Approximate Implementation In Markovian Environments

Ludovic Renou & Tristan Tomala

5th CMSS Summer Workshop

Introduction

- Mechanism design: find a game such that there exists ONE equilibrium that implement desired outcomes.

  Drawback: multiplicity of equilibria and of induced outcomes.

- Full implementation: find a game such that ALL equilibria implement the desired outcomes.

  Drawbacks: Complex mechanisms (integer games, dictatorships), complex conditions (Maskin, Jackson).

- Implementation in dynamic settings: multiplicity of equilibria, Folk theorem effect.

- Our approach: all equilibria implement the desired outcomes most of the time and along most histories.

Main result

Our main result is a sufficient condition for approximate implementation:

If a social choice function f is strictly efficient in the set of all social choice functions that

(i) have a smaller range than f and

(ii) satisfy an undetectability condition,

then f is approximately implementable.

Single-shot implementation

A static or *single-shot implementation problem* $\mathcal{P}$ is a tuple $\langle \mathcal{I}, X, (u_i, \Theta_i)_{i \in \mathcal{I}} \rangle$, where

- $\mathcal{I} := \{1, \ldots, I\}$ is the set of I agents,
- X is a finite set of alternatives,
- Θ_i is a finite set of types for agent i,
- and for each agent i, $u_i : X \times \Theta_i \to \mathbb{R}$ is a type-dependent utility function (private values).

A social choice function f maps $\Theta := \times_i \Theta_i$ to $\Delta(X)$, where $\Delta(X)$ is the set of lotteries over X

A static mechanism G (or game form) is a pair $\langle (M_i)_{i \in \mathcal{I}}, g \rangle$, with

- M_i the set of messages of player i, and
- $g : \times_{i \in \mathcal{I}} M_i \to \Delta(X)$ the allocation rule.

Repeated implementation problems

- An infinitely repeated implementation problem $\mathcal{P}^\infty$ is the infinite repetition of $\mathcal{P}$.
- At each period t, agent i is privately informed of his type $\theta_{i,t}$.
- Agent i's type follows a Markov chain with initial distribution p_i and transition matrix P_i.
- Types are independently distributed. Assume that the process (p_i, P_i) is ergodic with (unique) invariant distribution p_i. ($p = \otimes p_i$.)
- Player i's expected payoff of implementing f is $V_i(f) := \sum_\theta u_i(f(\theta), \theta_i)p(\theta)$.
- We assume that players evaluate streams of payoff according to the discounting criterion, with discount factor δ.

Dynamic mechanisms

- Let $\mathcal{G}$ be the collection of possible static mechanisms.
- Let $\mathcal{H}_D^t$ be the set of all possible (designer) histories of mechanisms and corresponding messages over $t-1$ periods.
- A generic history $h_D^t \in \mathcal{H}_D^t$ is a sequence of mechanisms and corresponding messages $(G_1, m_1, \ldots, G_\tau, m_\tau, \ldots, G_{t-1}, m_{t-1})$.
- The set of all possible (designer) histories is $\mathcal{H}_D = \cup_t \mathcal{H}_D^t$.
- A *dynamic mechanism* specifies a mechanism $\mathcal{M}(h_D) \in \mathcal{G}$ as a function of the designer history $h_D \in \mathcal{H}_D$.

Dynamic mechanisms and induced games

- We assume perfect monitoring.
- At the beginning of period t, each agent i knows:
  - the entire profile of mechanisms chosen up to period $t-1$,
  - the entire profile of messages played up to period $t-1$,
  - his types realized up to period $t-1$,
  - the period t's mechanism, and
  - his realized type for period t.
- Let σ be a profile of strategies and $U_i^{\delta}(\sigma)$ agent i's associated payoff.

Approximate implementation

Fix a dynamic mechanism, a terminal history h^∞ and the corresponding infinite sequence $(\theta_t, x_t)_t$ of realized types and implemented alternatives.

