Repeated Implementation: Incomplete Information

Hamid Sabourian

CSEF Capri 23
23 June 2014

Some of the work joint with Jihong Lee

Motivation and Overview

- Implementation theory is concerned with constructing a mechanism such that, for each state, *every* equilibrium of the mechanism implements the socially desired outcome(s) defined by the SCR/SCF.

- Until recently, 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?

- Necessary and "almost" sufficient condition for one-shot Nash implementation with complete information is Maskin monotonicity. With one-shot incomplete information the critical conditions are Bayesian monotonicity and IC.

- These are all strong conditions.

- Do we need something similar (or more) in repeated contexts?

Our infinitely repeated implementation setup

- State-dependent utility functions (not necessarily transferable)

- At each period, a state drawn i.i.d.

- At each period players learn the state (complete information) or have partial/incomplete information regarding the state

- Aim to repeatedly implement an SCF (an 'optimal' choice in each state of the world) in each period at every possible history

- At the beginning of the game the planner has to devise (commit to) a mechanism *for each period* of the repeated game

- At each date the planner knows the outcome of past mechanisms but not the current or past states

A number of applications naturally fit this description

- In repeated voting the voters' preferences could follow a stochastic process, with the planner's goal being for instance to always enact an outcome that is Condorcet-consistent.

- In repeated auctions, the bidders' valuations could follow a stochastic process, with the planner's goal being for instance to always enact an outcome sells each object to the bidder with highest valuation

- A community that collectively owns a technology could repeatedly face the problem of efficiently allocating resources under changing circumstances.

Fundamental differences between one-shot and repeated implementation

1. **History dependence of players' choices and "collusion":** Repeated framework allows players to condition their strategies on past histories and thereby can induce a large no of “collusive” equilibria:

   - E.g. if the same mechanism enforced at each date then agents play a repeated game with random states; then by the FT any IR outcome is an equilibrium with complete information if players are sufficiently patient.

   - Seems to make implementation a hard problem

2. **History dependence of the mechanisms:** The planner chooses/commits to a sequence of mechanisms (referred to as a “regime”); thus, can condition the choice of mechanism at each period on past observable outcomes.

   - This gives the planner additional leverage of altering the future mechanisms in a way that rewards desirable behaviour while punishing undesirable

   - May make the implementation problem easier

Lee and Sabourian, *Econometrica 2011 (*henceforth LS)

- Assuming complete information it is shown, with some minor qualifications, that if players are sufficiently patient, efficiency (in the range) of the SCF is both necessary and sufficient for repeated implementation in Nash equilibrium.
  - Efficiency is necessary because otherwise the players may collude and achieve a higher payoff (as in the Folk Theorem).
  - The sufficiency result is established by constructing a history-dependent sequence of mechanisms that give each agent the incentive to deviate earlier from any future outcome that does not guarantee his target payoffs given by the SCF. This means that, for each player, continuation payoffs are bounded below by his target payoff. Efficiency of the target payoff then implies that the continutaion payoffs coincide with the target payoff for all players.
    - Maskin Monotonicity is not needed; efficiency is all that is required.
    - The construction of history-dependent mechanisms is very simple and does not involve complicated punishment and reward schemes/mechanisms.
    - No credibility or specific belief specification off-the-eq needed to sharpen predictions.

Incomplete Information and Results

- In this work, we consider the case of agents having incomplete information.

Sufficiency

- Additional complications with incomplete information:
  - Existence and detecting deviations:
    In addition to designing a dynamic mechanism that avoids the multiplicity of "collusive" equilibria, the designer must ensure that the mechanism has an equilibrium. With complete information, existence is trivial as deviations by any player can be detected by others. Not so with incomplete information, as deviations is hard to detect and hence existence of an equilibrium, without additional assumptions, becomes a difficult task.
  - Private equilibria and beliefs:
    Players can condition their strategies on their past private information; hence, continuation payoffs may depend on the agent's ex post beliefs about the others' past private information.

- We first show that any IC and efficient (in the range) SCF can be repeatedly implemented.
  - IC solves the existence problem trivially.
  - We avoid the issue of agents' ex post private beliefs, by considering repeated implementation in terms of *expected* continuation payoffs, where the expectation is *computed at the beginning at* $t = 1$.
  - The result is for a fairly general set-up: private or interdependent values, independent or correlated information
  - Bayesian monotonicity (used in one-shot implementation) is not needed..
- But IC is a strong condition; particularly with interdependent values where the conflict between efficiency and IC can be very severe (Maskin,1992 and Jehiel-Moldovanu, 2001).
- So our second set of results involves doing away IC and using inter-temporal incentives to achieve the existence of an equilibrium.
- Very informally, we show in a general set-up that, with sufficiently patient players, any efficient in the range SCF can be implemented (approximately implemented) if deviations can be identified (statistically identified).

