Indirect Inference

Shuowen Chen

BW Econometrics Reading Group

Development of estimation methods

- Closed–form expressions
  - OLS, IV
- Maximum likelihood and moment–based estimation
  - development of mathematics (Jennrich, 1969)
  - Hansen (1982), Manski and Tamer (2002)
- Simulation–based method
  - i.i.d discrete choice (Pakes and Pollard, 1989)
  - consumption asset pricing (Duffie and Singleton, 1993)
  - indirect inference (Smith, 1993; Gourieoux et al., 1993)
- Deep learning
  - single layer neural network (Chen and White, 1999)

Uncertain demand (Collard–Wexler, 2013)

Motivation

- sunk entry costs, uncertainty in demand conditions creates a barrier to entry
- business-cycle fluctuations can have real effects on the structure of markets and welfare

Uncertain demand (Collard–Wexler, 2013)

Motivation

- sunk entry costs, uncertainty in demand conditions creates a barrier to entry
- business-cycle fluctuations can have real effects on the structure of markets and welfare

Ready–mix concrete industry

- homogeneous good
- local oligopolies: wet concrete hard to travel
- demand fluctuations from construction sector and gov.t

Uncertain demand (Collard–Wexler, 2013)

Motivation

- sunk entry costs, uncertainty in demand conditions creates a barrier to entry
- business-cycle fluctuations can have real effects on the structure of markets and welfare

Ready–mix concrete industry

- homogeneous good
- local oligopolies: wet concrete hard to travel
- demand fluctuations from construction sector and gov.t

How does demand volatility affect the organization of production?

- estimate a model of entry and discrete investment choices
- counterfactual: demand–smoothing policy

Dynamic oligopoly game

Firm i's state at time t:

$$s_i^t = \{x_i^t, \varepsilon_i^t\}$$

- x_i^t: current and past plant sizes (common knowledge)
- ε_i^t: i.i.d private information

Dynamic oligopoly game

Firm i's state at time t:

$$s_i^t = \{x_i^t, \varepsilon_i^t\}$$

- x_i^t: current and past plant sizes (common knowledge)
- ε_i^t: i.i.d private information

Firm action:

$$a_i^t = \{\emptyset, \text{small}, \text{medium}, \text{large}\}$$

Dynamic oligopoly game

Firm i's state at time t:

$$s_i^t = \{x_i^t, \varepsilon_i^t\}$$

- x_i^t: current and past plant sizes (common knowledge)
- ε_i^t: i.i.d private information

Firm action:

$$a_i^t = \{\emptyset, \text{small}, \text{medium}, \text{large}\}$$

Demand evolves with a transition probability: $D(M^{t+1}|M^t)$

Dynamic oligopoly game

Firm i's state at time t:

$$s_i^t = \{x_i^t, \varepsilon_i^t\}$$

- x_i^t: current and past plant sizes (common knowledge)
- ε_i^t: i.i.d private information

Firm action:

$$a_i^t = \{\emptyset, \text{small}, \text{medium}, \text{large}\}$$

Demand evolves with a transition probability: $D(M^{t+1}|M^t)$

Firm's payoff

$$r(x_i^{t+1}, x_{-i}^{t+1}|\theta) + \tau(x_i^{t+1} = a_i^t, x_i^t|\theta) + \varepsilon_i^t$$

- $r(\cdot)$: rewards from operation
- $\tau(\cdot)$: transition costs
- ε_i^t: logit shock

Timing of the game

1. Firms privately observe ε_i^t and publicly observe x^t
2. Firms simultaneously choose a_i^t
3. Demand evolves from M^t to M^{t+1}; firms states evolve
4. Payoffs are realized

What moments to match

Optimal conditional choice probabilities

$$\Psi(a_i|x, \Gamma, \theta)$$

- Γ: summarizes state evolution of choices and market demand

Natural way

- Given θ, simulate data, compute an equilibrium to the dynamic game, match to the data, repeat

Challenge

- computing an equilibrium for dynamic games is non−trivial

Structure of the presentation

- Key ingredients of indirect inference
- Connection with SMM
- Theoretical properties
- Bias corrections (in panel data analysis)

Indirect Inference: Set up

Dynamic model:

$$y_t = r(y_{t-1}, x_t, u_t, \theta)$$
$$u_t = \varphi(u_{t-1}, \varepsilon_t, \theta), \quad \theta \in \Theta \in \mathbb{R}^P$$

