Estimation Dynamic Decision Problems and Dynamic Games

Sorawoot Srisuma (NUS and University of Surrey)

PIER, 19/01/2024

Aim

- Provide an overview of some structural dynamic models that are used in industrial organizations applications and other fields such as labor, health, enviromental, urban, marketing etc.
- Prototypical dynamic discrete models will be introduced and the focus will be on how to estimate them.
- There will be discussions on extensions.

Why consider a dynamic structural model?

What does structural mean and why is it useful?

- Structural means data is assumed to be generated from an economic model.
- Structural parameters have clear interpretation.
- Allows for direct testing of economic theory or perform counterfactual studies.

Why dynamics?

- Today's actions may have consequences on the future.
- Examples: consumers purchasing durable goods, firms' investment levels, entry decisions with sunk costs.
- Static models can paint incomplete pictures and may not be able to answer some questions.

Motivating examples

In Sanches, Silva, and Srisuma (2018), we were interested in effects of privatizations of public banks in rural Brazilian markets. Particularly, what privatization would do to market structure.

- High sunk costs for banks to open branches in rural areas; banks are clearly forward looking.
- Parameters in the model include entry costs and operational costs (along with monopoly and duopoly profits etc).
- Our counterfactual study turns public banks into private ones. Our model suggests the number of bank branches would reduce dramatically in rural areas. We are able to make policy a recommendation that it is more cost effective for the government to provide subsidies towards reducing operational costs relative to entry costs.

Some useful surveys - SA

- Rust (1994): "Estimation of Dynamic Structural Models: Problems and Prospects Part I: Discrete Decision Processes," *Proceedings of the 6th World Congress of the Econometric Society*, *Cambridge University Press*.
- Aguirregabiria and Mira (2010): "Dynamic Discrete Choice Structural Models: A Survey," *Journal of Econometrics*, 156, 38-67.
- Arcidiacono and Ellickson (2011): "Practical Methods for Estimation of Dynamic Discrete Choice Models" Annual Review of Economics Volume 3, 363-394.

Some useful surveys - Games

- Ackerberg, Benkard, Berry, and Pakes (2005): "Econometric Tools for Analyzing Market Outcome," *Handbook of Econometrics*, vol. 6, eds. J. Heckman and E. Leamer. North-Holland.
- Aguirregabiria and Nevo (2013): "Recent Developments in Empirical IO: Dynamic Demand and Dynamic Games". Advances in Economics and Econometrics: Theory and Applications. Tenth World Congress of the Econometric Society.
- Bajari, Hong, and Nekipelov (2013): "Econometrics for Game Theory". Advances in Economics and Econometrics: Theory and Applications. Tenth World Congress of the Econometric Society.

Today's Talk

1. Estimating single-agent decision problems
2. Estimating dynamic games
3. Least squares estimation of dynamic games with unobserved heterogeneity

Single-agent decision problems

Markov decision problem

- a_t : control variable, *consumption, price, investment*
- s_t : state variable, *capital, wealth, productivity*
- u : payoff function, *utility, profit*
- β : discounting factor
- $P\left(s_{t+1} | s_t, a_t\right)$: transition law

$$\begin{array}{ccccccccc} s_t & \longrightarrow & s_{t+1} & \longrightarrow & s_{t+2} & \longrightarrow & s_{t+3} & \cdots \\ & \searrow & \uparrow & \searrow & \uparrow & \searrow & \uparrow & \\ & & a_t & & a_{t+1} & & a_{t+2} & \end{array}$$

$$u\left(a_t, s_t\right) + \beta u\left(a_{t+1}, s_{t+1}\right) + \beta^2 u\left(a_{t+2}, s_{t+2}\right) + \cdots$$

At time t, given s_t, the agent's problem is to choose $\{a_\tau\}_{\tau=t}^{\infty}$ in order to maximize $E\left[\sum_{\tau=t}^{\infty} \beta^{\tau-t} u\left(a_\tau, s_\tau\right) \middle| s_t\right]$.

