Markovian Implementation

Johannes Horner (Yale)

Xiaosheng Mu (Harvard)

Nicolas Vieille (HEC)

Symposium on Game Theory, May 2015

General motivation

- World with a stochastically changing state.
- One agent/player privately observes the state.
- Principal/other player would like the agent to report truthfully:
  - To make decisions that are not in the agent's best interests;
  - Needs to incentivize the agent, with monetary transfers.
- Is it possible, and how ?

General motivation

- World with a stochastically changing state.
- One agent/player privately observes the state.
- Principal/other player would like the agent to report truthfully:
  - To make decisions that are not in the agent's best interests;
  - Needs to incentivize the agent, with monetary transfers.
- Is it possible, and how ?
- If transfer is only conditioned on current report, this is a static problem, which is well-understood.
- Can it help to condition transfers on past reports ?
- Is memory helpful ? How much memory is helpful ?

Related literature

Statistics/Econometrics : identification of Hidden Markov chains.

- Typically, HMC are not Markovian of any order.
- Yet, Blackwell-Koopmans (1957): law of (a_n) determined by fdd of size $2|S|^2 + 1$.
- Connuault (2014): generic identification of Hidden dynamic choice models.

Economics

- Repeated Bayesian games with communication (Athey Bagwell (2001), Escobar Toikka (2014), Horner, Takahashi and V. (2014)).

  Much inspired by Renault Solan and V. (2013).
- Dynamic implementation and mechanism design (Renou and Tomala (2014), Athey Segal (2014)).

Review: the static/memory 0 case

Setup

- S is the (finite) set of states, A is a copy of S.
- $\phi : A \to D$ is a decision rule of the principal.
- $u(s, a)$ is the utility of the agent when reporting a in state s (normalized : $u(s, s) = 0$ for each s).

Review: the static/memory 0 case

Setup

- S is the (finite) set of states, A is a copy of S.
- $\phi : A \to D$ is a decision rule of the principal.
- $u(s,a)$ is the utility of the agent when reporting a in state s (normalized : $u(s,s) = 0$ for each s).

Question: does there exist $t : A \to \mathbf{R}$ s.t. truth-telling is optimal:

$$\forall s, a \mapsto u(s,a) + t(a) \text{ maximal when } a = s.$$

If answer positive, ϕ is said IC.

Review: the static/memory 0 case

Setup

- S is the (finite) set of states, A is a copy of S.
- $\phi : A \to D$ is a decision rule of the principal.
- $u(s, a)$ is the utility of the agent when reporting a in state s (normalized : $u(s, s) = 0$ for each s).

Question: does there exist $t : A \to \mathbf{R}$ s.t. truth-telling is optimal:

$$\forall s, a \mapsto u(s, a) + t(a) \text{ maximal when } a = s.$$

If answer positive, ϕ is said IC.

Theorem (Rochet)

ϕ is IC iff u is cyclically monotone: $\sum_s u(s, \pi(s)) \leq 0$ for every permutation π of S.

Review: the static/memory 0 case

Equivalent formulation. Let $m \in \Delta(S)$ with full support, and $M_0 \subset \Delta(S \times A)$ the set of joint distributions with marginals m.

Then (RSV, 2013),

$$\phi \text{ is IC } \iff \mathbf{E}_\mu[u(s,a)] \leq 0 \text{ for all } \mu \in M_0.$$

The dynamic case

Setup

- S is the (finite) set of states.
- (s_n) follows a Markov chain over S, with transition matrix p.
- In each round n, agent observes s_n, chooses a report a_n.
- utility is $u(s_n, a_n)$ (plus some monetary transfer).

The dynamic case

Setup

- S is the (finite) set of states.
- (s_n) follows a Markov chain over S, with transition matrix p.
- In each round n, agent observes s_n, chooses a report a_n.
- utility is $u(s_n, a_n)$ (plus some monetary transfer).

Assume restricted to transfers with memory k: $t : A^k \times A \to \mathbf{R}$.

