Duality and Estimation of Undiscounted MDPs

Nicola Rosaia

Equiprice Seminar

September 23, 2021

Why estimating undiscounted MDPs?

- Standard models with widespread applications
  - Engineering, operations management
  - i.e. people do use them... how do we estimate them?
- Standard approach in applications: impose calibrated discount factor β
  - Often impose β large but arbitrary
- Approximate discounted models as $\beta \to 1$
  - Can be analytically more convenient than their discounted counterparts

This paper

1. Convex duality framework for undiscounted MDPs with i.i.d. shocks
   - Primal problem: payoff system $\mapsto$ dynamic choice outcomes (computation)
   - Dual problem: dynamic choice outcomes $\mapsto$ payoff system (inversion)

This paper

1. Convex duality framework for undiscounted MDPs with i.i.d. shocks

- Primal problem: payoff system $\mapsto$ dynamic choice outcomes (computation)
- Dual problem: dynamic choice outcomes $\mapsto$ payoff system (inversion)

Idea: undiscounted MDP $\sim$ static choice over long-run state-action frequencies

- Static choice duality goes through

This paper

2. Implications

- Identification results: empirical content, identifying restrictions
- Novel inversion & estimation procedures

This paper

2. Implications

- Identification results: empirical content, identifying restrictions
- Novel inversion & estimation procedures

Not today

- Axiomatically characterize any undiscounted i.i.d. model
- Straightforward extensions
  - Mixed i.i.d. models
  - Models where certain actions and/or states are unobserved

Outline

1. Framework
2. Duality
3. Identification
4. Estimation

Framework

DDC Framework

The discounted model would be

$$V(x) = \mathrm{E}_F \max_{a \in A} [\boldsymbol{u}(a, x) + \epsilon(a) + \beta T(a, x) \cdot V]$$

$$\boldsymbol{\sigma}(a|x) = \Pr_F[a \in \arg\max_{a' \in A} [\boldsymbol{u}(a', x) + \epsilon(a') + \beta T(a', x) \cdot V]]$$

Assumptions.

- A and X are finite
- ***Conditional Independence***. $\Pr(x', \epsilon' | x, \epsilon, a) = \Pr(\epsilon' | x') \Pr(x' | x, a)$
- F is absolutely continuous with full support
- ***Accessibility***: $\nexists$ strict subset of states absorbing under all possible policies

$$\forall\ Y \subsetneq X\ \exists\ y \in Y, x \in X \setminus Y, a \in A \text{ s.t. } T(x|a, y) > 0$$

Optimality

Agents choose a ***stationary policy*** $\pi : x, \epsilon \mapsto a$

Optimality

Agents choose a ***stationary policy*** $\pi : x, \epsilon \mapsto a$. Define

- The (long-run) ***expected average payoff***

$$\boldsymbol{w}(\boldsymbol{\pi}, x_0) \equiv \lim_{T\to\infty} \frac{1}{T+1} \sum_{t=0}^{T} \mathrm{E}_{F,\boldsymbol{\pi}}[\boldsymbol{u}(a_t, x_t) + \epsilon(a_t) | x_0]$$

Optimality

Agents choose a ***stationary policy*** $\pi : x, \epsilon \mapsto a$. Define

- The (long-run) ***expected average payoff***

  $$\boldsymbol{w}(\pi, x_0) \equiv \lim_{T\to\infty} \frac{1}{T+1} \sum_{t=0}^{T} \mathrm{E}_{F,\pi}[\boldsymbol{u}(a_t, x_t) + \epsilon(a_t)|x_0]$$

- The (long-run) ***state-action frequencies***

  $$\mu(a, x|\pi, x_0) \equiv \lim_{T\to\infty} \frac{1}{T+1} \mathrm{E}_{\pi,F}[\sum_{t=0}^{T} 1\{a_t = a, x_t = x\}|x_0]$$

Optimality

Agents choose a ***stationary policy*** $\boldsymbol{\pi} : x, \epsilon \mapsto a$. Define

- The (long-run) ***expected average payoff***

  $$\boldsymbol{w}(\boldsymbol{\pi}, x_0) \equiv \lim_{T\to\infty} \frac{1}{T+1} \sum_{t=0}^{T} \mathrm{E}_{F,\boldsymbol{\pi}}[\boldsymbol{u}(a_t, x_t) + \epsilon(a_t)|x_0]$$

- The (long-run) ***state-action frequencies***

  $$\boldsymbol{\mu}(a, x|\boldsymbol{\pi}, x_0) \equiv \lim_{T\to\infty} \frac{1}{T+1} \mathrm{E}_{\boldsymbol{\pi},F}[\sum_{t=0}^{T} 1\{a_t = a, x_t = x\}|x_0]$$

- The CCP system $\boldsymbol{\sigma}(a, x|\boldsymbol{\pi}) = \Pr_F[\boldsymbol{\pi}(x, \epsilon) = a]$

Optimality

Agents choose a ***stationary policy*** $\boldsymbol{\pi} : x, \epsilon \mapsto a$. Define

- The (long-run) ***expected average payoff***

  $\boldsymbol{w}(\boldsymbol{\pi}, x_0) \equiv \lim_{T\to\infty} \frac{1}{T+1} \sum_{t=0}^{T} \mathrm{E}_{F,\boldsymbol{\pi}}[\boldsymbol{u}(a_t, x_t) + \epsilon(a_t)|x_0]$

- The (long-run) ***state-action frequencies***

  $\boldsymbol{\mu}(a, x|\boldsymbol{\pi}, x_0) \equiv \lim_{T\to\infty} \frac{1}{T+1} \mathrm{E}_{\boldsymbol{\pi},F}[\sum_{t=0}^{T} 1\{a_t = a, x_t = x\}|x_0]$

- The CCP system $\boldsymbol{\sigma}(a, x|\boldsymbol{\pi}) = \Pr_F[\boldsymbol{\pi}(x, \epsilon) = a]$

Definition. $\boldsymbol{\pi}$ is **optimal** if $\forall\ x_0$ it solves $\max_{\boldsymbol{\pi}} \boldsymbol{w}(\boldsymbol{\pi}|x_0)$

Definition. $\boldsymbol{w}(\boldsymbol{u}) \equiv \boldsymbol{w}(\pi|x_0)$, $\boldsymbol{\mu}(\boldsymbol{u}) \equiv \boldsymbol{\mu}(\pi, x_0)$, $\boldsymbol{\sigma}(\boldsymbol{u}) \equiv \boldsymbol{\sigma}(\pi)$ for some optimal $\boldsymbol{\pi}$

