the identified set Θ_I is chosen as a contour set of the sample criterion function $Q_N(\theta, \xi)$ of level c :*

$$\hat{\Theta}_I := \mathcal{C}_N(\hat{c}, \hat{\xi}) := \{\theta \in \Theta, \quad NQ_N(\theta, \hat{\xi}) \leq \hat{c}\}.$$

The contour level \hat{c} , possibly data dependent, is chosen such that $\hat{\Theta}_I$ contains the true set Θ_I with probability approaching one

$$\hat{c} \geq \sup_{\theta \in \Theta_I} NQ_N(\theta, \hat{\xi}) \text{ w.p. } \rightarrow 1. \quad (18)$$

I establish convergence and inference properties of the contour set estimator $\hat{\Theta}_I$. The first property is formulated in terms of the convergence rate of $\hat{\Theta}_I$ to Θ_I is based on the notion of Hausdorff distance $d_H(\hat{\Theta}_I, \Theta_I)$:

$$d_H(\hat{\Theta}_I, \Theta_I) = \left\{ \sup_{\theta \in \hat{\Theta}_I} d_H(\theta, \Theta_I), \sup_{\theta \in \Theta_I} d_H(\hat{\Theta}_I, \theta) \right\}.$$

The set $\hat{\Theta}_I$ is said to converge to Θ_I at rate ϵ_N if the Hausdorff distance between the sets converges at rate ϵ_N :

$$d_H(\hat{\Theta}_I, \Theta_I) = O_P(\epsilon_N). \quad (19)$$

Under mild regularity conditions it is possible to achieve the nearly efficient rate $\epsilon_N = O(\sqrt{\frac{\log N}{N}})$.

To conduct inference, I fix a confidence level $\tau \in (0, 1)$. A confidence region $C_N(c_\tau, \hat{\xi})$ of level τ is defined as a contour set $C_N(c_\tau; \hat{\xi})$ of level c_τ such that $C_N(c_\tau; \hat{\xi})$ contains Θ_I with probability at least $1 - \tau$:

$$P(\Theta_I \subseteq C_N(c_\tau; \hat{\xi})) \rightarrow 1 - \tau, N \rightarrow \infty. \quad (20)$$

I construct a confidence region $C_N(\hat{c}_\tau; \hat{\xi})$, where \hat{c}_τ is a consistent estimator of the τ -quantile, denoted c_τ , of the inferential statistic:

$$C_N := \sup_{\theta \in \Theta_I} NQ_N(\theta; \xi_0). \quad (21)$$

The consistent estimator of the τ -quantile is chosen by the subsampling algorithm defined below.

Definition 3 (Subsampling Algorithm). *Partition the sample $(D_i)_{i=1}^N$ into $B_N = o(\sqrt{N})$ subsamples V_j , $j \in \{1, 2, \dots, B_N\}$ of equal size $b := \lceil N/B_N \rceil$. Compute the sample criterion function $Q_{j,b}(\theta; \hat{\xi}) = \|\frac{1}{b} \sum_{i \in V_j} g(D_i, \theta, \hat{\xi}_i)\|_+^2$. Choose the level \hat{c} of the order $\hat{c} \sim \log N$. Report \hat{c}_τ as the τ -quantile of the sample of statistics*

$$\left\{ \sup_{\theta \in C_N(\hat{c}, \hat{\xi})} bQ_{j,b}(\theta, \hat{\xi}), j = 1, 2, \dots, B_N \right\}.$$

Section 4 establishes the asymptotic validity of the set estimator given in Definition 2 and subsampling algorithm of Definition 3

3 Dynamic Game of Imperfect Information

Consider the dynamic model of strategic interaction from Bajari et al. (2007). There are K players, denoted by $\{1, 2, \dots, K\}$. Each player k makes a decision $a_k \in \mathcal{A}_k$ from a finite set of discrete alternatives \mathcal{A}_k at time periods $t \in \{0, 1, \dots, \infty\}$. In each period t the players commonly observe a vector of state variables $w_t \in \mathcal{W} \subset \mathcal{R}^{d_w}$. Given the state variable $w_t = w$, players choose actions simultaneously. Before choosing his action, each player k observes a vector of private shocks $(\epsilon_k(j))_{j \in \mathcal{A}_k}$ corresponding to each discrete alternative j in his choice set \mathcal{A}_k . The transition between states follows a conditional probability distribution: $P(\cdot | a, w)$ conditional on the current state w and the action profile $a = (a_1, a_2, \dots, a_K)$.

I focus on the structural parameter describing the utility of the first player. I assume that his per-period utility function is given by

$$\pi(a, w, \epsilon) = \tilde{\pi}(a, w; \theta_0; \zeta_0) + \epsilon_1(a_1), \quad (22)$$

where $a = (a_1, a_2, \dots, a_K)$ is the profile of the players' actions, $a_1 \in \mathcal{A}_1$ is the action of the first player, w is the state variable, and θ_0 is the true value of the structural parameter θ . The per-period utility is presented as the sum of a deterministic component $\tilde{\pi}(a, w; \theta; \zeta)$ and the private shock $\epsilon_1(a_1)$. I allow $\tilde{\pi}(a, w; \theta; \zeta)$ to depend on identified nuisance parameter ζ whose true value is ζ_0 .

I assume that each player follows a pure Markov policy. The pure Markov policy for player one $\sigma_1(w, \epsilon_1) : \mathcal{W} \times \mathcal{E} \rightarrow \mathcal{A}_1$ maps the current state w and the private shock of player one, ϵ_1 , into the action space \mathcal{A}_1 . If the behavior of the players is described by a Markov policy profile $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_K)$, the value function is given by

$$V(w; \theta; \sigma) := \mathbb{E}\left[\sum_{t \geq 0} \beta^t \pi(\sigma(w, \epsilon), w_t, \epsilon_t) | w\right]. \quad (23)$$

A strategy profile σ^* is a Markov perfect equilibrium if each player k prefers its strategy σ_k^* to all alternative Markov strategies as long as the others follow the equilibrium strategy σ_{-k}^* . That is, the value function $V(w; \theta_0; \sigma^*; \eta_0)$ of the first player at the strategy σ^* is weakly larger than the value function of any other strategy profile $\sigma = (\sigma_1, \sigma_{-1}^*)$:

$$V(w; \theta_0; \sigma^*; \eta_0) \geq V(w; \theta_0; \sigma; \eta_0) \quad \forall w \quad \forall \sigma, \quad (24)$$

where $\sigma = (\sigma_1, \sigma_{-1}^*)$ consists of a feasible suboptimal alternative for player one, σ_1 , and the equilibrium profile for the other players σ_{-1}^* . I assume that the each observation comes from the same Markov perfect equilibrium σ^* , although I do not provide conditions for the existence of such equilibrium and recognize that there could be many such equilibria.

ASSUMPTION 2 (Equilibrium Selection). *The data are generated by a single Markov perfect equilibrium σ^* .*

3.1 Casting problem as a moment inequality with a first-stage nuisance parameter

Define the choice-specific value function of player one $v(a_1, w)$ as the expected present value conditional on the current state w and the choice a_1

$$v(a_1, w) := \tilde{\pi}(a_1, w; \theta_0, \zeta_0) + \beta \mathbb{E}[V(w_{next}; \theta_0; \sigma^*; \eta_0) | w, (a_1, \sigma_{-1}^*)].$$

Then, following the optimal strategy $\sigma^*(w, \epsilon)$, the first player chooses a_1 if and only if the total utility of a_1 is not smaller than the total utility of any other choice $a'_1 \in \mathcal{A}_1$

$$v(a_1, w) + \epsilon(a_1) \geq v(a'_1, w) + \epsilon(a'_1).$$

Therefore, the optimal strategy of player one is characterized as

$$\sigma^*(w, \epsilon_1) = \arg \max_{a'_1 \in \mathcal{A}_1} \{v(a'_1, w) + \epsilon_1(a'_1)\},$$

or, equivalently,

$$\sigma^*(a_1, w) = \arg \max_{a'_1 \in \mathcal{A}_1} \{v(a'_1, w) - v(1, w) + \epsilon_1(a'_1)\}.$$

As discussed in Hotz and Miller (1993), under Assumption 1 the vector of differences of the choice-specific value functions can be expressed as

$$v(a_1, w) - v(1, w) = \log \frac{P(a_1|w)}{P(1|w)}, \quad (25)$$

where $P(a_1|w)$ is the probability of the choice $a_1 \in \mathcal{A}_1$ conditional on the state variable w . Finally, as a suboptimal Markov policy σ , I consider a cutoff-type strategy

$$\sigma_1(w, \epsilon_1) = \arg \max_{a_1 \in \mathcal{A}} \{v(a_1, w) - v(1, w) + \text{dev}(a_1, w) + \epsilon_1(a_1)\}, \quad (26)$$

where I add a deviation function $\text{dev}(a_1, w)$ for each element $v(a_1, w) - v(1, w)$. As a normalization condition, I set $\text{dev}(1, w) = 0 \quad \forall w$. Therefore, the inequality (24) depends on the nuisance parameter

$$\eta := (\zeta, \{\gamma_{jk}(w), j \in \mathcal{A}_k\}_{k=1}^K, P(w_{next}|w, a)) \quad (27)$$

that consists of any original nuisance parameter ζ that may appear in the utility function (22), the conditional choice probabilities of all players

$$\gamma_{jk}(w) := P(a_k = j|w), \quad j \in \mathcal{A}_k, k \in \{1, 2, \dots, K\},$$

and the conditional distribution $P(w_{next}|w, a)$.

I construct the identified set Θ_I using a subset of inequalities implied by the equilibrium definition (i.e, (24)). Let $q(w)$ be an L -vector of non-negative weighting functions. Let $\text{dev}(a_1, w) : \mathcal{A}_1 \times \mathcal{W} \rightarrow \mathcal{R}^d$ be an L -vector of deviation functions whose coordinate $l \in \{1, 2, \dots, L\}$ corresponds to a deviation strategy $\text{dev}_l(a_1, w)$. Define the moment vector-valued function as

$$m(w, \theta, \eta) = q(w) \cdot (V(w; \theta; \sigma; \eta) - V(w; \theta; \sigma^*; \eta)), \quad (28)$$

whose component $m_l(D, \theta, \eta)$ is equal to the weighted difference $V(w; \theta; \sigma_l; \eta) - V(w; \theta; \sigma^*; \eta)$:

$$m_l(w, \theta, \eta) := q_l(w)(V(w; \theta; \sigma_l; \eta) - V(w; \theta; \sigma^*; \eta)), \quad l \in \{1, 2, \dots, L\}$$

evaluated at a strategy profile $\sigma_l = (\sigma_{1,l}, \sigma_{-1}^*)$. The suboptimal strategy of the first player $\sigma_{1,l}$ is given in (26) with the deviation function $\text{dev}_l(a_1, w)$. The identified set Θ_I is defined as the collection of parameter values θ that obey inequality restrictions in expectation

$$\Theta_I := \{\theta \in \Theta : \mathbb{E}m(w, \theta, \eta_0) \leq 0\}, \quad (29)$$

where the expectation is taken with respect to the unconditional distribution of the state w .

3.2 Simulation estimator of the value function

The value function $V(w; \theta; \sigma; \eta)$ appearing in the moment function (28) is not available in closed form. However, it can be approximated by Monte Carlo simulation as in Bajari et al. (2007). In the first stage, one constructs an estimate $\hat{\eta}$ of the nuisance parameter η that is defined in (27). In the second stage, one simulates the sequence of states and shocks from the estimated first stage parameter $\hat{\eta}$ and averages the realization of the value function over multiple simulation draws. A single simulation draw is given by Algorithm 1.

Algorithm 1 Simulation Estimator of the Expected Value Function

Input: initial state w ; parameter value θ ; estimated first-stage parameter $\hat{\eta}$; strategy profile $\sigma(w, \epsilon)$ that is a known function of the first-stage parameter $\hat{\eta}$. Initialize $V(w_0; \theta; \sigma; \eta) = 0$.

- 1: Draw a shock vector ϵ from the type 1 extreme value distribution and compute the action profile $a = \sigma(w, \epsilon)$.
- 2: Draw the state variable w_t from the conditional distribution $P(w_t | w_{t-1}, a_{t-1})$.
- 3: Draw a vector of shocks ϵ_t from the type 1 extreme value distribution. Compute the action $a_t = \sigma(w_t, \epsilon_t)$.
- 4: Add the time t discounted utility:

$$V(w_0; \theta; \sigma; \hat{\eta}) = V(w_0; \theta; \sigma; \hat{\eta}) + \beta^t (\tilde{\pi}(a_t, w_t; \theta; \hat{\zeta}) + \epsilon_1(a_{1,t})).$$

Return $\hat{V}(w_0; \theta; \sigma; \hat{\eta})$.

In order to evaluate the value function $V(w; \theta; \sigma; \eta)$ at different parameter values $\theta_1 \in \Theta$ and $\theta_2 \in \Theta$, I use the same simulation draws. As long as the number of simulation draws is large enough,

the simulation error does not affect the asymptotic properties of the estimator of the identified set $\hat{\Theta}_I$ as discussed in Pakes and Pollard (1989).

3.3 Bias Correction Term for the Expected Value Function

When the nuisance parameter η is high-dimensional and estimated by machine learning, the plug-in estimator of value function is biased. To make the moment function (28) insensitive to first-stage bias, I derive bias correction term. As discussed in Newey (1994), the bias correction term for a vector-valued nuisance parameter is equal to the sum of individual terms of individual components. Moreover, the bias correction term for the conditional probability of choice j has the product structure

$$\alpha_j^{CCP}(D; \theta; \gamma_j) = \Gamma(w; \theta)(1_{\{a_1=j\}} - \gamma_j(w)).$$

Furthermore, according to Newey (1994), the function $\Gamma(w; \theta)$ is implicitly defined by the orthogonality condition explained below.

Let $g(D, \theta, \xi)$ be a moment function. Define the Gateaux derivative map $D_r : \Xi \times \Theta \rightarrow \mathcal{R}^L$ as

$$\partial_r \left\{ \mathbb{E} \left[g(D, \theta, r(\xi - \xi_0) + \xi_0) \right], \quad \xi \in \Xi \right\},$$

for all $r \in [0, 1)$, which I assume exists. I also denote the pathwise derivative of the expected moment function at the true value ξ_0

$$\partial_\xi \mathbb{E} g(D, \theta, \xi_0)[\xi - \xi_0] = \partial_0 \mathbb{E} g(D, \theta, r(\xi - \xi_0) + \xi_0).$$

Definition 4 (Neyman orthogonality of moment function). *The moment function $g(D; \theta; \xi)$ obeys the orthogonality condition at ξ_0 with respect to the nuisance realization set $\Xi_N \subset \Xi$ if the pathwise derivative $D_r[\xi - \xi_0]$ exists for all $r \in [0, 1)$ and vanishes at $r = 0$ for each $\theta \in \Theta$*

$$\partial_\xi \mathbb{E} g(D, \theta, \xi_0)[\xi - \xi_0] = 0, \quad \xi \in \Xi, \theta \in \Theta. \quad (30)$$

The definition of Neyman orthogonality requires that the moment function be insensitive to the biased estimation of ξ . This condition is the generalization of the orthogonality condition for point-identified models in Chernozhukov et al. (2017a) that is required to hold only at the true value θ_0 of identified parameter θ . In contrast to the point-identified case, I require the equality in (30) to hold at each point θ of the parameter space Θ . In many relevant cases (e.g., if the moment function $g(D, \theta, \xi)$ is linear in θ), the orthogonality condition (30) on the set Θ follows from the

orthogonality (30) on a finite subset of Θ . Rewriting the orthogonality condition (30) for the bias correction term for the value function gives

$$\partial_0 \mathbb{E}[V(w; \theta; \sigma^*; r(\gamma - \gamma_0) + \gamma_0) - \Gamma(w; \theta)[\gamma(w) - \gamma_0(w)]] = 0.$$

I find the function $\Gamma(w; \theta)$ from the recursive definition of the value function (5). In the case of Example 1 the recursive definition can be rewritten in the unconditional form

$$\begin{aligned} \mathbb{E} \left[V(w; \theta; \sigma^*; \gamma_0) - \left(-R(1 - \gamma_0(w)) + (-\mu s)\gamma_0(w) + PS_{\sigma^*}(\gamma_0) \right. \right. \\ \left. \left. + \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma_0) | w, a = 0](1 - \gamma_0(w)) \right. \right. \\ \left. \left. + \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma_0) | w, a = 1]\gamma_0(w) \right) \right] = 0, \end{aligned} \quad (31)$$

where $\gamma_0(w)$ is the true value of the conditional choice probability and $\theta = (R, \mu)$ consists of the replacement and maintenance costs. The unknown function $\gamma(w)$ appears in Equation (31) both **outside** the value function $V(w; \theta; \sigma^*; \gamma)$ and *inside* of this function. I consider a local deviation of the choice probability $\gamma(\cdot)$ from its value $\gamma_0(\cdot)$ for each value of the parameters R and μ . Applying chain rule to the Equation (31) yields a pathwise derivative

$$\begin{aligned} \partial_0 \mathbb{E}V(w; \theta; \sigma^*; r(\gamma - \gamma_0) + \gamma_0) &= \frac{1}{1 - \beta} \mathbb{E} \left(-\mu s + R + \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma_0) | w, a = 1] \right. \\ &\quad \left. - \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma_0) | w, a = 0] + \frac{dPS_{\sigma^*}(\gamma_0)}{d\gamma} \right) \\ &= \frac{1}{1 - \beta} \mathbb{E} \left(-\mu s + R + \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma_0) | w, a = 1] \right. \\ &\quad \left. - \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma_0) | w, a = 0] - \frac{2}{\gamma(w)} - \log \frac{\gamma(w)}{1 - \gamma(w)} \right), \end{aligned}$$

where the last equality follows from the logistic distribution of the private shocks (Assumption 1).

