Estimation of Structural Optimization Models: A Note on Identification[*]

Sorawoot Srisuma[†]
London School of Economics and Political Science

[†] Department of Economics, London School of Economics, Houghton Street, London, WC2A 2AE, United Kingdom. E-mail address: s.t.srisuma@lse.ac.uk

Abstract

Bajari, Benkard and Levin (2007) propose an estimation methodology for a broad class of dynamic optimization problems. To carry out their procedure, one needs to select a set of alternative policy functions and compare the implied expected payoffs with that from the data. We show that this can generally lead to objective functions that are not capable of consistently estimating an identified model.

Keywords: Consistency, Identification, Optimization Models.

Bajari, Benkard and Levin (2007), henceforth BBL, propose a methodology to estimate a broad class of structural dynamic models. The motivation behind the construction of their estimator is conceptually appealing as it relies directly on necessary conditions of an economic equilibrium. They also propose a set estimator to estimate partially identified models.

The idea behind BBL's estimator relies on the notion of an economic equilibrium so that the optimal policy will lead to higher values in expectation than those based on alternative policies. In practice, only a strict subclass of alternative policies can be considered. This may lead to the loss of identification since we do not make use of all the relevant constraints implied by the equilibrium behavior.[1]

We use a simple optimization example, which belongs to the class of models considered in BBL, to show that the moment inequality approach of BBL can lead to criterion functions that are not capable of consistently estimate an identified model. In what follows we use the same notation in BBL wherever possible.

Example: Leaving out the observable state variables, we specify the payoff function to be

$$\pi(a, \nu; \theta) = -a^2 + 2\theta a \nu,$$

where a and ν take values from the support of the control and state variables, A and $\mathcal{V}$, respectively.[2] We assume the researcher observes a sequence of $\{a_t\}_{t=1}^T$, which is the maximizer of the payoff function above when $\theta = \theta_0$, whilst $\{\nu_t\}_{t=1}^T$ are not observed. We are interested in estimating θ_0. Let G be *any* known distribution for ν_t with zero mean. It is easy to see that the optimal policy function $\sigma(\cdot; \theta)$ satisfies

$$\sigma(\nu; \theta) = \theta\nu \quad \text{for all } \theta \in \Theta, \nu \in \mathcal{V}.$$

Let Θ be any subset of $\mathbb{R}^+$, this ensures that the policy function will be increasing in the state variable, satisfying the *monotone choice* assumption which is essential to BBL's simulation method. Notice that if $\theta \neq \theta'$ then $\sigma(\nu_t; \theta) \neq \sigma(\nu_t; \theta')$ *a.s.*, therefore this parametric model is identified since any $\theta \in \Theta$ implies a unique corresponding policy function $\sigma(\cdot; \theta)$. Then given the data, along with some standard regularity conditions, one can construct a consistent estimator for θ_0 by maximum likelihood or other minimum distance methods.

BBL Methodology: Let $x \in \mathcal{X}$ denote a particular inequality, see BBL (p. 1347). In a single agent problem without other state variables, x simply corresponds to an alternative policy $\sigma' \in \Sigma$,

---

[1]Since BBL assumes the model is defined through a set of conditional moment restrictions, identification here means that there is a unique parameter value in the parameter space that satisfies all the moment restrictions; hence the loss of identification here refers to the situation when an implied submodel loses this uniqueness property.

[2]Including the observable state variables, s_t, will not change our general conclusion even if we allow s_t and ν_t to be correlated.

where Σ is a set of functions mapping $\mathcal{V}$ to A. For any $\theta \in \Theta$, the difference in the implied expected payoffs from using the true policy relative to an alternative policy σ' is defined by

$$g(x; \theta) = \mathbb{E}_\nu \left[\pi \left(\sigma \left(\nu_t; \theta_0 \right), \nu_t, \theta \right) \right] - \mathbb{E}_\nu \left[\pi \left(\sigma' \left(\nu_t \right), \nu_t, \theta \right) \right].$$

