Honest Signaling in Zero-Sum Games Is Hard, and Lying Is Even Harder*†

Aviad Rubinstein

UC Berkeley, Berkeley, CA, USA
aviad@eecs.berkeley.edu

Abstract

We prove that, assuming the exponential time hypothesis, finding an ϵ-approximately optimal signaling scheme in a two-player zero-sum game requires quasi-polynomial time $(n^{\tilde{\Omega}(\lg n)})$. This is tight by [8] and resolves an open question of Dughmi [12]. We also prove that finding a multiplicative approximation is NP-hard.

We also introduce a new model where a dishonest signaler may publicly commit to use one scheme, but post signals according to a different scheme. For this model, we prove that even finding a $(1 - 2^{-n})$-approximately optimal scheme is NP-hard.

1998 ACM Subject Classification F.2 Analysis of Algorithms and Problem Complexity

Keywords and phrases Signaling, Zero-sum Games, Algorithmic Game Theory, birthday repetition

Digital Object Identifier 10.4230/LIPIcs.ICALP.2017.77

1 Introduction

Many classical questions in economics involve extracting information from strategic agents. Lately, there has been growing interest within algorithmic game theory in *signaling*: the study of how to *reveal information* to strategic agents (see e.g. [16, 13, 14, 12, 8] and references therein). Signaling has been studied in many interesting economic and game theoretic settings. Among them, Zero-Sum Signaling proposed by Dughmi [12] stands out as a canonical problem that cleanly captures the computational nature of signaling. In particular, focusing on zero-sum games clears away issues of equilibrium selection and computational tractability of finding an equilibrium.

▶ **Definition 1** (Zero-Sum Signaling [12])**.** Alice and Bob play a Bayesian zero-sum game where the payoff matrix M is drawn from a publicly known prior. The signaler Sam privately observes the state of nature (i.e. the payoff matrix), and then publicly broadcasts a signal $\varphi(M)$ to both Alice and Bob. Alice and Bob Bayesian-update their priors according to $\varphi(M)$'s and play the Nash equilibrium of the expected game; but they receive payoffs according to the true M. Sam's goal is to design an efficient signaling scheme φ (a function from payoff matrices to strings) that maximizes Alice's expected payoff.

Dughmi's [12] main result proves that assuming the hardness of the Planted Clique problem, there is no additive FPTAS for Zero-Sum Signaling. The main open question

* A full version of the paper is available at http://arxiv.org/abs/1510.04991.

† This research was supported by Microsoft Research PhD Fellowship, as well as NSF grant CCF1408635 and by Templeton Foundation grant 3966. This work was done in part at the Simons Institute for the Theory of Computing.

44th International Colloquium on Automata, Languages, and Programming (ICALP 2017).
Editors: Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl;
Article No. 77; pp. 77:1–77:13

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

left by [12] is whether there exists an additive PTAS. Here we answer this question in the negative: we prove that assuming the Exponential Time Hypothesis (ETH) [15], obtaining an additive-ϵ-approximation (for some constant $\epsilon > 0$) requires quasi-polynomial time $(n^{\tilde{\Omega}(\lg n)})$. This result is tight thanks to a recent quasi-polynomial $(n^{\frac{\lg n}{\text{poly}(\epsilon)}})$ time algorithm by Cheng et al. [8]. Another important advantage of our result is that it replaces the hardness of Planted Clique with a more believable worst-case hardness assumption (see e.g. the discussion in [7]).

▶ **Theorem 2** (Main Result). *There exists a constant $\epsilon > 0$, such that assuming ETH, approximating* Zero-Sum Signaling *with payoffs in $[-1, 1]$ to within an additive ϵ requires time $n^{\tilde{\Omega}(\lg n)}$.*

Using a similar construction, we also obtain NP-hardness for computing a multiplicative-$(1-\epsilon)$-approximation. Unfortunately, in our example Alice can receives both negative and positive payoffs, which is somewhat non-standard (but not unprecedented [9]) in multiplicative approximation. One main reason that multiplicative approximation with negative payoffs is problematic is that this is often trivially intractable for any finite factor: Start with a tiny additive gap, where Alice's expected payoff is c in the "yes" case, and $s = c - \epsilon$ in the "no" case; subtract $(c + s)/2$ from all of Alice's payoffs to obtain an infinite multiplicative hardness. We note, however, that the combination of negative and positive payoffs in our construction serves only to obtain structural constraints on the resulting equilibria; the hardness of approximation is not a result of cancellation of negative with positive payoffs: Alice's payoff can be decomposed as a difference of non-negative payoffs $U = U^+ - U^-$, such that it is hard to approximate Alice's optimal payoff to within $\epsilon \cdot \mathsf{E}[U^+ + U^-]$. Nevertheless, we believe that extending this result to non-negative payoffs could be very interesting.

▶ **Theorem 3.** *There exists a constant $\epsilon > 0$, such that it is* NP*-hard to approximate* Zero-Sum Signaling *to within a multiplicative $(1-\epsilon)$ factor.*

Finally, we note that since all our games are zero-sum, the hardness results for Zero-Sum Signaling also apply to the respective notions of additive- and multiplicative-ϵ-Nash equilibrium.

1.1 The computational complexity of lying

As a motivating example, consider the purchase of a used car (not a zero-sum game, but a favorite setting in the study of signaling since Akerlof's seminal "Market for Lemons" [2]), and let us focus on the information supplied by a third party such as a mechanic inspection. The mechanic (Sam) publishes a signaling scheme: report any problem found in a one-hour inspection. Unbeknownst to the buyer (Bob), the mechanic favors the seller (Alice), and chooses to use a different signaling scheme: always report that the car is in excellent condition. Notice that it is crucial that the buyer does not know that the mechanic is lying (and more generally, we assume that neither party knows that the signaler is lying).

