Implementation via Information Design in Binary-Action Supermodular Games

Stephen Morris, Daisuke Oyama, and Satoru Takahashi

Presented by Edward Plumb

14^{th} March 2025

Background (Partial Implementation)

- **Objective:** Fix a game. Analyse which joint distributions over action profiles and states can be implemented via information design (picking an information structure and an equilibrium).
- **Revelation Principle:** The mechanism designer offers players action recommendations that players have an incentive to follow.
- **Equilibrium:** Bayes Nash Equilibria- each player is best responding given their signal.

This Paper (Full Implementation)

- **Objective:** Fix a game. Analyse which joint distributions over action profiles and states can be implemented via picking an information structure but not the equilibrium (i.e. all equilibria induce the outcome).
- **Revelation Principle:** Cannot use the revelation principle.
- **Setting:** Binary-action supermodular (BAS) games.
- **Subgoal:** Smallest Equilibrium Implementation- What outcomes can be induced by the smallest equilibrium if you pick the information structure?

Two Player Two State Example

Payoffs

Good State (g)

	Invest	Not
Invest	$4, 5$	$1, 0$
Not	$0, 2$	$0, 0$

Bad State (b)

	Invest	Not
Invest	$-5, -4$	$-8, 0$
Not	$0, -7$	$0, 0$

- **Payoff of Investing:**
  - Base payoff in the bad state: -8
  - $+9$ bonus if the state is good
  - $+1$ bonus if you are player 2 (asymmetry)
  - $+3$ bonus if the other player invests
- Both players have a dominant strategy to invest in the good state and not invest in the bad state.
- There is a common prior over the state of the world: half-half.

Main Result

Under a dominant state property, an outcome can be smallest equilibrium implemented if and only if it satisfies not only obedience but also 'sequential obedience'.

- '**Sequential obedience**': Designer recommends players to switch to the high action according to a randomly chosen sequence. A player has a strict incentive to switch when told to do given that they think all players before them in the sequence have switched (without knowing the true state or realised sequence).
- Implementable outcomes characterised by a finite linear programme.
- Full implementation requires in addition a reverse sequential obedience condition, where designer recommends switches from high action to low action.

Binary-Action Supermodular (BAS) Games

- $I = \{1, ..., |I|\}$: finite set of players.
- Θ: finite set of states.
- $\mu \in \Delta(\Theta)$: a common prior. Without loss of generality, we assume $\mu(\theta) > 0$ for any $\theta \in \Theta$.
- $A_i = \{0, 1\}$: the binary action set for player i.
- $A = \{0, 1\}^I$.
- $u_i : A \times \Theta \to \mathbb{R}$: player i's payoff, supermodular s.t:

$$d_i(a_{-i}, \theta) = u_i(1, a_{-i}, \theta) - u_i(0, a_{-i}, \theta)$$

  is weakly increasing in a_{-i} (along comparable directions).
- Dominant state assumption: there exists $\bar{\theta} \in \Theta$ such that $d_i(\mathbf{0}_{-i}, \bar{\theta}) > 0$ for all i. *Even if all others are playing* 0, *each player wants to play* 1.

Information

- T_i: a countable set of signals for player i.
- $T = \prod_{i \in I} T_i$.
- $\pi \in \Delta(T \times \Theta)$: a common prior.
- Without loss of generality, we assume for any t_i, there exists $\theta \in \Theta$ and $t_{-i} \in T_{-i}$ such that $\pi((t_i, t_{-i}), \theta) > 0$.

Bayes Nash Equilibrium

Given information structure $\mathcal{T} = (T, \pi)$, a strategy for player i is a mapping $\sigma_i : T_i \to \Delta(A_i)$.

Definition: Bayes Nash Equilibrium

A strategy profile σ is a Bayes Nash equilibrium if for any player $i \in I$, any type $t_i \in T_i$, and any action $a_i \in A_i$, whenever $\sigma_i(t_i)(a_i) > 0$, we have:

$$\sum_{t_{-i} \in T_{-i}, \theta \in \Theta} \pi(t_{-i}, \theta | t_i) u_i((a_i, \sigma_{-i}(t_{-i})), \theta) \geq \sum_{t_{-i} \in T_{-i}, \theta \in \Theta} \pi(t_{-i}, \theta | t_i) u_i((a_i', \sigma_{-i}(t_{-i})), \theta)$$

for any action $a_i' \in A_i$.

Outcomes

- Information structure $\mathcal{T}$ and strategy profile σ induce an outcome $\nu \in \Delta(A \times \Theta)$:

$$\nu(a, \theta) = \sum_{t \in T} \pi(t, \theta) \prod_{i \in I} \sigma_i(t_i)(a_i).$$

Definition. Outcome ν is *partially implementable* if there exists an information structure $\mathcal{T}$ and an equilibrium σ such that $(\mathcal{T}, \sigma)$ induces ν.