- x_t: observable exogenous variables
- u_t: unobservable
- ε_t: white noise with known distribution G_0
- θ: parameter of interest

Simulating Data from the Model

1. Specify initial values $z_0 = (y_0, u_0)$ and a guess θ
2. Set seed and draw shocks $\{\widetilde{\varepsilon}_t\}_{t=1}^T$ from G_0
3. Given $\{x_t\}$, for $t = 0, \ldots, T$, compute $\widetilde{y}_t(\theta, z_0)$, where

$$\begin{aligned}
\widetilde{y}_0(\theta, z_0) &= y_0 \\
\widetilde{y}_t(\theta, z_0) &= r\Big(\widetilde{y}_{t-1}(\theta, z_0), x_t, \widetilde{u}_t(\theta, u_0), \theta\Big) \\
\widetilde{u}_t(\theta, u_0) &= \varphi\Big(\widetilde{u}_{t-1}(\theta, u_0), \widetilde{\varepsilon}_t, \theta\Big)
\end{aligned}$$

Can we do MLE with simulated/synthetic data?

- Yes, can compute the conditional density of $\widetilde{y}_1, \ldots, \widetilde{y}_T$ given $z_0, x_1, \ldots, x_T$ and construct the conditional likelihood function
- But in practice the likelihood may be intractable
- Consider an alternative way of estimating θ

Auxiliary Model

The goal of indirect inference is to choose θ so that the observed data and the simulated data look the same from the vantage point of the chosen window (or auxiliary model)

- auxiliary model is parameterized by β

Back to the Example

Outcome vector from the data

$$y_n = \begin{pmatrix} \mathbf{1}(a_n = \text{small}) \\ \mathbf{1}(a_n = \text{medium}) \\ \mathbf{1}(a_n = \text{big}) \end{pmatrix}$$

Outcome vector from the model

$$\widetilde{y}_n(\theta) = \begin{pmatrix} \Psi(\text{small}|x_n, \theta) \\ \Psi(\text{medium}|x_n, \theta) \\ \Psi(\text{big}|x_n, \theta) \end{pmatrix}$$

Auxiliary model: linear probability (OLS)

$$y_n = Z_n\beta + u_n$$

- Z_n: indicators for the firm's current state; # competitors in a market; the log of construction employment in the county.

Structure of the presentation

- Connection with SMM

SMM is a special case of indirect inference

Main idea of indirect inference:

- fit data to an auxiliary model and obtain statistics $\widehat{\beta}, \widetilde{\beta}(\theta)$
- solve the following optimization problem

$$\min_{\theta \in \Theta} (\widehat{\beta} - \widetilde{\beta}(\theta)) W (\widehat{\beta} - \widetilde{\beta}(\theta))'$$

Recall SMM main steps

- moments from data $\widehat{m}$, e.g., mean, standard deviation, ratio
- moments from simulated data $\widetilde{m}(\theta)$
- solve the following optimization problem

$$\min_{\theta \in \Theta} (\widehat{m} - \widetilde{m}(\theta)) W (\widehat{m} - \widetilde{m}(\theta))'$$

Structure of the presentation

- Theoretical properties

Consistency

Under three conditions

1. an invertible relationship between θ and $\beta(\theta)$ (binding function)
2. $\widehat{\beta}$ converges to β_0
3. $\widetilde{\beta}(\theta)$ uniformly converges to $\beta(\theta)$:

$$\sup_{\theta \in \Theta} \|\widetilde{\beta}(\theta) - \beta(\theta)\| \xrightarrow{p} 0$$

Asymptotic normality

$$\sqrt{T}(\widehat{\theta} - \theta_0) \xrightarrow{d} \mathcal{N}\left(0, \left(1 + \frac{1}{H}\right)\left[\frac{\partial b'}{\partial \theta}(\theta_0) \mathcal{J}_0 I_0^{-1} \mathcal{J}_0 \frac{\partial b}{\partial \theta'}(\theta_0)\right]^{-1}\right)$$

Structure of the presentation

- Bias corrections (in panel data analysis)

Motivating Example: Neyman and Scott (1948)

