Repeated Implementation with Incomplete Information

Jihong Lee
Seoul National

Hamid Sabourian
Cambridge

December 2015

Motivation & Overview

- Implementation problem: Does there exist a mechanism such that, for each state, *every* equilibrium of the mechanism implements the socially desired outcome(s) defined by the SCR/SCF?
- The literature has been almost entirely concerned with implementing a SCR in one-shot settings
- Many real world institutions are used repeatedly, e.g. markets, voting, contracts.
- What is generally implementable in repeated contexts?

  Maskin monotonicity is a very demanding condition, esp. with incomplete information.

Motivation & Overview

- Our setup:
  - infinitely-lived agents with state-dependent utility functions (not necessarily transferable)
  - at each (discrete) period, a state drawn i.i.d.
  - at each period players learn about the state
  - aim to repeatedly implement an SCF in each period at every possible history
- The planner *commits to* a mechanism *for each period*

  at each date, learns the outcome of past mechanisms but not the current or past states

This Paper

- Repeated implementation with incomplete information:
  - Agents' utilities may be interdependent and their signals correlated.
  - Bayesian Nash equilibrium as solution concept
- Conditions for efficient repeated implementation
  - Efficiency in the range + (one-shot) IC for the general information structure
  - Efficiency in the range + payoff identifiability with interdependent values
  - Approximate implementation: efficiency in the range + "pairwise" identifiability + Cremer-McLean (with interdependent values)

Related Literature

- Repeated implementation with complete information
  - Lee and Sabourian (2011)
  - Lee and Sabourian (2015)
  - Mezzetti and Renou (2015)
  - Azacis and Vida (2015)
- Repeated implementation with incomplete information
  - Jackson and Sonnenschein (2007)
  - Renou and Tomala (2015)

Basic Definitions and Notation

Implementation problem with incomplete information
$\mathcal{P} = [I, A, (\Theta_i)_i, p, (u_i)_{i \in I}]$

- I finite agents
- A finite set of social outcomes
- $\Theta = \Theta_1 \times \cdots \times \Theta_I$ finite set of type profiles/states
- p probability distribution on Θ
  $p_i(\theta_i)$ marginal probability of type θ_i
  $p_i(\theta_{-i}|\theta_i)$ conditional probability of θ_{-i} given θ_i
- $u_i : A \times \Theta \to \mathbb{R}$ agent i's state-dependent utility

Basic Definitions and Notation

- $f : \Theta \to A$ social choice function (SCF)

  F set of all possible SCFs

  $f(\Theta)$ is the *range* of f
- Mechanism (in normal form) $g = (M, \psi)$
  - $M = M_1 \times \cdots \times M_I$ messages
  - $\psi : M \to A$ outcome function.
- Define

$$
\begin{aligned}
v_i(a) &= \sum_\theta p(\theta) u_i(a, \theta) \\
v_i(f) &= \sum_\theta p(\theta) u_i(f(\theta), \theta) \\
v(f) &= (v_1(f), .., v_I(f)) \\
V &= \left\{ v(f) \in \mathbb{R}^I : f \in F \right\}
\end{aligned}
$$

Basic Definitions and Notation

- Efficiency

  An SCF f is *efficient* if there exists no $w \in co(V)$ that weakly Pareto dominates $v(f)$:

  $\nexists\ w \in co(V)$ s.t. $w_i \geq v_i(f)\ \forall i$ and $w_i > v_i(f)$ for some i

- *Strictly efficient* if efficient $+ \nexists\ f' \neq f$ s.t. $v(f') = v(f)$

- *Strongly* efficient if strictly efficient $+\ v(f)$ is an extreme point of $co(V)$.

Basic Definitions and Notation

- Efficiency *in the range* defined w.r.t. $co(V(f))$

$$V(f) = \{v \in \mathbb{R}^I :$$
$$v = \sum_{\theta_1,\ldots,\theta_I \in \Theta} p(\theta_1,\ldots,\theta_I) u\left(f(\lambda_1(\theta_1),\ldots,\lambda_I(\theta_I)),\theta_1,\ldots,\theta_I\right)$$
$$\text{for some } \lambda_i : \Theta_i \to \Theta_i\}$$

Bayesian Repeated Implementation

$\mathcal{P}^\infty$ represents infinite repetition of $\mathcal{P} = [I, A, \Theta, p, (u_i)_{i \in I}]$

- Period $t \in \mathbb{Z}_{++}$
- θ drawn independently according to p in each period
- Outcome sequence $a^\infty = \left(a^{t,\theta}\right)_{t,\theta}$
- Discounted average expected utility; common $\delta \in (0, 1)$

$$\pi_i(a^\infty) = (1 - \delta) \sum_{t=1}^{\infty} \sum_{\theta \in \Theta} \delta^{t-1} p(\theta) u_i(a^{t,\theta}, \theta)$$

- The structure of $\mathcal{P}^\infty$, including δ, is common knowledge among the agents and, if exists, the planner.

