

Robust Information Design for Multi-Agent Systems with Complementarities: Smallest-Equilibrium Threshold Policies

Farzaneh Farhadi
Aston University
Birmingham, United Kingdom
f.farhadi@aston.ac.uk

Maria Chli
Aston University
Birmingham, United Kingdom
m.chli@aston.ac.uk

ABSTRACT

We study information design in multi-agent systems (MAS) with binary actions and strategic complementarities, where an external designer influences behavior only through signals. Agents play the *smallest-equilibrium* of the induced Bayesian game, reflecting conservative, coordination-averse behavior typical in distributed systems. We show that when utilities admit a convex potential and welfare is convex, the robustly implementable optimum has a remarkably simple form: *perfect coordination* at each state: either everyone acts or no one does. We provide a *constructive threshold rule*: compute a one-dimensional score for each state, sort states, and pick a single threshold (with a knife-edge lottery for at most one state). This rule is an explicit optimal vertex of a linear program (LP) characterized by *feasibility* and *sequential obedience* constraints. Empirically, in both vaccination and technology-adoption domains, our constructive policy matches LP optima, scales as $O(|\Theta| \log |\Theta|)$, and avoids the inflated welfare predicted by *obedience-only* designs that assume the designer can dictate the (best) equilibrium. The result is a general, scalable recipe for robust coordination in MAS with complementarities.

KEYWORDS

Information design; Robust coordination; Multi-agent systems; Potential games; Strategic complementarities; Linear programming

ACM Reference Format:

Farzaneh Farhadi and Maria Chli. 2026. Robust Information Design for Multi-Agent Systems with Complementarities: Smallest-Equilibrium Threshold Policies. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 9 pages. <https://doi.org/10.65109/HWRI8919>

1 INTRODUCTION

Binary cooperation with complementarities. Large-scale multi-agent systems (MAS) often involve binary cooperation decisions (such as whether to vaccinate [4], comply with regulations [8], or adopt a common standard [19]), where incentives are complementary: the benefit of cooperating increases with the number of other cooperating agents. In such settings, a designer (e.g., public-health authority, regulator, platform operator) typically observes the environment more accurately than agents but cannot mandate actions, and thus must *shape beliefs via information sharing*.



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/HWRI8919>

Limitations of classical information design. Classical information design (e.g., Bayesian persuasion, Bayes-correlated equilibrium) studies *partial implementation*: the designer specifies an information disclosure policy, makes recommendations accordingly (inducing a Bayesian game among agents), and assumes that when multiple equilibria exist, agents select an equilibrium consistent with those recommendations [5, 18]. This optimistic assumption effectively allows the designer to select the *best* equilibrium once incentive constraints hold. While analytically convenient, it can overstate what is achievable in decentralized or coordination-averse systems, where agents reason cautiously and coordination must emerge gradually rather than being taken for granted [7, 24].

Robust implementation via smallest-equilibrium play. We instead study implementation under smallest-equilibrium play, where agents coordinate on the worst equilibrium consistent with their information. This captures fragile coordination settings in which cooperation is chosen only if strictly profitable despite pessimistic beliefs.

From LP characterization to constructive design. We characterize smallest-equilibrium implementability via a linear program (LP) with feasibility and sequential-obedience constraints. Under convex potential and welfare, the LP admits an explicit threshold solution, reducing design to scoring and sorting. The optimal policy achieves perfect coordination per state and is computable in $O(|\Theta| \log |\Theta|)$ time.

Empirical validation in MAS domains. We validate the framework in two contrasting MAS domains. **Case A (Vaccination)** uses a small discrete type space to make the mechanics of the threshold rule transparent. **Case B (Technology adoption)** approximates a continuous state space, showing scalability and generality. In both settings, our policy matches the LP optimum and remains *robust under smallest-equilibrium play*, while classical partial-implementation designs predict systematically higher but non-credible welfare.

Contributions.

- **From partial to robust implementation.** We contrast classical *partial implementation*, where the designer can assume play of the best equilibrium, with the adversarial *smallest-equilibrium* setting, and we formalize implementability through an LP with *feasibility* and *sequential obedience* constraints.
- **Constructive, closed-form optimum.** Under convexity, the optimal smallest-equilibrium outcome is *perfect coordination*, achieved by a *threshold rule* based on a one-dimensional state score; this rule is an explicit optimal vertex of the LP.

- **Scalable to large MAS.** The constructive rule reduces design to scoring and sorting, avoiding large-scale optimization while matching LP solutions and scaling well.
- **Empirical evidence in MAS domains.** In vaccination and technology adoption, we quantify how partial-implementation baselines overstate welfare and show the robustness of our threshold policy.

2 RELATED WORK

Why not use existing methods? From classical to robust information design. Classical information design, including *Bayesian persuasion* and *Bayes-correlated equilibrium (BCE)*, assumes that once obedience holds, agents coordinate on the designer’s preferred equilibrium [5, 18]. Information disclosure has also been studied under this optimistic view in other settings, including security games Rabinovich et al. [25], Bayesian Stackelberg signaling Xu et al. [28], and persuasion with mechanism design Castiglioni et al. [6]. These approaches rely on favorable equilibrium selection.

In many coordination problems, such optimism is implausible: cooperation must build gradually, and agents may anticipate deviation. This motivates *robust information design*, which plans against *worst-case / smallest-equilibrium selection*. Recent studies formalize implementability under such pessimistic play [21, 24] and develop robust disclosure for coordination, regime change, and panic-resilient systems [3, 17]. We adopt this robust perspective for *super-modular MAS* via *sequential obedience*, where agents expect only predecessors to comply. This strengthens implementability beyond simultaneous approaches and builds on dynamic and staged persuasion [11, 12, 14, 20, 29], aligning with recent work on *constructive*, *incentive-compatible*, and *computationally tractable* mechanisms [10].

Outcome- vs. policy-level implementability under smallest play. Morris et al. [24] provide necessary and sufficient *outcome-level* conditions for distributions over actions and states achievable under smallest-equilibrium play. We translate these to a *policy-level* characterization: an LP whose decision variables are conditional probabilities over *sequences* of recommendations. Our LP comprises *feasibility* and *sequential obedience* constraints for cooperation and non-cooperation, which are the policy-wise counterparts of the outcome conditions. Beyond characterization, we extend [24] by introducing a *constructive threshold rule policy*, which under convexity solves the LP in closed form and yields the exact optimal sequential design.