The set of optimality conditions from the equilibrium, represented by inequalities, leads to the following type of criterion functions

$$Q(\theta) = \int_{\mathcal{X}} \left(\min \left\{ g(x; \theta), 0 \right\} \right)^2 dH(x),$$

where H is some distribution on $\mathcal{X}$. By definition of an optimal policy, $Q(\theta_0)$ must be zero since

$$\pi \left(\sigma \left(\nu_t; \theta_0 \right), \nu_t, \theta_0 \right) - \pi \left(\sigma' \left(\nu_t \right), \nu_t, \theta_0 \right) \geq 0 \quad a.s.,$$

for any function $\sigma' \in \Sigma$, and this must also hold in expectation, i.e. $g(x; \theta_0) \geq 0$ for all $x \in \mathcal{X}$.[3] BBL's minimum distance estimator is defined to minimize the sample analogue of $Q(\theta)$. However, it is generally infeasible to construct the empirical counterpart of $Q(\theta)$ to incoperate all the inequality constraints implied by the equilibrium behavior. Since this set is uncountably large, ad hoc exclusions of alternative policies are required for implementation. In what follows, we explicitly use the subscript $\mathcal{E}$ on $(\mathcal{X}, H)$ to denote the subset of all alternative policies and its corresponding underlying measure.

Additive Perturbation (I): Additive perturbation is first suggested in BBL, and it is the most widely used method to construct alternative policies in practice.[4] In fact, we present *two* versions of such sets of inequalities. Intuitively, one simply perturbs the true policy by an additive random shock and compare the implied expected payoffs with that from the true. More formally, for some user-chosen distribution Δ, typically $N(0, 1)$, define

$$\mathcal{X}_{\mathcal{E}_1} = \left\{ \sigma'(\cdot) : \sigma'(\cdot) = \sigma(\cdot; \theta_0) + \epsilon \text{ for } \epsilon \in \textit{Support}(\Delta) \subset \mathbb{R} \right\}.$$

The measure $H_{\mathcal{E}_1}$ denotes the measure that generates the random variable ϵ from Δ, *independent of* ν_t, for convenience we denote this by F_Δ. Therefore we can construct a criterion function indexed by $(\mathcal{X}_{\mathcal{E}_1}, H_{\mathcal{E}_1})$

$$Q_{\mathcal{E}_1}(\theta) = \int_{\mathcal{X}_{\mathcal{E}_1}} \left(\min \left\{ g(x; \theta), 0 \right\} \right)^2 dH_{\mathcal{E}_1}(x),$$

---

[3] If the parametric model is identified, it is easy to provide primitive conditions on $\pi(a_t, \nu_t, \theta)$ and $G(d\nu_t)$ to ensure the implied conditional moment inequality restriction is uniquely satisfy only at $\theta = \theta_0$. As satisfied by our example, it will be sufficient if (i) $\pi(a, \nu, \theta)$ is concave on a convex set A and has a unique maximizer in the interior of A for all (ν, θ); (ii) $\pi(a, \nu; \theta)$ is once continuously differentiable on A for all (ν, θ); (iii) The distribution of ν_t has finite first moment.

[4] The dynamic version of this class of alternatives is informally described in BBL (p. 1348). It is easy to find papers which randomly perturbed their policies in this way, for example see the Monte Carlo experiments in BBL, and the empirical studies in Norton (2009), Ryan (2009), Santos (2009).

The empirical analogue of $Q_{\mathcal{E}_1}(\theta)$ can be constructed in practice by simulation, as outlined in BBL, by taking random samples drawn from $H_{\mathcal{E}_1}$. Then one can construct an estimator based on minimizing $Q_{\mathcal{E}_1}(\theta)$. In our example, it is easy to check that the difference between payoffs obtained from the true and an alternative policy in $\mathcal{X}_{\mathcal{E}_1}$ is

$$\pi\left(\sigma\left(\nu_t;\theta_0\right),\nu_t,\theta\right)-\pi\left(\sigma\left(\nu_t;\theta_0\right)+\epsilon,\nu_t,\theta\right)=\epsilon^2+2\epsilon\nu_t\left(\theta_0-\theta\right). \tag{1}$$

To obtain $g(x;\theta)$ we integrate out ν_t. Since ν_t has zero mean, we have

$$Q_{\mathcal{E}_1}(\theta)=\int\left(\min\left\{\epsilon^2,0\right\}\right)^2 dF_\Delta(\epsilon).$$

Clearly, $Q_{\mathcal{E}_1}(\theta)=0$ for all $\theta\in\Theta$ and any distribution Δ that is not degenerate at zero (which would then not represent alternative policies). That is, in this example, this class of alternative policies has no identifying power for θ_0.