Given t, agent is facing a MDP with state space $A^k \times S$, action set A, utility in state $(a^{-k}, \ldots, a^{-1}, s)$ given action a is

$$u(s, a) + t(a^{-k}, \ldots, a^{-1}, a).$$

The dynamic case

Setup

- S is the (finite) set of states.
- (s_n) follows a Markov chain over S, with transition matrix p.
- In each round n, agent observes s_n, chooses a report a_n.
- utility is $u(s_n, a_n)$ (plus some monetary transfer).

Assume restricted to transfers with memory k: $t : A^k \times A \to \mathbf{R}$.

Given t, agent is facing a MDP with state space $A^k \times S$, action set A, utility in state $(a^{-k}, \ldots, a^{-1}, s)$ given action a is

$$u(s, a) + t(a^{-k}, \ldots, a^{-1}, a).$$

Definition: ϕ is k-IC if for some $t : A^k \times A \to \mathbf{R}$, it is a (long-run) stationary optimal policy to report truthfully.

Obviously, more memory cannot hurt.

Benchmarks

A sufficient condition.

ϕ 0-IC implies ϕ is k-IC for each k.

Hence ϕ k-IC if $\mathbf{E}_\mu[u(s,a)] \leq 0$ for all $\mu \in M_0$.

Benchmarks

A sufficient condition.

ϕ 0-IC implies ϕ is k-IC for each k.

Hence ϕ k-IC if $\mathbf{E}_\mu[u(s,a)] \leq 0$ for all $\mu \in M_0$.

A necessary condition.

Let $\mu \in \Delta(S \times A)$.

Benchmarks

A sufficient condition.

ϕ 0-IC implies ϕ is k-IC for each k.

Hence ϕ k-IC if $\mathbf{E}_\mu[u(s,a)] \leq 0$ for all $\mu \in M_0$.

A necessary condition.

Let $\mu \in \Delta(S \times A)$.

Assume there is a reporting strategy σ such that

- $\mathcal{L}(s_n, a_n) = \mu$ for each n.
- $\mathcal{L}((s_n)_n) = \mathcal{L}((a_n)_n)$.

Benchmarks

A sufficient condition.

ϕ 0-IC implies ϕ is k-IC for each k.

Hence ϕ k-IC if $\mathbf{E}_\mu[u(s,a)] \leq 0$ for all $\mu \in M_0$.

A necessary condition.

Let $\mu \in \Delta(S \times A)$.

Assume there is a reporting strategy σ such that

- $\mathcal{L}(s_n, a_n) = \mu$ for each n.
- $\mathcal{L}((s_n)_n) = \mathcal{L}((a_n)_n)$.

If ϕ is k-IC, then $\mathbf{E}_\mu[u(s,a)] \leq 0$.

Benchmarks

A sufficient condition.

ϕ 0-IC implies ϕ is k-IC for each k.

Hence ϕ k-IC if $\mathbf{E}_\mu[u(s,a)] \leq 0$ for all $\mu \in M_0$.

A necessary condition.

Let $\mu \in \Delta(S \times A)$.

Assume there is a reporting strategy σ such that

- $\mathcal{L}(s_n, a_n) = \mu$ for each n.
- $\mathcal{L}((s_n)_n) = \mathcal{L}((a_n)_n)$.

If ϕ is k-IC, then $\mathbf{E}_\mu[u(s,a)] \leq 0$.

Otherwise indeed, no matter t, expected transfers will be equal under σ and under truth-telling, and utilities higher under σ.

Benchmarks

A sufficient condition.

ϕ 0-IC implies ϕ is k-IC for each k.

Hence ϕ k-IC if $\mathbf{E}_\mu[u(s,a)] \leq 0$ for all $\mu \in M_0$.

A necessary condition.

Let $\mu \in \Delta(S \times A)$.

Assume there is a reporting strategy σ such that

- $\mathcal{L}(s_n, a_n) = \mu$ for each n.
- $\mathcal{L}((s_n)_n) = \mathcal{L}((a_n)_n)$.

If ϕ is k-IC, then $\mathbf{E}_\mu[u(s,a)] \leq 0$.

Otherwise indeed, no matter t, expected transfers will be equal under σ and under truth-telling, and utilities higher under σ.

So, if these two sets of measures M_0 and M_∞ coincide, memory does not help.