Markov decision problem

Under some regularity conditions there is a stationary solution: (see Stokey and Lucas (1989), Pakes (1994) and Rust (1994))

1. Value function is the unique solution to the Bellman's equation,

$$V\left(s_t\right) = \max_{a \in A} \left\{ u\left(a, s_t\right) + \beta E\left[V\left(s_{t+1}\right) \middle| s_t, a_t = a\right]\right\}.$$

2. There exists a Markovian policy function $\alpha : S \rightarrow A$ s.t.,

$$\alpha\left(s_t\right) = \arg\max_{a \in A} \left\{ u\left(a, s_t\right) + \beta E\left[V\left(s_{t+1}\right) \middle| s_t, a_t = a\right]\right\}.$$

The primitives of the model are $\mathcal{P} := \left(u, \beta, P\right)$.

Econometrics problem

- Suppose $\mathcal{P} \to \mathcal{D}_T$ where $\mathcal{D}_T := \{d_t\}_{t=1}^T$ for $d_t \subseteq (a_t, s_t)$.
- The goal is to learn $\mathcal{P}$ from $\mathcal{D}_T$.
- Some traditional concepts:
    1. **Identification** - concerns the limits of what we can learn with infinite data.
    2. **Estimation and inference** - work with finite T (with $T \to \infty$).
- The general dynamic model is under-identified (Rust (1994), Magnac and Thesmar (2002)).
- Parametric and nonparametric assumptions as well as institution knowledge aids identification.
- (For the large parts, asymptotics are quite standard.)
- A particular consideration that arises when estimating dynamic models is **computation**.

Econometrics problem

- We assume to have observable and unobservable states denoted by x_t and ε_t respectively: $s_t = (x_t, \varepsilon_t) \in X \times \mathcal{E}$, and let $\mathcal{D}_T = \{(a_t, x_t)\}_{t=1}^{T}$.
- We assume enough about the primitives are known and the remaining parts are identified unless stated otherwise.
- Let's focus on parametric payoff function., i.e. $\mathcal{P}$ becomes : (u_θ, β, P) that reduces to θ.

Trouble with value function

- An agent's optimal decision rule depends on the **value function**.
- The value function is generally intractable and is the main source of the computational problem in estimating many dynamic models.
- There are exceptions to this rule. E.g., suppose a_t denotes consumption. Under some conditions, we have the standard consumption-saving model where the optimal solution is characterized by the *Euler equation*:

$$u'_\theta\left(\alpha\left(s_t\right), s_t\right) = \beta E\left[u'_\theta\left(\alpha\left(s_{t+1}\right), s_{t+1}\right) \middle| s_t\right].$$

In this case we can identify and estimate θ under various assumptions.

Trouble with value function

- We usually have to deal with the value function explicitly when estimating model primitives in IO applications.
- The nature of the control variable leads to different models. There are two main cases.
  1. Unordered discrete choice (UDC);
  2. Ordered choice (OC).
     (These can be layered so one can have sequential decisions. E.g. entry then invest.)
- We will focus on UDC given time limitations.

UDC - model assumptions

Suppose $A = \{1, \ldots, K\}$ and $\mathcal{E} = \mathbb{R}^K$. The following assumptions come from Rust's seminal work (Rust (1987)):

Assumption UDC

1. (*Additive separability*)
   $u_\theta (a, x, \varepsilon) = \pi_\theta (a, x) + \sum_{a' \in A} \varepsilon (a') \mathbf{1} [a' = a]$.
2. (*Conditional independence*)

   $$P (x_{t+1}, \varepsilon_{t+1} | x_t, \varepsilon_t, a_t) = Q (\varepsilon_{t+1} | x_{t+1}) G (x_{t+1} | x_t, a_t),$$

   with $Q (\cdot)$ and $G (\cdot)$ known.
3. (*Finite observable states*) $X = \{1, \ldots, J\}$.

AS and CI put the econometrics problem on familiar grounds (cf. McFadden (1974)).
Finiteness of X is an important simplification assumption.

UDC - an example

Harold Zucher's bus engine problem (Rust (1987)):

- a_t takes value $\{0, 1\}$ representing not replace or replace the engine respectively.
- x_t is the mileage of the bus.
- ε_t some unobserved utility shock.
- $u_\theta(a_t, x_t, \varepsilon_t) = -\theta_1 \cdot a_t - c_{\theta_2}\left(\left(1 - a_t\right) \cdot x_t\right) + \varepsilon_t$.
- $\left(\varepsilon_t\left(1\right), \varepsilon_t\left(0\right)\right)$ has i.i.d. Type 1 extreme value distribution
- $x_{t+1}|x_t$ follows an exponential distribution if $a_t = 0$ ow distributed as $x_{t+1}|x_t = 0$.