Much of the work in economics is motivated by selfish agents manipulating their private information. Here we introduce a natural extension of Dughmi's signaling model, where the signaler manipulates his private information. We formalize this extension in the Zero-Sum Lying problem, where the signaling scheme consists of two functions φ_{ALLEGED} ("report any problem found") and φ_{REAL} ("car is in excellent condition") from payoff matrices to signals. Sam promises Alice and Bob to use φ_{ALLEGED}, which is what Alice and Bob use to compute the posterior distribution given the signal (i.e. the seller and buyer look at the mechanic's

report and negotiate a price as if the state of the car is correctly reflected). But instead Sam signals according to φ_{REAL}.

We formally define the Zero-Sum Lying problem below; notice that the original Zero-Sum Signaling (Definition 1) corresponds to the special case where we restrict $\varphi_{\text{REAL}} = \varphi_{\text{ALLEGED}}$.

▶ **Definition 4** (Zero-Sum Lying)**.** Alice and Bob play a Bayesian, one-shot, zero-sum game where the payoff matrix is drawn from a publicly known prior. A *dishonest signaling scheme* consists of two (possibly randomized) functions $\varphi_{\text{ALLEGED}}, \varphi_{\text{REAL}}$ from payoff matrices to signals, that induce the following protocol:

- Nature draws a private payoff matrix $M \sim \mathcal{D}_{\text{NATURE}}$.
- Alice and Bob observe the scheme φ_{ALLEGED} and the signal $\sigma \triangleq \varphi_{\text{REAL}}(M)$. (But they don't know the scheme φ_{REAL}!)
- Alice and Bob choose a Nash equilibrium $(\mathbf{x}; \mathbf{y})$ for the zero-sum game with payoff matrix $\mathsf{E}[M' \mid \varphi_{\text{ALLEGED}}(M') = \sigma]$[1].
  - (We assume that the support of φ_{REAL} is contained in the support of φ_{ALLEGED}.)
- Alice and Bob receive payoffs $\mathbf{x}^\top M \mathbf{y}$ and $-\mathbf{x}^\top M \mathbf{y}$, respectively.

Sam's goal is to compute a pair $(\varphi_{\text{ALLEGED}}, \varphi_{\text{REAL}})$ that maximizes Alice's expected payoff.

In the toy-setting of a biased car inspection, the Sam's optimal strategy was very simple. In contrast, we show that for a general distribution over zero-sum games, it is NP-hard to find a pair $(\varphi_{\text{ALLEGED}}, \varphi_{\text{REAL}})$ that is even remotely close to optimal. Notice that this is very different from the honest case where, as we mentioned earlier, NP-hardness of additive approximation is unlikely given the additive quasi-PTAS of [8].

▶ **Theorem 5.** *Approximating* Zero-Sum Lying *with Alice's payoffs in* $[0, 1]$ *to within an additive* $(1 - 2^{-n})$ *is* NP*-hard.*

Further discussion of dishonest signaling

It is important to note that the dishonest signaling model has a few weaknesses:

- Alice and Bob must believe the dishonest signaler. (See also further discussion below.)
- In particular, Sam cheats in favor of Alice, but Alice doesn't know about it – so what's in it for Sam? Indeed, we assume that Sam has some intrinsic interest in Alice winning, e.g. because Sam loves Alice or owns some of her stocks.
- The game for which players' strategies are at equilibrium may be very different from the actual game. Note, however, that this is also the case for the honest signaling model (when the signaling scheme is not one-to-one).
- The players may receive different payoffs for different equilibria; this may raise issues of equilibrium selection.

Despite those disadvantages, we believe that our simple model is valuable because it already motivates surprising results such as our Theorem 5. On a higher level, we hope that it will inspire research on many other interesting aspects on dishonest signaling. For example, notice that in our model Sam lies without any reservation; if, per contra, the game was repeated infinitely many times, one would expect that Alice and Bob will eventually stop

[1] When $\varphi_{\text{ALLEGED}}, \varphi_{\text{REAL}}$ are randomized, we have $\sigma \sim \varphi_{\text{REAL}}(M)$ and expectation conditioned on $\mathsf{E}\left[M' \mid \sigma \sim \varphi_{\text{ALLEGED}}\left(M'\right)\right]$.

believing the signals, hence only honest signaling is possible. There is also a spectrum of intermediate situations, where Alice and Bob observe some partial information about past games (e.g. marginal distribution of signals) and may encounter questions about distribution testing.

Another related direction of potential future research is to think about Sam's incentives. When is honest signaling optimal for Sam? When is it approximately optimal? How should one design an effective "punishing" mechanism?

1.2 Concurrent work of Bhaskar et al.

In independent concurrent work by Bhaskar et al. [5], quasi-polynomial time hardness for additive approximation of Zero-Sum Signaling was obtained assuming the hardness of the Planted Clique problem (among other interesting results[2] about network routing games and security games). Although we are not aware of a formal reduction, hardness of Planted Clique is a qualitatively stronger assumption than ETH in the sense that it requires average case instances to be hard. Hence in this respect, our result is stronger.

1.3 Techniques

Our main ingredient for the quasi-polynomial hardness is the technique of "birthday repetition" coined by [1] and recently applied in game theoretic settings in [7, 4]: We reduce from a 2-ary constraint satisfaction problem (2-CSP) over n variables to a distribution over N zero-sum $N \times N$ games, with $N = 2^{\Theta(\sqrt{n})}$. Alice and Bob's strategies correspond to assignments to tuples of $\sqrt{n}$ variables. By the birthday paradox, the two $\sqrt{n}$-tuples chosen by Alice and Bob share a constraint with constant probability. If a constant fraction of the constraints are unsatisfiable, Alice's payoff will suffer with constant probability. Assuming ETH, approximating the value of the CSP requires time $2^{\tilde{\Omega}(n)} = N^{\tilde{\Omega}(\lg N)}$.