Denote

$$N^\delta_{h^\infty}(\theta, x) = \sum_{t\geq 1}(1-\delta)\delta^{t-1}\mathbb{1}\{\theta_t = \theta, x_t = x\}$$

the average discounted number of times along the history h^∞ where the realized profile of types is θ and the implemented alternative is x.

Similarly, denote

$$N^\delta_{h^\infty}(\theta) = \sum_{t\geq 1}(1-\delta)\delta^{t-1}\mathbb{1}\{\theta_t = \theta\}$$

the average discounted number of times along the history h^∞ where the realized profile of types is θ.

Approximation implementation: cont'd

Definition (Approximate Implementation)

A social choice function $f : \Theta \to \Delta(X)$ is approximatively implementable if for all $\varepsilon > 0$, there exist a mechanism $\mathcal{M}_\varepsilon$ and a discount factor δ_ε such that for every $\delta \geq \delta_\varepsilon$, the set of equilibrium is non-empty and for every equilibrium σ in the δ-discounted game,

$$\mathbb{P}_{\mathcal{M},\sigma}(\{h^\infty \in \mathcal{H}^\infty : |N^\delta_{h^\infty}(\theta, x) - N^\delta_{h^\infty}(\theta) f(x|\theta)| \leq \varepsilon, \ \forall \theta, x\}) \geq 1 - \varepsilon.$$

Comments

- Assume that f is deterministic. The definition states that the subset of (terminal) histories such that $|N^\delta(\theta, f(\theta)) - N^\delta(\theta)| \leq \varepsilon$ and $|N^\delta(\theta, x)| \leq \varepsilon$ for all (θ, x) with $x \neq f(\theta)$ has probability at least $1 - \varepsilon$ under any equilibrium for δ large enough.

- This implies that the average number of times when the profile of type is θ and the alternative implemented is different from $f(\theta)$ is $O(\varepsilon)$, i.e.,

$$\sum_{t\geq 1}(1-\delta)\delta^{t-1}\mathbb{1}\{x_t \neq f(\theta_t)\} \leq 2\varepsilon\sharp\Theta\sharp X.$$

- We have $\mathbb{E}_{\mathcal{M}_\varepsilon,\sigma}N(x,\theta) \approx p(\theta)f(\theta)[x]$ for all x, θ.

- Also, $U_i(\sigma) \approx V_i(f)$ for all equilibria.

Deceptions

- A correlated deception is a mapping π from Θ to $\Delta(\Theta)$.
- Intuitively, $\pi(\theta)[\hat{\theta}]$ is the probability with which agents report $\hat{\theta}$ to the designer when their profile of types is θ.
- A deception is undetectable if $p\pi = p$, i.e., $\sum_\theta p(\theta)\pi(\theta)[\hat{\theta}] = p(\hat{\theta})$ for all $\hat{\theta}$.
- If agents follow an undetectable deception, the law of reported types is the same as the law of realized types.

Undetectable efficiency

- Let $\Pi_p := \{\pi : p\pi = p\}$ be the set of undetectable deceptions.
- Write $\pi f : \Theta \to \Delta(X)$ the social choice function given by $\pi f(\theta) = \sum_{\hat{\theta}} \pi(\theta)[\hat{\theta}] f(\hat{\theta})$.
- It is the function that results when the agents pretend that the profile of type is $\hat{\theta}$ when it is θ.

Undetectable efficiency

Definition (Undetectable Efficiency, I)

A social choice function f is *undetectable efficient* if for all collections $(\pi_k)_k$ of undetectable deceptions and convex coefficients $(\alpha_k)_k$ such that $\sum_k \alpha_k V_i(\pi_k f) \geq V_i(f)$ for all $i \in I$, then $\pi_k f = f$ for all k with $\alpha_k > 0$.

This implies that $V(f)$ is an extreme point of $\{v : v = V(\pi f), \pi \in \Pi_p\}$, but it is stronger than that.

Why convex combinations?

