Journal of Economic Behavior and Organization 242 (2026) 107420

Contents lists available at ScienceDirect

Journal of Economic Behavior and Organization

journal homepage: www.elsevier.com/locate/jebo

Efficient communication in organizations

Federico Vaccari

Department of Economics, University of Bergamo, Via dei Caniana 2, 24127, Bergamo, Italy

ARTICLE INFO

JEL classification:
D23
D82
D83

Keywords:
Information
Communication
Organizations
Efficiency
Costly talk
Signaling

ABSTRACT

This paper studies the organization of communication between biased senders and a receiver in binary decision-making problems. Senders can misreport their private information at a cost. Efficiency is achieved by resolving information asymmetries without incurring these costs—i.e., by attaining the complete-information outcome. Only one communication protocol is efficient, robust to collusion, and free from unnecessary complexities. This protocol has a simple, adversarial, and sequential structure. It always induces efficient equilibria, for which a closed-form characterization is provided. The findings are relevant for the design of organizations seeking to improve decision-making while limiting wasteful influence activities.

1. Introduction

Communication is an essential part of organizations. Decision-makers (owners, top management) often rely on better-informed parties within the same firm (lower-level managers, division heads) to provide relevant information. An agency problem arises when the informed parties' objectives are not aligned with the decision-makers'. In these cases, there are two potential sources of inefficiency. First, informed agents may dissipate considerable resources in influence activities. Second, decision-makers may take wrong decisions due to being poorly informed or swayed. As a result, organizations must structure communication to minimize wasteful influence activities while maximizing decision-makers' accuracy. This paper is concerned with communication protocols that achieve this goal.

To analyze communication protocols, I study a costly signaling game between an uninformed receiver and one or more informed senders. The receiver must select one of two actions. Senders know which action is better for the receiver. Before making a decision, the receiver interacts with the senders in a way pre-determined by a communication protocol. When communicating, senders can misreport information at a cost that is tied to the size of the lie. These "misreporting costs" represent the resources senders allocate to influence the receiver's decision. For example, a manager who shifts subordinates' labor to fabricate data will have fewer resources for the organization's production activities. As a result, the manager may underperform or incur penalties. From the organization's standpoint, these influence activities are unproductive and wasteful.[1]

E-mail address: vaccari.econ@gmail.com, federico.vaccari@unibg.it

[1] For instance, a manager may divert resources from productive tasks to preparing misleading performance reports. While this may increase the likelihood of securing the manager's preferred decision, it simultaneously reduces organizational output (Milgrom and Roberts, 1988). This is an example of how influence activities dissipate resources without improving decision-making. For related examples see, e.g., Kartik et al. (2007), Kartik (2009).

https://doi.org/10.1016/j.jebo.2026.107420
Received 27 May 2025; Received in revised form 1 January 2026; Accepted 6 January 2026
Available online 12 January 2026
0167-2681/

The paper's central finding is that, among all possible arrangements, only one protocol has three desirable properties at the same time: it is efficient, robust to collusion, and minimal. This protocol, referred to as *public advocacy*, consists of the sequential and public consultation of agents with opposed interests. It is a remarkably simple mechanism requiring neither commitment nor mediation. The results concern uniqueness as well as existence: a full equilibrium characterization shows that there always exists an efficient and robust equilibrium induced by the public advocacy protocol.

The notion of efficiency considered here is natural: it requires that all players obtain their complete-information payoff. An efficient protocol is critical in organizations because it solves asymmetric information problems without dissipating resources in unproductive signaling. Likewise, collusion is a common concern in organizations where members can discuss their intentions with each other, and it is relevant for protocols with multiple senders. Minimality requires protocols to have a simple structure, defined solely by the senders' biases and confidentiality. All else being equal, organizations typically prefer adopting simpler rather than more complicated protocols.

Communication protocols, as defined in this paper, specify the organization's mode of communication: how many senders to consult, their relative standing over decision-making, and the confidentiality of their recommendations. For example, a protocol may instruct the receiver to consult only one sender that favors a particular action (such as the head of a division that would benefit from a specific investment). Alternatively, it may specify to consult in a private setting two senders with aligned goals (e.g., members of the same department). In this last case, it is crucial to ensure that communication is not compromised by senders' collusion.

To study collusion, I use the notion of "coalition-proofness" developed by Bernheim et al. (1987). This solution concept allows testing whether a protocol remains effective when players engage in non-binding pre-play communication. It allows senders to discuss their strategies before consultation, but not to make commitments. For example, managers may share their intentions with each other before filing a report to the CEO. Even though managers are unable to make credible and binding commitments, they may still be able to effectively coordinate their actions. As we shall see, coalition-proofness turns out to be crucial for the main result.

The first part of the paper establishes necessary conditions for efficiency and robustness to collusion. An absence of communication results in inaccurate decision-making, whereas single-sender arrangements always result in wasteful persuasion attempts. It follows that multi-sender protocols are necessary for efficiency. They are also sufficient because, in the absence of pre-play communication, senders cannot coordinate persuasion, allowing the receiver to exploit this inability by privately consulting more than one sender. However, private communication protocols become susceptible to collusion if senders have the opportunity to discuss their strategies with each other before being consulted. In such cases, senders can restore coordination, undermining the protocol's effectiveness. No private protocol can simultaneously achieve efficiency and resilience to collusion.

Public protocols prescribe consulting senders through a sequential and public procedure. These arrangements are neither efficient nor robust to collusion when senders have relatively aligned preferences over decision-making. In these cases, there are always contingencies where senders can coordinate persuasion through pre-play communication.

The second part of this paper focuses on the last type of minimal arrangement left to analyze: public advocacy, that is, the sequential and public consultation of senders with conflicting interests over decision-making. The main result shows that public advocacy is efficient and robust to collusion. Importantly, it is the only minimal communication protocol to have these desirable properties. A continuum of efficient equilibria is characterized, showing mechanisms through which the receiver achieves efficiency: the report delivered by the first speaker sets the burden of proof borne by the second speaker, who has to prove its case "beyond a reasonable doubt." The endogenously determined burden of proof ensures that both senders consistently report truthfully. As a result, the receiver learns their private information and makes fully informed decisions. No resources are wasted in the attempt to persuade the receiver. All players obtain the payoff they would get if there were no information asymmetries in the first place.

Finally, the model's assumptions and several extensions are discussed. The results continue to hold, with some caveats, in settings with more than two actions and under general message spaces that allow for withholding and vagueness. The formal analysis is preceded by an introductory example based on courtroom trials, highlighting its relevance for judicial decision-making.

The paper's main contribution is two-fold. First, it identifies signaling structures under which information asymmetries can be resolved efficiently. This stands in sharp contrast to the classical literature on Spence-like signaling, where separating equilibria inherently involve wasteful costs. Second, the results provide an efficiency-based rationale for adversarial systems of judicial decision-making, such as those typically used for trials in common law countries. These findings sharply contrast with criticisms of adversarial systems, which are often viewed as procedurally expensive due to the unnecessary and costly production of misleading information (Tullock, 1980, 1975).

The rest of the paper is organized as follows. Section 2 reviews the related literature, and Section 3 presents the model. Section 4 introduces an illustrative example that conveys the paper's core insights. The main results are in Section 5. Section 6 discusses the model's assumptions and the robustness of the results. Finally, Section 7 concludes.

2. Literature

This paper contributes to the literature on organization design. In this line of work, Milgrom (1988) recognizes that influence activities constitute a direct opportunity cost for organizations. Milgrom (1988) focuses on restrictions of decision makers' discretion as a tool to limit these costs. Differently, the current paper is concerned with the design of communication protocols that eliminate influence activities. Bolton and Dewatripont (1994), Jehiel (1999), and Garicano (2000) study optimal information structures in organizations. In contrast with the current paper, they consider settings where there is no disagreement between players. Without an agency problem, influence activities are not a concern.

Deimen and Szalay (2019) study information structures in organizations where a sender and a receiver have conflicting preferences. In their framework, the optimal structure aligns their interests, yielding fully revealing equilibria. In contrast, the present paper features costly communication, and thus inefficiencies may stem directly from resources dissipated through influence activities. The focus here is on the design of the *communication structure*, rather than on the sender's information.

A strand of literature finds that optimal organization design results in advocacy structures. Dewatripont and Tirole (1999) study the optimal provision of incentives for information gathering. They consider settings with verifiable information and no agency problem. Battaglini (2002) takes a different approach by studying strategic communication with biased senders in a cheap talk framework. Both Dewatripont and Tirole (1999) and Battaglini (2002) make a case for "static" types of advocacy. By contrast, Krishna and Morgan (2001) show the optimality of rebuttals in a cheap talk model where senders engage in an extended debate. The current work differs from these papers in three key aspects. First, it is a model of partially verifiable information.[2] Second, it studies settings where influence activities yield direct costs to the organization that are not informational. Third, the analysis includes a wider array of arrangements, showing that public advocacy is the only communication protocol with several desirable characteristics.[3]

In this model, influence activities are costly due to misreporting costs. The signaling structure considered here relates this paper to the "costly talk" literature (Kartik et al., 2007; Kartik, 2009; Ottaviani and Squintani, 2006). These papers are mainly concerned with the single-sender case. Vaccari (2023) considers the case where two conflicting senders communicate simultaneously. As in traditional signaling models, the equilibria of these settings involve wasteful signaling expenditures, and thus are inefficient. Differently from all these papers, the current work considers a larger class of protocols and equilibria. The role of collusion, which is central here, is absent in this literature.

