

IDENTIFICATION AND ESTIMATION OF DYNAMIC GAMES WITH CONTINUOUS STATES AND CONTROLS

PAUL SCHRIMPF

ABSTRACT. This paper analyzes dynamic games with continuous states and controls. There are two main contributions. First, we give conditions under which the payoff function is nonparametrically identified by the observed distribution of states and controls. The identification conditions are fairly general and can be expected to hold in many potential applications. The key identifying restrictions include that one of the partial derivatives of the payoff function is known and that there is some component of the state space that enters the policy function, but not the payoff function directly. The latter of these restrictions is a standard exclusion restriction and is used to identify the payoff function off the equilibrium path. By manipulating the first order condition, we can show that the payoff function satisfies an integro-differential equation. Due to the presence of the value function in the first order condition, this integro-differential equation contains a Fredholm integral operator of the second kind. Invertibility of this operator, and knowledge of one of the partial derivatives of the payoff function is used to ensure that the integro-differential equation has a unique solution.

The second contribution of this paper is to propose a two-step semiparametric estimator for the model. In the first step the transition densities and policy function are estimated nonparametrically. In the second step, the parameters of the payoff function are estimated from the optimality conditions of the model. Because the state and action space are continuous, there is a continuum of optimality conditions. The parameter estimates minimize the norm of these conditions. Hence, the estimator is related to recent papers on GMM in Hilbert spaces and semiparametric estimators with conditional moment restrictions. We give high-level conditions on the first step nonparametric estimates for the parameter estimates to be consistent and parameters to be \sqrt{n} -asymptotically normal. Finally, we show that a kernel based estimator satisfies these conditions.

Date: December 1, 2011, first posted December 1, 2010.

This paper was originally written while I was a graduate student at MIT. I am very grateful for the guidance of my advisors, Victor Chernozhukov and Whitney Newey. I thank Arun Chandrasekhar and Denis Chetverikov for useful comments. I also thank Brad Larsen, Anna Mikusheva, and other participants at the MIT econometrics lunch for their feedback.

1. INTRODUCTION

This paper analyzes dynamic games with continuous states and controls. There have been many recent papers about estimating dynamic games, but few, if any, of them allow for both continuous states and continuous controls. This is a useful gap to fill since many state and action variables in real applications are approximately continuous. A generic and pervasive example is investment as a decision variable and capital stock as a state variable. This paper gives a comprehensive econometric analysis of continuous dynamic games. There are two main results. First, we give sufficient conditions for nonparametric identification of the payoff function from the observed transition densities. Second, we propose a two-step semiparametric estimator for the payoff function and give conditions for consistency and asymptotic normality.

The first main contribution of this paper is to give sufficient conditions for the payoff function to be identified by the transition density of the state variables and controls. It is well known (see e.g. Rust (1994)), that without some restrictions, the payoff function is unidentified. However, in the case of dynamic games with discrete controls, Magnac and Thesmar (2002) and Bajari, Chernozhukov, Hong, and Nekipelov (2009) give plausible restrictions sufficient for identification. This paper develops an analogous result for dynamic games with continuous controls. The key conditions for identification of the payoff function are that one of the partial derivatives of the payoff function is known, there is some value of the control such that value of the payoff function is known for that value of the control and all possible values of the state, and a certain integral operator of the second kind is invertible. The first two of these three conditions can be expected to hold in potential applications. We illustrate their plausibility with an example. The last of these three conditions is difficult to verify analytically, but could be checked in applications.

Our proof of identification involves manipulating the first order condition for the continuous control. The first order condition is an integro-differential equation—it involves both a derivative and an integral of the payoff function. After some manipulation, we transform the first order condition into an integral equation of the second kind with the derivative of the payoff function with respect to the control as the unknown function. We state conditions sufficient for this integral equation to have a unique solution.

Identification conditions for dynamic games with discrete controls can be found in Magnac and Thesmar (2002) and Bajari, Chernozhukov, Hong, and Nekipelov (2009). Blevins (2009) gives identification conditions for dynamic games with both discrete and continuous controls. His result relies on the presence of discrete controls and does not apply to the purely continuous case considered in this paper.

Although our identification result is nonparametric, we develop a semiparametric estimator of the payoff function. Given the limited sample sizes typically available for dynamic games, a fully nonparametric estimator may not be informative. Therefore, like most papers on dynamic games, we assume that the payoff function is known up to a finite dimensional parameter, but leave the transition density and policy function nonparametrically specified. Our estimation procedure has two steps. First, the transition density and policy function are nonparametrically estimated. In the second step, the fact that at the true parameters the estimated policy function should maximize the value function is used to form an objective function. Since the state and action space are continuous, there is a continuum of optimality conditions. The parameter estimates minimize the norm of these conditions.

Our asymptotic results are stated at three levels of increasing specificity. First, we analyze generic two-step semiparametric minimum distance estimators with a continuum of estimating equations. We give high-level conditions on the initial nonparametric estimates and estimating equations sufficient for the finite dimensional parameters to be \sqrt{n} asymptotically normal. Next, we specialize this result to the estimating equations that arise from the optimality conditions of dynamic games, but remain agnostic about the specific form of the initial nonparametric estimates. Finally, we describe kernel estimates of the transition density and policy function that satisfy the conditions of the previous two theorems.

The remainder of this paper proceeds as follows. Section 2 describes the model. Section 3 gives conditions for identification. Section 4 describes the estimation procedure. Section 5 contains the asymptotic analysis. Section 6 concludes.

2. MODEL

We consider a dynamic decision process with state variables, x , and policy variable, i . We will refer to the decision maker as a firm. The state variables follow a controlled Markov process, and have a density which is absolutely continuous with respect to

Lebesgue measue, i.e.

$$f(x_{t+1}|\mathcal{I}_t) = f_{x'|x,i}(x_{t+1}|x_t, i_t),$$

where \mathcal{I}_t is all information available at time t . Each period a firm's payoff function depends on the current state, the firm's actions, and an iid private shock, η . We denote the payoff function by $\Pi(x, i, \eta)$. The firm's value function is:

$$V(x_0, \eta_0) = \max_{\{i_t=I_t(x_t, \eta_t)\}_{t=0}^{\infty}} E \left[\sum_{t=0}^{\infty} \delta^t \Pi(x_t, i_t, \eta_t) | x_0, i_0, \eta_0 \right].$$

where $I_t(x_t, \eta_t)$ denotes the policy function at time t . Under well-known regularity conditions, see e.g. Rust (1994), a generalized version of Blackwell's Theorem ensures that the value function is the unique solution to the Bellman equation,

$$V(x, \eta) = \max_i \Pi(x, i, \eta) + \delta E [V(x', \eta') | x, i] \quad (2.1)$$

and that there is a time invariant measurable optimal policy function, which we will denote by $I(x, \eta)$.

Although the problem above is written for a single agent making decisions in isolation, it can be applied to dynamic games with stationary Markov perfect equilibria as well. In this case, the state variables would include information on all firms in the same market. Suppose there are N firms, each indexed by f . The transition density for x_t given the actions of all firms can be written

$$F(x_{t+1}|x_t, \{i_{ft}\}_{f=1}^N) = F(x_{t+1}|x_t, \{I(x_t, \eta_{ft})\}_{f=1}^N).$$

However, since η_{ft} is private knowledge each firm must form expectations using

$$F(x_{t+1}|x_t, i_{ft}) = E[F(x_{t+1}|x_t, \{I(x_t, \eta_{\bar{f}t})\}_{\bar{f}=1}^N) | \eta_{ft}]. \quad (2.2)$$

Hence, each firm's decision problem takes the form written above. In equilibrium, the policy function must satisfy both (2.1) and (2.2). It would be necessary to take this into account if we want to perform any counterfactual analysis. However, our identification result and estimator will only rely on the individual optimality of each firm's actions. As such we will not be explicit about equilibrium considerations throughout most of the paper.

2.1. **Example.** This section describes a concrete example that fall under the above setup. Throughout the paper, we will return to this example to illustrate the plausibility of various assumptions.

Example 1 *Investment in natural gas pipelines.* In ongoing work, we apply the estimator developed in this paper to natural gas pipelines. For clarity, we describe a simplified version of this application. Investment in natural gas pipelines in the United States is subject to regulatory approval. Our goal is to recover the implicit cost of investment, including costs of regulatory compliance. The state variables are revenues, r_t , operating expenses, o_t , pipeline capacity, q_t , and pipeline utilization, u_t . Each firm tracks both its own values of these four variables, as well as those of other firms in the same market. We will use an f subscript to differentiate among firms when necessary. The control variable is pipeline investment measured in dollars, i_t . We specify the profit function of firm f as,

$$\Pi_f(x_t, i_{ft}, \eta_{ft}) = r_{ft} - o_{ft} - i_{ft} - c(i_{ft}, \eta_{ft}, q_t, u_t),$$

where the final term, $c(i_{ft}, \eta_{ft}, q_t, u_t)$ represents the cost of regulatory compliance. We include capacity and utilization in this function based on the assumption that the regulator primarily looks at the rate of capacity use when deciding how costly to make approval of new pipeline projects. Also, note that although the revenues and expenses of other firms do not enter the profit function directly, they will help the firm better forecast other firms' choices of investment. Therefore, revenues and expenses of other firms will still enter the policy function of this firm.

3. IDENTIFICATION

Given observations of $\{x_t, i_t\}$, our goal is to recover the payoff function, $\Pi(x, i, \eta)$. Rust (1994) shows that this is, in general, impossible because adding any function of x to the payoff function results in the same policy function. However, it is still possible to give a limited set of additional restrictions that can be used to fully identify the payoff function. With continuous states variables and discrete actions, Bajari, Chernozhukov, Hong, and Nekipelov (2009) show that if payoffs are additively separable in the private shock, η , and the distribution of η is known, then the payoff function can be identified up to a location normalization with an exclusion restriction. In the same setting, continuous states and discrete controls, Berry and Tamer (2006) give a converse result. They show that if the payoff function is known, then the

distribution of η can be identified. In a setting very similar to this paper's, Berry and Pakes (2001) propose a parametric estimator based on firm's first order conditions. Their estimator requires observations of realized profits. We will also consider an estimator based on firm's first-order conditions, but we will not require data on profits.

We begin by assuming that the transition density of the states and controls, $f_{x',i'|x,i}$, is identified. We then show that the optimality of observed actions imply that the payoff function must satisfy an integro-differential equation that depends on this transition density. We propose restrictions on the payoff function and transition density that guarantee that this integro-differential equation has a unique solution. We begin by stating our assumptions and main result. The assumptions will be discussed in more detail below.

Assumptions I (Payoff function identification).

- I1 The transition density, $f_{x_t, i_t | x_{t-1}, i_{t-1}}$, is identified.
- I2 The policy function, $I(x, \eta)$, is weakly increasing in η and $f_{\eta_t | \mathcal{I}_t} = f_{\eta_t}$ is known.
- I3 The policy function satisfies the firm's first order condition.
- I4 The discount factor, δ , is known.
- I5 For some value of the control, i_0 , the payoff function, $\Pi(x, i_0, \eta)$, is known for all x and η .
- I6 There exists a component of the state space, $x^{(k)}$ with compact support, such that:
 - $\frac{\partial \Pi}{\partial x^{(k)}}(x, i, \eta)$ is known and not identically zero for all x, η and,
 - The policy function is pseudo-invertible at i_0 with respect to $x^{(k)}$ in that for all $x^{(-k)}$ and η there exists a function $\chi_k(i_0, x^{(-k)}, \eta)$ such that

$$I(\chi_k(i_0, x^{(-k)}, \eta), x^{(-k)}, \eta) = i_0$$

and χ_k is measurable with respect to $x^{(-k)}$ and η .

- I7 Let $\frac{\partial \Pi}{\partial i} \in \mathcal{G}$. Define the following operators $\mathcal{G} \rightarrow \mathcal{G}$,

$$\begin{aligned} \mathcal{D}(g)(x, \eta) &= \frac{\partial}{\partial i_t} E \left[\sum_{\tau=0}^{\infty} \delta^\tau g(x_{t+\tau}, \eta_{t+\tau}) | x_t = x, i_t = I(x, \eta) \right] \\ \mathcal{L}(g)(x, \eta) &= \int_{\chi_k(i_0, x^{(-k)}, \eta)}^{x^{(k)}} g(\tilde{x}^{(k)}, x^{(-k)}, \eta) \frac{\partial I}{\partial x^{(k)}}(\tilde{x}^{(k)}, x^{(-k)}, \eta) d\tilde{x}^{(k)} \\ \mathcal{K}(g)(x, \eta) &= \mathcal{D}(\mathcal{L}(g))(x, \eta) \end{aligned}$$

The only solution in \mathcal{G} to

$$0 = g(x, \eta) + \mathcal{K}(g)(x, \eta) \quad (3.1)$$

is $g(x, \eta) = 0$.