Definition. Outcome ν satisfies *consistency* if

$$\nu(A \times \{\theta\}) = \mu(\theta).$$

Definition. Outcome ν satisfies *obedience* if

$$\sum_{a_{-i} \in A_{-i}, \theta \in \Theta} \nu(a_i, a_{-i}, \theta) \left(u_i(a_i, a_{-i}, \theta) - u_i(a_i', a_{-i}, \theta)\right) \geq 0$$

for any $i \in I$ and $a_i, a_i' \in A_i$.

Bayes Correlated Equilibria

Proposition. (Bergemann and Morris (2016) Outcome ν is partially implementable if and only if it satisfies consistency and obedience i.e. there exists an information structure $\mathcal{T}$ and an equilibrium σ such that $(\mathcal{T}, \sigma)$ induces ν if and only if

1. $\nu(A \times \{\theta\}) = \mu(\theta)$
2. for any $i \in I$ and $a_i, a_i' \in A_i$:

$$\sum_{a_{-i} \in A_{-i}, \theta \in \Theta} \nu(a_i, a_{-i}, \theta) \left(u_i(a_i, a_{-i}, \theta) - u_i(a_i', a_{-i}, \theta) \right) \geq 0$$

Example: Partial Implementation

Payoffs

Good State (g)

	Invest	Not
Invest	$4, 5$	$1, 0$
Not	$0, 2$	$0, 0$

Bad State (b)

	Invest	Not
Invest	$-5, -4$	$-8, 0$
Not	$0, -7$	$0, 0$

- The dominant state assumption is satisfied by the good state.
- The following outcome is partially implementable (and is the 'best' partially implementable outcome if you want to maximise the expected number of players investing).

Good State (g)

	Invest	Not
Invest	$\frac{1}{2}$	0
Not	0	0

Bad State (b)

	Invest	Not
Invest	$\frac{2}{5}$	0
Not	$\frac{1}{10}$	0

Smallest Equilibrium Implementation

- Because the game is BAS, there is a smallest (pure strategy) Bayes Nash equilibrium, $\underline{\sigma}$.
- Outcome ν is **smallest equilibrium implementable** (S-implementable) if there exists an information structure $\mathcal{T}$ such that $(\mathcal{T}, \underline{\sigma})$ induces ν.
- Call the set of such outcomes 'smallest equilibrium implementable outcomes' SI.
- Call the closure of S-implementable outcomes $\overline{SI}$.
- Can be thought of as having an information designer who favours the high action but anticipates adversarial equilibrium selection as a worst-case scenario.

Sequential Obedience

- For purely hypothetical exposition, suppose that players' default action was to play action 0 but the information designer recommends a subset of players to play action 1.
- The designer gives recommendations sequentially, according to some commonly known distribution on states and sequences of recommendations.
- When players are advised to play action 1, they will accept the recommendation only if it is a best response provided that all players who received the recommendation earlier than they did have switched.

Sequences

- Let Γ be the set of all finite sequences of distinct players; for example, if $I = \{1, 2, 3\}$, then

$$\Gamma = \{\emptyset, 1, 2, 3, 12, 13, 21, 23, 31, 32, 123, 132, 213, 231, 312, 321\}$$

- An 'ordered outcome' is a distribution over sequences and states $\nu_\Gamma \in \Delta(\Gamma \times \Theta)$.
- For $\gamma \in \Gamma$, $\bar{a}(\gamma)$ denotes the action profile where player i plays action 1 iff player i appears in γ, i.e.:

$$\bar{a} : \Gamma \to A; \gamma \mapsto (a_1, ..., a_{|I|}) \text{ where } a_i = \begin{cases} 1, & \text{if } i \in \gamma, \\ 0 & \text{if } i \notin \gamma \end{cases}$$

Sequences

- Each 'ordered outcome' $\nu_\Gamma \in \Delta(\Gamma \times \Theta)$ induces outcome $\nu \in \Delta(A \times \Theta)$ by forgetting the ordering, i.e.,

$$\nu(a,\theta) = \sum_{\gamma \in \Gamma : \bar{a}(\gamma) = a} \nu_\Gamma(\gamma, \theta).$$

- Thus, in the example, $\nu((1,1),\theta) = \nu_\Gamma(12,\theta) + \nu_\Gamma(21,\theta)$ and $\nu((0,1),\theta) = \nu_\Gamma(2,\theta)$.
- Let $\Gamma_i = \{\gamma \in \Gamma \mid \text{player } i \text{ appears in } \gamma\}$. So, $\Gamma_1 = \{1, 12, 21\}$.
- For $\gamma \in \Gamma_i$, $a_{-i}(\gamma)$ denotes the action profile of player i's opponents where player j plays action 1 iff player j appears in γ before player i.
- Hence, $a_{-1}(1) = (0)$, $a_{-1}(12) = (0)$ and $a_{-1}(21) = (1)$.
- This is the action profile that a player 'believes in' when they choose between their actions.