Efficient outcomes are extremely rare in settings where information is conveyed through costly actions. In canonical signaling models, separating outcomes clear information asymmetries, but involve wasteful signaling costs (Spence, 1973). There are a few exceptions. Bagwell and Ramey (1991) study a limit pricing model with two incumbent firms. They show the existence of an equilibrium that yields the complete-information outcome. Likewise, Emons and Fluet (2009) construct an efficient equilibrium in an adversarial model of costly communication.[4] Results in the current paper show that efficiency can always be obtained in a large class of games with costly communication. Moreover, in this setting, only one specific type of minimal mechanism can yield efficient and coalition-proof outcomes.

Finally, a literature studies mechanisms to elicit correlated information (Crémer and McLean, 1988; McAfee and Reny, 1992). Laffont and Martimort (2000) do so in settings where agents can enter collusive agreements. Similarly, the current paper focuses on the elicitation of perfectly correlated information from multiple senders that can collude. Differently, it studies efficiency in environments where communication is costly, and the cost of reports is type-dependent. Due to this last feature, the revelation principle and the mechanism design approach (Myerson, 1983, 1986) cannot be used here.[5] There are mechanism design papers addressing costly misreporting, such as those by Deneckere and Severinov (2022) and Kartik and Tercieux (2012). The latter focuses on full implementation while requiring no costs in equilibrium. In this paper, the receiver is not subject to the same level of commitment assumed by the mechanism design approach.

3. The model

There are $N \geq 1$ senders in the set $S = \{1, \ldots, N\}$, and one receiver (r). Nature selects a state θ according to some distribution F with density f and full support in $\Theta = \mathbb{R}$. Only the senders observe the realized state. Then, depending on the setting, communication takes place either sequentially or simultaneously. In a sequential protocol, the order of communication is determined by senders' indexes $j \in S$, with sender 1 reporting first, and so on. When communicating, sender $j \in S$ delivers a report $r_j \in \Theta$ with the literal or exogenous meaning "the state is $\theta = r_j$." After observing all the senders' reports, the receiver selects an alternative a in the binary set $\{a^+, a^-\}$.

Payoffs.— Player $i \in S \cup \{r\}$ obtains a payoff of $u_i(a, \theta)$ when the receiver selects action a in state θ. Define player i's payoff-difference in state θ as

$$\Delta u_i(\theta) := u_i(a^+, \theta) - u_i(a^-, \theta).$$

I assume that $\Delta u_i(\theta)$ is weakly increasing in θ for every $i \in S \cup \{r\}$. Under this assumption, the state can be interpreted as a vertical differentiation parameter measuring the relative appeal of a^+ with respect to a^-. I normalize the receiver's payoffs by setting $\Delta u_r(\theta) > 0$ for all $\theta > 0$, $\Delta u_r(\theta) < 0$ for all $\theta < 0$, and $\Delta u_r(0) \geq 0$. For every sender $j \in S$, either $\Delta u_j(\theta) > 0$ for some $\theta < 0$, or $\Delta u_j(\theta) < 0$ for some $\theta > 0$. I say that two senders j and i are opposed-biased if $\Delta u_i(0) \cdot \Delta u_j(0) < 0$, and are like-biased otherwise.

[2] As we shall see, the presence of costly communication implies that each sender j can reasonably deliver only reports within a set, which depends on the true state θ and always includes θ itself. This approach to partially verifiable information is analogous to that in Green and Laffont (1986).

[3] The organization's problem considered here cannot be appreciated by models where information is not partially verifiable: influence activities are impossible when information is fully verifiable, and they come at no cost when information is not verifiable as in cheap talk models.

[4] To the best of my knowledge, those are the only two examples of efficient equilibria in settings where information is conveyed through costly actions. By contrast, efficiency is easier to obtain in cheap talk and disclosure games, where costs are not a concern by default. For example, see Battaglini (2002) and Milgrom and Roberts (1986).

[5] I thank Patrick Rey for this observation.

What is relevant for the analysis is not the exact location of indifference, but rather the direction of the conflict between the sender and the receiver. The sign of $\Delta u_j(0)$ uniquely identifies whether sender j is biased in favor of action a^+ or a^- *relative to the receiver*. Likewise, the sign of $\Delta u_i(0) \cdot \Delta u_j(0)$ determines whether senders i and j are biased toward the same action or not.[6]

Costly talk.— Sender j incurs a finite cost $C_j(r_j, \theta)$ when reporting r_j in state θ. The cost function C_j is continuous in its arguments and strictly increasing in $|r_j - \theta|$, with $C_j(\theta, \theta) = 0$ for every $\theta \in \Theta$. That is, misreporting is increasingly costly in the size of the lie, whereas truthful reporting is costless. Moreover, $C_j(r_j, \theta) \to \infty$ as $|r_j - \theta| \to \infty$, reflecting that extreme lies are prohibitively costly. Additionally, for every $r_j \in \Theta$, $C_j(r_j, \theta) > C_j(r_j, \theta')$ if $|r_j - \theta| > |r_j - \theta'|$. That is, reports are cheaper when delivered from closer states.[7] Sender j's total utility is $w_j(r_j, a, \theta) = u_j(a, \theta) - C_j(r_j, \theta)$.

Strategies.— When communication is simultaneous, a pure strategy for sender j is a function $\rho_j : \Theta \to \Theta$. When communication is sequential, a pure strategy for sender j is a function $\rho_j : \Theta^j \to \Theta$. A profile of reports $\{r_j\}_{j=1}^N$ is off-path if, given the senders' strategies, $\{r_j\}_{j=1}^N$ will never be observed by the receiver. A posterior belief function for the receiver is a mapping $p : \Theta^N \to \Delta(\Theta)$ that, given the senders' reports, generates posterior beliefs $p\big(\theta \mid \{r_j\}_{j\in S}\big)$. Given the senders' reports and posterior beliefs p, the receiver selects an action in the sequentially rational set β, where

$$\beta\big(\{r_j\}_{j\in S}\big) = \underset{a\in\{a^+,a^-\}}{\arg\max}\ \mathbb{E}_p\big[u_r(a,\theta) \mid \{r_j\}_{j\in S}\big].$$

I assume that the receiver selects a^+ when indifferent between the two alternatives.

Equilibrium.— The equilibrium concept is the perfect Bayesian equilibrium (PBE). To test for the protocols' robustness against collusion, I use the two related concepts of strong Nash equilibrium (Aumann, 1959) and coalition-proof Nash equilibrium (Bernheim et al., 1987). An equilibrium is *strong* if no coalition of players can jointly deviate so that all players in the coalition get strictly better payoffs. An equilibrium is *coalition-proof* if it is robust to coalitional deviations that are *self-enforcing*. A coalition is self-enforcing if no proper sub-coalition exists whose members, while taking as fixed the actions of players in the original coalition but outside the deviating sub-coalition, can unanimously agree to deviate from the proposed coalitional deviation in a way that makes each of them strictly better off.[8] I refer to equilibria where senders always report truthfully as *truthful*, and to equilibria where the receiver always learn the state as *fully revealing*.

3.1. Definitions

Here I define concepts that are useful for the analysis that follows. A discussion of the model's assumptions ensues in Section 6.

Reach.— Misreporting costs significantly affect information transmission, as certain reports cannot be profitably delivered from specific states. The notion of "reach" captures how far a sender can misreport before incurring a sure loss. Since senders can misreport by either exaggerating or understating the realized state, it is necessary to define two sender-specific, state-dependent thresholds—one for each direction of the lie. Reports that exceed these thresholds at a given state are unprofitable.

I define the reach of sender j in state θ as the report which associated cost offsets j's potential gains. Formally, the reach of sender j in state θ and when $\Delta u_j(0) > 0$ is

$$\bar{r}_j(\theta) := \max\{r_j \in \mathbb{R} \text{ s.t. } |\Delta u_j(\theta)| = C_j(r_j, \theta)\},$$

while when $\Delta u_j(0) < 0$ it is

$$\underline{r}_j(\theta) := \min\{r_j \in \mathbb{R} \text{ s.t. } |\Delta u_j(\theta)| = C_j(r_j, \theta)\}.$$

The reach is computed under the condition that the sender's report is persuasive. Intuitively, in equilibrium sender j will never deliver reports higher than $\bar{r}_j(\theta)$ or lower than $\underline{r}_j(\theta)$, as these reports are strictly dominated by truthful reporting independently of the receiver's decision.

Protocols and minimality.— A communication protocol describes how communication between senders and the receiver takes place. I say that a protocol is *minimal* if it can be described just by the senders' relative bias (i.e., opposed- or like-biased) and by the confidentiality of their reports (i.e., public or private).[9] Single-sender protocols, analyzed in Proposition 1, do not need a description of relative bias and confidentiality. Table 1 illustrates all minimal multi-sender protocols and refers to their respective results. Section 6 and the Online Appendix discuss more complicated protocols.