In what follows I derive the bias correction terms for the other nuisance parameters that are present in the dynamic discrete choice in Bajari et al. (2007). Let $q(w)$ be a weighting function. Define

$$\lambda(w') := \mathbb{E}[q(w) | w_{next} = w']$$

as the expectation of the weighting function $q(w)$ evaluated at the current state w conditional on the future state w_{next} . Denote the expectation of the choice j made by player one conditional on the state w by $\gamma_j(w) := \mathbb{E}[1_{\{a_1=j\}} | w]$. Let

$$\gamma(w) := (\gamma_2(w), \gamma_3(w), \dots, \gamma_{A_1}(w))$$

be the $A_1 - 1$ vector of these probabilities. Define the conditional expectation of the current private shock evaluated at the equilibrium strategy $\sigma_1^* = \sigma_1^*(w, \epsilon, \gamma)$ as

$$PS_{\sigma_1^*}(\gamma) := \mathbb{E}[\epsilon_1(\sigma_1^*(w, \epsilon_1, \gamma))|w]. \quad (32)$$

Lemma 1 gives the bias correction term, $\alpha_j^{CCP}(D; \theta; \gamma_j)$, that makes the value function $V(w; \theta; \sigma^*; \eta)$ insensitive to the biased estimation of γ_j .

Lemma 1 (Bias correction term for own conditional choice probability). *Suppose the state variable w has a stationary distribution and Assumption 1 holds. Suppose the conditional probability of each choice $\gamma_j(w)$, $j \in \{2, 3, \dots, J\}$ is bounded away from zero and one. Then the bias correction term $\alpha_j^{CCP}(D; \theta; \gamma_j)$ is*

$$\begin{aligned} \alpha_j^{CCP}(D; \theta; \gamma_j) &= \frac{q(w)}{q(w) - \beta\lambda(w)} (\mathbb{E}_{\epsilon_{-1}}[\tilde{\pi}((j, \sigma_{-1}^*(w, \epsilon_{-1})); w; \theta) - \tilde{\pi}(1, \sigma_{-1}^*(w, \epsilon_{-1})); w; \theta]) \\ &\quad + \beta\mathbb{E}[V(w_{next}; \theta; \sigma^*; \eta_0)|w, (j, \sigma_{-1}^*(w, \epsilon_{-1}))] \\ &\quad - \beta\mathbb{E}[V(w_{next}; \theta; \sigma^*; \eta_0)|w, (1, \sigma_{-1}^*(w, \epsilon_{-1}))] \\ &\quad + \partial_{\gamma_j} PS_{\sigma_1^*}(\gamma_0)(1_{a_1=j} - \gamma_j(w)) \end{aligned} \quad (33)$$

where $\partial_{\gamma_j} PS_{\sigma_1^*}(\gamma)$ is the partial derivative of (32) with respect to γ_j .

Lemma 2 (Bias Correction Term for the Law of Motion of the State Variable). *Suppose the high-dimensional state variable w has a stationary distribution. Let $a = (a_1, a_2, \dots, a_K)$ be a given profile of actions. Then the bias correction term for the conditional quantile function $Q(u, w, a)$ of level u is*

$$\begin{aligned} \alpha_{\sigma}^{TRANS}(D; \theta; \eta) &= \frac{\beta l(w)}{l(w) - \beta\lambda(w)} \mathbb{E}[\nabla_{w_{next}} V(w_{next}; \theta; \sigma; \eta_0)|w, a] \prod_{k=1}^K \mathbb{P}(\sigma_k(w, \epsilon_k) = a_k|w) \\ &\quad \frac{1_{w_{next} \leq Q(u, w, j) - u}}{f(w_{next}|w, a = \sigma(w, \epsilon))}. \end{aligned} \quad (34)$$

Lemma 3 (Bias Correction Term for the opponent's conditional choice probability). *Suppose the state variable w has a stationary distribution. Let the number of players be 2 (i.e., $K = 2$). Then*

the bias correction term for the conditional probability of the choice j_2 made by my opponent is

$$\begin{aligned} \alpha_j^{CCP,op}(D; \theta; \gamma_{j_2}) &= \frac{q(w)}{q(w) - \beta\lambda(w)} \left(\sum_{j_1=1}^{A_1} (\tilde{\pi}((j_1, j_2), w; \theta) - \tilde{\pi}((j_1, 1), w; \theta)) \mathbb{P}(\sigma_1(w, \epsilon_1) = j_1 | w) \right. \\ &\quad + \beta \left[\sum_{j_1=1}^{A_1} \mathbb{E}[V(w_{next}; \theta; \sigma^*; \eta_0) | w, (j_1, j_2)] - \sum_{j_1=1}^{A_1} \mathbb{E}[V(w_{next}; \theta; \sigma^*; \eta_0) | w, (j_1, 1)] \right] \cdot \\ &\quad \cdot \mathbb{P}(\sigma_1(w, \epsilon_1) = j_1 | w) (1_{a_2=j_2} - \mathbb{P}(a_2 = j_2)) \Big). \end{aligned} \quad (35)$$

3.4 Using Linearity to Reduce Computation

Both the value function and the bias correction terms presented above are not available in closed form and must be simulated. When the value function is a linear function of θ , the simulation can be simplified. Suppose there exist basis functions $\Psi_1(w; \sigma; \eta)$, $\Psi_2(w; \sigma; \eta)$ such that

$$V(w; \theta; \sigma; \eta) = \theta \Psi_1(w; \sigma; \eta) + \Psi_2(w; \sigma; \eta) \quad (36)$$

is an affine function of θ . Then one can simulate the basis functions $\Psi_1(w; \sigma; \eta)$, $\Psi_2(w; \sigma; \eta)$ instead of simulating the value function for each $\theta \in \Theta$. Lemma 4 provides the sufficient conditions for the linearity of the value function and the individual bias correction terms provided in Lemmas 1, 2, 3.

Lemma 4 (Sufficient Conditions for Linearity). *The following conditions hold. (1) The per-period utility function given in (22) is a linear function of θ . (2) The distribution of the private shock for each player is known. Then there exists a vector of basis functions $\Psi(w; \sigma; \eta)$ such that (36) holds. Furthermore, each individual bias correction term given in (33), (34), and (35) is also a linear function θ .*

4 Asymptotic Theory

Suppose I have a collection of inequality restrictions on an economic model coming from the data structure and/or the assumptions about the data generating process. These restrictions are embodied into a moment function $g(D, \theta, \xi) : \mathcal{D} \times \Theta \times \Xi \rightarrow \mathcal{R}^L$, where $\theta \in \Theta \subset \mathcal{R}^d$ is the target parameter. In addition to the target parameter θ , the moment function $g(D, \theta, \xi)$ depends on a nuisance parameter $\xi = \xi(\theta)$ whose true value $\xi_0 = \xi_0(\theta)$ is an identified parameter of the data distribution P_D and belongs to a convex subset of a normed vector space Ξ . The object of interest is an identified set Θ_I defined as a collection of parameter values θ that satisfy the inequality

restrictions

$$\Theta_I := \{\theta \in \Theta : \mathbb{E}g(D, \theta, \xi_0) \leq 0\} \quad (37)$$

at the true value ξ_0 of the nuisance parameter.

The identified set Θ_I is characterized as the minimizer of the criterion function

$$\Theta_I := \arg \min_{\theta \in \Theta} Q(\theta, \xi_0) = \arg \min_{\theta \in \Theta} \|\mathbb{E}g(D, \theta, \xi_0)\|^2.$$

I assume that the following partial identification condition holds. There exist positive constant $C > 0$ and $\delta > 0$ such that

$$\|\mathbb{E}g(D, \theta, \xi_0)\|_+ \geq C d_H(\theta, \Theta_I) \wedge \delta. \quad (38)$$

This condition states that once θ is bounded away from Θ , the moment $\mathbb{E}g(D, \theta, \xi_0)$ is bounded away from Θ_I by a number that is proportional to the Hausdorff distance $d_H(\theta, \Theta_I)$ from θ to the identified set Θ_I . This condition ensures that the true moment function $g(D, \theta, \xi_0)$ distinguishes the boundary of the identified set.

Impact of the First-Stage Estimation. In the next condition I introduce a sequence of neighborhoods $\Xi_N^\theta \subset \Xi^\theta$ of $\xi_0(\theta)$ that contain the estimate $\hat{\xi}(\theta)$ of $\xi_0(\theta)$ w.p. approaching one. As the sample size N increases, the neighborhoods shrink. The quality of the estimation of the first-stage parameter is defined as the speed of shrinkage of the neighborhood $\Xi_N(\theta)$ around $\xi_0(\theta)$. I refer to it as the first-stage rate g_N . Finally, I assume that the second-order derivative of the functional $\mathbb{E}g(D, \theta, \xi)$ is well-behaved. Combined with the orthogonality condition and the upper bound on the first-stage rate g_N , this assumption ensures that I can ignore the impact of the estimation error $\hat{\xi}(\theta) - \xi_0(\theta)$ on the second and the higher-order derivatives of the moment function $\mathbb{E}g(D, \theta, \xi)$ with respect to ξ at ξ_0 .

Condition 1 (Orthogonality). *There exists a sequence Ξ_N^θ of subsets of Ξ^θ : $\Xi_N^\theta \subset \Xi^\theta$ such that the following conditions hold. (1) The true value $\xi_0(\theta)$ belongs to Ξ_N^θ for all $N \geq 1$. (2) There exists a sequence of numbers $\phi_N = o(1)$ such that with probability at least $1 - \phi_N$, $\hat{\xi}(\cdot)$ belongs to $\Xi := \times \Xi^\theta$ uniformly over $\theta \in \Theta$. (3) The set Ξ_N^θ shrinks around $\xi(\theta)$ at the following statistical rate $g_N = o(N^{-1/4})$ uniformly over θ :*

$$\sup_{\theta \in \Theta} \sup_{\xi(\theta) \in \Xi^\theta} \|\hat{\xi}(\theta) - \xi_0(\theta)\|_{P,2} \leq g_N.$$

The moment function $g(D, \theta, \xi)$ obeys the orthogonality condition at ξ_0 . (5) There exists a sequence $s_N = o(N^{-1/2})$ such that the second Gateaux derivative of the functional $G(\theta, \xi(\theta))$ with respect to ξ at ξ_0 is bounded:

$$\sup_{\theta \in \Theta} \sup_{\xi \in \Xi_N^\theta} \sup_{r \in [0,1]} \|\partial_r^2 \mathbb{E}g(D, \theta, r(\xi - \xi_0) + \xi_0)\| \leq s_N.$$

The next condition requires that the moment function $g(D; \theta; \xi(\theta))$ is sufficiently regular with respect to θ for each fixed element $\xi \in \Xi$ of the nuisance realization set Ξ . I consider a class

$$\mathcal{F}_\xi := \{g(\cdot, \theta, \xi(\theta)), \theta \in \Theta\}$$

and require the uniform covering entropy of this class to be bounded.

Condition 2 (Regularity of Moment Function). *The following conditions hold. (1) There exists a measurable envelope function $F_\xi = F_\xi(D)$ that bounds all elements in the function class almost surely*

$$\sup_{\theta \in \Theta} |g_l(D, \theta, \xi)| \leq F_\xi(D) \text{ a.s. , } l \in \{1, 2, \dots, L\}.$$

Moreover, the envelope F_ξ has a finite c -norm for some $c > 2$ $\|F_\xi\|_{P,c} := (\int_{D \in \mathcal{D}} |F_\xi(D)|^c)^{1/c} \leq c_1$.

(2) *There exist finite constants a and v such that the uniform covering entropy of the class \mathcal{F}_ξ is bounded*

$$\sup_{\tilde{Q}} \log N(\epsilon \|\mathcal{F}_\xi\|_{\tilde{Q},2}, F_\xi, \|\cdot\|_{\tilde{Q},2} \leq v \log(a/\epsilon), \text{ for all } 0 < \epsilon \leq 1. \quad (39)$$

(3) *There exists a sequence r'_N obeying $r'_N \log(1/r'_N) = o(1)$ that is an upper bound for the following quantity*

$$\sup_{\theta \in \Theta} \sup_{\xi \in \Xi^\theta} (\mathbb{E} \|g(D, \theta, \xi(\theta)) - g(D, \theta, \xi_0(\theta))\|^2)^{1/2} \leq r'_N.$$

Conditions 2(1)-(2) are the generalization of the regularity assumption in the point-identified moment problem of Chernozhukov et al. (2017a). Because $\xi(\theta)$ depends on θ , conditions 2(1)-(2) are non-standard and require verification in applications. However, this requirement is mild when the nuisance parameter $\xi(\theta)$ is a linear function of θ . Suppose that the true value of the nuisance parameter $\xi_0(\theta)$ is a linear function of θ

$$\xi_0(\theta) := \xi_0^a(D)' \theta + \xi_0^b(D), \quad (40)$$

where $\xi_0^a(D)$ and $\xi_0^b(D)$ are the identified parameters of the distribution P_D . Then conditions 2(1)-(2) can be reformulated in terms of the nuisance parameters $\{\xi_0^a, \xi_0^b\}$ that no longer depend

on θ . When (40) holds, Conditions 2 (1)-(2) are satisfied for many practical cases. In particular, the functions $\xi_0^a(D), \xi_0^b(D)$ can be estimated by ℓ_1 -regularized methods, random forests, and deep neural nets under plausible assumptions about their structure.

Donsker Property. Let Θ' be an open neighborhood of Θ . I require the moment function $g(D_i, \theta, \xi_0)$ to have a Donsker property defined as follows. In the metric space $L^\infty(\Theta')$,

$$\mathbb{G}_N g(D_i, \theta, \xi_0) := \sqrt{N}(\mathbb{E}_N g(D_i, \theta, \xi_0) - \mathbb{E}g(D_i, \theta, \xi_0)) \Rightarrow \Delta(\theta), \quad (41)$$

where $\Delta(\theta)$ is a mean zero Gaussian process on Θ with a.s. continuous paths and $\text{Var}(\Delta(\theta)) > 0$ for each $\theta \in \Theta'$. In addition, the probability space $(\Omega, \mathcal{F}, \mathcal{P})$ is rich enough to support the representation (41).