The additive perturbation method is motivated by the fact that we decompose the true policy function into the sum of an alternative policy σ' and a residual function, $\epsilon(\cdot)$, defined on $\mathcal{V}$ so that

$$\sigma\left(\nu_t;\theta_0\right)=\sigma'\left(\nu_t\right)+\epsilon\left(\nu_t\right)\quad a.s.$$

Note that the criterion function constructed from $\mathcal{X}_{\mathcal{E}_1}$ ignores the dependence on ν_t.[5] So $\mathcal{X}_{\mathcal{E}_1}$ is exactly the set of functions that are translation shifts from the true policy. However, it remains unclear how to construct a subset of $\mathcal{X}$ to depend on ν_t appropriately. We now describe the other method of additive perturbation, which is often seen as an adequate improvement on $\mathcal{X}_{\mathcal{E}_1}$ to ensure consistency, that has been implemented in practice. It is motivated by the following condition

$$\text{span}\left\{\sigma\left(\nu\right)+\epsilon:\epsilon\in\mathbb{R}\right\}=\mathbb{R}\text{ for all }\nu\in\mathcal{V}.$$

Additive Perturbation (II): An alternative policy in this class can be constructed from adding a continuum of i.i.d. shocks indexed by $\nu\in\mathcal{V}$, with some distribution Δ that is *independent of* ν_t, to the true policy function pointwise on $\mathcal{V}$.[6] To compare the difference between payoffs, first note that the analogous expression to (1) in this case is

$$\pi\left(\sigma\left(\nu_t;\theta_0\right),\nu_t,\theta\right)-\pi\left(\sigma\left(\nu_t;\theta_0\right)+\epsilon_t,\nu_t,\theta\right)=\epsilon_t^2+2\epsilon_t\nu_t\left(\theta_0-\theta\right),$$

---

[5]Some readers might have been misled by the notation in BBL (p. 1348), and consider this type of perturbation, since ϵ is not explicitly written to depend on the state variables.

[6]More formally,

$$\mathcal{X}_{\mathcal{E}_2}=\left\{\sigma'\left(\cdot\right):\sigma'\left(\cdot\right)=\sigma\left(\cdot;\theta_0\right)+\epsilon\left(\cdot\right)\text{ for }\epsilon\in\mathfrak{E}\left(\mathcal{V};\Delta\right)\right\},$$

where $\mathfrak{E}\left(\mathcal{V};\Delta\right)$ is a set of functions mapping $\mathcal{V}$ to $Support\left(\Delta\right)$. Here $H_{\mathcal{E}_2}$ is the underlying measure that defines a stochastic process $\left(R_\nu\left(\Delta\right)\right)_{\nu\in\mathcal{V}}$, which is independent of ν_t, such that $R_\nu\left(\Delta\right)$ is independent of $R_{\nu'}\left(\Delta\right)$ whenever $\nu\neq\nu'$, and $R_\nu\left(\Delta\right)$ has distribution Δ for all $\nu\in\mathcal{V}$.

where ϵ_t has distribution Δ and is independent of ν_t. In expectation, which is often obtained by simulation, for a generic $x \in \mathcal{X}_{\mathcal{E}_2}$ we have

$$g(x; \theta) = \mathbb{E}\left[\epsilon_t^2\right] + 2\mathbb{E}\left[\epsilon_t \nu_t\right](\theta_0 - \theta).$$

Assume that the second moment of ϵ_t exist. Since ϵ_t and ν_t are independent and ν_t has zero mean, once again we have $g(x; \theta) \geq 0$ for all $\theta \in \Theta$. So this class of alternative policies cannot be used to consistently estimate θ_0; notice that $g(x; \theta)$ takes the same value, $\mathbb{E}\epsilon_t^2$, for every x in $\mathcal{X}_{\mathcal{E}_2}$.