I refer to protocols that are not minimal as *complicated*. Complicated arrangements can prescribe, for example, third-party mediation, arbitration, rebuttals, or a complex mix of private and public communication. There are three reasons why the paper emphasizes minimal protocols. First, to avoid redundancies. Section 6 discusses how complicated protocols that yield robust and efficient outcomes can be simplified to a minimal protocol while preserving robustness and efficiency. The second reason is a cost-based rationale.

[6] The cutoff at $\theta = 0$ is a normalization adopted to simplify the exposition. Any alternative normalization would yield equivalent results under an appropriate relabelling of preferences.

[7] That C_j is increasing in $|r_j - \theta|$ does not imply $C_j(r_j, \theta) > C_j(r_j, \theta')$ if $|r_j - \theta| > |r_j - \theta'|$. The latter requirement is a sort of single-crossing property, which can be violated if the shape of C_j depends on θ.

[8] The notion of coalition-proofness considers group deviations that are consistent with the model's non-cooperative framework. For a formal definition of strong and coalition-proof Nash equilibrium, see Aumann (1959) and Bernheim et al. (1987), respectively. For a textbook definition of perfect Bayesian equilibrium, see Fudenberg and Tirole (1991).

[9] Similarly, in Goltsman et al. (2009, p. 1400) a protocol is "the sequence in which the parties can talk, and whether their messages are public or private."

Table 1
Minimal multi-sender protocols and related results.

	Biases	
Confident.	lLike-biased	Opposed-biased
Private	Proposition 2 – There are no coalition-proof and efficient PBE	Proposition 2 – There are no coalition-proof and efficient PBE
Public	Proposition 3 – There are no coalition-proof and efficient PBE	Proposition 4 and Corollary 1 – There exist efficient and strong PBE; characterization provided

In complicated protocols, the receiver *must* scrutinize at least three reports. Results in Section 5 show that efficiency can be obtained with two senders only, each delivering a single report. Third, minimality is appealing because of its simplicity. In light of this discussion and the existence result in Proposition 4, an analysis of complicated protocols, while potentially interesting, is necessarily of second-order importance.

More formally, a protocol with $N \geq 1$ senders is minimal if it is fully characterized by

i) The *bias configuration*: the protocol specifies the bias relationship, either like- or opposed-biased, for all pairs of senders $i, j \in S = \{1, \dots, N\}$, with $i \neq j$;

ii) The *confidentiality*: communication is either simultaneous (i.e., private) or sequential (i.e., public). In both cases, each sender reports only once.

  iia) When communication is simultaneous, all senders report without observing the other senders' reports, and no further specification is required;

  iib) When sequential, senders report publicly one after the other. In this case, the protocol explicitly specifies the sequence in which senders communicate.

Most results in Section 5 hold for any $N \geq 2$ and any communication sequence, so these characteristics of minimal protocols will not be emphasized.[10]

Efficiency.— In an efficient equilibrium, all players obtain their respective complete-information payoffs: the receiver selects action a^+ when $\theta \geq 0$ and action a^- otherwise; senders always report truthfully, i.e., $\rho_j(\cdot, \theta) = \theta$ for every $j \in S$ and $\theta \in \Theta$. A communication protocol is efficient if it allows for an efficient equilibrium. Since truthful equilibria are fully revealing while non-truthful ones are wasteful, an equilibrium is efficient if and only if it is truthful. Full revelation is neither necessary nor sufficient for efficiency.[11]

The term "efficiency" used here may differ from that employed in other works, as we typically refer to Pareto or utilitarian efficiency, which do not necessarily coincide with complete-information outcomes. Nevertheless, the main result outlined in Proposition 4 constitutes a Pareto efficient outcome. Utilitarian efficiency can be achieved by "public advocacy" protocols provided that the receiver's gains from accurate decision-making sufficiently outweigh the senders' gains from persuasion. This assumption is reasonable for most organizations and institutions.

4. An introductory example

This section provides a simple introductory example to clarify the main intuitions and help guide the interpretation of the formal analysis that follows. Readers primarily interested in the formal analysis and derivation of results may, without loss, skip this section.

Consider a judicial decision-making problem, where a judge represents the receiver and the attorneys serve as the senders. The justice system acts as a planner that decides how to structure the interaction between the judge and the attorneys.

The judge must acquit (a^+) or convict (a^-) a suspect. The state is a number, θ, representing the strength of exonerating evidence relative to incriminating evidence. When $\theta \geq 0$, the exonerating evidence outweighs the incriminating evidence, and the suspect is considered innocent. When $\theta < 0$, the incriminating evidence dominates, and the suspect is considered guilty. The judge seeks to acquit the innocent (i.e., $a = a^+$ when $\theta \geq 0$) and convict the guilty (i.e., $a = a^-$ when $\theta < 0$), and dislikes making judicial mistakes (i.e., $a = a^-$ when $\theta \geq 0$ or $a = a^+$ when $\theta < 0$).

Attorneys can be either prosecutors or defendants. Prosecutors seek conviction, whereas defendants aim for acquittal, independently of the state. Formally, a sender i is a prosecutor if $u_i(a^-, \theta) = 1$ and $u_i(a^+, \theta) = 0$ for every θ. Likewise, j is a defendant if $u_j(a^+, \theta) = 1$ and $u_j(a^-, \theta) = 0$ for every θ. Therefore, $\Delta u_j = 1$ and $\Delta u_i = -1$. Prosecutors and defendants are opposed biased as $\Delta u_j \cdot \Delta u_j < 0$. All attorneys have equal and perfect access to the strength of the evidence, whereas the judge does not observe θ.

[10] Sequentiality and public observability are distinct features in complicated protocols that allow the receiver to communicate with the senders. Such protocols are relevant when the receiver has information about previously received reports that can be conveyed to senders. However, they are equivalent to minimal sequential protocols whenever the receiver has an incentive to truthfully disclose information about past reports. This is precisely the case, e.g., in the protocol analyzed in Proposition 4. For a formal analysis of protocols with receiver-to-sender communication, see the Online Appendix.

[11] Full revelation may occur alongside misreporting (see, e.g., Ottaviani and Squintani (2006) and Kartik et al. (2007)). Therefore, full revelation is not sufficient for efficiency. To obtain her complete-information payoff in this model, the receiver does not need to learn every state. She just need to know if the state is positive or not. Therefore, full revelation is not even necessary for efficiency. However, every state is fully revealed in the efficient equilibrium characterized by Proposition 4.

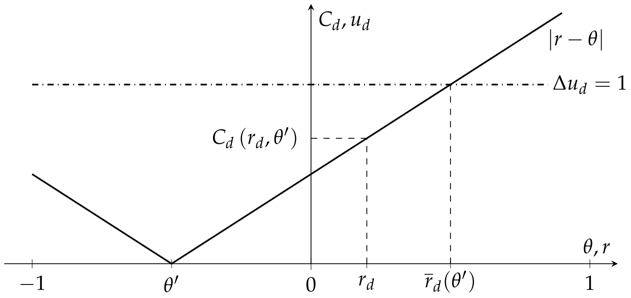

Fig. 1. The sender is a defendant with payoff difference $\Delta u_d = 1$ and cost function $C_d(r, \theta) = |r - \theta|$. The reach in state θ, denoted by $\bar{r}_d(\theta)$, is the highest report such that $|r - \theta| = 1$. A defendant in single-sender protocols, or a coalition of defendants in like-biased protocols, has an incentive to deviate from a truth-telling equilibrium. Reporting r_d in state θ' induces acquittal at a cost lower than the associated persuasion gains.

When called to communicate, attorney j delivers a report r_j, which is a literal statement or submission about the evidentiary record. Attorneys incur a misreporting cost of $|r_j - \theta|$ when claiming that $\theta = r_j$. These costs represent the expected burden or penalty incurred when presenting misleading, exaggerated, or false claims in court. They may include, for example, direct manipulation costs, such as data fabrication or tampering, reputational damage, and formal sanctions if caught misrepresenting evidence.

Misreporting costs incurred by attorneys extend beyond the individual and impose broader externalities on the justice system. While direct manipulation costs-such as fabricating evidence or tailoring arguments-may seem internal to the attorney, they often lead to resource misallocation, procedural inefficiencies, and delays in the delivery of justice. More importantly, when prosecutors or defense attorneys are found to have distorted facts or tampered with evidence, the justice system suffers reputational damage. Public trust in legal institutions is critical: when citizens perceive the system as biased or corrupt, the legitimacy of judicial outcomes erodes, with far-reaching societal consequences.[12]

The justice system is fully aligned with the judge in the sense that it seeks informed and correct decision-making. In light of the above observations, it also incurs *influence costs*. These costs are zero if and only if all attorneys report truthfully (i.e., $r_j = \theta$ for every $j \in S$), and are strictly positive otherwise. The justice system is tasked with arranging the communication between the judge and the attorneys and does so by choosing among minimal protocols with the goal of achieving accurate decision-making while limiting influence costs. There are virtually infinite ways to structure such communication, but as we shall see, only one arrangement proves desirable.

