Hierarchical Bayesian Persuasion: Importance of Vice Presidents

Majid Mahzoon*

Tepper School of Business, Carnegie Mellon University

February 9, 2022

We study strategic information transmission in a hierarchical setting where information gets transmitted through a chain of agents up to a decision maker whose action is of importance to every agent. This situation could arise whenever an agent can communicate to the decision maker only through a chain of intermediaries, for example, an entry-level worker and the CEO in a firm, or an official in the bottom of the chain of command and the president in a government. Each agent can decide to conceal part or all the information she receives. Proving we can focus on simple equilibria, where the only player who conceals information is the first one, we provide a tractable recursive characterization of the equilibrium outcome, and show that it could be inefficient. Interestingly, in the binary-action case, regardless of the number of intermediaries, there are a few pivotal ones who determine the amount of information communicated to the decision maker. In this case, our results underscore the importance of choosing a pivotal vice president for maximizing the payoff of the CEO or president.

*mmahzoon@andrew.cmu.edu

1 Introduction

Consider a hierarchical organization such as a firm or the government where the authority who takes action is the CEO or the president. Most of the times, an entry-level worker or an official in the bottom of the chain of command who has a request or suggestion, or a scandal to report, cannot communicate with the CEO/president directly. She can only communicate with her supervisor, who if desires, communicates with her own supervisor, and so on through the organization hierarchy. In other words, there is a chain of intermediaries between the entry-level worker and the CEO in the firm or between the low-ranked official and the president in the government. Knowing the communication strategy of their subordinates and the preferences of their superiors, each of these intermediaries will decide on whether to pass on her information to their supervisor, and if so, to what extent. How much information will be communicated to the CEO/president? How important is the hierarchy configuration, that is, the location and preferences of each intermediary? How essential is the role of vice presidents in hierarchical organizations?

In this paper, we are investigating the outcome of intermediated communication from a low-ranked agent to the CEO/president. It turns out regardless of the number of intermediaries, there are a few pivotal ones whose preferences and locations in the hierarchy determine the amount of information communicated to the CEO/president when facing a binary decision. In this case, our results underscore the importance of vice-presidents in hierarchical organizations: By choosing an appropriate vice president, the CEO/president can maximize his payoff and avoid inefficiencies in communication.

We model this kind of hierarchical communication in the framework of Bayesian persuasion developed by Kamenica and Gentzkow (2016). There is a chain of agents from an initial sender to the final receiver. The initial sender, who is the only agent with direct access to the state, intends to persuade the receiver to take some action in his favor by revealing (potentially partially) her information. However, she does not have direct access to the receiver and she can only reveal information to the next agent in the chain; the next agents in the chain do the same successively. The bottom line is that each agent, except for the first one, can only conceal all or part of the information she receives from the previous agent; she cannot reveal more information than what she receives as she does not have direct access to the state. The sequential nature of the game and the assumption that each agent, when choosing her revelation strategy, only observes the revelation strategies of the preceding players and not the information revealed by them leads us to use subgame perfect equilibrium as the solution concept.

First, we prove a version of the revelation principle for this setting: It is without loss of generality to focus on equilibria where the initial sender chooses how much information she will reveal to the receiver, as in the single-sender scenario in Kamenica and Gentzkow (2016), subject to that choice being incentive compatible for all intermediaries. This helps us formulate the problem recursively which implies that solving a hierarchical persuasion game is equivalent to solving a single-sender persuasion game subject to recursively-defined incentive compatibility constraints.

We then focus on the binary-action case where receiver will finally choose between two actions. In the special case where the state space is also binary, we fully characterize the equilibrium outcome. Existence of two agents with highly opposed preferences or an agent who would provide no information in a single-sender game would result in no information communicated to the receiver. Interestingly, if this is not the case, the relative location of two specific agents in the hierarchy determines the equilibrium outcome and its efficiency. The degree of information revelation (if any) in this case is determined by the preferences (bias) of one of these agents called the pivotal agent. One implication of this result is that if receiver could add an intermediary of his choice at the end of the hierarchy, i.e., the vice president, he would optimally make her the new pivotal agent. The optimal choice of the vice president is independent of the hierarchy configuration and eliminates any inefficiencies in communication.

The intuition behind the importance of vice presidents is as follows: Since agents move sequentially, every agent has to respect the preferences of the pivotal agent(s) as long as the pivotal agent is higher in the hierarchy, i.e., closer to the receiver. By making the vice president the new pivotal player, receiver, i.e., the CEO/president, forces all agents to respect vice president's preferences.

Next, we let the state space be the interval $[0, 1]$ keeping the action space binary. Focusing on simple equilibria implies that, without loss of generality, the game can be considered a binary-state, binary-action one from the intermediaries' point of view. By choosing how much information to reveal, the initial sender implicitly pins down the preferences of intermediaries in this subgame whose outcome was described above. In general, solving for the optimal choice of the initial sender is not straightforward. We partially characterize the solution in the general case, and fully characterize it in a special case where the prior is uniform, the differential payoff from the actions is linear in the state, and one of the actions yields zero payoff to the initial sender.

The results in this special case are very similar to the binary-state case. However, now the relative location and relative bias of four specific agents in the hierarchy determine the equilibrium outcome and its efficiency. Every equilibrium outcome is equivalent to one which simply distinguishes between two intervals $[0, x]$ and $[x, 1]$ where $x \in [0, 0.5]$ is determined by the preferences (bias) of two of these agents called the pivotal agents. In this general case, if receiver could add an intermediary of his choice at the end of the hierarchy, he would optimally make her one of the new pivotal agents where the optimal choice of the vice president depends on the hierarchy configuration. However, if receiver could add two players of his choice at the end of the hierarchy, he would optimally make them the new pivotal agents where the order does not matter; his choices only depend on the preferences of the current pivotal agents.

Finally, as mentioned before, each agent can only conceal all or part of the information she receives from the previous agent. Following the restriction to simple equilibria, one may wonder if the incentive compatibility constraints can be summarized in the condition that no intermediary prefers to conceal more information. We show that while this simple condition is sufficient for incentive compatibility, it is not necessary.

Related Literature. We contribute to the literature on information design and Bayesian persuasion; see Aumann, Maschler, and Richard (1995), Kamenica and Gentzkow (2016), and Bergemann and Morris (2016a, 2016b). We apply these tools to a setting with multiple senders who move sequentially and study how information revelation depends on senders' preferences and their order.

Ambrus et al. (2013) study the same problem under cheap talk communication while we model it in the framework of Bayesian persuasion. In fact, they extend the classic model of Crawford and Sobel (1982) to investigate intermediated communication. They focus on pure-strategy equilibria and show that the set of pure strategy equilibrium outcomes does not depend on the order of intermediaries and intermediation cannot improve information in these equilibria. These results do not hold in our setting.

Kamenica and Gentzkow (2107a, 2017b) study Bayesian persuasion with multiple senders. Their main assumption is that any sender can choose a signal that is arbitrarily correlated with signals of others. They focus on pure-strategy simultaneous-move equilibria and show that greater competition, e.g., adding senders, tends to increase the amount of information revealed. In our setting, senders move sequentially and design their experiments before observing any signal realizations and their result does not hold. In fact, Li and Norman (2018) show that adding senders can result in a loss of information if any of the following assumptions is violated: (i) information can be arbitrarily correlated, (ii) senders reveal information simultaneously, and (iii) senders play pure strategies.

The most closely related paper is Li and Norman (2021) who study the sequential version of the Bayesian persuasion game with multiple senders considered by Kamenica and Gentzkow (2107a, 2017b). The rich signal space they consider implies that each sender observes the signal realization of the preceding senders' experiments when designing her own experiment and the receiver observes all signals; that is, each sender may only decide whether to provide more information. This allows them to restrict attention to a finite set of vertex beliefs. This is not the case in our setting where each sender only observes the experiments designed by the preceding senders; that is, each player may only decide whether to provide less information. They introduce the notion of one-step equilibria which is the same as simple equilibria in our setting, and similarly formulate the incentive compatibility of intermediaries to characterize those equilibria. Investigating consultation with multiple experts, they show that adding a sender (expert) who moves first cannot reduce informativeness in equilibrium; in our setting, while this question is not relevant, the opposite holds. In fact, in our setting, which models communication in hierarchical organizations, it is more relevant to study the effect of adding a sender who moves last, i.e., the vice president; our results suggest that adding an appropriate such sender can increase informativeness and efficiency in equilibrium.

The rest of the paper is organized as follows. Section 2 sets up the model. In section 3, we prove that it is without loss of generality to focus on the set of simple equilibria where the initial sender is the only agent who may conceal information. Section 4 includes the recursive formulation of incentive compatibility constraints and restates the hierarchical persuasion problem as a single-sender one. In section 5, we fully characterize the equilibrium in the binary-state, binary-action case, and in a special case of the general binary-action case; we also discuss the efficiency of equilibria. Finally, section 6 investigates a simple sufficient condition for the incentive compatibility conditions of section 2.