1.3.0.1 The challenge

The main difficulty is that once the signal is public, the zero-sum game is tractable. Thus we would like to force the signaling scheme to output a satisfying assignment. Furthermore, if the scheme would output partial assignments on different states of nature (aka different zero-sum games in the support), it is not clear how to check consistency between different signals. Thus we would like each signal to contain an entire satisfying assignment. The optimal scheme may be very complicated and even require randomization, yet by an application of the Caratheodory Theorem the number of signals is, wlog, bounded by the number of states of nature [12]. If the state of nature can be described using only $\lg N = \tilde{\Theta}(\sqrt{n})$ bits[3], how can we force the scheme to output an entire assignment?

To overcome this obstacle, we let the state of nature contain a partial assignment to a random $\sqrt{n}$-tuple of variables. We then check the consistency of Alice's assignment with nature's assignment, Bob's assignment with nature's assignment, and Alice and Bob's assignments with each other; let $\tau^{A,Z}, \tau^{B,Z}, \tau^{A,B}$ denote the outcomes of those consistency checks, respectively. Alice's payoff is given by:

$$U = \delta\tau^{A,Z} - \delta^2\tau^{B,Z} + \delta^3\tau^{A,B}$$

[2] For zero-sum games, Bhaskar et al. also rule out an additive FPTAS assuming $\mathsf{P} \neq \mathsf{NP}$. This result follows immediately from our Theorem 14.

[3] In other words, N, the final size of the reduction, is an upper bound on the number of states of nature.

for some small constant $\delta \in (0,1)$. Now, both Alice and Bob want to maximize their chances of being consistent with nature's partial assignment, and the signaling scheme gains by maximizing $\tau^{A,B}$.

Of course, if nature outputs a random assignment, we have no reason to expect that it can be completed to a full satisfying assignment. Instead, the state of nature consists of N assignments, and the signaling scheme helps Alice and Bob play with the assignment that can be completed.

Several other obstacles arise; fortunately some can be handled using techniques from previous works on hardness of finding Nash equilibrium [3, 10, 4].

2 Preliminaries

Exponential Time Hypothesis

▶ **Hypothesis 6** (Exponential Time Hypothesis (ETH) [15]). *3SAT takes time* $2^{\Omega(n)}$.

PCP Theorem and CSP

▶ **Definition 7** (2CSP). 2-CSP (2-ary Constraint Satisfaction Problem) is a maximization problem. The input is a graph $G=(V,E)$, alphabet Σ, and a constraint $C_e \subseteq \Sigma \times \Sigma$ for every $e \in E$.

The output is a labeling $\varphi : V \to \Sigma$ of the vertices. Given a labeling, we say that a constraint (or edge) $(u,v) \in E$ is *satisfied* if $\varphi(u), \varphi(v) \in C_{(u,v)}$. The *value of a labeling* is the fraction of $e \in E$ that are satisfied by the labeling. The value of the instance is the maximum fraction of constraints satisfied by any assignment.

▶ **Theorem 8** (PCP Theorem [11]; see e.g. [6, Theorem 2.11] for this formulation). *Given a 3SAT instance ϕ of size n, there is a polynomial time reduction that produces a 2CSP instance ψ, with size $|\psi| = n \cdot \mathrm{polylog} n$ variables and constraints, and constant alphabet size, such that:*

Completeness. *If ϕ is satisfiable, then so is ψ.*

Soundness. *If ϕ is not satisfiable, then at most a $(1-\eta)$-fraction of the constraints in ψ can be satisfied, for some $\eta = \Omega(1)$.*

Balance. *Every variable in ψ participates in exactly $d = O(1)$ constraints.*

Finding a good partition

▶ **Lemma 9** (Essentially [4, Lemma 6]). *Let $G=(V,E)$ be a d-regular graph and $n \triangleq |V|$. We can partition V into n/k disjoint subsets $\{S_1, \ldots, S_{n/k}\}$ of size at most $2k$ such that:*

$$\forall i,j \quad |(S_i \times S_j) \cap E| \leq 8d^2k^2/n\,. \tag{1}$$

See full version for proof [17].

How to catch a far-from-uniform distribution

The following lemma due to [10] implies that:

▶ **Lemma 10** (Lemma 3 in the full version of [10]). *Let $\{a_i\}_{i=1}^n$ be real numbers satisfying the following properties for some $\theta > 0$: (1) $a_1 \geq a_2 \geq \cdots \geq a_n$; (2) $\sum a_i = 0$; (3) $\sum_{i=1}^{n/2} a_i \leq \theta$. Then $\sum_{i=1}^n |a_i| \leq 4\theta$.*

3 Additive hardness

▶ **Theorem 11.** *There exists a constant $\epsilon > 0$, such that assuming ETH, approximating* Zero-Sum Signaling *with payoffs in $[-1, 1]$ to within an additive ϵ requires time $n^{\tilde{\Omega}(\lg n)}$.*

Construction overview

Our reduction begins with a 2CSP ψ over n variables from alphabet Σ. We partition the variables into n/k disjoint subsets $\{S_1, \ldots, S_{n/k}\}$, each of size at most $2k$ for $k = \sqrt{n}$ such that every two subsets share at most a constant number of constraints.

Nature chooses a random subset S_i from the partition, a random assignment $\vec{u} \in \Sigma^{2k}$ to the variables in S_i, and an auxiliary vector $\hat{b} \in \{0,1\}^{\Sigma \times [2k]}$. As mentioned in Section 1.3, $\vec{u}$ may not correspond to any satisfying assignment. Alice and Bob participate in one of $|\Sigma|^{2k}$ subgames; for each $\vec{v} \in \Sigma^{2k}$, there is a corresponding subgame where all the assignments are XOR-ed with $\vec{v}$. The goal of the auxiliary vector $\hat{b}$ is to force Alice and Bob to participate in the right subgame, i.e. the one where the XOR of $\vec{v}$ and $\vec{u}$ can be completed to a full satisfying assignment. In particular, the optimum signaling scheme reveals partial information about $\hat{b}$ in a way that guides Alice and Bob to participate in the right subgame. The scheme also outputs the full satisfying assignment, but reveals no information about the subset S_i chosen by nature.