Duality

Static choice duality

Consider a static discrete choice model: $u \in \mathbb{R}^A$; $\sigma \in \Delta A$

Static choice duality

Consider a static discrete choice model: $u \in \mathbb{R}^A$; $\sigma \in \Delta A$

- u rationalizes σ if $\sigma(a) = \Pr_F[a \in \arg\max_{a'} [v(a') + \epsilon(a')]] \ \forall a$

Static choice duality

Consider a static discrete choice model: $u \in \mathbb{R}^A$; $\sigma \in \Delta A$

- u rationalizes σ if $\sigma(a) = \Pr_F[a \in \arg\max_{a'}[v(a') + \epsilon(a')]] \; \forall a$
- Define the inclusive value $w(u) = \mathrm{E}_F \max_a[u(a) + \epsilon(a)]$

Static choice duality

Consider a static discrete choice model: $u \in \mathbb{R}^A$; $\sigma \in \Delta A$

- u rationalizes σ if $\sigma(a) = \Pr_F[a \in \arg\max_{a'}[v(a') + \epsilon(a')]]\ \forall a$
- Define the inclusive value $w(u) = \mathrm{E}_F \max_a[u(a) + \epsilon(a)]$

Static choice duality

Consider a static discrete choice model: $u \in \mathbb{R}^A$; $\sigma \in \Delta A$

- u rationalizes σ if $\sigma(a) = \Pr_F[a \in \arg\max_{a'}[v(a') + \epsilon(a')]] \ \forall a$
- Define the inclusive value $w(u) = \mathrm{E}_F \max_a[u(a) + \epsilon(a)]$

Theorem (Chiong, Galichon & Shum 2016). TFAE:

1. u rationalizes σ
2. u solves $\max_{u \in \mathbb{R}^A}[\sigma \cdot u - w(u)]$ $(\equiv w^*(\sigma))$
3. σ solves $\max_{\sigma \in \Delta A}[\sigma \cdot v - w^*(\sigma)]$ $(= w(u))$

Static choice duality

Consider a static discrete choice model: $u \in \mathbb{R}^A$; $\sigma \in \Delta A$

- u rationalizes σ if $\sigma(a) = \Pr_F[a \in \arg\max_{a'}[v(a') + \epsilon(a')]]\ \forall a$
- Define the inclusive value $w(u) = \mathrm{E}_F \max_a[u(a) + \epsilon(a)]$

Theorem (Chiong, Galichon & Shum 2016). TFAE:

1. u rationalizes σ
2. u solves $\max_{u \in \mathbb{R}^A}[\sigma \cdot u - w(u)]$ $(\equiv w^*(\sigma))$
3. σ solves $\max_{\sigma \in \Delta A}[\sigma \cdot v - w^*(\sigma)]$ $(= w(u))$

Remark. $u \in \nabla w^*(\sigma)$ characterizes the identified set

Duality - intuition

Note that $\boldsymbol{w}(\boldsymbol{u}|\boldsymbol{\pi}, x_0) = \boldsymbol{\mu}(\boldsymbol{\pi}, x_0) \cdot \boldsymbol{u} + \sum_x \boldsymbol{\mu}_X(x|\boldsymbol{\pi}, x_0)\mathrm{E}_F[\epsilon(\boldsymbol{\pi}(x, \epsilon))]$

Duality - intuition

Note that $\boldsymbol{w}(\boldsymbol{u}|\boldsymbol{\pi}, x_0) = \boldsymbol{\mu}(\boldsymbol{\pi}, x_0) \cdot \boldsymbol{u} + \sum_x \boldsymbol{\mu}_X(x|\boldsymbol{\pi}, x_0)\mathrm{E}_F[\epsilon(\boldsymbol{\pi}(x, \epsilon))]$

Then can split optimization into nested problems

Duality - intuition

Note that $\boldsymbol{w}(\boldsymbol{u}|\boldsymbol{\pi}, x_0) = \boldsymbol{\mu}(\boldsymbol{\pi}, x_0) \cdot \boldsymbol{u} + \sum_x \boldsymbol{\mu}_X(x|\boldsymbol{\pi}, x_0) \mathrm{E}_F[\epsilon(\boldsymbol{\pi}(x, \epsilon))]$

Then can split optimization into nested problems

Let M be the set of all possible state-action frequencies

Let $\boldsymbol{\sigma}^{\boldsymbol{\mu}}(x) \in \Delta A$ be consistent with $\boldsymbol{\mu}$ at x (i.e. $\boldsymbol{\sigma}^{\boldsymbol{\mu}}(a, x) = \boldsymbol{\mu}(a, x)/\boldsymbol{\mu}_X(x)$)

Duality - intuition

Note that $\boldsymbol{w}(\boldsymbol{u}|\boldsymbol{\pi}, x_0) = \boldsymbol{\mu}(\boldsymbol{\pi}, x_0) \cdot \boldsymbol{u} + \sum_x \boldsymbol{\mu}_X(x|\boldsymbol{\pi}, x_0) \mathrm{E}_F[\epsilon(\boldsymbol{\pi}(x, \epsilon))]$

Then can split optimization into nested problems

Let M be the set of all possible state-action frequencies

Let $\boldsymbol{\sigma}^{\boldsymbol{\mu}}(x) \in \Delta A$ be consistent with $\boldsymbol{\mu}$ at x (i.e. $\boldsymbol{\sigma}^{\boldsymbol{\mu}}(a, x) = \boldsymbol{\mu}(a, x)/\boldsymbol{\mu}_X(x)$)

$$\max_{\boldsymbol{\pi}} \boldsymbol{w}(\boldsymbol{u}|\boldsymbol{\pi}, x_0)$$
$$= \max_{\boldsymbol{\mu} \in M}\{\boldsymbol{\mu} \cdot \boldsymbol{u} + \max_{\boldsymbol{\pi} \text{ inducing } \boldsymbol{\mu}}\{\textstyle\sum \boldsymbol{\mu}_X(x|\boldsymbol{\pi}, x_0) \mathrm{E}_F[\epsilon(\boldsymbol{\pi}(x, \epsilon))]\}$$

Duality - intuition

Note that $\boldsymbol{w}(\boldsymbol{u}|\boldsymbol{\pi}, x_0) = \boldsymbol{\mu}(\boldsymbol{\pi}, x_0) \cdot \boldsymbol{u} + \sum_x \boldsymbol{\mu}_X(x|\boldsymbol{\pi}, x_0) \mathrm{E}_F[\epsilon(\boldsymbol{\pi}(x, \epsilon))]$