Neccesity

- Extending LS result necessity result to incomplete information is made more complicated because it is difficult for players to detect deviations:
  - In LS the necessity argument is by contradiction: if a SCF were not efficient in the range and is implementable, one can construct another equilibrium that results in a Pareto superior outcome. Deviations from the new equilibrium is deterred by the threat of playing the original Pareto inferior equilibrium.
    - But with incomplete information detecting deviations is hard; particularly hard when private strategies are allowed.
  - I have only a necessity result with public strategies when limited to the special case private values and independent signals.
- The rest of this talk will be on the sufficiency part.

Basic definitions & notation

An implementation problem: $\mathcal{P} = [I, A, (\Theta_i)_{i \in I}, p, (u_i)_{i \in I}]$ where

- I finite agents
- A finite set of social outcomes
- Each i learns a private information $\theta_i \in \Theta_i$
- $\Theta = \Theta_1 \times \cdots \times \Theta_I$ is the set of type profiles/states
- Type space Θ_i is finite
- 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

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

- F set of all possible SCFs

- $v_i(a) \equiv \sum_\theta p(\theta) u_i(a, \theta)$ i's (ex ante) expected utility if a

- $v_i(f) \equiv \sum_\theta p(\theta) u_i(f(\theta), \theta)$ i's (ex ante) expected utility if f

- $v(f) = (v_1(f), .., v_I(f))$ expected utility profile if f

- $V = \left\{ v(f) \in \mathbb{R}^I : f \in F \right\}$ feasible (one-shot) payoffs

- SCF f satisfies (one-shot) IC if for any i, $\theta = (\theta_i, \theta_{-i})$ and θ'_i

$$\sum_{\theta_{-i}} p_{-i}(\theta_{-i} \mid \theta_i) u_i(f(\theta_i, \theta_{-i}), \theta) \geq \sum_{\theta_{-i}} p_{-i}(\theta_{-i} \mid \theta_i) u_i(f(\theta'_i, \theta_{-i}), \theta)$$

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

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

- SCF f is *strictly efficient* if (i) efficient and (ii) $\nexists\ f' \neq f$ s.t. $v(f') = v(f)$

- SCF f is *strongly efficient* if (i) strictly efficient and (ii) extreme point of the set $co(V)$.

- We also define a weaker concept of (individually) efficient in the range.

- This weaker notion of efficiency compares f only with SCFs that can be achieved through individual manipulations of f:

$$F(f) = \{f' : \forall\theta,\ f'(\theta) = f(s_1(\theta_1), .., s_I(\theta_i)) \text{ for some } s_1 : \Theta_1 \longrightarrow \Theta_1, .., s_I : \Theta_I \longrightarrow \Theta_I\}$$

- Let

$$V(f) = \{v(f') \in \mathbb{R}^I : f' \in F(f)\}$$

- A SCF f is *efficient in the range* if $\nexists\ w \in co(V(f))$ s.t. $w_i \geq v_i(f)\ \forall i$ and $w_i > v_i(f)$ for some i

- Similarly define SCF strictly/strongly efficient in the range

- Mechanism $g = (M, \psi)$
  - $M = M_1 \times \cdots \times M_I$ messages
  - $\psi : M \to A$ outcome function

- Trivial mechanisms:
  - constant outcome $a \in A$: $M_i = A$ for all i and $\psi(m) = a$ for all m
  - dictatorship of i: $M_i = A$ for all i and $\psi(m) = m_i$ for all m

- $v_i^i \equiv \sum_{\theta_i} p_i(\theta_i) \max_{a \in A} \sum_{\theta_{-i}} p_{-i}(\theta_{-i} \mid \theta_i) u_i(a, \theta)$ refers to i's (ex ante) expected payoff if i dictator
  (i's maximal expected utility given his own private information)

- $A^i(\theta_i) \equiv \arg\max_{a \in A} \sum_{\theta_{-i}} p_{-i}(\theta_{-i} \mid \theta_i) u_i(a, \theta)$ refers to the outcome set when i is the dictator and learns θ_i

- $v_i^j \equiv \sum_{\theta} p(\theta) \arg\max_{a \in A^j(\theta_j)} u_i(a, \theta)$ refers to i's maximal utility when j is the dictator

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
- At any t, each agent observes only his own type θ_i^t
- Outcome sequence $a^\infty = \left(a^{t,\theta}\right)_{t,\theta}$
- $\theta(t) = (\theta^1, .., \theta^{t-1}) \in \Theta^{t-1}$ finite history of states
- $q(\theta(t)) = \prod_{\tau=1}^{t-1} p(\theta^\tau)$

- H^t the set of all $(t-1)$ histories of mechanisms and corresponding actions/messages that are publicly observable (to agents and planner)

- $H^\infty = \cup_{t=1}^{\infty} H^t$ the set of all (finite) public histories

- $H_i^\infty = \cup_{t=1}^{\infty} (H^t \times \Theta_i^{t-1})$ the set of histories observable only to agent i (public histories and private information of i)

- Discounted average expected utility; common $\delta \in (0, 1)$

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

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

Regime

- G the set of all feasible mechanisms

- A "regime" (dynamic mechanism) R, is then a set of transition rules

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

- Note that the planner/society commits to a regime (a dynamic mechanism).