MAS, network effects, and domain bridges. Strategic *complementarities* pervade many MAS domains, where an agent’s incentive to act increases with the number of others who act. In public health, vaccination and sanitation decisions create strong positive spillovers, making herd immunity a coordination challenge [4, 16]. In *public-good provision*, contributions exhibit increasing returns as participation grows, leading to multiple self-fulfilling equilibria and free-riding traps [2]. *Compliance* with taxation or regulation also depends on peers’ behavior, shaping incentives to comply [1, 8]. Finally, *technology adoption* and standardization are classic cases of network externalities, where adoption value rises with others’ uptake, creating tipping points and lock-in effects [13, 15, 19].

Our analysis focuses on the *fully connected* benchmark, where complementarities are global. Extending to *network games* with local complementarities is natural and discussed in the conclusion, as network topology can shift critical thresholds and alter the knife-edge mixing of robust policies. The constructive design we propose (i.e., *score*, *sort*, *threshold*) is simple, interpretable, and well-suited to practical deployment in such structured MAS environments.

3 PROBLEM SETTING: BINARY-ACTION COMPLEMENTARITIES IN MAS DOMAINS

We study environments where each agent in a population chooses one of two actions: cooperation ($a_i = 1$) or non-cooperation ($a_i = 0$). The central feature is *strategic complementarity*: an agent’s incentive to cooperate rises with the number of other cooperating agents. Such complementarities appear in many multi-agent systems: individuals deciding whether to vaccinate [4], taxpayers deciding whether to comply with rules or fares [1], communities deciding whether to contribute to a public project [2], or firms deciding whether to adopt a technology standard [13, 19]. In each case, the information designer is an authority (e.g., a regulator, public health agency, tax administration, or standards consortium) that can influence outcomes only through information disclosure.

Our formulation abstracts away from domain specifics to capture three common primitives: a binary individual choice, payoff complementarities, and an external designer with superior information. We introduce a *domain adapter* (Table 1) that maps these primitives to diverse MAS applications, illustrating the scope of our analysis.

3.1 Environment

Agents. There are N agents, indexed by $i = 1, \dots, N$. Each agent chooses either the costly cooperative action ($a_i = 1$) or the default non-cooperative action ($a_i = 0$). An *action profile* is $a = (a_1, \dots, a_N)$, and $n(a) = \sum_{i=1}^N a_i$ denotes the total number of cooperating agents.

State. The payoff environment depends on an underlying state $\theta \in \Theta$, drawn from a commonly known prior distribution μ over the finite set Θ . This state captures exogenous conditions (such as infection prevalence, audit intensity, regulatory pressure, or market demand) that determine both the intrinsic value of cooperation and the strength of complementarities. Agents know only the prior but cannot observe the realized state.

Designer. Unlike the agents, a *designer* observes state θ and can send signals to influence the agents’ choices. The designer’s objective (formalized in Section 3.5) differs from the agents’ individual payoffs and seeks to steer the system toward desirable collective outcomes. This information asymmetry underpins the designer’s role: by shaping the flow of information, the designer can coordinate agents’ behavior and improve system-level performance.

3.2 Agents’ utilities and incentives

Utility function. Agent i ’s utility in action profile a and state θ is

$$u_i(a, \theta) = a_i \left(b_\theta + \lambda_\theta \frac{n(a_{-i})}{N-1} \right) - c a_i + \kappa_i(a_{-i}, \theta), \quad (1)$$

where a_{-i} denotes the action profile of all agents except i , and $n(a_{-i}) = \sum_{j \neq i} a_j$ is the number of other agents who choose the cooperative action.

Table 1: Domain adapter: mapping model primitives to MAS applications.

Domain	States θ	Cost c	Complementarity λ_θ	Objective V
Vaccination	Disease prevalence; vaccine effectiveness	Health risks; monetary/-time costs	Herd immunity (coverage \uparrow)	Social welfare from reduced infection
Tax / Fare Compliance	Enforcement level (audit rate, penalties)	Payment of taxes or fares	Compliance spillovers (norms/deterrence)	Revenue net of enforcement
Public Good Provision	Project productivity/value	Individual contribution cost (effort, money)	Thresholds / increasing returns	Surplus from funded project
Technology Adoption	Market demand; regulatory push; ecosystem readiness	Migration, retraining, integration costs	Interoperability / network effects	Industry efficiency / consumer surplus

Components. This specification has four essential components. First, b_θ is the baseline return to cooperation in state θ (e.g., private vaccination benefits or intrinsic gains from technology adoption). Second, the term $\lambda_\theta \frac{n(a_{-i})}{N-1}$ with $\lambda_\theta \geq 0$ captures *complementarities*: the payoff from cooperation increases with others’ participation. Third, cooperation incurs a fixed cost $c > 0$ (e.g., medical or migration costs). Fourth, $\kappa_i(a_{-i}, \theta)$ captures agent heterogeneity, such as legacy revenues, subsidies, or exposure to local risks.

Assumption (Dominance). We assume the existence of a *dominant state* $\theta^{\text{dom}} \in \Theta$ such that

$$b_{\theta^{\text{dom}}} - c > 0.$$

In this state, cooperation $a_i = 1$ is strictly profitable for every agent regardless of others’ actions, ensuring that universal cooperation is feasible in some environments.

Strategic complementarity. From (1), each agent i ’s *marginal gain* from cooperation is

$$G_i(a_{-i}, \theta) \triangleq u_i((1, a_{-i}), \theta) - u_i((0, a_{-i}), \theta) = b_\theta - c + \lambda_\theta \frac{n(a_{-i})}{N-1}. \quad (2)$$

The marginal gain is increasing in $n(a_{-i})$, establishing *strategic complementarity*. Moreover, the symmetry of marginal gains across agents implies the existence of a potential function.

3.3 Potential representation and convexity

A key property of the game defined in Section 3.2 is that it admits a *potential function*. A potential function Φ is a single scalar that tracks unilateral incentives: the change in any agent’s utility from flipping its action equals the corresponding change in Φ . Best responses therefore *climb* the potential, and many equilibrium and comparative-statics results follow directly from its shape.