I8 There is some component of the state x_i excluded from $\Pi(x, i, \eta)$, but entering $I(x, \eta)$.

Theorem 1 (Payoff function identification). *If conditions I1-I7 hold, then the equilibrium payoff function, $\Pi^*(x_t, \eta_t) = \Pi(x_t, I(x_t, \eta_t), \eta_t)$, and value function $V(x_t, \eta_t)$ are identified. If, additionally, condition I8 holds, then the payoff function, $\Pi(x, i, \eta)$ is identified.*

We now discuss the assumptions of this theorem. Its proof can be found in the next section. If the state vector, x_t is observed I1 is trivially satisfied. Hu and Shum (2008) give conditions that ensure I1 in the presence of unobserved components of the state vector. These conditions are discussed in detail in section 3.2 below.

Condition I2 is to allow $I(x, \eta)$ and the distribution of η to be recovered from the density of i and x . The requirement that the distribution of η is known is more a normalization than an assumption since η is an unobserved variable that enters unknown functions in an unrestricted manner.

Condition I3 is to ensure that the first order condition can be used to recover the payoff function. If it is not satisfied everywhere, such as with binding constraints on i or non-differentiable portions of $\Pi(\cdot)$, then the theorem can be adapted to apply for values of x and η where the first order condition is satisfied.

The first order condition can be manipulated to give an integral equation that can be solved for $\frac{\partial \Pi}{\partial i}$. Conditions I5, I6, and I7 guarantee that this integral equation has a unique solution. Conditions I5 and I6 are analogous to some of the conditions for identifying dynamic discrete games in Magnac and Thesmar (2002) or Bajari, Chernozhukov, Hong, and Nekipelov (2009). With discrete actions, it is necessary to normalize the payoff associated with one of the alternatives to zero. Condition I5 is the continuous analog of this normalization. With discrete actions, it is also common to assume that payoff function is additively separable in private shock. Here, the assumption that one of the partial derivatives of the payoff function is known, I6, takes the place of the assumption of additively separable private shocks. To make the analogy more clear, condition I6 could be placed on η instead of $x^{(k)}$.

Condition I7 is a key identifying assumption. It ensures that the integral equation for $\frac{\partial \Pi}{\partial i}$ has a unique solution. Similar conditions appear in the literature on nonparametric identification of simultaneous equations and are referred to as completeness conditions. In that literature, the completeness condition involves an integral equation of the first kind. Inversion of such equations is an ill-posed problem, which adds to the difficulty of estimation. Here, the completeness condition involves an integral equation of the second kind, so there is no ill-posed inverse problem.¹

It is difficult to verify condition I7 from more primitive conditions. One simple sufficient condition is that the operator $\mathcal{K} : \mathcal{G} \rightarrow \mathcal{G}$ defined above has operator norm less than one. However, there is no particular reason to believe that this condition would hold in applications. Fortunately, although it is difficult to determine a priori whether I7 holds, the condition can be checked for any given dataset. Condition I7 is an assumption about the transition density and the policy function. These two things can be identified without attempting to identify the payoff function or invoking assumption I7. Therefore, condition I7 can be checked. One could even develop a formal statistical test for I7, but we will leave that task for future work.

Finally, the exclusion restriction, I8, allows us to separate the structural payoff function from the equilibrium payoff function. Without the exclusion, we would only be able to identify the payoff function along the equilibrium path, that is only at $(x, I(x, \eta), \eta)$ rather than for any combination of (x, i, η) . While identifying the payoff function only along the equilibrium path is sufficient for some purposes, such as predicting payoffs in the current environment, identifying the payoff function everywhere is essential for counterfactual analysis.

We now discuss how these assumptions relate to our working example.

Example 1 *Investment in natural gas pipelines (continued)*. All of the state variables in this example are observed, so the transition density is identified, satisfying I1. Assumption I2 is simply a normalization. Condition I3, which says that first order condition holds, appears reasonable. It would follow from concavity in i of $c(i, \eta, q, u)$. One might think that investment involves a fixed adjustment cost, in which the first order condition need not hold when investment is zero. However, in this case we

¹Estimation still involves an ill-posed inverse problem in that it requires estimation of a conditional expectation operator. However, the inversion of that operator needed to solve for the value function is well-posed.

would observed a mass point at zero investment, and there does not appear to be any such mass point in the available data. Condition I5 can be met by assuming that when investment is zero, profits are simply revenues minus expenses. In other words, $c(0, \eta, q, u) = 0$. We have revenue data, so we satisfy the first part of I6 by assuming that profits are additively separable in revenues. The second part of the assumption is also very plausible. It is reasonable to believe that with sufficiently negative revenue, the firm would choose to shut down by setting i to a large negative value. Also, if revenues are sufficiently high, the firm should choose to invest a positive amount. By continuity, there should be some intermediate amount of revenue that leads to zero investment. As mentioned above, condition I7 is difficult to verify analytically, but it can be checked using estimates of the transition density and policy function. Finally, the exclusion restriction, I8, is met by assuming that each firm's payoffs depend on the firm's own revenue, but not the revenue of other firms. The revenue of others still enters the firm's policy function because the others' revenue helps each firm to predict others' actions, which can affect future payoffs.

3.1. Proof of identification (theorem 1). We first show that the policy function is identified. Assumption I1 implies that the conditional quantile function of actions given states is known. Denote it by $Q_{i|x}(\tau|x)$. Assumption I2 – that $I(x, \eta)$ is weakly increasing in η and F_η is known – imply that the policy function is

$$I(x, F_\eta^{-1}(\tau)) = Q_{i|x}(\tau|x).$$

Condition I3 says that the policy function satisfies the first order condition. As stated in the text, the first order condition is

$$0 = \frac{\partial \Pi}{\partial i}(x_t, I(x_t, \eta_t), \eta_t) + \frac{\partial}{\partial i} \sum_{\tau=1}^{\infty} \delta^\tau E[\Pi(x_{t+\tau}, I(x_{t+\tau}, \eta_{t+\tau}), \eta_{t+\tau}) | x_t, I(x_t, \eta_t)]. \quad (3.2)$$

We can express the payoff function in terms of its derivative,

$$\begin{aligned}
\Pi(x, I(x, \eta), \eta) &= \int_{x_0^{(k)}}^{x^{(k)}} \frac{\partial}{\partial x^{(k)}} [\Pi(x, I(x, \eta), \eta)] d\tilde{x}^{(k)} + \\
&\quad + \Pi\left(x_0^{(k)}, x^{(-k)}, I(x_0^{(k)}, x^{(-k)}, \eta), \eta\right) \\
&= \int_{\chi_k(i_0, x^{(-k)}, \eta)}^{x^{(k)}} \frac{\partial \Pi}{\partial i}(x, I(x, \eta), \eta) \frac{\partial I}{\partial x^{(k)}}(x, \eta) + \frac{\partial \Pi}{\partial x^{(k)}}(x, I(x, \eta), \eta) d\tilde{x}^{(k)} \\
&\quad + \Pi\left(\chi_k(i_0, x^{-k}, \eta), x^{-k}, i_0, \eta\right) \\
&= \mathcal{L}\left(\frac{\partial \Pi}{\partial i}\right)(x, \eta) + \Pi\left(\chi_k(i_0, x^{-k}, \eta), x^{-k}, i_0, \eta\right) + \\
&\quad + \int_{\chi_k(i_0, x^{(-k)}, \eta)}^{x^{(k)}} \frac{\partial \Pi}{\partial x^{(k)}}(x, I(x, \eta), \eta) d\tilde{x}^{(k)} \tag{3.3}
\end{aligned}$$

where the first line is just the fundamental theorem of calculus, and the second line comes from setting $x_0^{(-k)} = \chi_x(i_0, x^{(-k)}, \eta)$, and expanding $\frac{\partial}{\partial x^{(k)}} [\Pi(x, I(x, \eta), \eta)]$. Let

$$\begin{aligned}
\varphi(x, \eta) &= \Pi\left(\chi_k(i_0, x^{-k}, \eta), x^{-k}, i_0, \eta\right) + \\
&\quad + \int_{\chi_k(i_0, x^{(-k)}, \eta)}^{x^{(k)}} \frac{\partial \Pi}{\partial x^{(k)}}(x, I(x, \eta), \eta) d\tilde{x}^{(k)}
\end{aligned}$$

Note that by I5 and I6, $\varphi(x, \eta)$ is known. Substituting this and (3.3) into (3.2) gives,

$$\begin{aligned}
0 &= \frac{\partial \Pi}{\partial i}(x_t, I(x_t, \eta_t), \eta_t) + \frac{\partial}{\partial i} E \left[\sum_{\tau=1}^{\infty} \delta^\tau \mathcal{L}\left(\frac{\partial \Pi}{\partial i}\right)(x_{t+\tau}, \eta_{t+\tau}) + \varphi(x_{t+\tau}, \eta_{t+\tau}) \mid x_t, I(x_t, \eta_t) \right] \\
&= (\mathbf{1} + \mathcal{K}) \left(\frac{\partial \Pi}{\partial i} \right) + \mathcal{D}(\varphi)
\end{aligned}$$

We have assumed that $\frac{\partial \Pi}{\partial x^{(k)}}$ and $\Pi(x, i_0, \eta)$ are known, which implies that $\varphi(x, \eta)$ is known. Also, condition I7 ensures that the operator, $\mathbf{1} + \mathcal{K}$, in the above equation is invertible, so we can identify $\frac{\partial \Pi}{\partial i}(x, I(x, \eta), \eta)$. Integrating as in (3.3) gives $\Pi(x, I(x, \eta), \eta)$.

Finally, to recover the payoff function for at places where $i \neq I(x, \eta)$, we use the exclusion restriction I8. In particular this exclusion implies that the payoff function can be written as $\Pi(x^{(-i)}, i, \eta)$, so by varying $x^{(i)}$ while holding $x^{(-i)}$ and η fixed, we can identify the payoff function on the set $\{x^{(-i)}, i, \eta : i = I(x^{(i)}, x^{(-i)}, \eta)\}$. \square

3.2. Identifying transition densities with unobserved state variables . Theorem 1 above shows that given the transition density and some normalizations, we can identify the equilibrium payoff function, $\Pi^*(x, \eta)$. As stated above, when states

are fully observed, the transition density is easily identified. When some components of the state vector are unobserved, it can still be possible to identify the transition density. Hu and Shum (2008) give conditions to identify the transition density in the presence of unobserved state variables. Their identification argument relies on the eigendecomposition of an integral operator associated with the joint density of observed states.

Suppose the state vector can be written as $x_t = (w_t, \epsilon_t)$, where w_t is observed, but ϵ_t is not. Hu and Shum (2008) show that the transition density can be identified under the following assumptions:

Assumptions U (*Transition density identification with unobserved states*).

U1 The states and controls follow a stationary first-order Markov process,

$$f_{x_t, i_t | x_{t-1}, i_{t-1}, \mathcal{I}_{t-1}} = f_{x', i' | x, i}.$$

U2 Strategies are Markovian, so that

$$f_{i_t | w_t, \epsilon_t, w_{t-1}, \epsilon_{t-1}, i_{t-1}} = f_{i_t | w_t, \epsilon_t}.$$

U3 The unobserved states are not affected by the controls in that

$$f_{\epsilon_t | w_t, i_{t-1}, w_{t-1}, \epsilon_{t-1}} = f_{\epsilon_t | w_t, \epsilon_{t-1}}.$$

U4 There is sufficient variation in i_{t-1} with ϵ_{t-1} given fixed w_t and w_{t-1} , so that there exists a known function $\omega(i)$ such that $|E[\omega(i_{t-1}) | w_t, w_{t-1}, \epsilon_{t-1}]| < \infty$ and for any $\epsilon_{t-1} \neq \epsilon'_{t-1}$,

$$E[\omega(i_{t-1}) | w_t, w_{t-1}, \epsilon_{t-1}] \neq E[\omega(i_{t-1}) | w_t, w_{t-1}, \epsilon'_{t-1}].$$

U5 There is sufficient variation in i_t with i_{t-2} given fixed w_t and w_{t-1} so that the operators defined by

$$(L_{I_t, w_t, \omega | w_{t-1}, I_{t-2}} h)(i') = \int h(i_{t-2}) \int \omega(i_{t-1}) f_{i_t, w_t, i_{t-1} | w_{t-1}, i_{t-2}}(i', w_t, i_{t-1} | w_{t-1}, i_{t-2}) di_{t-1} di_{t-2}$$

are one-to-one for each w_t, w_{t-1} .