- Let $\pi = \sum_k \alpha_k \pi_k$. If all π_k are undetectable, then so is π.
- But we might have that $\pi f = f$, while $\pi_k f \neq f$ for some k with $\alpha_k > 0$.
- Idea: through communication or jointly controlled lotteries, the agents can implement $\pi_k f$ about α_k of the time. Since $\pi_k f \neq f$, we would have a contradiction with the requirement for approximate implementation.
- However, if $f(\theta) \notin \mathrm{co}\{f(\theta') : f(\theta') \neq f(\theta)\}$ for all θ, then $\pi f = f$ impies $\pi_k f = f$ for all k with $\alpha_k > 0$.

An alternative characterization

Lemma

A social choice function is undetectable efficient if and only if for all collections $(\rho_k)_k$ *of permutations and convex coefficients* $(\alpha_k)_k$ *such that* $\sum_\theta \alpha_k u_i(f(\rho_k(\theta)),\theta) \geq \sum_\theta u_i(f(\theta),\theta)$ *for all* i*, then* $f \circ \rho_k = f$ *for all* k *with* $\alpha_k > 0$*.*

To see this, suppose that there are only one agent and two types θ_L and θ_H. For all $\pi \in \Pi_p$, we can write

$$\begin{pmatrix} p(\theta_L)\pi(\theta_L)[\theta_L] & p(\theta_L)\pi(\theta_L)[\theta_H] \\ p(\theta_H)\pi(\theta_H)[\theta_L] & p(\theta_H)\pi(\theta_H)[\theta_H] \end{pmatrix} = \begin{pmatrix} p(\theta_L) & 0 \\ 0 & p(\theta_H) \end{pmatrix} - \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} + J,$$

where J is a bi-stochastic matrix (hence a convex combinations of permutation matrices).

Main result

Theorem

If a social choice function is undetectable efficient, then it is approximatively implementable.

An economic example: monopolistic screening

- There is one buyer and one seller.
- The buyer has two types θ_L and θ_H, with $\theta_H > \theta_L > 0$.
- If the buyer purchases q units at total price t, his payoff is $\theta q - t$.
- The cost of production is q^2.
- The reservation utility is normalized to zero.

Example cont'd

- Consider the contract $f := (q^*, t^*)$: $q^*(\theta_L) = \theta_L$, $q^*(\theta_H) = \theta_H$, $t^*(\theta_L) = \theta_L^2$ and $t^*(\theta_H) = \theta_H^2$, i.e., the principal implements the first best allocation and leaves no rent to the agent.
- This contract is approximatively implementable, because:

$$V(f) = 0 > \theta_H \theta_L - \theta_L^2 + \theta_L \theta_H - \theta_H^2 = -(\theta_L + \theta_H)^2 = V(f \circ \rho).$$

- The intuition for this result is simple. With the contract (q^*, t^*), the agent has a strict incentive to report θ_L when he is of type θ_H since his payoff is $\theta_H \theta_L - (\theta_L)^2 > 0$.
- However, for the deception to be undetectable, the agent would need to report θ_H when his type is θ_L, which generates a strictly negative payoff $(\theta_L \theta_H - \theta_H^2 < 0)$.
- The gains to misreport when the type is θ_H are offset by the losses to misreport when the type is θ_L.

Idea of the proof (for one agent and iid)

- The condition is equivalent to: if $p\pi = p$ and $\pi f \neq f$, then $V(f) > V(\pi f)$.
- Notice that the deception π^*, defined by $\pi^*(\theta)[\hat{\theta}] = p(\hat{\theta})$ for all $(\theta, \hat{\theta})$ is undetectable.
- Also, $\pi^* f$ is a constant social choice function.
- So, if $\pi^* f \neq f$, it can be used as a "punishment."
- And, if $\pi^* f = f$, then it is trivially implementable.