Each player has $\left(|\Sigma|^{2k} \times 2\right) \times \left(n/k \times \binom{n/k}{n/2k} \times |\Sigma|^{2k}\right) = 2^{\Theta\left(\sqrt{n}\right)}$ strategies. The first $|\Sigma|^{2k}$ strategies correspond to a Σ-ary vector $\vec{v}$ that the scheme will choose after observing the random input. The signaling scheme forces both players to play (w.h.p.) the strategy corresponding to $\vec{v}$ by controlling the information that corresponds to the next 2 strategies. Namely, for each $\vec{v}' \in \Sigma^{2k}$, there is a random bit $b(\vec{v}')$ such that each player receives a payoff of 1 if they play $(\vec{v}', b(\vec{v}'))$ and 0 for $(\vec{v}', 1 - b(\vec{v}'))$. The b's are part of the state of nature, and the optimal signaling scheme will reveal only the bit corresponding to the special $\vec{v}$. Since there are $|\Sigma|^{2k}$ bits, nature cannot choose them independently, as that would require $2^{|\Sigma|^{2k}}$ states of nature. Instead we construct a pairwise independent distribution.

The next n/k strategies correspond to the choice of a subset S_i from the specified partition of variables. The $\binom{n/k}{n/2k}$ strategies that follow correspond to a gadget due to Althofer [3] whereby each player forces the other player to randomize (approximately) uniformly over the choice of subset.

The last $|\Sigma|^{2k}$ strategies correspond to an assignment to S_i. The assignment to each S_i is XOR-ed entry-wise with $\vec{v}$. Then, the players are paid according to checks of consistency between their assignments, and a random assignment to a random S_i picked by nature. (The scheme chooses $\vec{v}$ so that nature's random assignment is part of a globally satisfying assignment.) Each player wants to pick an assignment that passes the consistency check with nature's assignment. Alice also receives a small bonus if her assignment agrees with Bob's; thus her payoff is maximized when there exists a globally satisfying assignment.

See formal construction below, or refer to summary table in full version [17].

Formal construction

Let ψ be a 2CSP-d over n variables from alphabet Σ, as guaranteed by Theorem 8. In particular, ETH implies that distinguishing between a completely satisfiable instance and $(1-\eta)$-satisfiable requires time $2^{\tilde{\Omega}(n)}$. By Lemma 9, we can (deterministically and efficiently) partition the variables into n/k subsets $\{S_1, \ldots, S_{n/k}\}$ of size at most $2k = 2\sqrt{n}$, such that every two subsets share at most $8d^2k^2/n = O(1)$ constraints.

States of nature. Nature chooses a state $\left(\hat{b}, i, \vec{u}\right) \in \{0,1\}^{\Sigma\times[2k]} \times [n/k] \times \Sigma^{2k}$ uniformly at random. For each $\vec{v}$, $b(\vec{v})$ is the XOR of bits from $\hat{b}$ that correspond to entries of $\vec{v}$:

$$\forall \vec{v} \in \Sigma^{2k} \;\; b(\vec{v}) \triangleq \left(\bigoplus_{(\sigma,\ell):\ [\vec{v}]_\ell = \sigma} \left[\hat{b}\right]_{(\sigma,\ell)} \right).$$

Notice that the $b(\vec{v})$'s are pairwise independent and each marginal distribution is uniform over $\{0,1\}$.

Strategies. Alice and Bob each choose a strategy $(\vec{v}, c, j, T, \vec{w}) \in \Sigma^{2k} \times \{0,1\} \times [n/k] \times \binom{[n/k]}{n/2k} \times \Sigma^{2k}$. We use $\vec{v}^A, c^A$, etc. to denote the strategy Alice plays, and similarly $\vec{v}^B, c^B$, etc. for Bob. For $\sigma, \sigma' \in \Sigma$, we denote $\sigma \oplus_\Sigma \sigma' \triangleq \sigma + \sigma' \pmod{|\Sigma|}$, and for vectors $\vec{v}, \vec{v}' \in \Sigma^{2k}$, we let $\vec{v} \oplus_\Sigma \vec{v}' \in \Sigma^{2k}$ denote the entry-wise $\oplus_\Sigma$.

Payoffs. Consider state of nature $\left(\hat{b}, i, \vec{u}\right)$ and players' strategies $\left(\vec{v}^A, c^A, j^A, T^A, \vec{w}^A\right)$ and $\left(\vec{v}^B, c^B, j^B, T^B, \vec{w}^B\right)$.

When $\vec{v}^A = \vec{v}^B = \vec{v}$, we set $\tau^{A,Z} = 1$ if assignments $\vec{w}^A$ and $(\vec{v} \oplus_\Sigma \vec{u})$ to subsets S_{j^A} and S_i, respectively, satisfy all the constraints in ψ that are determined by $\left(S_i \cup S_{j^A}\right)$, and $\tau^{A,Z} = 0$ otherwise. Similarly, $\tau^{B,Z} = 1$ iff $\vec{w}^B$ and $(\vec{v} \oplus_\Sigma \vec{u})$ satisfy the corresponding constraints in ψ; and $\tau^{A,B}$ checks $\vec{w}^A$ and $\vec{w}^B$. When $\vec{v}^A \neq \vec{v}^B$, we set $\tau^{A,Z} = \tau^{B,Z} = \tau^{A,B} = 0$.

We decompose Alice's payoff as:

$$U^A \triangleq U^A_b + U^A_{\text{Althofer}} + U^A_\psi,$$

where

$$U^A_b \triangleq \mathbf{1}\left\{c^A = b\left(\vec{v}^A\right)\right\} - \mathbf{1}\left\{c^B = b\left(\vec{v}^B\right)\right\},$$

$$U^A_{\text{Althofer}} \triangleq \mathbf{1}\left\{j^B \in T^A\right\} - \mathbf{1}\left\{j^A \in T^B\right\},$$

and

$$U^A_\psi \triangleq \delta\tau^{A,Z} - \delta^2\tau^{B,Z} + \delta^3\tau^{A,B}, \tag{2}$$

for a sufficiently small constant $0 < \delta \ll \sqrt{\eta}$.