U6 There is variation in i_t with ϵ_{t-1} given w_t, w_{t-1} so that the operators defined by

$$(L_{I_t | w_t, w_{t-1}, \epsilon_{t-1}} h)(i') = \int h(\epsilon_{t-1}) f_{i_t | w_t, w_{t-1}, \epsilon_{t-1}}(i' | w_t, w_{t-1}, \epsilon_{t-1}) d\epsilon_{t-1}$$

are one-to-one for each w_t, w_{t-1} .

U7 There is variation in i_t with ϵ_t given w_t so that the operators defined by:

$$(L_{i_t|w_t, \epsilon_t} h)(i) = \int h(\epsilon) f_{i_t|w_t, \epsilon_t}(i|w_t, \epsilon) d\epsilon$$

are one-to-one for each w_t .

U8 For each component of ϵ_{t-1} , there exists a known functional G_k such that

$$G_k [f_{i_t|w_t, w_{t-1}, \epsilon_{t-1}}(\cdot|w_t, w_{t-1}, \epsilon_{t-1})]$$

is monotonic in $\epsilon_{t-1, k}$.

Theorem 2 (Hu and Shum). *If assumptions U1-U8, then the observed density of $(i_t, w_t, i_{t-1}, w_{t-1}, i_{t-2}, w_{t-2})$ uniquely determines the equilibrium transition density $f_{i', w', \epsilon'|i, w, \epsilon}$.*

Assumptions U1 and U2, which state that the environment is stationary and Markovian, are standard for dynamic games and have already been made above. Assumption U3 restricts the nature of the unobserved state variables. It says that the unobserved state variables are independent of the control in the previous period conditional on the previous state and current observed state variable.

Assumptions U4-U7 ensure that there is a sufficiently strong relationship between observed and unobserved variables so that the distribution of observed variables can be used to construct the distribution of unobservables. Although assumptions U5-U7 can be difficult to fully verify, simple necessary conditions for each can be expected to hold. For example, U7 requires that conditional on each observed state, the unobserved state influences firms' actions. If this were not true for at least some observed state, then the unobserved state is irrelevant and should not be part of the model. Similarly, assumptions U5 and U6 require that actions be related to unobserved states and past and current unobserved states be related. Additionally, assumptions U6 and U7 require that the support of i_t be at least as "large" as the support of ϵ_t . For example, if ϵ_t is continuous with k dimensions, i_t must be continuous with at least k dimensions. If ϵ_t is discrete with k points of support, i_t must have at least k points of support.

Assumption U4 is easier than assumptions U5-U7 to verify. However, it also fails in some common models. In particular assumption U4 rules out models where there is a deterministic relationship between the current state, past action, and past state.

For example, if the state includes the capital stock, the action is investment, and depreciation is deterministic, then $i_t = w_t - (1 - \delta)w_{t-1}$ and there is no variation with ϵ in $E[\omega(i_t)|w_t, w_{t-1}, \epsilon_t]$. Nonetheless, even in models with deterministic accumulation, the above result can be useful. Since investment can be recovered from w_{t+1} and w_t , it can be excluded from the model while identifying the transition density. Then if remaining actions satisfy the above conditions, the transition density is identified. In practice, candidates for these remaining actions include anything that satisfy assumptions U2 and U3. That is, they must be functions of only the current state and should not influence the next unobserved state. Even actions that have no dynamic implications, such as variable inputs would be suitable. Additionally, depending on the model and interpretation of ϵ_t , outcome variables such as output, revenue, or costs might be suitable.

4. ESTIMATION

Although the above identification result is nonparametric, our estimation approach is semiparametric. We assume that the payoff function is known up to a finite dimensional parameter, but allow the transition distribution and policy function to be nonparametric. As in much of the dynamic game literature, we estimate the model in multiple steps. First, we estimate that transition distribution of state variables and the policy function. We then use the estimated policy function and transition distribution along with the optimality conditions of the model to estimate the payoff function.

4.1. Estimating policy functions. We begin by supposing we have some nonparametric estimates of the transition density, $\hat{f}_{x',i'|x,i}$, and the conditional distribution of actions given states, $\hat{F}_{i|x}$. Theorem 4 gives high level conditions on these estimates for the parameters of the payoff function to be consistent and \sqrt{n} asymptotically normal. In section 5.2 below, we show that certain kernel estimates satisfy these high level conditions. We expect that sieve estimates would as well.

Given $\hat{F}_{i|x}$, estimation of the policy function is straightforward. We have assumed that the policy function is weakly increasing in η , that η is independent of x , and that the distribution of η is known. Therefore, an estimate of the inverse policy function

can be formed from the following relationship:

$$F_\eta(\eta) = \hat{F}_{i|x}(i|x)$$

$$\hat{\eta}(i, x) = F_\eta^{-1} \left[\hat{F}_{i|x}(i|x) \right].$$

4.2. Estimating the payoff function. The payoff function is estimated from the optimality conditions of the model. To do this, we must solve for the value function given the above estimates and a candidate parameter value. This can be done by evaluating the Bellman equation at the estimated policy function using the estimated transition densities to evaluate the conditional expectation:

$$\hat{V}(x, i; \theta, \hat{f}, \hat{\eta}) = \Pi(x, i, \hat{\eta}(i, x); \theta) + \delta E_{\hat{f}} \left[\hat{V}(x', i'; \theta, \hat{f}, \hat{\eta}) | x, i \right]. \quad (4.1)$$

This functional equation can be used to solve for $\hat{V}(x, i; \theta)$, the value of being in state x with a private shock such that i is the estimated policy.² Let the value of deviating from the estimated policy for one period be denoted as

$$\mathcal{V}(x, \eta, \tilde{i}; \theta, \hat{f}, \hat{\eta}) = \Pi(x, \tilde{i}, \eta; \theta) + \delta E_{\hat{f}} \left[\hat{V}(x', \eta'; \theta, \hat{f}, \hat{\eta}) | x, \tilde{i} \right].$$

The optimality conditions of the model can be used to form moment conditions to estimate the profit function. At least three ways of using the optimality conditions of the model to estimate the payoff function have appeared in the literature. We will show that these three methods result in asymptotically equivalent estimators when the moments are weighted appropriately. Jenkins, Liu, Matzkin, and McFadden (2004) minimize the difference between observed actions and optimal actions. This action based moment condition is:

$$0 = m_a(i, x; \theta, \hat{h}) = i - \arg \max_{\tilde{i}} \mathcal{V}(x, \hat{\eta}(i, x), \tilde{i}; \theta, \hat{h}),$$

where $\hat{h} = (\hat{\eta}, \hat{f})$. Alternatively, the first order condition for investment can be used as a moment condition:

$$0 = m_f(i, x; \theta, \hat{h}) = \frac{\partial \mathcal{V}}{\partial i}(x, \hat{\eta}(i, x), i; \theta, \hat{h}).$$

²Since, given the state, investment and the private shock contain the same information, we could have instead defined a value function over x and η . In this case, the i in the right hand side of (4.1) would be replaced with $\hat{I}(x, \eta)$. Although defining the value function over x and η would be more standard, we chose not to do this because estimating the inverse policy function is slightly more straightforward than estimating the policy function.

Hong and Shum (2009) use this kind of first order condition based moment. A third candidate moment condition is the difference between the value of observed actions and the maximized value function. This type of moment condition is used by Macieira (2009):

$$0 = m_v(i, x; \theta, \hat{h}) = \hat{V}(x, \eta; \theta, \hat{h}) - \max_{\hat{i}} \mathcal{V}(x, \hat{\eta}(i, x), \tilde{i}; \theta, \hat{h}).$$

Note that each of these conditions hold for all i, x . An objective function can be formed by taking a combination of them. Since the estimates of the transition distribution and policy are estimated nonparametrically, we must use a continuum of moment conditions in the objective function to enable the payoff function parameters to be \sqrt{n} -consistent. Let $m(\cdot)$ be one of $m_a(\cdot)$, $m_f(\cdot)$, or $m_v(\cdot)$. The objective function is

$$\begin{aligned} Q_n(\theta; \hat{h}) &= \int_{\mathcal{I} \times \mathcal{X}} \int_{\mathcal{I} \times \mathcal{X}} m(i, x; \theta, \hat{h}) w_n(i, x; i', x') m(i', x'; \theta, \hat{h}) di dx di' dx' \\ &= \left\| B_n m(\cdot; \theta, \hat{h}) \right\|_n^2 \end{aligned}$$

where $w_n(i, x; i', x')$ is some weighting function, possibly data-dependent, and the second line is simply alternate notation for the first. In particular, $\|f\|^2 = \int f(i, x)^2 di dx$, and B_n is an integral operator with a kernel that satisfies

$$w_n(i, x; i', x') = \int b_n(i, x; \tilde{i}, \tilde{x}) b_n(i', x'; \tilde{i}, \tilde{x}) d\tilde{i} d\tilde{x}.$$

The parameters of the payoff function are estimated by minimizing the objective function,

$$\hat{\theta} = \arg \min_{\theta \in \Theta} Q_n(\theta; \hat{h}).$$

5. ASYMPTOTIC THEORY

The above estimator of θ is similar to the estimation frameworks of Carrasco and Florens (2000); Carrasco, Chernov, Florens, and Ghysels (2007); and Ai and Chen (2003). As a result, we can derive the asymptotic distribution of $\hat{\theta}$ using a suitable adaptation of their arguments. The current situation differs from the setup of Carrasco and Florens (2000) and Carrasco, Chernov, Florens, and Ghysels (2007) in that the moment conditions here include a preliminary nonparametric estimate. Ai and Chen (2003) study conditional moment conditions with nonparametric components. They derive their results for independent data, and so their weighting operator is

diagonal. Here, the data is not independent. Also, Ai and Chen (2003) consider simultaneous estimation of the nonparametric and parametric components of their model. Here, the nonparametric and parametric components are estimated in two steps.

The following theorem gives conditions on $m(\cdot; \theta, h)$ and \hat{h} sufficient for asymptotic normality. The conditions are stated for generic moment functions $m(\cdot; \theta, h)$ and nonparametric estimate \hat{h} . Theorem 4 gives more specific conditions on the payoff function and transition density that imply the generic conditions of Theorem 3 for each of the three versions of $m(\cdot; \theta, h)$ defined above. Section 5.2 describes kernel estimates of the transition density and inverse policy function that satisfy the generic conditions of 3. Let $\langle f, g \rangle$ denote the inner-product associated with the norm used to define Q_n , i.e.

$$\langle f, g \rangle = \int f(i, x)g(i, x)di dx.$$

For any functions, f and g , let $f(x) \lesssim g(x)$ mean that there exist a constant M such that $f(x) \leq Mg(x)$. Similarly, let $f(x) \lesssim_p g(x)$ mean that $f(x) \leq Mg(x) + o_p(1)$ uniformly in x .

Assumptions A (*Asymptotic normality for semiparametric minimum distance estimators in a Hilbert space*).