UDC - likelihood function

How do we estimate such model?

Consider the likelihood function of θ given $\mathcal{D}_T$ and some initial value (a_0, x_0):

$$
\begin{aligned}
L(\theta; \mathcal{D}_T) &= \prod_{t=1}^{T} \Pr_\theta [x_t, a_t | x_{t-1}, a_{t-1}, x_{t-2}, a_{t-2}, \ldots] \\
&= \prod_{t=1}^{T} \Pr_\theta [x_t, a_t | x_{t-1}, a_{t-1}] \\
&= \prod_{t=1}^{T} \Pr_\theta [a_t | x_t] \Pr [x_t | x_{t-1}, a_{t-1}] .
\end{aligned}
$$

We can estimate θ, for example, by maximizing the (log-)likelihood function wrt the parameters of the choice probabilities independently of $\Pr [x_t | x_{t-1}, a_{t-1}]$.

UDC - choice specific value function

The CCP looks like a standard multinomial probability:

$$\begin{aligned} & \Pr_{\theta}\left[a_t = a | x_t = x\right] \\ = \; & \Pr\left[v_\theta\left(a, x\right) + \varepsilon\left(a\right) > v_\theta\left(a', x\right) + \varepsilon\left(a'\right) \quad \text{for all} \quad a' \neq a\right], \end{aligned}$$

where $v_\theta\left(a, x\right) := \pi_\theta\left(a, x\right) + \beta \mathcal{V}_\theta\left(a, x\right)$ and

$$\mathcal{V}_\theta\left(a, x\right) := E\left[V_\theta\left(s_{t+1}\right) | x_t = x, a_t = a\right].$$

But $\mathcal{V}_\theta\left(\cdot\right)$ does not have a closed form and is only implictly defined.

UDC - NFP (Rust (1987))

Rust (1987) shows $\mathcal{V}_\theta(\cdot)$ is a fixed-point of some map and suggested a nested fixed-point MLE:

1. **Outer loop**: search over θ.
2. **Inner loop**: for each θ, perform fixed-point iterations to approximate $\mathcal{V}_\theta(\cdot)$.

- At the theoretical level, Rust's estimator is CAN and fully efficient.
- But it is computationally expensive and worsens in higher dimensions with additional issues for games.

UDC - Hotz and Miller (1993)

An alternative is to take the following *two-step* approach due to Hotz and Miller (1993).

Instead of solving the model for each θ (i.e. finding $\mathcal{V}_\theta(\cdot)$), consider the expected payoffs based on the observed decision rule. In particular, we know:

- $\mathcal{D}_T = \{(a_t, x_t)\}_{t=1}^{T}$ where $a_t = \alpha_{\theta_0}(s_t)$ for some θ_0.
- We know that $\mathcal{V}_{\theta_0}(a, x) = \sum_{\tau=1}^{\infty} \beta^{\tau-1} E\left[u_{\theta_0}(\alpha_{\theta_0}(s_\tau), s_\tau) \middle| x_0 = x, a_0 = a \right]$, replace $\mathcal{V}_\theta(a, x)$ with $\mathcal{W}_\theta(a, x) := \sum_{\tau=1}^{\infty} \beta^{\tau-1} E\left[u_\theta(\alpha_{\theta_0}(s_\tau), s_\tau) \middle| x_0 = x, a_0 = a \right]$.

UDC - Hotz and Miller (1993)

- Let $w_\theta(\cdot) := \pi_\theta(a,x) + \beta \mathcal{W}_\theta(a,x)$ where

$$\mathcal{W}_\theta(a,x) := \sum_{\tau=1}^{\infty} \beta^{\tau-1} E\left[u_\theta(a_\tau, s_\tau)\middle| x_0 = x, a_0 = a\right].$$

- HM show that w_θ can be written as a function of $\{P_{\theta_0}(a|x_t)\}_{a\in A}$ (and (β, Q, G)).
- This leads to model implied CCP,

$$P_\theta(a|x) := \Pr\left[w_\theta(a,x) + \varepsilon(a) > w_\theta(a',x) + \varepsilon(a') \quad \text{for all} \quad a' \neq a\right]$$

and a GMM estimator based on $E\left[\mathbf{1}[a_t = a] - P_\theta(a|x_t)\,|x_t\right] = 0$ when $\theta = \theta_0$,
- HM estimator is a standard two-step estimator where P_{θ_0} is estimated in the first stage.

