

Algorithmic Cheap Talk

Yakov Babichenko* Inbal Talgam-Cohen[†] Haifeng Xu[‡] Konstantin Zabarnyi[§]

November 15, 2023

Abstract

The literature on strategic communication originated with the influential *cheap talk* model, which precedes the Bayesian persuasion model by three decades. This model describes an interaction between two agents: *sender* and *receiver*. The sender knows some *state* of the world which the receiver does not know, and tries to influence the receiver’s action by communicating a cheap talk message to the receiver.

This paper initiates the algorithmic study of cheap talk in a finite environment (i.e., a finite number of states and receiver’s possible actions). We first prove that approximating the sender-optimal or the welfare-maximizing cheap talk equilibrium up to a certain additive constant or multiplicative factor is NP-hard. Fortunately, we identify three naturally-restricted cases that admit efficient algorithms for finding a sender-optimal equilibrium. These include a state-independent sender’s utility structure, a constant number of states or a receiver having only two actions.

*Technion – Israel Institute of Technology — *E-mail*: yakovbab@technion.ac.il.

[†]Tel Aviv University and Technion – Israel Institute of Technology — *E-mail*: inbaltalgam@gmail.com.

[‡]University of Chicago — *E-mail*: haifengxu@uchicago.edu.

[§]Technion – Israel Institute of Technology — *E-mail*: konstzab@gmail.com.

1 Introduction

What information should an informed and rational player (the *sender*) communicate to a decision maker (the *receiver*) so as to optimally influence the decision? Starting with the seminal 1982 work of Crawford and Sobel [CS82] (see also [Mil81; Gro81]), this question has been extensively studied in economics for four decades. The two leading branches of the economic literature on strategic information transmission are *cheap talk* and *Bayesian persuasion*. The fundamental difference between these two branches lies in the ability or disability of the sender to commit. The cheap talk model assumes no commitment power on the sender’s side, while in the Bayesian persuasion model, the sender can commit to a strategically chosen information revelation policy in a trustworthy way.

The work of [DX16] initiated the study of algorithmic aspects of Bayesian persuasion, including the complexity of computing the sender-optimal signaling policy. Since then, Bayesian persuasion and its various extensions have been extensively studied in the algorithmic game theory (AGT) and theory of computation (ToC) literature (see, e.g., [BB17; DX17; AB19; CDHW20; CCMG20; Xu20; GHHS21; HMP21; CMCG21; BMSW24] and an algorithmic survey in [Dug17]). Additionally, signaling policies with commitment have been addressed in many familiar contexts of the AGT literature, including auctions [DIR14; CCD⁺15; DPT16; BBX18; ABMW22]), routing games [BCKS16; ZNX22; GHKK22], abstract games [Dug14; Rub15; BCKS16] and recently in optimal stopping problems [DFH⁺23].

The complementary model of cheap talk is arguably more natural for strategic information communication; indeed, the economic literature has mainly focused on cheap talk for three decades (see, e.g., [Far87; MP89; AS90; FR96; Cra98; Bat02; AH03; OS06; KM16]), before Bayesian persuasion was introduced recently in 2011 by Kamenica and Gentzkow [KG11]. This early focus on cheap talk for modeling strategic communication stems from its minimalistic assumptions – in many scenarios, it is unreasonable to assume that the sender credibly commits to a signaling policy as in Bayesian persuasion, whereas “ordinary, informal talk” [FR96] more accurately describes the communication. The cheap talk model addresses such ordinary communication mathematically, via equilibrium concepts.

Despite being far earlier and more realistic in many scenarios, to the best of our knowledge, cheap talk has not yet been systematically approached from an algorithmic perspective. *The goal of this paper is to investigate the algorithmic aspects of cheap talk in its most basic form.*

1.1 Our Model and Results

We study the classic cheap talk model in its most basic form with discrete states and actions. In this model, a *state* of nature is drawn from a publicly known *prior* distribution with a finite support. The setting includes two players: a *sender* who eventually observes the state and a *receiver* who does not observe it. For clarity of presentation, we use the pronouns she/her for the sender and he/his for the receiver. The sender can “talk” to the receiver to convey information about the state at no cost – i.e., “cheaply”. The information – a.k.a. *signal* – does not necessarily have to be accurate or correct, and it can take any form. The receiver may then use or ignore the signal when choosing an action from a finite set of possibilities. The action coupled with the state determines the utilities of both the sender and the receiver. In some cases, it is natural to assume that the sender’s utility depends only on the action, in which case we say it is *state-independent* (see Subsection 4.1).

As is usual in equilibrium analysis, we assume the sender is playing according to a strategy, known in our context as a *signaling policy*. The signaling policy specifies the probability of sending different signals conditional on every possible state. Upon learning the true state, the sender draws a signal according to the marginal distribution induced by the signaling policy and sends it to the receiver. The receiver knows the prior distribution over the states, as well as the sender’s signaling policy. After getting the signal, he applies a Bayesian update to deduce his *posterior* belief (distribution) over the states. Based on this updated belief,

he then chooses his (possibly random) action according to his own strategy.

The two players' strategies form a *cheap talk equilibrium*, or more formally a *perfect Bayesian equilibrium* [FT91], if no agent has an incentive to deviate. Unlike general Nash equilibrium, finding a cheap talk equilibrium turns out to be a trivial task, since revealing no information always forms a cheap talk equilibrium (a.k.a. the *babbling equilibrium*). The common perspective adopted in the literature on strategic information transmission is the sender's point of view: The goal is to analyze the *sender-optimal equilibrium*, i.e., the equilibrium at which the sender's expected utility is maximized [CS82; KOS07; LR20; LL22]. Therefore, the sender-optimal equilibrium will also be the focus of our computational study.

Our main result shows that it is NP-hard to approximate the sender's utility value at the sender-optimal cheap talk equilibrium up to a certain multiplicative or additive constant (see Theorem 3.1). This result stands in contrast to the Bayesian persuasion model, in which a sender-optimal equilibrium and corresponding sender's utility can be computed efficiently by a simple linear program. This may serve as a complexity-theoretic explanation of the explosive interest in Bayesian persuasion in the recent economic literature after its introduction in [KG11]. Indeed, Bayesian persuasion is commonly considered to be more “workable” than cheap talk, and some recent work has focused on identifying special cases of cheap talk that have similar nice properties – one example is the class of state-independent cheap talk instances identified by [LR20]. Our results suggest a formalization of “workable” through the lens of computational tractability, by showing that finding a sender-optimal cheap talk equilibrium is NP-complete in finite environments (and even NP-hard to approximate), whereas the same problem in Bayesian persuasion is computationally tractable.

Our computational hardness result has implications for the economic analysis of cheap talk. It indicates that a simple characterization for the sender's optimal expected utility in cheap talk equilibria is unlikely to exist – or, at least, the characterization has to be complex enough to encode certain NP-hard problems. This again stands in contrast to the following fundamental results in strategic information transmission:¹

1. In the Bayesian persuasion model, the maximal sender's expected utility is precisely characterized as the concavification of the sender's *indirect utility function* – i.e., sender's utility as a function of the posterior distribution; see [KG11].
2. In the cheap talk model with the simplification of state-independent sender's utility, the sender's optimal expected utility at an equilibrium is precisely characterized as the quasi-concavification of the sender's indirect utility function; see [LR20].

On the positive side, we first show the existence of a sender-optimal equilibrium in which the sender uses no more signals than the number of states (Proposition 2.4). Besides being an interesting property of cheap talk equilibria, this result also implies the NP-membership of the problem of deciding whether the sender can ensure certain utility through cheap talk.

We also identify interesting tractable special cases. One important aforementioned special case of cheap talk is a state-independent sender's utility. We design a polynomial-time algorithm for computing a sender-optimal cheap talk equilibrium in this case (Theorem 4.1). In particular, we establish the tractability of computing the quasi-concavification of a high-dimensional piecewise-constant function.²

We further investigate the source for the hardness of cheap talk – is the problem hard due to a large number of states and/or due to a large number of actions? Our NP-hardness reduction uses cheap talk instances with many states and actions, and our positive results suggest that this hardness may reside in the combination of both. In particular, we show that for a constant number of states (and many actions), the sender-optimal equilibrium can be computed in polynomial time (Proposition 4.4). Also, for two actions (and many states), the sender-optimal equilibrium can be computed in polynomial time (Proposition 4.5) – this is known as the *binary-action* setting, a well-studied special case in the existing literature on strate-

¹The concavification of a function f is the pointwise-smallest concave function weakly above f ; similarly, the quasi-concavification of f is the pointwise-smallest quasi-concave function weakly above f .

²The sender's utility as a function of the receiver's posterior distribution is piecewise constant when the number of possible actions is finite.

gic communication [LRY16; DX17; KMZL17; GS19; BTCXZ22; FHT24]. An interesting remaining open question is whether there exists a polynomial-time algorithm for finding a sender-optimal cheap talk equilibrium with *any constant* number of actions.

1.2 Our Techniques

There are two key differences between cheap talk and Bayesian persuasion due to the lack of the sender’s commitment power in the former:

1. Unlike Bayesian persuasion, in the cheap talk model the sender in each state must only transmit signals that are optimal in that state (i.e., signals that yield the best possible sender’s utility given the receiver’s strategy).
2. In Bayesian persuasion, the sender can deviate to a strategy inducing a posterior distribution that is not induced by any signal under the original strategy. However, it *cannot* happen in cheap talk, as the receiver is assumed to correctly interpret only the signals that are sent with *strictly positive probability* according to the sender’s strategy. To induce a new posterior, the sender has to use a new signal that the receiver is unfamiliar with, and hence he will not be able to interpret the new signal in a meaningful way.

Each of the above differences is insufficient on its own to deduce a hardness result – the linear programming approach that is useful in Bayesian persuasion can be adjusted to tackle it. The combination of *both* these differences is the leading force of our hardness result. A key challenge is the continuous nature of (even the finite) cheap talk problem. Several techniques have been developed in the literature to reduce combinatorial problems to continuous game-theoretic ones. Unfortunately, we have not found the existing techniques applicable to cheap talk (for further discussion, see Subsection 1.3 below). Instead, we develop a novel technique that uses a reduction from a variation of 3SAT (see Section 3).

The positive result for state-independent sender’s utility starts from a characterization of [LR20]. Using this characterization, we reduce the cheap talk problem to a computational geometry problem which we refer to as *simultaneous separation*: Given a point and a collection of polytopes with given equations of facets, determine whether there exists a hyperplane that simultaneously separates the point from all the polytopes. If all the polytopes have polynomially-many vertices, this problem can be straightforwardly solved via linear programming. However, if the polytopes have an exponential number of vertices but admit compact representations, the problem is not straightforward. Using duality techniques, we show that the problem can be efficiently solved even when the number of vertices is exponential.

1.3 Related Literature

Connections to designing AI agents for strategic communication. Our theory work is of possible interest to more applied, cutting-edge research in AI: Driven by the recent maturity of generative language models, the AI/ML field is undergoing a shift of interest. From solving pure tactic-based games (such as Go [SSS⁺17], Poker [BS19] or StarCraft [VBC⁺19]), the focus is shifting to games that require strategic communication via natural language, often in the style of cheap talk. Notable examples include a recent influential work on developing a human-level AI agent for the game of *Diplomacy*, by combining language models with strategic reasoning [FBB⁺22]. *Diplomacy* starts with a communication phase to exchange agents’ private information – in the style of cheap talk without any commitment – followed by a decision-making phase that utilizes learned information from communication. Despite the short period since its publication, this work has spurred a rapidly-growing body of follow-up works on combining language communication with strategic reasoning, ranging from simple matrix games [GSG23] to board games like *Mafia* [O’G23] or *Werewolf* [XYF⁺23; XWL⁺23], and even to complex real-world negotiation [AGS⁺23a]. In all these applications, the bare bones of the problem are a cheap talk game in which the sender tries to

gain a strategic advantage by communicating information to another agent without commitment. An appropriate solution concept for such sequential Bayesian games is the perfect Bayesian equilibrium [FT91] we study. We view our work as a potential complexity-theoretic contribution to this trend, since it studies the algorithmics of computing “good” perfect Bayesian equilibria in arguably the simplest possible situations.

Complexity of equilibrium computation. The complexity of equilibrium computation has been one of the core questions in algorithmic game theory. In a series of important works [LMM03; DGP09; CDT09; Rub16], the complexity of finding a Nash equilibrium and the complexity of finding an approximate Nash equilibrium in a bi-matrix game was settled. From the computational perspective, the cheap talk game model we study has almost the same bi-matrix format, except that one of its dimensions represents the receiver’s action, whereas another dimension represents the state, drawn from a known prior distribution. Thus, the sender’s strategy cannot be distributed arbitrarily as in Nash equilibrium for a bi-matrix game. This subtle difference turns out to be fundamental – finding a cheap talk equilibrium immediately becomes a trivial task since revealing no information is always an equilibrium.

Complexity of “optimal” equilibrium computation. For Nash equilibria, it is known that the maximal clique problem [GZ89], the Monotone-1-in-3SAT problem [DFS17] and the Boolean formula satisfaction problem [CS03] can be reduced to computing the “optimal” equilibrium. However, the problem of computing the optimal Nash equilibrium turns out to significantly differ from the sender-optimal cheap talk equilibrium. One difference is the *non-local nature* of the sender’s strategy, which is not present at Nash equilibrium. In particular, the hardness of finding the optimal Nash equilibrium in almost all the reductions above lies in identifying a subset of player actions – often hinted by the solution to a maximal clique or the satisfying variable assignments – so that mixing them leads to desirable equilibrium utility. However, unlike such *selection* of a subset of actions, the sender’s signaling policy has to be a *decomposition* of the prior distribution over all the states. This is the major reason that we were unable to utilize existing reduction techniques to show the hardness of cheap talk. Instead, we introduce a new variant of the 3SAT problem and reduce to computing a sender-optimal cheap talk equilibrium from this variant, which better suits the above “decomposition” interpretation of sender policies.

