

Receiver-Oriented Cheap Talk Design

Itai Arieli

Ivan Geffner

Moshe Tennenholtz

Abstract

This paper considers the dynamics of cheap talk interactions between a sender and receiver, departing from conventional models by focusing on the receiver’s perspective. We study two models, one with transparent motives and another one in which the receiver can *filter* the information that is accessible by the sender. We give a geometric characterization of the best receiver equilibrium under transparent motives and prove that the receiver does not benefit from filtering information in this case. However, in general, we show that the receiver can strictly benefit from filtering and provide efficient algorithms for computing optimal equilibria. This innovative analysis aligns with user-based platforms where receivers (users) control information accessible to senders (sellers). Our findings provide insights into communication dynamics, leveling the sender’s inherent advantage, and offering strategic interaction predictions.

1 Introduction

The conventional sender-receiver model focuses on the interaction between an informed sender, who is fully aware of the state of the game, and an oblivious receiver. In these interactions, the sender must disclose some of the information to the receiver in order to influence its behavior, and the receiver must then play an action that affects the utility of both players. Relevant information disclosure models include cheap talk [10, 17], Bayesian persuasion [16], and their extensions to multiple senders and receivers [12, 7, 1, 2, 3, 4]. Traditionally, all these interactions are always viewed from the point of view of the sender. However, in order to get a complete picture of the communication model, it is also critical to consider the receiver’s perspective. Moreover, since the sender, as the informed side, has an inherent advantage in cheap talk models, analyzing the best outcomes for the receiver can level this advantage and provide a more reliable prediction of strategic interactions.

We focus on interactions where the sender and the receiver can communicate via cheap talk and analyze two settings. First, we consider a setting where the sender has transparent motives (i.e., where the sender has a state-independent utility) in a similar fashion as Lipnowski and Ravid [17], and we follow with a setting where there is a pre-play stage before the communication phase in which the receiver can *filter* the information that the sender receives about the realized state.

Our first result is providing a geometric characterization of the best sequential equilibrium for the receiver when the sender has transparent motives. It follows from this characterization that (a) the maximum utility that the receiver can get in equilibrium can be computed as a function of the prior by taking the maximum of finitely many linear functions, and (b) that the receiver cannot improve her utility by applying filters when the sender has transparent motives. We continue by showing that, in general (i.e., when the sender does not necessarily have transparent motives), the receiver can strictly benefit from filtering information. Moreover, if the receiver can filter, we provide an algorithm that outputs the best sequential equilibrium for the receiver when her action set is binary. Our algorithm runs in $O(k \log k)$ time, where k is the number of possible states. For our final result, we consider the same setting but with two senders instead, and reduce the problem of finding the best sequential equilibrium for the receiver to a linear programming instance with $O(k)$ variables

and $O(k)$ constraints. If there are three or more senders, it is easy to check that the receiver does not benefit from filtering information (see Section 4 for more details.)

A novelty in the second part of our analysis is giving the receiver the advantage of determining the information that the sender receives. Intuitively, we can see this setting as the *dual* of Bayesian persuasion [16]. More precisely, in Bayesian persuasion settings, the sender can freely choose the mechanism by which the information about the realized state is disclosed to the receiver. By contrast, in our setting, the receiver can freely choose the mechanism by which the information about the realized state is disclosed to the sender (note, however, that the sender has to disclose this information to the receiver via cheap talk, which means that the sender can still manipulate the information received if it is in her interests to do so.) This aspect is motivated by several scenarios, as for instance:

- Interacting with search engines like Google, Yahoo! or Bing. Here, the search engine plays the role of a sender that can access the necessary information about a given query on a large database and the user plays the role of a receiver that can impact the outcome of her interaction with the search engine by using specific keywords that influence how the information in the database is retrieved by the search engine.
- Scenarios in which the state of the game is a function of both the sender and the receiver’s information and the sender should propose an agreement to the receiver (as, for example, in big economic transactions or in diplomacy and politics). In these cases, the receiver should strategically decide how much information it should disclose to the sender in order to influence her proposal.
- Online platforms, such as Amazon, store a significant amount of information about their users. They have the ability to determine which information is disclosed to sellers, who, in turn, can utilize this data to target users through advertisements. Our model essentially examines the scenario in which the platform operates for the average user, revealing seller information to maximize the revenue for the average buyer in equilibrium.

Our analysis attempts to model these scenarios more accurately.

1.1 Related Literature

Our work on the setting with transparent motives is closely related to that of Lipnowski and Ravid [17], who studied a cheap talk model where the sender has transparent motives and showed that the best sender equilibrium is the quasi-concave envelope of the sender’s utility function. Our result is complementary since we give the geometric characterization of the best equilibrium for the receiver. In fact, we show that the best equilibrium for the receiver can be computed as the maximum of at most k^2 affine functions, where k is the number of states in the game.

Restricting the information available to the sender has been increasingly gaining traction in the community. Most notably, Bergemann, Brooks and Morris [6] considered a buyer-seller setting and studied possible ways to limit the seller’s information. In their work, they characterized the set of all pairs of possible utilities achievable in equilibrium. In particular, they provide an algorithm that achieves the optimal way to limit the seller’s information to maximize the expected buyers’ utility. Ichihashi [14] considered a Bayesian persuasion setting and studied how the outcome of the interaction is affected when the sender’s information is restricted. One of their results is that, if the receiver restricts sender information in a pre-play stage, the best utility that the receiver can get in this setting coincides with the one that the receiver would get in the “flipped game”, where the receiver persuades the sender. Other papers have studied similar questions in the context of common-interest coordination games [9, 11] and cheap talk games [15, 13]. The main differences

with our work are the fact that (a) we have no restrictions on the utility of the sender or that of the receiver, and (b) we focus on the best equilibrium for the receiver.

The rest of the paper is organized as follows. In Section 2 we introduce the cheap talk setting where the sender has transparent motives. We also provide a geometric characterization of the best equilibrium for the receiver. In Section 3 we introduce *filtered information aggregation games*, where the receiver can limit the information of the senders. In Section 4 we state our main results regarding the characterization of the best sequential equilibrium for the receiver in filtered information aggregation games, and these are proved in Sections 5, 6 and 7, where we also show how to extend the results of Section 2 to sequential equilibria. We end with a conclusion in Section 8.

2 Cheap Talk with Transparent Motives

2.1 Model

An *information transmission game* involves a *sender* s and a *receiver* r , and is defined by a tuple $\Gamma = (A, \Omega, p, M, u)$, where $A = \{a_1, \dots, a_\ell\}$ is the set of actions, $\Omega = \{\omega_1, \dots, \omega_n\}$ is the set of possible states, p is a commonly known prior distribution over Ω that assigns a strictly positive probability to each possible state, M is a finite set that contains the messages that the sender can send to the receiver (M is usually assumed to be the set of binary strings of length at most N , for some positive integer N), and $u : \{s, r\} \times \Omega \times A \rightarrow \mathbb{R}$ is a utility function such that $u(i, \omega, a)$ gives the utility of player i (where i is either the sender or the receiver) when action a is played at state ω . Each information transmission game instance is divided into three phases. In the first phase, a state ω is sampled from Ω following the prior distribution p , and this state is disclosed to the sender s . During the second phase, the sender sends a message $m \in M$ to the receiver, and in the third phase the receiver plays an action $a \in A$ and each player $i \in \{s, r\}$ receives $u(i, \omega, a)$ utility.

Given an information transmission game $\Gamma = (A, \Omega, p, M, u)$, a strategy profile $\vec{\sigma}$ for Γ consists of a pair of strategies (σ_s, σ_r) for the sender and the receiver, where σ_s is a map from Ω to distributions over M (which is denoted by $\Delta(M)$), and σ_r is a map from M to $\Delta(A)$. We say that $\vec{\sigma}$ is a Nash equilibrium if no player can increase its utility by defecting from $\vec{\sigma}$. More precisely, denote by $u_i(\vec{\sigma})$ the expected utility that i gets when players play $\vec{\sigma}$, then $\vec{\sigma}$ is a Nash equilibrium if

$$u_i(\vec{\sigma}_{-i}, \tau_i) \leq u_i(\vec{\sigma})$$

for all players $i \in \{s, r\}$ and all strategies τ_i for i .

For future reference, some of the results in this paper require a stronger notion of equilibrium known as *sequential equilibrium*. A sequential equilibrium consists of a strategy profile $\vec{\sigma}$ and a belief system β that (in this case) assigns a posterior distribution over Ω to each possible message in M . Intuitively, β represents the receiver's beliefs about the posterior distribution over Ω given the message that it receives from the sender. The pair $(\vec{\sigma}, \beta)$ form a sequential equilibrium if

- β is *consistent* with $\vec{\sigma}$, which means that there exists a sequence of completely mixed strategies $\vec{\sigma}^1, \vec{\sigma}^2, \dots$ that converge to $\vec{\sigma}$ such that the sequence of maps from messages to posterior distributions over Ω obtained by applying Bayes' rule assuming that the sender plays $\sigma_s^1, \sigma_s^2, \dots$ converges to β .
- σ_r is, according to β , a best-response for the sender for all messages $m \in M$.

In this section, we are interested in information transmission games in which the sender has *transparent motives*. This means that the sender's utility $u(s, \omega, a)$ only depends only on the action a played by the receiver, and not on the state ω . Thus, for simplicity, we will define functions $u_s : A \rightarrow \mathbb{R}$ and

$u_r : \Omega \times A \rightarrow \mathbb{R}$ that denote the utilities of the sender and the receiver respectively. More precisely, $u_s(a) := u(s, \omega, a)$ and $u_r(\omega, a) := u(r, \omega, a)$ for all $\omega \in \Omega$ and $a \in A$.

2.2 Characterizing the Best Equilibrium for the Receiver

In this section, we are interested in finding the best Nash equilibrium for the receiver given an information transmission game with transparent motives. We show later how to extend this equilibrium to a sequential equilibrium in Section 7.

We start by considering the case where the state space Ω is binary (i.e., $\Omega = \{\omega_0, \omega_1\}$). We can assume without loss of generality that there exist $\ell + 1$ real numbers $0 = b_1 < b_2 < \dots < b_{\ell+1} = 1$ such that action a_i is optimal for the receiver if and only if the posterior probability q of state ω_1 after her interaction with the sender is in the interval $[b_i, b_{i+1}]$. To see this, note that the expected utility of the receiver given some action a is a linear function over q that goes from $(0, u_r(\omega_0, a))$ to $(1, u_r(\omega_1, a))$. Therefore, the set of points in which an action a gives the maximal utility over all other actions is an interval in $[0, 1]$. If this interval is empty for some action a , it means that a will never be played under any circumstances since there will always be better options for the receiver. This implies that, without loss of generality, we can simply remove a from A and keep only those actions that are maximal at some point.

Using a slight abuse of notation, let $u_r(a, q)$ denote the expected utility of the receiver given action a and the posterior probability q of ω_1 . Note that $u_r(a, q)$ can be computed by the following expression:

$$u_r(a, q) := q \cdot u_r(a, \omega_1) + (1 - q) \cdot u_r(a, \omega_0).$$

For each $j \in \{1, 2, \dots, \ell + 1\}$, let $I(j) := [u_s(a_{j-1}), u_s(a_j)]$ (using the convention that $a_{-1} := a_0$ and $a_{\ell+1} := a_\ell$), and let C be the subset of pairs $(j, k) \in [l + 1]^2$ such that $I(j) \cap I(k) \neq \emptyset$. At a high level, $(j, k) \in C$ if the optimal actions with prior b_i and the optimal actions with prior b_j can be combined in such a way that the utility of the sender is equal in both combinations.

$$u_r(a, p) = pu_r(a, \omega_1) + (1 - p)u_r(a, \omega_0).$$

Let

$$\tilde{u}_r(q) := \max_{b_j \leq q \leq b_k \text{ s.t. } (j,k) \in C} (\alpha_{j,k} u_r(a_j, b_j) + (1 - \alpha_{j,k}) u_r(a_k, b_k)),$$


where $q \in [0, 1]$ and $\alpha_{j,k}$ is a value in $[0, 1]$ that satisfies that $\alpha_{j,k} b_j + (1 - \alpha_{j,k}) b_k = q$. The following theorem states that the best utility for the receiver in an information transmission game with transparent motives is given by \tilde{u}_r .