In our setting, the potential takes the form

$$\Phi(a, \theta) = F_\theta(n(a)), \quad F_\theta(n) = (b_\theta - c)n + \lambda_\theta \frac{n(n-1)}{2(N-1)}. \quad (3)$$

Indeed,

$$\Phi((1, a_{-i}), \theta) - \Phi((0, a_{-i}), \theta) = b_\theta - c + \lambda_\theta \frac{n(a_{-i})}{N-1}, \quad \forall i,$$

which coincides exactly with the marginal gain in (2). Thus the game is an *exact potential game* [23]: any unilateral deviation changes an agent’s utility by the same amount as it changes Φ .

LEMMA 3.1 (CONVEXITY OF THE POTENTIAL). *For each state $\theta \in \Theta$, the potential function $F_\theta(n)$ defined in (3) is (discretely) convex in n .*

PROOF. The discrete increments

$$\Delta F_\theta(k) = F_\theta(k+1) - F_\theta(k) = (b_\theta - c) + \lambda_\theta \frac{k}{N-1}$$

are weakly increasing in k because $\lambda_\theta \geq 0$. Hence F_θ is discretely convex. \square

3.4 Timing and information

The interaction unfolds in three stages. (i) The designer *commits* to an information disclosure policy π , which maps each possible state to a distribution of signals. These signals can be public or private and may be delivered either simultaneously or sequentially to the agents. (ii) Nature draws the state $\theta \sim \mu$. (iii) Signals are generated according to the committed policy; upon receiving them, agents update their beliefs and play a Bayesian Nash equilibrium (NE) of the resulting incomplete-information game.

Because the game exhibits *strategic complementarities*, best responses are monotone and the set of Bayesian NE can be ordered by the number of agents who choose the costly action. When multiple equilibria exist, we assume that agents coordinate on the one with the fewest cooperators, referred to as the *smallest-equilibrium*. Standard lattice results guarantee that this equilibrium exists and can be reached by iterated best responses starting from universal inaction [22, 26, 27]. This selection models coordination-averse systems: agents take the costly cooperative action only when it is strictly beneficial, which makes the resulting outcomes robust to coordination failure and to pessimistic interpretations of signals.

3.5 Designer’s problem

The designer evaluates collective performance through a welfare function $V : A \times \Theta \rightarrow \mathbb{R}$, normalized so that $V(\mathbf{0}, \theta) = 0$ for all θ , where $\mathbf{0}$ denotes universal non-cooperation. We assume that welfare is weakly increasing in the number of cooperating agents and satisfies the following convexity property.

Assumption (Convex welfare). For each $\theta \in \Theta$,

$$V(a, \theta) \leq \frac{n(a)}{N} V(\mathbf{1}, \theta),$$

where $\mathbf{1}$ denotes universal cooperation. This assumption captures the idea that the marginal social benefit of additional cooperating agents does not decrease as participation grows.

The designer’s goal is to choose an information disclosure policy π that maximizes expected welfare,

$$\max_{\pi} \mathbb{E}_{\theta \sim \mu} [V(a_{\pi}^*(\theta), \theta)], \quad (4)$$

where $a_{\pi}^*(\theta)$ is the smallest Bayesian NE induced by the policy π in state θ . The designer has no direct control over agents’ actions and can only shape their beliefs through the signals provided.

Solving the designer’s problem directly is challenging because the space of possible signaling schemes is vast: a policy may assign arbitrary public or private signals, possibly revealed in sequence, and induces complex belief updates and equilibria. To proceed, we next characterize the structure of information policies that are sufficient for optimality and derive the incentive constraints they must satisfy. This leads to a tractable *linear program*, which forms the foundation for our constructive solution in Section 5.

4 INFORMATION DESIGN AND IMPLEMENTABILITY

Classical information design makes an optimistic assumption: if several equilibria exist, agents coordinate on the one the designer prefers. With this assumption, the designer can restrict attention to *simultaneous recommendation policies*, where all agents receive private recommendations at once, and need only ensure that each recommendation is individually optimal to follow (obedience).

Our setting departs from this view by requiring robustness to the *smallest-equilibrium*: agents act conservatively and reason under the most cautious belief about how many others will cooperate, rather than optimistically believing that everyone will follow recommendations. In such environments, simultaneous recommendations are generally insufficient. The designer may need to disclose information *sequentially*, allowing cooperation to build gradually while remaining incentive compatible under these conservative beliefs. It has been shown in [24] that this conservative reasoning can be modeled by assuming that each agent expects only those recommended *before* them to comply with their recommendations. Ensuring that agents still follow their recommendations under this belief leads to a strengthened incentive condition, which we call *sequential obedience*, a stricter form of the classical obedience constraint.

In what follows, we formally define sequential information policies, introduce the sequential obedience condition, and show how these elements yield a linear programming formulation of the designer’s problem.

4.1 Sequential information policies

A *sequential policy* discloses recommendations in stages (e.g., staggered rollouts where only some agents are informed first). We model each staged recommendation process by an *ordered sequence* $\gamma = (i_1, \dots, i_m)$ of distinct agents, representing the order in which they are invited to cooperate (receive the recommendation $a_i = 1$). Agents not appearing in γ receive the default non-cooperation recommendation ($a_i = 0$). Let Γ denote the set of all such ordered subsequences of N agents, including the empty sequence \emptyset .

A *sequential information policy* is then a family of conditional distributions

$$\pi^{\text{seq}}(\gamma \mid \theta), \quad \gamma \in \Gamma, \theta \in \Theta,$$

where $\pi^{\text{seq}}(\gamma \mid \theta)$ is the probability that the designer, upon observing state θ , runs the staged recommendation process that invites agents exactly in the order specified by γ . For each θ , the policy must define a valid probability distribution:

$$\pi^{\text{seq}}(\gamma \mid \theta) \geq 0 \quad \text{and} \quad \sum_{\gamma \in \Gamma} \pi^{\text{seq}}(\gamma \mid \theta) = 1, \quad \forall \theta \in \Theta. \quad (\text{Feas})$$

These requirements are the *feasibility constraints*.