Complexity of Bayesian persuasion. Previous reductions for proving hardness in signaling with commitment do not apply to our setting without commitment. The hardness there often comes from the difficulty of coordinating multiple receivers’ actions, e.g., in zero-sum games [Dug14; Rub15] or routing [BCKS16; ZNX22]. The hardness in cheap talk, however, comes from a fundamentally different reason – indeed, computing a sender-optimal cheap talk equilibrium is NP-hard even when there is a single receiver; in this case, Bayesian persuasion can be solved via linear programming.

Other works on strategic communication. Besides the rich body of work on algorithmic Bayesian persuasion mentioned at the beginning, another related branch of literature on strategic information transmission studies the revelation of information that is *ex-post verifiable*, prohibiting the sender from sending a false message. As a consequence, the sender cannot “lie” in her signaling policy; see, e.g., [GH80; Gro81; GR04; HKP17; Tit22]. The work of [HIL⁺22] studies a novel variant in which the sender wishes both to persuade and to inform. In their model, the sender communicates with “anecdotes” – signals that are always truthful – with and without commitment. Their focus is on equilibrium analysis rather than computation. [HMP21] also study persuasion using evident signals, both with and without commitment. They focus on a special case with binary receiver actions *and* state-independent sender’s utility. Interestingly, in their setup with constrained communications, optimal commitment is hard, while cheap talk is easy.

Organization. In Section 2, we formally introduce cheap talk and provide a preliminary result that bounds the required number of signals in a sender-optimal cheap talk equilibrium; as a corollary, this implies that the decision version of computing the maximal sender’s expected utility at an equilibrium belongs to NP.

Section 3 shows that approximating a sender-optimal cheap talk equilibrium up to a certain multiplicative/additive constant is NP-hard. Section 4 analyzes three special cases in which computing such an equilibrium is tractable – a sender with state-independent utility, a constant number of states and a binary-action receiver. Section 5 summarizes, briefly discusses the computational complexity of social welfare maximization in cheap talk and mentions various directions for future research.

2 The Model

Normal form cheap talk games. We study the cheap talk model in arguably its simplest representation, which we refer to as *normal form*.³ In a normal form cheap talk game $\langle A, \Omega, \mu, u_S, u_R \rangle$, there are two players: *sender* (she) and *receiver* (he). The receiver is a decision maker choosing which *action* a to take from a given finite set of m actions $A = \{a_1, \dots, a_m\}$. Both players' utilities are determined by the receiver's action, as well as a random *state* of nature ω , supported on a finite set $\Omega = \{\omega_1, \dots, \omega_n\}$ of size n . We use $u_S(\omega, a)$ and $u_R(\omega, a)$ to denote the sender's and receiver's respective *utilities*. The players share a common *prior belief* $\mu = (\mu(\omega))_{\omega \in \Omega} \in \Delta(\Omega)$ about the random state, where $\mu(\omega)$ is the probability that the realized state is ω , and given some set S , the simplex $\Delta(S)$ is the set of all distributions over S . This makes cheap talk a Bayesian game.

Let Σ be an abstract, publicly known, sufficiently rich discrete set of *signals* (e.g., natural language messages up to a certain length).⁴ Specifically, we require $|\Sigma| \geq n$ (since by Proposition 2.4 below, n signals always suffice for some sender-optimal equilibrium). The game then proceeds as a two-step sequential interaction, as follows:

- Step 1 Nature draws a state $\omega \sim \mu$ from the prior distribution. The sender observes ω and transmits a (possibly random) signal $\sigma \in \Sigma$ to the receiver.
- Step 2 After observing σ , the receiver deduces a posterior distribution over Ω and chooses an action $a \in A$. The sender (resp., receiver) gets utility $u_S(\omega, a)$ (resp., $u_R(\omega, a)$).

Like the classic cheap talk literature, we analyze this two-step game via the standard equilibrium notion for such Bayesian sequential games, known as perfect Bayesian equilibrium [FT91].

Strategies and beliefs. The *sender's (mixed) strategy* in the game is a mapping $\pi : \Omega \rightarrow \Delta(\Sigma)$ that maps observed states to distributions over signals; it is called a *signaling policy*. Let $\pi(\sigma | \omega)$ denote the probability with which the sender transmits the signal σ conditional on observing the state ω . We use $\pi[\omega] \in \Delta(\Sigma)$ to denote the distribution of signals conditional on the state ω . We write $\sigma \sim \pi$ to denote the *unconditional* distribution of σ , i.e., this notation describes a two-step process in which we first draw a state $\omega \sim \mu$, and thereafter – a signal $\sigma \sim \pi[\omega]$.

Upon receiving a signal σ generated from the sender's signaling policy π , the receiver infers a *posterior belief* about the underlying state ω via a standard Bayesian update:

$$p_\sigma(\omega) := p_\sigma^\pi(\omega) := \mathbb{P}(\omega | \sigma) = \frac{\mathbb{P}(\omega, \sigma)}{\mathbb{P}(\sigma)} = \frac{\pi(\sigma | \omega) \cdot \mu(\omega)}{\sum_{\omega' \in \Omega} \pi(\sigma | \omega') \cdot \mu(\omega')}. \quad (1)$$

We denote by $p_\sigma := (p_\sigma(\omega))_{\omega \in \Omega}$ the n -dimensional vector representing the posterior distribution upon receiving signal σ . The *support* of a distribution ν over a discrete set X is $\text{supp}(\nu) := \{x \in X : \mathbb{P}_\nu[x] > 0\}$. In particular, the support of π is the set of all signals in Σ sent with positive probabilities. Notably, for $\sigma \notin \text{supp}(\pi)$, the posterior p_σ remains undefined.

³Many economic papers focus on cheap talk settings with continuum-sized state and action spaces. These settings are often structured; the utility functions are assumed to satisfy continuity and other regularity properties. See, e.g., [CS82; CKS08].

⁴The assumption that Σ is a discrete set is not essential; we only impose it to facilitate the presentation.

Remark 2.1. In the equilibrium notion suitable for cheap talk – perfect Bayesian equilibrium – the sender is not allowed to deviate to sending signals outside of $\text{supp}(\pi)$. More formally, one can define the receiver’s posterior belief outside of $\text{supp}(\pi)$ as being the worst possible for the sender in terms of her expected utility, which would disincentivize the sender from deviating to such signals.

The receiver’s (mixed) strategy is a mapping $s : \Sigma \rightarrow \Delta(A)$ that maps signals to distributions over the action set A . This strategy is set simultaneously with the sender’s strategy and determines how the receiver plays in Step 2. Let $s(a | \sigma)$ denote the receiver’s probability of taking action a conditional on receiving signal σ . We use $s[\sigma] \in \Delta(\Omega)$ to denote the distribution of actions conditional on signal σ .

Cheap talk equilibrium. In cheap talk, no players have commitment power, and thus they set their strategies simultaneously. At equilibrium, a player’s strategy must be a *best response* to the opponent’s strategy. Consider first the receiver. Given the sender’s signal σ and the sender’s strategy π (according to which the receiver performs the Bayesian update as in Equation (1) and deduces posterior p_σ), the receiver derives his expected utility from each action a :

$$\mathbb{E}_{\omega \sim p_\sigma}[u_R(\omega, a)] = \sum_{\omega \in \Omega} p_\sigma(\omega) u_R(\omega, a).$$

The receiver’s *best response action set* given σ and π is:

$$B_R(\sigma, \pi) = \left\{ a \in A : \mathbb{E}_{\omega \sim p_\sigma}[u_R(\omega, a)] \geq \mathbb{E}_{\omega \sim p_\sigma}[u_R(\omega, a')], \forall a' \in A \right\}. \quad (2)$$

The receiver’s distribution over actions $s[\sigma]$ is a *best response* to σ under π if it mixes over only best response actions, i.e., $\text{supp}(s[\sigma]) \subseteq B_R(\sigma, \pi)$. The receiver’s strategy s is a *best response to the sender’s strategy* π if $s[\sigma]$ is a best response to σ for every $\sigma \in \text{supp}(\pi)$.

The sender’s best response requires slightly more careful treatment. Upon observing a state ω , the sender may wish to deviate from sending a signal $\sigma \sim \pi[\omega]$ to sending an arbitrary signal $\sigma' \in \Sigma$. However, as noted in Remark 2.1, in cheap talk the sender can only deviate to signals $\sigma' \in \text{supp}(\pi)$. Thus, π is a best response if the sender does not benefit from any such deviations under any realized state ω . Formally, given the state ω and the receiver’s strategy s , the sender’s *best response signal set* is:

$$B_S(\omega, s) = \left\{ \sigma \in \text{supp}(\pi) : \mathbb{E}_{a \sim s[\sigma]} u_S(\omega, a) \geq \mathbb{E}_{a \sim s[\sigma']} u_S(\omega, a), \forall \sigma' \in \text{supp}(\pi) \right\}. \quad (3)$$

The sender’s signaling policy π is a *best response to receiver strategy* s if $\text{supp}(\pi[\omega]) \subseteq B_S(\omega, s)$ for any ω . That is, conditional on any realized state ω , the distribution of the signals $\pi[\omega]$ only mixes over the sender’s best response signals in the set $B_S(\omega, s)$.

Given the definitions of both players’ best responses above, we are now ready to define the standard equilibrium concept in the cheap talk model [CS82; FR96; LR20]:

Definition 2.2 (Cheap talk equilibrium). Consider a cheap talk game $\langle A, \Omega, \mu, u_S, u_R \rangle$. A pair of mixed strategies (π, s) is a perfect Bayesian equilibrium of the game, a.k.a. a cheap talk equilibrium, if:⁵

- π is a best response to s : $\text{supp}(\pi[\omega]) \subseteq B_S(\omega, s)$ for every $\omega \in \Omega$;
- s is a best response to π : $\text{supp}(s[\sigma]) \subseteq B_R(\sigma, \pi)$ for every $\sigma \in \text{supp}(\pi)$.

⁵Note that the outcome in the Bayesian persuasion model with sender’s commitment power does not require the first constraint on sender’s incentives.

While a cheap talk equilibrium always exists, the intractability of computing a Nash equilibrium naturally raises the computational concern of finding one. Interestingly, it turns out that computing one equilibrium in cheap talk is a trivial task. Indeed, revealing no information always yields an equilibrium, which is known as a *babbling equilibrium*.

Observation 2.3 (Babbling equilibrium). *Let $\Sigma = \{\sigma\}$ be a singleton set and $\pi(\sigma | \omega) = 1 \forall \omega$ be the no-revelation signaling policy. Let s be an action that best responds to the prior μ (that is played when the only possible signal σ is observed). Then the strategy pair (π, s) is a cheap talk equilibrium.*

Therefore, the interesting computational question in cheap talk is optimizing over any *non-trivial* equilibria that may exist in addition to the above no-information ones.

Example. Consider a cheap talk instance with a binary state space $\Omega = \{0, 1\}$ (i.e., $n = 2$), a uniform prior distribution, $m = 4$ actions and the utilities specified by Table 1.

u_S	a_1	a_2	a_3	a_4	u_R	a_1	a_2	a_3	a_4
$\omega = 0$	-1	2	-2	3	$\omega = 0$	3	2	-2	-5
$\omega = 1$	3	-2	2	-1	$\omega = 1$	-5	-2	2	3

Table 1: The utilities u_S and u_R .

Even though we are interested in the cheap talk model, it will be useful to first illustrate the Bayesian persuasion analysis in this example. Due to binary states, any posterior distribution p_σ is fully determined by the probability $p_\sigma(\omega_1)$, since $p_\sigma(\omega_0) + p_\sigma(\omega_1) = 1$. Thus, in the following discussion, we shall identify the posterior with $p_\sigma(\omega_1)$ to allow a more convenient geometric interpretation. Figure 1(a) shows that the action a_1 is a receiver's best response when the posterior is in the segment $[0, \frac{1}{4}]$, the action a_2 is a best response when the posterior is in $[\frac{1}{4}, \frac{1}{2}]$, a_3 is a best response in $[\frac{1}{2}, \frac{3}{4}]$ and a_4 is a best response in $[\frac{3}{4}, 1]$. The *sender's indirect utility* (i.e., sender's utility as a function of receiver's posterior) is as demonstrated in Figure 1(b).

The Bayesian persuasion solution is the concavification of the sender's indirect utility (see [KG11]) evaluated at the prior. In this case, it equals 1. The equilibrium (with sender's commitment) that arises is as follows: The sender splits the prior $\mu = \frac{1}{2}$ into the two posteriors $p_\sigma = \frac{1}{4}$ and $p_\sigma = \frac{3}{4}$ with equal probability $\frac{1}{2}$. As is standard in Bayesian persuasion, we assume that the receiver breaks ties in the sender's favor. Consequently, the receiver plays action a_2 when his posterior is $\frac{1}{4}$, and he plays action a_3 when his posterior is $\frac{3}{4}$.

Note, however, that the above profile does *not* constitute an equilibrium in the cheap talk model (without commitment). At state $\omega = 0$, the sender will prefer sending the signal inducing the posterior $\frac{1}{4}$ with probability 1 (instead of mixing between the two signals). Similar inconsistency arises at state $\omega = 1$ in which the receiver prefers the posterior $\frac{3}{4}$.