Theorem 5 (Estimation and Inference for Semiparametric Functional Inequalities). *Suppose Conditions (38), 1, 2, and (41) hold. Let $\hat{\Theta}_I$ be a contour set estimator of Definition 2. Let \hat{c} be such that*

$$\hat{c} \geq \sup_{\theta \in \Theta_I} NQ_N(\theta, \xi_0) \quad w.p. \rightarrow 1 \quad (42)$$

holds. Then the Hausdorff distance $d_H(\Theta_I, \hat{\Theta}_I)$ between the estimated set $\hat{\Theta}_I$ and Θ_I converges at rate $O_P(\sqrt{(1 \vee \hat{c})/N})$: $d_H(\hat{\Theta}_I, \Theta_I) = O_P(\sqrt{(1 \vee \hat{c})/N})$.

Theorem 5 is my first main result. It establishes the sufficient conditions on the moment function to deliver the rate of convergence of $\hat{\Theta}_I$ to Θ_I . It suggests that the contour level \hat{c} as small as possible subject to the constraint (42). Setting $\hat{c} = O_P(1)$ subject to (42) delivers the optimal rate, but this choice is infeasible. Setting $\hat{c} \sim \log N$ delivers a nearly efficient rate.

In many cases it is possible to establish convergence without the requirement (42). This is possible because the criterion function is degenerate (17).

Definition 5 (Degeneracy). *The following conditions hold. (1) There exists a sequence of subsets Θ_N of Θ , which cannot depend on ξ , such that the criterion function $Q_N(\theta, \xi)$ vanishes on these subsets w.p. approaching one. That is, $\forall p > 0$ there exists N_p such that for all $N \geq N_p$ $\inf_{\xi \in \Xi_N} \mathbb{P}(Q_N(\theta, \xi) - \inf_{\theta \in \Theta} Q_N(\theta, \xi) = 0 \quad \forall \theta \in \Theta) \geq 1 - p$. (2) These sets can approximate the identified set Θ_I in the Hausdorff distance sufficiently well: $d_H(\Theta_N, \Theta_I) \leq \epsilon_N$. (3) The sequence $\epsilon_N = O_P(N^{-1/2})$.*

Lemma 6 (Sufficient Conditions for Degeneracy). *Suppose Conditions (38), 1, 2, and (41) hold. In addition, there exist positive constants C, M, δ such that*

$$\begin{aligned} \max_l \mathbb{E} g_l(D, \theta, \xi_0) &\leq -C(\epsilon \wedge \delta) \quad \text{for all } \theta \in \Theta_I^{-\epsilon}, \\ d_H(\Theta_I^{-\epsilon}, \Theta_I) &\leq M\epsilon \quad \text{for all } \epsilon \in [0, \delta], \quad l \in \{1, 2, \dots, L\} \end{aligned} \quad (43)$$

Then the criterion function $Q_N(\theta, \xi)$ obeys the degeneracy condition in the sense of Definition 5. Suppose the contour level \hat{c} obeys

$$\hat{c} \geq \min_{\theta \in \Theta} Q_N(\theta, \hat{\xi}) \vee \frac{\log N}{\sqrt{N}} \quad \text{w.p.} \rightarrow 1. \quad (44)$$

Then the Hausdorff distance $d_H(\Theta_I, \hat{\Theta}_I)$ converges at rate $O_P(N^{-1/2})$, where $\hat{\Theta}_I$ is a contour level set as in Definition 2.

Lemma 6 is my second main result. It provides the sufficient conditions under which the identified set Θ_I can be estimated at the fastest possible rate $\epsilon_N = O_P(N^{-1/2})$. It requires that the sample criterion function $Q_N(\theta, \xi_0)$ be flat on the (possibly) data-dependent sets Θ_N that approximate the identified set Θ_I sufficiently well. If this requirements holds, the contour level $\hat{c}_0 := \arg \min_{\theta \in \Theta_0} Q_N(\theta, \xi_0)$ delivers the optimal rate for the contour level set $\mathcal{C}_N(\hat{c}_0, \xi_0)$ based on the true value of the nuisance parameter (see, e.g. Chernozhukov et al. (2007)). We show that a modified choice \hat{c} given in (44) delivers the optimal rate in the presence of the nuisance parameter $\hat{\xi}$ estimated in the first stage on an auxiliary sample.

Subsampling. I wish to construct the contour set $\mathcal{C}(\hat{c}, \hat{\xi})$ that has confidence region property (20). To do this, I must find the asymptotic distribution of the inferential statistic

$$\mathcal{C}_N := \sup_{\theta \in \Theta_I} Q(\theta; \xi_0)$$

that can be used to estimate the τ -quantile of \mathcal{C}_N . Define the random variable

$$\mathcal{C} := \sup_{\theta \in \Theta_I} \|I(\theta) + \Delta(\theta)\|_+^2, \quad (45)$$

where $I(\theta)$ is an L -vector of functions $I_l(\theta), l \in \{1, \dots, L\}$. The function $I_l(\theta)$ is defined as follows:

$$I_l(\theta) := \begin{cases} -\infty, & \mathbb{E} g_l(D, \theta, \xi_0) < 0, \\ 0, & \mathbb{E} g_l(D, \theta, \xi_0) = 0. \end{cases}$$

Chernozhukov et al. (2007) show that \mathcal{C}_N converges to \mathcal{C} in distribution, and that \mathcal{C} has a non-degenerate and continuous distribution function. The final requirement for the validity of the subsampling algorithm of Definition 3 is that the inferential statistic

$$\sup_{\theta \in \Theta_I^{\epsilon_N}} \mathcal{C}_N(\theta, \hat{\xi})$$

is well-behaved on the ϵ_N -expansion of the identified set Θ_I . A sufficient condition for this requirement is given below.

Condition 3 (Sufficient Conditions for Subsampling). (1) *There exists a constant $C_{\max} < \infty$ such that the moment function is bounded: $\sup_{\theta \in \Theta} |\mathbb{E}g(D, \theta, \xi_0)| \leq C_{\max} d_H(\theta, \Theta_I)$.* (2) *The rates s_N, r'_N and ϵ_N obey the following bound: $\sqrt{N}(s_N + r'_N \log(1/r'_N) + N^{-1/2+1/c})\epsilon_N = o(1)$.*

Theorem 7 (Validity of Subsampling for Moment Inequalities). *Suppose Conditions (38), 1, 2, (41), 3 hold. Suppose the number of subsamples satisfies $b = o(\sqrt{N})$, $b \rightarrow \infty$. Let τ be the desired coverage level. Then (1) the critical value \hat{c} of Definition 3 converges in probability to the τ -quantile of \mathcal{C} , where \mathcal{C} given in (45) and (2) $\mathbb{P}(\Theta_I \subseteq \mathcal{C}_N(\hat{c}, \hat{\xi})) = 1 - \tau$.*

5 Appendix

Notation. We use standard notation for numeric and stochastic dominance. For two numeric sequences denote $\{a_n, b_n\}, n \geq 1 : a_n \lesssim \sqrt{n}$ stands for $a_n = O(b_n)$. For two sequences of random variables denote $\{a_n, b_n, n \geq 1\} : a_n \lesssim_P \sqrt{n}$ stands for $a_n = O_P(b_n)$. Finally, let $a \wedge b = \min\{a, b\}$ and $a \vee b = \max\{a, b\}$. We use the standard notation for empirical processes. Denote the sample average $\mathbb{E}_N[f_i] := \frac{1}{N} \sum_{i=1}^N f(D_i)$ and the centered, scaled sample average as $\mathbb{G}_N[f_i] := \frac{1}{\sqrt{N}} \sum_{i=1}^N (f(D_i) - \mathbb{E}f(W_i))$. Denote the sample average within a partition I_k of size n as $\mathbb{E}_n[f(D_i)] := \frac{1}{n} \sum_{i \in I_k} f(D_i)$ and $\mathbb{G}_n[f(D_i)] := \frac{1}{\sqrt{n}} \sum_{i \in I_k} (f(D_i) - \mathbb{E}[f(D_i)])$. The Hausdorff distance between sets A and B is defined as $d_H(A, B) := \max\{\sup_{a \in A} d(a, B), \sup_{b \in B} d(b, A)\}$ where $d(b, A) := \inf_{a \in A} \|b - a\|$ and $d_H(A, B) = \infty$ if either of the sets is empty. The ϵ -expansion of the set Θ_I is defined as $\Theta_I^\epsilon := \{\theta \in \Theta : d(\theta, \Theta_I) \leq \epsilon\}$ and the ϵ -contraction of Θ_I is defined as $\Theta_I^{-\epsilon} := \{\theta \in \Theta_I : d(\theta, \Theta \setminus \Theta_I) \geq \epsilon\}$. Let $\|x\|_+ = \|\max(x, 0)\|$ and $\|x\|_- = \|\max(-x, 0)\|$, where the operation max is performed elementwise.

Proof of Remark 1. Step 1. Bias correction term for the conditional choice probability. I consider the value function $V(w; \theta; \sigma^*; \eta) = V(w; \theta; \sigma^*; \gamma)$ as a function of the current state variable

w , parameter θ , and the conditional choice probability $\gamma = \gamma(w)$ holding the other value of the nuisance component ρ at its true value ρ_0 . I wish to construct a bias correction term $\alpha^{CCP}(D; \theta; \gamma)$ such that orthogonality condition (30) holds:

$$\partial_\gamma \mathbb{E}[V(w; \theta; \sigma^*; \gamma_0) + \alpha^{CCP}(D; \theta; \gamma_0)] = 0 \quad \forall \theta \in \Theta.$$

Denote the expected value of the current private shock for the optimal strategy σ^* as

$$PS_{\sigma^*}(\gamma) := \mathbb{E}[\epsilon(\sigma^*(w, \epsilon, \gamma))|w],$$

where by logistic distribution of the private shocks I get:

$$PS_{\sigma^*}(\gamma) := \mathbb{E}(\epsilon(1) - \epsilon(0))1_{\{\epsilon(1) - \epsilon(0) \geq -\log \frac{\gamma(w)}{1-\gamma(w)}\}} = -\log \gamma(w) - (1 - \gamma(w)) \log \frac{1 - \gamma(w)}{\gamma(w)}. \quad (46)$$

I rewrite the recursive equation (5) for the Example 1:

$$\begin{aligned} V(w; \theta; \sigma^*; \gamma) &= -R(1 - \gamma(w)) + (-\mu s)\gamma(w) + PS_{\sigma^*}(\gamma) \\ &+ \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma)|w, a = 0](1 - \gamma(w)) + \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma)|w, a = 1]\gamma(w). \end{aligned} \quad (47)$$

I seek for the bias correction term $\alpha^{CCP}(D; \theta; \gamma)$ that has the following form:

$$\alpha^{CCP}(D; \theta; \gamma) = \Gamma(w; \theta)(1_{\{a=1\}} - \gamma(w)), \quad (48)$$

where $\Gamma(w; \theta)$ is determined by the following property:

$$\partial_0 \mathbb{E}V(w; \theta; \sigma^*; r(\gamma - \gamma_0) + \gamma_0) = \mathbb{E}\Gamma(w; \theta)[\gamma(w) - \gamma_0(w)]. \quad (49)$$

I find the function $\Gamma(w; \theta)$ by the application of the implicit function theorem to the recursive equation (47). I rewrite (47) in the unconditional form:

$$\begin{aligned} \mathbb{E} \left[\underbrace{V(w; \theta; \sigma^*; \gamma)}_{S_1(V(w; \sigma^*; \gamma))} - \underbrace{(-R(1 - \gamma(w)) + (-\mu s)\gamma(w) + PS_{\sigma^*}(\gamma))}_{S_2(\gamma)} \right. \\ \left. + \underbrace{\beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma)|w, a = 0](1 - \gamma(w))}_{S_3(V(w; \sigma^*; \gamma), \gamma)} + \underbrace{\beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma)|w, a = 1]\gamma(w)}_{S_4(V(w; \sigma^*; \gamma), \gamma)} \right] \\ = 0. \end{aligned}$$

Recognize that $\gamma(w)$ appears in each summand of the expression above: both inside the value function

$$S_1(V(w; \sigma^*; \gamma)), S_3(V(w; \sigma^*; \gamma), \gamma), S_4(V(w; \sigma^*; \gamma), \gamma)$$

and outside of it:

$$S_2(\gamma), S_3(V(w; \sigma^*; \gamma), \gamma), S_4(V(w; \sigma^*; \gamma), \gamma).$$

Taking the derivative with respect to the outside presence of γ gives:

$$\begin{aligned} \mathbf{S}_{out} &= \mathbb{E}(-\mu s + R + \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma)|w, a = 1] - \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma)|w, a = 0] \\ &\quad + \frac{dPS_{\sigma^*}(\gamma)}{d\gamma_0})[\gamma(w) - \gamma_0(w)]. \end{aligned}$$

Taking the derivative of the expected current shock $PS_{\sigma^*}(\gamma)$ defined in (46) with respect to γ gives:

$$\frac{dPS_{\sigma^*}(\gamma)}{d\gamma} = -\frac{2}{\gamma(w)} - \log \frac{\gamma(w)}{1 - \gamma(w)}.$$

Therefore, the pathwise (Gateaux) derivative with respect to the outside component is:

$$\mathbf{S}_{out} = \mathbb{E}\Gamma(w; \theta)[\gamma(w) - \gamma_0(w)].$$

Taking the derivative inside the value function requires some preparation. Recognize that:

$$\begin{aligned} \beta \mathbb{E}[S_3(V(w_{next}; \sigma^*; \gamma); \gamma_0)] &= \beta \mathbb{E}\mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma)|w, a = 0](1 - \gamma_0(w)) & (50) \\ &= \beta \mathbb{E}V(w_{next}; \theta; \sigma^*; \gamma) \frac{1_{\{a=1\}}}{1 - \gamma_0(w)} (1 - \gamma_0(w)) \\ &= \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma) 1_{\{a=1\}}] \\ &= \beta \mathbb{E}V(w_{next}; \theta; \sigma^*; \gamma) \lambda_0(w_{next}), \end{aligned}$$

where

$$\lambda_j(w_{next}) := \mathbb{E}[1_{\{a=j\}}|w_{next}], \quad j \in \mathcal{A}$$

is the expectation of the current choice $1_{\{a=j\}}$ conditional on the future state w_{next} . Assuming that the state variable has a stationary distribution, I get:

$$\mathbb{E}V(w_{next}; \theta; \sigma^*; \gamma) \lambda_0(w_{next}) = \mathbb{E}V(w; \theta; \sigma^*; \gamma) \lambda_0(w),$$

which implies:

$$\beta \mathbb{E}[S_3(V(w_{next}; \sigma^*; \gamma); \gamma_0)] = \mathbb{E}V(w; \theta; \sigma^*; \gamma) \lambda_0(w). \quad (51)$$

The same argument applied to S_4 delivers:

$$\beta \mathbb{E}[S_4(V(w_{next}; \sigma^*; \gamma); \gamma_0)] = \mathbb{E}V(w; \theta; \sigma^*; \gamma) \lambda_1(w). \quad (52)$$

Combining the (51) and (52) and recognizing that $\sum_{j \in \mathcal{A}} \lambda_j(w) = 1$ gives:

$$\beta \mathbb{E}[S_3(V(w_{next}; \sigma^*; \gamma); \gamma_0) + S_4(V(w_{next}; \sigma^*; \gamma); \gamma_0)] = \beta \mathbb{E}V(w; \theta; \sigma^*; \gamma).$$

Therefore,

$$\partial_\gamma \mathbb{E}[S_3(V(w_{next}; \sigma^*; \gamma); \gamma_0) + S_4(V(w_{next}; \sigma^*; \gamma); \gamma_0)] = \beta \partial_\gamma \mathbb{E}V(w; \theta; \sigma^*; \gamma).$$