2 Model

There are n players (she) interested in the action taken by a receiver R (he). Each player can try to influence the receiver's action by sending a message to the next player in a hierarchical manner. Players' and receiver's payoffs depend on the state of the world $\omega \in \Omega$ and the action taken by the receiver $a \in \mathcal{A}$. In other words, each player's and receiver's utility is given by a function $u_i(\omega, a)$ where $i = 1, \ldots, n, R$. All players are expected utility maximizers.

The players and the receiver are uncertain about the state of the world ω and their common prior belief is represented by a probability density function f[1]. However, they can obtain information in a hierarchical manner. Each player $i = 1, \ldots, n$, successively, gets to send a signal $s_i \in \mathcal{S}_i$ to player $i+1$, where $n+1 = R$; in other words, each player i, successively, designs (commits to) an experiment over $\mathcal{S}_{i-1}$, where $\mathcal{S}_0 = \Omega$, including a signal space $\mathcal{S}_i$ and a signal structure $\pi_i : \mathcal{S}_{i-1} \to \Delta(\mathcal{S}_i)$. In the remainder of the paper, I denote an experiment $(\mathcal{S}, \pi)$ simply by its signal structure π. Let Π and Π_i denote the set of all experiments over Ω and $\mathcal{S}_i$, respectively.

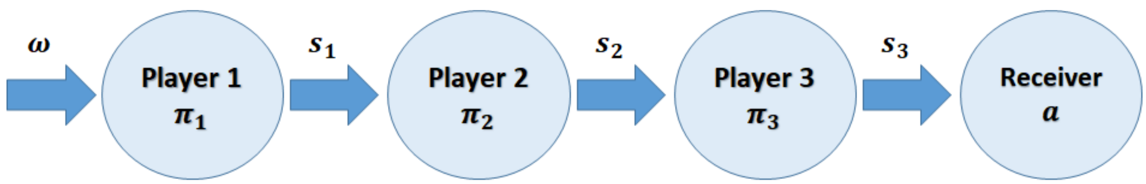

Figure 1: A hierarchy with $n = 3$

When designing her experiment, each player $i = 1, \ldots, n$ observes the experiments designed by the preceding players $\{\pi_j\}_{j=1}^{i-1}$, but not the signal realizations. After each player has designed an experiment, first the state of the world ω, and then, successively, signals $s_1, \ldots, s_n$ are realized according to the designed experiments. The receiver observes the experiments designed by all players and a signal s_n sent by player n; he updates his belief accordingly and chooses an action $a \in \mathcal{A}$ to maximize his expected utility. Note that player 1 is the only player with direct access to the state while player n is the only player with direct access to the receiver.

In summary, the game consists of three stages. In the first stage, each player,

[1] If prior belief is a finite-support distribution, then we simply assume f represents the probability mass function and replace all integrals with sums. The corresponding cumulative distribution function is represented by F.

successively, designs an experiment given the experiments chosen by the preceding players in order to maximize her expected utility. In the second stage, the state of the world is realized and the communication takes place in a hierarchical manner according to the designed experiments. In the last stage, the receiver updates his belief according to the designed experiments and the received signal, and decides what action to take in order to maximize his expected utility.

For the rest of the paper, we assume $\mathcal{S}_1, \mathcal{S}_2, \ldots, \mathcal{S}_n, \mathcal{A}$ are all finite. Therefore, each π_i for $i = 1, \ldots, n$ is a finite-support distribution and π_i represents the conditional probability mass function.

3 Equilibrium

Given the experiments designed by all the players $(\pi_1, \ldots, \pi_n)$ and the signal s_n received from player n, the receiver's expected utility, and thus his optimally chosen action, only depend on his posterior belief represented by a probability density function μ[2]. Therefore, we can, without loss of generality, represent receiver's optimal strategy by $a^* : \Delta(\Omega) \to \mathcal{A}$. Given the posterior belief of the receiver μ, let $v_i(\mu) = \mathbb{E}_\mu u_i(\omega, a^*(\mu))$ for $i = 1, 2, \ldots, n, R$ represent the expected utility of a player or the receiver.

The distribution of the signal observed by player i, that is, s_{i-1}, depends on the experiments designed by all the preceding players. The aggregate experiment observed by player i is denoted by $\pi^i : \Omega \to \Delta(\mathcal{S}_{i-1})$ where

$$\pi^i(s_{i-1}|\omega) = \sum_{s'_1} \ldots \sum_{s'_{i-2}} \pi_1(s'_1|\omega)\pi_2(s'_2|s'_1) \ldots \pi_{i-1}(s_{i-1}|s'_{i-2}).$$

For the ease of exposition, we write $\pi^i \equiv \pi_1 \circ \ldots \circ \pi_{i-1}$. After observing the preceding players' designed experiments, player i's expected utility, and thus her optimally chosen experiment, only depend on π^i, not the individual experiments designed by

[2]If posterior belief is a finite-support distribution, then we simply assume μ represents the probability mass function.

the preceding players $(\pi_1, \ldots, \pi_{i-1})$. Note that by designing an experiment, each player simply garbles the aggregate experiment designed by the preceding players, that is, $\pi^{i+1} \equiv \pi^i \circ \pi_i$ is less Blackwell-informative than π^i.

Given the receiver's optimal strategy a^*, every profile of experiments designed by the players $(\pi_1, \ldots, \pi_n)$ induces a conditional distribution over the actions taken by the receiver $\pi : \Omega \to \Delta(\mathcal{A})$ called the outcome of the game. Without loss of generality, let's assume $\mathcal{S}_n = \mathcal{A}$ and that the receiver's optimal action is equal to the signal s_n sent by player n[3]: $a^*(\mu_{s_n}) = s_n, \ \forall s_n \in \mathcal{S}_n$. Therefore, the outcome of the game is given by

$$\pi(a|\omega) = \sum_{s'_1} \ldots \sum_{s'_{n-1}} \pi_1(s'_1|\omega)\pi_2(s'_2|s'_1) \ldots \pi_n(a|s'_{n-1}), \ \forall a \in \mathcal{A}, \forall \omega \in \Omega,$$

which is the aggregate experiment observed by the receiver. For the ease of exposition, we write $\pi \equiv \pi_1 \circ \ldots \circ \pi_n$. It is now clear that the receiver's posterior belief is given by the probability density function μ where

$$\mu(\omega|\pi_1, \ldots, \pi_n, s) = \frac{f(\omega)\pi(s|\omega)}{\int_\Omega f(\omega')\pi(s|\omega')d\omega'}.$$

Furthermore, every outcome of the game π induces a distribution of posteriors for the receiver τ_π. We sometimes refer to τ_π as the outcome of the game. Since $\mathcal{S}_n$ is finite, τ_π is a finite-support distribution and we simply let τ_π represent the probability mass function. Note that the above assumption implies that each signal s_n sent by player n induces a distinct posterior belief for the receiver μ_{s_n}. Therefore,

$$\tau_\pi(\mu_a) = \int_\Omega \sum_{s'_1} \ldots \sum_{s'_{n-1}} f(\omega)\pi_1(s'_1|\omega)\pi_2(s'_2|s'_1) \ldots \pi_n(a|s'_{n-1})d\omega = \int_\Omega f(\omega)\pi(a|\omega)d\omega.$$

Definition 1. (Subgame Perfect Equilibria) A subgame perfect equilibrium of the game $\sigma^* = (\pi_1^*, \sigma_2^*, \ldots, \sigma_n^*)$ is defined by an experiment $\pi_1^* \in \Pi$ for player 1 and a function (strategy) $\sigma_i^* : \Pi \to \Pi_{i-1}$ for each player $i = 2, \ldots, n$ such that

[3]The proof is the same as that of Proposition 1 in Kamenica and Gentzkow (2016).

- given $\{\sigma_j^*\}_{j=2}^n$, π_1^* maximizes $\mathbb{E}v_1(\mu)$, and
- for $i = 2, \ldots, n$, given $\{\sigma_j^*\}_{j=i+1}^n$, $\sigma_i^*(\pi^i)$ maximizes $\mathbb{E}v_i(\mu)$, $\forall \pi^i \in \Pi$.

For a subgame perfect equilibrium of the game, let π_i^* denote the experiment chosen by player i on the equilibrium path. The corresponding equilibrium outcome and equilibrium distribution of posteriors for the receiver are denoted by $\pi^* \equiv \pi_1^* \circ \ldots \circ \pi_n^*$ and τ^*, respectively.

Now, we are ready to state the first result which is reminiscent of the revelation principle.

Proposition 1. (Simple Equilibria) Given a subgame perfect equilibrium of the game σ^* with equilibrium outcome π^*, there exists an outcome-equivalent subgame perfect equilibrium σ^{**} where player 1's experiment is given by $\pi_1^{**} = \pi^*$ and all other players use the following strategy:

$$\sigma_i^{**}(\pi^i) = \begin{cases} \mathcal{I} & \text{if } \pi^i = \pi^* \\ \sigma_i^*(\pi^i) & \text{Otherwise} \end{cases}$$

where $\mathcal{I}$ represents the full-revelation experiment, i.e., $\pi \circ I = \pi, \ \forall \pi \in \Pi$.

This equilibrium is simple in the sense that on the equilibrium path, all players except for the first one pass all the information they receive to the next player.

Proof: If we prove that the strategy profile σ^{**} is indeed an equilibrium strategy profile, then it is straightforward to see that this is outcome-equivalent to the equilibrium strategy profile σ^*.