Nevertheless, the above equilibrium with sender's commitment can be carefully modified to form a cheap talk equilibrium. Notice that for the posterior $\frac{1}{4}$, the receiver is indifferent between the actions a_1 and a_2 , and for the posterior $\frac{3}{4}$ the receiver is indifferent between the actions a_3 and a_4 . Can the receiver play mixed strategies under these posteriors causing the sender to be indifferent between the two signals in both states? In this example, the answer is affirmative. If at the posterior $\frac{1}{4}$ the receiver plays both actions a_1 and a_2 with equal probability $\frac{1}{2}$, and at the posterior $\frac{3}{4}$ he plays both a_3 and a_4 with equal probability $\frac{1}{2}$, then the sender is indifferent between inducing both posteriors in each state. Indeed, at state $\omega = 0$, the posterior $\frac{1}{4}$ will yield her an expected utility of $\frac{1}{2} \cdot 2 + \frac{1}{2} \cdot (-1) = \frac{1}{2}$, and the posterior $\frac{3}{4}$ will yield her an expected utility

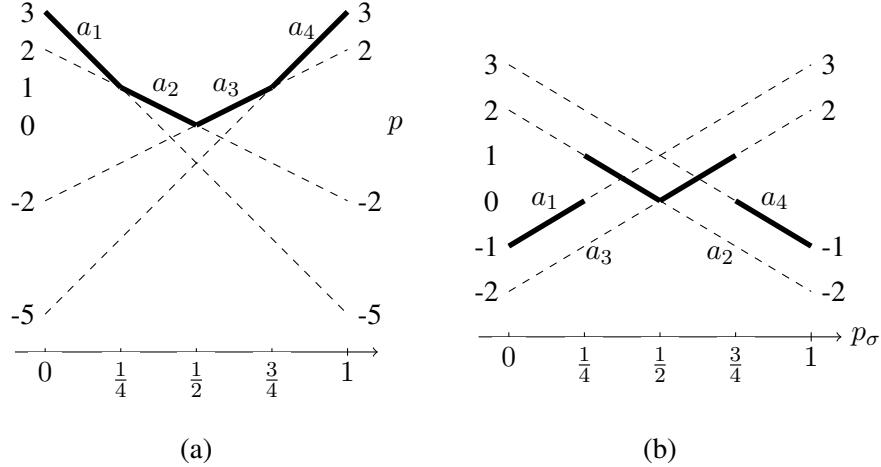


Figure 1: Figure (a) demonstrates the receiver's utility as a function of the posterior probability $p_\sigma(\omega_1)$ (the horizontal axis). Dashed lines capture the utilities of all actions; solid lines capture the utility of the receiver's best action. Figure (b) depicts the sender's indirect utility (i.e., her utility as a function of the receiver's posterior $p_\sigma(\omega_1)$). Again, dashed lines capture the utilities of all actions, while solid lines capture the sender's utility under a receiver's best response action.

of $\frac{1}{2} \cdot (-2) + \frac{1}{2} \cdot 3 = \frac{1}{2}$. Similarly, one may verify it for the state $\omega = 1$. It turns out that the aforementioned equilibrium is sender-optimal, with an expected sender's utility of $1/2$.⁶

2.1 A Tight Bound on the Number of Signals

We now show that, as far as the sender's equilibrium utility is concerned, any cheap talk equilibrium needs not to use more than n signals. The proof uses Carathéodory's Theorem. A similar result holds in the Bayesian persuasion setting as well, but in cheap talk, it requires slightly more delicate arguments.

Proposition 2.4. *For any expected sender's utility value v that can arise at some cheap talk equilibrium, there exists a cheap talk equilibrium (π, s) with $|\text{supp}(\pi)| \leq n$ that leads to expected sender's utility of v .*

The bound is obviously tight. Indeed, when the preferences of the sender and the receiver are perfectly aligned, the cheap talk equilibrium that achieves maximum sender's utility is to reveal the state; this requires n signals.

Proof. Given an equilibrium (π, s) with an arbitrary support size $|\text{supp}(\pi)| > n$ (or an infinite support $|\text{supp}(\pi)| = \infty$), we construct another equilibrium (π', s) with $|\text{supp}(\pi')| \leq n$ and with the same expected sender's utility. The receiver's strategy s will remain unchanged.

Note that the posterior beliefs p_σ , where $\sigma \sim \pi$, must form a split of the prior (this follows from the law of total probability). That is, we have $\mathbb{E}_{\sigma \sim \pi}[p_\sigma] = \mu$. By Carathéodory's Theorem, there exists another split ν of the prior with support size at most n that uses only posteriors from $\{p_\sigma\}_{\sigma \in \text{supp}(\pi)}$. We have $|\text{supp}(\nu)| \leq n$ and $\mathbb{E}_{p \sim \nu}[p] = \mu$. The well-known splitting Lemma of Aumann and Maschler [AMS95] implies that there exists a signaling policy π' whose induced distribution over posteriors is ν , and the signaling policy π' uses at most n signals.

To see that (π', s) is a cheap talk equilibrium, note that:

⁶To verify its optimality, one can apply an exhaustive search on all possible supports of the equilibrium, which by Proposition 2.4 proof can be assumed to be binary.

- Conditional on any state ω , sender's best response signals are $\{p_\sigma : \sigma \in \text{supp}(\pi), p_\sigma(\omega) > 0\}$. Under the strategy π' , the sender transmits a subset of these signals, as needed.
- In every posterior p_σ for $\sigma \in \text{supp}(\pi')$, the receiver best responds. This is because both the posterior and the receiver's strategy remain unchanged compared to the equilibrium (π, s) .

To see that the sender's expected utility remains unchanged, we observe that the sender is *indifferent* between all the posteriors $\{p_\sigma : \sigma \in \text{supp}(\pi), p_\sigma(\omega) > 0\}$ conditional on the state ω . She uses a subset of the signals inducing these posteriors at the new equilibrium (π', s) ; as all the posteriors induced by these signals, as well as the receiver's responses, are unchanged – the sender's expected utility remains unchanged. \square

By Proposition 2.4, there always exists a sender-optimal equilibrium in which the sender uses at most n signals. We immediately obtain a $\text{poly}(m, n)$ -sized witness for the ability of the sender to ensure a certain amount of expected utility at equilibrium. Therefore, we have the following corollary.

Corollary 2.5 (NP Membership of Cheap Talk). *Given a cheap talk instance $\langle A, \Omega, \mu, u_S, u_R \rangle$ and a value v , the problem of deciding whether there exists an equilibrium yielding an expected sender's utility of at least v belongs to NP.*

3 The Hardness of Cheap Talk

Following the approximation algorithm convention, when an equilibrium yields an expected sender's utility within an additive constant $c > 0$ of the best possible, we say it is *c-optimal* (or *c-sender-optimal*) *in the additive sense*. Our main hardness result is as follows.

Theorem 3.1. *There exists an absolute constant $c > 0$ such that it is NP-hard to compute a c-optimal equilibrium in the additive sense for normal form cheap talk games with normalized utilities in $[0, 1]$.*⁷

Theorem 3.1 immediately implies the same hardness result for multiplicative approximation. If an equilibrium yields expected sender's utility within a multiplicative factor $0 < c < 1$ of the best possible, we say it is *c-optimal in the multiplicative sense*. Since the utilities in Theorem 3.1 are normalized in $[0, 1]$, we immediately deduce the following corollary.

Corollary 3.2. *There exists an absolute constant $0 < c < 1$ such that it is NP-hard to compute a c-optimal equilibrium in the multiplicative sense for normal form cheap talk games with non-negative utilities.*

The proof of Theorem 3.1 is non-standard compared to existing hardness results on the complexity of equilibria – see Subsection 1.3. The remainder of this section is devoted to the formal proof. We start with an overview of our proof structure and then dive into the details of each proof step.

Theorem 3.1 proof overview. At a high level, our proof has two major steps: (1) identifying a suitable NP-hard problem instance to reduce from; (2) establishing the reduction. Regarding step (1), it appears challenging to establish an approximation-preserving reduction from classic NP-hard problems to cheap talk games. We thus introduced a new variant of the Max-3SAT problem, termed Max-Var-3SAT, defined below.

Recall that the classical Max-3SAT problem aims to find a Boolean variable assignment to maximize the number of satisfied clauses in a given 3CNF Boolean formula.⁸ The problem instance we shall use

⁷In our proof, the constant is $c = \frac{1}{113792}$. We did not try to optimize this constant, though this gap can be increased by using the inapproximability of other d -regular k SAT instances (larger d and smaller k are preferred). However, even though this may lead to a better (larger) constant, it is a highly non-trivial open question to identify the tight constant.

⁸A Boolean formula is 3CNF (i.e., 3-conjunctive normal form) if it is given by $C_1 \wedge \dots \wedge C_m$, where each clause C_j is a disjunction (\vee) of exactly 3 literals (a literal is a variable x_i or its negation $\neg x_i$), and no variable appears twice in the same clause.

is a variable-maximizing variant of this problem, hence the name **Max-Var-3SAT**, which maximizes the number of assigned Boolean variables without creating any contradictory clauses.

Formally, let ϕ be a 3CNF formula with n variables x_1, \dots, x_n and m clauses. A *partial assignment* x_S to k variables is specified by a subset of variables $S \subseteq [n] = \{1, 2, \dots, n\}$ of size k , as well as an assignment $x_i \in \{\text{True}, \text{False}\}$ for every $i \in S$. We say that a clause is *contradictory* (to satisfiability) under partial assignment x_S if x_S has assigned values to all the 3 variables associated with this clause and their values evaluate the clause to `False`. The partial assignment x_S is said to be *non-contradictory* if it does not create any contradictory clauses. We are now ready to define the **Max-Var-3SAT** problem.

Definition 3.3. *The Max-Var-3SAT problem takes any 3CNF formula ϕ as input and outputs the largest integer k such that there exists a non-contradictory partial assignment x_S of size k (i.e., $|S| = k$).*

Max-Var-3SAT is clearly NP-hard since its decision variant of determining whether there exists a non-contradictory partial assignment of size n or not is precisely the 3SAT problem. However, for our reduction, we shall need a stronger inapproximability result for a restricted version of Max-Var-3SAT. We say that a Max-Var-3SAT instance is *d-regular* if, in the input 3CNF formula, every variable appears in exactly d clauses. In Appendix A.1, we present a proof for the inapproximability of the following problem for 4-regular Max-Var-3SAT.

Proposition 3.4. *It is NP-hard to decide between the following two cases for a 4-regular Max-Var-3SAT instance, which is promised to be in one of the two cases: (1) $\text{Max-Var-3SAT}(\phi) \geq \frac{3047.6}{3048}n$; and (2) $\text{Max-Var-3SAT}(\phi) < \frac{3047.1}{3048}n$.*

The second, and also the most involved step, is to establish an approximation-preserving reduction from the above problem to cheap talk equilibrium computation. Our key conclusion is summarized in the following proposition.

Proposition 3.5. *Suppose $d \leq 6$ and let ϕ be any d -regular 3CNF formula with n variables and m clauses. Then there exists a $\text{poly}(m, n)$ -time algorithm that takes ϕ as input and outputs a cheap talk instance C with $7m$ states, $O(mn)$ receiver's actions and state-dependent sender's utilities in $[-7, 1]$, such that the best expected utility that the sender can achieve in C is $\frac{kd}{7m}$, where $k = \text{Max-Var-3SAT}(\phi)$ is the solution to the d -regular Max-Var-3SAT problem.*

The proof of Proposition 3.5 appears below. Before that, let us show how the two propositions 3.4 and 3.5 indeed complete the proof of Theorem 3.1. Recall that for approximation, one should normalize the sender's utilities to belong to $[0, 1]$. For the instance C obtained from the reduction in Proposition 3.5, this normalization can be done by first adding the constant 7 to each entry of the sender's utility and then dividing them by 8. Since such transformation would not change the set of equilibria, the corresponding sender-optimal cheap talk equilibrium now has expected sender's utility of $(\frac{kd}{7m} + 7)/8 = \frac{kd}{56m} + \frac{7}{8}$, where $k = \text{Max-Var-3SAT}(\phi)$. By Proposition 3.4, it is NP-hard to decide between the two cases

$$k \geq \frac{3047.6}{3048}n := q_1 n \quad \text{and} \quad k < \frac{3047.1}{3048}n := q_2 n.$$

Note that we have $n = \frac{3}{4}m$ for 4-regular Max-Var-3SAT instances ($d = 4$), which implies that the expected sender's utility difference between the above two situations is at least

$$\frac{q_1 n \times d}{56m} - \frac{q_2 n \times d}{56m} = \frac{(q_1 - q_2)4 \times 3/4}{56},$$

which equals $c = 1/113792$ for $d = 4$. This proves that it is NP-hard to compute a c -optimal equilibrium, in the additive sense, for the constructed instance. \square

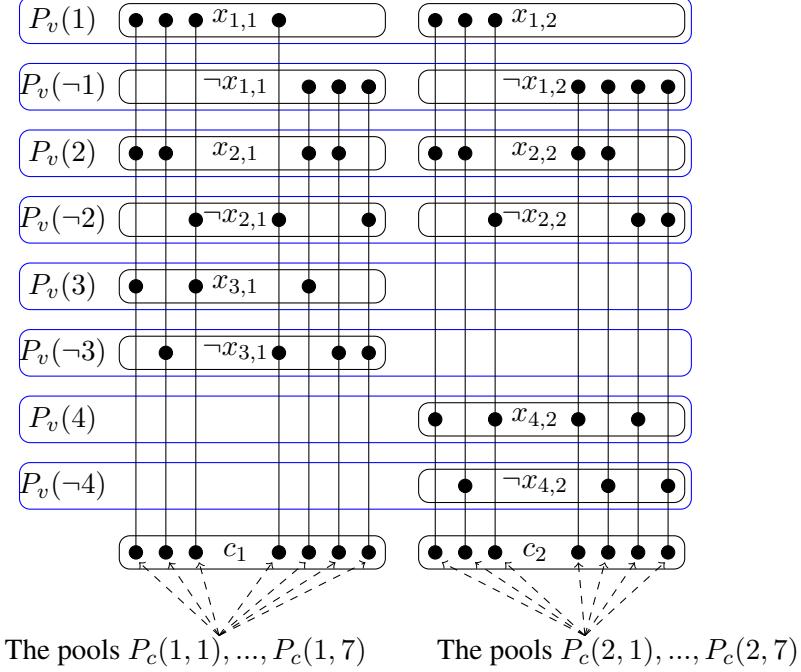


Figure 2: The states, the variable pools and the clause pools of the 3CNF formula with the variables x_1, x_2, x_3, x_4 and the two clauses $x_1 \vee x_2 \vee \neg x_3$ and $\neg x_1 \vee x_2 \vee x_4$. The states are depicted as black rectangles. The variable pools are depicted as blue rectangles. The clause pools are depicted by lines with dots, where a dot indicates that this state belongs to the pool.

3.1 Proof of Proposition 3.5: An Approximation-Preserving Reduction

In this subsection, we present the main part of our proof, which is an approximation-preserving reduction from d -regular Max-Var-3SAT to the computation of a sender-optimal cheap talk equilibrium. Due to the intricacy of the arguments, we present the proof in three steps.