Plugging this result into (47) gives:

$$\mathbf{S}_{ins} := (1 - \beta) \partial_0 \mathbb{E}V(w; \theta; \sigma^*; r(\gamma - \gamma_0) + \gamma_0).$$

According to the chain rule, the sum of the derivatives with respect to the inside and the outside presence of γ around γ_0 is equal to zero:

$$\mathbf{S}_{out} + \mathbf{S}_{ins} = 0,$$

which implies that

$$\partial_0 \mathbb{E}V(w; \sigma^*; r(\gamma - \gamma_0) + \gamma_0) = -\frac{1}{1 - \beta} \mathbb{E}\Gamma(w; \theta)[\gamma(w) - \gamma_0(w)].$$

Therefore, the function $\Gamma(w; \theta)$ is equal to:

$$\begin{aligned} \Gamma(w; \theta) &= -\mu s + R + \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma)|w, a = 1] - \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma)|w, a = 0] \\ &\quad - \frac{2}{\gamma(w)} - \log \frac{\gamma(w)}{1 - \gamma(w)} \end{aligned}$$

and the bias correction term for the conditional choice probability is equal to:

$$\alpha^{CCP}(D; \theta; \gamma) := -\frac{1}{1 - \beta} \Gamma(w; \theta) (1_{\{a=1\}} - \gamma(w)). \quad (53)$$

Bias correction term for the transition function. I consider the value function $V(w; \theta; \sigma; \eta)$ as a function of the current state w , the parameter θ and the transition function $\rho(\cdot)$ given in (34) holding the conditional probability of replacement $\gamma = \gamma_0$ fixed at its true value. Since the estimation of the nuisance parameter ρ affects the value function $V(w; \theta; \tilde{\sigma}; \eta)$ evaluated for both the optimal and a suboptimal strategy (i.e, $\tilde{\sigma} \in \{\sigma, \sigma^*\}$), I focus on a Markov policy $\tilde{\sigma}$ that includes both σ^* and σ as special cases. Let

$$PS_{\tilde{\sigma}}(w) := \mathbb{E}[\epsilon(\tilde{\sigma}(w, \epsilon))|w] \quad (54)$$

denote the expected current shock of the Markov policy $\tilde{\sigma}(w, \epsilon)$ conditionally on the current state w . Denote also by

$$\gamma_0^{\tilde{\sigma}}(w) := \mathbb{E}[\tilde{\sigma}(w, \epsilon)|w] \quad (55)$$

the conditional probability of choice $a = 1$ under the strategy σ . The recursive equation for the value function takes the form:

$$\begin{aligned} \mathbb{E} \left[\underbrace{V(w; \theta; \tilde{\sigma}; \rho)}_{R_1(V(w; \tilde{\sigma}; \rho))} - \left(-R(1 - \gamma_0^{\tilde{\sigma}}(w)) + (-\mu s)\gamma_0^{\tilde{\sigma}}(w) + PS_{\tilde{\sigma}}(w) \right. \right. \\ \left. \left. + \beta \underbrace{\mathbb{E}[V(w_{next}; \theta; \tilde{\sigma}; \rho)|w, \tilde{\sigma}(w, \epsilon) = 0]}_{R_3(V(w; \tilde{\sigma}; \rho))} (1 - \gamma_0^{\tilde{\sigma}}(w)) \right. \right. \\ \left. \left. + \beta \underbrace{\mathbb{E}_e[V(w_{next}; \theta; \tilde{\sigma}; \rho)|w, \tilde{\sigma}(w, \epsilon) = 1]}_{R_4(V(w; \tilde{\sigma}; \rho); \rho)} \gamma_0^{\tilde{\sigma}}(w) \right) \right] = 0 \end{aligned} \quad (56)$$

Recognize that the function ρ appears in each of the expressions

$$R_1(V(w; \tilde{\sigma}; \rho)), R_3(V(w; \tilde{\sigma}; \rho)), R_4(V(w; \tilde{\sigma}; \rho); \rho)$$

inside the value function. In the expression $R_4(V(w; \tilde{\sigma}; \rho); \rho)$, it also appears outside the value function, that is, through the conditional distribution of $P(w_{next}|w, a = 1)$. The derivative with respect to the outside presence of ρ is equal to:

$$\begin{aligned} \mathbf{R}_{out} &:= \partial_0 \mathbb{E} R_4(V(w; \tilde{\sigma}; \rho); r(\rho - \rho_0) + \rho_0) \\ &= \partial_0 \mathbb{E}_e V(r(\rho(w) - \rho_0(w)) + \rho_0(w) + e; \theta; \tilde{\sigma}; \rho_0) \gamma_0^{\tilde{\sigma}}(w) \\ &= \mathbb{E}_e \left[\frac{dV(x; \theta; \tilde{\sigma}; \rho_0)}{dx} \Big|_{x=\rho_0(w)+e} \right] \gamma_0^{\tilde{\sigma}}(w) \\ &= \mathbb{E} \left[\frac{dV(x; \theta; \tilde{\sigma}; \rho_0)}{dx} \Big|_{x=w_{next}|w, a=1} \right] \gamma_0^{\tilde{\sigma}}(w). \end{aligned}$$

The derivative with respect to the inside presence of ρ requires some preparation. For any Markov policy σ I see that:

$$\begin{aligned} \beta \mathbb{E} R_4(V(w_{next}; \theta; \tilde{\sigma}; \rho), \rho_0) &= \beta \mathbb{E}[V(w_{next}; \theta; \tilde{\sigma}; \rho)|w, \tilde{\sigma}(w, \epsilon) = 1] \gamma_0^{\tilde{\sigma}}(w) \\ &= \beta \mathbb{E} V(w_{next}; \theta; \tilde{\sigma}; \rho) \frac{\mathbf{1}_{\tilde{\sigma}(w, \epsilon)=1}}{\gamma_0^{\tilde{\sigma}}(w)} \gamma_0^{\tilde{\sigma}}(w) \\ &= \mathbb{E} V(w_{next}; \theta; \tilde{\sigma}; \rho) \lambda_1^{\tilde{\sigma}}(w_{next}) \\ &= \mathbb{E} V(w; \theta; \tilde{\sigma}; \rho) \lambda_1^{\tilde{\sigma}}(w), \end{aligned} \quad (57)$$

where in the third line I used $\lambda_j^{\tilde{\sigma}}(w_{next}) = \mathbb{E}[1_{\{\tilde{\sigma}(w,\epsilon)=j\}} | w_{next}]$, $j \in \mathcal{J}$, and in the fourth line I used the stationarity of the distribution. To sum up,

$$\beta \mathbb{E} R_4(V(w_{next}; \theta; \tilde{\sigma}; \rho), \rho_0) = \mathbb{E} V(w; \theta; \tilde{\sigma}; \rho) \lambda_1^{\tilde{\sigma}}(w). \quad (58)$$

The same argument can be applied to the term $\mathbb{E} R_3(V(w_{next}; \theta; \tilde{\sigma}; \rho))$:

$$\beta \mathbb{E} R_3(V(w_{next}; \theta; \tilde{\sigma}; \rho), \rho_0) = \mathbb{E} V(w; \theta; \tilde{\sigma}; \rho) \lambda_0^{\tilde{\sigma}}(w). \quad (59)$$

Combining (58) and (59) gives:

$$\mathbf{R}_{ins} := (1 - \beta) \partial_0 \mathbb{E} V(w; \theta; \tilde{\sigma}; r(\rho - \rho_0) + \rho_0) \quad (60)$$

for any Markov policy σ . According to the chain rule,

$$\mathbf{R}_{ins} + \mathbf{R}_{out} = 0.$$

Therefore,

$$\partial_0 \mathbb{E} V(w; \theta; \tilde{\sigma}; r(\rho - \rho_0) + \rho_0) = \frac{\beta}{1 - \beta} \mathbb{E}_e \left[\frac{dV(t; \theta; \tilde{\sigma}; \rho_0)}{dt} \Big|_{t=\rho_0(w)+e} \gamma_0^\sigma(w) \right],$$

and the bias correction term for the transition function is equal to:

$$\alpha_{\tilde{\sigma}}^{TRANS}(D; \theta; \xi) := \frac{\beta}{1 - \beta} \mathbb{E}_e \left[\frac{dV(t; \theta; \tilde{\sigma}; \rho_0)}{dt} \Big|_{t=\rho_0(w)+e} \right] \gamma_0^{\tilde{\sigma}}(w) (s_{next} - \rho(w)), \quad (61)$$

where ξ an unknown vector-valued function of the state variable w whose true value $\xi_0 = \xi_0(\theta)$ consists of the original nuisance parameter η_0 and

$$\xi_0(\theta) := \{\eta_0, \Pi_0(\cdot, \theta)\}, \quad (62)$$

where $\Pi_0(w, \theta) := \mathbb{E} \left[\frac{dV_{\tilde{\sigma}}(x; \theta; \eta_0)}{dx} \Big|_{x=w_{next}} | w \right]$ is the expectation of the derivative of the value function with respect to the state variable evaluated at the future state conditional on the current state w . To sum up, the bias correction term for Example 1 is equal to:

$$\alpha(D; \theta; \eta) := -\alpha^{CCP}(D; \theta; \gamma) + \alpha_{\sigma}^{TRANS}(D; \theta; \gamma) - \alpha_{\sigma^*}^{TRANS}(D; \theta; \gamma),$$

where $\alpha^{CCP}(D; \theta; \gamma)$ is given in (48) and $\alpha_{\sigma}^{TRANS}(D; \theta; \xi)$ is given in (61). ■

Proof of Remark 2. Step 1. Bias Correction Term for conditional choice probability of Apple. I seek for the bias correction term $\alpha^{CCP}(w; \gamma_A)$ such that takes the form:

$$\alpha^{CCP}(w; \theta; \gamma_A) = \Gamma_A(w; \theta) (1_{a_A=1} - \gamma_A(w))$$

and the function $\Gamma_A(w; \theta)$ satisfies the condition:

$$\partial_0 \mathbb{E}V(w; \theta; \sigma^*; r(\gamma_A(w) - \gamma_{A,0}(w)) + \gamma_{A,0}(w)) = \mathbb{E}\Gamma_A(w; \theta)[\gamma_A(w) - \gamma_{A,0}(w)].$$

I find the function $\Gamma_A(w; \theta)$ by the application of the implicit function theorem. I write $V(w; \theta; \sigma^*; \eta) = V(w; \theta; \sigma^*; \gamma_A)$ in order to focus on the nuisance parameter γ_A holding the other parameter γ_P at its true value. The recursive equation of the value function of Apple is:

$$\begin{aligned} \mathbb{E}V(w; \theta; \sigma^*; \gamma_A) &= \mathbb{E}[-R(1 - \gamma_A(w)) + (-\mu s)\gamma_A(w) + PS_{\sigma^*}(\gamma_A) \\ &\quad + \delta_0 \gamma_{P,0}(w) + (\delta_1 - \delta_0)\gamma_{P,0}(w)\gamma_A(w) \\ &\quad + \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma_A)|w, a = 0](1 - \gamma_A(w)) \\ &\quad + \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma_A)|w, a = 1](\gamma_A(w))]. \end{aligned}$$

The function γ_A appears in each summand of the equation above both inside and outside the value function $V(w; \theta; \sigma^*; \gamma_A)$. Let $\mathbf{S}_{out}, \mathbf{S}_{ins}$ be defined as in Example 1. Taking the derivative with respect to the outside presence of γ_A , I get:

$$\begin{aligned} \partial_0 \mathbf{S}_{out} &= -(-\mu s + R) + (\delta_1 - \delta_0)\gamma_{P,0}(w) + \frac{dPS_{\sigma^*}(\gamma)}{d\gamma}|_{\gamma_{A,0}} \\ &\quad + \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma_{A,0})|w, a = 1] - \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma_{A,0})|w, a = 0] \end{aligned}$$

where the second line follows from the definition of the choice-specific value functions and the third line follows from the logistic assumption of the shocks by (25). Taking the derivative with respect to the inside presence of γ_A , I get:

$$\partial_0 \mathbf{S}_{ins} = (1 - \beta)\partial_0 \mathbb{E}V(w; \theta; \sigma^*; r(\gamma_A - \gamma_{A,0}) + \gamma_{A,0}).$$

The function $\Gamma_A(w; \theta)$ follows:

$$\begin{aligned} \Gamma_A(w; \theta) &= (-\mu s + R) + (\delta_1 - \delta_0)\gamma_{P,0}(w) + \frac{dPS_{\sigma^*}(t)}{dt}|_{t=\gamma_{A,0}} \\ &\quad + \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma_A)|w, a = 1] - \beta \mathbb{E}[V(w_{next}; \theta; \sigma^*; \gamma_A)|w, a = 0] \end{aligned}$$

and

$$\alpha^{CCP}(D; \theta; \gamma_A) = \frac{1}{1 - \beta} \Gamma_A(w; \theta)(1_{a_A=1} - \gamma_A(w)).$$

I write $V(w; \theta; \sigma; \eta) = V(w; \theta; \sigma; \gamma_P)$ holding γ_A fixed at its true value $\gamma_{A,0}$. I find the bias correction term $\alpha_P^*(D; \theta; \gamma_P)$ such that:

$$\partial_{\gamma_P} \mathbb{E}[V(w; \theta; \sigma^*; \gamma_P) + \alpha_P^*(D; \theta; \gamma_P)] = 0.$$

The bias correction term takes the form:

$$\alpha_P^*(D; \theta; \gamma_P) := \Gamma_P(w; \theta)(1_{\{a_P=1\}} - \gamma_P(w)).$$

The function $\Gamma_P(w; \theta)$ is defined by the application of the implicit function theorem. Using the approach of Example 1, I find that

$$\partial_0 \mathbb{E}V(w; \theta; \sigma^*; r(\gamma_P - \gamma_{P,0}) + \gamma_{P,0})(1 - \beta) = \delta_0 + (\delta_1 - \delta_0)\gamma_{A,0}(w).$$

and

$$\Gamma_P(w; \theta) = \frac{1}{1 - \beta}(\delta_0 + (\delta_1 - \delta_0)\gamma_{A,0}(w)).$$

Similarly, the bias correction term $\alpha_P^\sigma(D; \theta; \gamma_P)$ takes the form:

$$\alpha_P^\sigma(D; \theta; \gamma_P) := \Gamma_P^\sigma(w; \theta)(1_{\{a_P=1\}} - \gamma_P(w)),$$

where

$$\Gamma_P^\sigma(w; \theta) = \frac{1}{1 - \beta}(\delta_0 + (\delta_1 - \delta_0)\gamma_0^\sigma(w)),$$

where $\gamma^\sigma(w) := \mathbb{P}(\sigma(w, \epsilon) = 1|w)$. To sum up, the bias correction term for the conditional probability of Player 2 entry is equal to:

$$\alpha_P^*(D; \theta; \gamma_P) := \frac{1}{1 - \beta}(\delta_1 - \delta_0)(\gamma_0^\sigma(w) - \gamma_{A,0}(w))(1_{\{a_P=1\}} - \gamma_P(w)). \quad (63)$$

■

Remark 3 (Bias correction term for point-identified problems). *Suppose the parameter of interest β_0 is the population average of the value function*

$$\beta_0 := \mathbb{E}[V(w; \theta_0; \sigma^*; \eta)].$$

Then the bias correction term for the conditional choice probability of player one is equal to:

$$\alpha^{CCP}(D; \theta; \gamma) := -\frac{1}{1 - \beta} \frac{2}{\gamma(w)} (1_{\{a_1=1\}} - \gamma(w)),$$

where the choice probability $\gamma_0(w)$ is defined in (25). Furthermore, if the parameter interest β_0 is equal to the weighted average of the value function

$$\beta_0 = \mathbb{E}q(w)V(w; \theta_0; \sigma^*; \eta), \quad (64)$$

the bias correction term is equal to:

$$\alpha_q^{CCP}(D; \theta; \gamma) = -\frac{q(w)}{q(w) - \beta\lambda(w)} \frac{2}{\gamma(w)} (1_{\{a_1=1\}} - \gamma(w)),$$

where $\lambda(w') := \mathbb{E}[q(w)|w_{next} = w']$ is the expectation of the weighting function $q(w)$ conditionally on the future state w_{next} .