Benchmarks

Corollary

Memory is useless (ϕ k-IC equivalent to ϕ 0-IC) in the following two cases:

- $|S| = 2$.
- (s_n) is an iid sequence.

In fact, memory is useless iff successive states are *quasi*-iid.

Benchmarks

Corollary

Memory is useless (ϕ k-IC equivalent to ϕ 0-IC) in the following two cases:

- $|S| = 2$.
- (s_n) is an iid sequence.

In fact, memory is useless iff successive states are *quasi*-iid.

To emphasize:

- $\mu \in M_0$ if there is a reporting strategy σ s.t. $\mathcal{L}(s_n, a_n) = \mu$.
- $\mu \in M_\infty$ if in addition, $\mathcal{L}((a_n)_n) = \mathcal{L}((s_n)_n)$.

Benchmarks

Corollary

Memory is useless (ϕ k-IC equivalent to ϕ 0-IC) in the following two cases:

- $|S| = 2$.
- (s_n) is an iid sequence.

In fact, memory is useless iff successive states are *quasi*-iid.

To emphasize:

- $\mu \in M_0$ if there is a reporting strategy σ s.t. $\mathcal{L}(s_n, a_n) = \mu$.
- $\mu \in M_\infty$ if in addition, $\mathcal{L}((a_n)_n) = \mathcal{L}((s_n)_n)$.

What's in between ?

M_k and the characterization of k-IC

Let M_k be the set of $\mu \in \Delta(S \times A)$ for which there is a reporting strategy σ s.t.

- $\mathcal{L}(s_n, a_n) = \mu$ for each n.
- $\mathcal{L}(a_n, a_{n+1}, \ldots, a_{n+k}) = \mathcal{L}(s_n, s_{n+1}, \ldots, s_{n+k})$ for each n.

That is, σ cannot be distinguished from truth-telling, with memory k tests.

M_k and the characterization of k-IC

Let M_k be the set of $\mu \in \Delta(S \times A)$ for which there is a reporting strategy σ s.t.

- $\mathcal{L}(s_n, a_n) = \mu$ for each n.
- $\mathcal{L}(a_n, a_{n+1}, \ldots, a_{n+k}) = \mathcal{L}(s_n, s_{n+1}, \ldots, s_{n+k})$ for each n.

That is, σ cannot be distinguished from truth-telling, with memory k tests.

Proposition

ϕ is k-IC iff $\mathbf{E}_\mu[u(s,a)] \leq 0$ for all $\mu \in M_k$.

Proof: The direct implication is straightforward.

For the reverse one, look at the zero-sum game where the principal chooses t, and the agent a measure in $\Delta((A^k \times S) \times A)$.

How does M_1 look like ? The cyclic case

Let $p = \begin{pmatrix} a & b & c \\ c & a & b \\ b & c & a \end{pmatrix}$, with $c < b < a \leq \frac{1}{2}$.

How does M_1 look like ? The cyclic case

Let $p = \begin{pmatrix} a & b & c \\ c & a & b \\ b & c & a \end{pmatrix}$, with $c < b < a \leq \frac{1}{2}$. Some μ's in M_1:

- Truth-telling: $\mu_{TT} = \begin{pmatrix} \frac{1}{3} & 0 & 0 \\ 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{3} \end{pmatrix}$. ($S \times A-$matrix.)

How does M_1 look like ? The cyclic case

Let $p = \begin{pmatrix} a & b & c \\ c & a & b \\ b & c & a \end{pmatrix}$, with $c < b < a \leq \frac{1}{2}$. Some μ's in M_1:

- Truth-telling: $\mu_{TT} = \begin{pmatrix} \frac{1}{3} & 0 & 0 \\ 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{3} \end{pmatrix}$. ($S \times A-$matrix.)
- Announcing $s+1$: $\mu = \begin{pmatrix} 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{3} \\ \frac{1}{3} & 0 & 0 \end{pmatrix}$.