- The choice of mechanisms (transition rules) is deterministic; removing this restriction can expand what can be implementable.

Strategies

- For any regime R, each agent i's strategy, σ_i, is described by

$$\sigma_i : H_i^\infty \times G \times \Theta_i \to \cup_{g \in G} M_i^g$$

s.t. $\sigma_i(h, \theta_i(t), g, \theta) \in M_i^g$ for all $(h, \theta_i(t), g, \theta) \in H_i^\infty(t) \times G \times \Theta_i$.

- The exposition is for pure strategies; the results extend to mixed strategies.

Fix regime R. For any $\sigma, t, \theta(t) = (\theta^1, .., \theta^{t-1})$, and $h \in H^t$

- $a^{\theta^t}(\theta(t))$ denotes the outcome implemented given past states $\theta(t)$, current state θ^t and all playing according to σ

- $\pi_i^\tau(h, \theta_i(t))$ denotes i's expected continuation payoff from period $\tau \geq t$ conditional on observing $(h, \theta_i(t))$ and on all playing according to σ.

- $E\pi_i^\tau = \pi_i^\tau(H^1, \theta(1))$ expected cont. payoff from period τ, where the expectation is calculated from the beginning at period 1.

Bayesian Nash repeated implementation

- $\Omega(R, \delta)$ set of (pure) Bayesian Nash equilibria in regime R with δ

Definition 1: An SCF f is *payoff-repeated-implementable* in Bayesian Nash equilibrium from period τ if $\exists$ a regime R s.t.

1. $\Omega(R, \delta)$ is non-empty
2. every $\sigma \in \Omega(R, \delta)$ is s.t. $E\pi_i^t = v_i(f)$ for all i, at any $t \geq \tau$.

Definition 2: An SCF f is (outcome) *repeated-implementable* in Bayesian Nash equilibrium from period τ if $\exists$ a regime R s.t.

1. $\Omega(R, \delta)$ is non-empty
2. every $\sigma \in \Omega(R, \delta)$ is s.t. $a^{\theta^t}(\theta(t)) = f(\theta^t)$ for any $t \geq \tau$, $\theta(t)$ and θ^t.

Sufficiency result with IC

- The planner can always implement
  - a mechanism that induces a payoff profile v^i (by making i the dictator)
  - a mechanism that induces the payoffs that results for any constant outcome $a \in A$.
- Hence, if the players are sufficiently patient it can induce any convex combination of the payoffs associated with dictatorships and constant outcomes.
- Therefore, as in LS, for any target f,under two minor conditions, for each player i, the planner can always implement a regime that induces the target payoff $v_i(f)$.
- Two minor conditions that ensure this are:

Condition η: An SCF f satisfies condition η if for each i, $v_i^i \geq v_i(f)$

- This condition is trivially satisfied with private values or perfect correlation.

Condition ω: An SCF f satisfies condition ω if for each i there exists some $\tilde{a}^i \in A$ s.t. $v_i(f) \geq v_i(\tilde{a}^i)$:

- Thus, the expected utility from the SCF for each agent is bounded below by that of some constant SCF.

- It is a weak condition as $\tilde{a}^i$ may be different for different players and it holds on average and not in each state. It is weaker than the standard 'bad outcome' assumption.

- It is satisfied for a very large no of applications.

Lemma 1 Consider f satisfying condition η and ω. Fix any i and any $\delta > \frac{1}{2}$. Then, $\exists$ a regime S^i s.t. $\pi_i(S^i) = v_i(f)$.

Proof:

1. If i is a dictator, he gets $v_i^i \geq v_i(f)$.

2. Given condition ω, there exists either some $\tilde{a}^i$ such that $v_i(f) \geq v_i(\tilde{a}^i)$.

3. Then using Sorin (1986), we can find a regime alternating i-dictatorship and trivial enforcement of $\tilde{a}^i$ s.t. agent i derives a payoff exactly $v_i(f)$.

- We need $\delta > \frac{1}{2}$ for this construction; assumed throughout in what follows.

- No restriction on δ is needed if either (i) the outcome space sufficiently rich so that $\exists$ some $\tilde{a}^i$ such that $v_i(\tilde{a}^i) = v_i(f)$ or (ii) if the choice of a regime allows for random transitions.

- In addition to conditions ω and η, we need a third minor assumption.

  Condition ν: An SCF f satisfies condition ν if the inequality in condition η is strict for at least two players:

$$\text{there exist two agents } i \text{ and } j \text{ s.t. } v_i^i > v_i(f) \text{ and } v_j^j > v_j(f)$$

- None of the three conditions are necessary for our results and can be replaced by alternative weaker conditions

Theorem 1: Suppose $\delta > 1/2$. If f is efficient and IC, and it satisfies conditions η, ω and ν, then f is payoff repeated-implementable in Bayesian Nash equilibrium from period 2.