6 Proofs from Section 3

Proof of Lemma 1. Let η be as in (27). Consider the value function $V(w; \theta; \sigma^*; \eta) = V(w; \theta; \sigma^*; \gamma_j)$ as a function of γ_j holding the other components of the nuisance parameter η fixed at the true value. The recursive equation for the value function takes the form:

$$\underbrace{\mathbb{E}q(w) V(w; \theta; \sigma^*; \gamma_j)}_{S_1(V(w; \sigma^*; \gamma_j))} - \underbrace{\sum_{j=2}^J \mathbb{E}_{\epsilon_{-1}} q(w) [\tilde{\pi}((j, \sigma_{-1}^*(w, \epsilon_{-1})); w; \theta_0)]}_{S_2(\gamma_j)} \quad (65)$$

$$- \underbrace{\tilde{\pi}((1, \sigma_{-1}^*(w, \epsilon_{-1})); w; \theta_0) | w] \gamma_j(w) + PS_{\sigma_1^*}(\gamma(w))}_{S_2(\gamma_j)} \quad (66)$$

$$+ \underbrace{\beta \mathbb{E}q(w) [V(w_{next}; \theta; \sigma^*; \gamma_j) | w, (1, \sigma_{-1}^*(w, \epsilon_{-1}))]}_{S_3(V(w; \sigma^*; \gamma_j))}$$

$$+ \beta \sum_{j=2}^J \mathbb{E}_{\epsilon_{-1}} [V(w_{next}; \theta; \sigma^*; \gamma_j) | w, (j, \sigma_{-1}^*(w, \epsilon_{-1}))] \quad (67)$$

$$- \mathbb{E}_{\epsilon_{-1}} [V(w_{next}; \theta; \sigma^*; \gamma_j) | w, (1, \sigma^*(w; \epsilon_{-1}))] \gamma_j(w) = 0.$$

The function $\gamma_j(w)$ appears in the equation both inside the value function $V(w; \theta; \sigma^*; \gamma_j)$ and outside of it. Taking the derivative with respect to the outside presence of γ_j gives:

$$\begin{aligned} \mathbf{S}_{out} &:= \mathbb{E}[\tilde{\pi}((j, \sigma^*(w; \epsilon_{-1})); w; \theta_0) - \tilde{\pi}((1, \sigma^*(w; \epsilon_{-1})); w; \theta_0) | w] + \partial_{\gamma_{j,0}} PS_{\sigma^*}(\gamma) \\ &+ \beta (\mathbb{E}_{\epsilon_{-1}} [V(w_{next}; \theta; \sigma^*; \gamma_j) | w, (j, \sigma^*(w; \epsilon_{-1}))] - \mathbb{E}_{\epsilon_{-1}} [V(w_{next}; \theta; \sigma^*; \gamma_j) | w, (1, \sigma^*(w; \epsilon_{-1}))]). \end{aligned}$$

Taking the derivative with respect to the inside presence gives:

$$\mathbf{S}_{ins} := (q(w) - \beta \lambda(w)) \partial_0 \mathbb{E}V(w; \theta; \sigma^*; r(\gamma_j - \gamma_{j,0}) + \gamma_{j,0})$$

using the same argument as in (50). Therefore, the bias correction term is

$$\alpha_j^{CCP}(D; \theta; \gamma_j) := -\frac{q(w)}{q(w) - \beta \lambda(w)} \Gamma_{\gamma,j}(w; \theta) (1_{\{a_1=j\}} - \gamma_j(w)),$$

where $\Gamma_{\gamma,j}(w, \theta)$ is given by

$$\begin{aligned} \Gamma_{\gamma,j}(w, \theta) &:= \mathbb{E}_{\epsilon_{-1}} [\tilde{\pi}((j, \sigma^*(w; \epsilon_{-1})); w; \theta_0) - \tilde{\pi}((1, \sigma^*(w; \epsilon_{-1})); w; \theta_0) | w] + \partial_{\gamma_{j,0}} PS_{\sigma^*}(\gamma) \\ &+ \beta (\mathbb{E}_{\epsilon_{-1}} [V(w_{next}; \theta; \sigma^*; \gamma_j) | w, (j, \sigma^*(w; \epsilon_{-1}))] \\ &- \mathbb{E}_{\epsilon_{-1}} [V(w_{next}; \theta; \sigma^*; \gamma_j) | w, (1, \sigma^*(w; \epsilon_{-1}))]). \end{aligned}$$

■

Proof of Lemma 2. Let a be a profile of actions. Let $Q(u, w, a)$ be the conditional quantile function of the conditional distribution $P(w_{next}|w, a)$ defined as

$$P(w_{next} \leq Q(u, w, a)|w, a) = u, \quad u \in [0, 1].$$

Thus I treat $\rho(u, w, a) := Q(u, w, a)$ as the nuisance parameter instead of $P(w_{next}|w, a)$. Let $\tilde{\sigma} \in \{\sigma, \sigma^*\}$ be a Markov policy. Consider the value function $V(w; \theta; \tilde{\sigma}; \eta) = V(w; \theta; \tilde{\sigma}; \rho)$ as a function of $\rho(u, w, a)$ holding the other components of the nuisance parameter η fixed at the true value. By the property of the conditional quantile function, the bias correction term $\alpha_{\tilde{\sigma}}^{TRANS}(D; \theta; \eta)$ is equal to:

$$\alpha_{\tilde{\sigma}}^{TRANS}(D; \theta; \eta) = \Pi(w; \theta) \cdot \frac{\mathbf{1}_{\{w_{next} \leq \rho(u, w, a)\}} - u}{f(w_{next}|w, a)},$$

where $f(w_{next}|w, a)$ is the conditional density of the future state w_{next} given the current state and the function $\Pi(w; \theta)$ is defined by the following equation:

$$\partial_0 \mathbb{E}q(w)V(w; \theta; \tilde{\sigma}; r(\rho - \rho_0) + \rho_0) = \mathbb{E}\Pi(w; \theta)[\rho - \rho_0].$$

I find the function $\Pi(w; \theta)$ by the application of the implicit function theorem to the recursive equation (5). The equation takes the form:

$$\begin{aligned} \mathbb{E}q(w) \left[V(w; \theta; \tilde{\sigma}; \rho) - \sum_{j=1}^J \mathbb{E}_{\epsilon_{-1}} q(w) [\tilde{\pi}((j, \sigma_{-1}^*(w, \epsilon_{-1})); w; \theta_0)] \prod_{k=1}^K P(\sigma_k(w, \epsilon_k) = a_k|w) \right. \\ \left. + PS_{\tilde{\sigma}_1}(w) + \beta \sum_{a' \in \prod_{k=1}^K \mathcal{A}_k} \mathbb{E}[V(w_{next}; \theta; \tilde{\sigma}; \eta_0)|w, a'] P(\tilde{\sigma}_k(w, \epsilon_k) = a'_k|w) \right] \\ = 0. \end{aligned}$$

The nuisance parameter ρ enters the equation above both inside and outside of the value function. The derivative with respect to the outside presence of ρ takes the form:

$$\begin{aligned} \beta \partial_0 \mathbb{E}q(w) [V(r(\rho(u_{next}, w, a) - \rho_0(u_{next}, w, a)) + \rho_0(u_{next}, w, a); \theta; \tilde{\sigma}; \eta_0)|w, a) \prod_{k=1}^K P(\tilde{\sigma}_k(w, \epsilon_k) = a_k|w) \\ = \beta \mathbb{E}q(w) \nabla_{w_{next}} V(w_{next}; \theta; \tilde{\sigma}; \eta_0) [\rho(u_{next}, w, a) - \rho_0(u_{next}, w, a)], \end{aligned}$$

where u_{next} is the random draw from the $U[0, 1]$. The derivative with respect to the inside presence is derived using the same argument of Lemma 1 that applies regardless of the equilibrium property of the strategy σ . The argument of Lemma 2 yields the bias correction term:

$$\alpha_{\tilde{\sigma}}^{TRANS}(D; \theta; \eta) = \frac{\beta q(w)}{q(w) - \beta \lambda(w)} \mathbb{E}[\nabla_{w_{next}} V(w_{next}; \theta; \tilde{\sigma}; \eta_0)|w, a] \cdot \frac{\mathbf{1}_{\{w_{next} \leq \rho(u, w, a)\}} - u}{f(w_{next}|w, a)}.$$

■

Proof of Lemma 3. Fix an action j_2 of player 2. Let $\gamma_{j_2} := \mathbb{P}(\sigma_2^*(w, \epsilon_2) = j_2 | w)$ be the conditional probability of the choice j_2 under the equilibrium strategy $\sigma_2^*(w, \epsilon_2)$ of player 2. I consider the value function $V(w; \theta; \tilde{\sigma}; \eta) = V(w; \theta; \tilde{\sigma}; \gamma_{j_2})$ holding the other nuisance parameters fixed at their true values. The bias correction term takes for $\gamma_{j,2}$ takes the form:

$$\alpha_{j_2}^{CCP,op}(D; \theta; \gamma_{j_2}) = \Gamma_{op}(w; \theta)(1_{\{a_2=j_2\}} - \gamma_{j_2}(w)),$$

where the function $\Gamma_{op}(w; \theta)$ is defined by the following equation:

$$\partial_0 \mathbb{E}q(w) V(w; \theta; \tilde{\sigma}; r(\gamma_{j_2} - \gamma_{j_2,0}) + \gamma_{j_2,0}) = \mathbb{E} \Gamma_{op}(w; \theta) [\gamma_{j_2}(w) - \gamma_{j_2,0}(w)].$$

■

I find the function $\Gamma_{op}(w; \theta)$ by the application of the implicit function theorem to the recursive equation (5). In the case $K = 2$ of two players (5) takes the form:

$$\begin{aligned} \mathbb{E}q(w) & \left[V(w; \theta; \tilde{\sigma}; \rho) - \sum_{j_1=1}^{A_1} \sum_{j_2=1}^{A_2} \tilde{\pi}((j_1, j_2), w, \theta) \mathbb{P}(\tilde{\sigma}(w, \epsilon_1) = j_1 | w) \gamma_{j_2}(w) \right. \\ & \left. + PS_{\tilde{\sigma}_1}(w) + \beta \sum_{j_1=1}^{A_1} \sum_{j_2=1}^{A_2} \mathbb{E}[V(w_{next}; \theta; \tilde{\sigma}; \eta_0) | w, (j_1, j_2)] \mathbb{P}(\tilde{\sigma}(w, \epsilon_1) = j_1 | w) \gamma_{j_2}(w) \right] \\ & = 0. \end{aligned}$$

The nuisance parameter $\gamma_{j_2}(w)$ enters the equation above both inside and outside of the utility function. The function $\gamma_{j_2}(w)$ enters linearly in its outside presence. Therefore, the derivative with respect to the outside presence takes the form:

$$\begin{aligned} \mathbb{E}q(w) & \left[\sum_{j_1=1}^{A_1} (\tilde{\pi}((j_1, j_2); w; \theta) - \tilde{\pi}((j_1, 1); w; \theta)) \mathbb{P}(\tilde{\sigma}(w, \epsilon_1) = j_1 | w) + \right. \\ & \left. \beta \sum_{j_1=1}^{A_1} (\mathbb{E}[V(w_{next}; \theta; \tilde{\sigma}) | w, (j_1, j_2)] - \mathbb{E}[V(w_{next}; \theta; \tilde{\sigma}) | w, (j_1, 1)]) \mathbb{P}(\tilde{\sigma}(w, \epsilon_1) = j_1 | w) \right] \\ & \cdot [\gamma_{j,2}(w) - \gamma_{j,2,0}(w)]. \end{aligned}$$

The derivative with respect to the inside presence is derived using the same argument of Lemma 1 that applies regardless of the equilibrium property of the strategy $\tilde{\sigma}$ and yields the bias correction

term:

$$\begin{aligned} \alpha_{j_2}^{CCP,op}(D; \theta; \gamma_{j_2}) &= \frac{q(w)}{q(w) - \beta\lambda(w)} \left[\sum_{j_1=1}^{A_1} (\tilde{\pi}((j_1, j_2); w; \theta) - \tilde{\pi}((j_1, 1); w; \theta)) \mathbb{P}(\sigma_1(w, \epsilon_1) = j_1 | w) + \right. \\ &\quad \left. \beta \sum_{j_1=1}^{A_1} (\mathbb{E}[V(w_{next}; \theta; \tilde{\sigma}) | w, (j_1, j_2)] - \mathbb{E}[V(w_{next}; \theta; \tilde{\sigma}) | w, (j_1, 1)]) \mathbb{P}(\sigma_1(w, \epsilon_1) = j_1 | w) \right] \cdot \\ &\quad \cdot (1_{\{a_2=j_2\}} - \gamma_{j_2}(w)). \end{aligned}$$

Proof of Lemma 4. Suppose there exists a function $B(a, w, \zeta)$ that is the basis function of the per-period utility function (22):

$$\tilde{\pi}(a; w; \theta; \zeta) = \theta \cdot B(a; w; \zeta).$$

Since ζ is an identified parameter of the distribution P_D , I drop it from the notation. The value function is a linear function of θ :

$$V(w; \theta; \sigma; \eta) = \underbrace{\mathbb{E}\left[\sum_{t \geq 0} \beta^t B(\sigma(w_t, \epsilon_t); w_t) | w\right]}_{=:\Psi_1(w, \sigma)} \cdot \theta + \underbrace{\mathbb{E}\left[\sum_{t \geq 0} \beta^t \epsilon_{1,t}(\sigma_1(w_t, \epsilon_t)) | w\right]}_{=:\Psi_2(w, \sigma)},$$

where $\Psi_1(w, \sigma)$ and $\Psi_2(w, \sigma)$ are equal to the expected discounted value of the deterministic per-period utility and the private shock, respectively. Therefore, $V(w; a; \sigma; \eta)$ is a linear function of θ . The profile-specific value function

$$\begin{aligned} \mathbb{E}[V(w_{next}; \theta; \sigma; \eta) | w, a] &= \mathbb{E}\left[\sum_{t \geq 1} \beta^t B(\sigma(w_t, \epsilon_t); w_t) | w_0 = w, a_0 = a\right] \cdot \theta \\ &\quad + \mathbb{E}\left[\sum_{t \geq 1} \beta^t \epsilon_{1,t}(\sigma_1(w_t, \epsilon_t)) | w_0 = w, a_0 = a\right] \end{aligned}$$

is also a linear function of θ . Therefore, the bias correction terms (33) and (35) are linear functions of θ . The expected derivative of the value function is also a linear function of θ :

$$\mathbb{E}[\nabla_{w_{next}} V(w_{next}; \theta; \sigma; \eta) | w_0 = w, a] = \mathbb{E}[\partial_{w_{next}} \Psi_1(w_{next}, \sigma) | w, a] \cdot \theta + \mathbb{E}[\partial_{w_{next}} \Psi_2(w_{next}, \sigma) | w, a].$$

Therefore, the bias correction term (34) is a linear function of θ . ■

7 Proofs from Section 4

7.1 Auxiliary Lemmas

Lemma 8 (Basic Inequalities). *The following inequalities hold for all vectors $x, y \in \mathcal{R}^d$:*

$$\begin{aligned} \|\|x\|^2 - \|y\|^2\| &\leq \|x - y\| \|x + y\| \\ \|\|x\|_+^2 - \|y\|_+^2\| &\leq \|x - y\|_+^2 + 2\|x - y\|_+ \|y\|_+. \end{aligned}$$

Proof. The first inequality holds:

$$\|\|x\|^2 - \|y\|^2\| = \left| \sum_{j=1}^d (x_j^2 - y_j^2) \right| \leq \sum_{j=1}^d |x_j - y_j| |x_j + y_j| \leq \|x - y\| \|x + y\|.$$

The second inequality holds:

$$\begin{aligned} \|\|x\|_+^2 - \|y\|_+^2\| &= \left| \sum_{j=1}^d ((x^+)_j^2 - (y^+)_j^2) \right| = \left| \sum_{j=1}^d (x_j^+ - y_j^+) (x_j^+ + y_j^+) \right| \\ &\leq \sum_{j=1}^d |(x_j^+ - y_j^+)| |x_j^+ + y_j^+| = \sum_{j=1}^d |(x_j^+ - y_j^+)|^2 + 2 \sum_{j=1}^d |(x_j^+ - y_j^+)| y_j^+ \\ &\leq \sum_{j=1}^d (|x_j - y_j|_+^2 + 2|x_j - y_j|_+ y_j^+) \leq \|x - y\|_+^2 + 2\|x - y\|_+ \|y\|_+. \end{aligned}$$

■

The assumptions below are the high-level assumptions on the population criterion function $Q(\theta, \xi)$. I will prove Theorem 5 under the assumptions below. I verify these assumptions from the conditions of Section 4 in the Subsection 7.3.