How does M_1 look like ? The cyclic case

Let $p = \begin{pmatrix} a & b & c \\ c & a & b \\ b & c & a \end{pmatrix}$, with $c < b < a \leq \frac{1}{2}$. Some μ's in M_1:

- Truth-telling: $\mu_{TT} = \begin{pmatrix} \frac{1}{3} & 0 & 0 \\ 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{3} \end{pmatrix}$. ($S \times A-$matrix.)
- Announcing $s+1$: $\mu = \begin{pmatrix} 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{3} \\ \frac{1}{3} & 0 & 0 \end{pmatrix}$.
- Announcing the state of the previous round.

How does M_1 look like ? The cyclic case

Let $p = \begin{pmatrix} a & b & c \\ c & a & b \\ b & c & a \end{pmatrix}$, with $c < b < a \leq \frac{1}{2}$. Some μ's in M_1:

- Truth-telling: $\mu_{TT} = \begin{pmatrix} \frac{1}{3} & 0 & 0 \\ 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{3} \end{pmatrix}$. ($S \times A-$matrix.)
- Announcing $s+1$: $\mu = \begin{pmatrix} 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{3} \\ \frac{1}{3} & 0 & 0 \end{pmatrix}$.
- Announcing the state of the previous round.
- That's not all: pick a reporting strategy $\sigma(a \mid s^{-1}, a^{-1}, s)$ such that for each (s^{-1}, a^{-1}) the conditional law of (s, a) has marginals $p(s \mid s^{-1})$ and $p(a \mid a^{-1})$.

How does M_1 look like ? The cyclic case

Let $p = \begin{pmatrix} a & b & c \\ c & a & b \\ b & c & a \end{pmatrix}$, with $c < b < a \leq \frac{1}{2}$. Some μ's in M_1:

- Truth-telling: $\mu_{TT} = \begin{pmatrix} \frac{1}{3} & 0 & 0 \\ 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{3} \end{pmatrix}$. ($S \times A-$matrix.)
- Announcing $s+1$: $\mu = \begin{pmatrix} 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{3} \\ \frac{1}{3} & 0 & 0 \end{pmatrix}$.
- Announcing the state of the previous round.
- That's not all: pick a reporting strategy $\sigma(a \mid s^{-1}, a^{-1}, s)$ such that for each (s^{-1}, a^{-1}) the conditional law of (s, a) has marginals $p(s \mid s^{-1})$ and $p(a \mid a^{-1})$. Then (a_n) is a Markov chain, hence any invariant measure is in M_∞.

How does M_1 look like ? The cyclic case

Let $p = \begin{pmatrix} a & b & c \\ c & a & b \\ b & c & a \end{pmatrix}$, with $c < b < a \leq \frac{1}{2}$. Some μ's in M_1:

- Truth-telling: $\mu_{TT} = \begin{pmatrix} \frac{1}{3} & 0 & 0 \\ 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{3} \end{pmatrix}$. ($S \times A-$matrix.)
- Announcing $s+1$: $\mu = \begin{pmatrix} 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{3} \\ \frac{1}{3} & 0 & 0 \end{pmatrix}$.
- Announcing the state of the previous round.
- That's not all: pick a reporting strategy $\sigma(a \mid s^{-1}, a^{-1}, s)$ such that for each (s^{-1}, a^{-1}) the conditional law of (s, a) has marginals $p(s \mid s^{-1})$ and $p(a \mid a^{-1})$. Then (a_n) is a Markov chain, hence any invariant measure is in M_∞.
- That's not all (but proving this is tricky).

How does M_1 look like ? The cyclic case

Let $p = \begin{pmatrix} a & b & c \\ c & a & b \\ b & c & a \end{pmatrix}$, with $c < b < a \leq \frac{1}{2}$. Some μ's in M_1:

- Truth-telling: $\mu_{TT} = \begin{pmatrix} \frac{1}{3} & 0 & 0 \\ 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{3} \end{pmatrix}$. ($S \times A$−matrix.)
- Announcing $s+1$: $\mu = \begin{pmatrix} 0 & \frac{1}{3} & 0 \\ 0 & 0 & \frac{1}{3} \\ \frac{1}{3} & 0 & 0 \end{pmatrix}$.
- Announcing the state of the previous round.
- That's not all: pick a reporting strategy $\sigma(a \mid s^{-1}, a^{-1}, s)$ such that for each (s^{-1}, a^{-1}) the conditional law of (s, a) has marginals $p(s \mid s^{-1})$ and $p(a \mid a^{-1})$. Then (a_n) is a Markov chain, hence any invariant measure is in M_∞.
- That's not all (but proving this is tricky).
- Here is one μ in $M_0 \setminus M_1$: $\mu = \begin{pmatrix} \frac{1}{3} & 0 & 0 \\ 0 & 0 & \frac{1}{3} \\ 0 & \frac{1}{3} & 0 \end{pmatrix}$.