Truthful reporting is key to efficiency. When $r_j = \theta$ for every $j \in S$ and every θ, the judge learns the state perfectly, and the justice system incurs no influence costs. This observation considerably simplifies the analysis, as we only need to test for the existence of equilibria with truthful state reporting. However, defendants may have an incentive to exaggerate the state ($r_j > \theta$), while prosecutors are incentivized to understate it ($r_j < \theta$). Given the attorneys' payoff structure, defendants have a reach of $\bar{r}_j(\theta) = \theta + 1$, while prosecutors have a reach of $\underline{r}_j(\theta) = \theta - 1$. As we shall see, the notion of reach is central to the existence of discriminating signals and, ultimately, to the main result. Fig. 1 illustrates the construction of the reach from the sender's payoff and cost functions, as assumed in this example.

Consider protocols with a single attorney, either a prosecutor or a defendant. Intuitively, the attorney's attempt to extract informational rents leads to misreporting, resulting in inefficient outcomes. Take, for example, a protocol with a single defendant and a truthful equilibrium. In this postulated equilibrium, the judge fully trusts the attorney's report, as no lying occurs. However, when the state lies between -1 and 0, the defendant can exploit this trust by misreporting a positive value within their reach of $\theta + 1$. Doing so triggers acquittal and yields a successful outcome for the defendant.

Because single-sender protocols are inefficient, the justice system must consider arrangements involving more than one attorney. In doing so, it must specify the order in which attorneys present their cases, or whether they submit them privately.

Consider protocols with $N \geq 2$ attorneys, all of whom are either defendants or prosecutors (i.e., like-biased senders). Let us focus on the case of defendants, and suppose they communicate privately or simultaneously. While this arrangement may seem excessively biased, it can trivially yield efficiency by exploiting the attorneys' lack of coordination. For example, with three defendants, a truthful equilibrium can be sustained because individual deviations are easily detected and do not alter the judge's beliefs.

The problem here lies in collusion. The defendants may meet before the trial begins and share their intentions with one another. By doing so, they can act as a single entity by forming a *deviating coalition*, bringing us back to the problem encountered in the single-sender protocol. When the state lies between -1 and 0, all defendants agree to deliver the same positive report within their reach of $\theta + 1$. Importantly, no defendant has an incentive to deviate from the deviating coalition, making the collective departure from truthful reporting self-enforcing. Indeed, any deviation from the collusive agreement would result in a heterogeneous profile of reports, which would arouse the judge's skepticism, as defendants are expected to truthfully report the same evidence.

[12] Reflecting this concern, most legal systems treat lying in court as a punishable offense, underscoring the institutional commitment to incentivize truth-telling and preserve credibility.

The same line of reasoning applies to protocols with like-biased attorneys who communicate sequentially, that is, one after the other. Thus, the justice system must turn to protocols involving a mix of prosecutors and defendants. Consider arrangements with $N \geq 2$ attorneys who communicate privately or simultaneously, of whom $P \geq 1$ are prosecutors and $D \geq 1$ are defendants. In a truthful equilibrium, all attorneys deliver the same report, which coincides with the state. The judge trusts profiles of identical reports and adjudicates accordingly. Unlike in the single-sender case, the judge must now have a choice rule-that is, beliefs—for how to respond when the reports are heterogeneous.

Intuitively, prosecutors and defendants will each form a coalition. All defendants form the coalition of defendants, while all prosecutors form the coalition of prosecutors. Attorneys within each coalition can coordinate their reporting of evidence. For simplicity, we may think of this situation as a two-sender game, where each sender is represented by its coalition. In a truthful equilibrium, the two opposing coalitions report the same evidence.

However, consider the following event. The state θ' is negative and lies between -1 and 0. The coalition of defendants may have an incentive to misreport by submitting $\theta_D > 0$, where $\theta_D < \theta' + 1$ lies within their reach. This coalitional deviation is profitable if and only if the judge acquits after receiving the conflicting reports θ_D and θ'. If not, i.e., if the judge convicts when faced with such discordance, then the deviation will not occur. In that case, the judge's belief system must interpret the discordant profile (θ_D, θ') as evidence favoring conviction.

However, this belief rule creates a new vulnerability: when the true state is positive and equal to θ_D, the coalition of prosecutors can profitably deviate from the conjectured truthful equilibrium by reporting $\theta_P = \theta' < 0$, since $\theta' > \theta_D - 1$ and thus lies within their reach. These coalitional deviations are self-enforcing because any heterogeneity within coalitions would trigger the judge's skepticism. They are also profitable for prosecutors, as $(\theta_D, \theta_P = \theta')$ yields conviction. One way or another, truthful equilibria break down, as there is always a coalition of attorneys that can profitably deviate in a self-enforcing way.

Once the justice system has ruled out all the previously discussed protocols, only sequential arrangements involving both prosecutors and defendants remain. Within this class, consider the simplest case: one prosecutor and one defendant. The prosecutor (π) speaks first, followed by the defendant (d). We analyze this situation using backward induction. Given r_π, the defendant must decide what to report. Suppose the state is positive, and r_π is also positive, indicating that acquittal is appropriate. In that case, the judge would acquit, since even the prosecutor explicitly supports that outcome. The defendant, therefore, has no reason to lie.

More interesting is the case where $r_\pi < 0$, indicating that conviction is appropriate. In this scenario, the judge needs to double-check the prosecutor's claim with the defendant, given the prosecutor's conflict of interest. Thanks to the protocol's sequential nature, the judge can now impose a burden of proof on the defendant that must be met to justify acquittal. Assuming the prosecutor is truthful and the true state is indeed negative and equal to r_π, it would not be in the defendant's interest to claim that the state is greater than or equal to their reach at $\theta = r_\pi$, which is $\bar{r}_d(r_\pi) = r_\pi + 1$. In that case, the misreporting cost would outweigh any benefit from persuading the judge to acquit.

The burden of proof is established: the defendant can convince the judge to acquit if and only if they deliver a report greater or equal than the maximum of zero and $\bar{r}_d(r_\pi)$. The former ensures that the state is actually positive; the latter implies that the prosecutor must have lied in the first place.[13]

Whenever the state is actually positive and the prosecutor claims otherwise, the endogenously determined burden of proof ensures that the defendant can always find a profitable way to secure acquittal. Because any lie by the prosecutor is costly and ultimately ineffective, the prosecutor always reports the state truthfully. The burden of proof also guarantees that, when the prosecutor is truthful, the defendant never finds it profitable to lie under any circumstance.

On the equilibrium path, both attorneys report the evidence truthfully, allowing the judge to learn the state and adjudicate correctly. Importantly, prosecutors and defendants cannot profitably form a coalition due to their opposing interests, making the equilibrium both efficient and resistant to collusion. Any other minimal arrangement is either inefficient or prone to collusion, whereas more complicated protocols add unnecessary complexities.

This result provides an efficiency-based foundation for organizations' choice of protocol. A justice system that values efficiency and seeks to achieve it in the simplest possible way has a compelling reason to implement sequential arrangements involving conflicting attorneys. By doing so, it can replicate complete-information outcomes with minimal interaction between players.

It is important to emphasize that signaling costs are central to achieving efficiency, as the burden of proof relies critically on the existence of discriminatory reports. This highlights a sharp contrast with the cheap talk framework. The remainder of the paper formally shows that the results and intuitions developed in this example extend to much more general environments, with broader payoff and signaling structures.

5. Communication protocols

The end goal of the first part of this section is to rule out communication protocols that are either inefficient or not robust to senders' collusion. A protocol with no senders is not efficient because it involves decision-making under risk. The following proposition shows that even single-sender protocols are not efficient.

Proposition 1 (Single-sender protocols)**.** *There are no efficient PBE if* $N = 1$.

[13] There are other outcome-equivalent equilibria with different burdens of proof. See Section 5.1.

Proof. Consider a protocol where $N = 1$ and $\Delta u_1(0) > 0$ (the proof is similar for $\Delta u_1(0) < 0$). Suppose that there exists an efficient PBE where $\rho_1(\theta) = \theta$ for every $\theta \in \Theta$. The receiver's posterior beliefs are degenerate on $\theta = r_1$. Sender 1 can profitably deviate by delivering $r_1 \in [0, \bar{r}_1(\theta'))$ in some $\theta' < 0$ where $\bar{r}_1(\theta') > 0$, contradicting the existence of a PBE in truthful strategies. □

Single-sender protocols are inefficient because there are always states in which the consulted sender misrepresents information. With no other source of advice, the receiver cannot cross-validate reports to spur truthful reporting. Proposition 1 implies that multi-sender protocols are necessary for efficiency. The following proposition considers arrangements where the receiver privately or simultaneously consults multiple senders with any bias configuration.

Proposition 2 (Simultaneous communication)**.** *Consider protocols with $N \geq 2$ senders that communicate simultaneously. Efficient PBE of these protocols are not coalition-proof.*

Proof. Consider first the case where all senders are like-biased, i.e., $\Delta u_j(0) > 0$ for all $j \in S$ (the proof is similar if $\Delta u_j(0) < 0$ for all $j \in S$). There exists an efficient PBE where $\rho_j(\theta) = \theta$ for every $j \in S$ and $\theta \in \Theta$, and beliefs p are such that $\beta(\{r_j\}_{j\in S}) = a^-$ if $r_i \neq r_k$ for some $i, k \in S$. Given strategies and beliefs, no sender has an incentive to deviate from truthful reporting. Posterior beliefs are pinned down by Bayes' rule only for the case $r_i = r_j$ for all $i, j \in S$. When all reports are identical, the receiver assigns probability 1 to $\theta = r_j$. Off-path beliefs ensure that individual deviations are not profitable.