- Mechanism $g^* = (M, \psi)$ defined:

  (i) for all i, $M_i = \Theta_i \times Z_+$ (i.e. report "private information + non-negative integer")

  (ii) for any $m = ((\theta_i, z^i))_{i \in I}$, $\psi(m) = f(\theta_1, \ldots, \theta_I)$.

- R^* represents any regime satisfying:

  (i) $R^*(\emptyset) = g^*$

  (ii) For any $h = \big((g^1, m^1), \ldots, (g^{t-1}, m^{t-1})\big) \in H^t$ s.t. $t > 1$ and $g^{t-1} = g^*$:

  a. if $m_i^{t-1} = (\theta_i, 0)$ for all i ("agreement"), then $R^*(h) = g^*$;

  b. if there exists some i s.t. $m_j^{t-1} = (\theta_j, 0)$ for all $j \neq i$ and $m_i^{t-1} = (\theta_i, z^i)$ with $z^i \neq 0$ ("the odd-one-out"), then $R^*|h = S^i$

  c. if m^{t-1} is of any other type ("disagreement") and i is the lowest-indexed agent among those who announce the highest integer, then i dictator henceforth.

- Therefore,
  - agents start by playing g^*
  - anything other than "agreement" (all announcing zero) ends strategic play in this regime
  - Integers do not affect the outcome implemented in the current period but they determine the continuation mechanism.

Existence:

- "Truth-telling" and declaring zero at every date/state generates $v(f)$ and is Markov. Hence, by IC, it constitutes a BNE:
  - By IC any unilateral deviation from the truth cannot improve the immediate payoff. Since others are doing Markov, it cannot improve future payoffs.
  - Any unilateral deviation in terms of announcing a positive integer induces the odd-one-out and therefore no gain in terms of the continuation payoff.

Characterisation

Fix any equilibrium and any date t.

- **Step 1:** *For any history* $(h, \theta(t)) \in H^t \times \Theta^{t-1}$ *on the equilibrium path at which mechanism* g^* *is played, each* i*'s expected cont. payoff from next period* $\pi_i^{t+1}(h, \theta_i(t)) \geq v_i(f)$.

  - Otherwise, he could make himself better off at date t by announcing a positive integer at t that is higher than any other possible reported integer at this date, *everything else the same*; this does not affect the payoff at t but results in $v_i(f)$ or v_i^i from the next period.

- **Step 2:** *If the mechanism at every equilibrium history* $(h, \theta(t)) \in H^t \times \Theta^{t-1}$ *at date* t *is* g^*, *then for all* i, $E\pi_i^{t+1} = v_i(f)$, *and hence* $\pi_i^{t+1}(h, \theta_i(t)) = v_i(f)$ *for all equilibrium* $(h, \theta_i(t)) \in H^t \times \Theta^{t-1}$ *at date* t.

  – By Step 1, $\pi_i^{t+1}(h, \theta_i(t)) \geq v_i(f)$ for each equilibrium $(h, \theta(t))$

  – $E\pi_i^{t+1} = \sum_{(h,\theta(t))} \pi_i^{t+1}(h, \theta_i(t)) \Pr(h, \theta_i(t)) \geq v_i(f)$ for each i

  – Since $E\pi_i^{t+1} \in Co(V)$ and f is efficient, $E\pi_i^{t+1} = v_i(f)$ for each i

  – $E\pi_i^{t+1} = v_i(f)$ and $\pi_i^{t+1}(h, \theta_i(t)) \geq v_i(f)$ imply that $\pi_i^{t+1}(h, \theta_i(t)) = v_i(f)$ for each equilibrium $(h, \theta_i(t))$.

.

- **Step 3:** *If the mechanism at any equilibrium history* $(h, \theta(t)) \in H^t \times \Theta^{t-1}$ *at date* t *is* g^* *then each player anounces zero integer after any* $(h, \theta(t))$ *and* θ^t*) for any* θ^t*.*

  - If this were not the case some agent j would be announcing a positive integer after some $(h, \theta(t))$ and θ^t.

  - By condition v there exists agent $i \neq j$ such that $v_i^i > v_i(f)$

  - Then i can make himself better off by announcing a higher integer than others at $(h, \theta(t))$ after any state, everything else the same:

    * such a deviation only affects continuation payoff from the next period
    * it either makes i the dictator forever, hence v_i^i, or induces S_i, hence $v_i(f)$
    * by Step 2 no deviation induces $v_i(f)$

- Since the regime begins with g^*, by induction and Step 3, g^* will be played on the equilibrium path at every date. Hence, by Step 2, $E\pi_i^{t+1} = v_i(f)$ for all t.

Outcome implementation

- For outcome implementation we need strong efficiency: efficiency + no combination of other SCFs can generate the same payoff profile as f.