Theorem 1. *Given an information transmission game $\Gamma = (A, \Omega, p, M, u)$ with $|\Omega| = 2$ and transparent motives, the maximum utility that the receiver can get in a Nash equilibrium is given by $\tilde{u}_r(q)$, where q is the prior probability that the realized state is ω_1 .*


Before proving Theorem 1, we illustrate its use via the following simple example. Consider an information transmission game $\Gamma = (A, \Omega, p, M, u)$ with transparent motives such that $|\Omega| = 2$ and $A = \{a_1, a_2, a_3, a_4, a_5\}$, in which the utilities are defined as follows:

	$u_r(a, \omega_0)$	$u_r(a, \omega_1)$	$u_s(a)$
a_1	$\frac{1}{2}$	$-\frac{3}{2}$	0
a_2	$\frac{1}{3}$	$-\frac{2}{3}$	2
a_3	0	0	-2
a_4	$-\frac{2}{3}$	$\frac{1}{3}$	1
a_5	$\frac{1}{2}$	$-\frac{3}{2}$	-1

If we plot the utility of the receiver in a bubbling equilibrium in which no information is revealed we get the following figure, where the dotted line represents the utility of the receiver as a function of the prior probability that the realized state is ω_1 .




If we follow the construction of Theorem 1, we have that $b_1 = 0$, $b_2 = \frac{1}{6}$, $b_3 = \frac{1}{3}$, $b_4 = \frac{2}{3}$, $b_5 = \frac{5}{6}$ and $b_6 = 1$, and the set C consists of all pairs of the form $(j, j + 1)$ in addition to $(1, 3)$, $(2, 5)$, $(3, 5)$ and $(4, 6)$. Thus, it follows from Theorem 1 that the function \tilde{u}_r that represents the receiver's utility in a Nash equilibrium is obtained by taking the maximum of the bubbling equilibrium and the following four segments: the segment s_1 that connects $(b_1, \frac{1}{2})$ to $(b_3, 0)$, the segment s_2 that connects $(b_2, \frac{1}{6})$ to $(b_5, \frac{1}{6})$, the segment s_3 that connects $(b_3, 0)$ to $(b_5, \frac{1}{6})$, and the segment s_4 that connects $(b_4, 0)$ to $(b_6, \frac{1}{2})$, which are represented in the figure below with a solid green line:



Thus, the maximum utility that the receiver can get in a Nash equilibrium of Γ is represented by

the solid black line in the figure below:



Next, we proceed with the proof of Theorem 1.

Proof of Theorem 1. First we show that the receiver can guarantee at least $\tilde{u}_r(q)$ utility in a Nash equilibrium. Let q be the prior probability that the realized state is ω_1 and suppose that $\tilde{u}_r(q) = \alpha_{j,k}u_r(a_j, b_j) + (1 - \alpha_{j,k})u_r(a_k, b_k)$, for some $(j, k) \in C$. Since $(j, k) \in C$, it follows that $I(b_j) \cap I(b_k) \neq \emptyset$. Therefore, by definition, we can find $\beta_j, \beta_k \in [0, 1]$ such that

$$\beta_j u_s(a_{j-1}) + (1 - \beta_j)u_s(a_j) = \beta_k u_s(a_{k-1}) + (1 - \beta_k)u_s(a_k). \quad (1)$$

We next describe a cheap talk equilibrium that gives the receiver a utility of $\alpha_{j,k}u_r(a_j, b_j) + (1 - \alpha_{j,k})u_r(a_k, b_k)$. Using the fact that $\alpha_{j,k}b_j + (1 - \alpha_{j,k})b_k = q$, the splitting lemma of Aumann and Maschler [5] implies that there exists a communication protocol such that

- The sender sends a message m_1 with probability $\alpha_{j,k}$ and another message m_2 with probability $1 - \alpha_{j,k}$.
- The posterior probability of ω_1 after the receiver receives message m_1 is b_j and the posterior probability of ω_1 after the receiver receives message m_2 is b_k .

Consider such a communication protocol in which, in the last phase, the receiver plays action a_{j-1} with probability β_j and action a_j with probability $1 - \beta_j$ whenever she receives message m_1 . If she receives message m_2 instead, she plays a_{k-1} with probability β_k and a_k with probability $1 - \beta_k$. If the receiver receives a message that is not m_1 or m_2 , she plays the action that gives the least possible utility for the sender. By construction, since the posterior probability of ω_1 after receiving message m_1 is b_j , both a_{j-1} and a_j are best responses for the receiver. Similarly, since the posterior probability of ω_1 after receiving message m_2 is b_k , so are a_{k-1} and a_k in this scenario. Altogether, this implies that the receiver best-responds to the sender's strategy. To see that this strategy profile is also incentive-compatible for the sender note that, by Equation 1, the sender is indifferent between sending m_1 or m_2 . Moreover, sending any other message m would give her less utility than sending m_1 or m_2 . This shows that the receiver can indeed achieve $\tilde{u}_r(q)$ utility in equilibrium.

We next show that the receiver's utility is at most $\tilde{u}_r(q)$ in any Nash equilibrium. It follows from standard extreme point considerations that if there exists a Nash equilibrium $\vec{\sigma}$ in which the sender sends three or more messages with strictly positive probability, there exists another Nash equilibrium

$\vec{\sigma}'$ in which the sender sends at most two different messages and both the sender and the receiver at least as much utility with $\vec{\sigma}'$ than with $\vec{\sigma}$. Thus, for the rest of the proof, we'll restrict our analysis to communication protocols in which the sender sends at most two different messages with positive probability.

Suppose that the sender only sends one message with strictly positive probability. This means that the receiver plays her action with no additional information about the realized state, except for the prior p . Let q be the prior probability that the realized state is ω_1 and let $j \in [\ell]$ be an index such that $q \in [b_j, b_{j+1}]$. Since, in this case, the posterior distribution is the same as the prior distribution, a_j is the receiver's best response, and her expected utility is given by $(1-q) \cdot u_r(\omega_0, a_j) + q \cdot u_r(\omega_1, a_j)$. We can rewrite this as an affine combination of $u_r(a_j, b_j)$ and $u_r(a_j, b_{j+1})$ with the following expression:

$$u_r(a_j, q) = \alpha u_r(a_j, b_j) + (1 - \alpha) u_r(a_j, b_{j+1}),$$

where α is the value in $[0, 1]$ such that $\alpha b_j + (1 - \alpha) b_{j+1} = q$. This shows that, by definition, $\tilde{u}_r(q) \geq u_r(a_j, q)$, and therefore the receiver cannot get more utility than $\tilde{u}_r(q)$ in Nash equilibria where the sender sends only one message.

Assume instead that the sender sends two messages m_1 and m_2 such that the posterior probability that the realized state is ω_1 after the receiver receives m_1 and m_2 are x_1 and x_2 , respectively. We first claim that we can assume without loss of generality that $x_1 = b_j$ and $x_2 = b_k$ for some $j, k \in [\ell + 1]$. Suppose that $b_j < x_1 < b_{j+1}$ for some $j \in [\ell + 1]$. Then, it means that the best response for the receiver when receiving message m_1 is to play action a_j with probability 1. Then, we can construct another equilibrium $\vec{\sigma}$ in which the posterior probability after receiving message m_1 is b_j while the posterior probability after receiving message m_2 remains unchanged (note that the conditional probabilities of sending m_1 and m_2 are uniquely determined by Bayes' plausibility constraints). By definition, playing a_j is still a receiver's best response to m_1 in $\vec{\sigma}$. Moreover, it is straightforward to check that the sender receives the same utility in the new equilibrium and the original equilibrium. It remains to check that the receiver also (weakly) increases her utility. To see this, note that the posterior distribution of $\vec{\sigma}$ is a mean preserving spread of the original one. Therefore the new posterior distribution Blackwell dominates the original distribution and by Blackwell theorem [8] the receiver is (weakly) better. We can therefore assume from now on that $x_1 = b_j$ and $x_2 = b_k$ for some $j, k \in [\ell + 1]$.

We next show that the posterior probabilities b_j and b_k described in the paragraph above satisfy that $(j, k) \in C$. To see this, note that the sender's expected utility must be equal when sending m_1 and m_2 (otherwise, she can increase her utility by sending one of these messages by probability 1 and the other with probability 0). By definition, the only best responses for the receiver given posterior probabilities b_j and b_k are $\{a_{j-1}, a_j\}$, and $\{a_{k-1}, a_k\}$, respectively. Suppose that the sender plays a_{j-1} with probability α_j when receiving message m_1 , and plays a_{k-1} with probability α_k when receiving message m_2 . Since the sender is indifferent between m_1 and m_2 , α_j and α_k must satisfy that

$$\alpha_j u_s(a_{j-1}) + (1 - \alpha_j) u_s(a_j) = \alpha_k u_s(a_{k-1}) + (1 - \alpha_k) u_s(a_k),$$

which implies that $I(j) \cap I(k) \neq \emptyset$, and therefore that $(j, k) \in C$. This implies that we can represent the receiver's utility as an affine combination of $u_r(a_j, b_j)$ and $u_r(a_k, b_k)$ using the following expression:

$$u_r(\vec{\sigma}) = \alpha_{j,k} u_r(a_j, b_j) + (1 - \alpha_{j,k}) u_r(a_k, b_k),$$

where $\alpha_{j,k}$ is the unique value in $[0, 1]$ such that $\alpha_{j,k} b_j + (1 - \alpha_{j,k}) b_k = q$ (note that this equality must hold because of Bayes' plausibility constraints). This completes the proof of Theorem 1. \square

2.3 General State space

In this section we generalize the proof of Theorem 1 to arbitrary information transmission games with transparent motives. At a high level, both the statement and the proof are quite similar to the case with binary states, although they require additional notation.

Let $\Gamma = (A, \Omega, p, M, u)$ be an information transmission game with transparent motives where $\Omega = \{\omega_1, \dots, \omega_n\}$ and $A = \{a_1, \dots, a_\ell\}$. As in the proof of Theorem 1, we identify $\Delta(\Omega)$ with Δ^{n-1} (the $n - 1$ -dimensional simplex) and assume without loss of generality that there exists no strictly dominated action. More generally, we assume without loss of generality that all actions are optimal in *some* posterior distribution, which is that, for all actions $a \in A$, the set

$$\left\{ q \in \Delta^{n-1} \quad : \quad \sum_{i=1}^n q_i u_r(a, \omega_i) > \max_{a' \in A, a' \neq a} \sum_{i=1}^n q_i u_r(a', \omega_i) \right\}$$

is non-empty.

Let D_1, D_2, \dots, D_ℓ be the polytopes contained in Δ^{n-1} such that D_j contains the set of points $q \in \Delta^{n-1}$ for which action a_j is a best response given that the posterior distribution is q . Given a point $q \in \Delta^{n-1}$, let $A(q)$ be the set of indices $j \in [\ell]$ such that $v \in D_j$ and let $I(q) := [\min_{j \in A(q)} u_s(a_j), \max_{j \in A(q)} u_s(a_j)]$. Intuitively, $I(q)$ is the interval ranging from the minimum to the maximum utility that the sender can get with a receiver's best-response for q . Let V be the union of the set of vertices of D_1, D_2, \dots, D_ℓ (i.e., V is the set of points that are vertices of at least one of the polytopes), and let C be the set of tuples of points v_1, \dots, v_k of V such that $k \leq n$ and the intersection of all of the intervals $I(v_i)$ is non-empty. More precisely,

$$C := \left\{ (v_1, v_2, \dots, v_k) \in V^k : k \leq n, \bigcap_{i=1}^k I(v_i) \neq \emptyset \right\}.$$

Moreover, given any vertex $\vec{v} = (v_1, \dots, v_k) \in C$, let $H(\vec{v}) \subseteq \Delta^{n-1}$ be the *convex hull* of $\{v_1, \dots, v_k\}$ (i.e., the set of all the affine combinations of points in \vec{v}). Let

$$\tilde{u}_r(p) := \max \left\{ \sum_{i=1}^k \alpha_i u_r(a(v_i), v_i) \mid \vec{v} = (v_1, \dots, v_k) \in C, p \in H(\vec{v}) \right\},$$

where $a(v)$ is the action in $A(v)$ that has the lowest index and $(\alpha_i)_{i=1}^k \in \Delta^{k-1}$ are some affine parameters such that $\sum_{i=1}^k \alpha_i v_i = p$. The generalization of Theorem 1 is as follows.