Step 1: Constructing the Cheap Talk Instance

We start by constructing the cheap talk instance $\langle A, \Omega, \mu, u_S, u_R \rangle$ from any given d -regular 3CNF formula ϕ with n variables and m clauses. A graphical illustration of our construction can be found in Figure 2.

States Ω and prior μ . For each variable x_i that appears (possibly as negation) in the j th clause, we create two *variable states* denoted as $x_{i,j}$ and $\neg x_{i,j}$. Additionally, we create a single *clause state* c_j for every clause $j \in [m]$. In total, there are $7m$ states – for every clause $j \in [m]$ there are 6 states for its three variables, as well as their negations, and one additional clause state for itself. Let Ω denote the set of all these $7m$ states. The prior μ over Ω is simply the uniform distribution. The constructions of the receiver's and sender's utilities are more delicate and hinge on a collection of special posterior distributions – i.e., points in the simplex $\Delta(\Omega)$ – that will play a special role in our reduction. We consider the simplex $\Delta(\Omega)$ to be embedded into \mathbb{R}^n , with the i -th coordinate representing the probability for the state being ω_i .

Special posteriors – pools. For every variable x_i , the uniform distribution over the states $\{x_{i,j}\}_j$ will be called the x_i *variable pool* and denoted by $P_v(i)$. Note that $P_v(i)$ is a uniform distribution over d states. Thus, we shall also refer to $P_v(i)$ as a point in the simplex $\Delta(\Omega)$. The name *pool* captures the idea that if the sender decides to pool together the states $\{x_{i,j}\}_j$ (namely, to send the same deterministic signal in all these

states only) then receiver's posterior belief will be exactly the distribution $U(\{x_{i,j}\}_j)$, where $U(S)$ denotes the uniform distribution over a subset S of states.⁹ Similarly, we define the $\neg x_i$ variable pool $P_v(\neg i)$ to be the uniform distribution over the states $\{\neg x_{i,j}\}_j$. Let $P_v := \{P_v(i), P_v(\neg i)\}_{i \in [n]}$ denote the set of all variable pools.

We now describe *clause pools*. For every clause $j \in [m]$, there are 7 True/False assignments for its variables $x_{i_1}, x_{i_2}, x_{i_3}$ that will satisfy the clause. For each one of these 7 assignments, we create a *clause pool* that is the uniform distribution over four states – the state c_j and the three corresponding states $x_{i_l,j}$ or $\neg x_{i_l,j}$ based on whether $x_{i_l} = \text{True}$ or $x_{i_l} = \text{False}$, respectively. For example, if $x_{i_1} = \text{True}$, $x_{i_2} = \text{False}$, $x_{i_3} = \text{True}$ is a satisfying assignment for the j th clause, then it induces the clause pool $U(\{c_j, x_{i_1,j}, \neg x_{i_2,j}, x_{i_3,j}\})$. These pools are denoted by $P_c(j, 1), \dots, P_c(j, 7) \in \Delta(\Omega)$ for every $j \in [m]$. We denote by $P_c := \{P_c(j, 1), \dots, P_c(j, 7)\}_{j \in [m]}$ the set of all clause pools.

We further introduce *singleton variable pools* that are the Dirac distributions over some $x_{i,j}$ and the Dirac distributions over some $\neg x_{i,j}$.¹⁰ These pools are denoted by $P_s(i, j) \in \Delta(\Omega)$ and $P_s(\neg i, j) \in \Delta(\Omega)$, respectively. Let $P_s := \{P_s(i, j), P_s(\neg i, j)\}_{i,j}$ (where i, j run over all $j \in [m]$ and $i \in [n]$ s.t. the j th clause contains the variable x_i or its negation) be the set of all singleton variable pools. Importantly, there are no singleton pools over the clause states c_j . The collection of all types of pools is denoted by $P := P_v \cup P_c \cup P_s$.

Receiver's actions A and utility u_R . At a high level, we would like to design Receiver's actions and utilities in a way that induces the following situation: Whenever the receiver's posterior p is not one of the pools – i.e., $p \notin P$ – then any receiver's best response action will be a “bad action” for the sender; yet for every posterior distribution $p \in P$, there will be a special action a_p that will serve as a best response only at the point p ; the actions a_p might not be bad for the sender, depending on the state. We shall show that it is indeed possible to achieve these requirements using polynomially many actions for the receiver. Since this result about a decision maker under uncertainty (the receiver) may be of independent interest, we state it as a proposition.

Proposition 3.6. *Given a finite state space Ω and any finite collection of points $P \subseteq \Delta(\Omega)$, there exists a polynomial-time algorithm that outputs a collection of $(|\Omega| + 1)|P|$ actions for the receiver $A = \{a_p\}_{p \in P} \cup \{a_{p,\omega}\}_{p \in P, \omega \in \Omega}$ and receiver's utilities $u_R : \Omega \times A \rightarrow \mathbb{R}$ with the following properties:*

1. *Under any posterior $y = p \in P$, the receiver's action a_p is a best response and any other action $a_{p'}$ for $p' \neq p$ is not a best response.¹¹*
2. *Under every posterior $y \notin P$, none of the actions a_p for $p \in P$ is a best response.*

In the proposition above, actions $a_{p,\omega}$ will play the role of “bad actions” for the sender, which will be used by the receiver when $y \notin P$. To construct the receiver's actions and utilities in our reduction, we apply Proposition 3.6 with P being the set of all the constructed *pools* above. It will incentivize the sender to induce only the receiver's posterior beliefs that belong to the pool set P , creating combinatorial structures over the sender's optimal signaling policy.

The proof of Proposition 3.6 leverages the geometry of convex functions to design the receiver's actions and utility values. Its main challenge is to use only polynomially-many receiver's actions to achieve the two desirable properties. We defer its technical argument to Appendix A.2. For the remainder of the proof, we shall utilize Properties (1) and (2) of Proposition 3.6 only, rather than the actual construction in the proof.

Sender's utility u_S . For every receiver's action a_{p,ω_i} ($p \in P, \omega_i \in \Omega$), we set sender's utility $u_S(\omega, a_{p,\omega_i}) = -7$ at every state ω . As mentioned previously, these are “bad actions” for the sender, giving her the least

⁹Note that the same posterior can arise in randomized signaling policies by pooling together fractions of the corresponding states.

¹⁰A *Dirac distribution* assigns probability 1 to some element.

¹¹Notably, some $a_{p',\omega}$ may also be best responses under posterior belief p . This property only ensures that $a_{p'}$ will not be a best response under $p \neq p'$, and this will suffice for our later argument.

possible utility regardless of the state. Now we turn to the actions a_p for $p \in P$. We remind that for any state ω , the statement $\omega \in \text{supp}(p)$ is equivalent to saying that the state ω appears in the pool p (i.e., a subset of states). We define the sender's utility by:

$$u_S(\omega, a_p) = \begin{cases} 1 & \text{if } p \in P_v \text{ and } \omega \in \text{supp}(p) \\ 0 & \text{if } p \in P_c \cup P_s \text{ and } \omega \in \text{supp}(p) \\ -7 & \text{otherwise.} \end{cases}$$

Namely, the sender can gain positive utility only under some *variable pool* $p \in P_v$, since this p is the only posterior belief under which the action a_p can be played. While aiming to put as high as possible mass on variable pools, in parallel the sender has the challenge of avoiding the high penalty of -7 . This is particularly challenging at a clause state c_j – unlike the variable states, which can form singleton pools in P_s , clause states must be pooled with some variable states to avoid the large penalty. However, this may contradict the sender's use of variable states to obtain positive utility under variable pools. Such conflicting reward-penalty tradeoff is the intrinsic source of hardness, as we shall formalize next.

Step 2: Upper-Bounding Sender's Utility via Equilibrium Analysis

The main goal of our second step is to establish an upper bound of $\frac{kd}{7m}$ for the sender's expected utility at any cheap talk equilibrium. We achieve this by proving several lemmas concerning the structure of the sender-optimal equilibrium. These lemmas also build the connection between our constructed instance and the Max-Var-3SAT problem.

Given any cheap talk equilibrium, a variable pool $p \in P_v$ will be called *attractive* if the posterior p is induced with positive probability by the sender, and under the posterior p the receiver plays the action a_p with a probability strictly above $\frac{7}{8}$. Note that this probability yields a strictly positive expected sender's utility in the relevant states. We observe that in an equilibrium, if a variable pool $p \in P_v$ is attractive, then the sender sends the signal that induces the posterior p at all relevant states (i.e., at states $x_{i,j}$ if $p = P_v(i)$ and at states $\neg x_{i,j}$ if $p = P_v(\neg i)$) with probability 1. This is because at these states, only the posterior p yields a strictly positive sender's utility; this follows from our instance construction, where variable pools form a partition of all variable states, and a variable state can enjoy strictly positive utility if and only if the variable pool containing it is attractive.

Our first lemma states that there cannot be two contradicting attractive pools at any sender-optimal equilibrium.

Lemma 3.7. *There is no sender-optimal equilibrium at which both pools $P_v(i)$ and $P_v(\neg i)$ are attractive.*

Proof. We shall show that any cheap talk equilibrium that contains contradicting attractive pools cannot be sender-optimal. Specifically, we can adjust it to a strictly better equilibrium for the sender.

Let $C \subseteq [n]$ denote the set of indices of contradicting attractive pools (i.e., $i \in C$ means both $P_v(i)$ and $P_v(\neg i)$ are attractive). For every clause j that contains a variable x_i from $i \in C$, we know that the sender's utility at the state c_j is necessarily -7 . This is because the only way to avoid the penalty of -7 is by pooling c_j into a clause pool; however, since in both states $x_{i,j}$ and $\neg x_{i,j}$ the sender sends their corresponding variable pool signal $P_v(i), P_v(\neg i)$ with probability 1, the mass of these states is taken and cannot be pooled with c_j .

We can apply the following modification to avoid this penalty (see Figure 3 for illustration). First, we break all the contradicting attractive pools in C . Second, for every variable state that currently is not part of any attractive variable pool, we separate out all its probability mass from the posteriors it belongs to. Note that after these two modifications, for every variable x_i (not necessarily in C), at least one of $x_{i,j}, \neg x_{i,j}$ is

fully available in the sense that its probability mass is not used in any of the remaining pools or posteriors. Third, for every clause j that contains a variable x_i from $i \in C$, we pool every clause state c_j with the currently available variable states into a clause pool. Since we have broken both attractive variable pools $P_v(i)$ and $P_v(\neg i)$, at least one of them can be pooled with c_j , together with the other two available variables (as for at least one of them we shall, indeed, get a clause pool). Fourth, all the remaining variable states that have not been pooled after this operation are set as singleton pools. Finally, we adjust the receiver's strategy to best respond to all the new pools, and in the sender's optimal manner when there is indifference.

To see that the modifications above still result in a cheap talk equilibrium, note that all the variable states in the modified strategies are deterministically set to one of the following three types of signals: (1) an attractive variable pool; (2) a singleton variable pool consisting of the variable state itself; (3) a clause pool containing this variable. Thus, the remaining posteriors can only be supported on some clause states that do not appear in any of the clause pools, and they will lead to the sender's utility of -7 under *any* signal. By construction, the sender's strategy after the above modification is a best response to the receiver's strategy.

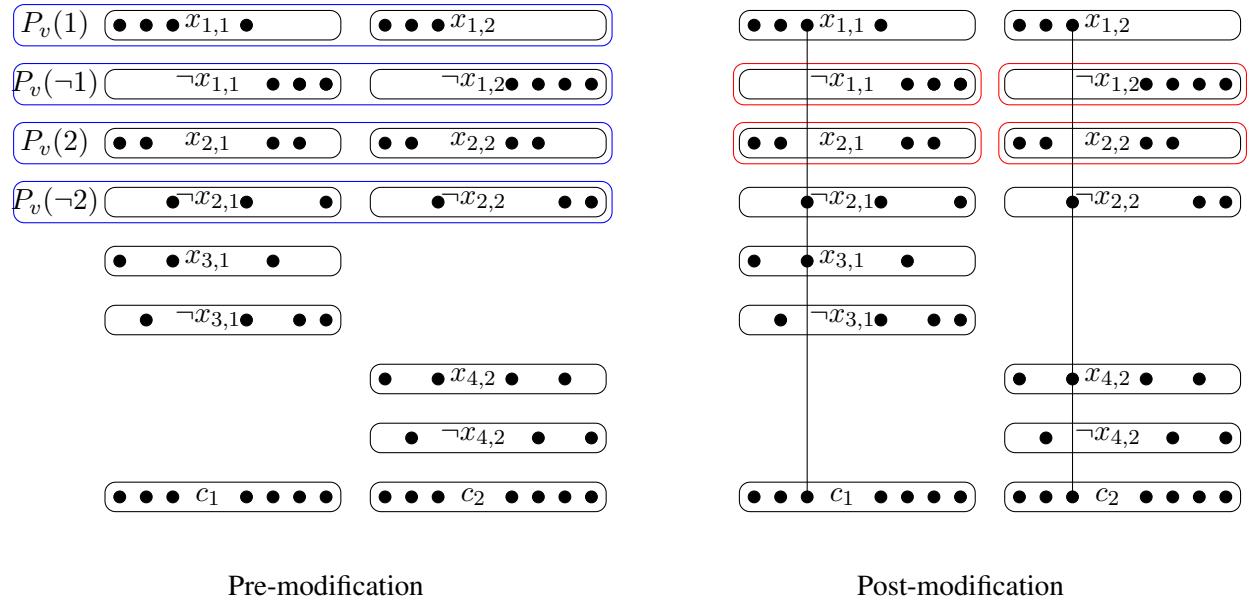


Figure 3: Illustration of a possible re-pooling in Lemma 3.7 proof for $n = 4, m = 2$, with the clauses being $c_1 = x_1 \vee x_2 \vee \neg x_3$ and $c_2 = \neg x_1 \vee x_2 \vee x_4$. The black rectangles represent states; the blue rectangles – attractive variable pools (note that $P_v(1)$ and $P_v(\neg 1)$ are contradicting, and so are $P_v(2)$ and $P_v(\neg 2)$); the black lines – clause pools (with dots marking the states in these pools); and the red rectangles – singleton pools. The modification breaks all contradicting attractive variable pools and creates new 2 clause and 4 singleton pools. The clause pools correspond to the assignment $(x_1, x_2, x_3, x_4) = (\text{True}, \text{False}, \text{True}, \text{True})$ that satisfies both clauses. The signaling policy in states $\neg x_{3,1}$ and $\neg x_{4,2}$ remains unchanged during the modification.