- Efficiency ensures that $\pi_i^{t+1}(h, \theta_i(t)) = \pi_i^{t+2}(h, \theta_i(t)) = v_i(f)$ for every t and i. If $v(f)$ cannot be generated by any combination of other SCFs, then the outcome at $t+1$ must implement f.

Efficiency in the range

- A similar set of results can be obtained for SCFs that are efficient *in the range* if the minor conditions are strengthened: E.g.

  - by restricting the dictatorships to choosing actions in the range of f and assuming the three conditions to this range

  - or by strengthening condition ν to: $\forall i, \exists\ j \neq i$ s.t. $v_j^j > \max\{v_j(f), v_j^i\}$

Ex post implementation

- *Ex post implementation:* The arguments for the incomplete information do not depend on the agents doing best responses to a given distribution of types. The results hold if we replace Bayesian eq by ex post equilibrium, appropriately defined.

Implementation without IC

- Without IC, we can show the following if agents are sufficiently patient:

A. Exact implementation

  - Any efficient SCF can be implemented if deviations can be detected by some player (this would require interdependent values).

B. Approximate implemetation

  - Any efficient SCF can be payoff approximately implemented for the following cases:
    - Private values and independent types
    - Private values with correlated types if Fudenberg-Maskin-Levine (1994) pairwise identifiabilty holds (i.e. deviations can be statistically distinguished)
    - Interdependent values if both Fudenberg-Levine-Maskin (1994) pairwise identifiabilty and if Cremer-Mclean (1988) condition hold.

Exact implementation with interdependent values

- Assume the agents know their utilities from the implemented outcomes at the end of each period.

- Identifiability assumption: For any i, any θ_i, θ_i' s.t. $\theta_i' \neq \theta_i$ and any θ_{-i}, $\exists$ some $j \neq i$ s.t.

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

  i.e. whenever there is one agent lying about his type while all others report their types truthfully, there exists another agent who obtains a (one-period) utility different from what he would have obtained under everyone behaving truthfully.

- Identifiability enable us to build a regime which admits a truth-telling equilibrium based on incentives of repeated play that involve punishments when someone misreports his type.

- Identifiability is in terms of payoffs; could be also written in terms of a general signal structure.

- To use intertemporal incentive we also need to strengthen condition ω to allow for the existence of a "bad outcome" (punishment) for *all* players:

  Condition ω': There exists $\widetilde{a} \in A$ s.t. $v_i(\widetilde{a}) < v_i(f), \forall i$

Theorem 2: Suppose $I \geq 2$ and δ is sufficiently close to 1. If f is efficient and identifiable and satisfies conditions η, ω' and υ, then f is expected payoff repeated-implementable in Bayesian Nash equilibrium from period 2.

- The idea behind the proof:
  - Modify the regime constructed with IC so that it is possible for players to "flag" a deviation by others and thereby induce a bad outcome.
  - Flaging can usually induce miscoordination type equilibria; here, this can be avoided by the possibility of making oneself "the-odd-out" before any miscoordination (this cannot be done at date 1). Otherwise, characterisation works as before.
  - Existence is obtained by constructing (credible) equilibrium strategies that anounce zero, tell the truth and flag deviations from truth-telling.

- Define Z as a mechanism for all i, $M_i = Z_+$ and for all $\psi(m) = a$ for some a

- Define $\widetilde{b}$ as the following extensive form mechanism:

  – Stage 1 - Each i announces θ_i, and $f(\theta_1, ..\theta_I)$ is implemented.
  – Stage 2 - Once agents learn their utilities, but before a new period begins, each announces a report belonging to $\{NF, F\} \times Z_+$.

- Stage 2 of $\widetilde{b}$ do not affect the outcome implemented in the current period but they determine the continuation play in the regime below

- Let $\widetilde{B}$ represent any regime satisfying the following transition rules:
  - $\widetilde{B}(\emptyset) = Z_+$;
  - For any $h = (g^1, m^1)$:
    - if $m_i^1 = 0$ for all i, then $\widetilde{B}(h) = \widetilde{b}$;
    - if there exists some i such that $m_j^1 = 0$ for all $j \neq i$ and $m_i^1 \neq 0$, then $\widetilde{B} \mid h = S^i$;
    - if $m^1$is of any other type and i is the lowest-indexed among those who announce the highest integer, then $\widetilde{B} \mid h = D^i$;
  - For any $h = ((g^1, m^1), ..(g^{t-1}, m^{t-1}) \in H^t$ s.t. $t > 2$ and $g^{t-1} = \widetilde{b}$:
    - if m_i^{t-1} is s.t. every agent reports NF and zero integer in Stage 2, then $\widetilde{B}(h) = \widetilde{b}$;
    - if m_i^{t-1} is s.t. at least one agent reports F in Stage 2, then $\widetilde{B} \mid h = \Phi^{\widetilde{a}}$;
    - if m_i^{t-1} is s.t. every agent reports NF and every agent except some i announces zero integer, then $\widetilde{B} \mid h = S^i$;
    - if m^{t-1} is of any other type and i is the lowest-indexed among those who announce the highest integer, then $\widetilde{B} \mid h = D^i$.