A general description of M_1

Let $\mu \in M_1$, and a corresponding reporting strategy σ.

Denote by $\sigma(s^{-1}, a^{-1}, s, a) \in \Delta(S \times A \times S \times A)$ the induced occupation measure.

A general description of M_1

Let $\mu \in M_1$, and a corresponding reporting strategy σ.

Denote by $\sigma(s^{-1}, a^{-1}, s, a) \in \Delta(S \times A \times S \times A)$ the induced occupation measure. Then:

- C1 $\sum_{s^{-1}, a^{-1}} \sigma(s^{-1}, a^{-1}, s, a) = \mu(s, a)$ (because $\mathcal{L}(s_n, a_n) = \mu$).

A general description of M_1

Let $\mu \in M_1$, and a corresponding reporting strategy σ.

Denote by $\sigma(s^{-1}, a^{-1}, s, a) \in \Delta(S \times A \times S \times A)$ the induced occupation measure. Then:

- C1 $\sum_{s^{-1},a^{-1}} \sigma(s^{-1}, a^{-1}, s, a) = \mu(s, a)$ (because $\mathcal{L}(s_n, a_n) = \mu$).
- C2 $\sum_{s^{-1},s} \sigma(s^{-1}, a^{-1}, s, a) = m(a^{-1})p(a \mid a^{-1})$ (because $\mathcal{L}(a_n, a_{n+1}) = \mathcal{L}(s_n, s_{n+1})$).

A general description of M_1

Let $\mu \in M_1$, and a corresponding reporting strategy σ.

Denote by $\sigma(s^{-1}, a^{-1}, s, a) \in \Delta(S \times A \times S \times A)$ the induced occupation measure. Then:

- C1 $\sum_{s^{-1},a^{-1}} \sigma(s^{-1}, a^{-1}, s, a) = \mu(s, a)$
  (because $\mathcal{L}(s_n, a_n) = \mu$).
- C2 $\sum_{s^{-1},s} \sigma(s^{-1}, a^{-1}, s, a) = m(a^{-1})p(a \mid a^{-1})$
  (because $\mathcal{L}(a_n, a_{n+1}) = \mathcal{L}(s_n, s_{n+1})$).
- C3 $\sum_{a} \sigma(s^{-1}, a^{-1}, s, a) = \mu(s^{-1}, a^{-1})p(s \mid s^{-1})$
  (because a_n does not anticipate on a_{n+1}).

A general description of M_1

Let $\mu \in M_1$, and a corresponding reporting strategy σ.

Denote by $\sigma(s^{-1}, a^{-1}, s, a) \in \Delta(S \times A \times S \times A)$ the induced occupation measure. Then:

- C1 $\sum_{s^{-1},a^{-1}} \sigma(s^{-1}, a^{-1}, s, a) = \mu(s, a)$ (because $\mathcal{L}(s_n, a_n) = \mu$).
- C2 $\sum_{s^{-1},s} \sigma(s^{-1}, a^{-1}, s, a) = m(a^{-1})p(a \mid a^{-1})$ (because $\mathcal{L}(a_n, a_{n+1}) = \mathcal{L}(s_n, s_{n+1})$).
- C3 $\sum_{a} \sigma(s^{-1}, a^{-1}, s, a) = \mu(s^{-1}, a^{-1})p(s \mid s^{-1})$ (because a_n does not anticipate on a_{n+1}).

Conversely, **C1-C3** (plus $\sigma \geq 0$) characterize M_1.

A similar characterization holds for M_k $(k < \infty)$.

Back to the cyclic case

Remember: $p = \begin{pmatrix} a & b & c \\ c & a & b \\ b & c & a \end{pmatrix}$, $c < b < a \leq \frac{1}{2}$.