Then can split optimization into nested problems

Let M be the set of all possible state-action frequencies

Let $\boldsymbol{\sigma}^{\boldsymbol{\mu}}(x) \in \Delta A$ be consistent with $\boldsymbol{\mu}$ at x (i.e. $\boldsymbol{\sigma}^{\boldsymbol{\mu}}(a, x) = \boldsymbol{\mu}(a, x)/\boldsymbol{\mu}_X(x)$)

$$\max_{\boldsymbol{\pi}} \boldsymbol{w}(\boldsymbol{u}|\boldsymbol{\pi}, x_0)$$

$$= \max_{\boldsymbol{\mu} \in M} \{\boldsymbol{\mu} \cdot \boldsymbol{u} + \max_{\boldsymbol{\pi} \text{ inducing } \boldsymbol{\mu}} \{\sum \boldsymbol{\mu}_X(x|\boldsymbol{\pi}, x_0) \mathrm{E}_F[\epsilon(\boldsymbol{\pi}(x, \epsilon))]\}$$

$$= \max_{\boldsymbol{\mu} \in M} \{\boldsymbol{\mu} \cdot \boldsymbol{u} + \sum_x \boldsymbol{\mu}_X(x) \max_{\pi:\, \mathbb{R}^A \to A \text{ inducing } \boldsymbol{\sigma}^{\boldsymbol{\mu}}(x)} \mathrm{E}_F[\epsilon(\pi(\epsilon))]\}$$

Duality - intuition

Note that $\boldsymbol{w}(\boldsymbol{u}|\boldsymbol{\pi}, x_0) = \boldsymbol{\mu}(\boldsymbol{\pi}, x_0) \cdot \boldsymbol{u} + \sum_x \boldsymbol{\mu}_X(x|\boldsymbol{\pi}, x_0) \mathrm{E}_F[\epsilon(\boldsymbol{\pi}(x,\epsilon))]$

Then can split optimization into nested problems

Let M be the set of all possible state-action frequencies

Let $\boldsymbol{\sigma}^{\boldsymbol{\mu}}(x) \in \Delta A$ be consistent with $\boldsymbol{\mu}$ at x (i.e. $\boldsymbol{\sigma}^{\boldsymbol{\mu}}(a,x) = \boldsymbol{\mu}(a,x)/\boldsymbol{\mu}_X(x)$)

$$\max_{\boldsymbol{\pi}} \boldsymbol{w}(\boldsymbol{u}|\boldsymbol{\pi}, x_0)$$

$$= \max_{\boldsymbol{\mu}\in M}\{\boldsymbol{\mu}\cdot\boldsymbol{u} + \max_{\boldsymbol{\pi} \text{ inducing } \boldsymbol{\mu}}\{\sum \boldsymbol{\mu}_X(x|\boldsymbol{\pi}, x_0)\mathrm{E}_F[\epsilon(\boldsymbol{\pi}(x,\epsilon))]\}$$

$$= \max_{\boldsymbol{\mu}\in M}\{\boldsymbol{\mu}\cdot\boldsymbol{u} + \sum_x \boldsymbol{\mu}_X(x) \max_{\pi:\, \mathbb{R}^A \to A \text{ inducing } \boldsymbol{\sigma}^{\boldsymbol{\mu}}(x)} \mathrm{E}_F[\epsilon(\pi(\epsilon))]\}$$

$$\text{(Galichon \& Salanie 2020)} = \max_{\boldsymbol{\mu}\in M}\{\boldsymbol{\mu}\cdot\boldsymbol{u} - \sum_x \boldsymbol{\mu}_X(x) w^*(\boldsymbol{\sigma}^{\boldsymbol{\mu}}(x))\}$$

Duality - statement

Theorem. Define $\boldsymbol{w}^*(\boldsymbol{\mu}) \equiv \sum_x \mu_X(x) w^*(\boldsymbol{\sigma}^{\boldsymbol{\mu}}(x))$. TFAE:

1. $\boldsymbol{\mu} = \boldsymbol{\mu}(\boldsymbol{u})$
2. $\boldsymbol{\mu}$ solves $\max_{\boldsymbol{\mu} \in M} [\boldsymbol{\mu} \cdot \boldsymbol{u} - \boldsymbol{w}^*(\boldsymbol{\mu})]$ $(= \boldsymbol{w}(\boldsymbol{u}))$
3. $\boldsymbol{u}$ solves $\max_{\boldsymbol{u} \in \mathbb{R}^{|A||X|}} [\boldsymbol{\mu} \cdot \boldsymbol{u} - \boldsymbol{w}(\boldsymbol{u})]$ $(= \boldsymbol{w}^*(\boldsymbol{\mu}))$

$\boldsymbol{w}$ and $\boldsymbol{w}^*$ are both convex and C^1; $\boldsymbol{w}^*$ is strictly convex

Duality - statement

Theorem. Define $\boldsymbol{w}^*(\boldsymbol{\mu}) \equiv \sum_x \mu_X(x) w^*(\boldsymbol{\sigma}^{\boldsymbol{\mu}}(x))$. TFAE:

1. $\boldsymbol{\mu} = \boldsymbol{\mu}(\boldsymbol{u})$
2. $\boldsymbol{\mu}$ solves $\max_{\boldsymbol{\mu} \in M} [\boldsymbol{\mu} \cdot \boldsymbol{u} - \boldsymbol{w}^*(\boldsymbol{\mu})]$ $(= \boldsymbol{w}(\boldsymbol{u}))$
3. $\boldsymbol{u}$ solves $\max_{\boldsymbol{u} \in \mathbb{R}^{|A||X|}} [\boldsymbol{\mu} \cdot \boldsymbol{u} - \boldsymbol{w}(\boldsymbol{u})]$ $(= \boldsymbol{w}^*(\boldsymbol{\mu}))$

$\boldsymbol{w}$ and $\boldsymbol{w}^*$ are both convex and C^1; $\boldsymbol{w}^*$ is strictly convex

Corollary. $\boldsymbol{\mu} = \boldsymbol{\mu}(\boldsymbol{u}) \Leftrightarrow \exists\, k \in \mathbb{R}$ s.t. $\boldsymbol{\nu} \cdot \boldsymbol{u} = \boldsymbol{\nu} \cdot \nabla \boldsymbol{w}^*(\boldsymbol{\mu}) + k \;\forall\; \boldsymbol{\nu} \in M$