However, efficient equilibria of this configuration are not coalition-proof. Consider a state $\vartheta_\epsilon < 0$ such that $\min\{\bar{r}_j(\vartheta_\epsilon)\}_{j\in S} > 0$. There is a coalition formed by all senders in S such that, when the state is ϑ_ϵ, each sender $j \in S$ deviates from truthful reporting by delivering $r_j = r' \in [0, \min\{\bar{r}_j(\vartheta_\epsilon)\}_{j\in S})$. Upon observing $\{r_j = r'\}_{j\in S}$, the receiver selects action a^+. Given p, if some sender delivers a report different than r', then the receiver selects a^-. Therefore, this coalitional deviation is mutually beneficial and self-enforcing: there is no proper sub-coalition that, taking fixed the action of its complement, can agree to deviate from the original deviation in a way that makes all of its members better off. As a result, every efficient PBE of this like-biased configuration is not coalition-proof.

Consider now the case where at least two senders are opposed-biased. That is, there are at least two senders $i, j \in S$, $i \neq j$, such that $\Delta u_i(0) \cdot \Delta u_j(0) < 0$. First, I show that misreporting occurs in every PBE when the number of senders is $N = 2$. Suppose there exists a PBE where misreporting never occurs, that is, where $\rho_1(\theta) = \rho_2(\theta) = \theta$ for every $\theta \in \Theta$. Consider such a truthful equilibrium, two opposed-biased senders with $\Delta u_2(0) < 0 < \Delta u_1(0)$, and a state $\theta = \vartheta_\epsilon > 0$, where ϑ_ϵ is small enough. We have that $\rho_1(\vartheta_\epsilon) = \rho_2(\vartheta_\epsilon) = \vartheta_\epsilon$. To discourage deviations, off path beliefs must be such that $\beta(\vartheta_\epsilon, -\vartheta_\epsilon) = a^+$. However, there always exists an $\vartheta_\epsilon > 0$ such that, when the state is $\theta = -\vartheta_\epsilon$, sender 1 can profitably deviate from the prescribed truthful strategy (i.e, $\rho_1(-\vartheta_\epsilon) = \rho_2(-\vartheta_\epsilon) = -\vartheta_\epsilon$) by reporting $r_1 = \vartheta_\epsilon$, as $\Delta u_1(-\vartheta_\epsilon) > C_1(\vartheta_\epsilon, -\vartheta_\epsilon)$. This contradicts that there exists an equilibrium where misreporting never occurs for the case $N = 2$.

For the case $N > 2$, efficient PBE can exist but are not coalition-proof. Consider protocols with more than two senders that communicate simultaneously. Define the sets of senders

$$Z := \{j \mid \Delta u_j(0) > 0\} \text{ and } Y := \{i \mid \Delta u_i(0) < 0\},$$

and the profile of reports $\tilde{r}_Z = \{r_j\}_{j\in Z}$ and $\tilde{r}_Y = \{r_i\}_{i\in Y}$. Say that $\tilde{r}_L = x$ when $r_j = x$ for all $j \in L$. The receiver's decision rule is $\beta(\tilde{r}_Z, \tilde{r}_Y)$. In a truthful equilibrium, $\tilde{r}_Z = \tilde{r}_Y = \theta$ for all $\theta \in \Theta$. Given a state $\theta = -\vartheta_\epsilon < 0$, beliefs are such that $\beta(-\vartheta_\epsilon, -\vartheta_\epsilon) = a^-$ and $\beta(\vartheta_\epsilon, \vartheta_\epsilon) = a^+$. Suppose that off path beliefs are such that $\beta(\vartheta_\epsilon, -\vartheta_\epsilon) = a^+$, and take ϑ_ϵ small enough[14] so that $\bar{r}_j(-\vartheta_\epsilon) \geq \vartheta_\epsilon$ for every $j \in Z$. Denote by $\tilde{r}^*_Z = \{r^*_j\}_{j\in Z}$ a profile of reports with the following features: $\beta(\tilde{r}^*_Z, -\vartheta_\epsilon) = a^+$, and $\beta(\{r_j\}_{j\in Z}, -\vartheta_\epsilon) = a^-$ for any $\{r_j\}_{j\in Z}$ such that $C_j(r_j, -\vartheta_\epsilon) \leq C_j(r^*_j, -\vartheta_\epsilon)$ for every $j \in Z$, with a strict inequality for some j. The report $\tilde{r}^*_Z$ exists because $\beta(\vartheta_\epsilon, -\vartheta_\epsilon) = a^+$, and therefore $r^*_j \leq \vartheta_\epsilon$ for every $j \in Z$. There is a coalition formed by all senders in Z such that, when the state is $\theta = -\vartheta_\epsilon$, each $j \in Z$ can deviate to $r^*_j \in \tilde{r}^*_Z$. This deviation is mutually beneficial and self-enforcing: by construction of $\tilde{r}^*_Z$, there is no proper sub-coalition that, taking fixed the action of its complement, can agree to deviate from the deviation in a way that makes all of its members better off. Equilibria in truthful strategies of these protocols are not coalition-proof (Bernheim et al., 1987). □

There are two reasons why simultaneous communication protocols may initially seem a promising avenue toward efficiency. First, fully revealing equilibria in truthful strategies exist in simultaneous cheap talk games with three or more senders with any bias type.[15] Second, the receiver can achieve efficiency even by simultaneously consulting two (or more) like-biased senders. In this last case, the receiver can induce truthful reporting by applying skepticism when reports do not coincide. Simultaneous communication protocols can exploit senders' lack of coordination to make any individual persuasion attempt futile. As we have seen, collusion effectively restores senders' ability to coordinate persuasion.

Proposition 2 shows that simultaneous protocols are ineffective when senders can engage in non-binding pre-play communication. The possibility of discussing strategies before consultation allows senders to coordinate persuasion in a self-enforcing way, even though they have no commitment power. This last result implies that sequential communication protocols are necessary to jointly achieve efficiency and robustness to collusion. The following proposition considers arrangements where the receiver sequentially consults multiple like-biased senders.

Proposition 3 (Sequential consultation of like-biased senders)**.** *There are no efficient and coalition-proof PBE if there are $N \geq 2$ like-biased senders that communicate publicly and sequentially.*

[14] If instead beliefs are such that $\beta(\vartheta_\epsilon, -\vartheta_\epsilon) = a^-$, then the proof is similar by considering a state $\theta = \vartheta_\epsilon > 0$ small enough and a deviation from the coalition of senders $i \in Y$.

[15] See, for example, the introductory section in Battaglini (2004). In these equilibria, individual deviations from truthful reporting are immediately detected. The same intuition carries over to the costly talk signaling structure considered here.

Proof. Suppose there is a PBE in truthful strategies where $\Delta u_j(0) > 0$ for all $j \in S$ (the proof is analogous if $\Delta u_j(0) < 0$ for all $j \in S$). In this equilibrium, $\rho_1(\theta) = \rho_n(\{r_j = \theta\}_{j=1}^{n-1}, \theta) = \theta$ for every $\theta \in \Theta$ and $n \in \{2, \dots, N\}$. On-path, the receiver's beliefs are degenerate on $\theta = r_j$, $j \in S$.

Consider a state $\theta' < 0$ such that $\Delta u_j(\theta') > 0$ for all $j \in S$. Given equilibrium beliefs p, construct the set containing all the profiles of reports $\{r'_j\}_{j\in S}$ which satisfy the following condition: they effectively induce $a = a^+$, and each sender j would profit from delivering r'_j conditional on $a = a^+$. That is, $u_j(\theta') > C_j(r'_j, \theta')$. Denote such a set by

$$\mathcal{K}_0(\theta', p) := \left\{ \{r'_j\}_{j\in S} \in \Theta^N \mid u_j(\theta') > C_j(r'_j, \theta') \; \forall\, j \in S \;\text{ and }\; \beta(\{r'_j\}_{j\in S}) = a^+ \right\}.$$

There is always a $\theta' < 0$ and a $r' > 0$ such that $\mathcal{K}_0(\theta', p)$ is non-empty, and where, e.g., $\{r_j = r'\}_{j\in S} \in \mathcal{K}(\theta', p)$. Next, construct the set $\mathcal{K}_1(\theta', p)$ as containing all the profiles of reports $\{r'_j\}_{j\in S} \in \mathcal{K}_0(\theta', p)$ that, compared to every other profile $\{r''_j\}_{j\in S} \in \mathcal{K}_0(\theta', p)$, they additionally satisfy $u_1(\theta') - C_1(r'_1, \theta') \geq u_1(\theta') - C_1(r''_1, \theta')$. Iteratively construct the set $\mathcal{K}_n$, where $n \in \{1, \dots, N\}$ and

$$\mathcal{K}_n(\theta', p) := \left\{ \{r'_j\}_{j\in S} \in \mathcal{K}_{n-1}(\theta', p) \mid C_n(r'_n, \theta') \leq C_n(r''_n, \theta') \; \forall \; \{r''_j\}_{j\in S} \in \mathcal{K}_{n-1}(\theta', p) \right\}.$$

Clearly, $\mathcal{K}_n(\theta', p) \subseteq \mathcal{K}_{n-1}(\theta', p)$. When $\mathcal{K}_0(\theta', p)$ is non-empty, then the set $\mathcal{K}_N(\theta', p)$ is by definition non-empty and a singleton. As noted before, there are always contingencies where $\mathcal{K}_0(\theta', p)$ is non-empty.