Theorem 2. *Given an information transmission game $\Gamma = (A, \Omega, p, M, u)$ with transparent motives, the maximum utility that the receiver can get in a Nash equilibrium is given by $\tilde{u}_r(p)$.*

Proof. Showing that $\tilde{u}_r(p)$ is achievable by the receiver in a Nash equilibrium is analogous to the case with two states. However, the converse direction requires a slightly different reasoning. As in the case with binary states, for every Nash equilibrium in which the sender sends $n + 1$ or more messages with positive probability (where n is the number of states), there exists a Nash equilibrium in which the sender sends n or fewer messages and gives the same utility for the sender and the receiver.

Suppose that $\vec{\sigma}$ is a Nash equilibrium where the sender sends $k \leq n$ different messages m_1, \dots, m_k with positive probability. For each $i \in [k]$, denote by v_i the posterior distribution of possible states after the receiver receives message m_i when the sender plays σ_s . As in the case of $n = 2$, we claim that, without loss of generality, we can assume that $v_i \in V$ for every $i \in [k]$. More precisely, we show next that, for every equilibrium $\vec{\sigma}$, there exists possibly a different equilibrium $\vec{\sigma}'$ in which

each message m_i induces a posterior v'_i such that (a) $v'_i \in V$ for all $i \in [k]$, (b) the sender gets the same utility in $\vec{\sigma}$ and $\vec{\sigma}'$, and (c) the receiver gets a (weakly) higher utility in $\vec{\sigma}'$ than in $\vec{\sigma}$.

Assume that $v_i \notin V$ for some $i \in [k]$. We claim that we can find a distribution $\lambda \in \Delta(V)$ such that (i) $\sum_{v \in V} \lambda(v)v = v_i$ and (ii) if $a_j \in A$ is a best response for the receiver given a posterior distribution v_i , then a_j is also a best response for the receiver in every $v \in V$ such that $\lambda(v) > 0$. To see this, let $M_i \subseteq A$ be the set of actions that are a best response for the receiver in v_i and let $T_i := \cap_{a_j \in M_i} D_j$. Since T_i is a convex polytope that contains v_i , we can define λ as an affine combination of the vertices of T_i (where the coefficients represent the probabilities of each vertex). All such vertices satisfy the above (ii) property.

Assume that, in $\vec{\sigma}$, each message m_j is sent with total probability probability q_j . We define $\vec{\sigma}'$ as follows. The strategy σ'_r of the receiver is equal to σ_r . The sender does the following. For all $j \neq i$, the sender sends message m_j with total probability q_j (the same probability as in σ_s). However, instead of sending message m_i with probability q_i , in σ'_s the sender sends each vertex v of T_i with probability $\lambda(v) \cdot q_i$. By property (i), the strategy σ'_s is well defined and, by property (ii), strategy σ'_r is a best-response for the receiver to σ'_s .

By construction, the sender has identical utility in $\vec{\sigma}$ and in $\vec{\sigma}'$. To see that the receiver (weakly) increases her utility, note that the posterior distribution in $\vec{\sigma}'$ is a mean preserving spread of $\vec{\sigma}$ since we split m_i to the vertices of T_i . Repeating this process for every $m_i \notin V$ gives the desired equilibrium.

To finish the proof, note that we can express the receiver's expected utility as an affine combination of $\{u_r(a(m_i), m_i)\}_{i=1}^k$. Moreover, by Bayes' plausibility assumption, the affine coefficients α_i of the combination are uniquely determined by the identity $\sum_{i=1}^k \alpha_i m_i = p$. This shows that $u_r(\vec{\sigma}) \leq \tilde{u}_r(p)$, as desired.

□

2.4 Efficiently computing $\tilde{u}_r(q)$ for a Binary State Space

In this section we provide an algorithm that, given an information transmission game $\Gamma = (A, \Omega, p, M, u)$ with $|A| = n$, computes the best utility of the receiver in a Nash equilibrium and runs in $O(n \log n)$ time. By Theorem 1, we know that the best utility of the receiver in a Nash equilibrium of Γ is a piece-wise linear function. More precisely, there exists a positive integer $k \leq n$, a sequence of $k + 1$ real numbers $0 = x_1 < x_2 < \dots < x_{k+1} = 1$, and k linear functions f_1, \dots, f_k such that, if q is the prior probability that the realized state is ω_1 , then

$$\tilde{u}_r(q) = \begin{cases} f_1(q) & \text{if } q \in [x_1, x_2] \\ f_2(q) & \text{if } q \in [x_2, x_3] \\ \vdots & \vdots \\ f_k(q) & \text{if } q \in [x_k, x_{k+1}]. \end{cases}$$

Before describing the algorithm, we need the following additional notation. Given two pairs of indices $(i, j) \in C$ and $(i', j') \in C$, we say that (i, j) is *dominated* by (i', j') if $i \geq i'$ and $j \leq j'$. Intuitively, (i, j) is dominated by (i', j') if they both belong to C and the interval $[i', j']$ contains the interval $[i, j]$. It is easy to check that, if (i', j') dominates (i, j) , then, for each point (x, y) in the segment $s_{i,j}$ that goes from $(b_i, u_r(a_i))$ to $(b_j, u_r(a_j))$, there exists a point $(x, y') \in s_{i',j'}$ with $y' \geq y$ (this follows directly from the fact that the utility of the receiver is a convex function over the prior). This means that, in order to find the maximum of all the segments $s_{i,j}$ with $(i, j) \in C$, we can restrict the search to segments that are not dominated by other segments. In particular, if we denote by $C_{\max}(i)$ the maximum index j such that $(i, j) \in C$, we can restrict our search to all segments of the form $s_{i, C_{\max}(i)}$.

Suppose that we have an oracle that gives $C_{\max}(i)$ for each i in constant time. We show next how to construct an $O(n \log n)$ algorithm $\pi^{C_{\max}}$ that outputs k and the sequences x_1, \dots, x_{k+1} and f_1, \dots, f_k using C_{\max} as a black box. At a high level, the algorithm computes all the segments of the form $s_{i, C_{\max}(i)}$ and runs a variant of the Bentley-Ottmann algorithm in which, (a) we only output the segment intersections that involve the segment with the highest y -coordinate, and (b) whenever two segments s and s' intersect, it eliminates the segment with the lowest slope. This guarantees that the number of segment intersections processed by the algorithm is at most n , and therefore that the total complexity of the algorithm is $O(n \log n)$. More precisely, $\pi^{C_{\max}}$ consists of running a vertical sweep-line L that moves across the x -axis from left to right while keeping a list of all the segments that L intersects, sorted by the y -coordinate of their intersection with L . In order to do so, we need two data structures: a *binary search tree* T whose purpose is to keep track of all the segments that L intersects, sorted by y -coordinate, and a *priority queue* Q whose purpose is to keep a list of all the *events* that occur as L moves from left to right, sorted by x -coordinate. There are three types of events:

- **A segment s begins.** In this case, we insert s in T and query from T the segments s' and s'' that are immediately above and below s , respectively. If s intersects s' or s intersects s'' , we add these intersections to the event queue.
- **A segment s ends.** If s is still in T , we remove s from T . Let s' and s'' be the segments in T that are immediately above and below s . If s' and s'' intersect, we add the intersection to the event queue.
- **There is an intersection P between two segments s and s' .** If either s or s' has been previously deleted from T , this event is ignored. Otherwise, we erase from T the topmost segment between s and s' and proceed as if the erased segment ended (i.e., we check the segments that are immediately above and below of the erased segment and, if they intersect, we add the intersection). If the segment removed was also the topmost segment of all T , we append P to the output.

The full algorithm is as follows. We initialize an empty binary search tree T and an empty priority queue Q and push $2n$ events to Q indicating the beginning and end of segments $s_{1, C_{\max}(1)}, s_{2, C_{\max}(2)}, \dots, s_{n, C_{\max}(n)}$. We then process the events in Q as described above until Q becomes empty and output the sequence of points obtained this way, in addition to $(0, u_r(a_1))$ and $(1, u_r(a_n))$. The values x_i correspond to the x -coordinates of the points in the output, and the linear functions f_i are the ones that interpolate each pair of consecutive points.

The proof of correctness is equivalent to that of the Bentley-Ottmann algorithm. Note that the only difference between the algorithms is the way that we process the intersection events in $\pi^{C_{\max}}$. The Bentley-Ottmann algorithm outputs all segment intersections and, to do so, it swaps s and s' inside T whenever s and s' intersect (instead of processing them the way that we do). However, since a segment that is surpassed in y -coordinate by another segment can never be maximal again, we can optimize the algorithm by removing it right after processing its intersection. Since each intersection event erases one of the segments involved, we process at most n intersection events, which means that we process in total $3n$ events (n starting points, n endpoints, and n intersections). Since each event is processed in $O(\log n)$ time, the total time complexity is $O(n \log n)$. To see that the output of $\pi^{C_{\max}}$ is correct, note that the vertices that define the piece-wise linear function \tilde{u}_r are either (a) intersections of different segments $s_{i,j}$ or (b) points in which a segment $s_{i,j}$ ends and another segment $s_{i',j'}$ begins. If we use an implementation trick where we always process starting point events before endpoint events, the points defined in (b) would also be considered as intersections. This means that, using this implementation trick, the desired set of vertices is the subset of all segment intersections that occur at the top layer of the segments, which is precisely what $\pi^{C_{\max}}$ outputs.

As a final note regarding this part of the algorithm, note that, as in Theorem 1, we are assuming

that Γ is given in a form such that there exist $n + 1$ real numbers $0 = b_1 < b_2 < \dots < b_{n+1} = 1$ such that action a_i is the receiver's optimal action in the interval $[b_i, b_{i+1}]$. However, even if we can assume that Γ has this form without loss of generality, the input of the algorithm may not be reduced to this form yet. However, we can perform the same variant of the Bentley-Ottmann algorithm in order to reduce Γ to such a form, and to compute b_1, b_2, \dots, b_{n+1} . To see this, note that the linear function that outputs the utility that the receiver gets by playing $a_i \in A$ as a function of the prior is a segment going from $(0, u_r(0, a_i))$ to $(1, u_r(1, a_i))$. Reducing Γ to the desired form means precisely to compute the upper envelope of this set of segments.

Now, we have an $O(n \log n)$ algorithm $\pi^{C_{\max}}$ that, given C_{\max} as a black box, computes the best utility for the receiver. It remains to show how to compute $\{C_{\max}(1), \dots, C_{\max}(n)\}$ in $O(n \log n)$ time as well.