Multiplicative Scale: We consider another natural class of policies, based on a multiplicative scale of the true policy, which has also been used in practice.[7] Formally, for some user-chosen distribution Δ, define

$$\mathcal{X}_{\mathcal{E}_3} = \{\sigma'(\cdot) : \sigma'(\cdot) = \epsilon\sigma(\cdot; \theta_0) \text{ for } \epsilon \in Support(\Delta)\}.$$

The measure $H_{\mathcal{E}_3}$ then denotes the measure that generates the random variable ϵ from Δ, *independent of ν_t*. To be more specific, suppose that Δ is the uniform distribution on a unit interval, the non-negative support is chosen to ensure that the alternative policies are also monotone on $\mathcal{V}$. For each $\theta \in \Theta$, the difference between payoffs from the true and an alternative x in $\mathcal{X}_{\mathcal{E}_3}$ are

$$\begin{aligned}
\pi\left(\sigma\left(\nu_t; \theta_0\right), \nu_t, \theta\right) - \pi\left(\epsilon\sigma\left(\nu_t; \theta_0\right), \nu_t, \theta\right) &= -\theta_0\left(1 - \epsilon\right)\left(\left(1 + \epsilon\right)\theta_0 - 2\theta\right)\nu_t^2, \\
g(x; \theta) &= -\theta_0\left(1 - \epsilon\right)\left(\left(1 + \epsilon\right)\theta_0 - 2\theta\right)\mathbb{E}_\nu\left[\nu_t^2\right].
\end{aligned}$$

Assuming $\mathbb{E}_\nu\left[\nu_t^2\right] < \infty$, from examining $g(x; \theta)$ we see that if $\theta \geq \theta_0$, then it follows that $g(x; \theta) \geq 0$. So any $\theta \in [\theta_0, \infty)$ will imply that $Q_{\mathcal{E}_2}(\theta) = 0$. On the other hand it is easy to show that any $\theta \in (0, \theta_0)$ will imply $Q_{\mathcal{E}_2}(\theta) > 0$. The criterion function based on this class of alternative policies can at best consistently estimate the set $\Theta \backslash (0, \theta_0)$. We now show that the class of multiplicative scale can lead to criterion functions that has a unique minimizer at θ_0 when Δ has larger support.

Without any exclusions, $\mathcal{X}$ is the set of measurable functions $\sigma \in \Sigma$. For a generic $x \in \mathcal{X}$, it follows from simple algebra that

$$g(x; \theta) = -\left(\theta - \theta_0\right)^2 \mathbb{E}_\nu\left[\nu_t^2\right] + \mathbb{E}_\nu\left[\left(\theta\nu_t - \sigma\left(\nu_t\right)\right)^2\right] \quad \text{for any } \theta \in \Theta.$$

In this case, the class of alternative policies which is a multiplicative scale from the truth can ensures that we can construct criterion functions with a unique minimizer at θ_0 so long as the support of ϵ is sufficiently large. To see this, note that we require $g(x; \theta) < 0$ to hold with positive measure (on some $\mathcal{X}_{\mathcal{E}}$) whenever $\theta \neq \theta_0$; this inequality is equivalent to

$$\mathbb{E}_\nu\left[\left(\theta\nu_t - \sigma\left(\nu_t\right)\right)^2\right] < \mathbb{E}_\nu\left[\nu_t^2\right]\left(\theta - \theta_0\right)^2.$$

---

[7]Jeziorski (2009) applies uniformly distributed multiplicative shocks to his choice probabilities in implementing BBL's methodology.

For any $\theta = \theta_0 + \delta$, note that, by letting $\sigma(\nu_t) = (\theta - \eta)\nu_t$, the inequality above holds whenever $|\eta| < |\delta|$. This means that when Θ is a compact subset of $\mathbb{R}^+$ containing θ_0, setting $\mathcal{X}_\mathcal{E} = \mathcal{X}_{\mathcal{E}_3}$ and letting $H_\mathcal{E}$ be any continuous distribution with full support on the positive half-line will be sufficient to ensure that $Q_\mathcal{E}(\theta)$ has a well separated minimum at θ_0.