Example 1. Consider a setting with $N = 3$ agents $\{1, 2, 3\}$ and a state space $\Theta = \{L, H\}$. A possible sequential policy might be

$$\begin{aligned} \pi^{\text{seq}}((1, 3) \mid L) &= 0.6, & \pi^{\text{seq}}((2, 3) \mid L) &= 0.4, \\ \pi^{\text{seq}}((3, 1, 2) \mid H) &= 1, \end{aligned}$$

meaning that in state L the designer invites agents to cooperate in the order $(1, 3)$ with probability 0.6 and $(2, 3)$ with probability 0.4, while in state H the designer always sends cooperative invitations in the order $(3, 1, 2)$ (first agent 3, then agent 1, then agent 2). \square

4.2 Sequential obedience

Having introduced sequential policies, we now strengthen the classical incentive requirement to capture the conservative reasoning of smallest-equilibrium play. Under this belief system, an invited agent acts as if only those invited *before* it will comply. For an information policy to be feasible, every agent must find it optimal to follow its recommendation, whether to cooperate or not, even under this most cautious belief. We refer to this requirement as *sequential obedience*.

Let $a_{-i}^{\text{pred}}(\gamma)$ denote the action profile in which each $j \neq i$ cooperates if and only if j appears *before* i in the sequence γ . (If $i \notin \gamma$, then $a_{-i}^{\text{pred}}(\gamma)$ simply has all agents in γ cooperating and all others not.) Using the marginal gain function $G_i(a_{-i}, \theta)$ defined in (2), we express sequential obedience as two families of linear constraints.

Definition 4.1 (Sequential obedience – cooperation). A sequential policy π^{seq} satisfies *sequential obedience for cooperation* if, for every agent i ,

$$\sum_{\theta \in \Theta} \mu(\theta) \sum_{\gamma \in \Gamma_i} \pi^{\text{seq}}(\gamma \mid \theta) G_i(a_{-i}^{\text{pred}}(\gamma), \theta) x \geq 0, \quad (\text{SO-C})$$

where $\Gamma_i = \{\gamma \in \Gamma : i \in \gamma\}$ is the set of sequences in which i is invited to cooperate.

Constraint (SO-C) requires that even a fully cautious agent, believing only earlier invitees will cooperate, still finds it individually rational to accept a recommendation to cooperate. The inequality in (SO-C) is weak because sequential obedience allows indifference: strict gains identify the smallest equilibrium, while obedience requires only a weak best response.

Definition 4.2 (Sequential obedience – non-cooperation). A sequential policy π^{seq} satisfies *sequential obedience for non-cooperation* if, for every agent i ,

$$\sum_{\theta \in \Theta} \mu(\theta) \sum_{\gamma \in \Gamma \setminus \Gamma_i} \pi^{\text{seq}}(\gamma \mid \theta) G_i(a_{-i}^{\text{pred}}(\gamma), \theta) \leq 0, \quad (\text{SO-N})$$

where $\Gamma \setminus \Gamma_i = \{\gamma \in \Gamma : i \notin \gamma\}$ is the set of sequences in which i is not invited to cooperate.

Constraint (SO-N) states that an agent who is not invited must not have an incentive to deviate and cooperate.

Example 2 (Checking sequential obedience). Consider Example 1 and suppose the primitives are $\mu(L) = \mu(H) = 0.5$, $b_L = 1$, $b_H = 2.4$, $\lambda_L = 0.1$, $\lambda_H = 0.5$, $c = 2$. We illustrate the check for sequential obedience for cooperation (SO-C) using player 3.

Player 3. Sequences that invite player 3 are (1, 3) and (2, 3) in state L , and (3, 1, 2) in state H . The marginal gains of cooperation in these cases are

$$G_3(a_{-3}^{\text{pred}}((1, 3)), L) = G_3(a_{-3}^{\text{pred}}((2, 3)), L) = b_L - c + \lambda_L \frac{1}{2} = -0.95,$$

$$G_3(a_{-3}^{\text{pred}}((3, 1, 2)), H) = b_H - c + \lambda_H \frac{0}{2} = 0.4.$$

Sequential obedience for cooperation (SO-C) for player 3 is

$$\sum_{\theta \in \{L, H\}} \mu(\theta) \sum_{\gamma \in \Gamma_3} \pi^{\text{seq}}(\gamma | \theta) G_3(a_{-3}^{\text{pred}}(\gamma), \theta) \geq 0,$$

which evaluates to

$$0.5 \left[0.6 * (-0.95) + 0.4 * (-0.95) \right] + 0.5 \left[1.0 * 0.4 \right] = -0.275 < 0.$$

Thus (SO-C) fails for player 3, and the policy in Example 1 is not sequentially obedient. \square

Sequential obedience is strictly stronger than the standard (simultaneous) obedience used in Bayesian persuasion and BCE. Standard obedience requires that recommended actions be optimal assuming all other agents follow their recommendations. In contrast, sequential obedience (SO-C and SO-N) strengthens these constraints by certifying that each recommendation is a best reply under the conservative “only predecessors comply” belief. A simple one-state example illustrates the gap.

Example 3 (Sequential obedience is stronger than ordinary obedience). Consider $N = 3$ agents and a single state $\Theta = \{K\}$ with primitives $(b_K, \lambda_K, c) = (1, 1.5, 2)$. Suppose the designer commits to a policy that always invites all agents to cooperate in the fixed order $\gamma = (1, 2, 3)$.

Standard obedience. If every other agent cooperates, one agent’s deviation gain is $G_i(1, K) = b_K - c + \lambda_K = 0.5 > 0$, so “all cooperate” is a Bayesian NE and satisfies the usual obedience constraints.

Sequential obedience. Under smallest–equilibrium reasoning, an invited agent expects only earlier invitees to cooperate. The first invitee faces $G_1(0, K) = b_K - c = -1 < 0$ and would deviate, breaking the cooperation chain. Hence, the policy violates (SO-C) at the very first invitation. \square

This simple case shows that a recommendation profile implementable under standard/BCE obedience may fail under the stricter sequential obedience required for smallest–equilibrium play.