An argument similar to that for simultaneous like-bias protocols (Proposition 2) shows that truthful PBE of this configuration are not coalition-proof. Through pre-play communication, all senders can collectively agree to report in state θ' according to $\{r'_j\}_{j\in S} \in \mathcal{K}_N(\theta', p)$, provided that all senders before them have adhered to this plan. Upon observing $\{r'_j\}_{j\in S}$, the receiver selects action a^+. After observing the tuple of reports $\{r'_j\}_{j=1}^{N-1}$, the last sender N can profitably deliver r'_N, as $\beta(\{r'_j\}_{j=1}^{N}) = \beta(\{r'_j\}_{j\in S}) = a^+$. By induction, the same is true for every sender $j < N$, and thus every $j \in S$.

Therefore, this coalitional deviation is mutually beneficial and self-enforcing: by definition of $\mathcal{K}_N(\theta', p)$, there is no proper sub-coalition that, taking fixed the action of its complement, can agree to deviate from the original deviation in a way that makes all of its members better off. Truthful equilibria of this configuration are not coalition-proof. □

An example helps illustrate the workings of Proposition 3. Consider the case where all senders prefer action a^+. Suppose there exists a truthful equilibrium in which off-path beliefs are such that $\beta = a^-$ whenever $r_i \neq r_j$ for some $i, j \in S$. These beliefs are reasonably motivated by the fact that all senders are equally informed and biased toward action a^+. Therefore, upon observing conflicting reports, the receiver becomes skeptical and selects the senders' least preferred action. In this case, there always exists $\theta' < 0 < r'$ such that all senders can profitably agree to deviate from the truthful equilibrium by delivering r' in state θ', thereby inducing a^+. This coalitional deviation is self-enforcing, as no sub-coalition can profitably deviate by reporting differently.

Sequential protocols with like-biased senders remain problematic because collusion can still be enforced through non-binding pre-play communication. Before consultation, senders can jointly agree to deviate from truth-telling to achieve persuasion. It follows that truthful equilibria and efficient outcomes cannot be supported by sequential communication protocols where all senders are like-biased. Next, we turn to analyze sequential protocols with opposed-biased senders.

5.1. Public advocacy

Altogether, Propositions 1, 2, and 3 rule out a large class of communication protocols. This first batch of results implies that efficient and robust outcomes may be achieved only through a specific arrangement: multi-sender protocols where communication is sequential and senders are opposed-biased. I hereafter refer to these arrangements as *public advocacy*, as they involve the sequential and public consultation of senders with conflicting interests. The following proposition shows that public advocacy yields efficient equilibria.

Proposition 4 (Public advocacy). *Consider a protocol with $N = 2$ senders that communicate sequentially and are opposed-biased, i.e., $\Delta u_2(0) < 0 < \Delta u_1(0)$. There exist efficient perfect Bayesian equilibria where, for some $\varphi \in [0, r_1]$,*

$$\rho_1(\theta) = \theta \quad \forall\, \theta \in \Theta,$$

$$\rho_2(r_1, \theta) = \begin{cases} \min\left\{\underline{r}_2(\varphi), \theta\right\} & \text{if } \theta < 0 \leq r_1, \\ \theta & \text{otherwise.} \end{cases}$$

$$p \text{ are s.t. } \beta(r_1, r_2) = \begin{cases} a^+ & \text{if } r_1 \geq 0 \text{ and } r_2 > \min\left\{\underline{r}_2(\varphi), 0\right\}, \\ a^- & \text{otherwise.} \end{cases}$$

Proof. Given beliefs p and sender 2's strategy, truthful reporting is strictly dominant for sender 1 in every state. Given beliefs p and sender 1's strategy: (i) if $r_1 < 0$, then truthful reporting is always strictly dominant for sender 2, as $\beta(r_1, \cdot) = a^-$ for every r_2; (ii) if $r_1 \geq 0$ and $\theta \geq 0$, then truthful reporting is strictly dominant for sender 2 because action a^- can be induced only by delivering a $r_2 \leq \underline{r}_2(\varphi)$ which, by definition of reach, is never profitable in states $\theta > \varphi \in [0, r_1]$; (iii) if $r_1 \geq 0$ and $\theta < 0$, then sender 2 can induce action a^- only by delivering some $r_2 \leq \underline{r}_2(\varphi)$. By definition of reach, $r_2 = \min\left\{\underline{r}_2(\varphi), \theta\right\}$ is strictly dominant in this case. Indeed,

$$\underline{r}_2(\theta) < \underline{r}_2(0) \leq \underline{r}_2(\varphi) \leq \underline{r}_2(r_1).$$

Given senders' strategies, beliefs p are pinned down by Bayes' rule only for $r_1 = r_2$, and are free otherwise. When reports are identical, the receiver assigns probability 1 to $\theta = r_1 = r_2$. Off-path beliefs are free and set as in Proposition 4. Since senders play truthful strategies, have no profitable individual deviations, and beliefs are according to Bayes' rule whenever possible, this is a continuum of efficient equilibria parameterized by φ. □

Proposition 4 characterizes efficient equilibria, which allows us to understand the mechanism supporting truthful reporting on the equilibrium path. The key to efficiency in public advocacy stands in how the receiver allocates the burden of proof between the two senders.[16] Beliefs must be consistent with Bayes' rule when senders deliver identical reports in a truthful equilibrium. In these cases, the receiver always follows the senders' recommendations. By contrast, beliefs are free in all those cases where senders disagree. The construction of suitable off path beliefs is crucial in sustaining truthful equilibria.

To illustrate the role of beliefs, consider the case where sender 1 prefers a^+, sender 2 prefers a^-, and $\varphi = 0$. Recall that sender 1 speaks first.[17] When sender 1 recommends its least favorite action a^- (i.e., $r_1 < 0$), the receiver selects a^- no matter what sender 2 reports. By contrast, when sender 1 suggests its favorite action a^+ (i.e., $r_1 \geq 0$), the receiver's decision depends on what sender 2 reports. In this case, sender 2 can convince the receiver to select a^- only by delivering a report dominated in non-negative states (i.e., $r_2 \leq \underline{r}_2(0)$). Intuitively, the receiver follows sender 1's advice to choose a^+ only if sender 2 fails to provide undeniable evidence that the state is negative. The burden of proof allocation ensures that senders have no incentive to deviate from truthful reporting.[18]

Alternatively, when $\varphi = r_1$, sender 2 can convince the receiver to select a^- only by delivering a report that is dominated in states higher than r_1 (i.e., $r_2 \leq \underline{r}_2(r_1)$) and consistent with the state being non-negative (i.e., $r_2 \leq 0$). In this case, sender 2 needs only to provide undeniable evidence that sender 1 lied. This alternative construction of the burden of proof still supports an efficient equilibrium, while requiring the second speaker to expend fewer resources to prove their case. In fact, there exists a continuum of efficient equilibria generated by public advocacy. These equilibria are parameterized by φ, which determines how the receiver's relative skepticism shapes the burden of proof. All of these equilibria are outcome-equivalent and, along the equilibrium path, strategy-equivalent.

The sequential structure of public advocacy is key for efficiency. To illustrate the significance of sequentiality, it is useful to compare the equilibria in Proposition 4 with simultaneous protocols involving two opposed-biased senders as studied in Proposition 2. We have observed that a truthful equilibrium where $\rho_1(\theta) = \rho_2(\theta) = \theta$ for every $\theta \in \Theta$ cannot be sustained by a simultaneous protocol because there are always states where one of the two senders can profitably deviate from the prescribed strategy. By contrast, the sequential protocol in Proposition 4 sustains truthful reporting on the equilibrium path, that is, $\rho_1(\theta) = \rho_2(\theta, \theta) = \theta$ for every $\theta \in \Theta$. Crucially, this is possible because whoever speaks second can condition their report on the one delivered by the first speaker. Given the receiver's beliefs, the second sender has always the option to profitably counteract any persuasion attempt by the first sender. At the same time, whoever speaks second cannot profit from persuasion, as it is always prohibitively expensive.[19]

The adversarial structure of public advocacy provides an additional benefit: the senders, having conflicting goals, cannot coordinate to influence the receiver's decision. Resilience to collusion is desirable in organizations where informed agents can discuss their intentions before being consulted by the receiver. The following corollary confirms that the protocol in Proposition 4 is robust to non-binding pre-play communication.