Given the experiments chosen by the preceding players and the strategies of the succeeding players $(\pi_1, \ldots, \pi_{i-1}, \sigma_{i+1}, \ldots, \sigma_n)$, by choosing experiment π_i, player i essentially chooses the outcome of the game where her choice set is a feasible subset of Π. More explicitly, her choice set includes all outcomes of the game π such that $\pi \equiv \pi_1 \circ \pi_{i-1} \circ \pi_i \circ \sigma_{i+1}(\pi^{i+1}) \circ \ldots \circ \sigma_n(\pi^n)$ for some $\pi_i \in \Pi_{i-1}$.

- Consider player n. The only change in her strategy compared to σ_n^* is when she encounters a profile of experiments chosen by the preceding players such that

$\pi^n = \pi^*$. In this case, she can induce any outcome π less Blackwell-informative than $\pi^* = \pi_1^* \circ \ldots \circ \pi_n^*$. In the equilibrium corresponding to σ^*, when she could induce any outcome π less Blackwell-informative than $\pi_1^* \circ \ldots \circ \pi_{n-1}^*$, including π^*, she chooses to induce π^*. Obviously, the choice set is now smaller but includes the optimal choice of a larger choice set[4]; therefore, it is still optimal for player n to induce π^*; that is, $\sigma_n^{**}(\pi^*) = \mathcal{I}$ is optimal.

- Consider player i for $i = 2, \ldots, n-1$. The only change in her strategy compared to σ_i^* is when she encounters a profile of experiments chosen by the preceding players such that $\pi^i = \pi^*$. In this case, she can induce any outcome π such that $\pi = \pi_1^* \circ \ldots \circ \pi_{i-1}^* \circ \pi_i \circ \sigma_{i+1}^{**}(\pi^{i+1}) \circ \ldots \circ \sigma_n^{**}(\pi^n)$ by choosing π_i where π_i is strictly less Blackwell-informative than $\pi_i^* \circ \ldots \circ \pi_n^*$; she can also induce $\pi = \pi^*$. In the equilibrium corresponding to σ^*, when she could induce any outcome π of the form $\pi_1^* \circ \ldots \circ \pi_{i-1}^* \circ \pi_i \circ \sigma_{i+1}^*(\pi^{i+1}) \circ \ldots \circ \sigma_n^*(\pi^n)$ by choosing π_i, she chooses to induce π^*. Obviously, the choice set is now smaller but includes the optimal choice of a larger choice set; therefore, it is still optimal for player i to induce π^*; that is, $\sigma_i^{**}(\pi^*) = \mathcal{I}$ is optimal.

- Consider player 1. By choosing π_1, she can either choose any outcome π such that $\pi = \pi_1 \circ \sigma_2^{**}(\pi^2) \circ \ldots \circ \sigma_n^{**}(\pi^n) \neq \pi^*$, or choose $\pi = \pi^*$. In the equilibrium corresponding to σ^*, when she could induce any outcome π of the form $\pi_1 \circ \sigma_2^*(\pi^2) \circ \ldots \circ \sigma_n^*(\pi^n)$ by choosing π_1, she chooses to induce π^*. Obviously, the choice set is the same; therefore, it is still optimal for player 1 to induce π^*; that is, $\pi_1^{**} = \pi^*$ is optimal.

■

Given a subgame perfect equilibrium of the game σ^*, let $\sigma_i^*(h_j)$ represent the equilibrium strategy of player $i \geq j$ following a history of the game $h_j = (\pi_1, \ldots, \pi_{j-1})$ (alternatively, in the subgame starting from player j). Similarly, let $\pi_i^*(h_j)$ denote

[4]Blackwell information ranking is transitive.

the experiment chosen by player $i \geq j$ on the equilibrium path of the continuation game (alternatively, subgame).

Corollary 1. Given a subgame perfect equilibrium of the game σ^* with equilibrium outcome π^*, there exists an outcome-equivalent subgame perfect equilibrium σ^{**} where following any history of the game $h_j = (\pi_1, \ldots, \pi_{j-1})$ (alternatively, in the subgame starting from player j), player j's strategy is given by $\sigma_j^{**}(h_j) = \pi_j^*(h_j) \circ \pi_{j+1}^*(h_j) \circ \ldots \circ \pi_n^*(h_j)$ and all succeeding players use the following strategy:

$$\sigma_i^{**}(h_j, \pi_j, \ldots, \pi_{i-1}) = \begin{cases} \mathcal{I} & \text{if } \pi_j \circ \pi_{j+1} \circ \ldots \circ \pi_{i-1} = \\ & \quad \pi_j^*(h_j) \circ \pi_{j+1}^*(h_j) \circ \ldots \circ \pi_n^*(h_j) \\ \pi_i^*(h_i) \circ \pi_{i+1}^*(h_i) \circ \ldots \circ \pi_n^*(h_i) & \text{Otherwise} \end{cases}$$

where $h_i = (h_j, \pi_j, \ldots, \pi_{i-1})$.

This equilibrium is simple in the sense that following any history of the game, all players except for the first one to play pass all the information they receive to the next player.

Corollary 2: It is without loss of generality to assume $\mathcal{S}_1 = \ldots = \mathcal{S}_n = \mathcal{A}$.

4 Incentive Compatibility Constraints and Recursive Formulation

Given Proposition 1, in order to find the equilibrium outcome of the game π^*, or equivalently, the equilibrium distribution of posteriors for the receiver τ^*, we can simply focus on the "simple" equilibria where the first player chooses the outcome and others simply pass on the information. However, the chosen outcome must be Bayes plausible and it must be incentive compatible for other players to pass on that information. In this section, we first formally characterize these conditions, and then, formulate the problem as a set of recursive optimization problems.

Consider a pair of experiments $\pi, \pi' \in \Pi$ and their corresponding distributions of posteriors $\tau, \tau' \in \Delta(\Delta(\Omega))$. An experiment π is more Blackwell-informative than π' if and only if τ' is smaller than τ in the convex order: $\tau' \leq_{cx} \tau$. This means for all convex functions $\phi : \Delta(\Omega) \to \mathbb{R}$,

$$\sum_{\mu \in supp(\tau')} \phi(\mu)\tau'(\mu) \leq \sum_{\mu \in supp(\tau)} \phi(\mu)\tau(\mu)^{5}.$$

Now, I formally characterize the necessary conditions for the equilibrium outcome of the game:

- The equilibrium outcome of the game τ^* must be Bayes plausible, i.e., $\sum_{\mu} \mu\tau^*(\mu) = f$[6]. Let Γ_0 denote the set of Bayes-plausible distributions of posteriors.

- As mentioned before, by designing an experiment, the last player simply garbles the aggregate experiment designed by the preceding players. Therefore, in addition to the above condition, the last player should not prefer any outcome less Blackwell-informative than the equilibrium outcome π^*. Equivalently, she should not prefer any distribution of posteriors τ' that is smaller than the equilibrium distribution of posteriors τ^* in the convex order, i.e.,

$$\tau^* \in \Gamma_0, \quad \sum_{\mu} \tau'(\mu)v_n(\mu) \leq \sum_{\mu} \tau^*(\mu)v_n(\mu), \ \forall \tau' \leq_{cx} \tau^*. \tag{1}$$

  Let $\Gamma_n \subseteq \Gamma_0$ denote the set of distributions of posteriors τ satisfying (1)[7].

  Accordingly, we can define a value function representing the highest attainable expected utility of player n as a function of the distribution of posteriors τ^n corresponding to the aggregate experiment π^n designed by the preceding players:

$$V_n(\tau^n) = \max_{\tau' \leq_{cx} \tau^n} \sum_{\mu} \tau'(\mu)v_n(\mu). \tag{2}$$

  This implies that $\Gamma_n = \{\tau \in \Gamma_0 | \sum_{\mu} \tau(\mu)v_n(\mu) = V_n(\tau)\}$.

[5] In the remainder of the paper, for the ease of exposition, we write $\sum_{\mu}$ instead of $\sum_{\mu \in supp(\tau)}$.

[6] The proof is the same as that of Proposition 1 in Kamenica and Gentzkow (2016).

[7] Note that $\tau^* \in \Gamma_0$ and $\tau' \leq_{cx} \tau^*$ imply that $\tau' \in \Gamma_0$.

The following assumption ensures that player n's problem in (2) is well-defined.

Assumption 1: If receiver is indifferent among a set of actions A_μ given his posterior belief μ, he would take player n's favorite action, i.e., $\mathbb{E}_\mu u_n(\omega, a^*(\mu)) \geq \mathbb{E}_\mu u_n(\omega, a),\ \forall a \in A_\mu$.

Moreover, the following assumption simplifies the analysis by focusing on the most informative equilibrium outcome if there exist multiple equilibria.

Assumption 2: When indifferent among some experiments (revelation strategies), a player chooses the most Blackwell-informative one (reveals as much information as possible).