4.3 Characterizing implementability and the designer’s problem

Combining these conditions yields the following.

THEOREM 4.3 (CHARACTERIZATION OF IMPLEMENTABLE SEQUENTIAL POLICIES). *A sequential policy π^{seq} induces play consistent with smallest–equilibrium behavior if and only if it satisfies:*

- (1) **Feasibility (Feas),**
- (2) **Sequential obedience – cooperation (SO-C),**
- (3) **Sequential obedience – non-cooperation (SO-N).**

PROOF SKETCH. [24] characterize, at the *outcome level*, the distributions over action profiles and states that are implementable under smallest–equilibrium play. Our conditions (Feas), (SO-C), and (SO-N) provide the corresponding *policy-level* characterization, implying that any smallest–equilibrium outcome achievable by an arbitrary information disclosure scheme can also be implemented by a sequential policy satisfying these constraints. \square

These three families of constraints are linear in the policy distributions $\pi^{\text{seq}}(\gamma | \theta)$. The designer’s optimization problem can therefore be written as the following linear program:

$$\begin{aligned} \max_{\pi^{\text{seq}}} \quad & \sum_{\theta \in \Theta} \mu(\theta) \sum_{\gamma \in \Gamma} \pi^{\text{seq}}(\gamma | \theta) V(a(\gamma), \theta), \\ \text{s.t.} \quad & \text{(Feas), (SO-C), (SO-N),} \end{aligned} \tag{5}$$

where $a(\gamma)$ is the action profile in which each player i plays 1 if and only if i is listed in γ . The objective function in (5) is exactly the designer’s expected welfare objective introduced in (4), expressed here in terms of the sequential policy variables.

Every optimal solution of (5) corresponds to a sequential information policy that is robust under smallest–equilibrium play. Although problem (5) is a *linear program* and can, in principle, be solved using standard LP solvers, its size grows extremely quickly. The number of decision variables (i.e., the probabilities $\pi^{\text{seq}}(\gamma | \theta)$ for all $\gamma \in \Gamma$ and $\theta \in \Theta$) is $|\Theta| \cdot |\Gamma|$, where

$$|\Gamma| = \sum_{k=0}^N \binom{N}{k} k! = O(N!),$$

denotes the number of ordered subsequences of the N agents. Since $N!$ grows super-exponentially with N , the problem becomes computationally infeasible even for moderate N . To overcome this combinatorial blow-up, the next section introduces a *constructive algorithm* that computes an *exact optimal sequential policy* for (5) in time $O(|\Theta| \log |\Theta|)$, independent of the number of agents N .

5 PERFECT COORDINATION UNDER CONVEXITY: A CONSTRUCTIVE CHARACTERIZATION

Section 4 characterized the feasible set of *sequential information policies* via three families of linear constraints (feasibility, sequential obedience for cooperation, and sequential obedience for non-cooperation) and showed that the designer’s problem reduces to the linear program (5). A natural question for MAS is whether the optimal robust policy must be combinatorially complex (e.g., mixing over many partial invitation sequences) or whether a simple, statewise structure suffices.

The answer is strikingly simple: there always exists an optimal smallest–equilibrium implementable sequential policy that achieves perfect coordination in every state: either all agents cooperate or none do. Convexity of the welfare function removes any advantage of partial adoption, while sequential obedience then enforces a

sharp, statewise “all-or-none” rule that remains robust under the most cautious agent beliefs.

5.1 A constructive threshold rule

Our approach is constructive: we explicitly *build* an optimal perfectly coordinated policy π^{seq^*} and then verify its optimality.

Step 1: Score each state. For each state θ , define

$$S(\theta) = \frac{F_\theta(N)}{V(\mathbf{1}, \theta)},$$

the ratio of the potential to welfare under full cooperation.

Step 2: Order states and identify a threshold. Order states so that $S(\theta)$ is weakly increasing.¹ We write $\theta' \succ \theta$ if θ' comes after θ in that ordering.

Let θ^* be the (unique) *threshold* state and $p^* \in [0, 1]$ the mixing weight that solve

$$\sum_{\theta: \theta \succ \theta^*} \mu(\theta)F_\theta(N) + p^* \mu(\theta^*)F_{\theta^*}(N) = 0, \quad (6)$$

i.e., the weighted sum of full-cooperation potentials across the “cooperate” states balances exactly at zero.

Step 3: Construct the perfectly coordinated policy. Define the sequential policy π^{seq^*} by assigning probability mass only to the two extreme invitation sets: the empty sequence (no one invited) and the full sequence (all invited). For any ordered sequence $\gamma \in \Gamma$, let $|\gamma|$ denote its length. Then

$$\pi^{\text{seq}^*}(\gamma \mid \theta) = \begin{cases} 1/N!, & |\gamma| = N, \theta \succ \theta^*, \\ p^*/N!, & |\gamma| = N, \theta = \theta^*, \\ 1 - p^*, & |\gamma| = 0, \theta = \theta^*, \\ 1, & |\gamma| = 0, \theta \prec \theta^*, \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

States above the threshold induce *universal cooperation* (any full ordering of agents chosen uniformly); states below induce *universal defection*; and the threshold state mixes between them in just the right proportion. This construction balances the incentives implied by sequential obedience: cooperation is sustained only where the environment is strong enough, while weaker states lead to coordinated defection. The optimal policy π^{seq^*} sets the (SO-C) constraints exactly at equality via the choice of p^* (6). An algorithmic description appears in Algorithm 1.

5.2 Feasibility of π^{seq^*}

LEMMA 5.1 (FEASIBILITY). *The sequential policy π^{seq^*} in (7) satisfies feasibility (Feas), sequential obedience for cooperation (SO-C), and sequential obedience for non-cooperation (SO-N).*

PROOF SKETCH. Feasibility follows from constructing valid distributions in each state. The choice of p^* in (6), together with the symmetric construction in (7), ensures the (SO-C) constraints hold with equality for all agents. Moreover, convexity of the potential implies that in any state where full cooperation is not recommended, no single agent can profitably deviate; hence (SO-N) is satisfied. \square

¹Any tie-breaking among states with the same $S(\theta)$ does not affect the construction.