Duality - statement

Theorem. Define $\boldsymbol{w}^*(\boldsymbol{\mu}) \equiv \sum_x \mu_X(x) w^*(\boldsymbol{\sigma}^{\boldsymbol{\mu}}(x))$. TFAE:

1. $\boldsymbol{\mu} = \boldsymbol{\mu}(\boldsymbol{u})$
2. $\boldsymbol{\mu}$ solves $\max_{\boldsymbol{\mu} \in M} [\boldsymbol{\mu} \cdot \boldsymbol{u} - \boldsymbol{w}^*(\boldsymbol{\mu})]$ $(= \boldsymbol{w}(\boldsymbol{u}))$
3. $\boldsymbol{u}$ solves $\max_{\boldsymbol{u} \in \mathbb{R}^{|A||X|}} [\boldsymbol{\mu} \cdot \boldsymbol{u} - \boldsymbol{w}(\boldsymbol{u})]$ $(= \boldsymbol{w}^*(\boldsymbol{\mu}))$

$\boldsymbol{w}$ and $\boldsymbol{w}^*$ are both convex and C^1; $\boldsymbol{w}^*$ is strictly convex

Corollary. $\boldsymbol{\mu} = \boldsymbol{\mu}(\boldsymbol{u}) \Leftrightarrow \exists\ k \in \mathbb{R}$ s.t. $\boldsymbol{\nu} \cdot \boldsymbol{u} = \boldsymbol{\nu} \cdot \nabla \boldsymbol{w}^*(\boldsymbol{\mu}) + k\ \forall\ \boldsymbol{\nu} \in M$

In words: $\boldsymbol{\mu}$ identifies the average payoff of each strategy up to a constant

Today's focus

Focus on implications for estimation/inversion

i.e. from observed choices ($\boldsymbol{\mu}$) to primitives ($\boldsymbol{u}$)

Remark. $\boldsymbol{\mu}$ assumed known to the analyst - same as knowing $\boldsymbol{\sigma}$

Today's focus

Focus on implications for estimation/inversion

i.e. from observed choices ($\boldsymbol{\mu}$) to primitives ($\boldsymbol{u}$)

Remark. $\boldsymbol{\mu}$ assumed known to the analyst - same as knowing $\boldsymbol{\sigma}$

Computation of $\boldsymbol{\mu}(\boldsymbol{u})$ given $\boldsymbol{u}$ is well studied for case without shocks

The paper has a small result for the case with shocks

Identification of payoffs

Empirical content of aggregate behavior

Definition. π is ***pure*** if $\pi(x,\epsilon) = \pi(x,\epsilon') \; \forall \; \epsilon, \epsilon' \in \mathbb{R}^A$

$\mu \in M_0$ if $\mu = \mu(\pi, x_0)$ for some x_0 and pure π

Proposition. M_0 is a linear basis of M

Empirical content of aggregate behavior

Definition. $\boldsymbol{\pi}$ is ***pure*** if $\boldsymbol{\pi}(x,\epsilon) = \boldsymbol{\pi}(x,\epsilon') \; \forall \; \epsilon, \epsilon' \in \mathbb{R}^A$

$\boldsymbol{\mu} \in M_0$ if $\boldsymbol{\mu} = \boldsymbol{\mu}(\boldsymbol{\pi}, x_0)$ for some x_0 and pure $\boldsymbol{\pi}$

Proposition. M_0 is a linear basis of M

Corollary. $\boldsymbol{\mu} = \boldsymbol{\mu}(\boldsymbol{u}) \Leftrightarrow \exists \; k \in \mathbb{R}$ s.t. $\boldsymbol{\nu} \cdot \boldsymbol{u} = \boldsymbol{\nu} \cdot \nabla \boldsymbol{w}^*(\boldsymbol{\mu}) + k \; \forall \; \boldsymbol{\nu} \in M_0$

Empirical content of aggregate behavior

Definition. π is ***pure*** if $\pi(x,\epsilon) = \pi(x,\epsilon') \; \forall \; \epsilon,\epsilon' \in \mathbb{R}^A$

$$\boldsymbol{\mu} \in M_0 \text{ if } \boldsymbol{\mu} = \boldsymbol{\mu}(\pi, x_0) \text{ for some } x_0 \text{ and pure } \pi$$

Proposition. M_0 is a linear basis of M

Corollary. $\boldsymbol{\mu} = \boldsymbol{\mu}(\boldsymbol{u}) \Leftrightarrow \exists \; k \in \mathbb{R}$ s.t. $\boldsymbol{\nu} \cdot \boldsymbol{u} = \boldsymbol{\nu} \cdot \nabla \boldsymbol{w}^*(\boldsymbol{\mu}) + k \; \forall \; \boldsymbol{\nu} \in M_0$

$$M_0' \boldsymbol{u} = M_0' \nabla \boldsymbol{w}^*(\boldsymbol{\mu}) + k$$

Empirical content of aggregate behavior

Definition. π is ***pure*** if $\pi(x, \epsilon) = \pi(x, \epsilon') \;\forall\; \epsilon, \epsilon' \in \mathbb{R}^A$

$$\boldsymbol{\mu} \in M_0 \text{ if } \boldsymbol{\mu} = \boldsymbol{\mu}(\pi, x_0) \text{ for some } x_0 \text{ and pure } \pi$$

Proposition. M_0 is a linear basis of M

Corollary. $\boldsymbol{\mu} = \boldsymbol{\mu}(\boldsymbol{u}) \Leftrightarrow \exists\; k \in \mathbb{R}$ s.t. $\boldsymbol{\nu} \cdot \boldsymbol{u} = \boldsymbol{\nu} \cdot \nabla \boldsymbol{w}^*(\boldsymbol{\mu}) + k \;\forall\; \boldsymbol{\nu} \in M_0$

$$M_0' \boldsymbol{u} = M_0' \nabla \boldsymbol{w}^*(\boldsymbol{\mu}) + k$$

In words: $\boldsymbol{\mu}$ identifies the average payoff of each pure strategy up to a constant

Identifying restrictions

Let C represent a set of linear restrictions on $\boldsymbol{u}$ (e.g. $C'\boldsymbol{u} = 0$)

Definition. C ***identifies*** $\boldsymbol{u}$ if $\mu(\boldsymbol{u}) = \mu(\boldsymbol{v}) \wedge C'\boldsymbol{u} = C'\boldsymbol{v} \Rightarrow \boldsymbol{u} = \boldsymbol{v}$