Completeness

▶ **Lemma 12.** *If ψ is satisfiable, there exists a signaling scheme and a mixed strategy for Alice that guarantees expected payoff $\delta - \delta^2 + \delta^3$.*

Proof. Fix a satisfying assignment $\vec{\alpha} \in \Sigma^n$. Given state of nature $\left(\hat{b}, i, \vec{u}\right)$, let $\vec{v}$ be such that $(\vec{v} \oplus_\Sigma \vec{u}) = [\vec{\alpha}]_{S_i}$. The scheme outputs the signal $(\vec{v}, b(\vec{v}), \vec{\alpha})$. Alice's mixed strategy sets $\left(\vec{v}^A, c^A\right) = (\vec{v}, b(\vec{v}))$, picks j^A and T^A uniformly at random, and sets $\vec{w}^A = [\vec{\alpha}]_{S_{j^A}}$.

Because Bob has no information about $b(\vec{v}')$ for any $\vec{v}' \neq \vec{v}$, he has probability $1/2$ of losing whenever he picks $\vec{v}^B \neq \vec{v}$, i.e. $\mathsf{E}\left[U^A_b\right] \geq \frac{1}{2}\Pr\left[\vec{v}^B \neq \vec{v}\right]$. Furthermore, because Alice chooses T^A and j^A uniformly, $\mathsf{E}\left[U^A_{\text{Althofer}}\right] = 0$.

Since $\vec{\alpha}$ completely satisfies ψ, we have that $\tau^{A,Z} = 1$ as long as $\vec{v}^B = \vec{v}$ (regardless of the rest of Bob's strategy). Bob's goal is thus to maximize $\mathsf{E}\left[\delta^2\tau^{B,Z} - \delta^3\tau^{A,B}\right]$. Notice that $\vec{w}^A$ and $(\vec{v} \oplus_\Sigma \vec{u})$ are two satisfying partial assignments to uniformly random subsets from the

partition. In particular, they are both drawn from the same distribution, so we have that for any mixed strategy that Bob plays, $\mathsf{E}\left[\tau^{B,Z}\right] = \mathsf{E}\left[\tau^{A,B}\right]$. Therefore Alice's payoff is at least

$$\left(\delta - \delta^2 + \delta^3\right) \Pr\left[\vec{v}^B = \vec{v}\right] + \frac{1}{2} \Pr\left[\vec{v}^B \neq \vec{v}\right] \geq \delta - \delta^2 + \delta^3. \qquad \blacktriangleleft$$

Soundness

▶ **Lemma 13.** *If at most a $(1-\eta)$-fraction of the constraints are satisfiable, Alice's maxmin payoff is at most $\delta - \delta^2 + (1 - \Omega_\eta(1))\,\delta^3$, for any signaling scheme.*

Proof. Fix any mixed strategy by Alice; we show that Bob can guarantee a payoff of at least $-\left(\delta - \delta^2 + (1 - \Omega_\eta(1))\,\delta^3\right)$. On any signal, Bob chooses $\left(\vec{v}^B, c^B\right)$ from the same distribution that Alice uses for $\left(\vec{v}^A, c^A\right)$. He chooses j^B uniformly, and picks T^B so as to minimize $\mathsf{E}\left[U^A_{\text{Althofer}}\right]$. Finally, for each j^B, he draws $\vec{w}^B$ from the same marginal distribution that Alice uses for $\vec{w}^A$ conditioning on $j^A = j^B$ (and uniformly at random if Alice never plays $j^A = j^B$). By symmetry, $\mathsf{E}\left[U^A_b\right] = 0$ and $\mathsf{E}\left[U^A_{\text{Althofer}}\right] \leq 0$.

In this paragraph, we use Althoefer's gadget to argue that, wlog, Alice's distribution over the choice of j^A is approximately uniform. In Althofer's gadget, Alice can guarantee an (optimal) expected payoff of 0 by randomizing uniformly over her choice of j^A and T^A. By Lemma 10, if Alice's marginal distribution over the choice of j^A is $8\delta^2$-far from uniform (in total variation distance), then Bob can guess that j^A is in some subset $T^B \in \binom{[n/k]}{n/2k}$ with advantage (over guessing at random) of at least $2\delta^2$. Therefore $\mathsf{E}\left[U^A_{\text{Althofer}}\right] \leq -2\delta^2$; but this would imply $\mathsf{E}\left[U^A\right] \leq -2\delta^2 + \mathsf{E}\left[U^A_\psi\right] \leq \delta - 2\delta^2 + \delta^3$. So henceforth we assume wlog that Alice's marginal distribution over the choice of j^A is $O\left(\delta^2\right)$-close to uniform (in total variation distance).

Since Alice's marginal distribution over j^A is $O\left(\delta^2\right)$-close to uniform, we have that Bob's distribution over $\left(j^B, \vec{w}^B\right)$ is $O\left(\delta^2\right)$-close to Alice's distribution over $\left(j^A, \vec{w}^A\right)$. Therefore $\mathsf{E}\left[\tau^{B,Z}\right] \geq \mathsf{E}\left[\tau^{A,Z}\right] - O\left(\delta^2\right)$, and so we also get:

$$\mathsf{E}\left[U^A\right] \leq \mathsf{E}\left[U^A_\psi\right] \leq \delta - \delta^2 + \delta^3 \mathsf{E}\left[\tau^{A,B}\right] + O\left(\delta^4\right). \tag{3}$$

Bounding $\mathsf{E}\left[\tau^{A,B}\right]$. To complete the proof, it remains to show an upper bound on $\mathsf{E}\left[\tau^{A,B}\right]$. In particular, notice that it suffices to bound the probability that Alice's and Bob's induced assignments agree. Intuitively, if they gave assignments to uniformly random (and independent) subsets of variables, the probability that their assignments agree cannot be much higher than the value of the 2CSP; below we formalize this intuition.