- Again, by designing an experiment, the second to last player simply garbles the aggregate experiment designed by the preceding players. Moreover, the only outcomes the second to last player can induce are those in Γ_n. Therefore, in addition to (1), the second to last player should not prefer any "induceable" outcome less Blackwell-informative than the equilibrium outcome π^*. Equivalently, she should not prefer any distribution of posteriors $\tau' \in \Gamma_n$ that is smaller than the equilibrium distribution of posteriors τ^* in the convex order, i.e.,

$$\tau^* \in \Gamma_n, \quad \sum_\mu \tau'(\mu) v_{n-1}(\mu) \leq \sum_\mu \tau^*(\mu) v_{n-1}(\mu),\ \forall \tau' \in \Gamma_n \text{ s.t. } \tau' \leq_{cx} \tau^*. \quad (3)$$

Let $\Gamma_{n-1} \subseteq \Gamma_n$ denote the set of distributions of posteriors τ satisfying (3).

Accordingly, given V_n, we can define a value function representing the highest attainable expected utility of player $n-1$ as a function of the distribution of posteriors τ^{n-1} corresponding to the aggregate experiment π^{n-1} designed by the preceding players:

$$V_{n-1}(\tau^{n-1}) = \max_{\tau' \leq_{cx} \tau^{n-1}} \sum_{\mu} \tau'(\mu) v_{n-1}(\mu) \text{ subject to} \tag{4}$$
$$\sum_{\mu} \tau'(\mu) v_n(\mu) = V_n(\tau')$$

This implies that $\Gamma_{n-1} = \{\tau \in \Gamma_n | \sum_{\mu} \tau(\mu) v_{n-1}(\mu) = V_{n-1}(\tau)\}$.

- Continuing like this, now consider player i for $i = 2, \ldots, n-2$. Yet Again, by designing an experiment, player i simply garbles the aggregate experiment designed by the preceding players. Moreover, the only outcomes player i can induce are those in Γ_{i+1}. Therefore, in addition to $\tau^* \in \Gamma_{i+1}$, player i should not prefer any "induceable" outcome less Blackwell-informative than the equilibrium outcome π^*. Equivalently, she should not prefer any distribution of posteriors $\tau' \in \Gamma_{i+1}$ that is smaller than the equilibrium distribution of posteriors τ^* in the convex order, i.e.,

$$\tau^* \in \Gamma_{i+1}, \quad \sum_{\mu} \tau'(\mu) v_i(\mu) \leq \sum_{\mu} \tau^*(\mu) v_i(\mu), \ \forall \tau' \in \Gamma_{i+1} \text{ s.t. } \tau' \leq_{cx} \tau^*. \tag{5}$$

  Let $\Gamma_i \subseteq \Gamma_{i+1}$ denote the distributions of posteriors τ satisfying (5).

  Accordingly, given V_{i+1}, we can define a value function representing the highest attainable expected utility of player i as a function of the distribution of posteriors τ^i corresponding to the aggregate experiment π^i designed by the preceding players:

$$V_i(\tau^i) = \max_{\tau' \leq_{cx} \tau^i} \sum_{\mu} \tau'(\mu) v_i(\mu) \text{ subject to} \tag{6}$$
$$\sum_{\mu} \tau'(\mu) v_{i+1}(\mu) = V_{i+1}(\tau')$$

  This implies that $\Gamma_i = \{\tau \in \Gamma_{i+1} | \sum_{\mu} \tau(\mu) v_i(\mu) = V_i(\tau)\}$.

Now, I formally formulate the problem as a set of recursive optimization problems.

Proposition 2. (Equilibrium Distribution of Posteriors) Under Assumption 2, the equilibrium distribution of posteriors for the receiver is given by

$$\tau^* \in \operatorname*{argmax}_{\tau \in \Gamma_2} \sum_{\mu} \tau(\mu) v_1(\mu). \tag{7}$$

Solving a hierarchical persuasion game is equivalent to solving a single-sender persuasion game subject to recursively-defined incentive compatibility constraints.

Proof: As mentioned before, Proposition 1 allows us to simply focus on the "simple" equilibria where the first player chooses the outcome τ^* and others simply pass on the information, as long as τ^* is Bayes plausible and it is incentive compatible for other players to pass on the information. Based on what we have discussed so far, these conditions are equivalent to $\tau^* \in \Gamma_2$.

■

5 Binary-Action Games

We now consider a common special case where the action space is binary: $\mathcal{A} = \{0, 1\}$.

Based on Corollary 2, player 1's experiment can be written as $\pi_1 : \Omega \to \mathbb{R}$. Similarly, player i's experiment can be written as $\pi_i : \{0, 1\} \to \mathbb{R}$ for $i = 2, \ldots, n$. Note that $\pi_i(\cdot)$ represents the conditional probability of sending signal 1 to the next player for $i = 1, \ldots, n$. Following Proposition 1 and Corollary 2, every subgame perfect equilibrium of the game is outcome-equivalent to the one where the first player recommends the receiver what action to take and all the succeeding players pass on the recommendation.

Player 1 chooses an experiment π_1, or equivalently, a Bayes-plausible distribution of posteriors $\tau_1 \in \Delta(\Delta(\Omega))$ where $supp(\tau_1) = \{q_0, q_1\}$. From this point on, the game is a binary-state, binary-action one. In other words, given the experiment chosen by player 1, the game played by players $2, 3, \ldots, n$ is a binary-state one where the state space is given by $\hat{\Omega} = supp(\tau_1)$.

We now consider the binary-state, binary-action game.

5.1 Binary-State Binary-Action Games

Let $\Omega = \mathcal{A} = \{0, 1\}$. Each player's and receiver's utility is represented by four numbers: $u_i^{00} = u_i(\omega = 0, a = 0)$, $u_i^{10} = u_i(\omega = 1, a = 0)$, $u_i^{01} = u_i(\omega = 0, a = 1)$, and $u_i^{11} = u_i(\omega = 1, a = 1)$ for $i = 1, \ldots, n, R$. Let $p, \mu \in \mathbb{R}$ represent the common prior belief and receiver's posterior belief that $\omega = 1$, respectively, and let μ_i represent receiver's posterior belief at which player i is indifferent between the two actions:

$$(1 - \mu_i)u_i^{00} + \mu_i u_i^{10} = (1 - \mu_i)u_i^{01} + \mu_i u_i^{11}$$
$$\mu_i = \frac{u_i^{00} - u_i^{01}}{(u_i^{00} - u_i^{01}) + (u_i^{11} - u_i^{10})}$$

All we need to know about player i is her indifference posterior belief μ_i[8] and whether her preferred action at state $\omega = 1$ is $a = 1$ or $a = 0$. As a result, players can be categorized into different types as follows:

1. 0-Extremists: If $\mu_i > 1$ or $\mu_i < 0$ and player i prefers $a = 0$ at $\omega = 1$, she will always prefer $a = 0$.

2. 1-Extremists: If $\mu_i > 1$ or $\mu_i < 0$ and player i prefers $a = 1$ at $\omega = 1$, she will always prefer $a = 1$.

3. Conformists: If $0 < \mu_i < 1$ and player i prefers $a = 1$ at $\omega = 1$, she will prefer $a = 0$ if $\mu < \mu_i$ and $a = 1$ if $\mu > \mu_i$.

   (a) If $0 < \mu_i < p$, player i is said to be biased[9] toward $a = 1$; the lower μ_i, the higher the bias.

   (b) If $p < \mu_i < 1$, player i is said to be biased toward $a = 0$; the higher μ_i, the higher the bias.

4. Contrarians: If $0 < \mu_i < 1$ and player i prefers $a = 0$ at $\omega = 1$, she will prefer $a = 0$ if $\mu > \mu_i$ and $a = 1$ if $\mu < \mu_i$.

[8] For simplicity, I assume $\mu_i \notin \{0, p, 1\}$ and $\mu_i \neq \mu_j$ if $i \neq j$.

[9] A player is biased toward action a if she prefers that action given the common prior belief p.

(a) If $0 < \mu_i < p$, player i is said to be biased toward $a = 0$; the lower μ_i, the higher the bias.

(b) If $p < \mu_i < 1$, player i is said to be biased toward $a = 1$; the higher μ_i, the higher the bias.

For example, consider the following utility function

$$u_i(\omega, a) = -(a - \omega)^2 - \alpha_i a,$$

where α_i represents player i's willingness (bias) to take action $a = 0$. Then we have $\mu_i = \frac{1+\alpha_i}{2}$. In this example, if $\alpha_i > 1$, player i is a 0-extremist, and if $\alpha_i < -1$, she is a 1-extremist; otherwise, she is a conformist.

Similarly, consider the following utility function

$$u_i(\omega, a) = (a - \omega)^2 - \alpha_i(1 - a),$$

where α_i represents player i's willingness (bias) to take action $a = 1$. Then we have $\mu_i = \frac{1+\alpha_i}{2}$. In this example, if $\alpha_i > 1$, player i is a 1-extremist, and if $\alpha_i < -1$, she is a 0-extremist; otherwise, she is a contrarian.

If receiver is a 0-extremist or a 1-extremist, he will always take action $a = 0$ or $a = 1$, respectively, regardless of the experiments chosen by the players, and every equilibrium of the game is outcome-equivalent to the one where player 1 provides no information. Therefore, all players are indifferent among all experiments.