Identifying restrictions

Let C represent a set of linear restrictions on $\boldsymbol{u}$ (e.g. $C'\boldsymbol{u} = 0$)

Definition. C ***identifies*** $\boldsymbol{u}$ if $\mu(\boldsymbol{u}) = \mu(\boldsymbol{v}) \wedge C'\boldsymbol{u} = C'\boldsymbol{v} \Rightarrow \boldsymbol{u} = \boldsymbol{v}$

if $M_0'\boldsymbol{u} = M_0'\boldsymbol{v} + k \wedge C'\boldsymbol{u} = C'\boldsymbol{v} \Rightarrow \boldsymbol{u} = \boldsymbol{v}$

Identifying restrictions

Let C represent a set of linear restrictions on $\boldsymbol{u}$ (e.g. $C'\boldsymbol{u} = 0$)

Definition. C ***identifies*** $\boldsymbol{u}$ if $\mu(\boldsymbol{u}) = \mu(\boldsymbol{v}) \wedge C'\boldsymbol{u} = C'\boldsymbol{v} \Rightarrow \boldsymbol{u} = \boldsymbol{v}$

if $M_0'\boldsymbol{u} = M_0'\boldsymbol{v} + k \wedge C'\boldsymbol{u} = C'\boldsymbol{v} \Rightarrow \boldsymbol{u} = \boldsymbol{v}$

Proposition. $|M_0| = (|A| - 1)\,|X| + 1$

Identifying restrictions

Let C represent a set of linear restrictions on $\boldsymbol{u}$ (e.g. $C'\boldsymbol{u} = 0$)

Definition. C ***identifies*** $\boldsymbol{u}$ if $\mu(\boldsymbol{u}) = \mu(\boldsymbol{v}) \wedge C'\boldsymbol{u} = C'\boldsymbol{v} \Rightarrow \boldsymbol{u} = \boldsymbol{v}$

if $M_0'\boldsymbol{u} = M_0'\boldsymbol{v} + k \wedge C'\boldsymbol{u} = C'\boldsymbol{v} \Rightarrow \boldsymbol{u} = \boldsymbol{v}$

Proposition. $|M_0| = (|A| - 1)\,|X| + 1$

Corollary. TFAE:

1. C identifies $\boldsymbol{u}$

Identifying restrictions

Let C represent a set of linear restrictions on $\boldsymbol{u}$ (e.g. $C'\boldsymbol{u} = 0$)

Definition. C ***identifies*** $\boldsymbol{u}$ if $\mu(\boldsymbol{u}) = \mu(\boldsymbol{v}) \wedge C'\boldsymbol{u} = C'\boldsymbol{v} \Rightarrow \boldsymbol{u} = \boldsymbol{v}$

if $M_0'\boldsymbol{u} = M_0'\boldsymbol{v} + k \wedge C'\boldsymbol{u} = C'\boldsymbol{v} \Rightarrow \boldsymbol{u} = \boldsymbol{v}$

Proposition. $|M_0| = (|A| - 1)\,|X| + 1$

Corollary. TFAE:

1. C identifies $\boldsymbol{u}$
2. i) $\exists \{c^1, ..., c^{|X|-1}\} \subseteq C$ s.t. $\text{Span}\{c^1, ..., c^{|X|-1}\} \cap \text{Span}\, M = \{0\}$

   ii) $\exists\, \boldsymbol{\nu} \in \text{Span}\, C \cap \text{Span}\, M$ s.t. $\sum_{a,x} \nu(a,x) \neq 0$

Identifying restrictions

Let C represent a set of linear restrictions on $\boldsymbol{u}$ (e.g. $C'\boldsymbol{u} = 0$)

Definition. C ***identifies*** $\boldsymbol{u}$ if $\mu(\boldsymbol{u}) = \mu(\boldsymbol{v}) \wedge C'\boldsymbol{u} = C'\boldsymbol{v} \Rightarrow \boldsymbol{u} = \boldsymbol{v}$

if $M_0'\boldsymbol{u} = M_0'\boldsymbol{v} + k \wedge C'\boldsymbol{u} = C'\boldsymbol{v} \Rightarrow \boldsymbol{u} = \boldsymbol{v}$

Proposition. $|M_0| = (|A| - 1)\,|X| + 1$

Corollary. TFAE:

1. C identifies $\boldsymbol{u}$
2. i) $\exists\ \{c^1, ..., c^{|X|-1}\} \subseteq C$ s.t. $\text{Span}\{c^1, ..., c^{|X|-1}\} \cap \text{Span}\, M = \{0\}$

   ii) $\exists\ \boldsymbol{\nu} \in \text{Span}\, C \cap \text{Span}\, M$ s.t. $\sum_{a,x} \boldsymbol{\nu}(a, x) \neq 0$
3. $\begin{bmatrix} M_0' & 1_{|M_0|} \\ C' & 0_{|C|} \end{bmatrix}$ has full column rank

Example: normalizing one payoff at each state

Consider imposing that, for some a, $\boldsymbol{u}(a, x) = 0$ for every x

Example: normalizing one payoff at each state

Consider imposing that, for some a, $\boldsymbol{u}(a,x) = 0$ for every x

Corollary. This identifies $\boldsymbol{u}$ if and only if one can find a state x such that

$$\forall Y \subseteq X \setminus \{x\} \; \exists \; y \in Y \text{ s.t. } T(Y|a,y) < 1$$

$\exists \; x$ reachable from any $x' \neq x$ under the state-transitions generated by a

Intuition. $\{\boldsymbol{u}(a,x') = 0 : \; x' \neq x\}$ does not span any stationary measure

Example: normalizing one payoff at each state

Consider imposing that, for some a, $\boldsymbol{u}(a,x) = 0$ for every x

Corollary. This identifies $\boldsymbol{u}$ if and only if one can find a state x such that

$$\forall Y \subseteq X \setminus \{x\} \; \exists \; y \in Y \text{ s.t. } \; T(Y|a,y) < 1$$

$\exists \; x$ reachable from any $x' \neq x$ under the state-transitions generated by a

Intuition. $\{\boldsymbol{u}(a,x') = 0 : \; x' \neq x\}$ does not span any stationary measure

Remark. Conditions for identification for discounted models are less strong

Estimation

Parametric restrictions

Equivalently, can write restrictions in parametric form: $\exists$ unknown θ s.t. $\boldsymbol{u} = C\theta$