The proofs below are defined for the simple sample splitting procedure. The sample $(D_i)_{i=1}^N$ is decomposed into the auxiliary sample J_1 and the main sample J_2 of size $n := \lfloor N/2 \rfloor$ each. Let the nuisance parameter $\hat{\xi}$ be constructed on the sample J_1 . For each $\xi \in \Xi$ define the sample criterion function as

$$Q_n(\theta, \xi) := \left\| \frac{1}{n} \sum_{i \in J_2} g(D_i, \theta, \xi) \widehat{W}(\theta) \right\|_+^2. \quad (68)$$

where $\widehat{W}(\theta)$ is a diagonal weighting matrix that converges to a diagonal weighting matrix $W(\theta)$ with strictly positive diagonal elements.

ASSUMPTION 3 (Consistency of $Q_n(\theta, \xi_0)$). (1) *Regularity of $Q(\theta, \xi_0)$. The function $Q(\theta, \xi_0)$ is a non-negative lower semicontinuous function such that $\forall \epsilon > 0, \exists \delta(\epsilon) > 0$ such that*

$$\inf_{\Theta \setminus \Theta_I^\epsilon} Q(\theta, \xi_0) \geq \delta(\epsilon) > 0. \quad (69)$$

(2) *Fast Convergence on Θ_I . The sample loss $Q_n(\theta, \xi_0)$ converges to the population loss $Q(\theta, \xi_0)$ uniformly over the identified set Θ_I at the rate n :*

$$\sup_{\Theta_I} nQ_n(\theta, \xi_0) = O_P(1). \quad (70)$$

(3) *Slow Convergence on Θ . The sample loss $Q_n(\theta, \xi_0)$ converges to the population loss $Q(\theta, \xi_0)$ uniformly over the whole set Θ in the semi-metric $\|\cdot\|_+$ at the rate \sqrt{n} :*

$$\sup_{\Theta} (Q(\theta, \xi_0) - Q_n(\theta, \xi_0))_+ = O_P(1/\sqrt{n}). \quad (71)$$

(4) *There exist positive constants $\delta, \kappa > 0$ such that for any $p \in (0, 1)$ there exist constants d_p, n_p such that for any $n \geq n_p$:*

$$\inf_{\xi(\cdot) \in \Xi_n} P(Q_n(\theta, \xi) \geq \kappa[d(\theta, \Theta_I) \wedge \delta]^2 \quad \forall \theta : d_H(\theta, \Theta_I) \geq (d_p/n)^{1/2}) \geq 1 - p, \quad (72)$$

where $\Xi_n := \times_{\theta \in \Theta} \Xi_n^\theta$, and Ξ_n^θ is defined in Condition 1.

ASSUMPTION 4 (No effect of the first stage estimation error). (1) *Slow Convergence on Θ . For any $p > 0$ there exist constants r_p, n_p such that $\forall n \geq n_p$ the difference of the sample losses $Q_n(\theta, \xi) - Q_n(\theta, \xi_0)$ evaluated at the nuisance value $\xi(\theta)$ and at the true value $\xi_0(\theta)$ converges uniformly over Θ at rate \sqrt{n} for any element $\xi(\theta) \in \Xi_n^\theta$:*

$$\inf_{\xi(\cdot) \in \Xi_n} P(\sqrt{n} \sup_{\Theta} |Q_n(\theta, \xi) - Q_n(\theta, \xi_0)| \leq r_p) \geq 1 - p.$$

(2) *Fast Convergence on Θ_I . For any $p, \epsilon > 0$ there exists $n_{p,\epsilon}$ such that $\forall n \geq n_{p,\epsilon}$ the difference of the sample losses $Q_n(\theta, \xi) - Q_n(\theta, \xi_0)$ converges uniformly over Θ_I for any element $\xi(\cdot) \in \Xi_n$:*

$$\inf_{\xi(\cdot) \in \Xi_n} P(n \sup_{\Theta_I} |Q_n(\theta, \xi) - Q_n(\theta, \xi_0)| \leq \epsilon) \geq 1 - p.$$

ASSUMPTION 5 (Fast Convergence on the ϵ_n - Expansion of the Identified Set). *Let $\epsilon_n = O_P((\frac{1 \vee \hat{c}}{n})^{1/2})$ be the convergence rate in Theorem 9. For any $p, \epsilon > 0$ there exists $n_{p,\epsilon}$ such that $\forall n \geq n_{p,\epsilon}$ the difference between the sample losses $Q_n(\theta, \xi) - Q_n(\theta, \xi_0)$ converges uniformly over $\Theta_I^{\epsilon_n}$ for any element $\xi \in \Xi_n$:*

$$\inf_{\xi \in \Xi_n} P(n \sup_{\Theta_I^{\epsilon_n}} |Q_n(\theta, \xi) - Q_n(\theta, \xi_0)| \leq \epsilon) \geq 1 - p.$$

Lemma 9 (Coverage, Consistency, and Rate of Convergence for Loss Functions). *Let $\hat{c}/n \rightarrow_p 0$ and $\inf_{\xi \in \Xi_n} \mathbb{P}(\sup_{\theta \in \Theta_I} nQ_n(\theta, \xi) \leq \hat{c}) = 1 - o(1)$ (containment $\Theta_I \subseteq \hat{\Theta}_I$). Then, Assumptions 3[1-3] and 4[1] imply that $\Theta_I \subseteq \hat{\Theta}_I$ w.p. 1 and $d_H(\hat{\Theta}_I, \Theta_I) = o_P(1)$. Assumptions 3[1-4] and 4[1] imply that $d_H(\hat{\Theta}_I, \Theta_I) = O_P((1 \vee \hat{c})/n)^{1/2}$. In the case $\Theta_I = \Theta$, Assumption 3 implies that $d_H(\hat{\Theta}_I, \Theta_I) = 0$. In the case Assumption 4[2] holds, \hat{c} can be chosen as $\hat{c} = O_P(1)$.*

Lemma 10 (Concentration of Estimated Moments). *W.p. $\rightarrow 1$,*

$$\sup_{\theta \in \Theta} |\mathbb{G}_n[g(D_i, \theta, \hat{\xi}) - g(D_i, \theta, \xi_0)]| \lesssim_P r'_n \log(1/r'_n) + n^{-1/2+1/s} \log n. \quad (73)$$

Lemma 11 (Sufficient Conditions for General Moment Problems). *Conditions 1, 2, (41) imply that Assumptions 4 and 5 are satisfied for the moment inequalities problem with the population loss $Q(\theta, \xi_0) = \|\mathbb{E}g(W_i, \theta, \xi_0(\theta))\|_+$ and its sample analog $Q_n(\theta, \hat{\xi}) = \|\mathbb{E}_n g(D_i, \theta, \hat{\xi})\|_+$, as well as the moment equalities problem with the population loss $Q(\theta, \xi_0) = \|\mathbb{E}g(D_i, \theta, \xi_0(\theta))\|$ and its sample analog $Q_n(\theta, \hat{\xi}) = \|\mathbb{E}_n g(W_i, \theta, \hat{\xi})\|$.*

ASSUMPTION 6 (Limit Distribution of C_n). *There exists a law \mathcal{C} such that*

$$\mathbb{P}(C_n \leq c) \rightarrow \mathbb{P}(\mathcal{C} \leq c) \quad \forall c \in \mathcal{R},$$

where the distribution function of \mathcal{C} is continuous on $[0, \infty)$.

ASSUMPTION 7 (Approximability of C_n). *For any sequence of random measurable sets Θ_n such that $d_H(\Theta_n, \Theta_I) = o_P(n^{-1/2})$, the sequence of the suprema of the sample loss over sets Θ_n , $C'_n := \sup_{\theta \in \Theta_n} nQ_n(\theta, \xi_0)$ satisfies uniform convergence:*

$$\sup_{c \in \mathcal{R}} |\mathbb{P}(C'_n \leq c) - \mathbb{P}(\mathcal{C} \leq c)| = o(1).$$

7.2 Proof of Theorems from Main Text

Proof of Theorem 5 . . . Assumption 3[1] holds:

$$\inf_{\Theta \setminus \Theta_I^\epsilon} Q(\theta, \xi_0) = \|\mathbb{E}g(D_i, \theta, \xi_0)\|_+^2 \geq C_{\min}^2 (\epsilon \wedge \delta_{\min})^2 > 0. \quad ((38))$$

Assumption 3[2] holds by Lemma 8 and the P -Donsker property of $g(D_i, \theta, \xi_0(\theta))$:

$$\begin{aligned}
\sup_{\theta \in \Theta_I} n \|\mathbb{E}_n g(D_i, \theta, \xi_0)\|_+^2 &\leq \sup_{\theta \in \Theta_I} \|\mathbb{G}_n g(D_i, \theta, \xi_0)\|_+^2 \\
&\quad + 2 \|\mathbb{G}_n g(D_i, \theta, \xi_0)\|_+ \underbrace{\|\sqrt{n} \mathbb{E} g(D_i, \theta, \xi_0)\|_+}_{=0 \text{ by Definition of } \Theta_I} \\
&\quad + \underbrace{\|\sqrt{n} \mathbb{E} g(D_i, \theta, \xi_0)\|_+^2}_{=0 \text{ by Definition of } \Theta_I} \\
&\Rightarrow^d \|\Delta(\theta)\|_+^2 = O_P(1).
\end{aligned}$$

Assumption 3[3] holds by Lemma 8 and the P -Donsker property of $g(D_i, \theta, \xi_0)$:

$$\sup_{\theta \in \Theta} \sqrt{n} \|\mathbb{E} g(D_i, \theta, \xi_0)\|_+^2 - \|\mathbb{E}_n g(D_i, \theta, \xi_0)\|_+^2 = O_P(1).$$

Assumption 3[4] holds. Fix an amount of probability $p > 0$ and let n be large enough so that minimal eigenvalue \widehat{W} is bounded below w.p. $1 - p/4$:

$$\mathbb{P}(\mathcal{W}_n) := \mathbb{P}(\{\lambda_{\max} \geq \max \text{ eig } \widehat{W}(\theta) \geq \min \text{ eig } \widehat{W}(\theta) \geq \lambda_{\min}\}) \geq 1 - p/4.$$

By Condition 1(1)

$$\mathbb{P}(\mathcal{B}_n) = \mathbb{P}(\widehat{\xi}(\cdot) \in \Xi_n) \geq 1 - p/4,$$

by Lemma 10

$$\mathbb{P}(\mathcal{F}_n) := \mathbb{P}(\sup_{\theta \in \Theta} |\mathbb{G}_n g(D_i, \theta, \widehat{\xi}) - g(D_i, \theta, \xi_0)| \leq C_{p/4}) \geq 1 - p/4,$$

and by (41),

$$\mathbb{P}(\mathcal{G}_n) := \mathbb{P}(\|\mathbb{G}_n g(D_i, \theta, \xi_0)\| \leq C_{p/4}) \geq 1 - p/4.$$

On the event $\mathcal{B}_n \cap \mathcal{W}_n \cap \mathcal{F}_n \cap \mathcal{G}_n$ for n sufficiently large, $nQ_n(\theta, \widehat{\xi})$ is bounded below:

$$\begin{aligned}
nQ_n(\theta, \widehat{\xi}) &= \left\| \left(\mathbb{G}_n g(D_i, \theta, \widehat{\xi}) + \sqrt{n} \mathbb{E} g(D_i, \theta, \widehat{\xi}) \right)^\top \widehat{W}(\theta) \right\|_+^2 \\
&\geq \lambda_{\min} \|\mathbb{G}_n g(D_i, \theta, \widehat{\xi}) + \sqrt{n} \mathbb{E} g(D_i, \theta, \widehat{\xi})\|_+^2 \\
&\geq \lambda_{\min} \|\sqrt{n} \mathbb{E} g(D_i, \theta, \xi_0)\|_+^2 \left(\frac{\|\mathbb{G}_n g(D_i, \theta, \xi_0)\|}{\|\sqrt{n} \mathbb{E} g(D_i, \theta, \xi_0)\|_+^2} \right. \\
&\quad + \frac{\|\mathbb{G}_n [g(D_i, \theta, \widehat{\xi}) - g(D_i, \theta, \xi_0)] + \sqrt{n} \mathbb{E} [g(D_i, \theta, \widehat{\xi}) - g(D_i, \theta, \xi_0)]\|}{\|\sqrt{n} \mathbb{E} g(D_i, \theta, \xi_0)\|_+^2} \\
&\quad \left. + \frac{\|\sqrt{n} \mathbb{E} g(D_i, \theta, \xi_0)\|}{\|\sqrt{n} \mathbb{E} g(D_i, \theta, \xi_0)\|_+^2} \right).
\end{aligned}$$

Consider the set

$$d_H(\theta, \Theta_I) \geq 6C_{p/4}/C_{\min}\sqrt{n}$$

and n_p large enough so that

$$\forall n \geq n_p, \quad \sqrt{n}s_n < C_{p/4}, \quad r'_n \log(1/r'_n) < C_{p/4}.$$

On this set, the following bounds apply:

$$\begin{aligned} \|x\|_+ &:= \inf_{\theta: d_H(\theta, \Theta_I) \geq 6C_{p/4}/C_{\min}\sqrt{n}} \|\sqrt{n}\mathbb{E}g(D_i, \theta, \xi_0)\|_+ \geq C_{\min}(d_H(\theta, \Theta_I) \wedge \delta_{\min}) \geq 6C_{p/4}, \\ \|y\| &:= \|\mathbb{G}_n[g(D_i, \theta, \xi_0)] + \mathbb{G}_n[g(D_i, \theta, \hat{\xi}) - g(D_i, \theta, \xi_0)] + \sqrt{n}\mathbb{E}[g(D_i, \theta, \hat{\xi}) - g(D_i, \theta, \xi_0)]\| \\ &\leq C_{p/4} + 2C_{p/4} \leq 3C_{p/4}. \end{aligned}$$

Therefore, $\|x\|_+ \geq 2\|y\|$. Plugging in x and y into the inequality below:

$$\begin{aligned} \frac{\|x+y\|_+}{\|x\|_+} &\geq 1 - \frac{\|y\|}{\|x\|_+} \geq 1 - \frac{1}{2} = \frac{1}{2}, \\ \left(\frac{\|x+y\|_+}{\|x\|_+}\right)^2 &\geq \frac{1}{4} \end{aligned}$$

implies that $\frac{\|x+y\|_+}{\|x\|_+}$ is greater than or equal to $\frac{1}{2}$ on the set $\mathcal{B}_n \cap \mathcal{W}_n \cap \mathcal{F}_n \cap \mathcal{G}_n$ for n sufficiently large. Setting $\kappa = C_{\min}^2/16$, $\delta = \delta_{\min}$, $\gamma = 1/2$, $\kappa_p := (6C_{p/4}/C_{\min})^2$ and n large enough implies:

$$\begin{aligned} \inf_{\xi \in \Xi_n} \mathbb{P}(Q_n(\theta, \xi) \geq C_{\min}^2/16[d(\theta, \Theta_I) \wedge \delta_{\min}]^2) \\ \forall \theta : d_H(\theta, \Theta_I) \geq 6C_{p/4}/C_{\min}\sqrt{n} \geq 1 - p. \end{aligned}$$

Therefore, Assumption 3(4) holds. Lemma 11 verifies Assumption 4. ■

Proof of Lemma 6. Step 1. Let us show that the degeneracy property and the choice \hat{c}' given in (44) suffice for the rate $O_P(n^{-1/2})$. Let \hat{c} be as in (18) and $0 \leq \hat{c}' \leq \hat{c}$ w.p. $\rightarrow 1$. Conditionally on the event \mathcal{B}_n for any $\xi \in \Xi_n$ the following inclusion relation holds:

$$\inf_{\xi \in \Xi_n} \mathbb{P}(\Theta_n \subseteq \mathcal{C}_n(\hat{c}', \xi) \subseteq \mathcal{C}_n(\hat{c}, \xi)) \geq 1 - p.$$

By condition (2) of the degeneracy property, $d_H(\Theta_n, \Theta_I) = O_P(n^{-1/2})$. It has been shown in Theorem 5 that $d_H(\mathcal{C}_n(\hat{c}, \hat{\xi}) = O_P(n^{-1/2})$. Therefore, $d_H(\mathcal{C}_n(\hat{c}', \hat{\xi}), \Theta_I) = O_P(n^{-1/2})$ conditionally on \mathcal{B}_n . Since $\mathbb{P}(\mathcal{B}_n) \rightarrow 1$, the bound holds unconditionally.