- This regime begins with Z. If all report zero, then the next period is $\widetilde{b}$; otherwise, either S^i or D^i for some i.

- $\widetilde{b}$ sets up two reporting stages:
  - first, report the private information
  - second, report detection of a lie by raising "flag" plus integer play

  Transitions being the same as the regime with IC above with "no flag". One "flag" overrules the integers and activate permanent implementation of the bad outcome

- Integers reported also in Stage 2 because otherwise flag can be used to stop the cont. play resulting from integers

Characterisation of equilibria

- As before, except that need to ensure no flagging on the equilibrium path at any $t > 1$
  - if flagging on the equilibrium path at any $t > 1$ then incentive to announce a positive integer at $t - 1$;
- Mechanism $\widetilde{b}$ can induce flagging at $b = 1$ (there is no date before). Hence, we have mechanism Z at $t = 1$.

Existence

- Truth-telling, zero integer and flag iff deviations for truth-telling by someone (including oneself) is a BNE if players are sufficiently patient:
  - Clearly, positive integers and flag cannot make a player better off
  - Not telling the truth will be detected and will induce the bad outcome
- The above strategy profile is also sequentially rational because if deviations from the truth then given that one other player flags, it is a best response to also flag.
- The mutual optimality of flagging after deviations,however, is supported by an indifference argument. The following modification makes it strictly optimal to flag given that there is another flag:
  - if a subset of agents $J \subset I$ flag, the continuation regime makes each player $j \in J$ the dictator in the first j-th period of the continuation regime while implementing $\widetilde{a}$ in every other period.

Approximate implementation

- Identifiability ensures that deviations are detected. Not always possible - private values.

- Without IC and identifiability, it may still be possible to provide intertemporal incentives by detecting deviations probabilistically; but then approximate implementation may be the best we can hope for.

Definition 2: An SCF f is approximately payoff-repeated-implementable in Bayesian Nash equilibrium if for any $\epsilon > 0$ there exists a regime R^ϵ and $\overline{\delta}$ s.t., for all $\delta > \overline{\delta}$, (i) $\Omega(R^\epsilon, \delta)$ is non-empty; and (ii) every $\sigma \in \Omega(R^\epsilon, \delta)$ is s.t, for every $t \geq 1$, $\mid E\pi_i^t - v_i(f) \mid < \epsilon$, for all i.

- By modifying the regime used with IC and by using linear programing arguments to construct appropriate continuation payoffs, we can establish approximate implementability of efficient SCFs without IC.

- In what follows, we establish the results restricting attention to public strategies in order to simplify the exposition. Similar set of results can be established with private strategies.

Private values with independent signals

- Let $V^* = \{v \in V : v \geq v(\widetilde{a})\}$.

Theorem 3: Suppose private values with independent signals. Assume also that f is efficient, and satisfies conditions ω' and ν (with private values η trivially holds). If the set CoV^* is full-dimensional (non-empty interior) then f is approximately payoff repeated-implementable in public strategies.

Sketch of the proof

- We need to modify the regime R^* used with IC so that we can appeal to some results by FLM (1994) on repeated games with imperfect public monitoring.

- Fix the degree of approximation $\epsilon > 0$. Then the modified regime R^ϵ has two new features:

  - when agents are called to announce a state, each can ensure that the bad outcome $\widetilde{a}$ is implemented in that period;

  - each agent i can guarantee himself $v_i(f) - \phi$, for some small $\phi > 0$ that depends on on the degree of approximation ϵ, when he is the odd one out.

- Specifically, define modified mechanism $b^* = (M, \psi)$:

  (i) for all i, $M_i = Y_i \times Z_+$, where $Y_i = \Theta_i \cup N$

  (ii) for any $m = ((y_i, z^i))_{i \in I}$, $\psi(m) = f(\theta_1, \ldots, \theta_I)$ if $y_i = \theta_i$ for all i, and $\psi(m) = \widetilde{a}$, otherwise.

- Modified regime R^ϵ satisfies the following:

  (i) $R^\epsilon(\emptyset) = b^*$

  (ii) For any $h = \big((g^1, m^1), \ldots, (g^{t-1}, m^{t-1})\big) \in H^t$ s.t. $t > 1$ and $g^{t-1} = b^*$:

  - A. if $m_i^{t-1} = (y_i, 0)$ for all i, then $R^\epsilon(h) = b^*$;

    B. if there exists some i s.t. $m_j^{t-1} = (y_j, 0)$ for all $j \neq i$ and $m_i^{t-1} = (y_i, z^i)$ with $z^i \neq 0$, then $R^\epsilon|h = S_\phi^i$, where S_ϕ^i induces a payoff $v_i(f) - \phi$ for some small $\phi > 0$ that depends on ϵ;

    C. if m^{t-1} is of any other type and i is the lowest-indexed agent among those who announce the highest integer, then i dictator forever

Characterisation: By a similar argument as before, we show that, for sufficiently small ϕ, any BNE is such that , at any date $t > 1$, each i's expected cont. payoff from period t is within ϵ of $v_i(f)$.