Suppose receiver is a conformist who is biased toward $a = 1$, i.e., $0 < \mu_R < p$. Define the following sets of players:

- $A = \{i : i \text{ is a conformist with } 0 < \mu_i < \mu_R\}$
- $B = \{i : i \text{ is a conformist with } p < \mu_i < 1\}$
- $C = \{i : i \text{ is a contrarian with } \mu_R < \mu_i < 1\}$
- $D = \{i : i \text{ is a contrarian with } 0 < \mu_i < \mu_R\}$

- $E_0 = \{i : i \text{ is a 0-extremist}\}$
- $E_1 = \{i : i \text{ is a 1-extremist}\}$

The next proposition implies that the relative location of two specific players is relevant in determining the equilibrium outcome of the binary-state, binary-action game:

- Let A^* represent the player with the highest bias among those in A: $\mu_{A^*} = \min_{i \in A} \mu_i$. If $A = \emptyset$, let $A^* = R$.
- If $D \cup E_0 \neq \emptyset$, let E^* represent the player closest to the receiver among those in $D \cup E_0$: $E^* = \max_{i \in D \cup E_0} i$.

Proposition 3: Let $\Omega = \mathcal{A} = \{0, 1\}$ and suppose receiver is a conformist who is biased toward $a = 1$[10] and Assumption 1 holds[11]. If there exists a 1-extremist or a contrarian with $\mu_i > \mu_{A^*}$[12], the equilibrium is characterized by *no-information* outcome[13]. Otherwise,

a. if all players are conformists[14], the equilibrium outcome τ^* is characterized by $supp(\tau^*) = \{0, 1\}$ (*full-information outcome*).

b. if there are players who are not conformists and $A^* < E^*$, the equilibrium is characterized by *no-information* outcome.

c. if there are players who are not conformists and $A^* > E^*$, the equilibrium outcome τ^* is characterized by $supp(\tau^*) = \{\mu_{A^*}, 1\}$[15].

[10] $0 < \mu_R < p$.

[11] Assumption 2 is not necessary in the binary-state, binary-action games.

[12] $C \cup E_1 \neq \emptyset$ or $\exists i \in D$ such that $\mu_i > \mu_{A^*}$.

[13] No information outcome is characterized by $supp(\tau^*) \subset (\mu_R, 1]$ if $a^*(\mu_R) = 0$ or $supp(\tau^*) \subset [\mu_R, 1]$ if $a^*(\mu_R) = 1$, which is determined by Assumption 1.

[14] $D \cup E_0 = \emptyset$.

[15] Note that Assumption 1 is needed here if $A^* = R$; otherwise, there would be no equilibrium.

In the equilibrium, if all players are conformists, receiver gets full information. Otherwise[16]*, location of the pivotal player A^* determines whether receiver gets any information, and if so, her bias μ_{A^*} determines the amount of information communicated to the receiver; the higher the bias of A^*, the more information communicated.*

No-information outcome could emerge because of three reasons:

1. Existence of players who would provide no information in a single-sender game: This is the case if there exists a 1-extremist or a contrarian with $\mu_R < \mu_i < 1$, i.e., $C \cup E_1 \neq \emptyset$.

2. Existence of players with highly opposed preferences: Assuming $C \cup E_1 = \emptyset$, this is the case if there exists $i \in D$ such that $\mu_i > \mu_{A^*}$.

3. Location of the pivotal player: Assuming $C \cup E_1 = \emptyset$ and there exists no $i \in D$ such that $\mu_i > \mu_{A^*}$, this is the case if $A^* < E^*$, i.e., E^* is closer to the receiver than the pivotal player P. This equilibrium is inefficient since all players prefer every outcome τ with $\min(supp(\tau)) \in [0, \mu_{A^*}]$. The equilibrium is inefficient if and only if this condition holds.

The next corollary shows that the configuration 3 above is the only one which causes inefficiency in communication.

Corollary 3: Let $\Omega = \mathcal{A} = \{0, 1\}$ and suppose receiver is a conformist who is biased toward $a = 1$. The equilibrium is inefficient if and only if the conditions of Proposition 3.b. hold.

The reason is that when designing her experiment, E^*, whose preferences are highly opposed to those of A^*, considers only the incentive compatibility constraints of the succeeding players, not those of the preceding players such as A^*. Knowing this, A^* preemptively provides no information.

[16] This is the case except for some rather uninteresting cases resulting in no information for the receiver, as explained in what follows.

One question that may arise is if the receiver, who could be the CEO of a firm or the president, can get more information and thus increase her payoff by assigning a vice president and if so, how. The next corollary shows that the answer is affirmative.

Let D^* represent the player with the lowest bias among those in D with $\mu_i < \mu_{A^*}$: $\mu_{D^*} = \max_{i \in D: \mu_i < \mu_{A^*}} \mu_i$. If $\{i \in D : \mu_i < \mu_{A^*}\} = \emptyset$, let $\mu_{D^*} = 0$.

Corollary 4: Let $\Omega = \mathcal{A} = \{0, 1\}$ and suppose receiver is a conformist who is biased toward $a = 1$. If receiver could add a player of his choice at the end of the hierarchy, he would choose a conformist biased toward $a = 1$ with higher bias than A^* but lower bias than D^*[17]. Receiver's utility would then be increasing in the bias of the added player and the new equilibrium would be efficient. Moreover, he does not benefit from adding more players.

Essentially, every player has to respect the preferences of the pivotal player as long as she is closer to the receiver. By behaving as prescribed in the corollary, receiver makes vice president the new pivotal conformist, and all players have to respect her preferences. As mentioned before, the higher the bias of the pivotal player, the more information communicated.

To generalize Proposition 3 to other types of the receiver, we call conformist the opposite type of contrarian and vice versa. Similarly, we call bias toward $a = 1$ the opposite of bias toward $a = 0$ and vice versa. Also, when two players are of opposite types and bias, we compare their bias in the same way as if they are of the same bias.

Define the following sets of players:

- $A = \{i : i$ is of the same type and bias as receiver but with higher bias$\}$
- $B = \{i : i$ is of the same type as receiver but with opposite bias$\}$
- $C = \{i : i$ is of opposite type to receiver but with the same bias or lower opposite bias$\}$
- $D = \{i : i$ is of opposite type to receiver but with higher opposite bias$\}$

[17] $\mu_{D^*} < \mu_i < \mu_{A^*}$.

- $E_0 = \{i : i \text{ is an extremist of opposite bias to receiver}\}$

- $E_1 = \{i : i \text{ is an extremist of the same bias as receiver}\}$

As in Proposition 3, the relative location of two specific players is relevant in determining the equilibrium outcome of the binary-state, binary-action game:

- Let A^* represent the player with the highest bias among those in A; if $A = \emptyset$, let $A^* = R$.

- If $D \cup E_0 \neq \emptyset$, let E^* represent the player closest to the receiver among those in $D \cup E_0$: $E^* = \max_{i \in D \cup E_0} i$.

Proposition 4: Let $\Omega = \mathcal{A} = \{0, 1\}$ and suppose receiver is not an extremist and Assumption 1 holds. If $C \cup E_1 \neq \emptyset$ or if there exists $i \in D$ such that i is less biased than A^*, the equilibrium is characterized by *no-information* outcome[18]. Otherwise,

a. if $D \cup E_0 = \emptyset$, the equilibrium outcome τ^* is characterized by $supp(\tau^*) = \{0, 1\}$ (*full-information outcome*).

b. if $D \cup E_0 \neq \emptyset$ and $A^* < E^*$, the equilibrium is characterized by *no-information* outcome.

c. if $D \cup E_0 \neq \emptyset$ and $A^* > E^*$, the equilibrium outcome τ^* is characterized by $supp(\tau^*) = \{\mu_{A^*_E}, a\}$ where $a = a^*(p)$ if receiver is a conformist and $a = 1 - a^*(p)$ if he is a contrarian[19].

In the equilibrium, if all players are of the same type as receiver, he gets full information. Otherwise[20]*, location of the pivotal player A^* determines whether receiver*

[18]No information outcome is characterized by (i) $supp(\tau^*) \subset (\mu_R, 1]$ or $supp(\tau^*) \subset [\mu_R, 1]$, or (ii) $supp(\tau^*) \subset [0, \mu_R)$ or $supp(\tau^*) \subset [0, \mu_R]$, whichever includes p; open or closed interval is determined by Assumption 1.

[19]Note that Assumption 1 is needed here if $A^* = R$; otherwise, there would be no equilibrium.

[20]This is the case except for some rather uninteresting cases resulting in no information for the receiver, as explained in what follows.

gets any information, and if so, her bias μ_{A^*} determines the amount of information communicated to the receiver; the higher the bias of A^*, the more information communicated.

No-information outcome could emerge because of three reasons:

1. Existence of players who would provide no information in a single-sender game: This is the case if there exists an extremist of the same bias as receiver or a player of opposite type to receiver but with the same bias or lower opposite bias $\mu_R < \mu_i < 1$, i.e., $C \cup E_1 \neq \emptyset$.

2. Existence of players with highly opposed preferences: Assuming $C \cup E_1 = \emptyset$, this is the case if there exists $i \in D$ such that i is more biased than A^*.