Step 2. Suppose the conditions of Lemma 6 holds. Conditionally on \mathcal{B}_n for any $\xi \in \Xi_n$ the following bound holds uniformly on $\theta \in \Theta)I$:

$$nQ_n(\theta, \xi) \leq \lambda_{\max} \|\mathbb{G}_n g(W_i, \theta, \xi) + \sqrt{n} \mathbb{E}g(W_i, \theta, \xi)\|_+^2 \quad (74)$$

$$\leq \lambda_{\max} \sum_{l=1}^L |\mathbb{G}_n g_l(W_i, \theta, \xi) + \sqrt{n} \mathbb{E}g_l(W_i, \theta, \xi)|_+^2 \quad (75)$$

$$\leq \lambda_{\max} \sum_{l=1}^L |o(1) + O_P(1) + \sqrt{n} \mathbb{E}g_l(W_i, \theta, \xi_0)|_+^2 \quad (76)$$

$$\leq \lambda_{\max} \sum_{l=1}^L |o(1) + O_P(1) - \sqrt{n} C(d(\theta, \Theta \subset \Theta_I) \wedge \delta)|^2. \quad (77)$$

Conclude that $Q_n(\theta, \xi) = 0 \quad \forall \theta \in \Theta_I^{-\epsilon_n}$ with $\epsilon_n := 2 \sum_{l=1}^L L |\mathbb{E}g(W_i, \theta, \xi_0)| / \frac{1}{C}$ satisfies $\epsilon_n = O_P(1/\sqrt{n})$. Fix any $p > 0$. Since ϵ_n does not depend on ξ , there exists R_p and N_p such that for all $n \geq N_p$

$$P(\sqrt{n} d_H(\Theta_I^{-\epsilon_n}, \Theta_I) \leq R_p) \geq 1 - p.$$

■

Proof of Theorem 7. Assumption 4[2] implies that $\hat{c} = O_P(1)$ satisfies the conditions of Theorem 9. Let $\epsilon_n \asymp (\frac{\log^2 n}{n})^{1/2}$ be a numerical sequence. Fix a particular subsample j and its corresponding objective $Q_{j,b}$ for that subsample. Define

$$\bar{C}_{j,b,\xi} := \sup_{\theta \in \Theta_I^{\epsilon_n}} bQ_{j,b}(\theta, \xi)$$

and

$$\underline{C}_{j,b,\xi} := \inf_{K \in K_n} \sup_{\theta \in K} bQ_{j,b}(\theta, \xi),$$

where $K_n = \{\Theta_n : d_H(\Theta_n, \Theta) \leq \epsilon_n\}$ is closed set of all fixed(non-random) sets within Hausdorff distance ϵ_n of Θ_I . Since K_n is a closed set, there exists a set Θ_b^* where the infimum above is achieved:

$$\underline{C}_{j,b,\xi} := \inf_{K \in K_n} \sup_{\theta \in K} bQ_{j,b}(\theta, \xi) = \sup_{\theta \in \Theta_b^*} bQ_{j,b}(\theta, \xi).$$

By Lemma 9(c), for any $p > 0$ there exists n_p^A such that $\forall n \geq n_p^A \quad \hat{c} < \epsilon_n$ holds with probability at least $1 - p$:

$$P(\mathcal{A}_n) := P(\hat{c} < \epsilon_n) \geq 1 - p.$$

Therefore, on the event \mathcal{A}_n for any $\xi \in \Xi_n^\theta$ the following inequality holds:

$$\underline{C}_{j,b,\xi} \leq \hat{C}_{j,b} := \sup_{\theta \in \mathcal{C}_n(\hat{c}, \xi)} bQ_{j,b}(\theta, \xi) \leq \bar{C}_{j,b,\xi}.$$

By Assumption 1, for $p > 0$ there exists n_p^B such that $\forall n \geq n_p^B$

$$P(\mathcal{B}_n) := P(\hat{\xi} \in \Xi_n^\theta) \geq 1 - p.$$

Define the event

$$D_{\epsilon, \xi} := \left\{ \sup_{\theta \in \Theta_{\hat{\xi}}^n} b |Q_{j,b}(\theta, \xi) - Q_{j,b}(\theta, \xi_0)| \leq \epsilon \right\}$$

By Assumption 5, for $p, \epsilon > 0$ there exists $n_{p,\epsilon}^D$ such that $\forall n \geq n_{p,\epsilon}^D$ and any $\xi \in \Xi_n$ the event $D_{\epsilon, \xi}$ hold with probability at least $1 - 2p$ (i.e, $P(D_{p,\hat{\xi}} \cap \mathcal{B}_n) \geq 1 - 2p$).

On the event $\mathcal{A}_n \cap \mathcal{B}_n \cap D_{p,\hat{\xi}}$ for $n \geq n_p^A \vee n_p^B \vee n_{p,\epsilon}^D$,

$$\underline{C}_{j,b,\xi_0} - \epsilon \leq^i \underline{C}_{j,b,\hat{\xi}} \leq^{ii} \hat{C}_{j,b} := \sup_{\theta \in \mathcal{C}_n(\hat{c}, \hat{\xi})} b Q_{j,b}(\theta, \hat{\xi}) \leq^{iii} \bar{C}_{j,b,\hat{\xi}} \leq^{iv} \bar{C}_{j,b,\xi_0} + \epsilon,$$

where i and iv hold by definition of $D_{p,\hat{\xi}}$ and the property of supremum; ii and iii hold by definition of the event \mathcal{A}_n and \mathcal{B}_n . Denote

$$\hat{G}_{\hat{\xi}}(x) := \frac{1}{B_n} \sum_{j=1}^{B_n} 1_{\{\hat{C}_{j,b,\hat{\xi}} \leq x\}}, \quad \bar{G}(x) := \frac{1}{B_n} \sum_{j=1}^{B_n} 1_{\{\bar{C}_{j,b,\xi_0} \leq x\}}, \quad \underline{G}(x) := \frac{1}{B_n} \sum_{j=1}^{B_n} 1_{\{\bar{C}_{j,b,\xi_0} \leq x\}}.$$

Since

$$\hat{C}_{j,b,\hat{\xi}} \leq \bar{C}_{j,b,\xi_0} + \epsilon$$

the event $\bar{C}_{j,b,\xi_0} \leq x - \epsilon$ implies $\hat{C}_{j,b} \leq x$. Therefore, $\frac{1}{B_n} \sum_{j=1}^{B_n} 1_{\{\hat{C}_{j,b,\hat{\xi}} \leq x\}} \geq \frac{1}{B_n} \sum_{j=1}^{B_n} 1_{\{\bar{C}_{j,b,\xi_0} \leq x - \epsilon\}}$, and using the notation above, I get:

$$\underline{G}(x - \epsilon) \leq \hat{G}_{\hat{\xi}}(x).$$

Similar argument gives:

$$\underline{G}(x - \epsilon) \leq \hat{G}_{\hat{\xi}}(x) \leq \bar{G}(x + \epsilon).$$

Step 2 shows that continuity of the c.d.f of $P(\mathcal{C} \leq x)$ (Assumption 6) implies that

$$\underline{G}(x - \epsilon) \rightarrow_p P(\mathcal{C} \leq x - \epsilon) \rightarrow P(\mathcal{C} \leq x), b \rightarrow \infty$$

and

$$\bar{G}(x + \epsilon) \rightarrow_p P(\mathcal{C} \leq x + \epsilon) \rightarrow P(\mathcal{C} \leq x), b \rightarrow \infty.$$

Step 2. Consider the function

$$h(W_1, W_2, \dots, W_n) := \frac{1}{B_n} \sum_{j=1}^{B_n} 1_{\{\bar{C}_{j,b,\xi_0} \leq x - \epsilon\}}.$$

Let us show that $h(\cdot)$ has bounded differences: replacement of observation W_k by W'_k changes at most one subsample j and results in at most a $\frac{1}{B_n}$ change in the function itself.

$$|h(W_1, W_k, \dots, W_n) - h(W_1, W'_k, \dots, W_n)| \leq \frac{1}{B_n}.$$

McDiarmid's inequality implies:

$$\sup_{\epsilon \in \mathcal{R}} \mathbb{P}(|\underline{\mathbf{G}}(x - \epsilon) - \mathbb{E}\underline{\mathbf{G}}(x - \epsilon)| > t) \leq 2 \exp^{-2t^2 B_n^2/n} := \delta$$

$$\underline{\mathbf{G}}(x - \epsilon) = \mathbb{E}\underline{\mathbf{G}}(x - \epsilon) + O_P\left(\left(\log \frac{2}{\delta}\right)^{1/2} \frac{\sqrt{n}}{B_n}\right)$$

uniformly over $\epsilon > 0$. Since the subsamples $j_1 \neq j_2$ are i.i.d,

$$\mathbb{E}\underline{\mathbf{G}}(x - \epsilon) = \mathbb{P}(\bar{C}_{j,b,\xi_0} \leq x - \epsilon).$$

Step 3. By Assumption 7,

$$\begin{aligned} |\mathbb{P}(\bar{C}_{j,b,\xi_0} \leq x - \epsilon) - \mathbb{P}(\mathcal{C} \leq x)| &\leq \sup_{x \in \mathcal{R}} |\mathbb{P}(\bar{C}_{j,b,\xi_0} \leq x) - \mathbb{P}(\mathcal{C} \leq x)| && \text{(Assumption 7)} \\ &+ |\mathbb{P}(\mathcal{C} \leq x - \epsilon) - \mathbb{P}(\mathcal{C} \leq x)| && \text{(Continuity of } C_{\nu_0}) \\ &= o(1) && \text{(78)} \end{aligned}$$

as $n \rightarrow \infty, b \rightarrow \infty$. Since the bound above holds for any $\epsilon > 0$, the statement is proved. To conclude, I have shown that $\bar{G}(x+\epsilon) = P(\mathcal{C} \leq c) + o_P(1)$. Similarly, I can show that $\underline{G}(x-\epsilon) = P(\mathcal{C} \leq c) + o_P(1)$. Since $P(\mathcal{C} \leq c)$ is continuous in c , the α - quantile of $\hat{G}_{\hat{\xi}}(x)$ converges to the α -quantile of \mathcal{C} . ■

7.3 Proof of Auxiliary Lemmas

Proof of Lemma 9. The proof relies on the following basic inequalities. For any set Θ and two functions $p(\theta), q(\theta)$ the following holds:

$$\sup_{\theta \in \Theta} p(\theta) \leq \sup_{\theta \in \Theta} q(\theta) + \sup_{\theta \in \Theta} (p(\theta) - q(\theta))_+ \quad (79)$$

$$\inf_{\theta \in \Theta} p(\theta) \geq \inf_{\theta \in \Theta} q(\theta) - \sup_{\theta \in \Theta} (q(\theta) - p(\theta))_+. \quad (80)$$

Consider a sequence of events $\mathcal{B}_n := \{\hat{\xi}(\cdot) \in \Xi_n\}$ whose probability approaches one:

$$\mathbb{P}(\mathcal{B}_n) = 1 - o(1).$$

Step 1. On an event \mathcal{B}_n

$$\mathbb{P}\left(\sup_{\theta \in \Theta_I} nQ_n(\theta, \hat{\xi}) \leq \hat{c} | \mathcal{B}_n\right) \geq \inf_{\xi \in \Xi_n} \mathbb{P}\left(\sup_{\theta \in \Theta_I} nQ_n(\theta, \xi) \leq \hat{c}\right) = 1 - o(1),$$

which implies

$$P(\Theta_I \subset \hat{\Theta}_I | \mathcal{B}_n) = 1 - o(1).$$

Since $P(\mathcal{B}_n) = 1 - o(1)$, with probability approaching one $P(\Theta_I \subset \hat{\Theta}_I)$ holds.

Step 2. Proof of convergence without guaranteed rate. Fix an $\epsilon > 0$. Let us show that

$$P(d_H(\hat{\Theta}_I, \Theta_I) \leq \epsilon) \rightarrow 1.$$

By Assumption 3 (Equation (69)), there exists $\delta(\epsilon)$ such that

$$\inf_{\Theta \setminus \Theta_I^\epsilon} Q(\theta, \xi_0) \geq \delta(\epsilon) > 0.$$

To see that $P(\hat{\Theta}_I \subset \Theta_I^\epsilon) = 1 - o(1)$, recognize that on the event \mathcal{B}_n conditionally on the subsample J_1 :

$$\sup_{\hat{\Theta}_I} Q(\theta, \xi_0) \leq \sup_{\hat{\Theta}_I} Q_n(\theta, \xi_0) + \sup_{\hat{\Theta}_I} (Q(\theta, \xi_0) - Q_n(\theta, \xi_0))_+ \quad (\text{Eq. (79)})$$

$$\leq \sup_{\hat{\Theta}_I} Q_n(\theta, \xi_0) + O_P(1/\sqrt{n}) \quad (\text{Eq. (71)})$$

$$\leq \sup_{\hat{\Theta}_I} Q_n(\theta, \hat{\xi}(\theta)) + \sup_{\hat{\Theta}_I} (Q_n(\theta, \xi_0) - Q_n(\theta, \hat{\xi}))_+ + O_P(1/\sqrt{n}) \quad (\text{Eq. (79)})$$

$$\leq \hat{c}/n + O_P(1/\sqrt{n} + 1/\sqrt{n}) = o_P(1) \quad (\text{Assumption 4 and the choice of } \hat{c})$$

Since $\sup_{\hat{\Theta}_I} Q(\theta, \xi_0) < \frac{\epsilon}{2} \Rightarrow \hat{\Theta}_I \subset \Theta_I^\epsilon$, $P(\hat{\Theta}_I \subset \Theta_I^\epsilon) = 1 - o(1)$.