The mechanism

The proof is constructive. Consider the following mechanism parameterized by T, L, η:

- Time is divided in consecutive blocks of length T, denoted $B_1, B_2, \dots, B_k, \dots$.
- At each block B_k, the agent is either active or is inactive.
- Whenever the agent is active, his set of messages is Θ, while it is a singleton when he is inactive.
- At the first block B_1, the agent is active.
- For each block B_k at which the agent is active, the designer implements x_t at period t with probability $f(\hat{\theta}_t)[x_t]$, where $\hat{\theta}_t$ is the agent's report at t.
- If the agent is inactive, the designer draws a fictitious type $\hat{\theta}_t$ with probability $p(\hat{\theta}_t)$ and then implements x_t at period t with probability $f(\hat{\theta}_t)[x_t]$ (i.e., it implements $\pi^* f$).

- At the end of each block B_k where the agent is active, the designer computes the empirical frequency μ_k over the block

$$\mu_k(\theta) = (1/T) \sum_t \mathbb{1}\{\hat{\theta}_t = \theta\},$$

and tests the agent as follows:

  - If for all θ,
$$\left| \mu_k(\theta) - p(\theta) \right| \leq \eta,$$
the agent passes the test and remains active in the next block B_{k+1}.

  - If not, the agent fails the test and remains inactive for the L successive blocks, B_{k+1} to B_{k+L}. Agent i is again active at block B_{k+L+1}.

Why this mechanism?

The mechanism serves several purposes:

- To test whether the agent's empirical distribution of reported types is consistent with the iid process p.
- Whenever the agent fails the test, the distribution of simulated ("reported") types over the consecutive L blocks is consistent with p. The mechanism thus guarantees that the overall distributions of reported types is close to p (for L large).
- Therefore, the agent's reporting (average) strategy must be close to an undetectable deception.
- The test also guarantees that a truthful agent gets a payoff arbitrarily closed to $V(f)$.

The roadmap of the proof

- Step 1: we bound the payoff $U(\sigma)$ to the agent associated with any strategy profile σ. In particular, $U(\sigma) \approx \mathbb{E} V(\pi_{h^\infty} f)$ for some deception π_{h^∞} such that $p\pi_{h^\infty} \approx p$ (almost undetectable) for a set of histories h^∞ with a large probability (under $\mathbb{P}_\sigma$).
- Step 2: we show that if the agent is truthful (strategy φ), then he guarantees himself $V(f) - C\varepsilon$ (where C is a constant and $\varepsilon > 0$).
- Step 3: If σ is optimal, it must be that $U(\sigma) \geq U(\varphi) \geq V(f) - C\varepsilon$. Use then the efficiency condition to complete the proof.

Step 1

- Fix any strategy profile σ and a terminal history h^∞.
- The probability of reporting $\hat{\theta}_t$ given the history h^t and current type θ_t is $\sigma(h^t, \theta_t)[\hat{\theta}_t]$. (The types reported when the agent is inactive are assumed to be the one simulated by the designer.)
- Write $\bar{\sigma}_k(h^\infty, \theta)[\hat{\theta}] = (1/T) \sum_{t \in B_k} \sigma(h^t, \theta)[\hat{\theta}]$ for the average deception over the block B_k.
- Write $\bar{\sigma}^\delta(h^\infty, \theta)[\hat{\theta}] = \sum_k (1-\delta^T)\delta^{T(k-1)} \bar{\sigma}_k(h^\infty, \theta)[\hat{\theta}]$ for the discounted average deception over the entire history.
- Similarly, write $\bar{\mu}^\delta_{h^\infty}(\theta, \hat{\theta}, x)$ for the discounted average number of times the type is θ, the type reported is $\hat{\theta}$, and the alternative implemented is x.

Step 1 cont'd

We have the following result:

Lemma

For each strategy profile σ*, each* $\alpha > 0$ *and discount factor* δ*,*

$$\mathbb{P}_\sigma\Big(\{h^\infty : \sum_{\theta,\hat{\theta},x} |\bar{\mu}^\delta_{h^\infty}(\theta,\hat{\theta},x) - p(\theta)\bar{\sigma}^\delta(h^\infty,\theta)[\hat{\theta}]f(\hat{\theta})[x]| \leq \alpha\}\Big)$$

$$\geq 1 - C(1-\delta)/\alpha^2,$$

with $C := (2\sharp\Theta + \sharp X)(\sharp\Theta^4.\sharp X^2)$.