3. Location of the pivotal player: Assuming $C \cup E_1 = \emptyset$ and there exists no $i \in D$ such that i is more biased than A^*, this is the case if $A^* < E^*$, i.e., E^* is closer to the receiver than the pivotal player P. This equilibrium is inefficient since all players prefer every outcome τ with $\min(supp(\tau)) \in [0, \mu_{A^*}]$ or $\max(supp(\tau)) \in [\mu_{A^*}, 1]$ whichever does not include p. The equilibrium is inefficient if and only if this condition holds.

Corollary 5: Let $\Omega = \mathcal{A} = \{0, 1\}$ and suppose receiver is not an extremist. The equilibrium is inefficient if and only if the conditions of Proposition 4.b. hold.

Let D^* represent the player with the lowest bias among those in D with higher bias than A^*. If $\{i \in D : i \text{ has lower bias than } A^*\} = \emptyset$, let $\mu_{D^*} = 1 - a^*(p)$ if receiver is a conformist and $\mu_{D^*} = a^*(p)$ if he is a contrarian.

Corollary 6: Let $\Omega = \mathcal{A} = \{0, 1\}$ and suppose receiver is not an extremist. If receiver could add a player of his choice at the end of the hierarchy, he would choose one of the same type and bias as himself but with higher bias than A^* and lower bias than D^*. Receiver's utility would then be increasing in the bias of the added player and the new equilibrium would be efficient. Moreover, he does not benefit from adding more players.

5.2 General Binary-Action Games

Let $\Omega = [0,1]$. Given the experiment π_1 chosen by player 1, Proposition 3 characterizes the outcome of the binary-state, binary-action subgame starting from player 2. By choosing π_1, player 1 not only chooses the state space of the following subgame $supp(\tau_1) = \{q_0, q_1\}$, but also implicitly the four numbers representing the utilities of the subsequent players[21], and thus their types.

In the general binary-action game, for each player, there is a utility function corresponding to each action: $u_i(\omega, 0)$ and $u_i(\omega, 1)$. Let $\Delta u_i(\omega) = u_i(\omega, 1) - u_i(\omega, 0)$ and suppose $\Delta u_i(\omega) = \alpha_i \omega + \beta_i$ for some $\alpha_i, \beta_i \in \mathbb{R}$. For example, it could be the case that $u_i(\omega, 0) = 0$ and $u_i(\omega, 1) = \alpha_i \omega + \beta_i$. Let $\omega_i = -\frac{\beta_i}{\alpha_i}$ represent the state at which i is indifferent between the two actions. Player i's action preference only depends on the sign of α_i and her posterior mean: whether it is larger or smaller than ω_i[22].

Generally, as mentioned before, player 1 chooses a Bayes-plausible distribution of posteriors $\tau_1 \in \Delta(\Delta(\Omega))$ with $supp(\tau_1) = \{q_0, q_1\}$. From this point on, the game is a binary-state, binary-action one. In the special case where $\Delta u_i(\omega) = \alpha_i \omega + \beta_i$, since action preferences only depend on the posterior mean, player 1 simply chooses a distribution of posterior means $G \in \Delta(\omega)$[23] where $supp(G) = \{m_0, m_1\}$[24], and G is a mean preserving contraction of F (implying $\mathbb{E}_G[\omega] = \mathbb{E}_F[\omega] = m$)[25]. For example, if F is the uniform distribution over $[0,1]$, every pair (m_0, m_1) such that $m_1 - m_0 \leq 0.5$ gives a mean preserving contraction of F with the Bayes-plausible probabilities (i.e.,

[21] $u_i^{lk} = \mathbb{E}_{q_l}[u_i(\omega, k)],\ \forall l, k \in \{0,1\}$

[22] Again, for simplicity, I assume $\omega_i \notin \{0, m, 1\}$ where $m = \mathbb{E}_F[\omega]$ and $\omega_i \neq \omega_j$ if $i \neq j$.

[23] Since G is a finite-support distribution, g represents the probability mass function.

[24] Clearly, $0 \leq m_0 \leq m$ and $m \leq m_1 \leq 1$. Recall that in the binary-state, binary-actions game, each player was represented by the posterior belief μ_i (or equivalently, posterior mean) at which she was indifferent between the two actions. In this special case, in the binary-state, binary-action game starting from player 2, we have $\mu_i = \frac{\omega_i - m_0}{m_1 - m_0}$.

[25] In the special case where $\Delta u_i(\omega) = \alpha_i \omega + \beta_i$, with some abuse of notation, the outcome τ represents a distribution of posterior means.

$g(m_0)m_0 + (1 - g(m_0))m_1 = m$).[26]

Generally, shape of the utility functions u_i and choice of the first player G determine the type of each subsequent player (conformist, contrarian, or extremist) in the following binary-state, binary-action subgame. In the special case where $\Delta u_i(\omega) = \alpha_i\omega + \beta_i$, if $\omega_i < 0$ or $\omega_i > 1$, player i is called an absolute extremist; that is, she will be an extremist regardless of G chosen by player 1. Otherwise, $\alpha_i > 0$ (< 0) implies that player i is a conformist (contrarian). However, given G, and thus, m_0 and m_1, each conformist or contrarian may turn into an extremist: this happens if $\omega_i < m_0$ or $\omega_i > m_1$; also, whenever G turns an i-biased conformist or contrarian into an extremist, she will become an i-extremist. The new extremists are related as follows: if player i with $\omega_i < m$ ($\omega_i > m$) turns into an extremist, all players j with $\omega_j < \omega_i$ ($\omega_j > \omega_i$) turn into extremists as well.

Consider the binary-state, binary-action game starting from player 2 where $\Omega = \{0, 1\}$ and $\mu_i = \omega_i$, for all $i = 2, \ldots, n, R$. We call this game the *reduced* binary-state game corresponding to the general binary-action game.

Lemma 1: Let $\Delta u_i(\omega) = \alpha_i\omega + \beta_i$ for $i = 2, \ldots, n, R$. If the equilibrium outcome of the reduced binary-state game corresponding to the binary-action game is the no-information outcome, the same holds for the original binary-action game[27].

In general, the optimal choice of player 1 is not clear just based on $\Delta u_1(\omega)$. To make analysis tractable, let's assume $u_1(\omega, 0) = 0$[28]. Consider the binary-state, binary-action game where $\Omega = \{0, 1\}$ and $\mu_i = \omega_i$, for all $i = 1, \ldots, n, R$. We call this game the binary-state game corresponding to the general binary-action game.

Lemma 2: Let $\Delta u_i(\omega) = \alpha_i\omega + \beta_i$ for $i = 1, \ldots, n, R$ and $u_1(\omega, 0) = 0$. If the

[26] Note that, generally, if the pair (m_0, m_1) gives a mean preserving contraction of F, any pair (m_0', m_1') also gives a mean preserving contraction of F if $m_0 \leq m_0' \leq m \leq m_1' \leq m_1$.

[27] Lemma 1 can be written more generally as follows: If the equilibrium outcome of the reduced binary-state game starting from player $i \geq 2$ is the no-information outcome, the same holds for the binary-action game starting from player $j < i$.

[28] This implies that $u_1(\omega, 1) = \alpha_1\omega + \beta_1$.

equilibrium outcome of the binary-state game corresponding to the a binary-action game is the no-information outcome, the same holds for the binary-action game[29].

Again, if receiver is a an absolute extremist, he will always take action $a = 0$ or $a = 1$ regardless of the experiments chosen by the players, and every equilibrium of the game is outcome-equivalent to the one where player 1 provides no information[30]. Therefore, all players are indifferent among all experiments.

Suppose receiver is a conformist who is biased toward $a = 1$, i.e., $\alpha_R > 0$ and $0 < \omega_R < m$. Define the sets of players (not including player 1) A, B, C, D, E_0, E_1 as in Proposition 3 by replacing μ_i and μ_R with ω_i and ω_R, respectively. Similarly, define A^* and E^*.

The next proposition implies that the relative location and relative bias of three specific players in addition to E^* are relevant in determining the equilibrium outcome of the binary-action game in the special case where $u_1(\omega, 0) = 0$[31] and the prior belief F is uniform:

- Let P represent the conformist with the highest bias toward $a = 1$[32]: $\omega_P = \min_{i:\alpha_i>0, 0<\omega_i<1} \omega_i$.
- Let B^* represent the player with the highest bias among those in B: $\omega_{B^*} = \max_{i \in B} \omega_i$. If $B = \emptyset$, let $\omega_{B^*} = m$.
- Let B_P^* represent the player with the highest bias among those in B and closer to the receiver than P: $\omega_{B_P^*} = \max_{i \in B: i > P} \omega_i$. If $\{i \in B : i > P\} = \emptyset$, let $\omega_{B_P^*} = m$.

[29] Lemma 2 can be written more generally as follows: If the equilibrium outcome of the binary-state game starting from player i is the no-information outcome, the same holds for the binary-action game starting from player $j \leq i$.

[30] This is clearly true for any binary-action game regardless of the assumptions we have made in this section.

[31] This implies that $u_1(\omega, 1) = \alpha_1 \omega + \beta_1$.

[32] This implies $P = 1$ or $P = A^*$.