By the premise, any assignment to all variables violates at least an η-fraction of the constraints. In particular, this is true in expectation for assignments drawn according to Alice's and Bob's mixed strategy. This is a bit subtle: in general, it is possible that Alice's assignment alone doesn't satisfy many constraints and neither does Bob's, but when we check constraints between Alice's and Bob's assignments everything satisfied (for example, think of the 3-Coloring CSP, where Alice colors all her vertices blue, and Bob colors all his vertices red). Fortunately, this subtlety is irrelevant for our construction since we explicitly defined Bob's mixed strategy so that conditioned on each set S_j of variables, Alice and Bob have the same distribution over assignments.

The expected number of violations between pairs directly depends on the value of the 2CSP. To bound the *probability* of observing at least one violations, recall that every pair of subsets shares at most a constant number of constraints, so this probability is within a

constant factor of the expected number of violations. In particular, an $\Omega(\eta)$-fraction of the pairs of assignments chosen by Alice and Bob violate ψ.

Finally, Alice doesn't choose j^A uniformly at random; but her distribution is $O(\delta^2)$-close to uniform. Therefore, we have $\mathsf{E}\left[\tau^{A,B}\right] \leq 1 - \Omega(\eta) + O(\delta^2)$. Plugging into (3) completes the proof. ◀

4 Multiplicative hardness

▶ **Theorem 14.** *There exists a constant $\epsilon > 0$, such that it is* NP*-hard to approximate* Zero-Sum Signaling *to within a multiplicative $(1-\epsilon)$ factor.*

Construction overview

Our reduction begins with a 2CSP ψ over n variables from alphabet Σ.

Nature chooses a random index $i \in [n]$, a random assignment $u \in \Sigma$ for variable x_i, and an auxiliary vector $\vec{b} \in \{0,1\}^{\Sigma}$. Notice that u may not correspond to any satisfying assignment. Alice and Bob participate in one of $|\Sigma|$ subgames; for each $v \in \Sigma$, there is a corresponding subgame where all the assignments are XOR-ed with v. The optimum signaling scheme reveals partial information about $\vec{b}$ in a way that guides Alice and Bob to participate in the subgame where the XOR of v and u can be completed to a full satisfying assignment. The scheme also outputs the full satisfying assignment, but reveals no information about the index i chosen by nature.

Alice has $(|\Sigma| \times 2) \times (n \times n \times |\Sigma|) = \Theta(n^2)$ strategies, and Bob has an additional choice among n strategies (so $\Theta(n^3)$ in total). The first $|\Sigma|$ strategies correspond to a value $v \in \Sigma$ that the scheme will choose after observing the state of nature. The signaling scheme forces both players to play (w.h.p.) the strategy corresponding to v by controlling the information that corresponds to the next 2 strategies. Namely, for each $v' \in \Sigma$, there is a random bit $b(v')$ such that each player receives a small bonus if they play $(v', b(v'))$ and not $(v', 1 - b(v'))$. The b's are part of the state of nature, and the signaling scheme will reveal only the bit corresponding to the special v.

The next n strategies correspond to a choice of a variable $j \in [n]$. The n strategies that follow correspond to a hide-and-seek gadget whereby each player forces the other player to randomize (approximately) uniformly over the choice of j. For Bob, the additional n strategies induce a hide-and-seek game against nature, which serves to verify that the scheme does not reveal too much information about the state of nature (this extra verification was unnecessary in the reduction for additive inapproximability).

The last $|\Sigma|$ strategies induce an assignment for x_j. The assignment to each x_j is XOR-ed with v. Then, the players are paid according to checks of consistency between their assignments, and a random assignment to a random x_i picked by nature. (The scheme chooses v so that nature's random assignment is part of a globally satisfying assignment.) Each player wants to pick an assignment that passes the consistency check with nature's assignment. Alice also receives a small bonus if her assignment agrees with Bob's; thus her payoff is maximized when there exists a globally satisfying assignment.

Formal construction

Let ψ be a 2CSP-d over n variables from alphabet Σ, as guaranteed by Theorem 8. In particular, it is NP-hard to distinguish between ψ which is completely satisfiable, and one

where at most a $(1-\eta)$-fraction of the constraints can be satisfied. We denote $(i,j) \in \psi$ if there is a constraint over variables (x_i, x_j).

States of nature. Nature chooses a state $\left(\vec{b}, i, u\right) \in \{0,1\}^{\Sigma} \times [n] \times \Sigma$ uniformly at random.

Strategies. Alice chooses a strategy $\left(v^A, c^A, j^A, t^A, w^A\right) \in \Sigma \times \{0,1\} \times [n] \times [n] \times \Sigma$, and Bob chooses $\left(v^B, c^B, j^B, t^B, q^B, w^B\right) \in \Sigma \times \{0,1\} \times [n] \times [n] \times [n] \times \Sigma$. For $\sigma, \sigma' \in \Sigma$, we denote $\sigma \oplus_{\Sigma} \sigma' \triangleq \sigma + \sigma' \pmod{|\Sigma|}$, and for a vector $\vec{\alpha} \in \Sigma^n$ we let $(\sigma \oplus_{\Sigma} \vec{\alpha}) \in \Sigma^n$ denote the $\oplus_{\Sigma}$ of σ with each entry of $\vec{\alpha}$.

Payoffs. Consider players' strategies $\left(v^A, c^A, j^A, t^A, w^A\right)$ and $\left(v^B, c^B, j^B, t^B, q^B, w^B\right)$ and state of nature $\left(\vec{b}, i, u\right)$.

When $v^A = v^B = v$, we set $\tau^{A,Z} = 1$ if ψ contains a constraint for variables $\left(j^A, i\right)$, and the assignments w^A and $(v \oplus_{\Sigma} u)$ to those variables, respectively, satisfy this constraint, and $\tau^{A,Z} = 0$ otherwise. Similarly, $\tau^{B,Z} = 1$ iff w^B and $(v \oplus_{\Sigma} u)$ satisfy a corresponding constraint in ψ; and $\tau^{A,B}$ checks w^A with w^B. When $v^A \neq v^B$, we set $\tau^{A,Z} = \tau^{B,Z} = \tau^{A,B} = 0$.