Conclusion: We show that applications of the moment inequality approach of BBL can lead to objective functions that cannot be used to consistently estimate an identified model. The practical implication is potentially serious as there is generally no economic prior nor statistical theory to help us select the set of inequalities to ensure that the corresponding objective function has a unique minimum in the limit. Since most applications of BBL use their point estimation method, in finite sample, various optimization routines will produce point estimates that may not be informative at all. Although we have not explicitly shown analogous identification problems in a set identified model, we also expect such examples to exist.

The intuition behind our findings is related to the identification issue studied in Domínguez and Lobato (2004). They show one can lose identification in an identified conditional moment restriction model by only relying on a finite number of unconditional moment conditions. They also show how to obtain an informationally equivalent unconditional model, that can be used for consistent estimation, from conditional moment equality restrictions. Khan and Tamer (2009) applies a similar idea to preserve the informational content under the moment inequality framework. However, the inequality restrictions considered in BBL is complicated by the fact that their conditional model is indexed by the set of alternative policy functions.[8] As our example shows, even if we begin with an identified unconditional model, we can lose identification by only considering a subset of all alternative policies. The issue of how we can generally preserve the information implied by the set of alternative policies remains an open problem.

However, it is clear that this identification problem can be alleviated by integrating over larger classes of policies. It is important to note that considering more alternative policies in the same class may not be informative. It is more crucial to consider *broader classes of policies*. We show that, even when the class of alternative policy functions spans A pointwise on $\mathcal{V}$, it does not necessarily grow dense in the set of all alternative policies. We illustrate this point with the most widely used method to construct alternatives, by adding *random noise* to the policy, where we exploit the fact that these alternatives do not use any information on the state variable. The intuition behind this is that the spanning property only implies we cover the same range as the set of all alternative policies. However, this is not sufficient, as the set of alternative policies are not determined by their range, but by how the state variable is transformed into an action. Randomly perturbing the true policy at

---

[8]Unlike $\mathbb{R}^n$, a set of functions has no obvious ordering property so we cannot simply extend the techniques used in Domínguez and Lobato (2004) and Khan and Tamer (2009).

each state does not allow the perturbation shock to depend *stochastically* on the state, it merely uses $\mathcal{V}$ as an index set. So it is not surprising that the set of randomly perturbed alternatives is not dense in the set of all alternatives. Therefore, the perturbations which are correlated to the states will contain valuable information for identifying the parameter of interest and these should be exploited in practice.

Lastly, we comment that most applications of BBL methodology focus on alternative policies which do not depend on the model. In the other extreme, one can construct inequalities from model implied policies. The advantage of using the model is that we can write down an interpretable identification condition, if this holds then we can consistently estimate the true parameter of interest under some standard regularity conditions. However, implementing this in the forward simulation framework may likely limit the computational advantages described in BBL.

References

[1] BAJARI, P., C. L. BENKARD, AND J. LEVIN (2007): "Estimating Dynamic Models of Imperfect Competition," *Econometrica*, **75**, 1331-1370.

[2] DOMÍNGUEZ, M.A., AND I.N. LOBATO (2004): "Consistent Estimation of Models defined by Conditional Moment Restrictions," *Econometrica*, **72**, 1601-1615.

[3] JEZIORSKI, P. (2009): "Dynamic Determinants of Mergers and Product Characteristics in the Radio Industry," *Working Paper, Stanford University.*

[4] KHAN, S. AND E. TAMER (2009): "Inference on Endogenously Censored Regression Models Using Conditional Moment Inequalities," *Journal of Econometrics*, **152**, 104-119.

[5] NOTON, C. (2009): "Structural Estimation of Price Adjustment Costs in the European Car Market," *Working Paper, University of California Berkeley.*

[6] RYAN, S. (2009): "The Costs of Environmental Regulation in a Concentrated Industry," *Working Paper, MIT.*

[7] SANTOS, C. (2009): "Recovering the Sunk Costs of R&D: The Moulds Industry Case," *Working Paper, University of Alicante.*