Algorithm 1 Constructing the perfectly coordinated optimal sequential policy π^{seq^*}

- 1: Compute $S(\theta) = F_\theta(N)/V(\mathbf{1}, \theta)$ for all θ .
- 2: Sort states by nondecreasing $S(\theta)$.
- 3: Find the smallest θ^* such that

$$\sum_{\theta: \theta \succ \theta^*} \mu(\theta)F_\theta(N) \leq 0 < \sum_{\theta: \theta \geq \theta^*} \mu(\theta)F_\theta(N).$$

- 4: Choose p^* so that

$$\sum_{\theta: \theta \succ \theta^*} \mu(\theta)F_\theta(N) + p^* \mu(\theta^*)F_{\theta^*}(N) = 0.$$

- 5: Build π^{seq^*} according to (7).
-

5.3 Optimality

THEOREM 5.2 (PERFECT COORDINATION IS OPTIMAL). *In any binary-action supermodular game with a convex potential and a designer whose welfare is convex, the sequential policy π^{seq^*} defined in (7) is an optimal solution of the linear program (5).*

PROOF SKETCH. The LP (5) is a convex optimization over a polytope. Convexity of both the potential and welfare ensures that extreme points of this feasible set are all-or-none policies. Balancing sequential obedience identifies the optimal extreme point, which coincides with π^{seq^*} . \square

Generality and Computational Efficiency. Theorem 5.2 relies on our standing assumptions: an exact-potential game and convexity of the potential and welfare. In contrast, the LP (5) applies to any binary-action supermodular game without these assumptions but becomes computationally intractable as the number of agents grows. Our constructive method instead yields a closed-form optimum under convexity and is computationally efficient: computing state scores requires $O(|\Theta|)$ time, sorting them takes $O(|\Theta| \log |\Theta|)$, and finding the optimal threshold and mixing weight is $O(|\Theta|)$. Overall, the entire policy can be computed in $O(|\Theta| \log |\Theta|)$ time, independent of the population size N .

5.4 MAS takeaways and implementation guidance

- **One-dimensional score.** States are ranked by $S(\theta) = F_\theta(N)/V(\mathbf{1}, \theta)$; sorting and balancing sequential obedience yields the exact optimum.
- **All-or-none robustness.** Under convexity, optimal robust policies never use partial adoption within a state: either *everyone cooperates* or *no one does*.
- **Closed-form scalability.** The threshold rule is an explicit LP vertex computable in $O(|\Theta| \log |\Theta|)$ time.

Figure 1 visually summarizes the transition from classical persuasion to robust implementability and the resulting threshold structure.

6 NUMERICAL RESULTS

We validate the constructive characterization of Section 5 through two multi-agent domains. The first revisits the running example

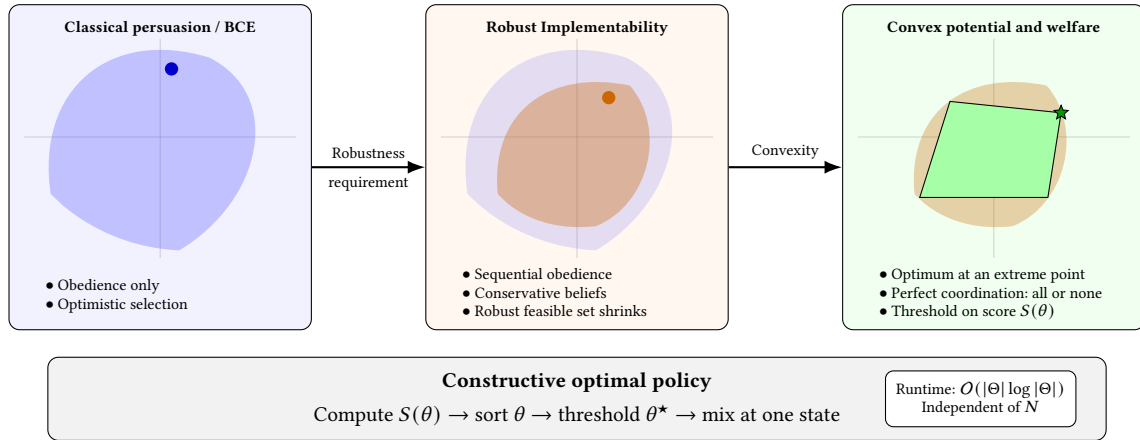


Figure 1: Robust information design under smallest-equilibrium play. Sequential obedience restricts implementability relative to classical persuasion. Under convex potential and welfare, the optimum is an extreme point, yielding a threshold policy computable in $O(|\Theta| \log |\Theta|)$ time.

to clarify the threshold mechanics and highlight the contrast with partial implementation. The second considers a nearly continuous state space, showing that the approach remains tractable and interpretable beyond small illustrative settings. All results are reproducible; code is available [9].

6.1 Case study 1: Vaccination-style environment

Setup. We revisit the running example as a vaccination scenario with $N = 3$ agents choosing $a_i \in \{0, 1\}$. States are $\Theta = \{L, H\}$ with prior $\mu(L) = \mu(H) = 0.5$, parameters $(b_L, \lambda_L) = (1, 0.1)$, $(b_H, \lambda_H) = (2.4, 0.5)$, and cost $c = 2$. Designer welfare is convex in coverage:

$$V(a, \theta) = \alpha_\theta \left(\frac{n(a)}{N} \right)^\beta, \quad \beta = 1.5,$$

with $\alpha_L = 6$ and $\alpha_H = 12$.

Constructive policy. Full-vaccination potentials are $F_L(N) = -2.85$ and $F_H(N) = 1.95$, giving scores $S(L) = -0.475$ and $S(H) = 0.1625$, hence $H \succ L$. Solving the sequential-obedience balance

$$\mu(H)F_H(N) + p^* \mu(L)F_L(N) = 0$$

yields $p^* = 0.684$ with threshold $\theta^* = L$.

The optimal policy π^{seq^*} is:

- in H , recommend vaccination to all agents;
- in L , recommend vaccination to all agents with probability 0.684, and to none otherwise.

When recommending vaccination, probability is split equally across all six orderings. This achieves expected welfare 8.05.