Sequential Obedience

Definition. Ordered outcome ν_Γ satisfies **sequential obedience** if

$$\sum_{\gamma \in \Gamma_i, \theta \in \Theta} \nu_\Gamma(\gamma, \theta) d_i(a_{-i}(\gamma), \theta) > 0$$

for all i such that Γ_i is non-empty. Ordered outcome ν_Γ satisfies **weak sequential obedience** if we replace inequalities with equalities.

Definition. Outcome $\nu \in \Delta(A \times \Theta)$ satisfies (weak) sequential obedience if there exists an ordered outcome $\nu_\Gamma \in \Delta(\Gamma \times \Theta)$ that induces ν and satisfies (weak) sequential obedience.
Warning: Do not get confused into thinking that there is actually a sequence- this is just a story/interpretation of the condition that comes from the proof.

Sequential Obedience in the Example

- Remember the example:

Good State

	Invest	Not
Invest	$4, 5$	$1, 0$
Not	$0, 2$	$0, 0$

Bad State

	Invest	Not
Invest	$-5, -4$	$-8, 0$
Not	$0, -7$	$0, 0$

- Invest is a best response of player 1 if they assigns probability $\frac{2}{3}$ to player 2 having invested and (independent) probability $\frac{2}{3}$ to the state being good.
- Invest is a best response of player 2 if they assigns probability $\frac{1}{3}$ to player 1 having invested and independent probability $\frac{2}{3}$ to the state being good.

Sequential Obedience in the Example

- Consider 'ordered outcome'

ν_Γ	$\{21\}$	$\{12\}$	$\emptyset$
Good	$\frac{1}{3}$	$\frac{1}{6}$	0
Bad	$\frac{1}{6}$	$\frac{1}{12}$	$\frac{1}{4}$

- This ordered outcome satisfies weak sequential obedience:
  - Conditional on being 'asked', Player 1 assigns probability $\frac{2}{3}$ to player 2 having been asked already and thus invested.
  - Furthermore, they also assign independent probability $\frac{2}{3}$ to the state being good.
  - Similarly, player 2 assigns probability $\frac{1}{3}$ to player 1 having invested and independent probability $\frac{2}{3}$ to the state being good.
- You can (and I did) perturb it to establish sequential obedience.

Result

Theorem 1.1. If an outcome is smallest equilibrium implementable, then it satisfies consistency, obedience and sequential obedience.
Note: It is immediate that it must satisfy consistency and obedience.

Definition. Outcome $\nu \in \Delta(A \times \Theta)$ satisfies grain of dominance if $\nu(1, \bar{\theta}) > 0$ (assigns positive probability to the dominant strategy state).

Theorem 1.2. If an outcome ν satisfies consistency, obedience, sequential obedience and grain of dominance, then ν is smallest equilibrium implementable.
Corollary 1. $\nu \in \overline{SI}$ if and only if it is satisfies consistency, obedience and weak sequential obedience.

Necessity of Sequential Obedience: Summary of Proof

- Suppose outcome ν is smallest equilibrium implementable via an information structure $\mathcal{T}$.
- Consider the best-response dynamics starting from the all-0 action profile.
- This will converge to the smallest equilibrium.
- The order of best-responses gives an ordered outcome which both induces ν and satisfies sequential obedience.

Necessity of Sequential Obedience

- Suppose that ν is smallest equilibrium implementable.
- Let $\mathcal{T} = (T, \pi)$ be an information structure whose smallest equilibrium induces ν.
- Start from constant 0 strategies and iteratively apply myopic best responses (say, round robin by player).
- This process converges to the smallest equilibrium.

Necessity of Sequential Obedience

- For each player $i \in I$, we define a function n_i on T_i such that, for each type $t_i \in T_i$, if type t_i changes from action 0 to action 1 in the n-th step, then $n_i(t_i) = n$; if they never change, then $n_i(t_i) = \infty$.
- Define for $\gamma = (i_1, ..., i_k)$, $T(\gamma)$ to be the set of type profiles t such that $(n_i(t_i))_{i \in I}$ is ordered according to γ, i.e. $n_{i_l}(t_{i_l}) < n_{i_m}(t_{i_m}) < \infty$ if and only if $l < m$ and $n_i(t_i) = \infty$ for all $i \notin \{i_1, ... i_k\}$.
- Suppose $n_1(g) = 1, n_1(b) = \infty, n_2(g) = 2, n_2(b) = 2$. Then:
  - $T(12) = \{(g, g), (g, b)\}$
  - $T(2) = \{(b, g), (b, b)\}$
  - $T(\emptyset) = T(1) = T(21) = \emptyset$
- Define the following order outcome:

