Sequential versus Simultaneous Disclosure

Keri (Peicong) Hu and Joel Sobel

December 29, 2019

MODEL

- I finite player set.
- Common action set X; X is a finite subset of the real line.
- $x = (x_1, \ldots, x_I)$, $x_i \in \mathbb{R}$, let $M(x) = \max\{x_1, \ldots, x_I\}$.
- Payoffs $\tilde{u}_i(x) = u_i(M(x))$, where $u_i(\cdot) : X \to \mathbb{R}$ are arbitrary.

Interpretation

- Unmodeled decision maker in organization.
- DM needs authorization from (at least) one player.
- Authorization of the form “you can take action x_i.” (Default is $\min X$.)
- DM prefers higher x.

Not Today

1. X not linearly ordered.
2. X not common to all agents.
3. DM's preferences not monotonic.
4. Authorization requires more than one agent.

Time permitting:
Alternative interpretations of model.

Basic Questions

1. How well does DM do?
2. How should DM consult agents?
3. What is the value of having additional agents?

Basic Answers

1. DM's outcome:
   - Simultaneous consultation has full disclosure equilibrium.
   - This equilibrium is silly.
   - Equilibria in full disclosure game are Pareto ranked.
   - Refinement picks agents' favorite; DM's least favorite.
2. DM with discretion (described later) does exactly as well with sequential mechanisms as with simultaneous.
3. Easy to quantify whether an agent agent helps.

Context

1. Literature on communication suggests adding just one agent can generate DM's favorite outcome.
2. Literature on communication suggests simultaneous superior to sequential.

These (stylized) assertions are "less true" in this model:

1. DM's favorite **is** an equilibrium in two-player game, but it is typically not an equilibrium if one removes weakly dominated strategies.
2. Sequential consultation cannot do better than simultaneous consultation (generically), but if DM has commitment power it will not do worse.

Pareto Efficiency

Definition

The **smallest strictly Pareto equilibrium outcome** is

$$\pi^* = \min\{\pi : u_i(\pi) > u_i(x_i) \text{ for all } x_i > \pi \text{ and all } i\}.$$

Definition

The **smallest weakly Pareto equilibrium outcome** is

$$\tilde{\pi}^* = \min\{\pi : u_i(\pi) \geq u_i(x_i) \text{ for all } x_i > \pi \text{ and all } i\}.$$

1. π^* and $\tilde{\pi}^*$ are well defined.
2. $\pi^* \geq \tilde{\pi}^*$.
3. Equality if $u_i(\cdot)$ is one-to-one for each player.
4. "Pareto" (efficient) from the point of view of agents (restrict to equilibria).

Simultaneous Disclosure

Agent i selects x_i. Payoffs $u_i(M(x))$

Simple Observations about NE

1. If $x = (x_1, \ldots, x_I)$ satisfies $x_i \leq \pi$ and at least two $x_j = \pi^*$ is a Nash Equilibrium for $\pi = \pi^*$ and $\tilde{\pi}^*$.
2. $\max X$ is always NE outcome.
3. Pure-strategy NE are Pareto ranked: If x^* and x^{**} are both Nash Equilibria and $M(x^*) \leq M(x^{**})$, then $\tilde{u}_i(x^*) \geq \tilde{u}_i(x^{**})$ for all i. This leads us to consider a more restrictive solution concept.

$\max X$ great for the (unmodeled) DM, but is the worst NE for agents.

Simple Observations about NE

1. If $x = (x_1, \ldots, x_I)$ satisfies $x_i \leq \pi$ and at least two $x_j = \pi^*$ is a Nash Equilibrium for $\pi = \pi^*$ and $\tilde{\pi}^*$.
2. $\max X$ is always NE outcome.
3. Pure-strategy NE are Pareto ranked: If x^* and x^{**} are both Nash Equilibria and $M(x^*) \leq M(x^{**})$, then $\tilde{u}_i(x^*) \geq \tilde{u}_i(x^{**})$ for all i. This leads us to consider a more restrictive solution concept.

$\max X$ great for the (unmodeled) DM, but is the worst NE for agents.

Look for refinements.

Warm Up: Single-Peaked Preferences

Assume each i has single-peaked preferences:

For each i there is m_i such that $u_i(\cdot)$ increases for $x_i < m_i$, decreases for $x_i > m_i$.

1. General properties of equilibria remain (Pareto ranked, full disclosure possible), but
2. "Obviously" it is dominant for Agent i to play m_i.
3. Properties:
   3.1 Easy characterization.
   3.2 (Refined) outcome is (typically) less than full disclosure.
   3.3 DM only needs one agent (maximum m_i).
   3.4 Informally, no loss or gain associated with sequential procedures.

Iterative Deletion of Weakly Dominated Strategies

Definition

Given subsets $X_i' \subset X_i$, with $X' = \Pi_{i \in I} X_i'$, Player i's strategy $x_i \in X_i'$ is weakly dominated relative to X' if there exists $z_i \in X_i'$ such that $\tilde{u}_i(x_i, x_{-i}) \leq \tilde{u}_i(z_i, x_{-i})$ for all $x_{-i} \in X'_{-i}$, with strict inequality for at least one $x_{-i} \in X'_{-i}$.