Proposition 5: Let $\Omega = [0,1]$, $\mathcal{A} = \{0,1\}$, $\Delta u_i(\omega) = \alpha_i \omega + \beta_i$ for $i = 1, \ldots, n, R$, $u_1(\omega, 0) = 0$ and the prior belief F be the uniform distribution over $[0,1]$. Suppose receiver is a conformist who is biased toward $a = 1$[33], Assumptions 1 and 2 hold and the no-information outcome is not the equilibrium outcome of the corresponding binary-state game. The equilibrium is characterized by *no-information outcome*[34] unless $\omega_{B_P^*} - \omega_P \leq 0.5$[35]; in this case,

a. if there are players who are not conformists[36], $supp(\tau^*) = \{\omega_P, \omega_P + 0.5\}$.

b. if all players are conformists and $\omega_{B^*} - \omega_P > 0.5$, $supp(\tau^*) = \{\omega_P, \omega_P + 0.5\}$.

c. if all players are conformists and $\omega_{B^*} - \omega_P \leq 0.5$[37], $supp(\tau^*) = \{m_0, m_0 + 0.5\}$ where $m_0 = \min\{\omega_P, \max(\omega_{B^*} - 0.5, 0.5\omega_1)\}$[38].

Assuming the equilibrium outcome of the corresponding binary-state game is not the no-information outcome, receiver gets some information if and only if $\omega_{B_P^*} - \omega_P \leq 0.5$*, i.e., there is not a conformist whose preferences are highly opposed to those of the pivotal player* P *and is closer to the receiver. In this case, players* P *and* B^* *are the pivotal players determining the information communicated to the receiver. Every equilibrium outcome is outcome-equivalent to one which simply distinguishes between the two intervals* $[0, m_0]$ *and* $[m_0, 1]$ *for some* $m_0 \in [0, \omega_P]$.

Proof: We first state two lemmas:

Lemma 3: Let $\Omega = [0,1]$, $\mathcal{A} = \{0,1\}$ and $\Delta u_i(\omega) = \alpha_i \omega + \beta_i$ for $i = 1, \ldots, n, R$. Suppose receiver is a conformist who is biased toward $a = 1$, Assumptions 1 and 2

[33] $\alpha_R > 0$ and $0 < \omega_R < m$.

[34] No information outcome is characterized by $supp(\tau^*) \subset (\mu_R, 1]$ if $a^*(\mu_R) = 0$ or $supp(\tau^*) \subset [\mu_R, 1]$ if $a^*(\mu_R) = 1$, which is determined by Assumption 1.

[35] That is, there would be no new 0-extremists closer to the receiver than P if information is provided: $\omega_i \leq \omega_P + 0.5$ for all players $i > P$.

[36] That is, there are players who are absolute 0-extremists or contrarians with $0 < \omega_1 < \omega_P$.

[37] That is, there would be no new 0-extremists if information is provided: $\omega_i \leq \omega_P + 0.5$ for all players.

[38] Equivalently, $m_0 = \max\{\omega_{B^*} - 0.5, \min(\omega_P, 0.5\omega_1)\}$.

hold and the no-information outcome is not the equilibrium outcome of the corresponding binary-state game[39]. Player 1's choice of distribution of posterior means (m_0, m_1) would result in a no-information equilibrium outcome in the following subgame unless:

a. (i) all remaining players are conformists[40], (ii) $m_0 \leq \omega_{A^*}$ and $m_1 \geq \omega_{B^*}$: In this case, player 1's choice of distribution of posterior means (m_0, m_1) would be passed on to the receiver without any change[41].

b. (i) all remaining players are conformists but $m_1 < \omega_{B^*}$, or there are players (except for player 1) who are not conformists, (ii) $m_0 \leq \omega_{A^*}$, (iii) all absolute 0-extremists, contrarians and the new 0-extremists $(\omega_i > m_1)$ have a smaller index than A^*: In this case, player 1's choice of distribution of posterior means (m_0, m_1) would reach the receiver as $(\mu_{A^*}m_1 + (1-\mu_{A^*})m_0 = \omega_{A^*}, m_1)$[42].

Lemma 4: Let $\Omega = [0,1]$, $\mathcal{A} = \{0,1\}$, $\Delta u_i(\omega) = \alpha_i\omega + \beta_i$ for $i = 1, \ldots, n, R$ and $u_1(\omega, 0) = 0$. Suppose receiver is a conformist who is biased toward $a = 1$, i.e., $\alpha_R > 0$ and $0 < \omega_R < m$. The preferences of different types of player 1 over outcomes τ with $supp(\tau) = \{m_0, m_1\}$ are as follows:

- If player 1 is an absolute 1-extremist or a contrarian $(\alpha_1 < 0)$ with $\omega_R < \omega_1 < 1$, she would prefer to provide no information[43].

- If player 1 is an absolute 0-extremist, she would prefer higher m_0 and m_1, as long as it is not outcome-equivalent to no-information equilibrium.

[39] This implies that there exists no absolute 1-extremists or contrarians with $\omega_{A^*} < \omega_i < 1$, and all absolute 0-extremists and contrarians with $0 < \omega_i < \omega_{A^*}$ have a smaller index than A^*.

[40] $C \cup D \cup E_0 \cup E_1 = \emptyset$.

[41] Without Assumption 2, if $A^* \neq R$ and $m_0 = \omega_{A^*}$, another possible equilibrium is no-information outcome. Similarly, if $m_1 = \omega_{B^*}$, in other possible equilibria, player 1's choice of distribution of posterior means (m_0, m_1) would reach the receiver as (m', ω_{B^*}) where $m_0 \leq m' \leq \omega_{A^*}$.

[42] Note that Assumption 1 is needed here if $A^* = R$; otherwise, there would be no equilibrium.

[43] Thus, the equilibrium outcome τ^* is characterized by $supp(\tau^*) \subset (\mu_R, 1]$ if $a^*(\mu_R) = 0$ or $supp(\tau^*) \subset [\mu_R, 1]$ if $a^*(\mu_R) = 1$, which is determined by Assumption 1 (*no-information outcome*). Note that $u_1(\omega, 0) = 0$ is not necessary for this part.

- If player 1 is a contrarian ($\alpha_1 < 0$) with $0 < \omega_1 < \omega_R$, she would prefer higher m_0 and m_1, as long as $m_0 \geq \omega_1$ and it is not outcome-equivalent to no-information equilibrium[44].

- If player 1 is a conformist with $0 < \omega_1 < \omega_R$, she would prefer lower m_0 and higher m_1, as long as $m_0 \leq \omega_1$[45]; otherwise, she would prefer to provide no information.

- If player 1 is a conformist ($\alpha_1 > 0$) with $\omega_R < \omega_1 < m$, she would prefer lower m_0 and higher m_1, as long as it is not outcome-equivalent to no-information equilibrium.

- If player 1 is a conformist ($\alpha_1 > 0$) with $m < \omega_1 < 1$, she would prefer higher m_1 and (i) lower m_0 if $m_1 \geq \omega_1$, (ii) higher m_0 if $m_1 < \omega_1$, as long as it is not outcome-equivalent to no-information equilibrium[46].

Combining Lemma 3 and Lemma 4, it is straightforward to prove Proposition 5.

■

No-information outcome could emerge because of four reasons:

1. *Condition 1*: No-information outcome is the equilibrium outcome of the corresponding binary-state game.

    a. Existence of players who would provide no information in a single-sender game: This is the case if there exists an absolute 1-extremist or a contrarian ($\alpha_1 < 0$) with $\omega_R < \omega_i < 1$.

    b. Existence of players with highly opposed preferences: If condition (a) does not hold, this is the case if there exists a contrarian ($\alpha_1 < 0$) with $\omega_P < \omega_i < \omega_R$.

[44]In this case, if $m_0 = \omega_1$, she would be indifferent about m_1 and Assumption 2 implies this part.

[45]In this case, if $m_0 = \omega_1$, she would be indifferent about m_1 and Assumption 2 implies this part.

[46]In this case, if $m_1 = \omega_1$, she would be indifferent among all m_0, as long as it is not outcome-equivalent to no-information equilibrium and Assumption 2 implies this part.

c. Location of the pivotal player P: If conditions (a) and (b) do not hold, this is the case if $P < E^*$, i.e., E^* is closer to the receiver than the pivotal player P. This equilibrium is inefficient since all players prefer every outcome τ with $\min(supp(\tau)) \in (0, \omega_P]$.

2. *Condition 2*: Location and bias of the pivotal player P: If condition (1) does not hold, this is the case if $\omega_{B_P^*} - \omega_P > 0.5$, i.e., there exists a conformist whose preferences are highly opposed to those of the pivotal player P and is closer to the receiver. This equilibrium is inefficient since all players prefer every outcome τ with $\min(supp(\tau)) \in (0, \omega_P]$.

Condition 1c and Condition 2 are similar: there exists a player with highly opposed preferences to those of the pivotal player P who is closer to the receiver. The next corollary shows that this configuration is the only one which causes inefficiency in communication.

Corollary 7: Let $\Omega = [0, 1]$, $\mathcal{A} = \{0, 1\}$, $\Delta u_i(\omega) = \alpha_i \omega + \beta_i$ for $i = 1, \ldots, n, R$, $u_1(\omega, 0) = 0$ and the prior belief F be the uniform distribution over $[0, 1]$. Suppose receiver is a conformist who is biased toward $a = 1$ and Assumptions 1 and 2 hold. The equilibrium is inefficient if and only if Condition 1c or Condition 2 holds.