Step 1 cont'd

Sketch of proof.

- Define the random variable:

$$X_k = \frac{1}{T} \sum_{t \in B_k} \left(\mathbb{1}\{\theta_t = \theta, \hat{\theta}_t = \hat{\theta}, x_t = x\} - \mathbb{E}_\sigma[\mathbb{1}\{\theta_t = \theta, \hat{\theta}_t = \hat{\theta}, x_t = x\} | h^t] \right).$$

- Observe that

$$\mathbb{E}_\sigma[\mathbb{1}\{\theta_t = \theta, \hat{\theta}_t = \hat{\theta}, x_t = x\} | h^t] = p(\theta)\sigma(h^t, \theta)[\hat{\theta}] f(\hat{\theta})[x].$$

- Note that

$$\sum (1 - \delta^T)\delta^{T(k-1)} X_k = \bar{\mu}^\delta_{h^\infty}(\theta, \hat{\theta}, x) - p(\theta)\bar{\sigma}^\delta(h^\infty, \theta)[\hat{\theta}] f(\hat{\theta})[x].$$

- Apply Bienaymé-Chebychev.

□

Step 1 cont'd

We have the following corollary:

Corollary

For each strategy profile σ, each $\alpha > 0$ and discount factor δ,

$$\mathbb{P}_\sigma \left(\{ h^\infty : \sum_{\hat{\theta}} \left| \bar{\mu}^\delta_{h^\infty}(\hat{\theta}) - \sum_\theta p(\theta) \bar{\sigma}^\delta(h^\infty, \theta)[\hat{\theta}] \right| \leq \alpha \} \right) \geq 1 - C(1-\delta)/\alpha^2.$$

It states that the discounted empirical distribution of reported types could be made arbitrarily closed to $p\bar{\sigma}^\delta$.

Step 1 cont'd

Lemma

For each strategy profile σ, each $\alpha > 0$ and discount factor δ,

$$\mathbb{P}_\sigma\left(\{h^\infty : \sum_{\hat\theta}\left|\bar\mu^\delta_{h^\infty}(\hat\theta) - p(\hat\theta)\right| \le \alpha\}\right) \ge 1 - C(\eta, T, L, \alpha)(1-\delta).$$

Sketch of proof.

- If the agent is inactive, his reported type is simulated according to the distribution p.
- If the agent is active and passes the test, the empirical distribution of his reported types is close to the distribution p.
- If the agent is active and fails the test, then the distribution of reported types may be very different from p. But this won't happen too often. For instance, for large δ, at worst, it would happen once every $L+1$ blocks.

□

Step 1 cont'd

Let us summarize what we have learned.

- Define π_{h^∞} the average deception $\theta \mapsto \bar{\sigma}(h^\infty, \theta)$.
- From the first lemma, we have that $U(\sigma) \approx \mathbb{E}V(\pi_{h^\infty} f)$.
- From the corollary and second lemma, we have that $p(\hat{\theta}) \approx \sum_\theta p(\theta)\pi_{h^\infty}(\theta)[\hat{\theta}]$ for a set of histories with high proba, i.e., π_{h^∞} is almost undetectable.
- Lastly, we have that $N^\delta(x, \theta) \approx p(\theta)(\pi_{h^\infty} f)(\theta)[x]$ for δ large enough and on a set of histories with high proba.
- It remains to show that $\pi_{h^\infty} f \approx f$, so that $N^\delta(x, \theta) \approx p(\theta) f(\theta)[x]$, as required.

Step 2

- If an agent is truthful, he guarantees himself approximatively $V(f)$.