Regimes and Strategies

- G the set of all feasible mechanisms
- H^∞ the set of all possible "public" histories of mechanisms and corresponding actions that are publicly observable
- $\mathbf{h} = (h, \theta(t)) \in \mathbf{H}^\infty$ the set of "full" histories

  where $\theta(t) = (\theta^1, \ldots, \theta^{t-1})$
- $\mathbf{h}_i = (h, \theta_i(t)) \in \mathbf{H}_i^\infty$

Regimes and Strategies

- A "regime", R, is a set of transition rules

$$R : H^{\infty} \to G$$

  Note that the planner/society commits to a regime

- For any regime R, each agent i's (private) strategy, σ_i, is

$$\sigma_i : \mathbf{H}_i^{\infty} \times G \times \Theta_i \to \Delta(\cup_{g \in G} M_i^g)$$

Nash Repeated Implementation

- $Q(R, \delta)$ set of Bayesian Nash equilibria in regime R with δ
- More notation: for any R, σ, let
  - $\mathbf{H}^t(\sigma, R)$ the set of full histories that occur with pos. prob. at date t.
  - $A^{h,\theta}(\sigma, R)$ the set of outcomes occurring with pos. prob. at $\mathbf{h} = (h, \theta(t))$
  - $E_{\mathbf{h}}\pi_i^\tau(\sigma, R)$ i's (expected) continuation payoff at period $\tau \geq t$ conditional on $\mathbf{h}$
    - This involves posterior "beliefs"
  - $E\pi_i^\tau(\sigma, R)$ is the expected payoff at period $\tau \geq t$ evaluated from $t = 1$

Nash Repeated Implementation

- An SCF f is *payoff-repeatedly implementable* in Bayesian Nash equilibrium from period τ if $\exists$ a regime R s.t.
  1. $Q(R, \delta)$ is non-empty
  2. every $\sigma \in Q(R, \delta)$ is s.t. $E\pi_i^t(\sigma, R) = v_i(f)$ for all i and $t \geq \tau$.
- An SCF f is *repeatedly implementable* in Bayesian Nash equilibrium from period τ if $\exists$ a regime R s.t.
  1. $Q(R, \delta)$ is non-empty
  2. every $\sigma \in Q(R, \delta)$ is s.t. $A^{h,\theta}(\sigma, R) = \{f(\theta)\}$ for any $t \geq \tau$, $\mathbf{h} \in \mathbf{H}^t(\sigma, R)$ and $\theta \in \Theta$.

Obtaining Target Payoffs

- Dictatorships
  - For $C \subseteq A$, $d^i(C)$ is i-dictatorship over C
  - $v_i^i(C)$ is i's dictatorial payoff; $v_j^i(C)$ is j's maximal payoff
  - with private values, $v_i^i(C) \geq v_i(f)$ if $f(\Theta) \subseteq C$ and $v_i^i(C) \geq v_i^j(C)$

Obtaining Target Payoffs

- Condition ω
  - For each $i \in I$, $\exists\ \tilde{a}^i \in A$ such that $v_i(\tilde{a}^i) \leq v_i(f)$
- Condition υ
  - For each $i \in I$,$\exists\ C^i \subseteq A$ s.t.

    (a) $v_i^i\left(C^i\right) \geq v_i(f)$.

    (b) $\exists\ i, j$ s.t. $v_i^i\left(C^i\right) > v_i(f)$ and $v_j^j\left(C^j\right) > v_j(f)$.
  - With private values, we can take, for instance, $C^i = f(\Theta)$ for part (a).

Obtaining Target Payoffs

- By Sorin (1986), we have the following:

  Lemma: Consider f satisfying conditions ω and υ. Then, if $\delta > 1/2$, $\exists$ a history-independent regime S^i for each i that generates a unique payoff to i exactly equal to $v_i(f)$.