Finally, let us count the sender's gains and losses from the above modification. For every clause j that contains a variable x_i for some $i \in C$, we have increased the utility at clause state c_j by 7. The total number of such clauses is at least $|C|d/3$. Meanwhile, there are $|C|d \times 2$ states at which we have decreased the utility by 1 due to breaking all the contradicting attractive variable pools in C . In all other states, the utility remained unchanged. Since $7 \times \frac{|C|d}{3} - |C|d \times 2 > 0$, such a modification strictly improves the sender's utility. \square

Lemma 3.7 implies that the attractive pools form a partial assignment for the variables. Formally, given

an equilibrium, we define the *induced (partial) assignment* as follows.

- If $P_v(i)$ is an attractive pool, the induced assignment will be $x_i = \text{False}$.
- If $P_v(\neg i)$ is an attractive pool, the induced assignment will be $x_i = \text{True}$.
- If neither $P_v(i)$ nor $P_v(\neg i)$ is an attractive pool, we do not assign any value to x_i .

Note that the fourth option of both being attractive is excluded by Lemma 3.7. The next lemma shows that in a sender-optimal equilibrium, contradictions cannot be created within a clause.

Lemma 3.8. *For every sender-optimal equilibrium, its induced assignment to the variables of ϕ does not contradict any clause.*

Proof. Given an equilibrium that contradicts a clause, we modify it to a better one for the sender. Let $j \in [m]$ be a contradictory clause that contains the variables $x_{i_1}, x_{i_2}, x_{i_3}$. For simplicity of notations, we assume that clause j is $x_{i_1} \vee x_{i_2} \vee x_{i_3}$ and the induced assignment is $x_{i_1} = x_{i_2} = x_{i_3} = \text{False}$, which leads to contradiction; same arguments apply for the other seven cases. This means that $P_v(i_1), P_v(i_2)$ and $P_v(i_3)$ are attractive variable pools, and that the pool that consists of the states $\{c_j, \neg x_{i_1,j}, \neg x_{i_2,j}, \neg x_{i_3,j}\}$ is not one of the clause pools of the j th clause.

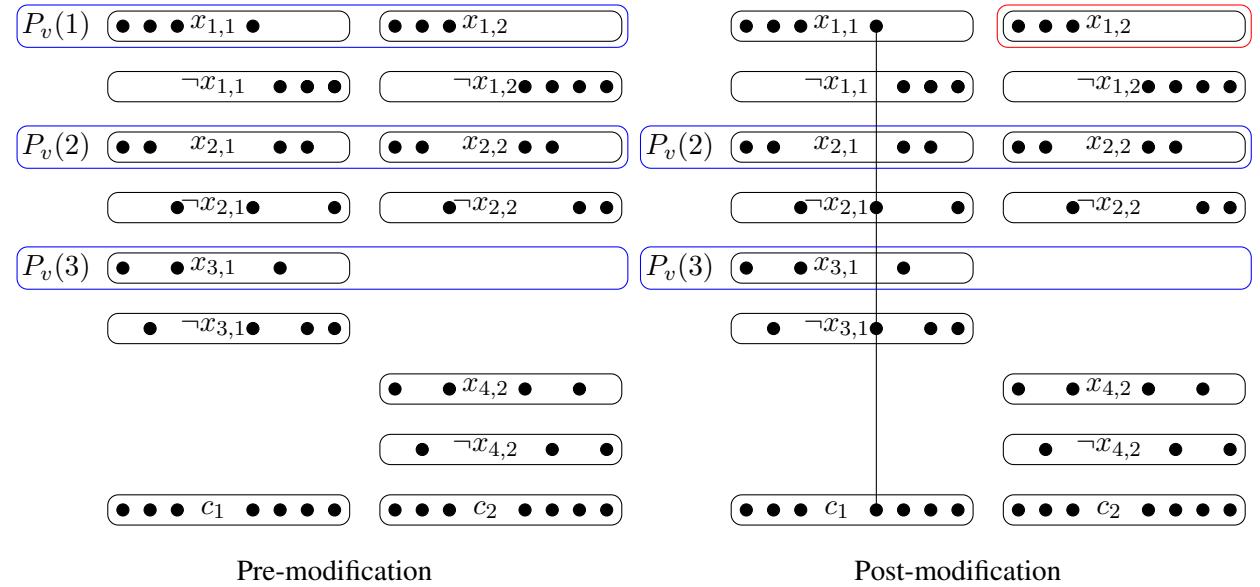


Figure 4: Illustration of re-pooling in Lemma 3.8 proof for $n = 4, m = 2$, demonstrated on the clause $x_1 \vee x_2 \vee x_3$ with the index $j = 1$. The black rectangles represent states, the blue rectangles – attractive variable pools, the black lines – clause pools (with dots in the states belonging to them) and the red rectangle – a singleton pool. Under the transformation, the attractive variable pool $P_v(1)$ is removed, and a new clause pool $\{c_1, x_{i_1,1}, \neg x_{i_2,1}, \neg x_{i_3,1}\}$ is created, together with a new singleton pool $\{x_{i_1,2}\}$. The signaling policy in states $\neg x_{i_1,1}, \neg x_{i_1,2}, \neg x_{i_2,2}, x_{i_4,2}, \neg x_{i_4,2}$ and c_2 remains unchanged during the modification.

Since in all three states $x_{i_1,j}, x_{i_2,j}$ and $x_{i_3,j}$ the sent signal is the one that induces the variable pool posterior, one cannot use these states to pool c_j together with them. Therefore, at state c_j , the sender's utility must be -7 . We apply the following modification to the equilibrium (see Figure 4 for illustration). We break the pool $P_v(i_1)$. Now we pool the states $\{c_j, x_{i_1,j}, \neg x_{i_2,j}, \neg x_{i_3,j}\}$ into a clause pool. Finally, for all the now un-pooled states $x_{i,j'} (j' \neq j)$, we pool them into singleton pools. The resultant sender's strategy still yields a cheap talk equilibrium, assuming that the receiver takes the corresponding action a_p for each newly created pool p .

The sender's utility gain from this modification is 7 for state c_j , whereas her losses from this modification appear in the d states of $P_v(i_1)$ in which the utility has been reduced from 1 to 0. Since we assume $d \leq 6$, such a modification strictly increases the sender's utility, as desired. \square

Lemma 3.8 implies that in a sender-optimal equilibrium, the induced partial assignment to the variables of ϕ must not contradict any clause. Moreover, the sender's utility is at most 0 for any posterior that is not an attractive pool. Therefore, we can bound from above the sender's utility in any equilibrium by $\frac{kd}{7m}$, where k is the maximum size of a non-contradicting variable assignment. Note that kd counts the total number of states appearing in all attractive pools, and $\frac{1}{7m}$ is a normalization factor due to the uniform prior over the $7m$ states.

Step 3: Lower-Bounding Sender's Utility via Equilibrium Construction

Finally, we construct an equilibrium that indeed achieves the $\frac{kd}{7m}$ expected sender's utility upper bound proved above. Given a partial assignment x_S for k variables with $|S| = k$ and $x_i = \text{True/False}$ for every $i \in S \subseteq [n]$, we construct a cheap talk equilibrium as follows:

- First, we create k variable pools based on the partial assignment for the k variables. If $x_i = \text{True}$, we create the pool $P_v(\neg x_i)$, whereas if $x_i = \text{False}$ – we create the pool $P_v(x_i)$.
- Second, for every clause $j \in [m]$, the sender pools c_j with the corresponding states $x_{i,j}$ or $\neg x_{i,j}$ in a way that is consistent with the assignment to create a clause pool for every clause. That is, if $x_i = \text{True}$, the sender pools c_j with $x_{i,j}$ (since $\neg x_i$ was used to create the pool $P_v(\neg x_i)$ above); and if $x_i = \text{False}$, then the sender pools c_j with $\neg x_{i,j}$. If x_i is not assigned a value, then we include in c_j 's pool the state corresponding to that literal from the pair $\{x_{i,j}, \neg x_{i,j}\}$ which appears in the j th clause (we can consider this as allowing to give the variable x_i different values to satisfy different clauses). In this way, we guarantee that each clause belongs to some clause pool.
- Third, all the remaining variable states that have not been pooled yet are pooled into singleton pools. This completes the description of the sender's signaling policy.
- Finally, since the sender's signaling policy only leads to posteriors in the pool set P , the receiver's response is simply set to take action a_p deterministically for every created pool p .

To see that the construction above is a cheap talk equilibrium, note that the receiver is obviously best responding by the construction of his utility as in Proposition 3.6. Moreover, the receiver's deduction of the posterior beliefs is Bayesian given the sender's strategy. To see that the sender cannot improve her utility by sending a different signal at any state, note that her utility at any state is either the largest possible utility 1 (at a variable pool) or 0 (at any other pools). Whenever her utility is 0, she cannot induce a strictly positive utility by deviating to a different pool by our construction. Finally, note that the sender's utility at this equilibrium is precisely $\frac{kd}{7m}$, concluding Proposition 3.5 proof.

4 Tractable Cases

Following the negative result from the previous section, we turn to analyzing well-motivated special cases. We note that the previous intractability result holds when both the number of states and the number of actions are variable, and moreover – sender's utility is state-dependent. In this section, we show that in the following three cases, computing a sender-optimal cheap talk equilibrium is tractable: (i) the sender's utility is independent of the state; (ii) there are $n = O(1)$ states; (iii) there are $m = 2$ actions for the receiver.

4.1 State-Independent Sender's Utility

In this subsection, we show how to compute a sender-optimal cheap talk equilibrium, provided that the sender's utility does not depend on the state. Throughout the subsection, we shall assume $u_S(\omega, a) = u_S(a)$ for every $\omega \in \Omega$. We remark that the case of state-independent sender's utility has been extensively studied in the literature [CH07; LR20; DK21; AGS23b] and is an important special case with many applications. As noted by [LR20]: “Salespeople want to sell products [...]; politicians want to get elected; lawyers want favorable rulings; and so on”.

Theorem 4.1. *If the sender's utility is state-independent, then a sender-optimal cheap talk equilibrium can be computed in $\text{poly}(m, n)$ -time.*

We devote the remainder of this subsection to proving Theorem 4.1. Interestingly, our algorithm turns out to hinge on an efficient solution to the following geometric problem, for which we shall develop an efficient algorithm first.

Definition 4.2 (Simultaneous Polytope Separation (SPS)). *Let $P_0, P_1, \dots, P_m \subseteq \mathbb{R}^d$ be $m + 1$ polytopes, with each $P_i = \{x : A_i x \leq b_i\}$ given in its canonical representation for some matrix-vector pair (A_i, b_i) . The simultaneous polytope separation (SPS) problem looks to determine whether there exists a hyperplane of form $\alpha x + \beta = 0$ that simultaneously separates polytope P_0 from all other polytopes P_i for $i \in [m]$, and if so – identify one such hyperplane.*

A special case of SPS is when each P_i degenerates to a single point. In this situation, the SPS can be solved efficiently via a linear feasibility problem. This approach easily generalizes to the situations in which *every* P_i has polynomially-many vertices. We shall show that SPS admits an algorithm whose running time is polynomial in the canonical representation of the polytopes P_i given by matrix-vector (A_i, b_i) ; thus, it can be solved even when every P_i might have exponentially-many vertices. Moreover, this algorithm is practically efficient and only requires solving a single max-min problem that we derive next.¹²

We consider the following max-min problem, where “max” is over the parameter $\alpha \in B[0, 1]$ (the closed unit ball) and “min” is over the parameter $\lambda \in \Delta_m$ (the m -dimensional simplex):

$$\max_{\alpha: \|\alpha\|_2 \leq 1} \min_{\lambda \in \Delta_m} \mathcal{L}(\alpha; \lambda), \quad \text{where } \mathcal{L}(\alpha; \lambda) := \min_{x_0 \in P_0} (\alpha \cdot x_0) - \sum_{i \in [m]} \lambda_i \cdot \max_{x_i \in P_i} (\alpha \cdot x_i). \quad (4)$$

Note that the function is homogeneous in α , which is why we restrict α to be within the unit ball to guarantee the existence of a solution. Since minima of linear functions are concave, whereas maxima of linear functions are convex, we know that $\mathcal{L}(\alpha; \lambda)$ is concave in α and linear in λ . Thus, the max-min problem can be solved efficiently in polynomial time. The following lemma shows that the max-min problem (4) is all we need to solve the SPS problem.

Lemma 4.3. *Let (α^*, λ^*) be an optimal solution to $\max_{\alpha: \|\alpha\|_2 \leq 1} \min_{\lambda \in \Delta_m} \mathcal{L}(\alpha; \lambda)$. Then $\mathcal{L}(\alpha^*; \lambda^*) \geq 0$ if and only if there exists a hyperplane that separates P_0 from all P_i for $i \in [m]$. Moreover, whenever such a hyperplane exists, it is given explicitly as follows: $\alpha^* x + \beta^* = 0$, where $\beta^* := -\min_{x_0 \in P_0} (\alpha^* \cdot x_0)$.*

¹²We acknowledge that during private communication, Dirk Bergemann and Yang Cai observed the tractability for a special case of SPS with P_0 as a *singleton point* through the lens of the *ellipsoid method*. Specifically, since we can efficiently optimize every linear function over polytope P_i for each $i \in [m]$, we can also optimize over the union of these polytopes, and thus also over their convex hull. The equivalence between optimization and separation, as a consequence of the ellipsoid method, implies an efficient separation oracle for determining whether the point P_0 is in their convex hull or not. Despite learning this elegant proof, we choose to present a different proof here based on a clean, yet non-trivial max-min re-formulation of SPS, which leads to a more general (applicable to polytope P_0), and also a more practical algorithm.