Corollary 1. *The PBE characterized by Proposition 4 are strong.*

Proof. The proof follows from the observation that there is no profitable coalitional deviation involving two opposed-biased senders. First, any deviation from the prescribed equilibrium entails a cost for at least one sender. Second, a coalitional deviation that makes one sender strictly better off must make the other sender strictly worse off, as they are opposed-biased. Furthermore, the receiver cannot gain from a coalitional deviation because the equilibrium is already efficient. Therefore, the equilibria in Proposition 4 are strong (Aumann, 1959) and coalition-proof (Bernheim et al., 1987). □

An important implication of Proposition 4, together with Corollary 1, concerns the *existence* of efficient outcomes that are robust to collusion. Moreover, when combined with Propositions 1 to 3, we obtain a *uniqueness* result: public advocacy is the only minimal protocol capable of achieving efficiency while remaining resilient to collusion. Of course, organizations do not always have access to two opposed-biased experts who can be arranged to communicate according to the public advocacy structure. The uniqueness result therefore implies that, in such cases, the organization cannot attain complete-information outcomes.

---

16 Since the receiver fully learns the state after sequentially consulting two opposed-biased senders, efficiency can be achieved by public advocacy protocols with $N \geq 2$ senders. The focus on $N = 2$, based on a minimality principle, is therefore without loss of generality.

17 The order in which senders communicate is irrelevant for the result in Proposition 4.

18 As set in Proposition 4, the burden of proof is reminiscent of the "onus" applied in trials before tribunals. In legal disputes, one party is initially presumed to be correct. In contrast, the other party is burdened by providing sufficiently persuasive evidence to prove its case "beyond a reasonable doubt."

19 One might think that we can use the receiver's beliefs as outlined in Proposition 4 to attain truthful equilibria in simultaneous protocols. However, this approach would prove ineffective: in a slightly negative state, say $\theta = \vartheta_\epsilon < 0$, sender 1 can profitably deviate by reporting some r_1' such that $\underline{r}_2(r_1') < \vartheta_\epsilon$ and $C_1(r_1', \vartheta_\epsilon) < \Delta u_1(\vartheta_\epsilon)$. To preserve efficiency, sender 2 would be compelled to misreport and deliver some $r_2 \leq \underline{r}_2(\varphi) < \vartheta_\epsilon$ to correctly induce $\beta = a^-$. This strategy is inherently non-truthful, and thus inefficient.

6. Discussion

Findings in the previous section show that there is a unique communication protocol that is efficient, minimal, and resilient to collusion. This protocol, termed public advocacy, requires the receiver to consult sequentially and publicly two senders with conflicting interests. This section first discusses the model's assumption. Then, it proceeds with a discussion of the findings' robustness, and highlights further differences with respect to related work.

6.1. Discussion of the model

For simplicity, it is assumed that the state space coincides with the real line. However, for the results to carry through, it is sufficient that the state space is large enough to allow the delivery of reports that unambiguously signal the state's sign. Specifically, it is sufficient that

$$\Theta \supseteq \bar{\Theta} := \left[\min\{\underline{r}_j(0)\}_{j\in S}, \max\{\bar{r}_j(0)\}_{j\in S}\right].$$

Intuitively, there is no equilibrium where senders misreport by delivering reports that are strictly beyond their own reach at zero. Every such a report is always equilibrium dominated.[20] It is qualitatively irrelevant for the paper's results if the state space's size is larger than $\bar{\Theta}$.

Second, the model assumes that payoff-differences $\Delta u_i(\theta)$ are increasing in θ. The results do not hinge on this assumption. All findings hold true as long as the payoffs are such that: (i) the sets of states in which a player prefers a particular alternative are convex; (ii) each report remains cheaper when delivered from states that are closer to the report. More formally, for every player j, the sets

$$\Theta_j^+ := \{\theta \in \Theta \mid \Delta u_j(\theta) \geq 0\} \text{ and } \Theta_j^- := \{\theta \in \Theta \mid \Delta u_j(\theta) \leq 0\}$$

must be convex. In addition, for every sender j, report $r_j \in \Theta$, action $a \in \{a^+, a^-\}$, and pair of states θ', θ'' such that $|r_j - \theta'| < |r_j - \theta''|$, the payoffs must satisfy $w_j(r_j, a, \theta') > w_j(r_j, a, \theta'')$. Together, these conditions maintain the results while allowing $\Delta u_j(\theta)$ to be sometimes decreasing in its argument.[21]

To better understand the role of partial verifiability, two benchmarks are relevant. First, if we set $C_j(r_j, \theta) = 0$ for every $r_j, \theta \in \Theta$ and $j \in S$, the model becomes one of cheap talk. In this case, every coalition-proof equilibrium results in babbling in states where some senders have a conflict of interest with the receiver.[22] Senders can credibly transmit information only when recommending the action they are biased against. As a result, there are no efficient and robust protocols. Second, misreporting is prohibitively expensive when $C_j(r_j, \theta) > |\Delta u_j(\theta)|$ for every $r_j \neq \theta$ and $j \in S$. In this case, every communication protocol results in an efficient and coalition-proof outcome. This paper is concerned with the intermediate case of partially verifiable information where misreporting is possible at a non-prohibitive cost.

6.2. Discussion of the results

The model assumes that senders can neither withhold information nor deliver vague reports. The results do not rely on these limitations. Consider an extension of the model where senders can omit or muddle information at a cost.[23] Every equilibrium of the main model is also an equilibrium of this extended model where the receiver interprets unexpected omissions or vagueness unfavorably. Inefficient protocols remain inefficient in this extended model with additional costly actions, as taking a costly action is by itself inefficient. Public advocacy remains the only efficient, minimal, and robust protocol.[24]

The analysis further presumes that the receiver has no commitment power and cannot communicate with the senders. These modeling assumptions are inconsequential for the results. All findings continue to hold if the receiver is allowed to make binding commitments on the decision rule, to communicate with senders about previously received reports, or both. These extensions are formalized in the Online Appendix.

The analysis is focused on binary decision problems, which are prevalent in numerous social and economic contexts extensively examined in the literature: e.g., voting, judicial decision-making, market entry decisions.[25] It is important to recognize the implications on the results when extending the model beyond the binary-action framework. Emons and Fluet (2009) offer an example that indicates

[20] By definition of reach, reports $r_j > \bar{r}_j(0)$ (resp. $r_j < \underline{r}_j(0)$) are strictly dominated when the realized state is negative (resp. positive). They are unambiguous signals of the state's sign. In some equilibria, reports that are equal or close to a sender's reach at zero are delivered on-path. The state space size should allow for the delivery of such reports. For a detailed discussion, see the analysis in Vaccari (2023).

[21] The analysis in Section 5 requires that, for every sender, the reaches are finite, unique, and monotonic in the state. The misreporting cost functions C_j do not need to be everywhere continuous as long as the reaches are well-defined for every sender.

[22] In the cheap talk version, credible information transmission remains possible in equilibrium, e.g., in extreme states where all players agree on the best course of action.

[23] Vagueness and omissions are easily detectable, whereas misreporting is not. Organizations can replace or decrease the budget of managers known to be perfectly informed and yet purposefully refuse to provide accurate information.

[24] See the Online Appendix for the formal result.

[25] Other examples of widely studied binary decision settings are: medical treatment choices (e.g., opting for surgery or not), hiring decisions, technology adoption (e.g., whether to implement a new technology), and consumer purchase decisions (e.g., whether to buy a product).

the findings are applicable to settings with a continuous action space.[26] However, an element of *limited liability*, which holds naturally under the binary-action framework, is crucial for the uniqueness result. The Online Appendix extends the analysis to a three-action framework. It shows that the presence of a third, "safe" alternative can yield efficient outcomes under private advocacy protocols. This extension suggests that, in frameworks with more than two actions, a "liability condition" is necessary to preserve the uniqueness result outlined in Section 5. Without liability constraints and with more than two actions, the appendix shows that two opposed-biased senders suffice to achieve efficiency, regardless of whether they are consulted simultaneously or sequentially.