- A1 $\hat{\theta}$ approximately minimizes $Q_n(\theta)$, so that $Q_n(\hat{\theta}) \leq \inf_{\theta \in \Theta} Q_n(\theta) + o_p(n^{-1/2})$.
- A2 Let $\|B\|_{op}$ be the usual operator norm, $\|B\|_{op} = \sup_{\|f\|=1} \|Bf\|$, assume that $\|B_n - B\|_{op} \xrightarrow{p} 0$.
- A3 The nonparametric estimates of the transition density and policy function converge at a rate faster than $n^{-1/4}$ with respect to some norm $\|\cdot\|_{\mathcal{H}}$, in that

$$\left\| B(\hat{h} - h) \right\|_{\mathcal{H}} = o_p(n^{-1/4}).$$

A4 The following derivatives exist and satisfy the following conditions:

- A4.i For all θ with $\|\theta - \theta_0\| \leq \delta_n$, $m(\cdot; \theta, h)$ is pathwise differentiable with respect to h at h_0 . We denote this derivative as $D_h^m(\theta, h_0)$. This derivative is such that

$$\sup_{\theta: \|\theta - \theta_0\| \leq \delta_n} \left\| B_n \left(m(\cdot; \theta, \hat{h}) - m(\cdot; \theta, h_0) - D_h^m(\theta, h_0)(\hat{h} - h_0) \right) \right\| \lesssim_p \left\| \hat{h} - h_0 \right\|_{\mathcal{H}}^2$$

and

A4.ii

$$\sup_{\theta: \|\theta - \theta_0\| \leq \delta_n} \left\| B_n \left(D_m^h(\theta, h_0)(\hat{h} - h_0) - D_m^h(\theta_0, h_0)(\hat{h} - h_0) \right) \right\| \lesssim_p \left\| \hat{h} - h_0 \right\|_{\mathcal{H}} \|\theta - \theta_0\|.$$

A4.iii In a neighborhood of θ_0 , $m(\cdot; \theta, h_0)$ is continuously differentiable with respect to θ . This derivative is bounded away from zero in that

$$1 / \left(\inf_{\theta: \|\theta - \theta_0\| \leq \delta_n} \left\| B_n \frac{\partial m}{\partial \theta}(\cdot; \theta, h_0) \right\| \right) \lesssim_p 1,$$

and this derivative is Lipschitz continuous, so that

$$\sup_{\theta_1, \theta_2: \|\theta_i - \theta_0\| \leq \delta_n} \left\| B_n \left(\frac{\partial m}{\partial \theta}(\cdot; \theta_1, h_0) - \frac{\partial m}{\partial \theta}(\cdot; \theta_2, h_0) \right) \right\| \lesssim_p \|\theta_1 - \theta_2\|.$$

A4.iv In a neighborhood of θ_0 , $\|B_n m(\cdot; \theta, h_0)\|^2$ is continuously differentiable with respect to θ with derivative

$$\frac{\partial}{\partial \theta} \|B_n m(\cdot; \theta, h_0)\|^2 = 2 \left\langle B_n \frac{\partial m}{\partial \theta}(\cdot; \theta, h_0), B_n m(\cdot; \theta, h_0) \right\rangle.$$

A5 \hat{h} and B_n are bounded in the sense that there is some constant M such that with probability approaching one

$$\left\| B_n m(\cdot; \theta, \hat{h}) \right\| \leq M$$

and

$$\left\| B_n \left(m(\cdot; \theta, h_0) + D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\| \leq M.$$

A6

$$\sqrt{n} \left\langle B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), B D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle \xrightarrow{d} N(0, \Omega_B)$$

Theorem 3 (Asymptotic distribution for semiparametric minimum distance estimators in a Hilbert space). *If A1-A6 then $\sqrt{n}(\hat{\theta} - \theta_0) \xrightarrow{d} N(0, M_B^{-1} \Omega_B M_B^{-1})$ with $M_B = \langle B \frac{\partial m}{\partial \theta}, B \frac{\partial m}{\partial \theta} \rangle^{-1}$.*

Proof. We will show that $\hat{\theta}$, which minimizes

$$Q_n(\theta) = \left\| B_n m(\cdot; \theta, \hat{h}) \right\|^2$$

is close to the minimizer of the linearized objective function,

$$L_n(\theta) = \left\| B_n \left(m(\cdot; \theta, h_0) + D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\|^2.$$

First, we must show that $Q_n(\theta)$ is close to $L_n(\theta)$ in that

$$\sup_{\|\theta - \theta_0\| \leq \delta_n} |Q_n(\theta) - L_n(\theta)| = o_p(n^{-1/2}).$$

Note that

$$\begin{aligned} |Q_n(\theta) - L_n(\theta)| &= \left| \left\| B_n m(\cdot; \theta, \hat{h}) \right\|^2 - \left\| B_n \left(m(\cdot; \theta, h_0) + D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\|^2 \right| \\ &= \left| \left(\left\| B_n m(\cdot; \theta, \hat{h}) \right\| - \left\| B_n \left(m(\cdot; \theta, h_0) + D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\| \right) \times \right. \\ &\quad \left. \times \left(\left\| B_n m(\cdot; \theta, \hat{h}) \right\| + \left\| B_n \left(m(\cdot; \theta, h_0) + D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\| \right) \right| \\ &\lesssim \left| \left\| B_n m(\cdot; \theta, \hat{h}) \right\| - \left\| B_n \left(m(\cdot; \theta, h_0) + D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\| \right| \\ &\lesssim \left\| B_n \left(m(\cdot; \theta, \hat{h}) - m(\cdot; \theta, h_0) - D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\| \end{aligned}$$

where we used assumption A5 and the reverse triangle inequality. Note that

$$\begin{aligned} \left\| B_n \left(m(\cdot; \theta, \hat{h}) - m(\cdot; \theta, h_0) - D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\| &\leq \\ &\leq \left\| B_n \left(m(\cdot; \theta, \hat{h}) - m(\cdot; \theta, h_0) - D_h^m(\theta, h_0)(\hat{h} - h_0) \right) \right\| + \\ &\quad + \left\| B_n \left(D_h^m(\theta, h_0)(\hat{h} - h_0) - D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\| \end{aligned}$$

Invoking assumptions A4.i and A4.ii gives

$$\begin{aligned} \left\| B_n \left(m(\cdot; \theta, \hat{h}) - m(\cdot; \theta, h_0) - D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\| &\lesssim_p \left\| \hat{h} - h_0 \right\|_{\mathcal{H}}^2 + \\ &\quad + \left\| \hat{h} - h_0 \right\|_{\mathcal{H}} \|\theta - \theta_0\| \end{aligned}$$

Using assumption A3, we have

$$\left\| B_n \left(m(\cdot; \theta, \hat{h}) - m(\cdot; \theta, h_0) - D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\| \lesssim_p o_p(n^{-1/2}) + \quad (5.1)$$

$$+ o_p(n^{-1/4}) \left\| \hat{\theta} - \theta_0 \right\|. \quad (5.2)$$

We will now show that $\left\| \hat{\theta} - \theta_0 \right\| \leq O_p(n^{-1/2})$. By assumption A4.iii, we can take a mean value expansion in θ around θ_0 , so that

$$m(\cdot; \hat{\theta}, h_0) = m(\cdot; \theta_0, h_0) + \frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}, h_0)(\hat{\theta} - \theta_0)$$

Therefore,

$$m(\cdot; \hat{\theta}, \hat{h}) - m(\cdot; \hat{\theta}, h_0) + \left(m(\cdot; \theta_0, h_0) + \frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}, h_0)(\hat{\theta} - \theta_0) \right) = 0$$

Rearranging and using the fact that $m(\cdot; \theta_0, h_0) = 0$, we have

$$\left\| B_n \frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}, h_0)(\hat{\theta} - \theta_0) \right\| = \left\| B_n \left(m(\cdot; \hat{\theta}, \hat{h}) - m(\cdot; \hat{\theta}, h_0) \right) \right\|$$

By A4.i we have

$$\begin{aligned} \left\| B_n \frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}, h_0)(\hat{\theta} - \theta_0) \right\| &= \left\| B_n D_h^m(\hat{h} - h_0) \right\| + o_p(n^{-1/2}) \\ \left\| \hat{\theta} - \theta_0 \right\| &\lesssim \left\| B D_h^m(\hat{h} - h_0) \right\| + o_p(n^{-1/2}) \end{aligned}$$

Thus, by assumptions A4.iii and A6, we can conclude that $\|\hat{\theta} - \theta_0\| \leq O_p(n^{-1/2})$.

Combining the just obtained rate bound for $\hat{\theta}$ with equation (5.2) gives:

$$|Q_n(\theta) - L_n(\theta)| = o_p(n^{-1/2}).$$

Now we show that the minimizer of $Q_n(\theta) = \|B_n m(\cdot; \theta, \hat{h})\|$ is close to the minimizer of $L_n(\theta) = \left\| B_n \left(m(\cdot, \theta, h_0) + D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\|$. Let $\hat{\theta}$ denote the former and $\tilde{\theta}$ denote the later. By assumption A1,

$$Q_n(\hat{\theta}) \leq Q_n(\tilde{\theta}) + o_p(n^{-1/2}).$$

Using the expansion results above we then have,

$$L_n(\hat{\theta}) \leq L_n(\tilde{\theta}) + o_p(n^{-1/2}).$$

On the other hand, $\tilde{\theta}$ is defined as the minimizer of $L_n(\theta)$, so $L_n(\tilde{\theta}) \leq L_n(\hat{\theta})$ and we can conclude that

$$L_n(\hat{\theta}) = L_n(\tilde{\theta}) + o_p(n^{-1/2}).$$

By assumption A4.iv we can expand both sides around θ_0 to get

$$\begin{aligned} \left\langle B_n \frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}_1, h_0), B_n D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle (\hat{\theta} - \theta_0) &= \\ &= \left\langle B_n \frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}_2, h_0), B_n D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle (\tilde{\theta} - \theta_0) + o_p(n^{-1/2}) \end{aligned}$$

Rearranging gives

$$\begin{aligned} \left\langle B_n \frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}_2, h_0), B_n D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle (\tilde{\theta} - \hat{\theta}) &= \\ &= \left\langle B_n \left(\frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}_1, h_0) - \frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}_2, h_0) \right), B_n D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle (\hat{\theta} - \theta_0) + o_p(n^{-1/2}) \end{aligned}$$

By assumption A4.iii,

$$\left\| \frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}_1, h_0) - \frac{\partial m}{\partial \theta}(\cdot; \bar{\theta}_2, h_0) \right\| \lesssim \|\bar{\theta}_1 - \bar{\theta}_2\|.$$

Also we have already shown that $\|\hat{\theta} - \theta_0\| \leq O_p(n^{-1/2})$, which implies that $\|\bar{\theta}_1 - \bar{\theta}_2\| \leq O_p(n^{-1/2})$. Thus, we can conclude that $\|\tilde{\theta} - \hat{\theta}\| = o_p(n^{-1/2})$.

Finally, expanding the first order condition for the linearized objective function gives,

$$\begin{aligned} 0 &= \left\langle B_n \frac{\partial m}{\partial \theta}(\cdot, \tilde{\theta}, h_0), B_n \left(m(\cdot, \tilde{\theta}, h_0) + D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\rangle \\ 0 &= \left\langle B_n \frac{\partial m}{\partial \theta}(\cdot, \tilde{\theta}, h_0), B_n \left(m(\cdot, \theta_0, h_0) + \frac{\partial m}{\partial \theta}(\cdot, \tilde{\theta}, h_0)(\tilde{\theta} - \theta_0) + \right. \right. \\ &\quad \left. \left. + D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right) \right\rangle \\ (\tilde{\theta} - \theta_0) &= - \left\langle B_n \frac{\partial m}{\partial \theta}(\cdot, \tilde{\theta}, h_0), B_n \frac{\partial m}{\partial \theta}(\cdot, \tilde{\theta}, h_0) \right\rangle^{-1} \times \\ &\quad \times \left\langle B_n \frac{\partial m}{\partial \theta}(\cdot, \tilde{\theta}, h_0), B_n D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle \\ &= - \left\langle B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0) \right\rangle^{-1} \times \\ &\quad \times \left\langle B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), B D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle + o_p(n^{-1/2}) \end{aligned}$$

To conclude, we have

$$\begin{aligned} \sqrt{n}(\hat{\theta} - \theta_0) &= - \sqrt{n} \left\langle B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0) \right\rangle^{-1} \times \\ &\quad \times \left\langle B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), B D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle + o_p(1) \end{aligned}$$

and the statement in the theorem follows from assumption A6. \square

We now give conditions on the payoff function and other model primitives such that the conditions of theorem 3 are satisfied. We begin by establishing differentiability of $m(\cdot; \theta, h)$ for each of the three candidates for $m(\cdot; \theta, h)$ introduced above. Each $m(\cdot; \theta, h)$ is simple function of $\mathcal{V}(x, \eta, i; \theta, f, \boldsymbol{\eta})$, so Lemma 1 about the derivatives of \mathcal{V} is useful. Before stating this lemma, we define some additional notation. Let $\mathcal{E}_f(g)(x, i)$ denote the following conditional expectation operator:

$$\mathcal{E}_f(g)(x, i) = E[g(x', i') | x, i]$$

where $g : X \times I \rightarrow \mathbb{R}$ and the expectation is taken with respect to the transition density, $f_{x',i'|x,i}$. Also, let \mathbf{I} denote the identity operator, $g \circ \boldsymbol{\eta} = g(x, i, \boldsymbol{\eta}(i, x))$ for $g : X \times I \times \mathcal{N} \rightarrow \mathbb{R}^k$, and let $\Pi_\theta(x, i, \eta) = \Pi(x, i, \eta; \theta)$. Note that $(\mathbf{I} - \delta\mathcal{E}_f)^{-1}$ exists because $\|\delta\mathcal{E}_f\|_{op} = \delta < 1$. Using this notation, we can write \mathcal{V} as

$$\mathcal{V}(x, \eta, i; \theta, \hat{f}, \hat{\boldsymbol{\eta}}) = \Pi_\theta(x, i, \eta) + \delta\mathcal{E}_{\hat{f}} \left((\mathbf{I} - \delta\mathcal{E}_{\hat{f}})^{-1} \Pi_\theta \circ \hat{\boldsymbol{\eta}} \right) (x, i, \eta)$$