The proof of Lemma 4.3 is relegated to Appendix B.1. We now employ it to prove Theorem 4.1.

Proof of Theorem 4.1. At a high level, the core reason underlying this tractability is that when the sender’s utility is state-independent, every signal sent at an equilibrium signaling policy must have the same sender’s utility, irrespective of the realized state. This property enables a characterization of the possible sender-expected utilities that can arise at an equilibrium. Specifically, [LR20] shows that there exists a cheap talk equilibrium yielding an expected sender’s utility of at least v if and only if there exists a signaling policy π and a receiver’s best response strategy $s[\sigma]$ that *secure* utility v for the sender – i.e., the expected sender’s utility when $a \sim s[\sigma]$ and $\omega \sim p_\sigma$ is at least v for every $\sigma \in \text{supp}(\pi)$. To secure the sender’s utility, the receiver’s strategy s is simple – always best responding and breaking ties in favor of the sender whenever there are ties. Thus, the main difficulty is to identify the signaling policy π . Moreover, given such π , [LR20] shows a simple way to modify π to become a cheap talk equilibrium with utility v for the sender. The equilibrium construction is polynomial (in the number of states and actions).

We shall prove that there exists a $\text{poly}(m, n)$ -time algorithm that outputs a signaling policy π securing the sender’s utility of v whenever it exists, or asserts that such a π does not exist. The core idea of our proof is to relate the above condition to the SPS problem defined above. Let us consider the $(n - 1)$ -dimensional simplex of posteriors $\Delta(\Omega)$. Let region B_a denote the set of all posteriors under which the action a is a receiver’s best response. Region B_a is a polytope specified by $m - 1$ linear constraints on the posterior p . Specifically, $B_a = \{p \in \Delta(\Omega) : \sum_{\omega \in \Omega} p(\omega)u_R(\omega, a) \geq \sum_{\omega \in \Omega} p(\omega)u_R(\omega, a') \quad \forall a' \in A \setminus a\}$. For any value v , let $A_v \subseteq A$ denote the set of receiver actions yielding a sender’s expected utility of at least v . Note that A_v can be found in $O(m)$ -time by enumerating the m actions.

We argue that there exists a signaling policy π that secures the sender’s utility of v if and only if there is no hyperplane separating the prior distribution μ from all the polytopes $\{B_a\}_{a \in A_v}$ (i.e., an SPS problem). To see the “only if” direction, suppose some π secures the sender’s utility of v . Then there is a best response receiver’s action $a(\sigma)$ that yields the sender’s utility of at least v for every signal σ used by π . That is, $a(\sigma) \in A_v$ for every $\sigma \in \text{supp}(\pi)$. Consequently, $\mu = \mathbb{E}_{\sigma \sim \pi}[p_\sigma]$ must be in the convex hull of the polytopes $\{B_a\}_{a \in A_v}$, and thus cannot be separated from all of these polytopes. The “if” direction follows a similar argument. Suppose μ cannot be separated from all the polytopes $\{B_a\}_{a \in A_v}$. Then μ must be within their convex hull, and thus there exists a signaling policy under which each signal yields best responses in A_v . Moreover, such a signaling policy can be computed efficiently by solving linear systems.

Finally, we apply Lemma 4.3 to determine whether there is a hyperplane separating the prior distribution μ from all the polytopes $\{B_a\}_{a \in A_v}$, concluding the proof. \square

4.2 A Constant Number of States

In this subsection, we prove that finding a sender-optimal cheap talk equilibrium is computationally tractable when the number of states n is constant. The proof is based on the bound of $n = O(1)$ on $|\text{supp}(\pi)|$ at a sender-optimal equilibrium, as ensured by Proposition 2.4. This bound allows a translation of the cheap talk problem to a normal form game between the sender and the receiver in which the sender has $O(1)$ actions. The main result of this subsection is as follows.

Proposition 4.4. *Suppose that the number of states n is constant. Then one can compute a sender-optimal equilibrium (π, s) in $\text{poly}(m)$ -time.*

Proof. Fix a set of signals Σ of size n . By Proposition 2.4, this set is rich enough to include the support of the sender’s signaling policy at some sender-optimal equilibrium.

The cheap talk interaction can be viewed as a two-player interaction in which:

- The sender has n^n pure strategies (i.e., mappings from Ω to Σ), and the signaling policies are mixtures over these.

- The receiver has m^n pure strategies (i.e., mappings from Σ to A), and his mixed strategies are mixtures over these.

Notice that the sender has $n^n = 0(1)$ actions, and the receiver has $m^n = \text{poly}(m)$ actions. Therefore, one can apply an exhaustive search over all supports of size $n^n = 0(1)$ for both players to find all the possible equilibrium outcomes (in terms of the two agents' utilities), including the sender-optimal one – see, e.g., the *support testing* algorithm in [vS07].¹³ \square

4.3 A Binary-Action Receiver

In this subsection, we show that when the number of actions is $m = 2$, finding a sender-optimal cheap talk equilibrium is computationally tractable. Moreover, just two signals are required; these signals separate the states in which the sender prefers action a_1 over a_2 from the states in which she prefers a_2 over a_1 .

Proposition 4.5. *When the receiver has $m = 2$ actions, there exists a $\text{poly}(n)$ -time algorithm for computing a sender-optimal cheap talk equilibrium (π, s) . Furthermore, in the outputted (π, s) , the sender uses at most two signals, and the receiver uses only pure strategies upon getting any signal – i.e., $|\text{supp}(\pi)| \leq 2$ and $|\text{supp}(s[\sigma])| = 1$ for every $\sigma \in \text{supp}(\pi)$.*

The proof outline relies on the possibility of reducing the analysis of mixtures over two actions to a single-dimensional optimization problem (and, therefore, does not generalize to more than two actions). Specifically, let σ_1 be a signal s.t. $s[\sigma_1]$ assigns the highest probability to the receiver taking action a_1 . Since there are just two actions, the sender weakly betters off deviating to deterministically transmitting σ_1 in any state in which she prefers action a_1 over a_2 . Therefore, in all the states in which the sender strictly prefers a_1 over a_2 , she must deterministically induce some posterior p (the same for all these states); this posterior may further be induced in some states in which the sender is indifferent between the two actions. A simple analysis shows that if under p the receiver weakly prefers a_1 over a_2 , then he must deterministically play a_1 as a response to p or we get a contradiction to sender-optimality;¹⁴ and if it is not possible to induce such p , then the sender cannot do better than a babbling equilibrium. For the formal proof, see Appendix B.2.

5 Discussion

Computational complexity of social welfare maximization. The *social welfare* is defined as the sum of the sender's and receiver's utilities $u_S + u_R$. It immediately follows from Theorem 3.1 that computing a certain additive or multiplicative constant approximation for the social welfare-maximizing equilibrium is NP-hard. This is because one can scale down the receiver's utility function to make it negligible (e.g., smaller than $c/2$ in Theorem 3.1 notations) compared to the sender's utility. We are yet to find a positive computational result on maximizing social welfare in a cheap talk equilibrium. In particular, the characterization of [LR20] that allowed us to deduce a positive computational result for sender-optimal equilibria when the sender's utility is state-independent no longer applies.

Further research on algorithmic cheap talk. Our results lead to many interesting opportunities for future research. One immediate question following our hardness result is whether there is a quasi-polynomial time algorithm for finding a sender-best or welfare-maximizing cheap talk equilibrium, and what the tight inapproximability constant c is. Moreover, the sender's utility is state-independent in many possible real-life

¹³If the matrix of receiver's utilities does not have the full rank of $\min\{m, n\}$ then the exhaustive search might miss some equilibria, but it is still done over every equilibrium outcome in terms of utilities [vS07]. More precisely, for any mixed Nash equilibrium, there exists an equilibrium yielding the same respective sender's and receiver's expected utilities s.t. each player mixes over at most n^n pure strategies.

¹⁴Similar considerations can also be applied to the states in which the receiver prefers a_2 over a_1 .

applications, and we show the computational tractability of finding a sender-optimal cheap talk equilibrium in this case. It remains to be explored whether it can be extended to multiple or even just two receivers, with either private or public signaling.

We showed that finding a sender-optimal cheap talk equilibrium is intractable. An interesting follow-up question is whether determining if a cheap talk instance admits a non-trivial equilibrium (i.e., not a babbling equilibrium) is computationally tractable. What about finding such a non-trivial equilibrium? Another question is an appropriate notion of approximate-equilibrium for cheap talk (cf. [Rub16]), and the complexity of finding an optimal such solution. Finally, we prove that when there are constantly-many states, it is possible to compute a sender-optimal equilibrium in a polynomial time. However, it is an interesting open question on whether an analogous result holds for constantly many actions, even for the constant being three.

Zooming out, given the rich body of algorithmic studies of Bayesian persuasion, it is not difficult to see the various possible follow-up algorithmic directions pertaining to cheap talk. Examples include algorithmic cheap talk with many receivers, privately or publicly, with or without receiver externalities. It will also be interesting to investigate computational applications of cheap talk in domains such as entertainment games with communications [FBB⁺22], security games and negotiations, in which commitment is often impossible.

Alternative frameworks for strategic information transmission. Another interesting research direction is to analyze the computational aspects of intermediate models between the two extreme settings – Bayesian persuasion (full sender’s commitment power) and cheap talk (no commitment at all). A natural question is whether the sender’s optimization problem is tractable for persuasion with partial sender’s commitment – e.g., in the credible persuasion framework [LL22].

Acknowledgements. Yakov Babichenko is supported by the Binational Science Foundation BSF grant No. 2018397. This work is funded by the European Union (ERC, ALGOCONTRACT, 101077862, PI: Inbal Talgam-Cohen). Inbal Talgam-Cohen is supported by a Google Research Scholar award and by the Israel Science Foundation grant No. 336/18. Yakov and Inbal are supported by the Binational Science Foundation BSF grant No. 2021680. Haifeng Xu is supported by NSF award No. CCF-2303372, Army Research Office Award No. W911NF-23-1-0030 and Office of Naval Research Award No. N00014-23-1-2802. Konstantin Zabarnyi is supported by a PBC scholarship for Ph.D. students in data science.

References

- [AB19] Itai Arieli and Yakov Babichenko. Private Bayesian persuasion. *Journal of Economic Theory*, 182:185–217, 2019.
- [ABMW22] Reza Alijani, Siddhartha Banerjee, Kamesh Munagala, and Kangning Wang. The limits of an information intermediary in auction design. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 849–868, 2022.
- [AGS⁺23a] Sahar Abdehnabi, Amr Gomaa, Sarath Sivaprasad, Lea Schönherr, and Mario Fritz. LLM-deliberation: Evaluating LLMs with interactive multi-agent negotiation games. *arXiv preprint arXiv:2309.17234*, 2023.
- [AGS23b] Itai Arieli, Ronen Gradwohl, and Rann Smorodinsky. Informationally robust cheap-talk. *arXiv preprint arXiv:2302.00281*, 2023.
- [AH03] Robert J Aumann and Sergiu Hart. Long cheap talk. *Econometrica*, 71(6):1619–1660, 2003.

[AMS95] Robert J Aumann, Michael Maschler, and Richard E Stearns. *Repeated games with incomplete information*. MIT press, 1995.

[AS90] David Austen-Smith. Information transmission in debate. *American Journal of political science*, pages 124–152, 1990.

[Bat02] Marco Battaglini. Multiple referrals and multidimensional cheap talk. *Econometrica*, 70(4):1379–1401, 2002.

[BB17] Yakov Babichenko and Siddharth Barman. Algorithmic aspects of private Bayesian persuasion. In *8th Innovations in Theoretical Computer Science Conference (ITCS 2017)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017.

[BBX18] Ashwinkumar Badanidiyuru, Kshipra Bhawalkar, and Haifeng Xu. Targeting and signaling in ad auctions. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 2545–2563. SIAM, 2018.

[BCKS16] Umang Bhaskar, Yu Cheng, Young K Ko, and Chaitanya Swamy. Hardness results for signaling in Bayesian zero-sum and network routing games. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 479–496, 2016.

[BKS04] Piotr Berman, Marek Karpinski, and Alexander Scott. Approximation hardness of short symmetric instances of MAX-3SAT. Technical report, Institut des Hautes Etudes Scientifiques, 2004.

[BMSW24] Siddhartha Banerjee, Kamesh Munagala, Yiheng Shen, and Kangning Wang. Fair price discrimination. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*. SIAM, 2024.

[BS19] Noam Brown and Tuomas Sandholm. Superhuman ai for multiplayer poker. *Science*, 365(6456):885–890, 2019.

[BTCXZ22] Yakov Babichenko, Inbal Talgam-Cohen, Haifeng Xu, and Konstantin Zabarnyi. Regret-minimizing Bayesian persuasion. *Games and Economic Behavior*, 136:226–248, 2022.

[CCD⁺15] Yu Cheng, Ho Yee Cheung, Shaddin Dughmi, Ehsan Emamjomeh-Zadeh, Li Han, and Shang-Hua Teng. Mixture selection, mechanism design, and signaling. In *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pages 1426–1445. IEEE, 2015.

[CCMG20] Matteo Castiglioni, Andrea Celli, Alberto Marchesi, and Nicola Gatti. Online Bayesian persuasion. *Advances in Neural Information Processing Systems*, 33:16188–16198, 2020.

[CDHW20] Rachel Cummings, Nikhil R Devanur, Zhiyi Huang, and Xiangning Wang. Algorithmic price discrimination. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2432–2451. SIAM, 2020.

[CDT09] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player nash equilibria. *Journal of the ACM (JACM)*, 56(3):1–57, 2009.

[CH07] Archishman Chakraborty and Rick Harbaugh. Comparative cheap talk. *Journal of Economic Theory*, 132(1):70–94, 2007.

[CKS08] Ying Chen, Navin Kartik, and Joel Sobel. Selecting cheap-talk equilibria. *Econometrica*, 76(1):117–136, 2008.

[CMCG21] Matteo Castiglioni, Alberto Marchesi, Andrea Celli, and Nicola Gatti. Multi-receiver online Bayesian persuasion. In *International Conference on Machine Learning*, pages 1314–1323. PMLR, 2021.

