ESTIMATING DYNAMIC MODELS USING TWO-STEP METHODS

KIET

March 26/27, 2024

Stephen P. Ryan

WUSTL, NBER, and CESifo

10000m Motivation

- **Dynamics are important in a variety of IO settings**
- **Firms make many decisions with an eye to the future:**
- Research and development
- Technology adoption
- Capacity investment
- Entry and exit of markets, product lines, variety
- Advertising
- **Key point: decisions today change the future**
- **Key question: how to estimate models of conduct with dynamics?**

How to Approach?

- **For a long time, this was considered a very difficult task**
- **Why?**
- **Let's think about how we estimate parameters in models**
- First, posit a theoretical model (even for reduced-form approaches)
- Second, form an estimator
- Third, solve for parameters that minimize some objective function
- **What does that look like in the case of dynamic games?**

Pieces of a Dynamic Model

- Dynamic model consists of three main components:

1. State space
2. Payoff functions
3. Transitions

Pieces of a Dynamic Model

- **Dynamic model consists of three main components:**

1. **State space**

- This a description of the economic environment
- Examples: how many firms are active in a market, their sizes, some measure of their technology/productivity, period of time, etc.

2. **Payoff functions**

3. **Transitions**

Pieces of a Dynamic Model

- **Dynamic model consists of three main components:**

1. **State space**
2. **Payoff functions**

- Players have actions available to them in each period
- Actions generate costs/benefits
- Flow payoffs as a function of the state

3. **Transitions**

Pieces of a Dynamic Model

- **Dynamic model consists of three main components:**

1. State space

2. Payoff functions

3. Transitions

- Between periods, as a result of endogenous and exogenous forces, state variables change
- Example of exogenous: movement in interest rate
- Example of endogenous: competitor decides to exit the industry

Forming an Estimator

- What problem is the firm solving?

$$max_a \pi(a, s; \theta) + \beta \int V(s')dP(s'; s, a)$$

- Ok, what is going on here?
- The per-period payoff function is $\pi(a, s; \theta)$: function of both state and action taken (and maybe actions of other firms!) Note θ!
- The discount factor is β
- The future is represented by a value function, $V(s)$, that represents the expected value of being active at that state (conditional on equilibrium)
- Transition function between states is given by $dP(s'; s, a)$

Empirical Content

- **Key question: what is the empirical content of this equation?**
- **Answer: the optimal policy function, $a^*(s)$**
- **What is our goal as an econometrician?**
- **We want to recover the full set of parameters that govern, along with the model, agent behavior**
- **Why? So that we can understand *what* happened, *why*, and, most importantly, what *would* happen in a different economic environment**
- **Our objective: given the structure of the model, what parameters rationalize the behavior of optimizing agents?**

Putting it Together

- **Two major empirical approaches in the literature**
- Maximum likelihood
- Generalized method of moments
- **Maximum likelihood is efficient, requires fully-specified model and analytic formulas (aside: SML)**
- **This is used in Rust (1994), Harold Zurcher engine replacement problem**
- **Moments don't require full specification, consistent under simulation (MSM)**
- **So let's think about GMM**

GMM Dynamic Model

- **We’ll form some moments from the data**
- **Several candidates:**
- Probability of entry or exit given state
- Probability of investment given state
- Size of investment conditional on investing
- Probability of introducing new product line / redesign
- Amount of advertising given opponent’s advertising, prior history
- **Key point: all of these are observable**
- **How do we use these to find parameters?**
- **Take a guess of the parameters, solve the prior Bellman equation, find optimal policy function, see how well it matches against the empirical moments in the data**

Sounds Good, Right?

- **One small problem with this approach**
- **It is generally infeasible**
- **The reason is that solving the dynamic programming problem is computationally expensive**
- The Bellman equation is a functional equation
- You must solve for the value function at every point in the state space, for all agents, jointly!
- **Also, there may be multiple equilibria (i.e. solutions to $V(s)$).**

Two-Step Methods to the Rescue

- One solution to this problem comes via Bajari, Benkard, and Levin (2004)
- Following the general ideas of Hotz and Miller (1993), let the agents solve the problem for us
- Then we try to find parameters that rationalize their behavior
- A table may help explain...