Lemma 1 (Differentiability of $\mathcal{V}(x, \eta, \tilde{i}; \theta, f, \boldsymbol{\eta})$). *If*

- (1) *The payoff function is twice continuously differentiable with respect to i , η , and θ for all (i, x, η) , and for all θ in a neighborhood of θ_0 . The derivative with respect to θ is bounded by a function with finite expectation with respect to f , i.e.*

$$\left| \frac{\partial \Pi}{\partial \theta} \right| \leq m(i, x)$$

with

$$\int m(i', x') f_0(i', x' | i, x) di' dx' < \infty$$

uniformly in (i, x) . The derivatives with respect to η are Lipschitz continuous in that for all θ in a neighborhood of θ_0 ,

$$\left| \frac{\partial \Pi_\theta}{\partial \eta}(x, i, \eta_1) - \frac{\partial \Pi_\theta}{\partial \eta}(x, i, \eta_2) \right| \leq M_1(i, x) |\eta_1 - \eta_2|$$

with $\|B_n M_1\|$ bounded. Similarly, for all θ in a neighborhood of θ_0 ,

$$\left| \frac{\partial^2 \Pi_\theta}{\partial \eta \partial i}(x, i, \eta_1) - \frac{\partial^2 \Pi_\theta \partial i}{\partial \eta}(x, i, \eta_2) \right| \leq M_2(i, x) |\eta_1 - \eta_2|$$

with $\|B_n M_2\|$ bounded.

- (2) *The conditional expectation operator, \mathcal{E}_f is twice continuously differentiable respect to i .*

then \mathcal{V} is twice continuously differentiable with respect to i and θ at $h = h_0$, for all i , and for all θ in a neighborhood of θ_0 . Moreover, these derivatives are given by:

$$\frac{\partial \mathcal{V}}{\partial \theta}(x, i, \eta; \theta, f_0, \boldsymbol{\eta}_0) = \frac{\partial \Pi_\theta}{\partial \theta}(x, i, \eta) + \delta \mathcal{E}_{f_0} \left((I - \delta \mathcal{E}_{f_0})^{-1} \frac{\partial \Pi_\theta}{\partial \theta} \circ \boldsymbol{\eta}_0 \right) (x, i) \quad (5.3)$$

$$\frac{\partial \mathcal{V}}{\partial i}(x, i, \eta; \theta, f_0, \boldsymbol{\eta}_0) = \frac{\partial \Pi_\theta}{\partial i}(x, i, \eta) + \delta \frac{\partial \mathcal{E}_{f_0}}{\partial i} \left((I - \delta \mathcal{E}_{f_0})^{-1} \Pi_\theta \circ \boldsymbol{\eta}_0 \right) (x, i) \quad (5.4)$$

$$\frac{\partial^2 \mathcal{V}}{\partial i^2}(x, i, \eta; \theta, f_0, \boldsymbol{\eta}_0) = \frac{\partial^2 \Pi_\theta}{\partial i^2}(x, i, \eta) + \delta \frac{\partial^2 \mathcal{E}_{f_0}}{\partial i^2} \left((I - \delta \mathcal{E}_{f_0})^{-1} \Pi_\theta \circ \boldsymbol{\eta}_0 \right) (x, i) \quad (5.5)$$

$$\frac{\partial^2 \mathcal{V}}{\partial i \partial \theta}(x, i, \eta; \theta, f_0, \boldsymbol{\eta}_0) = \frac{\partial^2 \Pi_\theta}{\partial i \partial \theta}(x, i, \eta) + \delta \frac{\partial \mathcal{E}_{f_0}}{\partial i} \left((I - \delta \mathcal{E}_{f_0})^{-1} \frac{\partial \Pi_\theta}{\partial \theta} \circ \boldsymbol{\eta}_0 \right) (x, i). \quad (5.6)$$

Also, \mathcal{V} is pathwise differentiable with respect to f and $\boldsymbol{\eta}$ in a neighborhood of f_0 and $\boldsymbol{\eta}_0$ with derivatives

$$D_{\boldsymbol{\eta}}^{\mathcal{V}}(\theta, f, \boldsymbol{\eta})(\boldsymbol{\eta}_1 - \boldsymbol{\eta}) = \delta \mathcal{E}_f \left((I - \delta \mathcal{E}_f)^{-1} \left(\left(\frac{\partial \Pi_\theta}{\partial \boldsymbol{\eta}} \circ \boldsymbol{\eta} \right) (\boldsymbol{\eta}_1(\cdot) - \boldsymbol{\eta}(\cdot)) \right) \right) \quad (5.7)$$

$$\begin{aligned} D_f^{\mathcal{V}}(\theta, f, \boldsymbol{\eta})(f_1 - f) &= \delta \mathcal{E}_f \left[(I - \delta \mathcal{E}_f)^{-1} \left(\delta \int \left((I - \delta \mathcal{E}_f)^{-1} \Pi_\theta \circ \boldsymbol{\eta} \right) (x', i') \times \right. \right. \\ &\quad \left. \left. \times (f_1(x', i'|\cdot) - f(x', i'|\cdot)) dx' di' \right) \right] + \\ &\quad + \delta \int \left((I - \delta \mathcal{E}_f)^{-1} \Pi_\theta \circ \boldsymbol{\eta} \right) (x', i') (f_1(x', i'|\cdot) - f(x', i'|\cdot)) dx' di'. \end{aligned} \quad (5.8)$$

Finally, $D_{\boldsymbol{\eta}}^{\mathcal{V}}$ and $D_f^{\mathcal{V}}$ are continuously differentiable with respect to i these derivatives are equal to the pathwise derivatives of $\frac{\partial \mathcal{V}}{\partial i}$ with respect to $\boldsymbol{\eta}$ and f , which also exist.

These derivatives are given by

$$\frac{\partial D_{\boldsymbol{\eta}}^{\mathcal{V}}}{\partial i}(\theta, f, \boldsymbol{\eta})(\boldsymbol{\eta}_1 - \boldsymbol{\eta}) = \delta \frac{\partial \mathcal{E}_f}{\partial i} \left((I - \delta \mathcal{E}_f)^{-1} \left(\left(\frac{\partial \Pi_\theta}{\partial \boldsymbol{\eta}} \circ \boldsymbol{\eta} \right) (\boldsymbol{\eta}_1(\cdot) - \boldsymbol{\eta}(\cdot)) \right) \right) \quad (5.9)$$

$$\begin{aligned} \frac{\partial D_f^{\mathcal{V}}}{\partial i}(\theta, f, \boldsymbol{\eta})(f_1 - f) &= \delta \frac{\partial \mathcal{E}_f}{\partial i} \left[(I - \delta \mathcal{E}_f)^{-1} \left(\delta \int \left((I - \delta \mathcal{E}_f)^{-1} \Pi_\theta \circ \boldsymbol{\eta} \right) (x', i') \times \right. \right. \\ &\quad \left. \left. \times (f_1(x', i'|\cdot) - f(x', i'|\cdot)) dx' di' \right) \right] + \\ &\quad + \delta \int \left((I - \delta \mathcal{E}_f)^{-1} \Pi_\theta \circ \boldsymbol{\eta} \right) (x', i') \times \\ &\quad \times \left(\frac{\partial f_1}{\partial i}(x', i'|\cdot) - \frac{\partial f}{\partial i}(x', i'|\cdot) \right) dx' di', \end{aligned} \quad (5.10)$$

and these derivatives are Lipschitz continuous with respect to $\boldsymbol{\eta}$, f , and θ uniformly in a neighborhood of $\boldsymbol{\eta}_0$, f_0 , and θ_0 .

Proof. Assumption 1 allows us to apply the dominated convergence theorem to show that

$$\frac{\partial}{\partial \theta} \mathcal{E}_f \Pi(\cdot; \theta) = \mathcal{E}_f \frac{\partial \Pi}{\partial \theta}(\cdot; \theta)$$

and

$$\frac{\partial}{\partial \theta} (\mathbf{I} - \delta \mathcal{E}_f)^{-1} \Pi(\cdot; \theta) = (\mathbf{I} - \delta \mathcal{E}_f)^{-1} \frac{\partial \Pi}{\partial \theta}(\cdot; \theta)$$

Given this, (5.3)-(5.6) follow directly from the assumed differentiability of the payoff function and conditional expectation.

We now show (5.7). Recall that

$$\mathcal{V}(x, i, \eta; \theta, f, \boldsymbol{\eta}) = \Pi_\theta(x, i, \eta) + \delta \mathcal{E}_f \left[(\mathbf{I} - \delta \mathcal{E}_f)^{-1} \Pi_\theta \circ \boldsymbol{\eta} \right]$$

Let $\mathcal{V}_j = \mathcal{V}(x, i, \eta; \theta, f, \boldsymbol{\eta}_j)$ for $j = 1, 2$. Note that,

$$\begin{aligned} \mathcal{V}_1 - \mathcal{V}_2 &= \delta \mathcal{E}_f \left[(\mathbf{I} - \delta \mathcal{E}_f)^{-1} \Pi_\theta \circ \boldsymbol{\eta}_1 \right] - \delta \mathcal{E}_f \left[(\mathbf{I} - \delta \mathcal{E}_f)^{-1} \Pi_\theta \circ \boldsymbol{\eta}_2 \right] \\ &= \delta \mathcal{E}_f \left[(\mathbf{I} - \delta \mathcal{E}_f)^{-1} (\Pi_\theta \circ \boldsymbol{\eta}_1 - \Pi_\theta \circ \boldsymbol{\eta}_2) \right] \end{aligned}$$

Also,

$$\begin{aligned} \mathcal{V}_1 - \mathcal{V}_2 - D_{\boldsymbol{\eta}}^{\mathcal{V}}(\cdot; \theta, f, \boldsymbol{\eta}_1) \\ = \delta \mathcal{E}_f \left[(\mathbf{I} - \delta \mathcal{E}_f)^{-1} \left(\Pi_\theta \circ \boldsymbol{\eta}_1 - \Pi_\theta \circ \boldsymbol{\eta}_2 - \left(\frac{\partial \Pi_\theta}{\partial \boldsymbol{\eta}} \circ \boldsymbol{\eta}_1 \right) (\boldsymbol{\eta}_1 - \boldsymbol{\eta}_2) \right) \right] \end{aligned}$$

Taking a mean value expansion in $\boldsymbol{\eta}$ of $\Pi_\theta \circ \boldsymbol{\eta}_2$, along with the assumed Lipschitz condition on $\frac{\partial \Pi}{\partial \boldsymbol{\eta}}$ gives

$$\left\| \mathcal{V}_1 - \mathcal{V}_2 - D_{\boldsymbol{\eta}}^{\mathcal{V}}(\cdot; \theta, f, \boldsymbol{\eta}_1) \right\| \leq \frac{\delta}{1 - \delta} \|f\|_{\mathcal{F}} \|M_1\| \|\boldsymbol{\eta}_1 - \boldsymbol{\eta}_2\|^2$$

Thus, $D_{\boldsymbol{\eta}}^{\mathcal{V}}$ is the pathwise derivative as claimed. A similar argument would also show that (5.9) is the pathwise derivative of $\frac{\partial \mathcal{V}}{\partial i}$ with respect to $\boldsymbol{\eta}$.