[Cra98] Vincent P Crawford. A survey of experiments on communication via cheap talk. *Journal of Economic theory*, 78(2):286–298, 1998.

[CS82] Vincent P Crawford and Joel Sobel. Strategic information transmission. *Econometrica: Journal of the Econometric Society*, pages 1431–1451, 1982.

[CS03] Vincent Conitzer and Tuomas Sandholm. Complexity results about nash equilibria. In *Proceedings of the 18th international joint conference on Artificial intelligence*, pages 765–771, 2003.

[DFH⁺23] Bolin Ding, Yiding Feng, Chien-Ju Ho, Wei Tang, and Haifeng Xu. Competitive information design for Pandora’s box. In *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 353–381. SIAM, 2023.

[DFS17] Argyrios Deligkas, John Fearnley, and Rahul Savani. Computing constrained approximate equilibria in polymatrix games. In *Algorithmic Game Theory: 10th International Symposium, SAGT*, pages 93–105. Springer, 2017.

[DGP09] Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou. The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.

[DIR14] Shaddin Dughmi, Nicole Immorlica, and Aaron Roth. Constrained signaling in auction design. In *Proceedings of the 2014 ACM-SIAM Annual Symposium on Discrete Algorithms, SODA*, pages 1341–1357, 2014.

[DK21] Christoph Diehl and Christoph Kuzmics. The (non-) robustness of influential cheap talk equilibria when the sender’s preferences are state independent. *International Journal of Game Theory*, 50(4):911–925, 2021.

[DPT16] Constantinos Daskalakis, Christos H Papadimitriou, and Christos Tzamos. Does information revelation improve revenue? In *Proceedings of the 2016 ACM Conference on Economics and Computation, EC*, pages 233–250. ACM, 2016.

[Dug14] Shaddin Dughmi. On the hardness of signaling. In *2014 IEEE 55th Annual Symposium on Foundations of Computer Science*, pages 354–363. IEEE, 2014.

[Dug17] Shaddin Dughmi. Algorithmic information structure design: A survey. *ACM SIGecom Exchanges*, 15(2):2–24, 2017.

[DX16] Shaddin Dughmi and Haifeng Xu. Algorithmic Bayesian persuasion. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 412–425, 2016.

[DX17] Shaddin Dughmi and Haifeng Xu. Algorithmic persuasion with no externalities. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 351–368, 2017.

[Far87] Joseph Farrell. Cheap talk, coordination, and entry. *The RAND Journal of Economics*, 18(1):34–39, 1987.

[FBB⁺22] Meta Fundamental AI Research Diplomacy Team (FAIR)[†], Anton Bakhtin, Noam Brown, Emily Dinan, Gabriele Farina, Colin Flaherty, Daniel Fried, Andrew Goff, Jonathan Gray, Hengyuan Hu, et al. Human-level play in the game of diplomacy by combining language models with strategic reasoning. *Science*, 378(6624):1067–1074, 2022.

[FHT24] Yiding Feng, Chien-Ju Ho, and Wei Tang. Rationality-robust information design: Bayesian persuasion under quantal response. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*. SIAM, 2024.

[FR96] Joseph Farrell and Matthew Rabin. Cheap talk. *The Journal of Economic Perspectives*, 10(3):103–118, 1996.

[FT91] Drew Fudenberg and Jean Tirole. Perfect Bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory*, 53(2):236–260, 1991.

[GH80] Sanford J Grossman and Oliver D Hart. Disclosure laws and takeover bids. *The Journal of Finance*, 35(2):323–334, 1980.

[GHHS21] Ronen Gradwohl, Niklas Hahn, Martin Hoefer, and Rann Smorodinsky. Algorithms for persuasion with limited communication. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 637–652. SIAM, 2021.

[GHKK22] Svenja M Griesbach, Martin Hoefer, Max Klimm, and Tim Koglin. Public signals in network congestion games. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 736–736, 2022.

[GR04] Jacob Glazer and Ariel Rubinstein. On optimal rules of persuasion. *Econometrica*, 72(6):1715–1736, 2004.

[Gro81] Sanford J Grossman. The informational role of warranties and private disclosure about product quality. *The Journal of Law and Economics*, 24(3):461–483, 1981.

[GS19] Yingni Guo and Eran Shmaya. The interval structure of optimal disclosure. *Econometrica*, 87(2):653–675, 2019.

[GSG23] Kanishk Gandhi, Dorsa Sadigh, and Noah D Goodman. Strategic reasoning with language models. *arXiv preprint arXiv:2305.19165*, 2023.

[GZ89] Itzhak Gilboa and Eitan Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93, 1989.

[HIL⁺22] Nika Haghtalab, Nicole Immorlica, Brendan Lucier, Markus Mobius, and Divyarthi Mohan. Communicating with anecdotes. Working paper available at <https://doi.org/10.48550/arXiv.2205.13461>, 2022.

[HKP17] Sergiu Hart, Ilan Kremer, and Motty Perry. Evidence games: Truth and commitment. *American Economic Review*, 107(3):690–713, 2017.

[HMP21] Martin Hoefer, Pasin Manurangsi, and Alexandros Psomas. Algorithmic persuasion with evidence. In *12th Innovations in Theoretical Computer Science Conference (ITCS 2021)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2021.

[KG11] Emir Kamenica and Matthew Gentzkow. Bayesian persuasion. *American Economic Review*, 101(6):2590–2615, 2011.

[KM16] Vijay Krishna and John Morgan. Cheap talk. In *The New Palgrave Dictionary of Economics*, pages 1–8. Palgrave Macmillan UK, London, 2016.

[KMZL17] Anton Kolotilin, Tymofiy Mylovanov, Andriy Zapecelnyuk, and Ming Li. Persuasion of a privately informed receiver. *Econometrica*, 85(6):1949–1964, 2017.

[KOS07] Navin Kartik, Marco Ottaviani, and Francesco Squintani. Credulity, lies, and costly talk. *Journal of Economic theory*, 134(1):93–116, 2007.

[LL22] Xiao Lin and Ce Liu. Credible persuasion. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 469–469, 2022.

[LMM03] Richard J Lipton, Evangelos Markakis, and Aranyak Mehta. Playing large games using simple strategies. In *Proceedings of the 4th ACM Conference on Electronic Commerce*, pages 36–41, 2003.

[LR20] Elliot Lipnowski and Doron Ravid. Cheap talk with transparent motives. *Econometrica*, 88(4):1631–1660, 2020.

[LRY16] Zhuozheng Li, Heikki Rantakari, and Huanxing Yang. Competitive cheap talk. *Games and Economic Behavior*, 96:65–89, 2016.

[Mil81] Paul R Milgrom. Good news and bad news: Representation theorems and applications. *The Bell Journal of Economics*, pages 380–391, 1981.

[MP89] Steven A Matthews and Andrew Postlewaite. Pre-play communication in two-person sealed-bid double auctions. *Journal of Economic theory*, 48(1):238–263, 1989.

[O'G23] Aidan O'Gara. Hoodwinked: Deception and cooperation in a text-based game for language models. *arXiv preprint arXiv:2308.01404*, 2023.

[OS06] Marco Ottaviani and Peter N Sørensen. Reputational cheap talk. *The Rand journal of economics*, 37(1):155–175, 2006.

[Rub15] Aviad Rubinstein. ETH-hardness for signaling in symmetric zero-sum games. *CoRR abs/1510.04991*, 2015.

[Rub16] Aviad Rubinstein. Settling the complexity of computing approximate two-player nash equilibria. In *2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 258–265. IEEE Computer Society, 2016.

[SSS⁺17] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, et al. Mastering the game of go without human knowledge. *Nature*, 550(7676):354–359, 2017.

[Tit22] Maria Titova. Persuasion with verifiable information. Technical report, working paper, 2022.

[VBC⁺19] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *Nature*, 575(7782):350–354, 2019.

[vS07] Bernhard von Stengel. Equilibrium computation for two-player games in strategic and extensive form. *Algorithmic game theory*, pages 53–78, 2007.

[Xu20] Haifeng Xu. On the tractability of public persuasion with no externalities. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*, pages 2708–2727. SIAM, 2020.

[XWL⁺23] Yuzhuang Xu, Shuo Wang, Peng Li, Fuwen Luo, Xiaolong Wang, Weidong Liu, and Yang Liu. Exploring large language models for communication games: An empirical study on werewolf. *arXiv preprint arXiv:2309.04658*, 2023.

[XYF⁺23] Zelai Xu, Chao Yu, Fei Fang, Yu Wang, and Yi Wu. Language agents with reinforcement learning for strategic play in the werewolf game. *arXiv preprint arXiv:2310.18940*, 2023.

[ZNX22] Chenghan Zhou, Thanh H Nguyen, and Haifeng Xu. Algorithmic information design in multi-player games: Possibilities and limits in singleton congestion. In *Proceedings of the 23rd ACM Conference on Economics and Computation*, pages 869–869, 2022.

A Omitted Proofs in Section 3

A.1 Proof of Proposition 3.4

Our reduction starts from the following known hard problem.

Lemma A.1. *[adapted from Theorem 2 of [BKS04]] For every $\epsilon \in (0, 1/2)$, there exists a family of 4-regular 3SAT instances such that each instance ϕ consists of $m = 1016M$ clauses for some M , and it is NP-hard to distinguish between the following two situations for a ϕ that is promised to be in one of them:*

- (A1) *at least $(1016 - \epsilon)M$ clauses of ϕ can be satisfied simultaneously;*
- (A2) *at most $(1015 + \epsilon)M$ clauses in ϕ can be satisfied simultaneously.*

To avoid carrying cumbersome numbers, let $q_1 := \frac{1016 - \epsilon}{1016}$ be the fraction of the $m = 1016M$ clauses that case **A1** can satisfy, whereas $q_2 := \frac{1015 + \epsilon}{1016}$ is the fraction for case **A2**. Note that $3m = 4n$. We reduce the problem in Lemma A.1 to distinguishing between the following two situations for our 4-regular Max-Var-3SAT problem, when promised that ϕ is in one of the following two cases:

- (B1) ϕ admits a non-contradictory partial assignment to at least $\frac{4q_1 - 1}{3}n = \frac{3048 - 4\epsilon}{3048}n$ variables;
- (B2) any partial assignment to strictly more than $\frac{2+q_2}{3}n = \frac{3047 + \epsilon}{3048}n$ variables must lead to contradictory clauses in ϕ .

The proposition proof follows by taking $\epsilon = 0.1$. The reduction goes by proving that, for any 4-regular 3SAT formula ϕ , case **A1** implies case **B1** and case **A2** implies case **B2**. Consequently, if there is an algorithm for solving the above distinguishing problem for 4-regular Max-Var-3SAT problem, this algorithm – when used for the instances from Lemma A.1 – can also distinguish case **A1** (which satisfies **B1**) from case **A2** (which satisfies **B2**). It, thus, implies the NP-hardness of distinguishing **B1** from **B2** for 4-regular Max-Var-3SAT instances, when promised that the instance is at one of the two cases.

We first argue that case **A1** implies case **B1**. Suppose that there is a full variable assignment that can satisfy at least $q_1 m$ clauses. This means there are at most $(1 - q_1)m$ unsatisfied clauses. We can modify this full assignment by removing any one variable from each of the unsatisfied clauses and make its value unassigned. This results in a partial assignment with at least $n - (1 - q_1)m = n - \frac{4(1 - q_1)n}{3} = \frac{4q_1 - 1}{3}n$ assigned variables. Therefore, if the instance ϕ belongs to case **A1**, it must belong to case **B1**.

Next, we argue that case **A2** implies case **B2**. We prove its contrapositive – that is, if there is a non-contradictory partial assignment of size strictly larger than $\frac{2+q_2}{3}n$, then strictly more than $q_2 m$ clauses can be satisfied in ϕ . Indeed, this non-contradictory partial assignment leaves less than $\frac{1-q_2}{3}n$ variables unassigned, and these variables can span less than $\frac{4(1-q_2)}{3}n$ clauses. Therefore, more than $m - \frac{4(1-q_2)}{3}n = q_2 m$ clauses are satisfied by this partial assignment, as desired.

A.2 Proof of Proposition 3.6

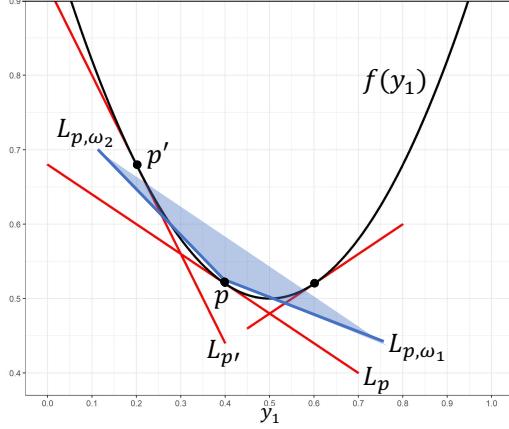


Figure 5: Geometric Illustration for the Proof of Proposition 3.6.

The proof is geometric. We start with considering some strictly convex function $f : \Delta(\Omega) \rightarrow \mathbb{R}$. For instance, one can pick $f(y) = \sum_i y_i^2$, though any strictly convex function will be equally good. An illustration for the two-state situation is provided in Figure 5, where the input variable y_1 to $f(y_1)$ denotes the probability of the first state.

For every $p \in P$, let $L_p : \Delta(\Omega) \rightarrow \mathbb{R}$ be the tangent plane to the plot of f at the point p . Notice that L_p uniquely defines an action a_p for the receiver in terms of receiver's utility. Indeed, given some state ω_i ($i \in [n]$), consider the vector e_i of length n with the i th coordinate being 1 and the other coordinates being 0. Setting $u_R(\omega_i, a_p) = L_p(e_i)$ ensures, by linearity of receiver's expected utility as a function of the posterior $y \in \Delta(\Omega)$, that $L_p(y)$ exactly equals the expected receiver's utility at the posterior y .