Consider independent $\{X_{it}\}$, $i = 1, \dots, n$ and $t = 1, \dots, T$ where

$$X_{it} \sim \mathcal{N}(\alpha_{i0}, \theta_0)$$

Goal: estimate common variance θ_0 with unknown individual α_{i0}'s

FE estimator of common variance

- Fixed Effects Estimator:

$$\widehat{\alpha}_i = \frac{1}{T}\sum_{t=1}^{T} X_{it}, \quad \widehat{\theta}_{FE} = \frac{1}{nT}\sum_{i=1}^{n}\sum_{t=1}^{T}(X_{it} - \widehat{\alpha}_i)^2$$

- Incidental parameters: α_i's
  1. nuisance parameters
  2. dimension increases with n

Incidental parameter problem

- Let $n \to \infty$ while fix T, $\widehat{\theta}_{FE} \xrightarrow{p} \theta_T$, where

$$\theta_T = \theta_0 - \frac{\theta_0}{T}$$

- Source of the problem: each $\widehat{\alpha}_i$ is estimated using T observations, estimation errors don't vanish as $n \to \infty$ and T fixed

Indirect Inference FE Estimator (Chen, 2022)

- Given θ, simulate $\{X_{it}^h(\widehat{\alpha}_i, \theta), i = 1, \ldots, n; t = 1, \cdots, T\}$ as

$$X_{it}^h = \widehat{\alpha}_i + \sqrt{\theta} u_{it}^h, \quad u_{it}^h \sim \mathcal{N}(0, 1)$$

  h denotes a simulation path ($h = 1, \cdots, H$)
- From simulated data, can obtain $\widehat{\theta}_{FE}^h(\widehat{\alpha}_1, \cdots, \widehat{\alpha}_n, \theta)$

Indirect Inference FE Estimator (Chen, 2022)

- Given θ, simulate $\{X_{it}^h(\widehat{\alpha}_i, \theta), i = 1, \ldots, n; t = 1, \cdots, T\}$ as

$$X_{it}^h = \widehat{\alpha}_i + \sqrt{\theta} u_{it}^h, \quad u_{it}^h \sim \mathcal{N}(0, 1)$$

  h denotes a simulation path ($h = 1, \cdots, H$)
- From simulated data, can obtain $\widehat{\theta}_{FE}^h(\widehat{\alpha}_1, \cdots, \widehat{\alpha}_n, \theta)$
- If θ is close to θ_0, then $\{X_{it}(\alpha_{i0}, \theta_0)\}$ and $\{X_{it}^h(\widehat{\alpha}_i, \theta)\}$ should look similar in terms of statistics of data:

$$\widehat{\theta}_{FE}, \quad \widehat{\theta}_{FE}^h(\widehat{\alpha}_1, \cdots, \widehat{\alpha}_n, \theta)$$

- Calibrate θ so that

$$\widehat{\theta}_{FE} = \frac{1}{H} \sum_{h=1}^{H} \widehat{\theta}_{FE}^h(\widehat{\alpha}_1, \cdots, \widehat{\alpha}_n, \theta)$$

Monte Carlo for Neyman Scott

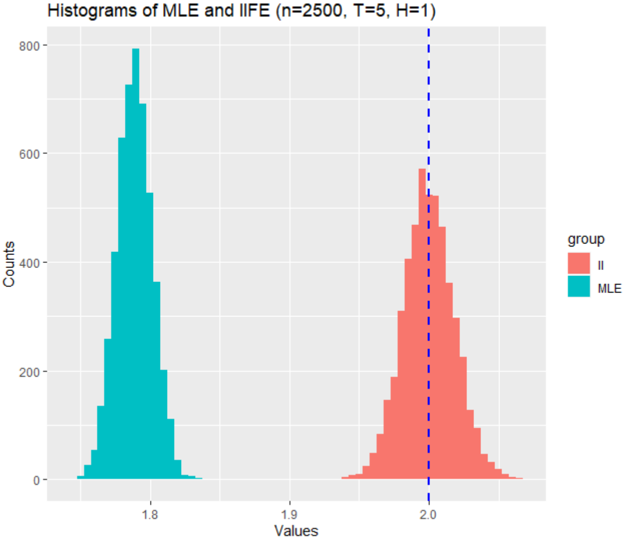

- $\widehat{\theta}_{FE}$ is severely biased
- $\widehat{\theta}_{II}$ centers around θ_0

Intuition for bias correction

- the auxiliary model is incorrectly specified
- whether $\widehat{\beta}$ is a correct estimator does not matter
- observed data and simulated data are fit to the same auxiliary model
- will mimic the bias structure