Efficient Computation of C_{\max}

By definition, $(j, k) \in C$ if $I(j) \cap I(k) \neq \emptyset$, where $I(j) = [u_s(a_{j-1}), u_s(a_j)]$. For simplicity, denote by s_j and e_j the minimum and maximum of $I(j)$, respectively. Consider the following algorithm ρ :

1. **Step 1:** Initialize an array C of n elements in which $C[i] = i$ for all $i \in [n]$.
2. **Step 2:** Initialize an empty array A . For each $j \in [n]$, append $(s_j, 0, j)$ and $(e_j, 1, j)$ to A .
3. **Step 3:** Sort A in lexicographical order (i.e., sort by the first coordinate, break ties with the second coordinate and then with the third one).
4. **Step 4:** Initialize an empty binary search tree T and two empty arrays Pos_1, Pos_2 of n elements each. Iterate through the elements $A[i] := (a, b, j)$ of A and do the following. If $b = 0$, set $Pos_1[j] = i$, insert j in T , and set $C[a] = T_{\max}$, where T_{\max} is the maximal element of T . Otherwise, set $Pos_2[j] = i$, and remove j from T .
5. **Step 5:** For each $i \in [n]$, set $C[i]$ to the maximum between $C[i]$ and the highest third coordinate between all the elements of A in the interval $[Pos_1[i], Pos_2[i]]$.

We claim that ρ outputs C_{\max} and can be implemented in $O(n \log n)$ time. To see that it gives the correct output, note that, given $k \in [n]$, the highest j such that $I(k) \cap I(j) \neq \emptyset$ is either (i) the highest j such that s_j or e_j appear between s_i and e_i (after sorting all the starting points and endpoints of each interval as in Step 3), or (ii) the highest j such that s_j appears before s_i and e_j does not appear before s_i . The values described in (i) and (ii) are calculated in step 4 and step 5, respectively. To see that ρ can be implemented in $O(n \log n)$ time, note that steps 1,2,3 and 4 can be implemented in $O(n \log n)$ time. A naive implementation of step 5 would be quadratic in n , but it can be performed in $O(n \log n)$ using a Segment Tree data structure.

3 Filtered Information Aggregation Games

3.1 Model and Analysis

In this section we analyze the effects of filtering on *information aggregation games*. At a high level, information aggregation games are information transmission games with multiple senders. For a rigorous definition, an information aggregation game consists of a receiver r , and a tuple $\Gamma = (S, A, \Omega, p, M, u)$ $S = \{1, 2, \dots, k\}$ is the set of senders and $A = \{a_1, \dots, a_\ell\}$, $\Omega = \{\omega_1, \dots, \omega_k\}$, $p \in \Delta(\Omega)$, M and u are defined as in Information Transmission games (see Section 2.1), with the only exception that u is a function from $S \cup \{r\} \times \Omega \times A$ to \mathbb{R} instead of a function from $\{s, r\} \times \Omega \times A$ to \mathbb{R} .

We are interested in a setting where the receiver can filter potential information from the senders. We model this as follows: at the beginning of the game, there is an additional phase (phase 0) where the receiver publicly chooses a function $X : \Omega \rightarrow \Delta(\{0, 1\}^*)$ that maps each state ω to a distribution over possible signals (we assume that signals are binary strings of arbitrary length). We call this function *filter*. The rest of the game follows exactly in the same way as in a (standard) information aggregation game, except that, if the realized state is ω , the senders receive a signal x sampled from $X(\omega)$ instead of ω itself. We call this a *filtered information aggregation game* (FIAG). More formally, a FIAG $\Gamma^F = (S, A, \Omega, p, M, u)$ runs as follows:

- **Phase 0:** The receiver chooses a filter $X : \Omega \rightarrow \Delta(\{0, 1\}^*)$. The filter is disclosed to the senders.
- **Phase 1:** A state $\omega \in \Omega$ is sampled according to the distribution p and is disclosed to the senders.
- **Phase 2:** Each sender $i \in S$ sends a message $m_i \in M$ to the receiver r .
- **Phase 3:** The receiver plays an action $a \in A$.
- **Phase 4:** Each player $i \in S \cup \{r\}$ receives $u(i, \omega, a)$ utility.

For an illustration of how the receiver may benefit from filtering, consider the following example.

Example 1. Let $\Gamma = (S, A, \Omega, p, M, u)$ be an information aggregation game with $S = \{s\}$, $A = \{0, 1\}$, $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, p is the uniform distribution over Ω and u is defined by the following table, where the first component of each cell is the utility of the sender and the second component is the utility of the receiver:

	$u(0, \omega)$	$u(1, \omega)$
ω_1	0, 0	1, 1
ω_2	1, 0	0, 1
ω_3	1, 1	0, 0
ω_4	0, 1	1, 0

We show next that, in Example 1, the receiver can double her maximal utility by filtering the sender's information.

Proposition 1. *If Γ is the game defined in Example 1,*

- (a) *The utility of the receiver in all sequential equilibria of Γ is $\frac{1}{2}$.*
- (b) *There exists a sequential equilibrium in Γ^F where the receiver ends up playing her preferred action in every single state.*

Proof. We claim that in a standard cheap talk equilibrium, the receiver utility is always $\frac{1}{2}$ and thus the receiver cannot get more than his utility at the bubbling equilibrium where information is not revealed. To see this consider without loss of generality an equilibrium with only two messages m_1, m_2 that are sent with positive probability. If, by way of contradiction, the equilibrium induces a utility that is strictly higher than $\frac{1}{2}$ to the receiver, then both actions are played with positive probability and the receiver is strictly better off playing one of the actions, say action 1, for at least one of the messages, say m_1 . This means that the sender is strictly better off sending m_1 at states $\Omega = \omega_1, \omega_4$ and message m_2 at states ω_2, ω_3 . But in this equilibrium, the receiver's utility is $\frac{1}{2}$. A contradiction.

In contrast, consider the following equilibrium in the filtered game. First, the receiver filters the sender's information by partitioning Ω as follows $\{\{\omega_1, \omega_2\}, \{\omega_3, \omega_4\}\}$ so that the sender only observes the chosen partition element. Then, at the second stage, the sender has two messages that are used to fully reveal his information to the receiver who plays action $a = 1$ if the chosen partition element is $\{\omega_1, \omega_2\}$ and action $a = 0$ if the chosen partition element is $\{\omega_3, \omega_4\}$. It is easy to see that the sender maximizes his utility given the realized partition element as he is indifferent among the two possible messages. The receiver in contrast receives his maximal utility under this equilibrium. \square

However, even if filtering can greatly improve the receiver's utility in in general, the following result complements Theorems 1 and 2 by stating that the receiver cannot improve her utility by filtering information if the sender has transparent motives.

Proposition 2. *Let $\Gamma^F = (S, A, \Omega, p, M, u)$ be a FIAG where the sender has transparent motives. Then, the receiver cannot increase her utility by filtering information.*

Proof. Given a sequential equilibrium $\vec{\sigma}$ for Γ^F , consider the following strategy profile $\vec{\sigma}'$ where the receiver does not filter. In Phase 0, the receiver selects the filter that simply forwards the realized state to the senders (i.e., the information that the senders receive is not filtered). In Phase 2, the sender first simulates what signal would she get with $\vec{\sigma}$ and what message m would she send to the receiver in $\vec{\sigma}$, and then she sends m to the receiver. In Phase 3, the receiver plays whatever it would have played in $\vec{\sigma}$ after receiving message m .

It is straightforward to check that $\vec{\sigma}'$ is also sequential equilibrium that gives each player the same utility as in $\vec{\sigma}$. \square

In the next section we give an efficient protocols that computes the optimal filters for the receiver in arbitrary games with binary actions and one or two senders.

4 Best Sequential Equilibria with Filters and Binary Actions

We next address the question of best receiver equilibrium from a computational aspect. Both of our results restrict attention to the case where the action set of the receiver is binary.

Theorem 3. *Let $\Gamma^F = (S, A, \Omega, p, u)$ be a filtered information aggregation game with $S = \{s\}$ and $A = \{0, 1\}$. Then, there exists an algorithm π that outputs the best sequential equilibrium for the receiver in Γ^F that runs in $O(k \log k)$ time, where $k = |\Omega|$.*

Theorem 4. *Let $\Gamma^F = (S, A, \Omega, p, u)$ be a filtered information aggregation game with $|\Omega| = k$, $S = \{s_1, s_2\}$, and $A = \{0, 1\}$. Then, finding the best sequential equilibrium for the receiver in Γ^F reduces to a linear programming instance with k variables and $2k + 2$ constraints.*

Note that we only provide results for the cases of one and two senders. If Γ^F is a filtered information aggregation game with $|S| \geq 3$ there is a trivial strategy profile $\vec{\sigma}_{MAX}$ that gives the receiver the maximal possible utility of the game. Using this strategy, the receiver puts no filters (more precisely, it selects the identity filter), the senders send the signal they get to the receiver, and then the receiver computes the signal x that was sent by the majority of the senders and plays the action that gives her the most utility conditioned that the senders got signal x . It is straightforward to check that this is indeed a Nash equilibrium: the senders do not get any additional utility by defecting since the other senders will still send the true signal (which means that the receiver will be able to compute the true signal as well), and the receiver does not get any additional utility either since she is always playing the optimal action for each possible state. In Section 7 we show that this strategy profile can be extended to a sequential equilibrium $(\vec{\sigma}_{MAX}^*, b)$ by defining an appropriate belief assessment b . This gives the following result:

Theorem 5. *Let $\Gamma^F = (S, A, \Omega, p, u)$ be a filtered information aggregation game with such that $|S| \geq 3$. Then, $\vec{\sigma}_{MAX}$ can be extended to a sequential equilibrium that gives the receiver the maximal possible utility of the game in Γ^F .*

In particular, we have the following Corollary.

Corollary 1. *If there are three or more senders, the receiver does not get any additional utility by filtering information.*

In the following sections, we give algorithms that output the best Nash equilibria for the settings described in Theorems 3, 1 and 4, respectively. Later, in section 7, we show how to extend these Nash equilibria to sequential equilibria. Section 7 also includes the missing details in the proof of Theorem 5.

5 Proof of Theorem 3

In this section we provide an algorithm that outputs the best Nash equilibrium for the receiver in any filtered information game $\Gamma^F = (S, A, \Omega, p, u)$ with $S = \{s\}$ and $A = \{0, 1\}$. In Section 7 we show how to extend the resulting Nash equilibrium to a sequential equilibrium.

Consider a game Γ_X^F that is identical to Γ^F except for the fact that the filter X is fixed beforehand (which means that the receiver doesn't get to choose the filter in Phase 0). The following proposition shows that there exists a Pareto-optimal Nash equilibrium with a relatively simple form.

Proposition 3. *There exists a Pareto-optimal Nash equilibrium $\vec{\sigma}$ in Γ_X^F such that, in $\vec{\sigma}$, either:*

- *The receiver always plays 0.*
- *The receiver always plays 1.*
- *The receiver always plays the best action for the sender.*

Before proving Proposition 3 we need additional notation. First, given a filtered information game $\Gamma_X^F = (S, A, \Omega, p, u)$, let $u_i(x, a)$ be the expected utility of player i on signal x and action a . This expected utility can be computed by the following equation:

$$u_i(x, a) = \sum_{\omega \in \Omega} \Pr[\omega | x] \cdot u_i(\omega, a),$$

where $\Pr[\omega | x]$ is the probability that the realized state is ω conditional on the fact that the sender received signal x . Similarly, $\Pr[\omega | x]$ has the following expression:

$$\Pr[\omega | x] = \frac{\Pr[\omega \leftarrow p, x \leftarrow X(\omega)]}{\sum_{\omega \in \Omega} \Pr[\omega \leftarrow p, x \leftarrow X(\omega)]}.$$

Let X_0 be the set of signals x in $\{X(\omega)\}_{\omega \in \Omega}$ such that $u_s(x, 0) > u_s(x, 1)$ (i.e., the set of signals in which the sender prefers 0), let X_1 be the set of signals such that $u_s(x, 0) < u_s(x, 1)$, and $X_=\$ be the set signals such that $u_s(x, 0) = u_s(x, 1)$. Given a strategy profile $\vec{\sigma}$ for Γ_X^F , denote by $\sigma_s(x, m)$ the probability that the sender sends message m given signal x and denote by $\sigma_r(a, m)$ the probability that the receiver plays action a given message m . Moreover, let $M_0^{\vec{\sigma}}$ denote the set of messages that have a strictly positive probability to be sent by the sender on at least one signal x in which the sender prefers 0 (i.e., $M_0^{\vec{\sigma}}$ is the set of messages m such that there exists $x \in X_0$ such that $u_s(x, m) > 0$). We define $M_1^{\vec{\sigma}}$ and $M_=\^{\vec{\sigma}}$ analogously.