Step 3. Proof of convergence at rate $\epsilon_n := (\frac{d_{p/3} \kappa \vee \hat{c}}{n \kappa})^{1/2}$. Fix a probability level $p \in (0, 1)$ and let the constants κ, δ be as specified in Assumption 3 (Equation (72)). I have to show:

$$\forall p > 0 \quad \exists d_p, n_p : \quad \forall n \geq n_p \quad P(d_H(\hat{\Theta}_I, \Theta_I) \leq \epsilon_n) \geq 1 - p.$$

Since $\hat{c}/n \rightarrow_p 0$, there exists $n_{p/3}^A$ such that for a sufficiently large $n : n \geq n_{p/3}^A$,

$$P(\mathcal{A}_n) := P(\hat{c}/(n\kappa) < (\delta/2)^2) = 1 - p/3.$$

By the definition of $\hat{\xi}(\theta)$, for a sufficiently large $n : n \geq n_{p/3}^B$,

$$P(\mathcal{B}_n) := P(\hat{\xi}(\cdot) \in \Xi_n) \geq 1 - p/3.$$

Define a set

$$\mathcal{D}_{\hat{\xi}} := \left\{ \inf_{\theta: d_H(\theta, \Theta_I) \geq \epsilon_n} nQ_n(\theta, \hat{\xi}) \geq n\kappa(\epsilon_n \wedge \delta)^2 \right\}$$

Since $\epsilon_n \geq (d_{p/3}/n)^{1/2}$ holds absolutely surely by the choice of ϵ_n , Assumption 3[4] implies that for $n \geq n_{p/3}$:

$$\mathbb{P}(\mathcal{D}_{\hat{\xi}}|\mathcal{B}_n) \geq \inf_{\xi \in \Xi_n} \mathbb{P}\left(\inf_{\theta: d_H(\theta, \Theta_I) \geq \epsilon_n} nQ_n(\theta, \xi) \geq n\kappa(\epsilon_n \wedge \delta)^2\right) \geq 1 - p/3.$$

Therefore, for a sufficiently large $n : n \geq n_{p/3}^A$,

$$\begin{aligned} \inf_{\theta: d_H(\theta, \Theta_I) \geq \epsilon_n} nQ_n(\theta, \hat{\xi}) \geq n\kappa(\epsilon_n \wedge \delta)^2 &\Rightarrow \\ \inf_{\theta: d_H(\theta, \Theta_I) \geq \epsilon_n} nQ_n(\theta, \hat{\xi}) \geq n\kappa\epsilon_n^2 = d_p\kappa \vee \hat{c} \geq \hat{c}. \end{aligned} \quad (81)$$

Therefore, for a sufficiently large $n : n \geq n_{p/3} \vee n_{p/3}^A \vee n_{p/3}^B$,

$$\begin{aligned} \mathbb{P}\left(\inf_{\theta: d_H(\theta, \Theta_I) \geq \epsilon_n} nQ_n(\theta, \hat{\xi}) \geq \kappa(\epsilon_n \wedge \delta)^2 | \mathcal{B}_n\right) \\ \geq \inf_{\xi \in \Xi_n} \mathbb{P}\left(\inf_{\theta: d_H(\theta, \Theta_I) \geq \epsilon_n} nQ_n(\theta, \xi) \geq \kappa(\epsilon_n \wedge \delta)^2\right). \end{aligned} \quad (82)$$

Combining (82) and (81) gives

$$\begin{aligned} \mathbb{P}\left(\inf_{\theta: d_H(\theta, \Theta_I) \geq \epsilon_n} nQ_n(\theta, \hat{\xi}) \geq \kappa(\epsilon_n \wedge \delta)^2 \cap \mathcal{B}_n \cap \mathcal{A}_n\right) \\ \geq 1 - \mathbb{P}((\mathcal{B}_n)^c) + \mathbb{P}((\mathcal{A}_n)^c) + \mathbb{P}((\mathcal{D}_{\hat{\xi}})^c) \\ \geq 1 - (p/3 + p/3 + p/3). \end{aligned}$$

Since $\mathbb{P}(\sup_{\hat{\Theta}_I} nQ_n(\theta, \hat{\xi}) \leq \hat{c}) = 1 - o(1)$, this implies that $\mathbb{P}(\hat{\Theta}_I \subset \Theta_I^{\epsilon_n}) = 1 - o(1)$. By Step 1, I conclude that $d_H(\hat{\Theta}_I, \Theta_I) = O_P(\epsilon_n)$.

Step 4. That $\hat{c} = O_P(1)$ satisfies the conditions of Lemma 9, follows from

$$\begin{aligned} \sup_{\theta \in \Theta_I} nQ_n(\theta, \xi) &\leq \sup_{\theta \in \Theta_I} nQ_n(\theta, \xi_0) + \sup_{\theta \in \Theta_I} n|Q_n(\theta, \xi) - Q_n(\theta, \xi_0)| \\ &\leq O_P(1) + \sup_{\theta \in \Theta_I} n|Q_n(\theta, \xi) - Q_n(\theta, \xi_0)| \quad (\text{Assumption 3[2]}) \\ &\leq O_P(1) + o_P(1) \quad (\text{Assumption 4[2]}) \end{aligned}$$

for any $\xi \in \Xi_n$. ■

Proof of Lemma 10. Conditionally on the auxiliary sample, $\hat{\xi}(\cdot)$ can be treated as fixed, and w.p. approaching one, $\hat{\xi}(\cdot) \in \Xi_n$. Consider the function class:

$$\mathcal{F}_2 = \{g_l(D_i, \theta, \hat{\xi}) - g_l(D_i, \theta, \xi_0), l = 1, 2, \dots, L, \theta \in \Theta\} \subset \mathcal{F}_{\hat{\xi}} - \mathcal{F}_{\xi_0}.$$

Let a function $F_2 := F_{\hat{\xi}} + F_{\xi_0}$ be the envelope function for the class \mathcal{F}_2 . This function satisfies the envelope requirements of Lemma 6.2 of Chernozhukov et al. (2017a) since $\|F_2\|_{P,q} \leq \|F_{\hat{\xi}}\|_{P,q} + \|F_{\xi_0}\|_{P,q} \leq 2C_1$ and

$$\log \sup_{\tilde{Q}} N(\epsilon \|F_2\|_{Q,2}, \mathcal{F}_2, \|\cdot\|_{\tilde{Q},2}) \leq 2v \log(2a/\epsilon).$$

On the event \mathcal{B}_n ,

$$\begin{aligned} \sup_{\theta \in \Theta} \mathbb{E}[g_j(D_i, \theta, \hat{\xi}) - g_j(D_i, \theta, \xi_0)]^2 &\leq \sup_{\theta \in \Theta, \xi \in \Xi_n} \mathbb{E}[\|g_j(D_i, \theta, \xi(\theta)) - g_j(D_i, \theta, \xi_0)\|]^2 \\ &\leq (r'_n)^2 \end{aligned}$$

The application of Lemma 6.2 of Chernozhukov et al. (2017a) conditionally on the auxiliary sample, with the function class \mathcal{F}_2 , envelope F_2 , and $\sigma^2 := \sqrt{n}^2$ yields:

$$\begin{aligned} &\sup_{f \in \mathcal{F}} |\mathbb{G}_n[g(D_i, \theta, \hat{\xi}) - g(D_i, \theta, \xi_0)]| \\ &\lesssim_P r'_n \log^{1/2}(1/r'_n) + n^{-1/2+1/s} \log n = o_P(1) \end{aligned} \tag{83}$$

■

Proof of Lemma 11. Let the events $\mathcal{W}_n, \mathcal{B}_n, \mathcal{F}_n, \mathcal{G}_n$ be as defined at in the Proof of Theorem 5. The application of Lemma 8 on the event $\mathcal{W}_n \cap \mathcal{B}_n \cap \mathcal{F}_n \cap \mathcal{G}_n$ gives:

$$\begin{aligned} |Q_n(\theta, \xi) - Q_n(\theta, \xi_0)| &= \left| \underbrace{\|\mathbb{E}_n g(D_i, \theta, \hat{\xi}) \widehat{W}(\theta)\|_+^2}_x - \underbrace{\|\mathbb{E}_n g(D_i, \theta, \xi_0) \widehat{W}(\theta)\|_+^2}_y \right| \\ &\leq \|x - y\|_+^2 + 2\|x - y\|_+ \|y\|_+. \end{aligned}$$

The terms $\|x - y\|_+$ and $\|y\|_+$ admit the following bound:

$$\begin{aligned} \|x - y\|_+ &= \|\mathbb{E}_n[g(D_i, \theta, \xi) - g(D_i, \theta, \xi_0)] \widehat{W}(\theta)\| \\ &\leq \lambda_{\max} \|\mathbb{E}[g(D_i, \theta, \xi) - g(D_i, \theta, \xi_0)] + \mathbb{G}_n[g(D_i, \theta, \xi) - g(D_i, \theta, \xi_0)]/\sqrt{n}\| \\ &\leq \lambda_{\max}(s_n + (r'_n \log(1/r'_n) + n^{-1/2+1/s})), \\ \|y\|_+ &= \|\mathbb{G}_n g(D_i, \theta, \xi_0)/\sqrt{n} + \mathbb{E}g(D_i, \xi_0)\|_+. \end{aligned}$$

Therefore, Assumption 4(a) holds, and

$$r_p = (2\lambda_{\max} C_{p/4})^2 + 2\lambda_{\max} C_{p/4} 2(C_{p/4} + \sup_{\theta \in \Theta} \|\mathbb{E}g(D_i, \theta, \xi_0)\|)$$

satisfies

$$\inf_{\xi \in \Xi_n} \mathbb{P}(\sqrt{n} \sup_{\Theta} |Q_n(\theta, \xi) - Q_n(\theta, \xi_0)| \leq r_p) \geq 1 - p.$$

and for n large enough $\sup_{\theta \in \Theta_I} \|C_{p/4} + \sqrt{n} \mathbb{E}g(D_i, \theta, \xi_0)\|_+ = 0$, $r'_p = (2\lambda_{\max} C_{p/4})^2$

$$\inf_{\xi \in \Xi_n} \mathbb{P}(n \sup_{\Theta_I} |Q_n(\theta, \xi) - Q_n(\theta, \xi_0)| \leq r'_p) \geq 1 - p.$$

If Assumption 3(2) holds,

$$\begin{aligned} \sup_{\Theta_I^\epsilon} n|Q_n(\theta, \xi) - Q_n(\theta, \xi_0)| &\leq n\|x - y\|_+^2 + 2n\|x - y\|_+\|y\|_+ \\ \sqrt{n}\|x - y\|_+ &= \|\sqrt{n}\mathbb{E}_n[g(D_i, \theta, \xi) - g(D_i, \theta, \xi_0)]\widehat{W}(\theta)\| \\ &\leq \lambda_{\max}\|\sqrt{n}\mathbb{E}[g(D_i, \theta, \xi) - g(D_i, \theta, \xi_0)] + \mathbb{G}_n[g(D_i, \theta, \xi) - g(D_i, \theta, \xi_0)]\| \\ &\leq \lambda_{\max}(\sqrt{n}s_n + r'_n \log(1/r'_n) + n^{-1/2+1/s}) = o(1). \\ \sqrt{n}\|y\|_+ &= \|\mathbb{G}_n g(D_i, \theta, \xi_0) + \sqrt{n}\mathbb{E}g(D_i, \xi_0)\|_+ \\ \sup_{\theta \in \Theta_I^{\epsilon_n}} \sqrt{n}\|y\|_+ &\leq C_{p/4} + \sqrt{n}\epsilon_n \leq C_{p/4} + O_P(1). \end{aligned}$$

Therefore, for any $\epsilon, p > 0$ for n sufficiently large

$$\inf_{\xi \in \Xi_n} \mathbb{P}(\sup_{\Theta_I^{\epsilon_n}} n|Q_n(\theta, \xi) - Q_n(\theta, \xi_0)| \leq \epsilon) \geq 1 - p.$$

■

References

- Andrews, D. (1994). Asymptotics for semiparametric econometric models via stochastic equicontinuity. *Econometrica*, 1(62):43–72.
- Arcidiacono, P., Bayer, P., Bugni, F., and James, J. (2013). Approximating high-dimensional dynamic models: Sieve value function iteration. *Advances in Econometrics*, 31:45 – 95.
- Athey, S. (2017). Beyond prediction: Using big data for policy problems. *Science*, 355(6324):483–485.
- Bajari, P., Benkard, C. L., and Levin, J. (2007). Estimating dynamic models of imperfect competition. *Econometrica*, 75(5):1331–1370.

- Bajari, P., Chernozhukov, V., Hong, H., and Nekipelov, D. (2015). Identification and efficient semiparametric estimation of dynamic models.
- Bajari, P., Hong, H., Krainer, J., and Nekipelov, D. (2010). Estimating static models of strategic interaction. *Journal of Business and Economic Statistics*, 28(4):469–482.
- Belloni, A., Chernozhukov, V., and Wei, Y. (2016). Post-selection inference for generalized linear models with many controls. *Journal of Business & Economic Statistics*, 34(4):606–619.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., Newey, W., and Robins, J. (2017a). Double/debiased machine learning for treatment and causal parameters.
- Chernozhukov, V., Escanciano, J. C., Ichimura, H., and Newey, W. (2017b). Locally robust semi-parametric estimation.
- Chernozhukov, V., Goldman, M., Semenova, V., and Taddy, M. (2017c). Orthogonal machine learning for demand estimation: High dimensional causal inference in dynamic panels. <https://arxiv.org/abs/1712.09988>.
- Chernozhukov, V., Hong, H., and Tamer, E. (2007). Estimation and confidence regions for parameter sets in econometric models. *Econometrica*, 75(5):1243–1284.
- Chernozhukov, V., Nekipelov, D., Semenova, V., and Syrgkanis, V. (2018). Plug-in regularized estimation of high-dimensional sparse parameters in nonlinear econometric models.
- Ciliberto, F. and Tamer, E. (2009). Market structure and multiple equilibria in airline markets. *Econometrica*, 77(6).
- Hotz, V. J. and Miller, R. A. (1993). Conditional choice probabilities and the estimation of dynamic models. *Review of Economic Studies*, 60(3):497–529.
- Ibbragimov, I. R. and Hasminskii, R. Z. (1981). *Statistical Estimation: Asymptotic Theory*.
- Kaido, H. and White, H. (2014). A two-stage procedure for partially identified models. *Journal of Econometrics*, 1(182):5–13.
- Lanjouw, J. O., Pakes, A., and Putnam, J. (1998). How to count patents and value of intellectual property: The uses of patent renewal and application data. *Journal of Industrial Economics*, 46(4):405–432.

- Newey, W. (1994). The asymptotic variance of semiparametric estimators. 62:245–271.
- Neyman, J. (1959). Optimal asymptotic tests of composite statistical hypotheses. *Probability and Statistics*, 213(57).
- Pakes, A. (1986). Patents as options: Some estimates of the value of holding european patent stocks. *Econometrica*, 54(4):755–784.
- Pakes, A. and Pollard, D. (1989). Simulation and the asymptotics of optimization estimators. *Econometrica*, 5(57):1027–1057.
- Pakes, A. and Schankerman, M. (1984). The rate of obsolescence of patents, research gestation lags, and the private rate of return to research resources. *RD, Patents and Productivity, NBER Conference Series. Chicago, The University of Chicago Press.*
- Robins, J. and Rotnitzky, A. (1995). Semiparametric efficiency in multivariate regression models with missing data. *Journal of American Statistical Association*, 90(429):122–129.
- Robins, J., Rotnitzky, A., and Zhao, L. P. (1994). Estimation of regression coefficients when some regressors are not always observed. *Journal of American Statistical Association*, 89(427):846–866.
- Romano, J. and Shaikh, A. (2010). Inference for the identified set in partially identified econometric models. *Econometrica*, 1(78):169–211.
- Rosen, A. (2006). Confidence sets for partially identified econometric models.
- Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica*, 55(5):999–1033.
- Ryan, S. (2012). The costs of environmental regulation in a concentrated industry. *Econometrica*, 80(3):10191061.
- Scott, P. (2013). Dynamic discrete choice estimation of agricultural land use.
- van der Vaart, A. (1998). Asymptotic statistics.
- Wager, S. and Athey, S. (2016). Estimation and inference of heterogeneous treatment effects using random forests. <https://arxiv.org/abs/1510.04342>.