We decompose Alice's payoff as:

$$U^A \triangleq U^A_b + U^A_{\text{seek}} + U^A_{\psi},$$

where

$$U^A_b \triangleq \mathbf{1}\left\{c^A = \left[\vec{b}\right]_{v^A}\right\}/n - \mathbf{1}\left\{c^B = \left[\vec{b}\right]_{v^B}\right\}/n,$$

$$U^A_{\text{seek}} \triangleq 2 \cdot \mathbf{1}\left\{j^B = t^A\right\} - \mathbf{1}\left\{j^A = t^B\right\} - \mathbf{1}\left\{i = q^B\right\},$$

and[4]

$$U^A_{\psi} \triangleq \delta^3 \tau^{A,Z} - \delta^4 \tau^{B,Z} + \delta^5 \tau^{A,B},$$

for a sufficiently small constant $0 < \delta \ll \sqrt{\eta}$.

Completeness

▶ **Lemma 15.** *If ψ is satisfiable, there exists a signaling scheme, such that for every signal $\mathbf{s}$ in the support, Alice can guarantee an expected payoff of $\frac{d}{n}\left(\delta^3 - \delta^4 + \delta^5\right)$.*

Notice that the *for every signal in the support* qualification is different than the corresponding Lemma 12 (and there is a similar difference between Lemma 16 and Lemma 13). Indeed, this is stronger than we need for proving Theorem 14, but will come handy in Section 5.

Proof. Fix a satisfying assignment $\vec{\alpha} \in \Sigma^n$. Given state of nature $\left(\hat{b}, i, u\right)$, let v be such that $(v \oplus_{\Sigma} u) = [\vec{\alpha}]_i$. The scheme outputs the signal $\mathbf{s} \triangleq \left(v, \vec{b}_v, \vec{\alpha}\right)$. Alice's mixed strategy sets $\left(v^A, c^A\right) = \left(v, \vec{b}_v\right)$; picks j^A and t^A uniformly at random; and sets $w^A = [\vec{\alpha}]_{j^A}$. See full version for details [17]. ◀

[4] We use $\delta^3\tau^{A,Z} - \delta^4\tau^{B,Z} + \delta^5\tau^{A,B}$ instead of $\delta^1\tau^{A,Z} - \delta^2\tau^{B,Z} + \delta^3\tau^{A,B}$ as in 2, because the square of the first coefficient appears in the proof. We have $\left(\delta^3\right)^2 \ll \delta^5$, but $\delta^2 \gg \delta^3$.

Soundness

▶ **Lemma 16.** *If at most a* $(1-\eta)$*-fraction of the constraints are satisfiable, then for any signaling scheme and every signal* $\mathbf{s}$ *in the support, Alice's maxmin payoff is at most* $\frac{d}{n}\left(\delta^3-\delta^4+\left(1-\Omega\left(1\right)\right)\delta^5\right)$.

Proof. On any signal, Bob chooses $\left(v^B,c^B\right)$ from the same distribution that Alice uses for $\left(v^A,c^A\right)$. He draws j^B uniformly at random, and picks t^B and q^B so as to minimize $\mathsf{E}\left[U^A_{\text{seek}}\mid\mathbf{s}\right]$. Finally, for each j^B, Bob draws w^B from the same distribution that Alice uses for w^A conditioning on $j^A=j^B$ (and uniformly at random if Alice never plays $j^A=j^B$). By symmetry, $\mathsf{E}\left[U^A_b\mid\mathbf{s}\right]=0$ and $\mathsf{E}\left[U^A_{\text{seek}}\mid\mathbf{s}\right]\leq 0$. See full version for details [17]. ◀

5 Lying is even harder

▶ **Theorem 17.** *Approximating* Zero-Sum Lying *with Alice's payoffs in* $[0,1]$ *to within an additive* $(1-2^{-n})$ *is* NP*-hard.*

Construction

Consider the construction from Section 4 for the honest signaling problem. Lemmata 15 and 16 guarantee that there exists a distribution $\mathcal{D}_{\text{HONEST}}$ of $n\times n$ zero-sum games and constants $c_1>c_2$ such that it is NP-hard to distinguish between the following:

Completeness. If ψ is satisfiable, there exists a signaling scheme φ_{HONEST}, such that for any signal in φ_{HONEST}'s support, Alice's maxmin payoff is at least c_1/n.

Soundness. If ψ is $(1-\eta)$-unsatisfiable, for every signaling scheme φ'_{HONEST} and every signal in the support, Alice's maxmin payoff is at most c_2/n.

For Zero-Sum Lying, we construct a hard distribution of $n\times(n+1)$ zero-sum games as follows. With probability 2^{-n} Alice's payoffs matrix is of the form:

$$\begin{pmatrix} & -\left(c_1+c_2\right)/2n \\ -A^{\top}_{\text{HONEST}} & \vdots \\ & -\left(c_1+c_2\right)/2n \end{pmatrix}, \tag{4}$$

where Alice chooses a row (Bob chooses a column), and A_{HONEST} is an $n\times n$ matrix drawn from $\mathcal{D}_{\text{HONEST}}$. In other words, Bob has to choose between receiving payoff $\left(c_1+c_2\right)/2n$, or playing a game drawn from $\mathcal{D}_{\text{HONEST}}$, but with the roles reversed.

Otherwise (with probability $1-2^{-n}$), Alice's payoff depends only on Bob: it is 1 if Bob chooses any of his first n actions, and 0 otherwise; we call this the *degenerate game*.

Notice that we promised payoffs in $[0,1]$, whereas (4) has payoffs in $[-1,0]$. $[0,1]$ payoffs can be obtained, without compromising the inapproximability guarantee, by scaling and shifting the entries in (4) in a straightforward manner.