With this notation, the following lemma describes all strategy profiles in Γ_X^F that are incentive-compatible for the sender.

Lemma 1. *A strategy profile $\vec{\sigma}$ for Γ_X^F is incentive-compatible for the sender if and only if the following is satisfied:*

- (a) *If $m \in M_0^{\vec{\sigma}}$, then $\sigma_r(0, m) \geq \sigma_r(0, m')$ for all messages m' .*
- (b) *If $m \in M_1^{\vec{\sigma}}$, then $\sigma_r(1, m) \geq \sigma_r(1, m')$ for all messages m' .*

Lemma 1 states that, for a strategy profile to be incentive-compatible for the sender, the receiver should play 0 with maximal probability on all messages that could be sent on signals in which the sender prefers 0, and the receiver should play 0 with minimal probability on all messages that could be sent on signals in which the sender prefers 1. In particular, we have the following Corollary.

Corollary 2. *A strategy profile $\vec{\sigma}$ for Γ_X^F is incentive-compatible for the sender if and only if there exist two real numbers $a, b \in [0, 1]$ with $a \geq b$ such that:*

- (a) $m \in M_0^{\vec{\sigma}} \implies \sigma_r(0, m) = a.$
- (b) $m \in M_1^{\vec{\sigma}} \implies \sigma_r(0, m) = b.$
- (c) $m \in M_{\pm}^{\vec{\sigma}} \implies b \leq \sigma_r(0, m) \leq a.$

Proof of Lemma 1. Clearly, if (a) and (b) are satisfied, then $\vec{\sigma}$ is incentive-compatible for the sender. Conversely, suppose that $\vec{\sigma}$ is incentive-compatible for the sender but it doesn't satisfy (a). This means that there exists a signal $x \in X_0$ and a message m that satisfies $u_s(x, m) > 0$ and such that $u_r(0, m) < u_r(0, m')$ for some other message m' . Therefore, if the sender sends m' instead of m whenever it receives signal x , it could increase its expected utility. This contradicts the fact that $\vec{\sigma}$ is incentive-compatible for the sender. The proof of the case in which $\vec{\sigma}$ doesn't satisfy (b) is analogous. \square

Corollary 2 characterizes the necessary and sufficient conditions for a strategy profile $\vec{\sigma}$ in Γ_X^F to be incentive-compatible for the sender. We show next that Proposition 3 follows from adding the receiver incentive-compatibility constraints into the mix. Denote by $u_r^{\vec{\sigma}}(m, 0)$ the receiver's expected utility when playing action 0 conditioned on the fact that it received message m and that the sender plays σ_s . Then, we have the following cases:

Case 1: There exists $m \in M_0^{\vec{\sigma}}$ such that $u_r^{\vec{\sigma}}(m, 0) < u_r^{\vec{\sigma}}(m, 1)$: Since $\vec{\sigma}$ is incentive-compatible for the receiver, it must be the case that $a = 0$ and therefore that $\sigma_r(0, m) = 0$ for all messages m .

Case 2: There exists $m \in M_1^{\vec{\sigma}}$ such that $u_r^{\vec{\sigma}}(m, 0) > u_r^{\vec{\sigma}}(m, 1)$: This time it must be the case that $b = 1$ and therefore that $\sigma_r(0, m) = 1$ for all messages m .

Case 3: $u_r^{\vec{\sigma}}(m, 0) \geq u_r^{\vec{\sigma}}(m, 1)$ for all $m \in M_0^{\vec{\sigma}}$ and $u_r^{\vec{\sigma}}(m, 0) \leq u_r^{\vec{\sigma}}(m, 1)$ for all $m \in M_1^{\vec{\sigma}}$: In this case, consider a strategy profile $\vec{\sigma}'$ that is identical to $\vec{\sigma}$ except that, whenever the receiver receives a message $m \in M_0^{\vec{\sigma}}$, it plays action 0 with probability 1 (as opposed to probability a), and when the receiver receives a message $m \in M_1^{\vec{\sigma}}$, it plays action 1 with probability 1 (as opposed to $1 - b$). It is easy to check that, by construction, $\vec{\sigma}'$ Pareto-dominates $\vec{\sigma}$. Even so, $\vec{\sigma}'$ can be further improved: consider a strategy profile $\vec{\sigma}^s$ such that the sender sends message 0 in all subsets $C \in X_0$ and in all subsets $C \in X_{\pm}$ such that $u_r(C, 0) > u_r(C, 1)$, and the sender sends message 1 in all other subsets. Additionally, the receiver plays action 0 when receiving message 0 and plays action 1 when receiving message 1. It is easy to check that $\vec{\sigma}^s$ satisfies that $u_r^{\vec{\sigma}^s}(0, 0) \geq u_r^{\vec{\sigma}^s}(0, 1)$ and $u_r^{\vec{\sigma}^s}(1, 0) \leq u_r^{\vec{\sigma}^s}(1, 1)$. Therefore, $\vec{\sigma}^s$ is a Nash equilibrium of Γ_X^E . Moreover, $\vec{\sigma}^s$ Pareto-dominates $\vec{\sigma}'$. To check this, note that both $\vec{\sigma}^r$ and $\vec{\sigma}'$ play the best action for the sender at every possible signal x . The only difference

is that, in signals in which the sender is indifferent, the receiver plays the best action for herself in $\vec{\sigma}^s$ while it doesn't necessarily do so in $\vec{\sigma}'$.

This analysis provides a refinement of Proposition 3. In fact, consider the following two strategy profiles:

- Strategy $\vec{\sigma}'$: Regardless of the signal received, the sender sends an empty \perp message. The receiver plays the action that gives her the most utility with no information.
- Strategy $\vec{\sigma}^s$: After receiving the signal, the sender sends its preferred action. The receiver plays the action sent by the sender.

We have the following characterization of Pareto-optimal Nash equilibria:

Proposition 4 (Refinement of Proposition 3). *If $\vec{\sigma}^s$ is incentive-compatible for the receiver, then $\vec{\sigma}^s$ is a Pareto-optimal Nash equilibrium of Γ_X^F . Otherwise, $\vec{\sigma}'$ is a Pareto-optimal Nash equilibrium of Γ_F^X .*

Proposition 4 shows that there is a simple algorithm that outputs a Pareto-optimal Nash equilibrium of Γ_X^F once the filter X has been fixed beforehand. This reduces the problem of finding the best Nash equilibrium of Γ^F for the receiver to that of finding the filter X^{OPT} such that the Pareto-optimal Nash equilibrium of Γ_X^F gives the best possible utility for the receiver. In fact, the best Nash equilibrium $\vec{\sigma}^{OPT}$ goes as follows:

1. The receiver chooses the filter X^{OPT} such that $\vec{\sigma}^s$ is incentive-compatible for the receiver and gives the receiver the maximum utility in Γ_X^F . If there is no filter such filter X^{OPT} , the receiver chooses an arbitrary filter.
2. Given the chosen filter X and the signal x , the sender computes the Pareto-optimal Nash equilibrium $\vec{\sigma}$ of Γ_X^F and plays σ_s with signal x .
3. If the receiver sent an arbitrary filter in step 1, it plays the best action assuming no information. Otherwise, if the receiver receives 0, it plays action 0. If it receives 1, it plays action 1. If it receives something else it plays the best action assuming no information.

The following proposition shows that $\vec{\sigma}^{OPT}$ satisfies the desired properties.

Proposition 5. *$\vec{\sigma}^{OPT}$ is the Nash equilibrium of Γ^F that gives the most utility to the receiver.*

Proof. Once the sender chooses a filter X , the receiver and the sender both maximize their utilities if they play the Pareto-optimal strategy profile described in Proposition 4. Therefore, it is optimal for the receiver to choose X^{OPT} . \square

5.1 Computation of X^{OPT}

In this section we show how to compute the optimal filter X^{OPT} for the receiver. Our aim is to find a filter X^{OPT} such that $\vec{\sigma}^s$ is incentive-compatible for the receiver in $\Gamma_{X^{OPT}}^F$ and such that the expected utility for the receiver with $\vec{\sigma}^s$ is as large as possible. We begin by showing an algorithm that runs in $O(k \log k)$ time, where k is the number of states (i.e., the size of Ω), and then we show the proof of correctness for the algorithm provided.

5.1.1 An $O(k \log k)$ Algorithm

For our algorithm, we restrict our search to *binary* filters (i.e., filters that only send signals in $\{0, 1\}$) and later, in Lemma 2, we show that this is done without loss of generality. With this restriction, we can describe possible candidates X by a function X_* that maps each state ω to the probability that $X(\omega)$ outputs 0. For simplicity, we'll also say that X is incentive-compatible for the sender (resp., for the receiver) if the strategy profile in which the sender forwards the signal to the receiver and the receiver plays whatever it gets from the sender is incentive-compatible for the sender (resp., for the receiver). We will also say that the sender (resp., the receiver) gets a expected utility in X , if a is the expected utility it gets when running the previous strategy in Γ_X^F . Note that $\bar{\sigma}^s$ is always incentive-compatible for the sender (for all filters X), but a binary filter X might not always be since we are forcing the sender to send the signal received instead of its true preference. However, if $\bar{\sigma}^s$ is incentive-compatible for the sender (resp., for the receiver) in some game Γ_X^F for some binary filter X , and the sender does not always have the same preference in Γ_X^F , then either X or its *complementary* (i.e., the filter X' that sends 0 whenever X sends 1, and sends 1 whenever X sends 0) is incentive-compatible for the sender (resp., the receiver) and gives the same utility as $\bar{\sigma}^s$ in Γ_X^F . This shows that, without loss of generality, we can simply search for the binary filter that (a) is incentive-compatible for both the sender and the receiver, and (b) that gives the most utility for the receiver.

Before we start, let Ω^0 (resp., Ω^1) denote the set of states in which both the sender and the receiver prefer 0 and let Ω_0^1 (resp., Ω_1^0) denote the set of states in which the sender strictly prefers 0 and the receiver strictly prefers 1 (resp., the sender strictly prefers 1 and the receiver strictly prefers 0). The following algorithm outputs the optimal filter X^{OPT} for the receiver.

- Step 1 Set $X_*^{OPT}(\omega) = 1$ for all $\omega \in \Omega^0$.
- Step 2 Set $X_*^{OPT}(\omega) = 0$ for all $\omega \in \Omega^1$.
- Step 3 Sort all states $\omega \in \Omega_0^1 \cup \Omega_1^0$ according to the value $\frac{u_r(\omega, 0) - u_r(\omega, 1)}{u_s(\omega, 1) - u_s(\omega, 0)}$. Let $\omega^1, \omega^2, \dots, \omega^{k'}$ be the resulting list.
- Step 4 Set $X_*^{OPT}(\omega) = 1$ for all $\omega \in \Omega_1^0$ and $X_*^{OPT}(\omega) = 0$ for all $\omega \in \Omega_0^1$. If the resulting filter is incentive-compatible for the sender, output X^{OPT} .
- Step 5 For $i = 1, 2, \dots, k'$ do the following. If $X_*^{OPT}(\omega^i)$ is set to 1 (resp., to 0), set it to 0 (resp., to 1). If the resulting filter is incentive-compatible for the sender, find the maximum (resp., the minimum) q such that setting $X_*^{OPT}(\omega^i)$ to q is incentive-compatible for the sender (we show below that it can be computed with a linear equation with amortized cost). If the resulting filter is incentive-compatible for the receiver, output X^{OPT} , otherwise output the constant filter $X_*^{OPT} \equiv 0$.