Definition

The set $S = S_1 \times \cdots \times S_I \subset X$ survives iterated deletion of weakly dominated strategies (IDWDS) if for $k = 0, 1, 2 \ldots$ there are sets $S^k = S_1^k \times \cdots \times S_I^k$, such that $S^0 = X$, $S^k \subset S^{k-1}$ for $k > 0$; S_i^k is obtained by (possibly) removing strategies in S_i^{k-1} that are weakly dominated relative to S^{k-1}; $S^k = S^{k-1}$ if and only if for each i no strategy in S_i^k is weakly dominated relative to S^{k-1}; and each S_i can be written in the form $\cap_{k=1}^{\infty} S_i^k$.

Comments

1. Process stops after finitely many steps (finite game).
2. Order generally matters (but not in generic cases and not much in general).

First Result

Proposition

If x is a strategy profile that survives IDWDS, then $M(x) \in [\tilde{\pi}^*, \pi^*]$. If x is a Nash equilibrium strategy profile that survives IDWDS, then $\tilde{u}_i(x) \geq u_i(\pi^*)$ for all i.

Corollary

If $\pi^* = \tilde{\pi}^*$, then for all x that survives IDWDS, $M(x) = \pi^*$.

Comments

1. Bounds on payoffs that survive refinement (typically strict reduction).
2. Corollary follows directly from Proposition.
3. Corollary says "generically" IDWDS selects Senders' favorite equilibrium.
4. Compared to single-peaked case:
   - 4.1 Generally need full power of iterated deletion.
   - 4.2 Still typically get less than $\max X$.
   - 4.3 Extra agents help if they increase π^*.

Idea of Proof

1. "Low disclosures stay."

   There always remain strategy profiles x such that $\max\{x_1, \ldots, x_I\} \leq \pi^*$.

   By definition, if other agents are below π^*, best response is below π^*.
2. "Very low disclosures leave."

   There exists no strategy profile $x \in S$ such that $M(x) < \tilde{\pi}^*$.

   Disclosing $\tilde{\pi}^*$ (or higher) eventually dominates for someone.
3. "High strategies leave."

   No strategy $z_i > \pi^*$ survives IDWDS.

Example: Second Result Needs Assumptions

1. Five information structures, 1, 2, 3, 4, 5; higher numbers better for DM.
2. Two Senders.
3. Sender 1 has strict preferences: $2 \succ 4 \succ 1 \succ 5 \succ 3$.
4. Sender 2 has strict preferences: $1 \succ 3 \succ 2 \succ 5 \succ 4$.
5. Unique outcome that survives IDWDS in the simultaneous game is full disclosure: $\pi^* = \tilde{\pi}^* = 5$.

More

Four possible sequences without returning to an agent: consulting exactly one agent, or consulting both in either order. Without commitment, the possible outcomes are:

Sequence	Disclosure
Agent 1	2
Agent 2	1
Agent 1, then 2	1
Agent 2, then 1	2

More claims

Returning to an agent will not lead to higher x. Without commitment, sequential process need not lead to π^*.

Commitment, no return:

- Start with E_1 and sometimes asks E_2, then more is possible with commitment.
- If the *DM* stops after 1, 2, or 3, E_1 will announce 2, which will be the final outcome. If the *DM* stops after 4 or 5, the outcome will be 4.
- If the *DM* consults E_1 first, then he would do best by committing not to consult E_2 if E_1 plays at least 4.
- If *DM* consults E_2 first, he gets 1 if the *DM* stops after 1; 2 if the DM stops after 2, 4 or 5, and 3 if the DM stops at 3. *DM* can obtains 3 by consulting first E_2, then E_1 (with commitment).

Conclude

- The best the *DM* can do with commitment but without returning to agents is 4.
- Hence commitment increases DM's value, but does not generate 5.
- But: there is a sequential disclosure protocol that generates $\pi^* = \tilde{\pi}^*$.
- The protocol involves asking E_1, then E_2, and then going back to E_1, with the commitment to stop if E_2 discloses 3.
- Why? If you start with E_2 and promise to stop after a disclosure of 3 (or more), then the disclosure will be either 3 or 5 (depending on the starting point). This leaves E_1 no choice but to disclose everything.

Assertion

If DM can:

1. Pick order of agents consulted;
2. Return to agents;
3. Commit to ending process (after an appropriate action by an agent)

DM can induce π^*. (Never more.)

Comparison

1. Simultaneous and Sequential are equivalent (same equilibrium value).
2. Simultaneous better: DM does not need to know preferences.
3. Simultaneous better: no need for commitment assumption.
4. Sequential better: only need one agent (with the threat of calling in others).

More Commitment, Less Information Needed

Sequential protocol to obtain π^*:

1. Ask Agent 1: "Will you permit me to take $\max X$?"

   If yes, stop. If no, continue.

2. General step (no permission granted): Keep track of last agent asked (i) and which x asked.

   2.1 If $i < I$, ask agent $i + 1$ "Will you permit me to take x?"

   If yes, stop. If no, continue.

   2.2 If $i = I$, ask agent 1: "Will you permit me to take $L(x)$?"

   If yes, stop. If no, continue.
   [$L(x)$ is the element of X just below x.]

Comparison

Stronger commitment power versus lower information requirement.

Alternative Interpretation I

Agents provide vital inputs:

1. x_i is input to production process (could be information or physical item).
2. DM processes $x = (x_1, \dots, x_I)$ and takes an action.
3. $M(x)$ is sufficient statistic for DM. (Makes more sense in the informational interpretation.)
4. DM's preferences are increasing in x.
5. Agents' preferences are arbitrary.

Alternative Interpretation II

Bayesian Persuasion

1. Agents actions are "experiments"
2. Here it is essential to study multi-dimensional case.