Step 2 cont'd

- If the agent is truthful, he passes the statistical test in block B_k with (arbitrarily) high probability (law of large numbers). Note that the probability ξ to pass the test is constant.
- We need to show that the expected average number N of blocks a truthful agent fails is arbitrarily small.
- N can be computed recursively: Suppose that the agent is active at block k.
  - At the next block, the agent is again active with probability ξ and the process starts anew.
  - With probability $1 - \xi$, the agent is inactive in the next L blocks and the process starts anew at block B_{k+L+1}.
- So, we have

$$N = (1 - \delta^T) + \delta^T \xi N + \delta^{T(L+1)}(1 - \xi)N.$$

- This is close to 1 if T is sufficiently large (ξ depends on T).

Step 3

So, we have that

- $\mathbb{E}_\sigma V(\pi_{h^\infty} f) \geq V(f) - C\varepsilon$, since σ is optimal.
- Since $V(f) \geq V(\pi f)$ for all $\pi \in \Pi_p$ and $p\pi_{h^\infty} \approx p$ for a set of histories with high probability, we have $V(f) \geq V(\pi_{h^\infty} f) - C\varepsilon$.
- We now argue that $V(\pi_{h^\infty} f) \geq V(f) - C\varepsilon$ for a large set of histories, so that undetectable strict efficiency tells us that $\pi_{h^\infty} f \approx f$, and we are done.

Step 3 cont'd

Consider the random variable X with

$$X(h^\infty) := V(f) - V(\pi_{h^\infty} f).$$

We have the following:

- $\mathbb{E}_\sigma X \leq C\varepsilon$.
- $\mathbb{P}_\sigma(\{h^\infty : X(h^\infty) \leq -C\varepsilon\}) \leq \zeta$, with ζ small (in fact, can be made arbitrarily small).
- $X \leq K$, a positive constant (say twice the maximal payoff).

X then satisfies: For any $\nu > 0$,
$\mathbb{P}_\sigma(\{h^\infty : X(h^\infty) \geq \nu\}) \leq (2C\varepsilon + K\zeta)/\nu$.

Idea: Write $X^+ = X + X^-$, note that $\mathbb{E}_\sigma X^- \leq C\varepsilon + \zeta K$, and apply Markov to X^+.

Modifications required in general

- Change the test so as to test the transitions and independence. This insures that a truthful agent can guarantee himself about $V_i(f)$, regardless of the strategies of the others.
- This also insures that step 1 remains valid (since if the empirical transition is close to P, the empirical frequency is close to p, the invariant distribution). This is all we need for step 1.
- From the efficiency condition, there exists $\lambda \in \Delta(\mathcal{I})$ such that $\langle \lambda, V(f) \rangle \geq \langle \lambda, V(\pi_{h^\infty} f) \rangle - C\varepsilon$.
- Repeat the step 3 as above with $X(h^\infty) := \langle \lambda, V(f) \rangle - \langle \lambda, V(\pi_{h^\infty} f) \rangle$.

Related literature

- Repeated implementation problems: Lee and Sabourian (2011), Mezzetti and Renou (2012).
- Linking implementation problems: Townsend (1982), Jackson and Sonnensheim (2007), Matsushima and al. (2010).
- "Review" mechanisms: Radner (1985), Gossner (1995), Renault, Solan and Vieille (2012), Escobar and Toikka (2012), among others.
- Undetectability: Lehrer (1989), Renault and Tomala (2004), Rahman (2010), among others.

Concluding remarks

- Necessity: For iid processes, the weak efficiency of f in the set $\{\pi f : \pi \in \Pi_p^\times\}$, with $\Pi_p^\times := \{\pi : \pi = \otimes_i \pi_i, p_i\pi_i = p_i, \forall i\}$ the set of undetectable independent deceptions, is necessary.

- For persistent processes, one can add the condition of Renault, Solan, and Vieille (2012) on the transitions.

- Open issues: more general stochastic processes, interdependent values.

- General issue: how to design a test that a truthful person can pass, regardless of the behaviors of the others. And simultaneously guaranteeing that an untruthful player is appropriately punished, unless his behavior is indistinguishable from truth-telling.