We now show (5.8). Recall that $\mathcal{V}(x, i, \eta; \theta, f, \boldsymbol{\eta})$ denotes the value of choosing i when the state is x and the private shock is η and $\hat{V}(x, i; \theta, f, \boldsymbol{\eta})$ denotes the value of being in state x and choosing i when $\eta = \boldsymbol{\eta}(i, x)$, i.e. the value of choosing i when η is such that i is optimal. \mathcal{V} can be written as

$$\mathcal{V}(x, i, \eta; \theta, f, \boldsymbol{\eta}) = \Pi_\theta(x, i, \eta) + \delta \mathcal{E}_f \hat{V}(\cdot; \theta, f, \boldsymbol{\eta}),$$

where \hat{V} can be defined as the solution to the following equation:

$$\hat{V}(i, x; \theta, f, \boldsymbol{\eta}) = \Pi_\theta(x, i, \boldsymbol{\eta}(i, x)) + \delta \mathcal{E}_f \hat{V}(\cdot; \theta, f, \boldsymbol{\eta}).$$

Consider the difference between \mathcal{V} for two values of f .

$$\begin{aligned}\mathcal{V}_1 - \mathcal{V}_2 &= \delta \mathcal{E}_{f_1} \hat{V}_1 - \delta \mathcal{E}_{f_2} \hat{V}_2 \\ &= \delta \int \hat{V}_1(i', x') (f_1(i', x'|\cdot) - f_2(i', x'|\cdot)) di' dx' + \delta \mathcal{E}_{f_2} (\hat{V}_1 - \hat{V}_2)\end{aligned}$$

where now $\mathcal{V}_j = \mathcal{V}(\cdot; \theta, f_j, \boldsymbol{\eta})$ and $\hat{V}_j = \hat{V}(\cdot; \theta, f_j, \boldsymbol{\eta})$. Similarly,

$$\hat{V}_1 - \hat{V}_2 = \delta \int \hat{V}_1(i', x') (f_1(i', x'|\cdot) - f_2(i', x'|\cdot)) di' dx' + \delta \mathcal{E}_{f_2} (\hat{V}_1 - \hat{V}_2)$$

Although this equation is nearly identical to the previous one, here we can subtract $\delta \mathcal{E}_{f_2} (\hat{V}_1 - \hat{V}_2)$ from both sides and apply $(\mathbf{I} - \delta \mathcal{E}_{f_2})^{-1}$ to obtain

$$\hat{V}_1 - \hat{V}_2 = (\mathbf{I} - \delta \mathcal{E}_{f_2})^{-1} \left[\delta \int \hat{V}_1(i', x') (f_1(i', x'|\cdot) - f_2(i', x'|\cdot)) di' dx' \right].$$

Now note that,

$$\begin{aligned}\mathcal{V}_1 - \mathcal{V}_2 - D_f^{\mathcal{V}}(\cdot; \theta, f_1, \boldsymbol{\eta}) &= \\ &= \mathcal{E}_{f_2} \left((\mathbf{I} - \delta \mathcal{E}_{f_2})^{-1} \left[\delta \int \hat{V}_1(i', x') (f_1(i', x'|\cdot) - f_2(i', x'|\cdot)) di' dx' \right] \right) - \\ &\quad - \mathcal{E}_{f_1} \left((\mathbf{I} - \delta \mathcal{E}_{f_1})^{-1} \left[\delta \int \hat{V}_1(i', x') (f_1(i', x'|\cdot) - f_2(i', x'|\cdot)) di' dx' \right] \right) \\ &= \mathcal{E}_{f_2} \left([(\mathbf{I} - \delta \mathcal{E}_{f_2})^{-1} - (\mathbf{I} - \delta \mathcal{E}_{f_1})^{-1}] \left[\delta \int \hat{V}_1(i', x') (f_1(i', x'|\cdot) - f_2(i', x'|\cdot)) di' dx' \right] \right) + \\ &\quad + (\mathcal{E}_{f_2} - \mathcal{E}_{f_1}) \left((\mathbf{I} - \delta \mathcal{E}_{f_1})^{-1} \left[\delta \int \hat{V}_1(i', x') (f_1(i', x'|\cdot) - f_2(i', x'|\cdot)) di' dx' \right] \right).\end{aligned}$$

It is trivial that

$$\begin{aligned}& \left\| (\mathcal{E}_{f_2} - \mathcal{E}_{f_1}) \right. \\ & \left. \left((\mathbf{I} - \delta \mathcal{E}_{f_1})^{-1} \left[\delta \int \hat{V}_1(i', x') (f_1(i', x'|\cdot) - f_2(i', x'|\cdot)) di' dx' \right] \right) \right\| \lesssim \\ & \lesssim \|f_2 - f_1\|_{\mathcal{F}}^2.\end{aligned}$$

Also, for any function $g(i, x)$,

$$\|(\mathbf{I} - \delta \mathcal{E}_{f_2})^{-1} g - (\mathbf{I} - \delta \mathcal{E}_{f_1})^{-1} g\| \leq \frac{\delta}{1 - \delta} \|g\| \|f_1 - f_2\|_{\mathcal{F}}$$

Thus, we can conclude that

$$\lim_{f_2 \rightarrow f_1} \frac{\|\mathcal{V}_1 - \mathcal{V}_2 - D_f^{\mathcal{V}}(\cdot; \theta, f_1, \boldsymbol{\eta})(f_1 - f_2)\|}{\|f_1 - f_2\|_{\mathcal{F}}} = 0$$

and $D_f^{\mathcal{V}}$ is the derivative as claimed. It is elementary to show that $D_f^{\mathcal{V}}$ has the derivative with respect to i claimed in (5.10). Showing that $\frac{\partial \mathcal{V}}{\partial i}$ has (5.10) pathwise derivative involves the same argument just used to show (5.8), so we omit the details. Also, showing that the pathwise derivatives are Lipschitz can be done in similar manner, so we omit the details. \square

We now give conditions on the payoff function that are sufficient for the conditions theorem 3. Let \mathcal{F} be the Sobolev like space of functions $f : (X \times \mathcal{I}) \times (X \times \mathcal{I})$ that are differentiable with respect to their last argument. Define a norm on this space by

$$\|f\|_{\mathcal{F}} = \sup_{\|Bg\|=1} \left\| B \int g(x', i') f(x', i' \cdot) dx' di' \right\| + \sup_{\|Bg\|=1} \left\| B \int g(x', i') \frac{\partial f}{\partial i}(x', i' \cdot) dx' di' \right\|.$$

Theorem 4 (Asymptotic distribution of the payoff function parameters). *If assumptions A1, A2, and the conditions of lemma 1 hold and,*

A'1 The estimates of $\boldsymbol{\eta}$, f , and $\frac{\partial f}{\partial i}$ converge faster than $n^{-1/4}$, i.e.

$$\begin{aligned} \|B(\hat{\boldsymbol{\eta}} - \boldsymbol{\eta}_0)\| &= o_p(n^{-1/4}) \\ \|\hat{f} - f_0\|_{\mathcal{F}} &= o_p(n^{-1/4}) \end{aligned}$$

A'2 $\frac{\partial \Pi}{\partial \theta}$ and $\frac{\partial^2 \Pi}{\partial i \partial \theta}$ are Lipschitz continuous with respect to θ in a neighborhood of θ_0 in that

$$\left\| B \left(\frac{\partial \Pi}{\partial \theta}(\cdot; \theta_1) - \frac{\partial \Pi}{\partial \theta}(\cdot; \theta_2) \right) \right\| \lesssim |\theta_1 - \theta_2|$$

and

$$\left\| B \left(\frac{\partial^2 \Pi}{\partial i \partial \theta}(\cdot; \theta_1) - \frac{\partial^2 \Pi}{\partial i \partial \theta}(\cdot; \theta_2) \right) \right\| \lesssim |\theta_1 - \theta_2|$$

A'3 $\left| \frac{\partial^2 \mathcal{V}}{\partial i \partial \theta}(x, i) \right| \leq C(x, i)$ with $(BC)(x, i)$ and $\|BC\|$ finite.

A'4 \hat{f} and $\hat{\boldsymbol{\eta}}$ are bounded with probability approaching one

$$\|B_n \hat{\boldsymbol{\eta}}\| \leq M$$

and

$$\|\hat{f}\|_{\mathcal{F}} \leq M.$$

A'5

$$\sqrt{n} \begin{bmatrix} \left\langle B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), BD_{\boldsymbol{\eta}}^m(\theta_0, h_0)(\hat{\boldsymbol{\eta}} - \boldsymbol{\eta}_0) \right\rangle \\ \left\langle B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), BD_f^m(\theta_0, h_0)(\hat{f} - f_0) \right\rangle \end{bmatrix} \xrightarrow{d} N(0, \Omega_B)$$

Proof. We will verify the conditions of theorem 3. We have assumed A1 and A2. Also, condition A'1 is basically a restatement of A3. To verify A4, first note that the derivatives of m_a , m_v , and m_f are simply functions of the derivatives of \mathcal{V} . For m_f , it is immediate that

$$\begin{aligned}\frac{\partial m_f}{\partial \theta} &= \frac{\partial^2 \mathcal{V}}{\partial i \partial \theta} \\ D_{\boldsymbol{\eta}}^{m_f}(\boldsymbol{\eta}_1 - \boldsymbol{\eta}) &= \left(\frac{\partial \Pi}{\partial \boldsymbol{\eta}} \circ \boldsymbol{\eta} + \frac{\partial D_{\boldsymbol{\eta}}^{\mathcal{V}}}{\partial i} \right) (\boldsymbol{\eta}_1 - \boldsymbol{\eta}) \\ D_f^{m_f}(f_1 - f) &= \frac{\partial D_f^{\mathcal{V}}}{\partial i} (f_1 - f).\end{aligned}$$

For m_a , a simple application of the implicit function theorem gives,

$$\begin{aligned}\frac{\partial m_a}{\partial \theta} &= - \left(\frac{\partial^2 \mathcal{V}}{\partial i^2} \right)^{-1} \frac{\partial^2 \mathcal{V}}{\partial i \partial \theta} \\ D_{\boldsymbol{\eta}}^{m_a}(\boldsymbol{\eta}_1 - \boldsymbol{\eta}) &= - \left(\frac{\partial^2 \mathcal{V}}{\partial i^2} \right)^{-1} \left(\frac{\partial \Pi}{\partial \boldsymbol{\eta}} \circ \boldsymbol{\eta} + \frac{\partial D_{\boldsymbol{\eta}}^{\mathcal{V}}}{\partial i} \right) (\boldsymbol{\eta}_1 - \boldsymbol{\eta}) \\ D_f^{m_a}(f_1 - f) &= - \left(\frac{\partial^2 \mathcal{V}}{\partial i^2} \right)^{-1} \frac{\partial D_f^{\mathcal{V}}}{\partial i} (f_1 - f).\end{aligned}$$

For m_v , using the envelope and mean value theorems gives,

$$\begin{aligned}\frac{\partial m_a}{\partial \theta} &= \frac{\partial^2 \mathcal{V}}{\partial i \partial \theta} (i - i^*) \\ D_{\boldsymbol{\eta}}^{m_a}(\boldsymbol{\eta}_1 - \boldsymbol{\eta}) &= \left(\frac{\partial \Pi}{\partial \boldsymbol{\eta}} \circ \boldsymbol{\eta} + \frac{\partial D_{\boldsymbol{\eta}}^{\mathcal{V}}}{\partial i} \right) (\boldsymbol{\eta}_1 - \boldsymbol{\eta}) (i - i^*) \\ D_f^{m_a}(f_1 - f) &= \frac{\partial D_f^{\mathcal{V}}}{\partial i} (f_1 - f) (i - i^*),\end{aligned}$$

where $i^*(i, x; \theta, f, \boldsymbol{\eta}) = \arg \max_{\tilde{i}} \mathcal{V}(x, \boldsymbol{\eta}(i, x), \tilde{i}, i; \theta, f, \boldsymbol{\eta})$. Conditions A4.i and A4.ii then follow from the Lipschitz continuity of $\frac{\partial D_{\boldsymbol{\eta}}^{\mathcal{V}}}{\partial i}$ and $\frac{\partial D_f^{\mathcal{V}}}{\partial i}$ shown in lemma 1. Similarly, condition A4.iii follows directly from lemma 1 and assumption A'2.