Comparison of Methods

$$max_a \pi(a, s; \theta) + \beta \int V(s')dP(s'; s, a)$$

	Method of Moments / Likelihood	Bajari, Benkard, and Levin
Optimal policy functions	Derived from computed full solution	Plug in observed empirical actions
State transitions	Derived from computed full solution	Plug in observed empirical transitions
Discount factor	Assumed	Assumed
Payoff function	Parameterized	Parameterized
Value function	Derived from computed full solution	Interesting…

Computation of Value Function

- Q: What does $V(s)$ tell us?
- A: The value of being at a given state
- Several variations on a theme of how to recover that without computing the full solution, but all the same idea:

$$V(s) = \sum_{t=0} \beta^t \pi(s_t; a(s_t))$$

- We can replace the value function with an infinite sum of payoffs in the future
- Note we need optimal action at each state, transitions between states
- BBL insight: we *observe* those objects in the data!

BBL Estimator

- Under optimizing behavior, agent can’t deviate from observed strategy and do better:

$$V(s; a^*(s); \theta) \geq V(S; a'(s); \theta), \forall a'(s) \neq a^*(s)$$

- Note that we can simulate both the left- and right-hand sides of that inequality!
- We just read off policy functions, empirical transitions, sum up payoffs
- How to get alternative policies? Anything works (in principle)
- BBL: Under true θ, inequality satisfied

Ryan (2012)

- My 2012 paper applies the BBL framework to the US Portland cement industry
- Characterized by very slow technological progress, regional markets, capacity constraints, infrequent entry/exit/investment
- Policy question: environmental regulation -> changes in market structure
- Specifically: 1990 Amendments to the Clean Air Act

Ingredients

- **Who are the players?**
- Cement plants
- **What are their actions?**
- Enter, exit, invest, divest
- **What is the state space?**
- Number of active firms, their capacities
- **What are the transitions?**
- Change in capacity (to zero for exits)

What Are We Trying to Do?

- **What are the unknown parameters?**
- The distribution of fixed costs to entry, exit, capacity adjustment, marginal cost of production, demand
- **I simplified the problem tremendously through assumption/exploiting institutional details**
- Regional markets
- Fixed demand, marginal costs, productivity
- **Static per-period output payoffs pin the value function**

What Needs to Be Estimated?

- **Two steps**
- First step: optimal policy functions (entry, exit, investment), transitions (degenerate in my case)
- Second step: project policy functions onto underlying model, find parameters that rationalize observed behavior

What Needs to Be Estimated?

- **Two steps**
- First step: optimal policy functions (entry, exit, investment), transitions (degenerate in my case)
- *Entry: probit*
- *Exit: probit*
- *Investment: (s,S) rule (Attanasio, 2000)*
- Second step: project policy functions onto underlying model, find parameters that rationalize observed behavior

What Needs to Be Estimated?

- **Two steps**
- First step: optimal policy functions (entry, exit, investment), transitions (degenerate in my case)
- Second step: project policy functions onto underlying model, find parameters that rationalize observed behavior

$$\boldsymbol{V}(\boldsymbol{s};\boldsymbol{a}^{*}(\boldsymbol{s});\boldsymbol{\theta}) \geq \boldsymbol{V}(\boldsymbol{S};\boldsymbol{a}'(\boldsymbol{s});\boldsymbol{\theta}), \forall \boldsymbol{a}'(\boldsymbol{s}) \neq \boldsymbol{a}^{*}(\boldsymbol{s})$$

Details / Discussion

- To recover a distribution, you have to take expectations conditional on the probability of the action, since there is selection
- You must trade off bias and variance in the reduced form policy functions (ML may help with this a lot!)
- Some alternative policies are not going to be informative
- Inference is still a bit tricky here
- Once you have the parameters in hand, you still face the same problem that BBL was designed to avoid to compute counterfactuals

Some General Takeaways

- It was not something I knew beforehand, but having independent regional markets was a huge boon to the estimation
- It is tremendously helpful to have something non-dynamic that can be estimated independently of the dynamic parameters to pin down the value function
- ML may really help in the first stage
- The literature has advanced next-to-zero on the computation of counterfactuals
- The empirical literature hasn’t made much progress, either
- You need credible estimates of the first-stage estimates to really make the whole thing work
- The bounds on profitability help you perform sanity check on other parameters

Thank you!

- That’s an overview of the general methodology
- Happy to discuss your specific application / question