Step1: At any history at which the mechanism is b^*, expected continuation payoffs from next period is bounded below by $v_i(f) - \phi$

Step 2: For any t, if the mechanism at every equilibrium history at date t is b^*, the expected continuation payoffs from next period cannot differ from $v(f)$ by ϵ.

(This follows from the previous step, f being efficient and by making ϕ sufficiently small.)

Step 3: Mechanism b^* is played on the equilibrium path at any date.

Existence:

Consider first a history-independent regime R s.t. at every stage the players are required to play the mechanism b^* (no integer game).

- Since in such a regime, at each stage the strategy is a mapping $s_i : \Theta_i \longrightarrow Y_i$, the regime is simply a repeated Bayesian (adverse selection) game.

- A repeated Bayesian game can be represented as a repeated game with imperfect monitoring.

- FLM provide a series of Folk Theorems for repeated games with imperfect public monitoring.

- In particular, for an imperfect public monitoring repeated game obtained from a repeated Bayesian game with private values and independent signal, FLM show that if the players are sufficiently patient,

  - any efficient payoff that Pareto dominates a stage Nash equilibrium can be approximately sustained by a public perfect equilibrium (PPE) and *all the continuation payoffs* of the equilibrium can be made arbitrarily close to the efficient payoff.

- By construction the one shot-game of b^* has a NE that involves announcing N by each player and induces a payoff $v(\widetilde{a})$ that is Pareto dominated by $v(f)$. This means that we can appeal to FLM to show the following:

  - for any $\phi < v_i(f) - v_i(\widetilde{a})$, regime R has a PPE σ s.t. every continuation payoff is strictly within ϕ of $v(f)$, if the players are sufficiently patient.

- The above demonstrates the existence of an equilibrium σ for regime R s.t. every continuation payoff $w \in (v(f) - \phi, v(f))$. How about regime R^ϵ?

- For the regime R^ϵ, consider a strategy profile σ^ϵ that at any history at which b^* is the mechanism it plays according to σ and announces zero.

- σ is an equilibrium of R implies that σ^ϵ is an equilibrium of R^ϵ:

  - Since σ_i is a BR to σ_{-i} in R and since σ^ϵ_{-i} prescribes playing zero integer at every history, σ^ϵ_i must be a best response to σ^ϵ_{-i} amongst all strategies that choose zero integers at all histories in regime R^ϵ.

  - Since σ^ϵ_{-i} prescribes playing zero integer, choosing a positive integer by i induces a payoff of $v_i(f) - \phi$. Choosing σ^ϵ_i induces a payoff that exceeds $v_i(f) - \phi$. Therefore, σ^ϵ_i is better than any strategy that announces a positive integer at some history.

- Above strategies are credible (form a PPE) and not just Nash.

Private values with correlated signals and interdependent values

- The characterisation result for regime R^{ϵ} above also works with correlated signals and interdependent values.

- With correlated signals and private values, the existence result can also be extended except that FLM method needs an additional assumption: any efficient strategy profiles for the one shot game are "pairwise identifiability".

  - A strategy profile satisfies pairwise identifiability if, for every pair of players, the distributions over outcomes induced by one player's unilateral deviations are distinct from those induced by the other's deviations

- Since the assumption implies that the two players' deviations can be distinguished statistically, the appropriate incentives can be constructed to ensure that the regime has an equilibrium

Theorem 4: Suppose preferences have private values and that the SCF f is efficient, and satisfies conditions ω' and ν. Assume also that the set CoV^* is full-dimensional and the information structure in mechanism b^* is pairwise identifiable for any one-shot strategy profile that generates an efficient payoff in the set V^*. Then f is approximately payoff repeated-implementable in public strategies.

- If types are independent then pairwise identifiability is trivially satisfied.

- For a generic probability distribution over types, truth-telling is pairwise identifiability if $I \geq 3$ and no one player has "too many" more possible types than any other (for all $i \neq j$, $\mid \Theta_i \mid \leq \prod_{k \neq i,j} \mid \Theta_k \mid$).

- Pairwise identifiability:
  - Ignoring the choice of of N, a strategy in the one-shot game is a mapping $s_i : \Theta_i \longrightarrow \Theta_i$
  - Let the set of such strategies be $\{s_i^1, .., s_i^{K_i}\}$. Also, let the set of states be denoted by $\Theta = \{\theta^1, .., \theta^L\}$.
- For any s_{-i}, $\Gamma_i(s_{-i})$ refers to a matrix whose (k,l) elements corresponds to the probability of θ^l being reported when (s_i^k, s_{-i}) is chosen: $p(\theta^l \mid s_i = k, s_{-i})$.
- Thus it has ${K_i}^{K_i}$ rows and L columns.
- For any s_{-i},
  - $$\Gamma_{ij}(s) = \begin{pmatrix} \Gamma_i(s_{-i}) \\ \Gamma_j(s_{-j}) \end{pmatrix}.$$