- Fudenberg and Maskin (1991): We can also make the regime s.t. continuation payoffs always approximate $v_i(f)$ if δ close enough to 1.

- Fix $\delta > 1/2$; also $\tilde{a}^i$ and C^i for each i.

Sufficiency Results with IC

- An SCF f is *incentive compatible* if, for any i, θ_i and θ_i',

$$v_i(f|\theta_i) \geq \sum_{\theta_{-i}\in\Theta_{-i}} p_i(\theta_{-i}|\theta_i)u_i(f(\theta_{-i},\theta_i'),(\theta_{-i},\theta_i))$$

- **Theorem 1:** Fix any $I \geq 2$. If f is efficient, incentive compatible and satisfies conditions ω and υ, f is payoff-repeatedly implementable in Bayesian Nash equilibrium from period 2.
- With strong efficiency, we obtain repeated implementation in terms of outcomes (from period 2).
- Efficiency in the range with slight strengthening of ω & ν.

Regime Construction

- Mechanism $b^* = (M^*, \psi^*)$:
  (i) For all i, $M_i^* = \Theta_i \times \mathbb{Z}_+$, where $\mathbb{Z}_+$ non-negative integers.
  (ii) For any $m = ((\theta_i, z^i))_{i \in I}$, $\psi^*(m) = f(\theta_1, \ldots, \theta_I)$.
- Regime B^* with the following transition rules:
  1. $B^*(\emptyset) = b^*$
  2. At any period $t \geq 1$

     [A] if all announce zero ("agreement"), then play b^*

     [B] if all but one ("odd-one-out") i announce zero then the cont. regime is S^i (with payoff $v_i(f)$ to i)

     [C] otherwise, permanent dictatorship to the agent i with highest integer D^i (C^i).

Steps of the Proof: Existence

IC ensures existence of Markov eqm in which agents always tell the truth and 0.

- IC deters deviation to lying
- Rule [B] of the regime construction deters deviation to a pos. integer.

Steps of the Proof: Characterization

Fix any equilibrium.

1. By Rule [B], at any date t and history $\mathbf{h}$, if b^* is to be played,

   $$E_{\mathbf{h}}\pi_i^{t+1} \geq v_i(f) \quad \forall i$$

   $\Rightarrow$ Otherwise, "pre-emptive deviation" to the highest integer.

2. Suppose b^* is played at all possible histories up to t. Then, by the previous step, "efficiency" implies, $\forall i$,

   $$E\pi_i^{t+1} = v_i(f)$$

   which, in turn, implies

   $$E_{\mathbf{h}}\pi_i^{t+1} = v_i(f).$$

   We circumvent the issue of tracking posterior beliefs

Steps of the Proof: Characterization

3 By the previous step, condition υ (part (b)) and Rule [C], induction shows agents always announce 0 and hence b^* is always played on path.

$\Rightarrow$ Otherwise, some agent would deviate to the highest integer.

- With complete information, we can show this step first, but here, incomplete information makes things trickier.
- Lee ans Sabourian (2015) replace the integer arguments with "finite" mechanism under complete information.

Repeated Implementation without IC

- Our eqm characterization above does not depend on the particular information structure

  Or agents' precise knowledge about others' information, or posterior beliefs

- Existence poses an issue:

  Either (one-shot) IC is satisfied, or there has to be some avenue via which deviation can be detected and subsequently punished

- Two approaches:

  1. Interdependent values and payoff identifiability
  2. Approximate implementation

Interdependent values and Payoff identifiability

- An SCF f is payoff-identifiable if, for any i, any $\theta_i, \theta_i' \in \Theta_i$ and any $\theta_{-i} \in \Theta_{-i}$, $\exists$ some $j \neq i$ s.t.

$$u_j(f(\theta_i', \theta_{-i}), \theta_i', \theta_{-i}) \neq u_j(f(\theta_i', \theta_{-i}), \theta_i, \theta_{-i}).$$

  Thus, if an agent deviates from a truth-telling eqm then there will be at least one other agent who can detect the lie at the end of the period (assuming utilities are learned at the end of the period)

  One can then build intertemporal punishment

Interdependent values and Payoff identifiability

- Condition ω^*:

  There exists $\tilde{a} \in A$ ("bad outcome") such that $v_i(\tilde{a}) < v_i(f)$ for all i.