Definition. C ***identifies*** θ if $\mu(C\theta) = \mu(C\delta) \Rightarrow \theta = \delta$

Parametric restrictions

Equivalently, can write restrictions in parametric form: $\exists$ unknown θ s.t. $\boldsymbol{u} = C\theta$

Definition. C ***identifies*** θ if $\mu(C\theta) = \mu(C\delta) \Rightarrow \theta = \delta$

Corollary. C identifies θ if and only if $\begin{bmatrix} M_0'C & \mathbf{1}_{|M_0|\times 1} \end{bmatrix}$ has full column rank

Parametric restrictions

Equivalently, can write restrictions in parametric form: $\exists$ unknown θ s.t. $\boldsymbol{u} = C\theta$

Definition. C ***identifies*** θ if $\boldsymbol{\mu}(C\theta) = \boldsymbol{\mu}(C\delta) \Rightarrow \theta = \delta$

Corollary. C identifies θ if and only if $\begin{bmatrix} M_0'C & \mathbf{1}_{|M_0|\times 1} \end{bmatrix}$ has full column rank

Definition. C ***just-identifies*** θ if C identifies θ and

$$\forall\ \boldsymbol{\mu} \in M \cap \mathbb{R}_{++}^{|A||X|}\ \exists\ \theta \text{ s.t. } \boldsymbol{\mu}(C\theta) = \boldsymbol{\mu}$$

Parametric restrictions

Equivalently, can write restrictions in parametric form: $\exists$ unknown θ s.t. $\boldsymbol{u} = C\theta$

Definition. C ***identifies*** θ if $\boldsymbol{\mu}(C\theta) = \boldsymbol{\mu}(C\delta) \Rightarrow \theta = \delta$

Corollary. C identifies θ if and only if $\begin{bmatrix} M_0' C & \mathbf{1}_{|M_0| \times 1} \end{bmatrix}$ has full column rank

Definition. C ***just-identifies*** θ if C identifies θ and

$$\forall\ \boldsymbol{\mu} \in M \cap \mathbb{R}_{++}^{|A||X|}\ \exists\ \theta \text{ s.t. } \boldsymbol{\mu}(C\theta) = \boldsymbol{\mu}$$

Corollary. C just-identifies θ if and only if C identifies θ and $|\theta| = (|A| - 1)\,|X|$

Inversion

Take C s.t. C just-identifies θ

Problem. Given $\boldsymbol{\mu} \in M \cap \mathbb{R}_{++}^{|A||X|}$ find the unique θ s.t. $\boldsymbol{\mu}(C\theta) = \boldsymbol{\mu}$

Inversion

Take C s.t. C just-identifies θ

Problem. Given $\boldsymbol{\mu} \in M \cap \mathbb{R}_{++}^{|A||X|}$ find the unique θ s.t. $\boldsymbol{\mu}(C\theta) = \boldsymbol{\mu}$

Assumption. F is s.t. $\forall\ a \neq a'$ the density of $\epsilon_{a'} - \epsilon_a$ is bounded above

Two alternative inversion algorithms based on static v.s. dynamic duality

Inversion

Take C s.t. C just-identifies θ

Problem. Given $\boldsymbol{\mu} \in M \cap \mathbb{R}_{++}^{|A||X|}$ find the unique θ s.t. $\boldsymbol{\mu}(C\theta) = \boldsymbol{\mu}$

Assumption. F is s.t. $\forall\ a \neq a'$ the density of $\epsilon_{a'} - \epsilon_a$ is bounded above

Two alternative inversion algorithms based on static v.s. dynamic duality

When θ is over-identified, they suggest alternative two-steps estimators

Inversion from static duality

1. Exploit static duality to compute $\nabla \boldsymbol{w}^*(\boldsymbol{\mu})$ (compare with CGS 2016)

Inversion from static duality

1. Exploit static duality to compute $\nabla \boldsymbol{w}^*(\boldsymbol{\mu})$ (compare with CGS 2016)

Proposition. $\nabla \boldsymbol{w}^*(\boldsymbol{\mu})$ is the unique $\boldsymbol{u} \in \mathbb{R}^{|A||X|}$ s.t. $\forall\ x$

$$\boldsymbol{u}(x) \in \arg\max_{u \in \mathbb{R}^{|A|}} [\boldsymbol{\sigma}^{\boldsymbol{\mu}}(x) \cdot u - w(u)] \text{ and } w(\boldsymbol{u}(x)) = 0$$

Inversion from static duality

1. Exploit static duality to compute $\nabla \boldsymbol{w}^*(\boldsymbol{\mu})$ (compare with CGS 2016)

Proposition. $\nabla \boldsymbol{w}^*(\boldsymbol{\mu})$ is the unique $\boldsymbol{u} \in \mathbb{R}^{|A||X|}$ s.t. $\forall\, x$

$$\boldsymbol{u}(x) \in \arg\max_{u \in \mathbb{R}^{|A|}} [\boldsymbol{\sigma}^{\boldsymbol{\mu}}(x) \cdot u - w(u)] \text{ and } w(\boldsymbol{u}(x)) = 0$$

$\boldsymbol{u}(x)$ rationalizes $\boldsymbol{\sigma}^{\boldsymbol{\mu}}(x)$ (in the static sense)

Inversion from static duality

1. Exploit static duality to compute $\nabla \boldsymbol{w}^*(\boldsymbol{\mu})$ (compare with CGS 2016)

Proposition. $\nabla \boldsymbol{w}^*(\boldsymbol{\mu})$ is the unique $\boldsymbol{u} \in \mathbb{R}^{|A||X|}$ s.t. $\forall\ x$

$$\boldsymbol{u}(x) \in \arg\max_{u \in \mathbb{R}^{|A|}} [\boldsymbol{\sigma}^{\boldsymbol{\mu}}(x) \cdot u - w(u)] \text{ and } w(\boldsymbol{u}(x)) = 0$$

$$\boldsymbol{u}(x) \text{ rationalizes } \boldsymbol{\sigma}^{\boldsymbol{\mu}}(x) \text{ (in the static sense)}$$

Algorithm. $\forall\ x$ gradient descent with constant step size γ is

$$u^{n+1} = u^n + \gamma[\boldsymbol{\sigma}^{\boldsymbol{\mu}}(x) - \sigma^n] - w(u^n + \gamma[\boldsymbol{\sigma}^{\boldsymbol{\mu}}(x) - \sigma^n])$$

$$\text{where } \sigma^n(a) = \mathrm{P}_F[u^n(a) + \epsilon(a) = \max_{a'}[u^n(a') + \epsilon(a')]]$$