M_0 is the set of copulas.

- Can be viewed as a subset of $\mathbf{R}^4$.

Back to the cyclic case

Remember: $p = \begin{pmatrix} a & b & c \\ c & a & b \\ b & c & a \end{pmatrix}$, $c < b < a \leq \frac{1}{2}$.

M_0 is the set of copulas.

- Can be viewed as a subset of $\mathbf{R}^4$.
- It is a polytope with 6 vertices.

Back to the cyclic case

Remember: $p = \begin{pmatrix} a & b & c \\ c & a & b \\ b & c & a \end{pmatrix}$, $c < b < a \leq \frac{1}{2}$.

M_0 is the set of copulas.

- Can be viewed as a subset of $\mathbf{R}^4$.
- It is a polytope with 6 vertices.

M_1 is also a polytope, but a complex one.

- It has 57 vertices. Most of the 3-D faces are (topologically) cubes.
- some of the 3-D faces have 12 vertices and 10 faces. Six of them have four vertices, two have five, and two have three.

Back to the cyclic case

Remember: $p = \begin{pmatrix} a & b & c \\ c & a & b \\ b & c & a \end{pmatrix}$, $c < b < a \leq \frac{1}{2}$.

M_0 is the set of copulas.

- Can be viewed as a subset of $\mathbf{R}^4$.
- It is a polytope with 6 vertices.

M_1 is also a polytope, but a complex one.

- It has 57 vertices. Most of the 3-D faces are (topologically) cubes.
- some of the 3-D faces have 12 vertices and 10 faces. Six of them have four vertices, two have five, and two have three.

Theorem

(In the cyclic case, $|S| = 3$), one has $M_1 = M_\infty$.

In such cyclic cases, memory one is useful and enough.

This is a fragile result

- Take $p = \begin{pmatrix} 1-\varepsilon_1 & \varepsilon_1 & 0 \\ 0 & 1-\varepsilon_2 & \varepsilon_2 \\ \varepsilon_3 & 0 & 1-\varepsilon_3 \end{pmatrix}$.
  If $\varepsilon_1, \varepsilon_2$ and ε_3 are distinct, the sequence of sets (M_k) is strictly decreasing.

This is a fragile result

- Take $p = \begin{pmatrix} 1-\varepsilon_1 & \varepsilon_1 & 0 \\ 0 & 1-\varepsilon_2 & \varepsilon_2 \\ \varepsilon_3 & 0 & 1-\varepsilon_3 \end{pmatrix}$.
  If $\varepsilon_1, \varepsilon_2$ and ε_3 are distinct, the sequence of sets (M_k) is strictly decreasing.
- Every neighborhood of the uniform iid $p = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{pmatrix}$ contains a transition matrix such that $M_\infty \neq M_1$.

This is a fragile result

- Take $p = \begin{pmatrix} 1-\varepsilon_1 & \varepsilon_1 & 0 \\ 0 & 1-\varepsilon_2 & \varepsilon_2 \\ \varepsilon_3 & 0 & 1-\varepsilon_3 \end{pmatrix}$.
  If $\varepsilon_1, \varepsilon_2$ and ε_3 are distinct, the sequence of sets (M_k) is strictly decreasing.
- Every neighborhood of the uniform iid $p = \begin{pmatrix} \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \\ \frac{1}{3} & \frac{1}{3} & \frac{1}{3} \end{pmatrix}$
  contains a transition matrix such that $M_\infty \neq M_1$.

We suspect that generically, $M_\infty \neq M_1$.

Conclusion 1: from a “statistical” perspective, a large memory is typically useful.

The “economist’s perspective

Remember:

- ϕ 1-IC iff $\mathbf{E}_\mu[u(s,a)] \leq 0$ for all $\mu \in M_1$.
- ϕ 2-IC iff $\mathbf{E}_\mu[u(s,a)] \leq 0$ for all $\mu \in M_2$.

Hence $M_1 = M_2$ implies equivalence between $1 - IC$ and $2 - IC$.

The “economist’s perspective

Remember:

- ϕ 1-IC iff $\mathbf{E}_\mu[u(s,a)] \leq 0$ for all $\mu \in M_1$.
- ϕ 2-IC iff $\mathbf{E}_\mu[u(s,a)] \leq 0$ for all $\mu \in M_2$.