It is important to note that the algorithm always terminates, since if we get to k' in step 5, the resulting filter will always output the signal that the sender prefers. It remains to show how to compute q in amortized linear time. For this purpose, let d_i^s (resp., d_i^r) denote the difference in utility for the sender (resp., for the receiver) between playing 0 and playing 1. More precisely,

$$\begin{aligned} d_i^s &:= u_s(\omega_i, 0) - u_s(\omega_i, 1) \\ d_i^r &:= u_r(\omega_i, 0) - u_r(\omega_i, 1). \end{aligned}$$

Then, given a binary filter X , $\bar{\sigma}^s$ is incentive-compatible for the sender in Γ_X^F if and only if

$$\begin{aligned} \sum_i p(\omega_i) \cdot d_i^s \cdot X_*(\omega_i) &\geq 0 \\ \sum_i p(\omega_i) \cdot d_i^s \cdot (1 - X_*(\omega_i)) &\leq 0. \end{aligned} \tag{2}$$

The first and second equation state that the sender prefers 0 on signal 0 and prefers 1 on signal 1, respectively. Similarly, $\vec{\sigma}^s$ is incentive-compatible for the receiver in Γ_X^F if and only if

$$\begin{aligned} \sum_i p(\omega_i) \cdot d_i^r \cdot X_*(\omega_i) &\geq 0 \\ \sum_i p(\omega_i) \cdot d_i^r \cdot (1 - X_*(\omega_i)) &\leq 0. \end{aligned} \quad (3)$$

It is straightforward to check that the incentive-compatibility equations for the sender monotonically increase and decrease, respectively, whenever we iterate in the fifth step of the algorithm. Therefore, we can simply pre-compute the sums

$$\begin{aligned} Y_0^s &:= \sum_{\omega \in \Omega_0^0} p(\omega) \cdot (u_s(\omega, 0) - u_s(\omega, 1)) \\ Y_1^s &:= \sum_{\omega \in \Omega_1^1} p(\omega) \cdot (u_s(\omega, 0) - u_s(\omega, 1)) \\ W_{i,0}^{(s,s)} &:= \sum_{j < i, \omega^j \in \Omega_0^1} p(\omega^j) \cdot (u_s(\omega^j, 0) - u_s(\omega^j, 1)) \\ W_{i,1}^{(s,s)} &:= \sum_{j < i, \omega^j \in \Omega_1^0} p(\omega^j) \cdot (u_s(\omega^j, 0) - u_s(\omega^j, 1)) \\ W_{i,0}^{(s,r)} &:= \sum_{j > i, \omega^j \in \Omega_0^1} p(\omega^j) \cdot (u_s(\omega^j, 0) - u_s(\omega^j, 1)) \\ W_{i,1}^{(s,r)} &:= \sum_{j > i, \omega^j \in \Omega_1^0} p(\omega^j) \cdot (u_s(\omega^j, 0) - u_s(\omega^j, 1)) \end{aligned}$$

and analogues for the receiver:

$$\begin{aligned} Y_0^r &:= \sum_{\omega \in \Omega_0^0} p(\omega) \cdot (u_r(\omega, 0) - u_r(\omega, 1)) \\ Y_1^r &:= \sum_{\omega \in \Omega_1^1} p(\omega) \cdot (u_r(\omega, 0) - u_r(\omega, 1)) \\ W_{i,0}^{(r,s)} &:= \sum_{j < i, \omega^j \in \Omega_0^1} p(\omega^j) \cdot (u_r(\omega^j, 0) - u_r(\omega^j, 1)) \\ W_{i,1}^{(r,s)} &:= \sum_{j < i, \omega^j \in \Omega_1^0} p(\omega^j) \cdot (u_r(\omega^j, 0) - u_r(\omega^j, 1)) \\ W_{i,0}^{(r,r)} &:= \sum_{j > i, \omega^j \in \Omega_0^1} p(\omega^j) \cdot (u_r(\omega^j, 0) - u_r(\omega^j, 1)) \\ W_{i,1}^{(r,r)} &:= \sum_{j > i, \omega^j \in \Omega_1^0} p(\omega^j) \cdot (u_r(\omega^j, 0) - u_r(\omega^j, 1)). \end{aligned}$$

Note that these sums can be computed in amortized linear time over the total number of states. Once these sums are computed, if we are in the i th iteration of Step 5, to find q we can check in constant time if the solution of any of the following linear equations is in $[0, 1]$:

$$\begin{aligned} Y_0^s + W_{i,0}^{(s,s)} + p(\omega^i) \cdot q \cdot u_s(\omega, 0) - u_s(\omega, 1) + W_{i,0}^{(s,r)} &= 0 \\ Y_1^s + W_{i,1}^{(s,s)} + p(\omega^i) \cdot (1 - q) \cdot u_s(\omega, 0) - u_s(\omega, 1) + W_{i,1}^{(s,r)} &= 0. \end{aligned}$$

If there exists such solution q , to check that it is incentive-compatible for both players, we should simply check if the following inequalities are satisfied:

$$\begin{aligned} Y_0^s + W_{i,0}^{(s,s)} + p(\omega^i) \cdot q \cdot u_s(\omega, 0) - u_s(\omega, 1) + W_{i,0}^{(s,r)} &\geq 0 \\ Y_1^s + W_{i,1}^{(s,s)} + p(\omega^i) \cdot (1 - q) \cdot u_s(\omega, 0) - u_s(\omega, 1) + W_{i,1}^{(s,r)} &\leq 0. \\ Y_0^r + W_{i,0}^{(r,s)} + p(\omega^i) \cdot q \cdot u_r(\omega, 0) - u_r(\omega, 1) + W_{i,0}^{(r,r)} &\geq 0 \\ Y_1^r + W_{i,1}^{(r,s)} + p(\omega^i) \cdot (1 - q) \cdot u_r(\omega, 0) - u_r(\omega, 1) + W_{i,1}^{(r,r)} &\leq 0. \end{aligned}$$

All of these computations can be performed in constant time if the necessary sums are pre-computed. This completed the description of the algorithm, note that it runs in $O(k \log k)$ time since it takes $O(k \log k)$ operations to sort the elements of $\Omega_0^1 \cup \Omega_1^0$, $O(k)$ operations to compute the sums, and $O(1)$ operations in each iteration of Step 5. In the next section, we show that the algorithm's output is correct.

5.1.2 Proof of Correctness

In this section, we show that the Algorithm presented in Section 5.1 is correct. We begin by showing that we can indeed restrict our search to binary filters filters.

Lemma 2. *Let X be a filter such that $\bar{\sigma}^s$ is incentive-compatible for the receiver in Γ_X^F . Then, there exists a binary filter X' such that*

- (a) $\bar{\sigma}^s$ is incentive-compatible for the receiver in $\Gamma_{X'}^F$,
- (b) With strategy profile $\bar{\sigma}^s$, the expected utility of the receiver in Γ_X^F and $\Gamma_{X'}^F$ is identical.

Proof. Recall that $\bar{\sigma}^s$ is a strategy profile in which, for each possible signal x , the sender sends its preferred action and then the receiver plays whatever is sent by the sender. This means that, if $\bar{\sigma}$ is incentive-compatible for the receiver in Γ_X^F , then we can merge all signals in which the sender prefers 0, and also merge all signals in which the sender prefers 1. More precisely, given filter X , let X_0 and X_1 be the sets of signals in which the sender prefers 0 and 1 respectively. Consider a filter X' that sends signal 0 whenever X would send a signal in X_0 , and sends signal 1 whenever X would send a signal in X_1 . In $\Gamma_{X'}^F$, by construction, both the sender and the receiver prefer action 0 on signal 0 and action 1 on signal 1. This means that $\bar{\sigma}^s$ is a Nash equilibrium of $\Gamma_{X'}^F$. Moreover, again by construction, the expected utility of the receiver (and the sender) when playing $\bar{\sigma}^s$ in $\Gamma_{X'}^F$ is identical to the one they'd get in Γ_X^F . \square

Recall that binary filters X can be described by a function $X_* : \Omega \rightarrow [0, 1]$ that maps each state $\omega \in \Omega$ to the probability that $X(\omega)$ assigns to 0. Because of Lemma 1, our aim is to find which values in $[0, 1]$ we should assign to each element in Ω . The following lemmas characterizes these values.

Lemma 3. *There exists a filter X^{OPT} that maximizes the utility for the receiver such that*

- (a) $X_*^{OPT}(\omega) = 1$ for all $\omega \in \Omega_0^0$.
- (b) $X_*^{OPT}(\omega) = 0$ for all $\omega \in \Omega_1^1$.
- (c) At most one state $\omega \in \Omega_0^1 \cup \Omega_1^0$ satisfies that $X_*^{OPT}(\omega) \notin \{0, 1\}$.

Proof. Given a binary filter X , if we increase $X_*(\omega_i)$ by ϵ , the expected utility of the receiver increases by $\epsilon p(\omega_i) d_i^r$. This means that, if $d_i^s \geq 0$ and $d_i^r \geq 0$, setting $X_*(\omega_i)$ to 1 both increases the receiver's utility and preserves the incentive-compatibility constraints (see Equations 2 and 3). Analogously, the same happens by setting $X_*(\omega_i)$ to 0 when $d_i^s \leq 0$ and $d_i^r \leq 0$. This proves (a) and (b).

To prove (c) suppose that there exist two states ω_i and ω_j such that $X_*^{OPT}(\omega_i), X_*^{OPT}(\omega_j) \notin \{0, 1\}$. Then, because of (a) and (b) we can assume without loss of generality that $\omega_i, \omega_j \in \Omega_0^1 \cup \Omega_1^0$. Therefore, if we increase $X_*(\omega_i)$ by ϵd_j^s and decrease $X_*(\omega_j)$ by ϵd_i^s , we'd have that the sender's utility and incentive-compatibility constraints remain unchanged, but the receiver's utility increases by $\epsilon(d_j^s d_i^r - d_i^s d_j^r)$ (note that this value can be negative). This means that if we choose an ϵ that is small enough and is of the same sign as $d_j^s d_i^r - d_i^s d_j^r$, not only the expected utility of the receiver increases, but also the values $\sum_i p(\omega_i) \cdot d_i^r \cdot X_*(\omega_i)$ and $\sum_i p(\omega_i) \cdot d_i^r \cdot (1 - X_*(\omega_i))$ increase and decrease respectively. This contradicts the fact that the filter X^{OPT} is optimal for the receiver. \square

Lemma 3 shows that we can assign $X_*^{OPT}(\omega) = 1$ for all $\omega \in \Omega_0^0$, $X_*^{OPT}(\omega) = 0$ for all $\omega \in \Omega_0^0$, and we have to assign either probability 1 or probability 0 to all but at most one of the remaining states. Ideally, we would like to assign probability 1 to all states in Ω_0^1 and probability 0 to all states in

Ω_1^0 , since this guarantees that the receiver gets the maximum possible utility. However, this may not always be incentive-compatible for the sender, which means that we may have to assign to some of the states the probability that the sender prefers, as opposed to the probability that the receiver prefers. The following lemma characterizes these states.

Lemma 4. *Given a filtered information game $\Gamma^F = (S, A, \Omega, p, u)$, let $\omega^1, \omega^2, \dots, \omega^\ell$ be the states in $\Omega_0^1 \cup \Omega_1^0$ sorted by $\frac{u_r(\omega, 0) - u_r(\omega, 1)}{u_s(\omega, 1) - u_s(\omega, 0)}$. Then, there exists a binary filter X^{OPT} that is optimal for the receiver and a value $\ell' \leq \ell$ such that:*

- (a) $X_*^{OPT}(\omega) = 1$ for all $\omega \in \Omega_0^0$.
- (b) $X_*^{OPT}(\omega) = 0$ for all $\omega \in \Omega_1^1$.
- (c) For $i < \ell'$, $X_*^{OPT}(\omega^i) = 0$ if $\omega \in \Omega_1^0$ and $X_*^{OPT}(\omega^i) = 1$ if $\omega \in \Omega_0^1$.
- (d) For $i > \ell'$, $X_*^{OPT}(\omega^i) = 0$ if $\omega \in \Omega_0^1$ and $X_*^{OPT}(\omega^i) = 1$ if $\omega \in \Omega_1^0$.