- **Theorem 2:** Consider the case of interdependent values, and suppose f satisfies efficiency, payoff-identifiability, and conditions ω^* and υ. Then f is payoff-repeatedly implementable from $t = 2$ if δ is sufficiently large.

Regime Construction

- Integer-only mechanism Z
- Two-stage mechanism $\tilde{b}^*$
  - Stage 1 - Each i reveals θ_i; $f(\theta_1, \ldots, \theta_I)$ implemented
  - Stage 2 - Each i chooses from $\{NF, F\} \times \mathbb{Z}_+$ (i.e. whether to "flag" $+$ integer)

Regime Construction

- Regime $\widetilde{B}^*$:
  1. Starts with Z
  2. At any period $t \geq 1$
     - If all 0 then play $\tilde{b}^*$
     - If "odd-one-out" i in integer + all "NF" (from $t \geq 2$) then S^i
     - If at least one "F" then "bad outcome" $\tilde{a}$ forever
     - Otherwise, dictatorship to the highest integer $D^i(C^i)$

Equilibria

- Truth-telling, zero integer and no flag for all is a BNE if players are sufficiently patient
- Punishment: upon deviation from truth-telling, the "deviator" and the "detector" (payoff identifiability) flag to activate the bad outcome forever
  - This is mutually optimal; we can modify the continuation regime to make it also "strictly"' BR
- No incentive to play positive integer or flag.

Equilibria

- For characterization, we follow similar inductive steps as before.
- Additionally, we have show that "F" cannot occur on eqm path – again, such coordination failure can be prevented by the possibility of "pre-emptive" deviation.
- Some remarks

  \- Z in $t = 1$ to ensure no mis-coordination on "F"

  \- Integer play cannot occur before flagging

Approximate Implementation

- Without IC or (payoff) identifiability, "approximate" implementation may be the best we can hope for.
- f is ϵ-payoff-repeatedly implementable in BNE if, for any $\epsilon > 0$, $\exists$ regime R and $\bar{\delta} \in (0,1)$ s.t., for any $\delta \in (\bar{\delta}, 1)$, $Q^\delta(R)$ is non-empty and every $\sigma \in Q^\delta(R)$ is s.t., for every $t \geq 1$ and i,

$$\mid E\pi_i^t(\sigma, R) - v_i(f) \mid < \epsilon$$

Regime Construction

- Assume condition ω^*
- Mechanism b^{**}
  - For each i, $M_i = \Theta_i \cup \{N\} \times \mathbb{Z}_+$
  - If anyone reports N then implement $\tilde{a}$ ("bad outcome"); otherwise, same as b^*
- Regime B^ϵ for $\epsilon > 0$
  1. Starts with b^{**}
  2. At any period $t \geq 1$

     - if all 0 then play b^* next period

     - if "odd-one-out" i then $S^i(\eta)$ s.t. i gets payoff $v_i(f) - \eta$, with $\eta(\epsilon) \in (0, \epsilon)$ set precisely

     - otherwise, permanent dictatorship to i with highest integer $D^i\left(C^i\right)$.

Equilibria

- Restrict attention to *public* strategies
- Characterization:

  Efficiency $\Rightarrow$ As before, equilibrium continuation payoffs are within ϵ of $v(f)$
- Existence:
  - We adapt the techniques of Fudenberg, Levine and Maskin (1994)
  - Sufficiency conditions depend on the information structure: private vs. interdependent values