    Thus it has ${K_i}^{K_i} + {K_j}^{K_j}$ rows and L columns.
- s is pairwise identifiability if for all i and j,
  - $$\text{rank of } \Gamma_{ij}(s) = \text{rank of } \Gamma_i(s_{-i}) + \text{ rank of } \Gamma_j(s_{-j}) - 1.$$

- In demonstrating their Folk Theorem results for repeated imperfect monitoring games, FLM appeal to the following enforceability property:
  - for any efficient one-shot strategies, there exist continuation payoffs that enforces the strategies (i.e. makes them mutual best responses given the continuation payoffs).

- When a repeated Bayesian game is represented as a repeated game with imperfect monitoring, as we have done above, the above enforceability condition holds if the Bayesian game has private values.

- Private values has no other role.

- But private values is not necessary for such enforceability.

- Cremer and McLean condition on conditional beliefs, a standard condition in mechanism design, ensures that enforceability of efficient outcomes is indeed the case.

- CM condition: For any i, θ_i and any non-negative coefficients $\{\mu(\theta_i')\}_{\theta_i' \neq \theta_i}$

  $p_{-i}(\theta_{-i} \mid \theta_i) \neq \sum_{\theta_i' \neq \theta_i} p_{-i}(\theta_{-i} \mid \theta_i')\mu(\theta_i')$

- CM rules out the possibility that player i of type θ_i could generate the same conditional probabilities on the types of the others through a random untruthful reporting.

- For a generic probability distribution over types, CM holds if $I \geq 3$ and no one player has "too many" more possible types than any other (for all $i \neq j$, $\mid \Theta_i \mid \leq \prod_{k \neq i,j} \mid \Theta_k \mid$)

- In the context of one-shot mechanism design with transfers, CM condition ensures that there exists transfers for which truthtelling is optimal (in our context transfers correspond to continuation payoffs).

- In our set-up CM condition allows us to show that any strategy profile s such that $s_i : \Theta_i \to \Theta_i$ is 1-to-1, for every i, is enforceable.

- Let $V^{**} = \{v : v = \sum p(\theta)u(f(s(\theta)), \theta)$ for some 1-to-1 s and $v \geq v(\widetilde{a})\}$

Theorem 5: Consider the case with interdependent values. Suppose that f is efficient and satisfies conditions η, ω' and ν. Assume CM, the set $Co(V^{**})$ is full-dimensional and the information structure in mechanism b^* is pairwise identifiable for any 1-to-1 strategy profile that generates an efficient payoff in the set $Co(V^{**})$. Then f is approximately payoff repeated-implementable in public strategies.

Some recent related works

- Complete information
  - Mezzetti and Renou (2013) try to unify the results from one shot implementation with those in LS by introducing the concept of dynamic monotonicity.
  - By appealing to a bounded rationality type reasoning, Lee and Sabourian (2013) demonstrate that LS results with complete information holds with with finite mechanisms (no integer games). The approach pursues the implications of agents having a preference for less complex strategies (at the margin) on the mechanism designer's ability to discourage undesired equilibrium outcomes.

- Incomplete information
  - Jackson and Sonnenschein (*Econometrica* 2008) budgeted mechanism setup can be interpreted as a finitely repeated implementation problem.
    - They show an approximate implementation (as the horizon becomes longer) if SCF is efficient, for private values with independent signals case.
    - They use budgeting only to derive their characterization result. For existence, they just appeal to standard arguments for existence of NE in mixed strategies
  - Matsushima, Miyazaki and Nobuyuki (*JET* 2010) uses the linking mechanism of JS to characterize the virtually implementable SCFs in adverse selection settings where a principal delegates multiple tasks to an agent.

- In a concurrent, independent study, Renou and Tomala (2013):
  * Consider independent private values infinite horizon with Markovian evolution of private information
  * Propose an alternative concept of approximate implementation: finding regimes such that the probability of infinite histories in which the desired SCF is almost repeatedly implemented approaches 1 as the players become perfectly patient
  * Our notion requires approximating the target payoffs at every equilibrium history. Hence, in RT, it is possible with a small probability that the players may still end up in equilibrium paths along which the continuation payoffs diverge from the desired payoffs after certain histories.
  * RT obtain a sufficiency result for SCFs that are efficient only among the space of payoffs given by undetectable deceptions
  * They use review strategy type construction and their results seem to depend on the assumption of independent private values. Our methods are different and holds for the case of interdependent values as well as correlated types.

Conclusion

- With some qualifications, we can achieve repeated (approximate) implementation of "efficient" (in the range) SCRs for incomplete information.

- Conjectures, questions and future research

  - While the results are for iid across time, we conjecture that the techniques used can be applied to obtain similar results when states do not evolve according to an iid process.

  - Regimes with imperfect monitoring.

  - Applications.