Hence $M_1 = M_2$ implies equivalence between $1 - IC$ and $2 - IC$.

<u>The reverse is not true</u>.

- The reason is that $\mu_{TT} \in M_k$ for all k, and $\mathbf{E}_{\mu_{TT}}[u(s,a)] = 0$.

The “economist’s perspective

Remember:

- ϕ 1-IC iff $\mathbf{E}_\mu[u(s,a)] \leq 0$ for all $\mu \in M_1$.
- ϕ 2-IC iff $\mathbf{E}_\mu[u(s,a)] \leq 0$ for all $\mu \in M_2$.

Hence $M_1 = M_2$ implies equivalence between $1 - IC$ and $2 - IC$.

<u>The reverse is not true</u>.

- The reason is that $\mu_{TT} \in M_k$ for all k, and $\mathbf{E}_{\mu_{TT}}[u(s,a)] = 0$.
- In fact, if M_2 is obtained by shrinking M_1 towards μ_{TT}, then equivalence between 1-IC and 2-IC.

So what ?

So what ?

Let $C_k = \{\lambda\mu + (1-\lambda)\mu_{TT},\ \lambda \geq 0 \text{ and } \mu \in M_k\}$ be the cone with vertex μ_{TT} spanned by M_k. Then

$$C_1 = C_2 \Leftrightarrow \text{(1-IC is equivalent to 2-IC).}$$

For implementation, the relevant sets are the C_k's, not the M_k's.

So what ?

Let $C_k = \{\lambda\mu + (1-\lambda)\mu_{TT},\ \lambda \geq 0 \text{ and } \mu \in M_k\}$ be the cone with vertex μ_{TT} spanned by M_k. Then

$$C_1 = C_2 \Leftrightarrow \text{(1-IC is equivalent to 2-IC)}.$$

For implementation, the relevant sets are the C_k's, not the M_k's.

Theorem

Assume $|S| = 3$ and $p(s' \mid s) \leq \frac{1}{2}$. Assume moreover that no column of p contains three distinct entries. Then one has
$C_1 = C_\infty$.
(And $C_0 \neq C_1$ typically).

It looks likely that the additional condition on p can be dropped.

That is, the sets (M_k) are always shrinking towards M_∞ along the extreme rays of C_1.

So what ?

Let $C_k = \{\lambda\mu + (1-\lambda)\mu_{TT},\ \lambda \geq 0 \text{ and } \mu \in M_k\}$ be the cone with vertex μ_{TT} spanned by M_k. Then

$$C_1 = C_2 \Leftrightarrow (\text{1-IC is equivalent to 2-IC}).$$

For implementation, the relevant sets are the C_k's, not the M_k's.

Theorem

Assume $|S| = 3$ and $p(s' \mid s) \leq \frac{1}{2}$. Assume moreover that no column of p contains three distinct entries. Then one has $C_1 = C_\infty$.
(And $C_0 \neq C_1$ typically).

It looks likely that the additional condition on p can be dropped.

That is, the sets (M_k) are always shrinking towards M_∞ along the extreme rays of C_1.

Conclusion 2: from an "economist" perspective, memory one is *always* enough.

Going beyond

Remember: we assume

- A1 $p(s' \mid s) \leq \frac{1}{2}$.
- A2 three states.
- A3 no discounting.

Going beyond

Remember: we assume

- A1 $p(s' \mid s) \leq \frac{1}{2}$.
- A2 three states.
- A3 no discounting.

If **A1** is dropped, then:

- In the cyclic case, $M_1 \neq M_\infty$ (always) and $C_1 = C_\infty$ (always).
- We have examples in which $C_1 \neq C_2$.

Going beyond

Remember: we assume

- A1 $p(s' \mid s) \leq \frac{1}{2}$.
- A2 three states.
- A3 no discounting.

If **A1** is dropped, then:

- In the cyclic case, $M_1 \neq M_\infty$ (always) and $C_1 = C_\infty$ (always).
- We have examples in which $C_1 \neq C_2$.

The extension to more than three states is on the to-do list.