In minimal protocols, senders speak only once, and are all subject to the same type of confidentiality. Naturally, these rules do not have to apply to more general arrangements. I refer to non-minimal protocols as *complicated*. Public advocacy is one among many efficient and coalition-proof protocols when considering also complicated ones. For example, we can construct a multi-stage protocol where public advocacy is used in the first stage. Knowing that the public advocacy stage induces truthful reporting, senders communicating in all other stages optimally economize in costs by reporting truthfully. Likewise, they cannot form profitable and self-enforcing deviating coalitions. This class of complicated protocols is large, efficient, and robust to collusion. Similarly, the paper does not analyze protocols with rebuttals, that is, where senders report more than once. The main reason for omitting rebuttals is that they are unnecessary to achieve efficiency.[27] This is in contrast with Krishna and Morgan (2001)'s result that rebuttals are necessary for full revelation in their cheap talk setup.[28] The focus on minimal protocols avoids this type of redundancies with the idea that, all else equal, organizations prefer using simpler to complicated protocols.[29]

There are other differences between the results obtained here and those in Krishna and Morgan (2001)'s cheap talk model, which explores a more complex setup and addresses different questions. Most importantly, while this paper examines both information transmission and influence costs, the latter—captured by the sum $\sum_{i\in N} C_j(r_j,\theta)$—are not a focus in Krishna and Morgan (2001), as such costs are set to zero in cheap talk frameworks. This distinction highlights how the intuition for the main result in this paper cannot be similar to that in Krishna and Morgan (2001), as efficiency in public advocacy arises from the ability to deliver reports that, due to their costs, serve as unequivocal signals of the state's sign.[30] Both approaches contribute to understanding the interplay between communication structures and outcomes in strategic settings.[31]

7. Concluding remarks

This paper shows that there exist a simple institution and a signaling structure that can solve information asymmetries at no cost. At least since Akerlof (1970), it is well known that the presence of asymmetric information can yield inefficient outcomes. Subsequent work by Spence (1973) shows that signaling can resolve information asymmetries but does not eliminate inefficiencies, as signals are wasteful and unproductive. This paper departs from the canonical one-sender-one-receiver setting and uses a costly talk signaling structure. When considering multi-sender protocols, the concern for collusion naturally emerges. Yet, in this setting, it is possible to structure communication in a way that fully restores efficiency without using wasteful signaling expenditures or commitment power. Likewise, there is no need for mediation, arbitration, or other complex arrangements.[32]

The main result has potentially significant implications for the understanding of organizational design. It shows that only one minimal protocol can achieve efficiency under the threat of senders' collusion. This protocol prescribes the sequential and public consultation of two informed agents with conflicting interests. The proposed arrangement has a plain structure and does not require commitment power as ex-ante, and in the interim, the organization adheres to the protocol. Importantly, such an arrangement *always* yields an efficient outcome for any configuration permitted by the model described in Section 3. This finding provides a rationale for using public advocacy structures.

Public advocacy is widespread. As an example, consider the justice system. Trials take place with an adversarial procedure of judicial decision-making, whereby two advocates—prosecutors and defendants—engage in public debates in the courtroom. Other examples of public advocacy include, e.g., managers and ministries competing for budget allocation. Budgeting processes are typically sequential and adversarial. When the contending parties disagree, the receiver—judge, CEO, or prime minister—adjudicates.

---

26 See Section 4.2, p. 147 therein. In a similar model of costly communication with a continuous action space, Emons and Fluet (2009) show that there exist a truthful equilibrium when two biased senders with opposed goals communicate sequentially. Even though it is not discussed, their equilibrium is naturally coalition-proof because of the conflict of interest between senders.

27 Introducing rebuttals would require additional and non-trivial assumptions on the misreporting costs. For example, if such costs are purely psychological, then repeating the same lie may be more costly than lying once. By contrast, direct manipulation costs such as effort and time may not duplicate.

28 See Proposition 1, p. 756 Krishna and Morgan (2001), where full revelation is necessary for efficiency.

29 Proposition 2 applies also to complicated private protocols where some pairs of senders are like-biased and others are opposed-biased. Similarly, Proposition 4 and Corollary 1 equally apply to complicated public advocacy protocols with more than two senders, of which at least two are like-biased.

30 Intuitively, the introduction of misreporting costs should (weakly) increase both information transmission and expenditures in influence activities. Public advocacy is never efficient when talk is cheap because not enough information is transmitted in equilibrium.

31 Other differences are in the protocols examined and results obtained. Krishna and Morgan (2001) show that their rebuttal protocol can fully reveal information when senders are sufficiently aligned. By contrast, this paper demonstrates that efficiency can persist even under state-independent sender preferences, extending the analysis to scenarios where senders' biases are extreme.

32 Mechanisms that involve transfers are inefficient because, compared to the outcome under complete information, at least one player incurs a cost when participating in a transfer. Therefore, allowing for transfers does not affect the paper's results.

The model has two distinguishing features central to the main result. First, it considers organizations which dislike wasteful influence activities. Public advocacy is uniquely optimal only if this is the case. Fully revealing equilibria exist in simultaneous protocols, and full revelation is sufficient for optimality when the organization does not incur direct influence costs. Second, senders can engage in non-binding, pre-play communication. Public advocacy is not the uniquely efficient protocol when collusion is not a concern: the proof of Proposition 2 shows that, in these cases, simultaneous and like-bias protocols can also yield efficient outcomes. It is reasonable to jointly assume a distaste for wasteful activities and concern for collusion for a wide range of organizations and central planners.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

Data availability

No data was used for the research described in the article.

Declaration of interests

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

I am grateful to Yair Antler, Daniele Condorelli, Harry Di Pei, Edoardo Grillo, Navin Kartik, Wataru Kitano, Matthew Mitchell, Santiago Oliveros, Federico Quartieri, and Patrick Rey for helpful comments and suggestions. I also thank the Editor, Sudipta Sarangi, and anonymous referees for their valuable comments. I thank participants at the 2022 YETI Workshop, 2023 Oligo Workshop, 2023 Spring Workshop on Economic Theory at LUISS, APSA's Formal Theory Virtual Workshop, the 2025 PRIN Workshop "Digital Markets: Design and Policy," and the University of Chieti-Pescara. All errors are mine.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at 10.1016/j.jebo.2026.107420

References

Akerlof, G.A., 1970. The market for "lemons": quality uncertainty and the market mechanism. Q. J. Econ. 84 (3), 488–500.
Aumann, R.J., 1959. Acceptable Points in General Cooperative n-Person Games. Princeton University Press. pp. 287–324.
Bagwell, K., Ramey, G., 1991. Oligopoly limit pricing. RAND J. Econ. 22 (2), 155–172.
Battaglini, M., 2002. Multiple referrals and multidimensional cheap talk. Econometrica 70 (4), 1379–1401.
Battaglini, M., 2004. Policy advice with imperfectly informed experts. The B.E. J. Theor. Econ. 4 (1), 1–32.
Bernheim, B.D., Peleg, B., Whinston, M.D., 1987. Coalition-proof nash equilibria I. concepts. J. Econ. Theory 42 (1), 1–12.
Bolton, P., Dewatripont, M., 1994. The firm as a communication network. Q. J. Econ. 109 (4), 809–839.
Crémer, J., McLean, R.P., 1988. Full extraction of the surplus in bayesian and dominant strategy auctions. Econometrica 56 (6), 1247–1257.
Deimen, I., Szalay, D., 2019. Information and communication in organizations. AEA Papers and Proceedings 109, 545–49.
Deneckere, R., Severinov, S., 2022. Signalling, screening and costly misrepresentation. Canadian J. Econ. 55 (3), 1334–1370.
Dewatripont, M., Tirole, J., 1999. Advocates. J. Pol. Econ. 107 (1), 1–39.
Emons, W., Fluet, C., 2009. Accuracy versus falsification costs: the optimal amount of evidence under different procedures. J. Law Econ. Organ. 25 (1), 134–156.
Fudenberg, D., Tirole, J., 1991. Game Theory. MIT Press.
Garicano, L., 2000. Hierarchies and the organization of knowledge in production. J. Pol. Econ. 108 (5), 874–904.
Goltsman, M., Hörner, J., Pavlov, G., Squintani, F., 2009. Mediation, arbitration and negotiation. J. Econ. Theory 144 (4), 1397–1420.
Green, J.R., Laffont, J.-J., 1986. Partially verifiable information and mechanism design. Rev. Econ. Stud. 53 (3), 447–456.
Jehiel, P., 1999. Information aggregation and communication in organizations. Manage. Sci. 45 (5), 659–669.
Kartik, N., 2009. Strategic communication with lying costs. Rev. Econ. Stud. 76 (4), 1359–1395.
Kartik, N., Ottaviani, M., Squintani, F., 2007. Credulity, lies, and costly talk. J. Econ. Theory 134 (1), 93–116.
Kartik, N., Tercieux, O., 2012. Implementation with evidence. Theor. Econ. 7 (2), 323–355.
Krishna, V., Morgan, J., 2001. A model of expertise. Q. J. Econ. 116 (2), 747–775.
Laffont, J.-J., Martimort, D., 2000. Mechanism design with collusion and correlation. Econometrica 68 (2), 309–342.
McAfee, R.P., Reny, P.J., 1992. Correlated information and mechanism design. Econometrica 60 (2), 395–421.
Milgrom, P., 1988. Employment contracts, influence activities, and efficient organization design. J. Polit. Econ. 96 (1), 42–60.
Milgrom, P., Roberts, J., 1986. Relying on the information of interested parties. RAND J. Econ. 17 (1), 18–32.
Milgrom, P., Roberts, J., 1988. An economic approach to influence activities in organizations. Am. J. Sociol. 94, S154–S179.
Myerson, R.B., 1983. Mechanism design by an informed principal. Econometrica 51 (6), 1767–1797.
Myerson, R.B., 1986. Multistage games with communication. Econometrica 54 (2), 323–358.
Ottaviani, M., Squintani, F., 2006. Naive audience and communication bias. Int. J. Game Theory 35 (1), 129–150.
Spence, M., 1973. Job market signaling. Q. J. Econ. 87 (3), 355–374.
Tullock, G., 1975. On the efficient organization of trials. Kyklos 28 (4), 745–762.
Tullock, G., 1980. Trials on Trial: The Pure Theory of Legal Procedure. Columbia University Press.
Vaccari, F., 2023. Competition in costly talk. J. Econ. Theory 213, 105740.