Why private sequential signaling matters. If the same randomization (0.684 in L , 1 in H) were implemented via a public signal, observing “vaccinate” yields

$$\Pr(H \mid \text{vac}) = \frac{0.5}{0.5 + 0.5 * 0.684} \approx 0.595.$$

They would then play a simultaneous-move game under this common posterior, where the expected marginal gain from vaccinating

is

$$G_i(a_{-i}, \text{posterior}) = 0.595 G_i(a_{-i}, H) + 0.405 G_i(a_{-i}, L).$$

In particular, if no one vaccinates, each agent’s deviation payoff is

$$G_i(0, \text{posterior}) = 0.595 * 0.4 + 0.405 * (-1) = -0.167 < 0,$$

so universal non-vaccination is a Bayesian NE and the *smallest-equilibrium* of the simultaneous game. Under public signaling, the intended vaccination policy collapses to no adoption. The robust optimum, therefore, relies on private sequential invitations.

Contrast with partial-implementation (optimistic) design.

A designer checking only classical obedience would recommend vaccination in both states, predicting welfare 9. Since the signal is uninformative, agents retain the prior 0.5, yielding

$$G_i(a_{-i}, 0.5) = 0.5 G_i(a_{-i}, H) + 0.5 G_i(a_{-i}, L).$$

Then, we can verify that both *no vaccination* and *full vaccination* are Bayesian NE:

- **No vaccination is a Bayesian NE:**

$$G_i(0, 0.5) = 0.5 * (0.4) + 0.5 * (-1) = -0.3 < 0,$$

so no agent wishes to vaccinate when everyone declines.

- **Full vaccination is a Bayesian NE:**

$$G_i(1, 0.5) = 0.5 * (0.4 + 0.5) + 0.5 * (-1 + 0.1) = 0 \geq 0,$$

so no agent has an incentive to deviate from vaccination when all others vaccinate.

Thus, under partial implementation, this policy is *optimal* because full vaccination is a Bayesian NE in every state. Yet under smallest-equilibrium play, agents coordinate on no vaccination, yielding welfare 0. This illustrates why sequential obedience is essential: ignoring it can produce policies that appear optimal but collapse under conservative reasoning.

Take-home message. The threshold rule (vaccinate in H , mix in L with $p^* = 0.684$) is simple, exact, and robust. Sequential obedience is essential: public or optimistic designs fail once agents reason conservatively.

6.2 Case study 2: Technology adoption

Setup. The vaccination example illustrated how the robust policy works in a simple two-state setting and why smallest-equilibrium reasoning matters. We now consider a richer environment that approximates a continuous type space: a technology adoption problem with $N = 10$ agents and nearly continuous uncertainty. States are $\theta \in \{0.01, 0.02, \dots, 1.00\}$, representing *ecosystem readiness*. As readiness rises, both the private benefit b_θ (from 0.5 to 2) and the complementarity λ_θ (from 0.1 to 0.8) increase linearly. This fine discretization approximates a continuous domain within our finite-state framework. Designer welfare uses the same convex coverage function as in Case Study 1, with α_θ increasing from 6 to 12.

Baselines and evaluation. We compare three policies:

- (1) **Robust (ours).** The policy π^{seq^*} , which satisfies (SO-C) and (SO-N) and is robust under smallest-equilibrium play.
- (2) **BCE-Optimistic.** Solves the designer’s problem under classical obedience (Bayesian persuasion/BCE), assuming coordination on the designer-preferred equilibrium.
- (3) **BCE-Realized.** Takes the BCE policy from (2) but *evaluates it* under smallest-equilibrium play.

Results. With $c = 2$, Fig. 2 shows that the score $S(\theta) = F_\theta(N)/V(1, \theta)$ orders the states smoothly, yielding an almost continuous threshold. The robust policy prescribes full adoption at high readiness, none at low readiness, and mixes at $\theta^* = 0.56$ with $p^* = 0.24$. In contrast, BCE-Optimistic recommends adoption for all $\theta > 0.28$ but collapses to no adoption under smallest-equilibrium play.

Fig. 3 compares welfare under the three policies: **Robust** (blue), **BCE-Optimistic** (orange), and **BCE-Realized** (red). The orange curve shows welfare *predicted by classical BCE analysis*, assuming agents coordinate on the designer-preferred equilibrium. The red curve shows the *realized welfare* when agents instead play the smallest-equilibrium: whenever “no adoption” is also a BNE ($1.5 \leq c \leq 2.95$), agents select that equilibrium, and welfare drops to zero. Hence, BCE-Optimistic overstates achievable welfare, while BCE-Realized reveals what the same policy delivers with cautious agents. The robust policy (blue) achieves the true coordination frontier under smallest-equilibrium play.

At low costs ($c \leq 1.5$), all policies coincide (universal adoption); at high costs ($c \geq 2.95$), all converge to zero (universal non-adoption). Overall, the results highlight the practical value of smallest-equilibrium robustness: it identifies credible, implementable coordination outcomes where classical persuasion fails.

Discussion. This case shows the constructive threshold rule extends beyond small finite settings. Even though our analysis assumes finitely many states, the technology-adoption example demonstrates that with fine discretization the method handles nearly continuous domains, yielding stable, interpretable policies and welfare. Together, the two case studies illustrate both the rule’s clarity in simple environments and its robustness in richer ones.




Figure 2: Technology adoption with nearly continuous states. Scores $S(\theta)$ induce a near-continuous threshold. The robust policy π^{seq^*} adopts only where coordination is sustainable, while BCE policy (orange) recommends broader adoption but collapses to no adoption under smallest-equilibrium play.




Figure 3: Welfare under Robust, BCE-Optimistic (classical obedience only), and BCE-Realized (same BCE policy evaluated under smallest-equilibrium). For medium costs, the BCE-Optimistic curve sits above Robust, but BCE-Realized drops to zero, revealing that the optimistic gains are not achievable under conservative play.

7 CONCLUSION AND FUTURE WORK

We develop a constructive threshold framework for robust information design under smallest-equilibrium play. Unlike classical obedience-based approaches, our method ensures robustness to coordination failure. The resulting policies are simple, interpretable, and scalable to nearly continuous domains while capturing the true limits of sustainable coordination.