$$\nu_\Gamma(\gamma, \theta) = \sum_{T(\gamma)} \pi(t, \theta).$$

Necessity of Sequential Obedience

- Because this process converges to the smallest equilibrium, we know that ν_Γ induces ν: for each $(a, \theta) \in A \times \Theta$

$$\begin{aligned}\sum_{\gamma:\bar{a}(\gamma)=a} \nu_\Gamma(\gamma, \theta) &= \sum_{\gamma:\bar{a}(\gamma)=a} \sum_{t \in T(\gamma)} \pi(t, \theta) \\ &= \sum_{t:n_i(t_i)<\infty \Leftrightarrow a_i=1} \pi(t, \theta) \\ &= \sum_{t:\underline{\sigma}(t)=a} \pi(t, \theta) = \nu(a, \theta)\end{aligned}$$

Necessity of Sequential Obedience

- To show sequential obedience, note that for each $t_i \in T_i$ with $n_i(t_i) < \infty$, we have

$$\sum_{t_{-i},\theta} \pi((t_i, t_{-i}), \theta) d_i(a_{-i}(t), \theta) > 0,$$

where $a_{-i}(t)$ is the action profile of player i's opponents in the myopic best response process when i switches; so player j of type t_j would be playing action 1 iff $n_j(t_j) < n_i(t_i)$.

- By adding up these inequalities over all such t_i, we have

$$\begin{aligned} 0 &< \sum_{t_i : n_i(t_i) < \infty} \sum_{t_{-i} \in T_{-i}, \theta \in \Theta} \pi((t_i, t_{-i}), \theta) d_i(a_{-i}(t), \theta) \\ &= \sum_{\gamma \in \Gamma_i} \sum_{t \in T(\gamma)} \sum_{\theta \in \Theta} \pi(t, \theta) d_i(a_{-i}(\gamma), \theta) \\ &= \sum_{\gamma \in \Gamma_i, \theta \in \Theta} \nu_\Gamma(\gamma, \theta) d_i(a_{-i}(\gamma), \theta) \end{aligned}$$

for any $i \in I$ such that $\Gamma_i \neq \emptyset$.

Sufficiency of Sequential Obedience: Summary of Proof

Inductive Step:

- Under sequential obedience, when a player is recommended to switch, they believe that all players before them have already switched to action 1.
- This belief ensures a strict incentive for the player to also switch, so if every player with an earlier switching time chooses 1, then any player scheduled later finds it optimal to switch as well.

Proof:

- Grain of Dominance 'kick-starts' the induction, and the inductive step then guarantees that every player with a finite switching time will eventually choose action 1.

Sufficiency of Sequential Obedience and Grain of Dominance

- Generalisation of email game in Rubinstein (1989)
- We construct information structure no.1 as follows
- Let ν_Γ be an ordered outcome establishing sequential obedience
- Draw γ according to ν_Γ and draw integer m from $\mathbb{Z}_+$ with almost 'uniform' probability $\eta(1-\eta)^m$, i.e. for small $\eta > 0$.
- Let the type of player i be given by

$$t_i = \begin{cases} m + \text{ranking of } i \text{ in } \gamma, & \text{if } \gamma \in \Gamma_i, \\ \infty, & \text{otherwise.} \end{cases}$$

Inductive Step

CLAIM: Suppose that for some $k \geq |I|$, we knew that all types $t_i \leq k$ choose action 1. Then types $k+1$ of all players must choose action 1.

- Consider type $k+1$ of player i. They will know that all players before him in the realised sequence γ are playing 1. As $\eta \to 0$, their belief over sequences will approximate the belief in the sequential obedience condition and his payoff to action 1 will approach

$$\sum_{\gamma \in \Gamma_i, \theta \in \Theta} \nu_\Gamma(\gamma, \theta) d_i(a_{-i}(\gamma), \theta) > 0$$

- So claim holds for sufficiently small η.

Perturb Payoffs

- Now construct information structure no.2 by re-arranging payoffs in information structure no.1 so that types $1, \ldots, |I|$ have a dominant strategy to play action 1.
- Possible because we assumed $\nu\left(1, \bar{\theta}\right) > 0$ and we can choose η sufficiently close to 0 so that $\nu\left(1, \bar{\theta}\right)$ is much larger than

$$\eta\left(1 + (1-\eta) + \cdots + (1-\eta)^{|I|-1}\right)$$

- Have types $1, \ldots, |I|$ assign probability 1 to state $\bar{\theta}$ and rearrange beliefs for higher types accordingly.
- Thus types $1, \ldots, |I|$ have a dominant strategy to play action 1.
- Now inductive argument implies all types $t_i < \infty$ choose action 1.