A4.iv can be verified through repeated application of the dominated convergence theorem. In particular if $|\frac{\partial m}{\partial \theta}(x, i; \theta)| \leq C(x, i)$ with $BC(x, i) < \infty$ and $|BC| < \infty$, then

$$\begin{aligned}\frac{\partial}{\partial \theta} Bm &= \frac{\partial}{\partial \theta} \int m(x', i') b(x', i'; x, i) dx' di' \\ &= \int \frac{\partial m}{\partial \theta}(x', i') b(x', i'; x, i) dx' di' .\end{aligned}$$

Also,

$$\begin{aligned} \int \frac{\partial}{\partial \theta} (Bm)(x, i)^2 dx di &= \int 2(B \frac{\partial m}{\partial \theta})(x, i)(Bm)(x, i) dx di \\ &\leq \|BC\| \|Bm\| \end{aligned}$$

so

$$\frac{\partial}{\partial \theta} \|Bm\|^2 = 2 \left\langle B \frac{\partial m}{\partial \theta}, Bm \right\rangle.$$

Thus, we just have to show for each of the three possibilities of m that $|\frac{\partial m}{\partial \theta}(x, i; \theta)| \leq C(x, i)$ with $BC(x, i) < \infty$ and $|BC| < \infty$.

Condition A5 follows from A'4 and the following inequalities. To show the first part of A5 note that

$$\begin{aligned} \|B_n m(\cdot; \theta, \hat{h})\| &\lesssim \left\| B_n \frac{\partial \Pi}{\partial i} \circ \hat{\eta} \right\| + \left\| B_n \delta \frac{\partial \mathcal{E}_f}{\partial i} \left(\mathbf{I} - \delta \mathcal{E}_f \right)^{-1} \Pi \circ \hat{\eta} \right\| \\ &\lesssim \|B_n(\hat{\eta} - \eta_0)\| + \frac{\delta}{1 - \delta} \left(\|B_n\|_{op} \|\hat{f} - f\|_{\mathcal{F}} + \|B_n(\hat{\eta} - \eta_0)\| \right). \end{aligned}$$

Similarly, the second part of A5 follows from

$$\begin{aligned} \|B_n D_h^m(\hat{h} - h)\| &\lesssim \left\| B_n \frac{\partial D_{\eta}^{\mathcal{V}}}{\partial i}(\hat{\eta} - \eta_0) \right\| + \left\| B_n \frac{\partial D_f^{\mathcal{V}}}{\partial i}(\hat{f} - f_0) \right\| \\ &\lesssim \frac{\delta}{1 - \delta} \left(\|B_n(\hat{\eta} - \eta_0)\| + \|\hat{f} - f_0\|_{\mathcal{F}} \right). \end{aligned}$$

Finally, A'5 implies A6. □

5.1. Asymptotic equivalence. In this subsection we show that the three types of moment conditions result in asymptotically equivalent estimators when appropriate weighting operators are used. This fact follows directly from results derived in the proofs of Theorems 3 and 4.

Corollary 1 (Asymptotic equivalence). *Suppose the conditions of Theorem 4 hold. Let $\hat{\theta}_f(B)$ be the estimate from using moment conditions m_f and weighting operator B . Similarly, define $\hat{\theta}_a(B)$ and $\hat{\theta}_v(B)$. Then*

$$\left\| \hat{\theta}_f(B) - \hat{\theta}_a(B \frac{\partial^2 \mathcal{V}}{\partial i}) \right\| = o_p(n^{-1/2}) \left\| \hat{\theta}_f(B) - \hat{\theta}_v(B(i^* - i)^{-1}) \right\| = o_p(n^{-1/2}).$$

Proof. The proof of Theorem 3 shows that

$$\begin{aligned} \sqrt{n}(\hat{\theta} - \theta_0) &= -\sqrt{n} \left\langle B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0) \right\rangle^{-1} \times \\ &\quad \times \left\langle B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), BD_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle + o_p(1). \end{aligned}$$

Substituting in the formulas for D_h^m and $\frac{\partial m}{\partial \theta}$ given in the proof of theorem 4 verifies the proposition. \square

5.2. A kernel estimator . This section describes kernel estimates of the transition density and policy function that satisfy the conditions for asymptotic normality given in theorem 4. We assume that our data comes from a balanced panel of length T . There are M independent markets, each with F firms. We will state results for M and F fixed with $T \rightarrow \infty$, but our proofs could easily be adapted to situations with $M \rightarrow \infty$ and T either growing or fixed, provided that at least $MT \rightarrow \infty$. We will let $n = MTF$. When the distinction among M , T , and F is unimportant, we will let

$$\sum_{j=1}^n x_j = \sum_{t=1}^T \sum_{m=1}^M \sum_{f=1}^F x_{t,m,f}.$$

Our estimate of the policy function is

$$\hat{\eta}(x, i) = F_\eta^{-1} \left(\frac{\sum_{j=1}^n \mathbf{1}(i_j \leq i) k((x_j - x)/h_n)}{\sum_{j=1}^n k((x_j - x)/h_n)} \right),$$

where k is a kernel satisfying the conditions states below and h_n is the bandwidth. We will let

$$F(\widehat{i|x})f(x) = \frac{1}{h_n n} \sum_{j=1}^n \mathbf{1}(i_j \leq i) k((x_j - x)/h_n)$$

and

$$\hat{f}(x) = \frac{1}{h_n n} \sum_{j=1}^n k((x_j - x)/h_n).$$

Let $z = (x, i)$. The transition density is estimated by

$$\hat{f}(z'|z) = \frac{\sum_{m=1}^M \sum_{f=1}^F \sum_{t=1}^{T-1} k((z_{t+1,m,f} - z', z_{t,m,f} - i)/h_n)}{\sum_{m=1}^M \sum_{f=1}^F \sum_{t=1}^{T-1} k((z_{t,m,f} - z)/h_n)},$$

and the estimate of its derivative is simply the derivative of the estimate of the transition density,

$$\frac{\partial \hat{f}}{\partial i}(z'|z) = \frac{\partial}{\partial i} \frac{\sum_{m=1}^M \sum_{f=1}^F \sum_{t=1}^{T-1} k((z_{t+1,m,f} - z', z_{t,m,f} - z)/h_n)}{\sum_{m=1}^M \sum_{f=1}^F \sum_{t=1}^{T-1} k((z_{t,m,f} - z)/h_n)}.$$

We will denote the joint density of z_{t+1} and z_t by f_{zz} , the marginal density of z by f_z , and the marginal density of x by f_x . We denote kernel estimates of these densities by \hat{f}_{zz} , \hat{f}_z , and \hat{f}_x .

To apply Theorem 4 to conclude that $\hat{\theta}$ is \sqrt{n} asymptotically normal; we must verify that the kernel estimates above satisfy the rate condition A'1 and the asymptotic normality condition A'5. The rate condition can be verified by using available results on uniform convergence rates for kernel estimators. We employ the results of Hansen (2008) for this purpose. Showing condition A'5 is more involved. In independent work, Srisuma and Linton (2010) propose a similar kernel estimator for the case with continuous states and discrete actions, and Srisuma (2010) develops a similar estimator for models with discrete states and continuous controls. These papers employ U-statistics to show \sqrt{n} asymptotic normality of the parameters of the payoff function. We take a different approach. We extend the uniform central limit theorem for smoothed empirical processes of van der Vaart (1994) and Giné and Nickl (2007) to allow for dependent data. We then utilize this theorem to verify condition A'5.

Condition A'5 is that

$$\sqrt{n} \left\langle B^* B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle \xrightarrow{d} N(0, \Omega_B).$$

To verify this condition, we first observe that this inner-product has the following integral form:

$$\begin{aligned} & \left\langle B^* B \frac{\partial m}{\partial \theta}(\cdot, \theta_0, h_0), D_h^m(\theta_0, h_0)(\hat{h} - h_0) \right\rangle = \\ & = \int c_1(z', z) \left(\hat{f}_{z',z}(z', z) - f_{z',z}(z', z) \right) dz' dz + \\ & + \int c_2(z', z) \left(\hat{f}_z(z) - f_z(z) \right) dz' dz + \\ & + \int c_3(z', z) \left(\frac{\partial \hat{f}_{z',z}(z', z)}{\partial i} - \frac{\partial f_{z',z}(z', z)}{\partial i} \right) dz' dz + \\ & + \int c_4(z', z) \left(\frac{\partial \hat{f}_z(z)}{\partial i} - \frac{\partial f_z(z)}{\partial i} \right) dz' dz + \\ & + \int c_5(z', z) \left(\hat{f}_x(x) - f_x(x) \right) dz' dz + \\ & + \int c_6(z', z) \left(F(\widehat{i|x})f(x) - F(i|x)f(x) \right) dz' dz + o_p(n^{-1/2}), \end{aligned} \tag{5.11}$$

where $c_j(\cdot)$ are functions that depend on B , $\frac{\partial m}{\partial \theta}$ and D_h^m . The exact form of $c_j(\cdot)$ and formal verification of (5.11) are given in Lemma 2 below. Given (5.11), we can then apply the following theorem, which is an extension of van der Vaart (1994) and Giné and Nickl (2007) to allow for dependent data and estimates of derivatives of densities as well as densities themselves. This result may be of independent interest.

Theorem 5 (Extension of van der Vaart (1994) and Giné and Nickl (2007)). *Let $X = (X_i)_{i \in \mathbb{Z}}$ be a stationary sequence with marginal distribution \mathbb{P} and β -mixing coefficients $\beta(k)$. Let $\mathcal{F} \subset \mathcal{L}^2(\mathbb{P})$ be a translation invariant class of functions. Suppose there exists $2 < p < \infty$ such that $\sum_{k=1}^{\infty} k^{2/(2-p)}\beta(k) < \infty$ and \mathcal{F} is not too complex, namely the bracketing entropy integral is finite, $J_{[]}(\infty, \mathcal{F}, \mathcal{L}_p(\mathbb{P})) < \infty$. Let $\|f\|_{LR} = \sum_{i \in \mathbb{Z}} E[f(z_0)f(z_i)]$. Let $\{\mu_n\}_{n=1}^{\infty}$ converge weakly to the Dirac measure at zero, δ_0 , or one of its derivatives, $\partial^\alpha \delta_0$. Also assume that for all n , $\mu_n(\mathbb{R}^d) \leq C$, $\mathcal{F} \subseteq \mathcal{L}^1(|\mu_n|)$, and $\int \|f(\cdot - y)\|_{LR} d|\mu_n|(y) < \infty$ for all $f \in \mathcal{F}$. If*

$$\sup_{f \in \mathcal{F}} \sum_{i \in \mathbb{Z}} E \left[\left(\int f(X_0 + y) - f(X_0) d\mu_n(y) \right) \left(\int f(X_i + y) - f(X_i) d\mu_n(y) \right) \right] \rightarrow 0 \quad (5.12)$$

and

$$\sup_{f \in \mathcal{F}} \left| E \int f(X + y) - f(X) d\mu_n(y) \right| \rightarrow 0 \quad (5.13)$$

then

$$\sqrt{n} (\mathbb{P}_n * \mu_n - \mathbb{P} * \mu_\infty) \rightsquigarrow \mathcal{G} \text{ in } \ell^\infty(\mathcal{F}),$$

where \mathcal{G} is the \mathbb{P} -Brownian bridge indexed by \mathcal{F} and μ_∞ is either δ_0 or $\partial^\alpha \delta_0$.

Proof. This is Theorem 2 of Giné and Nickl (2007), except it allows for dependent data and $\mu_n \rightarrow \partial^\alpha \delta_0$ in addition to $\mu_n \rightarrow \delta_0$. Allowing for $\partial^\alpha \delta_0$ requires no modification of the proof. Allowing for dependent data simply requires substituting the empirical process results for independent data used by van der Vaart (1994) and Giné and Nickl (2007) with a suitable empirical process result for dependent data. A variety of such results are available; we use Theorem 11.22 of Kosorok (2008). This theorem says that given the condition on β -mixing coefficients and bracketing entropy above

$$\sqrt{n} (\mathbb{P}_n - \mathbb{P}) \rightsquigarrow \xi \text{ in } \ell^\infty(\mathcal{F}),$$

where ξ is a tight, mean zero Gaussian process with covariance

$$V(f, g) = \sum_{i \in \mathbb{Z}} \text{cov}(f(X_0)g(X_i)).$$

Then, by following the arguments used to prove Theorem 1 and Lemmas 1 and 2 of Giné and Nickl (2007), but with $\|\cdot\|_{LR}$ in place of $\|\cdot\|_{2, \mathbb{P}}$, we can show that

$$\mathbb{G}_n(f, \mu) = \sqrt{n}(\mathbb{P}_n - \mathbb{P}) \int f(X + y)d\mu(y)$$

converges in distribution in $\ell^\infty(\mathcal{F} \times \mathcal{M})$. Finally, the proof of van der Vaart (1994), with his semi-metric d redefined as

$$d^2((f, \mu), (g, \nu)) = \sum_{i \in \mathbb{Z}} E \left[\left(\int f(X_0 + y)d\mu(y) - \int g(X_0 + y)d\nu(y) \right) \times \right. \\ \left. \times \left(\int f(X_i + y)d\mu(y) - \int g(X_i + y)d\nu(y) \right) \right],$$

leads to the conclusion. \square

We now state regularity conditions that are sufficient for (5.11) to hold and Theorem 5 to apply.