Future work. One direction is to study agents embedded in social networks, where complementarities are local rather than global, and to analyze how network topology affects optimal disclosure. Another extension is to apply the framework to real vaccination data, enabling empirical validation of smallest-equilibrium effects and threshold-based policies. Finally, one could shift the designer’s objective from maximizing compliance to promoting diversity (e.g., avoiding monopolies); in this case the perfect-coordination reduction no longer applies, but the potential-based scoring approach may be adapted to reward heterogeneous outcomes.

REFERENCES

- [1] Michael G. Allingham and Agnar Sandmo. 1972. Income tax evasion: A theoretical analysis. *Journal of Public Economics* 1, 3-4 (1972), 323–338.
- [2] James Andreoni. 1988. Privately provided public goods in a large economy: the limits of altruism. *Journal of Public Economics* 35, 1 (1988), 57–73.
- [3] Deepal Basak and Zhen Zhou. 2025. Panics and Early Warnings. *Journal of Political Economy* 133, 7 (2025), 2089–2138.
- [4] Chris T. Bauch and David JD Earn. 2004. Vaccination and the theory of games. *Proceedings of the National Academy of Sciences* 101, 36 (2004), 13391–13394.
- [5] Dirk Bergemann and Stephen Morris. 2016. Bayes correlated equilibrium and the comparison of information structures in games. *Theoretical Economics* 11, 2 (2016), 487–522.
- [6] Matteo Castiglioni, Alberto Marchesi, and Nicola Gatti. 2022. Bayesian persuasion meets mechanism design: Going beyond intractability with type reporting. In *Proceedings of the 21st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 226–234.
- [7] Maria Chli and Michael Winsper. 2015. Using the Max-Sum Algorithm for Supply Chain Emergence in Dynamic Multiunit Environments. *IEEE Transactions on Systems, Man, and Cybernetics: Systems* 45, 3 (2015), 422–435. <https://doi.org/10.1109/TSMC.2014.2351782>
- [8] Roberto Dell'Anno. 2009. Tax evasion, tax morale and policy maker's effectiveness. *The Journal of Socio-Economics* 38, 6 (2009), 988–997.
- [9] Farzaneh Farhadi and Maria Chli. 2026. Reproducibility Code for “Robust Information Design for Multi-Agent Systems with Complementarities: Smallest-Equilibrium Threshold Policies” (AAMAS 2026). <https://doi.org/10.5281/zenodo.18613403>
- [10] Farzaneh Farhadi, Maria Chli, and Nicholas R Jennings. 2023. Optimal and efficient auctions for the gradual procurement of strategic service provider agents. *Journal of Artificial Intelligence Research* 76 (2023), 959–1018.
- [11] Farzaneh Farhadi and Demosthenis Teneketzis. 2022. Dynamic information design: A simple problem on optimal sequential information disclosure. *Dynamic Games and Applications* 12, 2 (2022), 443–484.
- [12] Farzaneh Farhadi, Demosthenis Teneketzis, and S Jamaloddin Golestani. 2018. Static and dynamic informational incentive mechanisms for security enhancement. In *2018 European control conference (ECC)*. IEEE, 1048–1055.
- [13] Joseph Farrell and Garth Saloner. 1985. Standardization, compatibility, and innovation. *RAND Journal of Economics* 16, 1 (1985), 70–83.
- [14] Jiarui Gan, Rupak Majumdar, Goran Radanovic, and Adish Singla. 2022. Bayesian persuasion in sequential decision-making. In *Proceedings of the AAI Conference on Artificial Intelligence*, Vol. 36. 5025–5033.
- [15] Michael Garlick and Maria Chli. 2010. Agent-based simulation of lock-in dynamics in a duopoly. In *9th International Joint Conference on Autonomous Agents and Multiagent Systems*. 1545–1546.
- [16] Raymond Guiteras, James Levinsohn, and Ahmed Mushfiq Mobarak. 2019. Encouraging sanitation investment in the developing world: A cluster-randomized trial. *Science* 365, 6454 (2019), 67–98.
- [17] Nicolas A. Inostroza and Alessandro Pavan. 2025. Adversarial Coordination and Public Information Design. *Theoretical Economics* 20, 2 (2025), 763–813.
- [18] Emir Kamenica and Matthew Gentzkow. 2011. Bayesian persuasion. *American Economic Review* 101, 6 (2011), 2590–2615.
- [19] Michael L. Katz and Carl Shapiro. 1985. Network externalities, competition, and compatibility. *The American Economic Review* 75, 3 (1985), 424–440.
- [20] Fei Li and Peter Norman. 2021. Sequential persuasion. *Theoretical Economics* 16, 2 (2021), 639–675.
- [21] Laurent Mathevet, Jacopo Perego, and Ina Taneva. 2020. On Information Design in Games. *Journal of Political Economy* 128, 4 (2020), 1370–1404.
- [22] Paul Milgrom and John Roberts. 1990. Rationalizability, Learning, and Equilibrium in Games with Strategic Complementarities. *Econometrica* 58, 6 (1990), 1255–1277.
- [23] Dov Monderer and Lloyd S. Shapley. 1996. Potential Games. *Games and Economic Behavior* 14, 1 (1996), 124–143.
- [24] Stephen Morris, Daisuke Oyama, and Shota Takahashi. 2024. Implementation via Information Design in Binary-Action Supermodular Games. *Econometrica* 92, 3 (2024), 775–813.
- [25] Zinovi Rabinovich, Albert Xin Jiang, Manish Jain, and Haifeng Xu. 2015. Information disclosure as a means to security. In *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 645–653.
- [26] Donald M Topkis. 1979. *Equilibrium points in nonzero-sum n-person submodular games*. *SIAM Journal on Control and Optimization*, 17(6): 773–787.
- [27] Xavier Vives. 2005. Complementarities and games: New developments. *Journal of Economic Literature* 43, 2 (2005), 437–479.
- [28] Haifeng Xu, Rupert Freeman, Vincent Conitzer, Shaddin Dughmi, and Milind Tambe. 2016. Signaling in Bayesian Stackelberg games. In *Proceedings of the 15th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 150–158.
- [29] Yunfei Yao. 2024. Dynamic persuasion and strategic search. *Management Science* 70, 10 (2024), 6778–6803.