Inversion from static duality

1. Exploit static duality to compute $\nabla \boldsymbol{w}^*(\boldsymbol{\mu})$ (compare with CGS 2016)

Proposition. $\nabla \boldsymbol{w}^*(\boldsymbol{\mu})$ is the unique $\boldsymbol{u} \in \mathbb{R}^{|A||X|}$ s.t. $\forall\ x$

$$\boldsymbol{u}(x) \in \arg\max_{u \in \mathbb{R}^{|A|}} [\boldsymbol{\sigma}^{\boldsymbol{\mu}}(x) \cdot u - w(u)] \text{ and } w(\boldsymbol{u}(x)) = 0$$

$$\boldsymbol{u}(x) \text{ rationalizes } \boldsymbol{\sigma}^{\boldsymbol{\mu}}(x) \text{ (in the static sense)}$$

Algorithm. $\forall\ x$ gradient descent with constant step size γ is

$$u^{n+1} = u^n + \gamma[\boldsymbol{\sigma}^{\boldsymbol{\mu}}(x) - \sigma^n] - w(u^n + \gamma[\boldsymbol{\sigma}^{\boldsymbol{\mu}}(x) - \sigma^n])$$

$$\text{where } \sigma^n(a) = \mathrm{P}_F[u^n(a) + \epsilon(a) = \max_{a'}[u^n(a') + \epsilon(a')]]$$

Proposition. For γ small enough, u^n converges linearly to $\nabla \boldsymbol{w}^*(\boldsymbol{\mu})(x)$

Proof is standard (can prove that w is smooth)

Estimation from static duality

Given $\nabla \boldsymbol{w}^*(\boldsymbol{\mu})$, get θ (and k) from $M_0' C\theta = M_0' \nabla \boldsymbol{w}^*(\boldsymbol{\mu}) + k$

Estimation from static duality

Given $\nabla \boldsymbol{w}^*(\boldsymbol{\mu})$, get θ (and k) from $M_0' C\theta = M_0' \nabla \boldsymbol{w}^*(\boldsymbol{\mu}) + k$

If C over-identifies θ, solution might not exist

Estimation from static duality

Given $\nabla \boldsymbol{w}^*(\boldsymbol{\mu})$, get θ (and k) from $M_0' C\theta = M_0' \nabla \boldsymbol{w}^*(\boldsymbol{\mu}) + k$

If C over-identifies θ, solution might not exist

Consider $\hat{\theta}$ solving $\min_{\theta \in \mathbb{R}^{|C|}, \boldsymbol{u} \in \mathbb{R}^{|A||X|}} \|\boldsymbol{u} - C\theta\|^2$ s.t. $\boldsymbol{\mu}(\boldsymbol{u}) = \boldsymbol{\mu}$

i.e. minimize distance between $\boldsymbol{\mu}$ and $\boldsymbol{\mu}(C\theta)$ in the space of payoffs

Estimation from static duality

Given $\nabla \boldsymbol{w}^*(\boldsymbol{\mu})$, get θ (and k) from $M_0' C\theta = M_0' \nabla \boldsymbol{w}^*(\boldsymbol{\mu}) + k$

If C over-identifies θ, solution might not exist

Consider $\hat{\theta}$ solving $\min_{\theta \in \mathbb{R}^{|C|}, \boldsymbol{u} \in \mathbb{R}^{|A||X|}} \|\boldsymbol{u} - C\theta\|^2$ s.t. $\boldsymbol{\mu}(\boldsymbol{u}) = \boldsymbol{\mu}$

i.e. minimize distance between $\boldsymbol{\mu}$ and $\boldsymbol{\mu}(C\theta)$ in the space of payoffs

This is equivalent to the linear Constrained least squares

$\min_{\theta \in \mathbb{R}^{|C|}, \boldsymbol{u} \in \mathbb{R}^{|A||X|}, k \in \mathbb{R}} \|\boldsymbol{u} - C\theta\|^2$ s.t. $M_0' \boldsymbol{u} = M_0' \nabla \boldsymbol{w}^*(\boldsymbol{\mu}) + k$

which admits a closed form solution

Inversion from dynamic duality

2. Directly compute the unique solution $\hat{\theta}$ of $\max_{\theta \in \mathbb{R}^{|C|}} [\boldsymbol{\mu} \cdot C\theta - \boldsymbol{w}(C\theta)]$

   i.e. solve the dual imposing the identifying restrictions $\boldsymbol{u} = C\theta$

   If θ is just-identified then $\boldsymbol{\mu} = \boldsymbol{\mu}(C\hat{\theta})$

Inversion from dynamic duality

2. Directly compute the unique solution $\hat{\theta}$ of $\max_{\theta \in \mathbb{R}^{|C|}} [\boldsymbol{\mu} \cdot C\theta - \boldsymbol{w}(C\theta)]$

i.e. solve the dual imposing the identifying restrictions $\boldsymbol{u} = C\theta$

If θ is just-identified then $\boldsymbol{\mu} = \boldsymbol{\mu}(C\hat{\theta})$

Algorithm. Gradient descent with constant step size: $\theta^{n+1} = \theta^n + \gamma C'[\boldsymbol{\mu} - \boldsymbol{\mu}(C\theta^n)]$

Proposition. For γ small enough, θ^n converges linearly to $\hat{\theta}$

Inversion from dynamic duality

2. Directly compute the unique solution $\hat{\theta}$ of $\max_{\theta \in \mathbb{R}^{|C|}} [\boldsymbol{\mu} \cdot C\theta - \boldsymbol{w}(C\theta)]$

   i.e. solve the dual imposing the identifying restrictions $\boldsymbol{u} = C\theta$

   If θ is just-identified then $\boldsymbol{\mu} = \boldsymbol{\mu}(C\hat{\theta})$

Algorithm. Gradient descent with constant step size: $\theta^{n+1} = \theta^n + \gamma C'[\boldsymbol{\mu} - \boldsymbol{\mu}(C\theta^n)]$

Proposition. For γ small enough, θ^n converges linearly to $\hat{\theta}$

Proof idea. $\boldsymbol{w}$ is not smooth.