Assumptions K (Kernel estimator).

K1 The densities $f(z', z)$ and $f(z)$ are a ≥ 2 -times continuously differentiable.

K2 The data is comes from M markets with F firms and $T \rightarrow \infty$ periods. $\{z_{t,m,f}\}_{t \in \mathbb{Z}}$ is stationary and strongly mixing with strong mixing coefficients α_m that decay at rate β , i.e. $\alpha_m \lesssim m^{-\beta}$ for some $m > 1$. There exists $q > 0$ such that $\beta > 1 + \frac{d}{q} + d$ and

$$\frac{\beta - 1 - \frac{d}{q} - d}{d(\beta + 3 - d)} > \frac{1}{4r}.$$

K3 With probability approaching one,

$$\int \left| \left[B_n^* B_n D_h^m (\hat{h}(z', z) - h(z', z)) \right] (z_0) \right| dz' dz < C$$

uniformly in z_0 .

K4 The dimension of $z = (x, i)$ is d .

K5 $c_j(z', z)$ is $s > d/2$ times differentiable with uniformly bounded derivatives.

K6 The kernel is of order $r = a + s - k > a + d/2$.

K7 The bandwidth is such that $h_n^{a+s-k} n^{d/2} \rightarrow \infty$

K8 The weighting operator is such that $\int b_n(z_0, z_1) b_n(z, z_1) dz_1$ has compact convex support and vanishes at the boundary of its support.

Lemma 2 shows that these conditions are sufficient for (5.11).

Lemma 2 (Verification of (5.11)). *If K1-K7 hold, then the following functions exist*

$$\begin{aligned} d_1(z', z; z_0) &= \delta \left[\frac{\partial \mathcal{E}_f}{\partial i} \left((\mathbf{I} - \delta \mathcal{E}_f)^{-1} \right) \right] (z|z_0) \left[\delta (\mathbf{I} - \delta \mathcal{E}_f)^{-1} \Pi_{\theta_0} \circ \boldsymbol{\eta}_0 \right] (z') \\ d_2(z', z; z_0) &= \left[\delta (\mathbf{I} - \delta \mathcal{E}_f)^{-1} \Pi_{\theta_0} \circ \boldsymbol{\eta}_0 \right] (z') \\ d_\eta(z', z; z_0) &= \frac{\partial^2 \Pi_{\theta_0}}{\partial \eta \partial i} \circ \boldsymbol{\eta}_0(z_0) + \delta \left[\mathcal{E}_f \left((\mathbf{I} - \delta \mathcal{E}_f)^{-1} \left(\frac{\partial \Pi_{\theta_0}}{\partial \eta} \circ \boldsymbol{\eta}_0 \right) \right) \right] (z), \end{aligned}$$

and equation (5.11) holds with

$$\begin{aligned} c_j(z', z) &= \left\langle \frac{\partial m}{\partial \theta}, (B^* B \tilde{c}_j(z', \cdot))(z, \cdot) \right\rangle \\ &= \int \frac{\partial m}{\partial \theta}(z_0) \left(\int b(z_0, z_1) b(z, z_1) dz_1 \right) \tilde{c}(z', z; z_0) dz_0 \end{aligned}$$

where

$$\begin{aligned} \tilde{c}_1(z', z; z_0) &= d_1(z', z; z_0) \frac{1}{f_z(z)} - d_2(z', z; z_0) \frac{\partial f_{zz} / \partial i(z', z)}{f_z(z)^2} \\ \tilde{c}_2(z', z; z_0) &= -d_1(z', z; z_0) \frac{f_{zz}(z', z)}{f_z(z)^2} \\ &\quad - d_2(z', z; z_0) \left(\frac{\partial f_{zz} / \partial i(z', z)}{f_{zz}(z', z)} + 2 \frac{f_{zz}(z', z) \partial f_z / \partial i(z)}{f_z(z)^2} \right) \\ \tilde{c}_3(z', z; z_0) &= d_2(z', z; z_0) \frac{1}{f_z(z)} \\ \tilde{c}_4(z', z; z_0) &= d_2(z', z; z_0) \frac{f_{zz}(z', z)}{f_z(z)} \\ \tilde{c}_5(z', z; z_0) &= d_\eta(z', z; z_0) \frac{1}{f_x(x)} \\ \tilde{c}_6(z', z; z_0) &= -d_\eta(z', z; z_0) \end{aligned}$$

Proof. Equation (5.11) is derived by repeatedly applying the mean value theorem to write $\hat{h} - h$ as a sum of $\hat{f}_{zz} - f_{zz}$, $\hat{f}_z - f_z$, $\hat{f}_x - f_x$, and their derivatives. Condition K3 ensures that we can freely interchange the order of integration while doing so. The definitions of d_1 , d_2 , and d_η are then verified by straightforward calculation. Condition K7 and K2, along with the uniform convergence rate results of Hansen (2008), are sufficient to show that the remainder in (5.11) is $o_p(n^{-1/2})$. \square

Combining this lemma with Theorems 5 and 4 shows that the kernel estimator of $\hat{\theta}$ is asymptotically normal. Theorem 6 states this result.

Theorem 6 (Asymptotic distribution for kernel estimator). *If conditions A'2-A'4 and K1-K7 hold, then*

$$\sqrt{n}(\hat{\theta} - \theta_0) \xrightarrow{d} N(0, M_B^{-1} \Omega_B M_B^{-1}),$$

where

$$\Omega_b = \sum_{t=-\infty}^{\infty} \sum_{f_2, f_1=1}^F \left[\begin{array}{l} \text{cov}(c_1(z_{0+1, m_1, f}, z_{0, m, f_1}), c_1(z_{t+1}, m_2, f, z_{t, m, f_2})) + \\ \quad + \text{cov}\left(\int c_2(z', z_{0, m, f_1}) dz', \int c_2(z', z_{t, m, f_2}) dz'\right) + \\ \text{cov}\left(\frac{\partial c_3}{\partial i}(z_{0+1, m_1, f}, z_{0, m, f_1}), \frac{\partial c_3}{\partial i}(z_{t+1}, m_2, f, z_{t, m, f_2})\right) + \\ \quad + \text{cov}\left(\int \frac{\partial c_4}{\partial i}(z', z_{0, m, f_1}) dz', \int \frac{\partial c_4}{\partial i}(z', m_2, f, z_{t, m, f_2}) dz'\right) + \\ \quad + \text{cov}\left(\int c_5(z', (x_{0, m, f_1}, i)) dz' di, \int c_5(z', (x_{t, m, f_2}, i)) dz' di\right) + \\ \text{cov}\left(\int \frac{\partial c_6}{\partial i}(z', z_{0, m, f_1}) dz', \int \frac{\partial c_6}{\partial i}(z', m_2, f, z_{t, m, f_2}) dz'\right) \end{array} \right].$$

Proof. As in the proof of Lemm 2, conditions K1 and K7, along with the uniform convergence rate results of Hansen (2008), are sufficient to show that the rate assumption, A'1, holds. We have assumed A'2-A'4. All that remains is to show A'5 holds with the stated Ω_B . By Lemma 2, equation (5.11) holds. Applying Theorem 5 to (5.11) yields the desired conclusion. All that remains is to show that the conditions of Theorem 5 are satisfied.

Let $\mathcal{F} = \{c_j(z', z)\}_{j=1}^6$. This is finite class and trivially satisfies $J_{\square}(\infty, \mathcal{F}, \mathcal{L}_p(\mathbb{P})) < \infty$. Conditions K1 and K5 are sufficient to apply Proposition 1 of Giné and Nickl (2007) to show that (5.12) holds for c_1 - c_6 . Conditions K5-K7 are sufficient to apply Theorem 6 of Giné and Nickl (2007) to show that (5.13) holds for c_1 , c_2 , c_5 , and c_6 . These same conditions along with K8 and a slight modification of Theorem 6 of Giné and Nickl (2007) shows (5.13) for c_3 and c_4 . \square

6. CONCLUSION

We have shown that the payoff function in dynamic games with continous states and controls is nonparametrically identified by the observed distribution of states and controls. The key identifying restrictions include that one of the partial derivatives of the payoff function is known, that there is some component of the state space that enters the policy function, but not the payoff function directly, and that a certain integral operator of the second kind is invertible. We have also developed a semi-parametric estimator for the model. In the first step the transition densities and policy function are estimated nonparametrically. In the second step, the parameters

of the payoff function are estimated from the optimality conditions of the model. we gave high-level conditions on the first step nonparametric estimates for the parameter estimates to be consistent and parameters to be \sqrt{n} asymptotically normal, and we have shown that a kernel based estimator satisfies these conditions.

REFERENCES

- AI, C., AND X. CHEN (2003): “Efficient Estimation of Models with Conditional Moment Restrictions Containing Unknown Functions,” *Econometrica*, 71(6), 1795–1843.
- BAJARI, P., V. CHERNOZHUKOV, H. HONG, AND D. NEKIPELOV (2009): “Nonparametric and Semiparametric Analysis of a Dynamic Discrete Game,” Discussion paper.
- BERRY, S., AND A. PAKES (2001): “Estimation from the Optimality Conditions for Dynamic Controls,” .
- BERRY, S., AND E. TAMER (2006): “Identification in Models of Oligopoly Entry,” in *Advances in economics and econometrics: theory and applications, ninth world congress*, ed. by R. Blundell, W. K. Newey, and T. Persson. Cambridge University Press.
- BLEVINS, J. (2009): “Nonparametric Identification of Dynamic Games with Discrete and Continuous Choices,” Discussion paper, Duke University.
- CARRASCO, M., M. CHERNOV, J.-P. FLORENS, AND E. GHYSELS (2007): “Efficient estimation of general dynamic models with a continuum of moment conditions,” *Journal of Econometrics*, 140(2), 529 – 573.
- CARRASCO, M., AND J. FLORENS (2000): “Generalization of Gmm to a Continuum of Moment Conditions,” *Econometric Theory*, 16(06), 797–834.
- GINÉ, E., AND R. NICKL (2007): “Uniform central limit theorems for kernel density estimators,” *Probability Theory and Related Fields*, 141(3-4), 333–387.
- HANSEN, B. E. (2008): “Uniform Convergence Rates for Kernel Estimation with Dependent Data,” *Econometric Theory*, 24(03), 726–748.
- HONG, H., AND M. SHUM (2009): “Pairwise-Difference Estimation of a Dynamic Optimization Model,” *Review of Economic Studies*, 77(1), 273–304.
- HU, Y., AND M. SHUM (2008): “Identifying Dynamic Games with Serially-Correlated Unobservables,” Discussion paper, The Johns Hopkins University, Department of Economics.
- JENKINS, M., P. LIU, R. L. MATZKIN, AND D. L. MCFADDEN (2004): “The Browser War — Econometric Analysis of Markov Perfect Equilibrium in Markets with Network Effects,” Discussion paper.
- KOSOROK, M. R. (2008): *Introduction to empirical processes and semiparametric inference*. Springer Science.
- MACIEIRA, J. (2009): “A Dynamic Model of Innovation and Technological Competetion in the Supercomputer Industry,” Discussion paper.
- MAGNAC, T., AND D. THESMAR (2002): “Identifying Dynamic Discrete Decision Processes,” *Econometrica*, 70(2), 801–816.
- RUST, J. (1994): “Chapter 51 Structural estimation of markov decision processes,” vol. 4 of *Handbook of Econometrics*, pp. 3081–3143. Elsevier.
- SRISUMA, S. (2010): “Minimum distance estimation for a class of Markov decision processes,” Discussion paper, London School of Economics and Political Science.

SRISUMA, S., AND O. B. LINTON (2010): “Semiparametric Estimation of Markov Decision Processes with Continuous State Space,” *SSRN eLibrary*.

VAN DER VAART, A. (1994): “Weak Convergence of Smoothed Empirical Processes,” *Scandinavian Journal of Statistics*, 21(4), 501–504.

E-mail address: `schrumpf@mail.ubc.ca`