Sketch of the Existence Proof

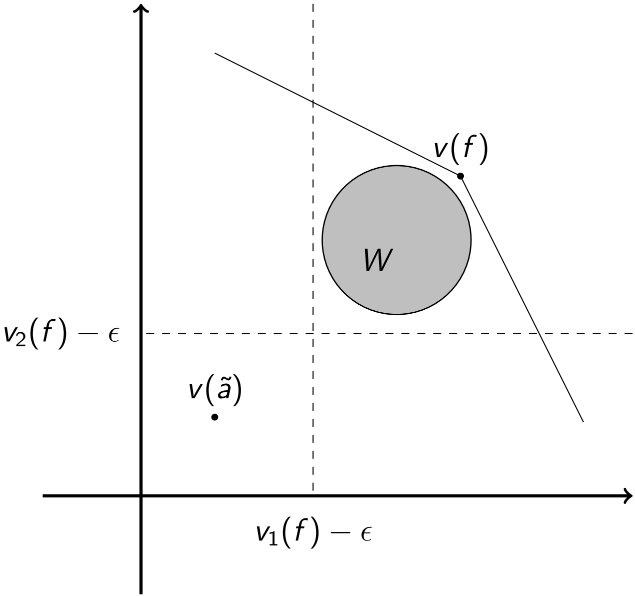

Sketch of the Existence Proof

- First consider an auxiliary stationary regime such that:
  - Mechanism $\tilde{b}^{**}$ identical to b^{**} except that no integer
  - One NE of this stage game is everyone choosing "N" with bad outcome.
  - $\overline{s}_i : \Theta_i \rightarrow \Theta_i$ "revelation strategy"
- FLM obtain a folk theorem of the corresponding repeated Bayesian game for the case of private values
- We will use an eqm of this game to construct an eqm of our regime B^ϵ

Sketch of the Existence Proof

- Nash reversion: $\tilde{a}$ is a Nash outcome of the stage game.
- Full dimensionality:

  $V^*(f) = \{v \in co(V(f)) : v > v(\tilde{a})\} \cup v(\tilde{a})$ has non-empty interior.

- We need "enforceability" of the efficient revelation profiles. This comes for free with "private values" (FLM).
- Information structure is "pairwise identifiable" w.r.t. $V(f)$ if every revelation profile $\bar{s}$ inducing efficient payoffs in $V(f)$ is pairwise identifiable for every pair of players.[1]

  This comes for free with IPV (FLM).

[1] A revelation profile is pairwise identifiable if, for every pair of players, the report distributions from one player's unilateral deviations are distinct from the other's deviations

Sketch of the Existence Proof

- With the conditions listed above, FLM's construction applies to the auxiliary regime (repeated adverse selection)
- Eqm profile of the aux. regime + always announcing 0 is a PPE of our original regime B^ϵ
  - Deviating to a positive integer by player i leads to a continuation payoff $v_i(f) - \eta$.
  - But, this is less than the eqm cont. payoff (with δ very close to 1).

Results: Private Values

- **Theorem 3:** Consider private values and f that satisfies efficiency and conditions ω^* and υ. Suppose:
  - information structure is pairwise identifiable w.r.t. $V(f)$
  - $V^*(f) = \{v \in co(V(f)) : v > v(\tilde{a})\} \cup v(\tilde{a})$ has non-empty interior.

  Then, f is ϵ-payoff-repeatedly implementable in public Bayesian Nash equilibrium.

- **Corollary:** With *independent* private values, the same is true without the pairwise identifiability condition.

Results: Interdependent Values

- Cremer and McLean (1988)

  The information structure satisfies condition CM if, for each i, any θ_i and any $\mu_i : \Theta_i \to \mathbb{R}_+$,

  $$p_i(\theta_{-i} \mid \theta_i) \neq \sum_{\theta_i' \neq \theta_i} \mu_i(\theta_i') p_i(\theta_{-i} \mid \theta_i').$$

  i.e. No player i of type θ_i could generate the same conditional probabilities on the types of the other players through a random untruthful reporting strategy. (Correlated types)

- With condition CM, we can extend FLM's arguments:

  Any 1-to-1 revelation profile is *enforceable*.

Results: Interdependent Values

- $\widetilde{V}(f) = \{v : v = \sum_\theta u(f(\lambda_1(\theta_1), \ldots, \lambda_I(\theta_I)), \theta)p(\theta)$ for some 1-to-1 functions $\lambda_i : \Theta_i \to \Theta_i\}$
- **Theorem 4:** Consider interdependent values, and f that satisfies efficiency and conditions ω^* and υ.

  Suppose the information structure satisfies condition CM and pairwise identifiability w.r.t. $\widetilde{V}(f)$;

  $\widetilde{V}^*(f) = \{v \in co(\widetilde{V}(f)) : v > v(\tilde{a})\} \cup v(\tilde{a})$ has non-empty interior.

  Then, f is ϵ-payoff-repeatedly implementable in public Bayesian Nash equilibrium.

Conclusion

- This paper sets up the general problem of repeated implementation with incomplete information
  and demonstrates the extent to which "efficient" social choices can be repeatedly implemented in BNE.
- Is efficiency (in the range) necessary?
  - Difficult to detect deviation from collusive behavior
  - The regime may possess complex transitional structures off-the-equilibrium which make any collusion difficult to sustain
- For exact implementation, we also need IC or (payoff) identifiability.