Similar to the binary-state case, the reason is that when designing her experiment, E^* or B_P^*, whose preferences are highly opposed to those of P, considers only the incentive compatibility constraints of the succeeding players, not those of the preceding players such as P. Knowing this, P preemptively provides no information.

Similar to the binary-state case, as the next corollary shows, receiver can get more information and thus increase her payoff by assigning a vice-president.

Let D^* represent the player with the highest bias among those in D with $\omega_i < \omega_{A^*}$: $\omega_{D^*} = \min_{i \in D: \omega_i < \omega_{A^*}} \omega_i$. If $\{i \in D : \omega_i < \omega_{A^*}\} = \emptyset$, let $\omega_{D^*} = 0$.

Corollary 8: Let $\Omega = [0, 1]$, $\mathcal{A} = \{0, 1\}$, $\Delta u_i(\omega) = \alpha_i \omega + \beta_i$ for $i = 1, \ldots, n, R$, $u_1(\omega, 0) = 0$ and the prior belief F be the uniform distribution over $[0, 1]$. Suppose receiver is a conformist who is biased toward $a = 1$ and Assumptions 1 and 2 hold.

If receiver could add a player of his choice at the end of the hierarchy, depending on the hierarchy configuration, he would either choose a conformist biased toward $a = 1$ with $\omega_{n+1} = \min\{\omega_P, \max(0.5\omega_R, \omega_{D^*})\}$[47], or a conformist biased toward $a = 0$ with $\omega_{n+1} = \min(0.5\omega_R, \omega_P) + 0.5$. The new equilibrium would be efficient[48].

Similar to the binary-state case, every player is better off by respecting the preferences of the pivotal player P as long as she is closer to the receiver. By behaving as prescribed in the corollary, most of the times, receiver makes vice president the new pivotal player P, and all players are better off by respecting her preferences. He chooses the bias of the vice president such that the outcome is as close as possible to the full-information outcome[49].

However, if all players have similar preferences[50], player 1 is better off by respecting the preferences of (i) not only the pivotal player P, (ii) but also the pivotal player B^*. By behaving as prescribed in the corollary, receiver makes vice president the new pivotal player P or B^*, and player 1 is better off by respecting her preferences. He chooses the bias of the vice president such that the outcome is as close as possible to the full-information outcome. He follows (i) and chooses $\omega_{n+1} = \min(\omega_P, 0.5\omega_R)$ if $\min(supp(\tau^*)) > 0.5\omega_R$ or $\min(supp(\tau^*)) = \omega_P$; otherwise, he follows (ii) and chooses $\omega_{n+1} = \min(0.5\omega_R, \omega_P) + 0.5$.

Corollary 9: Let $\Omega = [0, 1]$, $\mathcal{A} = \{0, 1\}$, $\Delta u_i(\omega) = \alpha_i\omega + \beta_i$ for $i = 1, \ldots, n, R$, $u_1(\omega, 0) = 0$ and the prior belief F be the uniform distribution over $[0, 1]$. Suppose receiver is a conformist who is biased toward $a = 1$ and Assumptions 1 and 2 hold. If receiver could add two players of his choice at the end of the hierarchy, he would choose a conformist biased toward $a = 1$ with $\omega_i = \min\{\omega_P, \max(0.5\omega_R, \omega_{D^*})\}$ and a conformist biased toward $a = 0$ with $\omega_{n+1} = \min(0.5\omega_R, \omega_P) + 0.5$. The order

[47]Equivalently, $\omega_{n+1} = \max\{\omega_{D^*}, \min(0.5\omega_R, \omega_P)\}$.

[48]To determine the new equilibrium outcome using Proposition 5, if there are two players with indifference belief ω_P, we consider the closest one to the receiver as the pivotal player P.

[49]Note that, from receiver's point of view, $supp(\tau) = \{0.5\omega_R, 0.5\omega_R + 0.5\}$ is equivalent to the full-information outcome.

[50]That is, all are conformists and $\omega_{B^*_P} - \omega_P \leq 0.5$.

does not matter and the new equilibrium would be efficient. Moreover, he does not benefit from adding more players.

Similar to the binary-state case, Proposition 5 can be generalized to other types of the receiver.

6 Can Incentive Compatibility Constraints be Simplified?

When formulating the incentive compatibility constraints, we mentioned that the last player should not prefer any outcome less Blackwell-informative than the equilibrium outcome, while player $i = 2, \ldots, n-1$ should not prefer any "induceable" outcome less Blackwell-informative than the equilibrium outcome. One may wonder whether we can simplify the incentive compatibility constraints and say player $i = 2, \ldots, n$ should not prefer any outcome less Blackwell-informative than the equilibrium outcome. To that end, let $\tilde{\Gamma}$ denote the Bayes-plausible distributions of posteriors τ_π such that no player (not considering player 1) prefers any outcome less Blackwell-informative than π:

$$\tilde{\Gamma} = \{\tau \in \Gamma_0 | \sum_\mu \tau(\mu) v_i(\mu) \geq \sum_\mu \tau'(\mu) v_i(\mu), \forall \tau' \leq_{cx} \tau, \forall i = 2, \ldots, n\}.$$

The next proposition illustrates that while this simple condition implies incentive compatibility and is thus a sufficient condition, it is more strict than needed and thus not a necessary condition.

Proposition 6: Under Assumption 2, $\tilde{\Gamma} \subseteq \Gamma_2$, but $\tilde{\Gamma} \neq \Gamma_2$.

Proof: The first part is trivial. The following simple counterexample proves the second part.

Consider a binary-state, binary-action game and let $n = 3$. Suppose $0 < \mu_3 < \mu_R < p < 1 < \mu_2$.

First of all, Γ_0 includes all outcomes τ with $supp(\tau) = \{q_0, q_1\}$ such that $0 \leq q_0 \leq$

$p \leq q_1 \leq 1$ and $\mathbb{E}_\tau[q] = p$. Note that receiver is a conformist who is biased toward $a = 1$, and suppose he takes action $a = 0$ at $\mu = \mu_R$.

In the binary case, for any experiment with the corresponding distribution of posteriors τ, let q_0 and q_1 represent posterior beliefs induced by signals 0 and 1, respectively. The set of distributions of posteriors smaller in the convex order than τ includes all τ' with $supp(\tau') = \{q'_0, q'_1\}$ such that $q_0 \leq q'_0 \leq p \leq q'_1 \leq q_1$ and $\mathbb{E}_{\tau'}[q'] = \mathbb{E}_\tau[q]$. In other words, π is more Blackwell-informative, the lower q_0 is and the higher q_1 is.

Starting with player 3, and considering (1), Γ_3 includes all outcomes $\tau \in \Gamma_0$ such that either $0 \leq q_0 \leq \mu_3$ or $\mu_R < q_0 \leq p$. Note that player 3 is a conformist who is more biased toward $a = 1$ than the receiver.

Now, continuing to player 2, and considering (3), Γ_2 includes all outcomes $\tau \in \Gamma_0$ such that either $q_0 = \mu_3$ or $\mu_R < q_0 \leq p$. Note that player 2 is a 0-extremist.

However, $\tilde{\Gamma}$ only includes all outcomes $\tau \in \Gamma_0$ such that $\mu_R < q_0 \leq p$. This is due to the fact that while player 2 prefers experiments with $q_0 = \mu_3$ to all less Blackwell-informative ones in Γ_3, i.e., those with $\mu_R < q_0 \leq p$, there exist less-Blackwell informative experiments in Γ_0, namely those with $\mu_3 < q_0 \leq \mu_R$, which she prefers even more.

■

References

Ambrus, Attila, Eduardo M. Azevedo, and Yuichiro Kamada (2013), "Hierarchical cheap talk." *Theoretical Economics*, 8(1): 233–261.

Aumann, Robert J., Michael Maschler, and Richard E. Stearns (1995), Repeated games with incomplete information. MIT press.

Bergemann, Dirk, and Stephen Morris (2016), "Bayes correlated equilibrium and the comparison of information structures in games." *Theoretical Economics*, 11(2): 487–522.

Bergemann, Dirk, and Stephen Morris (2016), "Information design, Bayesian persuasion, and Bayes correlated equilibrium." American Economic Review, 106(5): 586–591.

Crawford, Vincent P., and Joel Sobel (1982), "Strategic information transmission." *Econometrica: Journal of the Econometric Society*: 1431–1451.

Kamenica, Emir, and Matthew Gentzkow (2011), "Bayesian persuasion." *American Economic Review*, 101(6): 2590–2615.

Gentzkow, Matthew, and Emir Kamenica (2017), "Bayesian persuasion with multiple senders and rich signal spaces." *Games and Economic Behavior*, 104: 411–429.

Kamenica, Emir, and Matthew Gentzkow (2017), "Competition in persuasion." Review of economic studies, 84(1): 300–322.

Li, Fei, and Peter Norman (2018), "On Bayesian persuasion with multiple senders." *Economics Letters*, 170: 66–70.

Li, Fei, and Peter Norman (2021). "Sequential persuasion." *Theoretical Economics*, 16(2): 639–675.