UDC - Full solution vs two-step

Summary of the main idea.
Full solution:

- NFP solves for the *optimal* expected payoffs for each θ.
- This requires fixed-point iteration for each θ.

Two-step:

- Two-step approach constructs the expected payoffs based on observed policy function.
- This uses a *pseudo-optimal* expected payoffs for each θ that does not require solving of a dynamic programming problem even once. We will come back to this when discussing games.

The price for taking the two-step approach are loss of efficiency (relative to MLE) and first-step estimation can induce substantial finite sample bias (Aguirregabiria and Mira (2002), Kasahara and Shimotsu (2008)0.

Comments on assumptions - discount factor

- The discount factor is usually assumed to be known and often presumed to not be identified.
- Magnac and Thesmar (2002) show the discount factor can be identified in a simple two-period model.
- Komarova, Sanches, Silva, and Srisuma (2018) identify the discount factor under parametric assumptions.
- Abbring and Djalford (2020) identify the discount factor under nonparametric exclusion restrictions.

Comments on assumptions - finite X

- Finiteness of X is often done (sometimes manually) in applied work.
- Even if CCP and transition laws are not known they have parametric convergence rate.
- Ow having continuous components in X makes the model semiparametric.
- Altug and Miller (1998) use finite-dependence to ensure root-N/T consistency of the model parameters.
- Srisuma and Linton (2012) show FD is not necessary using theory from inverse problems.

Comments on assumptions - unobserved states

- The distribution of the unobserved state variable is partially identified with limited support (Norets and Tang (2013)).
- Nonparametric identification of the unobserved state is possible under exclusion restrictions together with a large support condition (Chen (2017)).
- Recent works allow for additional form of unobserved heterogeneity (Arcidiacono and Miller (2013)) or correlated unobservables (Norets (2009), Hu and Shum (2012)). More on this later.

Dynamic games

Elements of the game

- Players: $\mathcal{I} = \{1, \ldots, I\}$
- States: $s_{it} = (x_t, \varepsilon_{it}) \in S_i \; (:= X \times \mathcal{E}_i)$ *e.g. capital, productivity, number of active firms*
- Actions: $a_{it} \in A_i, a_{-it} \in A_{-i}, a \in A$ *e.g. consumption, investment, entry decision*
- Payoff functions: $u_i : A \times S_i \to \mathbb{R}$
- Discount factor: $\{\beta_i\}_{i=1}^{I}$
- Transition law: $P\left(s_{t+1} | s_t, a_t\right)$

Game progression:
players observe $\{s_{it}\}$ $\to$ choose actions $\{a_{it}\}$ $\to$ yield $u_i\left(a_t, s_{it}\right)$ $\to$ state evolves $\{s_{it+1}\}$ $\to$ game repeats...

Elements of the game

Game progression:

$$
\begin{array}{ccccccccc}
 & (s_{1t}) & a_{1t} & (s_{1t}) & a_{1t+1} & (s_{1t+1}) & \\
 & \nearrow & \downarrow & \nearrow & \downarrow & \nearrow & \\
\{s_{it}\}_{i=1}^{I} & \longrightarrow & \{s_{it+1}\}_{i=1}^{I} & \longrightarrow & \{s_{it+2}\}_{i=1}^{I} & \longrightarrow & \cdots \\
 & \searrow & \uparrow & \searrow & \uparrow & \searrow & \\
 & (s_{It}) & a_{It} & (s_{It}) & a_{It+1} & (s_{It+1}) &
\end{array}
$$

Player i's problem is to choose a_i to maximize

$$
\begin{aligned}
\Lambda_i\left(a_i, s_{it}; \sigma_i\right) \quad = \quad & E_{\sigma_i}\left[\left. u_i\left(a_{it}, a_{-it}, s_{it}\right)\right| s_{it}, a_{it} = a_i\right] \\
& + \beta_i E_{\sigma_i}\left[\left. V_i\left(s_{it+1}; \sigma_i\right)\right| s_{it}, a_{it} = a_i\right]
\end{aligned}
$$

based on beliefs σ_i, where
$V_i\left(s_{it}; \sigma_i\right) := \sum_{\tau=t}^{\infty} \beta_i^{\tau-t} E_{\sigma_i}\left[\left. u_i\left(a_{i\tau}, a_{-i\tau}, s_{i\tau}\right)\right| s_{it}\right]$.