Lemma 4 states that, if we sort the states in which the sender and the receiver have different preferences by the ratio between how much the receiver prefers 0 with respect to 1 and how much the sender prefers 1 with respect to 0 (note that these ratios are always positive), we get that we can split these states into two contiguous blocks in which, in the first block, we assign them the probability that is most convenient for the sender, and for the second block we assign them the probability that is most convenient for the receiver. The proof follows the lines of that of Lemma 3 part (c).

Proof of Lemma 4. Given two states $\omega, \omega' \in \Omega_0^1 \cup \Omega_1^0$, we say that $\omega < \omega'$ if $\frac{u_r(\omega, 0) - u_r(\omega, 1)}{u_s(\omega, 1) - u_s(\omega, 0)} < \frac{u_r(\omega', 0) - u_r(\omega', 1)}{u_s(\omega', 1) - u_s(\omega', 0)}$. Using the same argument as in the Proof of Lemma 3, we can assume w.l.o.g. that (a) and (b) are satisfied. Thus, it only remains to show (c) and (d), which follow from the fact that there exists an optimal binary filter X^{OPT} for the receiver such that:

- (s1) If $\omega, \omega' \in \Omega_1^0$, $\omega < \omega'$ and $X_*^{OPT}(\omega) > 0$, then $X_*^{OPT}(\omega') = 1$.
- (s2) If $\omega, \omega' \in \Omega_0^1$, $\omega < \omega'$ and $X_*^{OPT}(\omega) < 1$, then $X_*^{OPT}(\omega') = 0$.
- (s3) If $\omega \in \Omega_1^0$, $\omega' \in \Omega_0^1$, $\omega < \omega'$ and $X_*^{OPT}(\omega) > 0$, then $X_*^{OPT}(\omega') = 0$.
- (s4) If $\omega \in \Omega_0^1$, $\omega' \in \Omega_1^0$, $\omega < \omega'$ and $X_*^{OPT}(\omega) < 1$, then $X_*^{OPT}(\omega') = 1$.

We will prove (s1) and (s3), the proofs of (s2) and (s4) are analogous.

Proof of (s1): Suppose that there exist $\omega, \omega' \in \Omega_1^0$ such that $\omega < \omega'$ and $X_*^{OPT}(\omega) > 0$, but $X_*^{OPT}(\omega') < 1$. If this happens, we can set

$$\begin{aligned} X_*^{OPT}(\omega) &\leftarrow X_*^{OPT}(\omega) + \epsilon(u_s(\omega', 0) - u_s(\omega', 1)) \\ X_*^{OPT}(\omega') &\leftarrow X_*^{OPT}(\omega') + \epsilon(u_s(\omega, 1) - u_s(\omega, 0)). \end{aligned}$$

Since $u_s(\omega', 0) - u_s(\omega', 1) < 0$ and $u_s(\omega, 0) - u_s(\omega, 1) < 0$, there exists some $\epsilon > 0$ that is small enough so that $X_*^{OPT}(\omega)$ and $X_*^{OPT}(\omega')$ stay between 0 and 1. Moreover, as in the proof of Lemma 3 part (c), if we perform this change the sender's utility remains unchanged, but the receiver's utility increases by

$$\Delta := \epsilon((u_s(\omega', 0) - u_s(\omega', 1))(u_r(\omega, 0) - u_r(\omega, 1)) + (u_s(\omega, 1) - u_s(\omega, 0))(u_r(\omega', 0) - u_r(\omega', 1))).$$

By assumption, we have that

$$\frac{u_r(\omega, 0) - u_r(\omega, 1)}{u_s(\omega, 1) - u_s(\omega, 0)} < \frac{u_r(\omega', 0) - u_r(\omega', 1)}{u_s(\omega', 1) - u_s(\omega', 0)}$$

which, together with the fact that $u_s(\omega, 1) - u_s(\omega, 0) > 0$ and $u_s(\omega', 1) - u_s(\omega', 0) > 0$, implies that $\Delta \geq 0$ whenever $\epsilon > 0$.

Proof of (s3): The proof is almost identical to that of (s1). Suppose that there exist $\omega \in \Omega_1^0$, $\omega' \in \Omega_0^1$ such that $\omega < \omega'$, $X_*^{OPT}(\omega) > 0$ and $X_*^{OPT}(\omega') > 0$. In this case, we can again set

$$\begin{aligned} X_*^{OPT}(\omega) &\leftarrow X_*^{OPT}(\omega) + \epsilon(u_s(\omega', 1) - u_s(\omega', 0)) \\ X_*^{OPT}(\omega') &\leftarrow X_*^{OPT}(\omega') + \epsilon(u_s(\omega, 0) - u_s(\omega, 1)). \end{aligned}$$

The only difference with the proof of (s1) is that, in this case, $u_s(\omega, 0) - u_s(\omega, 1) < 0$ and $u_s(\omega', 0) - u_s(\omega', 1) > 0$. Again, there exists a sufficiently small $\epsilon > 0$ such that $X_*^{OPT}(\omega)$ and $X_*^{OPT}(\omega')$ remain between 0 and 1. Moreover, a straightforward computation shows that the sender's utility remains unchanged, but that the receiver's utility changes by

$$\Delta := \epsilon((u_s(\omega', 1) - u_s(\omega', 0))(u_r(\omega, 0) - u_r(\omega, 1)) + (u_s(\omega, 0) - u_s(\omega, 1))(u_r(\omega', 0) - u_r(\omega', 1))).$$

Using that

$$\frac{u_r(\omega, 0) - u_r(\omega, 1)}{u_s(\omega, 1) - u_s(\omega, 0)} < \frac{u_r(\omega', 1) - u_r(\omega', 0)}{u_s(\omega', 0) - u_s(\omega', 1)},$$

$u_s(\omega, 1) - u_s(\omega, 0) > 0$, and $u_s(\omega', 0) - u_s(\omega', 1) > 0$, we get that $\Delta \geq 0$ whenever $\epsilon > 0$. \square

Lemma 4 shows that the algorithm provided in Section 5.1 outputs the correct solution as long as it gets the right value of q . The next lemma shows that this value is precisely the one that the algorithm finds. This completes the proof of correctness.

Lemma 5. *Let X^r be the filter that assigns $X_*^r(\omega) = 1$ for $\omega \in \Omega_0^0 \cup \Omega_1^0$ and $X_*^r(\omega) = 0$ for $\omega \in \Omega_1^1 \cup \Omega_0^1$. If X^r is incentive-compatible for the sender, X^r is the optimal filter for the receiver that is incentive-compatible for both players. Otherwise, all filters X that are incentive-compatible for both players and are optimal for the receiver satisfy at least one of the following equations:*

$$\begin{aligned} \sum_i p(\omega_i) \cdot d_i^s \cdot X_*(\omega_i) &= 0 \\ \sum_i p(\omega_i) \cdot d_i^s \cdot (1 - X_*(\omega_i)) &= 0. \end{aligned}$$

Proof. Filter X^r gives the maximum utility to the sender of all possible filters. Therefore, if it is incentive-compatible for the sender, it is the optimal for the receiver. Suppose instead that X^r is not incentive-compatible for the sender but that an optimal filter X^{OPT} satisfies

$$\begin{aligned} \sum_i p(\omega_i) \cdot d_i^s \cdot X_*(\omega_i) &> 0 \\ \sum_i p(\omega_i) \cdot d_i^s \cdot (1 - X_*(\omega_i)) &< 0. \end{aligned} \tag{4}$$

Since X^r is not IC for the sender, there exists a state $\omega \in \Omega_{0,1}$ such that $X_*^{OPT} > 0$ or a state $\omega \in \Omega_{1,0}$ such that $X_*^{OPT} < 1$. In the first case, we can decrease X_*^{OPT} by a small value $\epsilon > 0$ such that Equation 4 is still satisfied. By doing this, we increase the receiver's expected utility while obtaining a new filter that is still incentive-compatible for both players. The latter case is analogous except that we increase X_*^{OPT} instead of decreasing it. \square

6 Proof of Theorem 4

The proof of Theorem 4 follows from the following result of Arieli et al. [3] describing the characterization of the best Nash equilibrium in information aggregation games with two senders, one receiver, binary actions and a mediator.

Theorem 6 ([3]). *Let $\Gamma = \{S, A, \Omega, p, u\}$ be an information aggregation game (without filters) with $S = \{s_1, s_2\}$ and $A = \{0, 1\}$ in which the players can communicate with a third-party mediator. Given $i, j, k \in \{0, 1\}$, let $\Omega_{i,j}^k$ be the set of states in which s_1 prefers action i , s_2 prefers action j , and the receiver prefers action k . Let also $\Omega_{i,j} := \Omega_{i,j}^0 \cup \Omega_{i,j}^1$. Consider the following six maps $M_1, M_2, M_3, M_4, M_5, M_6$ from Ω to $[0, 1]$:*

- (a) $M_2(\omega) = 1$ for all $\omega \in \Omega_{0,0}^0$ and $M_2(\omega) = 0$ otherwise.
- (b) $M_1(\omega) = 0$ for all $\omega \in \Omega_{1,1}^1$ and $M_1(\omega) = 1$ otherwise.
- (c) $M_3(\omega) = 1$ for all $\omega \in \Omega_{0,0} \cup \Omega_{0,1}$ and $M_3(\omega) = 0$ otherwise.
- (d) $M_4(\omega) = 1$ for all $\omega \in \Omega_{0,0} \cup \Omega_{1,0}$ and $M_4(\omega) = 0$ otherwise.
- (e) $M_5(\omega) = 1$ for all $\omega \in \Omega$.
- (f) $M_6(\omega) = 0$ for all $\omega \in \Omega$.

Then, there exists an $i \in \{1, 2, 3, 4, 5, 6\}$ and a Nash equilibrium of Γ that maximizes the receiver's utility in which the function that maps each state to the probability that the receiver ends up playing 0 is equal to M_i .

Intuitively, Theorem 6 states that if there were no filters and the two senders had access to a third-party mediator, there exists a Nash equilibrium that is optimal for the receiver in which the either (a) the receiver only plays 0 if all three players prefer 0, (b) the sender plays 1 if and only if all three players prefer 1, (c) the receiver always plays what the first sender prefers, (d) the receiver always plays what the second sender prefers, (e) the receiver plays always 0, or (f) the receiver plays always 1. In the setting described in Section 3.1, players have no access to a mediator. However, we can easily check if the outcomes described in (a), (b), (c), (d), (e) or (f) are incentive-compatible for the receiver (i.e., if the receiver gets more utility with the outcome than when playing with no information), there exists a strategy in Γ (without the mediator) that implements these outcomes.