Inversion from dynamic duality

2. Directly compute the unique solution $\hat{\theta}$ of $\max_{\theta \in \mathbb{R}^{|C|}} [\boldsymbol{\mu} \cdot C\theta - \boldsymbol{w}(C\theta)]$

   i.e. solve the dual imposing the identifying restrictions $\boldsymbol{u} = C\theta$

   If θ is just-identified then $\boldsymbol{\mu} = \boldsymbol{\mu}(C\hat{\theta})$

Algorithm. Gradient descent with constant step size: $\theta^{n+1} = \theta^n + \gamma C'[\boldsymbol{\mu} - \boldsymbol{\mu}(C\theta^n)]$

Proposition. For γ small enough, θ^n converges linearly to $\hat{\theta}$

Proof idea. $\boldsymbol{w}$ is not smooth.

1. Pick γ so that $\boldsymbol{\mu}_X(C\theta^n)$ remain bounded away from zero

Inversion from dynamic duality

2. Directly compute the unique solution $\hat{\theta}$ of $\max_{\theta \in \mathbb{R}^{|C|}} [\boldsymbol{\mu} \cdot C\theta - \boldsymbol{w}(C\theta)]$

i.e. solve the dual imposing the identifying restrictions $\boldsymbol{u} = C\theta$

If θ is just-identified then $\boldsymbol{\mu} = \boldsymbol{\mu}(C\hat{\theta})$

Algorithm. Gradient descent with constant step size: $\theta^{n+1} = \theta^n + \gamma C'[\boldsymbol{\mu} - \boldsymbol{\mu}(C\theta^n)]$

Proposition. For γ small enough, θ^n converges linearly to $\hat{\theta}$

Proof idea. $\boldsymbol{w}$ is not smooth.

1. Pick γ so that $\mu_X(C\theta^n)$ remain bounded away from zero
2. Then $\left\| \boldsymbol{\mu}(C\theta^{n+1}) - \boldsymbol{\mu}(C\theta^n) \right\|_2$ bounded by factor of $\left\| \boldsymbol{\sigma}(C\theta^{n+1}) - \boldsymbol{\sigma}(C\theta^n) \right\|_2$

Inversion from dynamic duality

2. Directly compute the unique solution $\hat{\theta}$ of $\max_{\theta \in \mathbb{R}^{|C|}} [\boldsymbol{\mu} \cdot C\theta - \boldsymbol{w}(C\theta)]$

   i.e. solve the dual imposing the identifying restrictions $\boldsymbol{u} = C\theta$

   If θ is just-identified then $\boldsymbol{\mu} = \boldsymbol{\mu}(C\hat{\theta})$

Algorithm. Gradient descent with constant step size: $\theta^{n+1} = \theta^n + \gamma C'[\boldsymbol{\mu} - \boldsymbol{\mu}(C\theta^n)]$

Proposition. For γ small enough, θ^n converges linearly to $\hat{\theta}$

Proof idea. $\boldsymbol{w}$ is not smooth.

1. Pick γ so that $\boldsymbol{\mu}_X(C\theta^n)$ remain bounded away from zero
2. Then $\left\|\boldsymbol{\mu}(C\theta^{n+1}) - \boldsymbol{\mu}(C\theta^n)\right\|_2$ bounded by factor of $\left\|\boldsymbol{\sigma}(C\theta^{n+1}) - \boldsymbol{\sigma}(C\theta^n)\right\|_2$
3. Then smoothness of w yields the progress bounds

Bregman projection

If θ is over-identified can show that $\hat{\theta}$ solves $\min_\theta D_{\boldsymbol{w}^*}(\mu, \mu(C\theta))$

$D_{\boldsymbol{w}^*}$ is ***Bregman divergence*** associated with $\boldsymbol{w}^*$: $D_{\boldsymbol{w}^*}(\mu, \nu) = \boldsymbol{w}^*(\mu) - \mu \cdot \nabla \boldsymbol{w}^*(\nu)$

Bregman projection

If θ is over-identified can show that $\hat{\theta}$ solves $\min_\theta D_{\boldsymbol{w}^*}(\mu, \mu(C\theta))$

$D_{\boldsymbol{w}^*}$ is ***Bregman divergence*** associated with $\boldsymbol{w}^*$: $D_{\boldsymbol{w}^*}(\mu, \nu) = \boldsymbol{w}^*(\mu) - \mu \cdot \nabla \boldsymbol{w}^*(\nu)$

Example. When F is logit, $D_{\boldsymbol{w}^*}$ is the Kullback–Leibler divergence

If $\mu(a, x) = \sum_{i=1}^N \frac{1\{a_i = a, x_i = x\}}{N}$ then $\hat{\theta}$ is Maximum Likelihood estimator

Further results (sketch)

Independence of Irrelevant Alternatives

Classic IIA for static discrete choice:

Relative frequency of choosing two alternatives is independent of the choice set

Luce (1959). IIA $\Leftrightarrow$ static choice rationalized by the logit model

Independence of Irrelevant Alternatives

Classic IIA for static discrete choice:

Relative frequency of choosing two alternatives is independent of the choice set

Luce (1959). IIA $\Leftrightarrow$ static choice rationalized by the logit model

For given F, formulate dynamic analogue of IIA:

Relative distance of any pair of state-action frequencies from the observed one
independent of which other state-action frequencies are available
where distance is Bregman divergence associated with F

Independence of Irrelevant Alternatives

Classic IIA for static discrete choice:

Relative frequency of choosing two alternatives is independent of the choice set

Luce (1959). IIA $\Leftrightarrow$ static choice rationalized by the logit model

For given F, formulate dynamic analogue of IIA:

Relative distance of any pair of state-action frequencies from the observed one
independent of which other state-action frequencies are available
where distance is Bregman divergence associated with F

Theorem. Dynamic IIA $\Leftrightarrow$ dynamic choice rationalized by undiscounted MDP with i.i.d. shocks $\sim F$

Extensions

The analyst observes a linear function of state-action frequencies

e.g. **Mixed models**. Heterogeneous agents. $\boldsymbol{u} \sim G$. Analyst observes $\bar{\boldsymbol{\mu}} = \int \boldsymbol{\mu}(\boldsymbol{u}) \mathrm{d}G$

Results apply to the estimation of $\bar{\boldsymbol{u}} = \int \boldsymbol{u} \mathrm{d}G$ (for known G)

Conclusion

Some results on estimation of undiscounted MDPs

- Convenient mapping to static discrete choice

Would be interesting to explore

- Estimation exploring cyclic monotonicity (Shi, Shum & Song 2018)
- Estimating random coefficient dynamic models from "market level" variation
- Beyond Conditional Independence (correlated unobservables)