Elements of the game

- In an equilibrium, for all i:
  1. For all i, $a_{it} = \arg\max_{a_i} \Lambda_i \left(a_i, s_{it}; \sigma_i\right)$;
  2. The beliefs are consistent with the distribution of $a_{it} | s_{it}$.
- The *primitives* of the game are $\left(\{u_i\}_{i=1}^{I}, \{\beta_i\}_{i=1}^{I}, P\right)$.
- We can parameterize payoffs by $\{\theta_i\}$.
- Suppose we have: $\mathcal{D}_T := \{d_t\}_{t=1}^{T}$, where $d_t = (a_t, x_t)$.
- The econometrics goal is to learn $\mathcal{P}$ from $\mathcal{D}_T$.

Problems when estimating games

1. Multiple equilibria leading to incomplete model (Tamer (2003)).

- Analyzed in static games, see Bajari, Hong, Krainer and Nekipelov (2009), Aradillas-Lopez (2010), Lewbel and Tang (2012).

2. Even if there is no indeterminacy issue, repeated *full-solution* approach is numerically infeasible.

- Computing equilibrium expected payoffs in games is a much harder task than in the single-agent's problem. E.g., see Pakes and McGuire (2001).

The two-step approach avoids these problems if one assumes the data comes from a single equilibrium.

Two-step estimators

Focusing on UDC games:

- There are several estimators: Aguirregabiria and Mira (2007), Pakes et al. (2008), Pesendorfer and Schmidt-Dengler (2008) and Bajari et al. (2012).
- PSD propose a class of *asymptotic least squares estimator* that is based on the equilibrium condition and encompasses the non-iterated estimator of AM and POB.
- Efficient PSD's estimator dominates estimators of AM and POB.
- I advocate the PSD approach, particularly based on Sanches, Silva and Srisuma (2016) who showed estimation of single agent and game models can be very simple and have a closed form solution in many cases.
- PSD and SSS estimators are asymptotically equivalent.

Maintaining assumptions

Assumption G

1. (*Additive separability*)
   $u_{i,\theta}\left(a_i, a_{-i}, x, \varepsilon\right) = \pi_{i,\theta}\left(a_i, a_{-i}, x\right) + \sum_{a_i' \in A} \varepsilon_i\left(a_i'\right) \mathbf{1}\left[a_i = a_i'\right]$
2. (*Conditional independence*)

   $$P\left(x_{t+1}, \varepsilon_{t+1} \middle| x_t, \varepsilon_t, a_t\right) = Q\left(\varepsilon_{t+1} \middle| x_{t+1}\right) G\left(x_{t+1} \middle| x_t, a_t\right),$$

   with $Q(\cdot)$ and $G(\cdot)$ known.
3. (*Independent private values*) $Q\left(\varepsilon_t | x_t\right) = \prod_{i=1}^{I} Q_i\left(\varepsilon_{it} | x_t\right)$.
4. (*Finite observable states*) $X = \{1, \ldots, J\}$.

These are standard assumptions in the literature that we will assume throughout.

Markovian framework

(E.g. see Maskin and Tirole (1994, 2001))

- **Markov strategies**
  $a_{it} = \alpha_i(s_{it}) = \alpha_i(s_{it'}) = a_{it'}$ whenever $s_{it} = s_{it'}$ for any t, t'.
- **Markov belief** (σ_i) is distribution of a_t condition on x_t.

Definition. *(Markov Perfect Equilibrium)*
A collection (α, σ) is a MPE if

1. *for all i, α_i is a best response to α_{-i} given the beliefs σ_i;*
2. *all players use Markov strategies;*
3. *for all i, the beliefs σ_i are consistent with the strategies α_i.*

It can be shown pure strategy MPE exists under Assumption G (Aguirregabiria and Mira (2007), Pesendorfer and Schmidt-Dengler (2008)).