For instance, suppose that we want to implement the outcome described in (a). Consider the strategy profile $\bar{\sigma}^{(0,0)}$ in which each sender sends a binary signal $m \in \{0, 1\}$ to the receiver that is equal to 0 if and only if the realized state is in $\Omega_{0,0}^0$, and the receiver plays 0 only if both signals are 0. It is easy to check that this is incentive-compatible for the senders: if the realized state is indeed in $\Omega_{0,0}^0$, it is a best response for both to send 0 since it guarantees that the receiver will play 0. Moreover, if the realized state is not in $\Omega_{0,0}^0$, none of the senders gets any additional utility by defecting from the main strategy since the other sender will always send signal 1 (which implies that the receiver will play 1). An analogous reasoning gives a strategy profile $\bar{\sigma}^{(1,1)}$ that implements (b) and is incentive-compatible for the senders. To implement (c), consider the strategy profile $\bar{\sigma}^{s_1}$ in which both senders send a binary signal $m \in \{0, 1\}$ that is equal to 0 if and only if the realized state is in $\Omega_{0,0} \cup \Omega_{0,1}$, and the receiver plays 0 if and only if the signal from the first sender is 0. Again, it is straightforward to check that this is incentive-compatible for the senders: it is a best-response for sender 1 to send its preference, while it doesn't matter what sender 2 sends since it will be ignored. Implementing (d) with a strategy profile $\bar{\sigma}^{s_2}$ that is incentive-compatible for the senders is analogous. To implement (e) (resp., (f)) consider the strategy profile $\bar{\sigma}^0$ (resp., $\bar{\sigma}^1$) in which the senders send signal 0 regardless of the realized state and the receiver always plays 0 (resp., 1). It is easy to check that $\bar{\sigma}^0$ (resp., $\bar{\sigma}^1$) are incentive-compatible for the senders. Since all outcomes that are implementable without a mediator are also implementable with a mediator, this implies the following proposition:

Proposition 6. *Let $\Gamma = \{S, A, \Omega, p, u\}$ be an information aggregation game with $S = \{s_1, s_2\}$ and $A = \{0, 1\}$. Then, either $\bar{\sigma}^0$, $\bar{\sigma}^1$, $\bar{\sigma}^{(0,0)}$, $\bar{\sigma}^{(1,1)}$, $\bar{\sigma}^{s_1}$ or $\bar{\sigma}^{s_2}$ is a Nash equilibrium that is optimal for the receiver.*

Proposition 6 implies that we can break the problem of finding the best Nash equilibrium for the receiver in a filtered information aggregation game into four sub-problems:

- (a) Finding a filter $X_{(0,0)}^{OPT}$ such that $\bar{\sigma}^{(0,0)}$ gives the maximal utility for the receiver in $\Gamma_{X_{(0,0)}^{OPT}}^F$.
- (b) Finding a filter $X_{(1,1)}^{OPT}$ such that $\bar{\sigma}^{(1,1)}$ gives the maximal utility for the receiver in $\Gamma_{X_{(1,1)}^{OPT}}^F$.
- (c) Finding a filter $X_{s_1}^{OPT}$ such that $\bar{\sigma}^{s_1}$ gives the maximal utility for the receiver in $\Gamma_{X_{s_1}^{OPT}}^F$.
- (d) Finding a filter $X_{s_2}^{OPT}$ such that $\bar{\sigma}^{s_2}$ gives the maximal utility for the receiver in $\Gamma_{X_{s_2}^{OPT}}^F$.

Note that these filters might not always exist (for instance, $X_{(0,0)}^{OPT}$ does not exist when the receiver always prefers 0 and the senders always prefer 1). Moreover, we are not including the optimal filters for $\bar{\sigma}^0$ and $\bar{\sigma}^1$ since all filters give the same utility with these strategies. Finding $X_{s_1}^{OPT}$ and $X_{s_2}^{OPT}$ (whenever they exist) reduce to the case of one sender (by ignoring s_2 and s_1 , respectively), and thus can be solved with the $O(k \log k)$ algorithm presented in Section 5.1. We next show how to compute $X_{(0,0)}^{OPT}$ using a linear program with k variables and $2k + 2$ constraints. Computing $X_{(1,1)}^{OPT}$ is analogous.

Suppose that X is a filter that maximizes the receiver utility when playing strategy $\bar{\sigma}^{(0,0)}$. Denote by M_0 the set of signals of X in which both senders and the receiver prefer 0 to 1. Consider a filter X' that samples a signal m according to X and does the following: if $m \in M_0$, it sends signal 0. Otherwise, it sends signal 1. It is straightforward to check that if $\bar{\sigma}^{(0,0)}$ is a Nash equilibrium in Γ_X^F , it is also a Nash equilibrium in $\Gamma_{X'}^F$. Moreover, players get the same utilities with X and X' , which means that X' also maximizes the receiver's utility. This implies that we can restrict our search to binary filters in which all players prefer action 0 on signal 0 and the receiver prefers action 1 on signal 1.

Our aim then is finding a binary filter X such that both senders and the receiver prefer action 0 on signal 0, the receiver prefers action 1 on signal 1, and such that the receiver gets the maximal possible utility by playing 0 on signal 0 and 1 on signal 1. Let $\omega_1, \omega_2, \dots, \omega_k$ be the elements of Ω and define $A_i := u(s_1, \omega_i, 0) - u(s_1, \omega_i, 1)$, $B_i := u(s_2, \omega_i, 0) - u(s_2, \omega_i, 1)$ and $C_i := u(r, \omega_i, 0) - u(r, \omega_i, 1)$. A binary filter X over Ω can be described by a sequence of k real numbers x_1, \dots, x_k between 0 and 1 such that x_i denotes the probability that the filter sends signal 0 on state ω_i . The condition that both senders prefer action 0 on signal 0 translates into

$$\begin{aligned} \sum_{i=1}^k A_i x_i &\geq 0 \\ \sum_{i=1}^k B_i x_i &\geq 0. \end{aligned}$$

Moreover, the utility of the receiver with filter X and strategy $\bar{\sigma}^{(0,0)}$ is given by $\sum_{i=1}^k (x_i u(r, \omega_i, 0) + (1 - x_i) u(r, \omega_i, 1))$. This sum can be rearranged into $\sum_{i=1}^k u(r, \omega_i, 1) + \sum_{i=1}^k C_i x_i$. Therefore, finding a filter $X_{(0,0)}^{OPT}$ such that $\bar{\sigma}^{(0,0)}$ gives the maximal utility for the receiver in $\Gamma_{X_{(0,0)}^{OPT}}^F$ reduces to solving the following linear programming instance:

$$\begin{aligned} \max \quad & \sum_{i=1}^k C_i x_i \\ \text{s.t.} \quad & \sum_{i=1}^k A_i x_i \geq 0 \\ & \sum_{i=1}^k B_i x_i \geq 0 \\ & 0 \leq x_i \leq 1 \quad \forall i \in [k]. \end{aligned}$$

Note that, even though we are maximizing the receiver's utility, it may be the case that $\bar{\sigma}^{(0,0)}$ is not incentive-compatible for the receiver in $\Gamma_{X_{(0,0)}^{OPT}}^F$. For instance, the receiver might prefer playing 0

whenever the senders send 1. If such a thing happens (which can be tested in linear time), it means that there is no filter satisfying the desired conditions.

7 Extending Nash to sequential equilibria

For simplicity, we showed that the strategies provided in the proofs of Theorems 1, 2, 3 and 4 are Nash equilibria. In this section we show how to extend them to sequential equilibria. This means that, for each of these strategy profiles, we have to provide a compatible belief assessment b and argue that the behavior of the players defined off the equilibrium path is a best-response according to their beliefs.

Note that, by construction, the only time that a player can be off the equilibrium path is when the receiver receives a message that the sender was not supposed to send under any circumstances (i.e., a message that had probability 0 of being sent). In the proofs of the Theorems stated above, the strategy profiles that we constructed either had the receiver play the best action for her in the bubbling equilibrium (i.e., the best action for the receiver without any information), or to play the action that gives the least utility to the sender between all the *relevant* actions (i.e., the actions that are not dominated by combinations of other actions). We show next that there exist belief assessments that are compatible with the given strategy profiles in which playing these actions is justified. More generally, we show that, for all relevant actions a , there exists a belief assessment that justifies playing a whenever the receiver receives an impossible message.

We prove the claim above for the case of one sender, the extension to Theorem 4 is straightforward. Denote by $\Omega = \{\omega_1, \omega_2, \dots, \omega_n\}$ the set of states and suppose that, with the proposed strategy σ_s , the sender sends messages in M following a certain function $msg : \Omega \rightarrow \Delta(M)$. Let a be an action that is optimal for the receiver at some posterior distribution $\lambda \in \Delta(\Omega)$ and let m be a message such that $msg(\omega)(m) = 0$ for all $\omega \in \Omega$. Consider a sequence of strategies $\sigma_s^1, \sigma_s^2, \dots$ for the sender in which, in each strategy σ_s^k , the sender computes the message that she would send to the receiver using function msg . However, before sending the message, the sender tosses a coin that lands tails with probability $\frac{\lambda(\omega)}{k}$. If the coin lands heads, she sends the computed message. Otherwise, she sends message m . Clearly, the sequence $\sigma_s^1, \sigma_s^2, \dots$ converges to σ_s . Moreover, the posterior distribution over states induced by each of these strategies conditioned to the fact that the receiver received message m is precisely λ . This means that playing action a is always a best-response for the receiver when receiving message m . Therefore, in the belief assessment induced by $\sigma_s^1, \sigma_s^2, \dots$, a is also a best response to message m .

8 Conclusion

This paper diverges from traditional sender-receiver models, emphasizing the receiver's perspective in cheap talk interactions. We explore two models: one, with transparent motives and state-independent utility, provides a geometric profile of the optimal receiver equilibrium; the second introduces a pre-play stage, empowering the receiver to influence information available to the sender. Our approach aligns with user-centric platforms, where receivers control information, similar to platforms like Amazon. Our findings yield nuanced insights into communication dynamics, enhancing our comprehension of strategic interactions. An intriguing avenue not explored in this paper is the integration of a third intermediary controlling senders' information to maximize the intermediary's utility.

References

- [1] Ricardo Alonso and Odilon Câmara. Persuading voters. *American Economic Review*, 106(11):3590–3605, 2016.
- [2] Itai Arieli and Yakov Babichenko. Private bayesian persuasion. *Journal of Economic Theory*, 182:185–217, 2019.
- [3] Itai Arieli, Ivan Geffner, and Moshe Tennenholtz. Mediated cheap talk design. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pages 5456–5463, 2023.
- [4] Itai Arieli, Ivan Geffner, and Moshe Tennenholtz. Resilient information aggregation. In Rineke Verbrugge, editor, *Proceedings Nineteenth conference on Theoretical Aspects of Rationality and Knowledge, TARK 2023, Oxford, United Kingdom, 28-30th June 2023*, volume 379 of *EPTCS*, pages 31–45, 2023.
- [5] Robert J Aumann, Michael Maschler, and Richard E Stearns. *Repeated games with incomplete information*. MIT press, 1995.
- [6] Dirk Bergemann, Benjamin Brooks, and Stephen Morris. The limits of price discrimination. *American Economic Review*, 105(3):921–957, 2015.
- [7] Sourav Bhattacharya and Arijit Mukherjee. Strategic information revelation when experts compete to influence. *The RAND Journal of Economics*, 44(3):522–544, 2013.
- [8] David Blackwell. Equivalent comparisons of experiments. *The annals of mathematical statistics*, pages 265–272, 1953.
- [9] Andreas Blume and Oliver Board. Language barriers. *Econometrica*, 81(2):781–812, 2013.
- [10] Vincent P Crawford and Joel Sobel. Strategic information transmission. *Econometrica: Journal of the Econometric Society*, pages 1431–1451, 1982.
- [11] Kris De Jaegher. A game-theoretic rationale for vagueness. *Linguistics and Philosophy*, 26(5):637–659, 2003.
- [12] Matthew Gentzkow and Emir Kamenica. Competition in persuasion. *The Review of Economic Studies*, 84(1):300–322, 2016.
- [13] Jeanne Hagenbach and Frédéric Koessler. Cheap talk with coarse understanding. *Games and Economic Behavior*, 124:105–121, 2020.
- [14] Shota Ichihashi. Limiting sender’s information in bayesian persuasion. *Games and Economic Behavior*, 117:276–288, 2019.
- [15] Gerhard Jäger, Lars P Metzger, and Frank Riedel. Voronoi languages: Equilibria in cheap-talk games with high-dimensional types and few signals. *Games and economic behavior*, 73(2):517–537, 2011.
- [16] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- [17] Elliot Lipnowski and Doron Ravid. Cheap talk with transparent motives. *Econometrica*, 88(4):1631–1660, 2020.