Fix $p' \in \Delta(\Omega) \setminus p$. The strict convexity of f implies that $L_{p'}(y) < f(y)$ for every $y \neq p'$. In particular, for $y = p \neq p'$, we get $L_{p'}(p) < f(p) = L_p(p)$, where $L_{p'}(p)$ is receiver's utility of action $a_{p'}$ at the posterior belief p . This implies that a_p is the unique best response among the actions $\{a_{p'}\}_{p' \in P}$ at the posterior p , which ensures Property (1).

Now we turn to define the additional actions $a_{p,\omega}$ to ensure Property (2). Let us set

$\epsilon := \min_{p,p' \in P: p \neq p'} \{L_p(p) - L_p(p')\}$. Since the set P is finite, ϵ is well-defined and strictly positive. We want to make the action a_p sub-optimal for every $y \neq p$. Geometrically, we create a reverse pyramid whose “button vertex” is p (the shadow region of Figure 5). The facets of the pyramid correspond to these additional actions. However, additional carefulness is needed when defining this pyramid. We must avoid the situations in which one the new facets will strictly exceed the value of $L_{p'}(p')$ for some point $p' \neq p$, which will ruin Property (1). Under the illustration of Figure 5, we must have $L_{p,\omega}(p') \leq L_{p'}(p')$ for every ω and p' . It turns out that such actions can be explicitly constructed.

Indeed, for every $p = (p_{\omega_i})_{i \in [n]} \in P$ and $\omega \in \Omega$, let us define the linear function $L'_{p,\omega}(y) := y_{\omega} - p_{\omega}$. Note that $L'_{p,\omega}(y) \leq 1$ and $L'_{p,\omega}(p) = 0$. Moreover, for every $y \neq p$, there exists $\omega \in \Omega$ such that $L'_{p,\omega}(y) > 0$. This is because $\sum_{\omega \in \Omega} L'_{p,\omega}(y) = \sum_{\omega \in \Omega} (y_{\omega} - p_{\omega}) = 1 - 1 = 0$, and thus we know that at least one of $\{L'_{p,\omega}(y)\}_{\omega}$ is strictly positive when $y \neq p$.

We now define the linear function $L_{p,\omega} := L_p + \epsilon L'_{p,\omega}$. For a fixed p and variable values of ω , these linear functions can be thought of as pyramid facets (see Figure 5). Denote the action that corresponds to some $L_{p,\omega}$ by $a_{p,\omega}$ – i.e., $u_R(\omega, a_{p,\omega}) = L_{p,\omega}(e_i)$ ($i \in [n]$). We argue that these additional actions $a_{p,\omega}$ ensure Property (2) without violating Property (1).

To see that Property (2) holds, we observe that for every $y \notin P$ and every $p \in P$, we have $L_{p,\omega}(y) =$

$L_p(y) + \epsilon L'_{p,\omega}(y) > L_p(y)$ for $\omega \in \Omega$ satisfying $L'_{p,\omega}(y) > 0$. Namely, any a_p is not a best response under any posterior belief $y \notin P$. To see that Property (1) is not violated, we observe the following two facts:

- For any $p' \neq p$, we have $L_{p',\omega}(p) = L_{p'}(p) + \epsilon L'_{p',\omega}(p) \leq L_{p'}(p) + \epsilon \leq L_p(p)$, where the last inequality follows from the definition of ϵ .
- Additionally, we have $L_{p,\omega}(p) = L_p(p) + \epsilon L'_{p,\omega}(p) = L_p(p)$.

Thus, even with introduction of the $a_{p',\omega}$ actions, a_p remains a best response under the posterior p (though it is not the unique best response, since all the $a_{p,\omega}$ actions are also best responses). This concludes the proof.

B Omitted Proofs in Section 4

B.1 Proof of Lemma 4.3

We first prove the “if” direction. Suppose $\alpha x + \beta = 0$ is a hyperplane that separates P_0 from all P_i ($i \in [m]$). By definition, we have $\alpha x_0 + \beta \geq 0$ for any $x_0 \in P_0$, whereas $\alpha x_i + \beta \leq 0$ for any $x_i \in P_i$, $i \in [m]$. Thus, $\min_{x_0 \in P_0} (\alpha \cdot x_0 + \beta) \geq 0$ and $\max_{x_i \in P_i} (\alpha \cdot x_i + \beta) \leq 0$. For any non-negative convex coefficients $\lambda \in \Delta_m$, we must have:

$$\begin{aligned} 0 &\leq \min_{x_0 \in P_0} (\alpha \cdot x_0 + \beta) - \sum_{i \in [m]} \lambda_i \times \max_{x_i \in P_i} (\alpha \cdot x_i + \beta) \\ &= \beta + \min_{x_0 \in P_0} (\alpha \cdot x_0) - \sum_{i \in [m]} \lambda_i \times \max_{x_i \in P_i} (\alpha \cdot x_i) - \sum_{i \in [m]} \lambda_i \cdot \beta \\ &= \min_{x_0 \in P_0} (\alpha \cdot x_0) - \sum_{i \in [m]} \lambda_i \times \max_{x_i \in P_i} (\alpha \cdot x_i) \end{aligned}$$

Note that the above derivation holds for any $\lambda \in \Delta_m$. Thus, the α from the separating hyperplane definition already satisfies $\min_{\lambda \in \Delta_m} \mathcal{L}(\alpha; \lambda) \geq 0$. Therefore, $\max_{\alpha: \|\alpha\|_2 \leq 1} \min_{\lambda \in \Delta_m} \mathcal{L}(\alpha; \lambda) \geq 0$, as desired.

It remains to prove the “only if” direction. Assume that the max-min problem has optimal objective greater or equal to 0. We must show that there exists a hyperplane that separates P_0 from all P_i for $i \in [m]$. Let (α^*, λ^*) be an optimal solution to the problem $\max_{\alpha: \|\alpha\|_2 \leq 1} \min_{\lambda \in \Delta_m} \mathcal{L}(\alpha; \lambda)$, and set $\beta^* := -\min_{x_0 \in P_0} (\alpha^* \cdot x_0)$. We argue that the hyperplane $\alpha^* x + \beta^* = 0$ separates P_0 from every P_i with $i \in [m]$.

A key observation about (α^*, λ^*) – being a max-min solution to $\max_{\alpha: \|\alpha\|_2 \leq 1} \min_{\lambda \in \Delta_m} \mathcal{L}(\alpha; \lambda)$ – is that $\lambda_i^* > 0$ only if $\max_{x_i \in P_i} (\alpha^* \cdot x_i) = \max_j \{ \max_{x_j \in P_j} (\alpha^* \cdot x_j) \}$ – i.e., it takes the maximum value among the coefficients of λ_i^* . Consequently, for any $\bar{x}_j \in P_j$ for some $j \in [m]$, we have:

$$\begin{aligned}
\alpha^* \bar{x}_j + \beta^* &= \alpha^* \bar{x}_j - \min_{x_0 \in P_0} (\alpha^* \cdot x_0) \\
&= \alpha^* \bar{x}_j - \left[\sum_{i \in [m]} \lambda_i^* \times \max_{x_i \in P_i} (\alpha^* \cdot x_i) + \mathcal{L}(\alpha^*; \lambda^*) \right] \\
&= \alpha^* \bar{x}_j - \max_i \left\{ \max_{x_i \in P_i} (\alpha^* \cdot x_i) \right\} - \mathcal{L}(\alpha^*; \lambda^*) \\
&\leq \left[\alpha^* \bar{x}_j - \max_{x_j \in P_j} (\alpha^* \cdot x_j) \right] - \mathcal{L}(\alpha^*; \lambda^*) \\
&\leq 0,
\end{aligned}$$

where the last transition holds since the choice of $\bar{x}_j \in P_j$ implies $\alpha^* \bar{x}_j - \max_{x_j \in P_j} (\alpha^* \cdot x_j) \leq 0$, whereas $\mathcal{L}(\alpha^*; \lambda^*) \geq 0$ is by assumption.

On the other hand, for any $\bar{x}_0 \in P_0$, we have the following:

$$\alpha^* \bar{x}_0 + \beta^* = \alpha^* \bar{x}_0 - \min_{x_0 \in P_0} (\alpha^* \cdot x_0) \geq 0,$$

as desired.

B.2 Proof of Proposition 4.5

By Proposition 2.4, one can restrict attention to cheap talk equilibria in which the sender uses finitely many signals. Note that we can assume w.l.o.g. that there is no state $\omega \in \Omega$ with $u_S(\omega, a_1) = u_S(\omega, a_2)$ and $u_R(\omega, a_1) = u_R(\omega, a_2)$. Indeed, the utility of both agents in such a state is independent of π and s , implying that such a state can be ignored in the equilibrium analysis. Let Ω^1, Ω^2 and $\Omega^{1,2}$ be, respectively, the sets of states in which the sender strictly prefers action a_1 over action a_2 , strictly prefers a_2 over a_1 , and is indifferent between the two actions.

Consider a sender-optimal equilibrium (π, s) with $\text{supp}(\pi)$ being the smallest among all possibilities for a sender-optimal equilibrium. Note that no two signals $\sigma^1, \sigma^2 \in \text{supp}(\pi)$ might have $s[\sigma^1] = s[\sigma^2]$, as replacing them by a single signal leads to an equivalent equilibrium in terms of utilities, contradicting the minimality of $\text{supp}(\pi)$. Indeed, given two such signals σ^1 and σ^2 – one can define a new strategy profile (π', s') s.t.: (i) π' is obtained from π by replacing σ^2 with σ^1 ; (ii) $s' = s|_{\text{supp}(\pi')}$ is obtained from s by setting $s'[\sigma] = s[\sigma]$ for $\sigma \neq \sigma^2$. Since the posterior $p_{\sigma^1}^{\pi'}$ is a convex combination of the posteriors $p_{\sigma^1}^{\pi}$ and $p_{\sigma^2}^{\pi}$, we get that $s[\sigma^1]$ is a best response to getting the signal σ_1 under π' .¹⁵ Moreover, as the sender's expected utility upon transmitting any signal under π' is the same as under π , and the receiver's mixed strategy upon getting a specific signal under π' is the same as upon getting this signal under π , we get that π' is a best response to s' . To sum up, (π', s') is a cheap talk equilibrium that has the same expected sender's and receiver's utilities as (π, s) , but $\text{supp}(\pi') < \text{supp}(\pi)$, a contradiction.

Let $\sigma_1 \in \text{supp}(\pi)$ be the signal upon getting which the receiver takes action a_1 with the highest probability over all signals in $\text{supp}(\pi)$, say with probability q_1 . Note that σ_1 is uniquely defined by the previous paragraph. Similarly, let σ_2 be the unique signal resulting in the receiver taking action a_2 with the highest probability among the signals in $\text{supp}(\pi)$, say with probability q_2 .

¹⁵Here we use the well-known fact that the set of posteriors for which a specific action is (weakly) best for the receiver is convex (see, e.g., [KG11]).

We claim that for any $\omega \in \Omega^1$, $\pi(\sigma_1 | \omega) = 1$. Indeed, if that condition is not satisfied for some $\omega_i \in \Omega^1$, then the sender will better off deviating from π to $\tilde{\pi}$ with $\tilde{\pi}[\omega] = \pi[\omega]$ for $\omega \neq \omega_i$ and $\tilde{\pi}(\sigma_1 | \omega_i) = 1$, a contradiction. Thus, σ_1 is sent with probability 1 for any $\omega \in \Omega^1$. Similarly, σ_2 is sent with probability 1 for any $\omega \in \Omega^2$.

Consider some $\omega \in \Omega^{1,2}$ and let σ_3 be some signal sent with a positive probability in the state ω . If $\sigma_3 \neq \sigma_1, \sigma_2$, then σ_3 is sent with a positive probability exclusively for $\omega \in \Omega^{1,2}$. Note that it is w.l.o.g. to assume $|\text{supp}(s[\sigma_3])| = 1$, since if both agents are indifferent between both actions upon σ_3 , then making the receiver choose action a_2 upon observing σ_3 does not break the equilibrium conditions, and the expected utilities of both agents are unchanged. Assume, w.l.o.g., that $\text{supp}(s[\sigma_3]) = \{a_2\}$. Then sending σ_2 whenever the sender should have transmitted σ_3 (under π) still yields a posterior for which a_2 is a best response for the receiver (see Footnote 15). Therefore, modifying π by sending σ_2 whenever one should have sent σ_3 yields a sender-optimal equilibrium using a smaller number of signals, a contradiction. Therefore, σ_1 and σ_2 are the only signals used by the sender.

Moreover, let $\Omega_2^1 := \{\omega \in \Omega^{1,2} : u_R(\omega, a_1) > u_R(\omega, a_2)\}$ and $\Omega_1^2 := \Omega^{1,2} \setminus \Omega_2^1$. Then, by the discussion above, σ_1 is sent with probability 1 for $\omega \in \Omega^1 \cup \Omega_2^1$ and σ_2 is sent with probability 1 for $\omega \in \Omega^2 \cup \Omega_1^2$. Note that the expected sender's utility is linear in q_1, q_2 and non-decreasing in each of them.

Recall that in each state, the signal sent is chosen deterministically. Therefore, we are only left with the receiver's incentive-compatibility constraints. Let us normalize the receiver's utility for action a_2 to be 0 in any state.¹⁶ Note that if $q_1 > 0$ and $q_2 > 0$, we must have: $\sum_{\omega \in \Omega^1 \cup \Omega_2^1} u_R(\omega, a_1) \cdot \mu(\omega) \geq 0$ and $\sum_{\omega \in \Omega^2 \cup \Omega_1^2} u_R(\omega, a_1) \cdot \mu(\omega) \leq 0$. If both these conditions hold, then in the sender-optimal equilibrium we have $q_1 = q_2 = 1$. If at least one of them does not hold, then $q_1 q_2 = 0$, implying that the receiver takes one action regardless of the signal. In this case, the sender cannot benefit from cheap talk, and the sender-optimal equilibrium is a babbling equilibrium using just one signal.

¹⁶It is possible since subtracting the same number from $u_R(\omega, a_1)$ and $u_R(\omega, a_2)$ for a fixed $\omega \in \Omega$ does not affect the set of equilibria and sender's preferences over them.