An entry game example

From PSD:

- 2 players and binary actions, $a_{it} \in \{0, 1\}$.
- Public information, $x_t = (a_{1t-1}, a_{2t-1})$.
- Private values, ε_{it} are i.i.d. normal or logit .
- Parameterize payoffs by $\theta = (\pi_0, \pi_1, F, W)^\top$:

$$\begin{aligned}\pi_{i,\theta}(a_{it}, a_{-it}, x_t) &= (\pi_0 + \pi_1 a_{-it}) \cdot \mathbf{1}[a_{it} = 1] \\ &\quad + F \cdot \mathbf{1}[a_{it} = 1, a_{it-1} = 0] \\ &\quad + W \cdot \mathbf{1}[a_{it} = 0, a_{it-1} = 1].\end{aligned}$$

An entry game example

In equilibrium (omit σ), under G

$$\alpha_{i,\theta}(x_t, \varepsilon_{it}) = \arg\max_{a_i \in \{0,1\}} \{w_{i,\theta}(a_i, x_t) + \varepsilon_{it}(a_i)\},$$

$$w_{i,\theta}(a_i, x_t) := E[\pi_{i,\theta}(a_{it}, a_{-it}, x_t) | x_t, a_{it} = a_i] + \beta \sum_{\tau=0}^{\infty} \beta^{\tau} E[\pi_{i,\theta}(a_{it+\tau}, a_{-it+\tau}, x_{t+\tau}) + \varepsilon_{it+\tau}(a_{it+}$$

$w_{i,\theta}$ resembles w_θ in the single agent problem, and is in fact just a function of $(\theta, \beta, Q, G, P_{i,\theta_0})$. Various estimators can be defined based on $\alpha_{i,\theta}$ or $P_{i,\theta}$ where

$$P_{i,\theta}(a_i|x) := \Pr\left[w_{i,\theta}(a_i, x) + \varepsilon(a_i) > w_{i,\theta}(a_i', x) + \varepsilon(a_i') \quad \text{for all} \quad a_i' \neq\right.$$

In practice, (β, Q) are often assumed to be known and (G, P_{i,θ_0}) are estimated from data.

Estimating games

- Much of insights from a single agent model applies with the two-step approach.
- Once $\widehat{\mathcal{P}}$ is available change to structural parameters or game structures can be applied for counterfactual simulations.
- Recent econometric research are focusing on issues of multiple equilibria but has seen very limited applications in the dynamic setting.

Least Squares Estimation of Dynamic Games with Unobserved Heterogeneity

G. Guo (NUS) A. Gupta (Essex)

F. Sanches (FGV) S. Srisuma (NUS and Surrey)

Overview

- We propose a new estimator for dynamic discrete games.
- Goal is to provide a simple estimator that can handle time varying unobserved heterogeneity.
- Idea is:
  - estimate CCP in each time period and do k-mean clustering to identify types.
  - structural parameters and unobserved heterogeneity effects are estimated by LS (e.g., by OLS).
- Aim of this talk is illustrate the main ideas for estimating games, and of our estimator and its numerical performance.

Game Setup

Similar setup as before. The focus is on an additional state variable and linear payoff function:

- $s_{mt} = (x_{mt}, d_t, \varepsilon_{mt})$ only x_{mt} is observed
- $u_{i,\theta}(a_{mt}, x_{mt}, d_t, \varepsilon_{mt}) =$ $\theta^\intercal \pi_i(a_{mt}, x_{mt}, d_t) + \sum_{a_i' \in A} \varepsilon_{imt}(a_i') \cdot \mathbf{1}[a_{imt} = a_i']$
- $P(x_{mt+1}, d_{t+1}, \varepsilon_{mt+1} | x_{mt}, d_t, \varepsilon_{mt}, a_{mt}) =$ $\prod_{i=1}^{I} Q_i(\varepsilon_{imt+1}) H(d_{t+1}|d_t) G(x_{mt+1}|x_{mt}, a_{mt})$
- x_{mt} and d_t have finite support

The general idea of the estimator applies without the linear payoff but no closed-form estimator.
Assume data from a single equilibrium

$$(\{u_{i,\theta_0}\}_{i=1}^{I}, \beta, P) \mapsto \{a_{imt}, x_{mt}\}_{m=1.t=1}^{M,T}$$

Want to learn about θ_0 as $M, T \to \infty$ with $T = o(M)$.

Related Literature

- Two-step game estimators without unobserved heterogeneity (UH) have been proposed by Aguirregabiria and Mira (2007), Bajari, Benkard and Levin (2007), Pakes, Ostrovsky and Berry (2007), Pesendorfer and Schmidt-Dengler (2008), Srisuma (2013), Sanches, Silva, and Srisuma (2016, SSS).
- UH is modelled using finite mixture, which is identified and estimated under some conditions (Kasahara and Shimotsu (2009), Hu and Shum (2012), Arcidiacono and Miller (2011)). Estimation involves nonlinear optimization routines and typically relies on further assumptions.
- Focusing on time varying UH, Kaloupstidi, Scott, and Souza-Rodrigues (2021) showed easy estimation (IV-type) of θ_0 is possible under *finite dependence* and *completeness* assumptions, but they can't be used for counterfactual studies.
- Our proposal, builds on SSS, is easy to compute and does not need FD or completeness + can use for counterfactuals.

ALS - Estimation Idea

Consider a model of *binary actions* based on

$$a_t(\theta) = \mathbf{1}\left[w_\theta(x_t) - \varepsilon_t \geq 0\right] \quad \text{for} \quad \theta \in \Theta \subset \mathbb{R}^p,$$

where x_t and ε_t are obs and unobs *state variables* s.t. w_θ and distribution of $\varepsilon_t | x_t$ are known (call its cdf Q).
For all x,

$$P_\theta(x) := \Pr\left[a_t(\theta) = 1 | x_t = x\right] = Q\left(w_\theta(x)\right).$$

Suppose we observe a random sample $\{a_t, x_t\}_{t=1}^T$ where $a_t = a_t(\theta_0)$ for some $\theta_0 \in \Theta$. θ_0 can be estimated based on minimizing the distance between $P(\cdot) := \Pr\left[a_t = 1 | x_t = \cdot\right]$ and $P_\theta(\cdot)$.

ALS - Estimation Idea

- When x_t is discrete, vectorizing $P(x)$ and $P_\theta(x)$ leads to

$$\widehat{\theta}_p(\mathcal{A}) = \arg\min_{\theta\in\Theta} \left(\widetilde{\mathbf{P}} - \mathbf{P}_\theta\right)^\top \mathcal{A} \left(\widetilde{\mathbf{P}} - \mathbf{P}_\theta\right), \tag{1}$$

  $\widetilde{\mathbf{P}}$ is a nonparametric estimator for $\mathbf{P}$.
- We prefer to minimize expected payoffs/rewards,

$$\widehat{\theta}_w(\mathcal{B}) = \arg\min_{\theta\in\Theta} \left(\widetilde{\mathbf{w}} - \mathbf{w}_\theta\right)^\top \mathcal{B} \left(\widetilde{\mathbf{w}} - \mathbf{w}_\theta\right), \tag{2}$$

  $\widetilde{\mathbf{w}}$ estimates $\mathbf{w}$ (vectorized $Q^{-1}(P(x))$), because when $\mathbf{w}_\theta = \mathbf{X}\theta$,

$$\widehat{\theta}_w(\mathcal{B}) = \left(\mathbf{X}^\top \mathcal{B}\mathbf{X}\right)^{-1} \mathbf{X}^\top \mathcal{B}\widehat{\mathbf{v}}$$

  N.B. dynamic games have similar structures as $w_\theta(x)$ becomes $E[u_{i,\theta}(a_t, s_t) | a_{it} = 1, x_t = x] + \beta E[...] + ...$
- SSS showed $\widehat{\theta}_w$ and $\widehat{\theta}_p$ are asymptotically equivalent (cf. Pesendorfer and Schmidt-Dengler (2008)).

ALS - Infeasible Estimation

- Our game assumptions give:

$$\mathcal{Y} = \mathcal{X}\theta_0,$$

  where $\mathcal{X}$ and $\mathcal{Y}$ are known functions of (β, Q, H, G).
- Under a full rank condition $\theta_0 = \left(\mathcal{X}^{\intercal}\mathcal{X}\right)^{-1} \mathcal{X}\mathcal{Y}$.
- Given estimators $\left(\widehat{\mathcal{Y}}, \widehat{\mathcal{X}}\right)$, eqv $\left(\widehat{H}, \widehat{G}\right)$, under suitable regularity conditions

$$\begin{aligned}
\widehat{\theta}(\widehat{\mathcal{W}}) &= \arg\min_{\theta\in\Theta}(\widehat{\mathcal{Y}} - \widehat{\mathcal{X}}\theta)^{\top}\widehat{\mathcal{W}}(\widehat{\mathcal{Y}} - \widehat{\mathcal{X}}\theta) \\
&= (\widehat{\mathcal{X}}^{\top}\widehat{\mathcal{W}}\widehat{\mathcal{X}})^{-1}\widehat{\mathcal{X}}^{\top}\widehat{\mathcal{W}}\widehat{\mathcal{Y}} \xrightarrow{d} N\left(0, \Sigma_{\mathcal{W}}\right).
\end{aligned}$$

- Properties of ALSE are driven by properties of $\left(\widehat{H}, \widehat{G}\right)$.

ALSE - Feasible Estimation

- Without observing d_t, how to estimate H?
- Our proposal:

  1. Np estimate $P_t(a_i|x) = \Pr[a_{imt} = a_i | x_{mt} = x] \rightarrow \left\{\widehat{\mathbf{P}}_t\right\}_{t=1}^{T}$.
  2. Perform k-mean clustering on $\left\{\widehat{\mathbf{P}}_t\right\}_{t=1}^{T}$ to obtain $\left\{\widehat{d}_t\right\}_{t=1}^{T}$.
  3. Feasible estimators of H and $(\mathcal{X}, \mathcal{Y})$ can then be obtained, leading to closed-form GLS/OLS estimators that are asymptotically normal.

- Under suitable conditions centers of clusters can be consistently estimated and have limiting normal distribution and $\left\{\widehat{d}_t\right\}_{t=1}^{T}$ can be consistently estimated with known cvg rate.
- Our theory draws on Pollard (1981,1982) and Bonhomme and Manresa (2015).

Simulations

Extend the 2-firm entry game in Pesendorfer and Schmidt-Dengler (2008)

- $a_{it} \in \{0, 1\}$ and $x_t = (a_{1t-1}, a_{2t-1})$
- Firm $1'$s period payoffs is

$$\begin{aligned}\theta^{\intercal}\pi_1(a_t, x_t, d_t) &= a_{1t}(1 - a_{2t})\mu_1 + \mu_2 a_{1t} a_{2t} \\ &\quad + a_{1t}(1 - a_{1t-1})F + (1 - a_{1t})a_{1t-1}W \\ &\quad + a_{1t}\sum_{s=1}^{\tau}\omega_s d_{st},\end{aligned}$$

  where $d_{st} = 1$ iff in type s (ow 0), and impose $\omega_\tau > \omega_{\tau-1} > \cdots > \omega_1 = 0$.

- Performed 1,000 simulations for different combinations of (M, T) and τ
- Estimation by OLS and use parametric bootstrap

Concluding Remarks on Paper

- Propose how to estimate games with unobserved time varying heterogeneity.
- Estimation is very simple:
  1. use k-mean clustering to classify CCPs into types
  2. extend SSS to do OLS/GLS accordingly
- Idea readily extendable to ordered discrete games (e.g., Gowrisankaran, Lucarelli, Schmidt-Dengler and Town (2018)).
- Allowing for continuous observables is an interesting and unsolved theoretical challenge.
- We also have estimate of a model of traders buying/selling Blur tokens.

Summary of Talk

- Dynamic models offer opportunities to analyze interesting problems.
- *But one has to be aware of the rigidity a structural framework can impose on the data.*
- The presentation focused on dynamic discrete choice models. Analogous methods exist for decision problems and games with ordered discrete or continuous choice (Bajari, Benkard and Levin (2007) and Srisuma (2013)). Actions can also layered, e.g., entry first then decide on investment level.