Completeness

▶ **Lemma 18.** *If* ψ *is satisfiable, there exists a dishonest signaling scheme, such that Alice's expected payoff is at least* $1-2^{-n}$.

Proof. We first construct φ_{ALLEGED} as follows. Whenever nature samples a payoff matrix as in (4), φ_{ALLEGED} outputs the signal that φ_{HONEST} would output for A_{HONEST}. Whenever Alice and Bob play the degenerate game, φ_{ALLEGED} outputs a special symbol $\perp$.

When Bob observes any symbol from the support of φ_{HONEST}, he can guarantee a payoff of $c_1/n > (c_1 + c_2)/2n$ by playing a mix of his first n strategies. Therefore he only uses his last strategy when observing the special symbol $\perp$.

Our true signaling scheme φ_{REAL} always outputs an (arbitrary) signal from the support of φ_{HONEST}, regardless of the state of nature. With probability $1 - 2^{-n}$, Alice and Bob are actually playing the degenerate game, so Alice's payoff is 1. ◀

Soundness

▶ **Lemma 19.** *If ψ is $(1-\eta)$-unsatisfiable, then for any dishonest signaling scheme $\left(\varphi'_{\text{ALLEGED}}, \varphi'_{\text{REAL}}\right)$, Alice's expected payoff is negative.*

Proof. Any signal in the support of $\varphi'_{\text{ALLEGED}}$ corresponds to a mixture of the degenerate game, and the distribution induced by some signal $\mathbf{s}'$ in the support of some honest signaling scheme φ'_{HONEST} for $\mathcal{D}_{\text{HONEST}}$. In the degenerate game, Bob always prefers to play his last strategy. For any $\mathbf{s}'$, Bob again prefers a payoff of $(c_1 + c_2)/2n$ for playing his last strategy over a maxmin of at most c_2/n when playing any mixture of his first n strategies. Therefore, Bob always plays his last strategy, regardless of the signal he receives, which guarantees him a payoff of $(c_1 + c_2)/2^{n+1}n > 0$. ◀

Acknowledgements. I thank Shaddin Dughmi for explaining [8]. I thank Jonah Brown-Cohen, Rishi Gupta, Christos Papadimitriou, Tselil Schramm, and anonymous reviewers for helpful comments on earlier drafts.

References

1 Scott Aaronson, Russell Impagliazzo, and Dana Moshkovitz. AM with multiple merlins. In *Computational Complexity (CCC), 2014 IEEE 29th Conference on*, pages 44–55. IEEE, 2014.

2 George Akerlof. The market for lemons: Qualitative uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970.

3 Ingo Althofer. On sparse approximations to randomized strategies and convex combinations. *Linear Algebra and its Applications*, 199:339–355, 1994.

4 Yakov Babichenko, Christos H. Papadimitriou, and Aviad Rubinstein. Can almost everybody be almost happy? In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, Cambridge, MA, USA, January 14-16, 2016*, pages 1–9, 2016. doi:10.1145/2840728.2840731.

5 Umang Bhaskar, Yu Cheng, Young Kun 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, EC'16, Maastricht, The Netherlands, July 24-28, 2016*, pages 479–496, 2016. doi:10.1145/2940716.2940753.

6 Mark Braverman, Young Kun-Ko, Aviad Rubinstein, and Omri Weinstein. ETH hardness for densest-k-subgraph with perfect completeness. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16-19*, pages 1326–1341, 2017. doi:10.1137/1.9781611974782.86.

7 Mark Braverman, Young Kun-Ko, and Omri Weinstein. Approximating the best Nash Equilibrium in $n^{o(\log n)}$-time breaks the Exponential Time Hypothesis. In *Proceedings of the Twenty-Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA*

2015, San Diego, CA, USA, January 4-6, 2015, pages 970–982, 2015. doi:10.1137/1.9781611973730.66.

8 Yu Cheng, Ho Yee Cheung, Shaddin Dughmi, Ehsan Emamjomeh-Zadeh, Li Han, and Shang-Hua Teng. Mixture selection, mechanism design, and signaling. In *FOCS*, 2015. To appear. URL: http://arxiv.org/pdf/1508.03679v1.pdf.

9 Constantinos Daskalakis. On the Complexity of Approximating a Nash Equilibrium. *ACM Transactions on Algorithms*, 9(3):23, 2013. doi:10.1145/2483699.2483703.

10 Constantinos Daskalakis and Christos H. Papadimitriou. On oblivious PTAS's for Nash equilibrium. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009, Bethesda, MD, USA, May 31 – June 2, 2009*, pages 75–84, 2009. Full version available at http://arxiv.org/abs/1102.2280. doi:10.1145/1536414.1536427.

11 Irit Dinur. The PCP theorem by gap amplification. *J. ACM*, 54(3):12, 2007. doi:10.1145/1236457.1236459.

12 Shaddin Dughmi. On the hardness of signaling. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 354–363, 2014. doi:10.1109/FOCS.2014.45.

13 Shaddin Dughmi, Nicole Immorlica, and Aaron Roth. Constrained signaling for welfare and revenue maximization. *SIGecom Exchanges*, 12(1):53–56, 2013. doi:10.1145/2509013.2509022.

14 Yuval Emek, Michal Feldman, Iftah Gamzu, Renato Paes Leme, and Moshe Tennenholtz. Signaling schemes for revenue maximization. *ACM Trans. Economics and Comput.*, 2(2):5, 2014. doi:10.1145/2594564.

15 Russell Impagliazzo and Ramamohan Paturi. On the complexity of k-sat. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001. doi:10.1006/jcss.2000.1727.

16 Peter Bro Miltersen and Or Sheffet. Send mixed signals: earn more, work less. In *ACM Conference on Electronic Commerce, EC'12, Valencia, Spain, June 4-8, 2012*, pages 234–247, 2012. doi:10.1145/2229012.2229033.

17 Aviad Rubinstein. Honest signaling in zero-sum games is hard, and lying is even harder. *CoRR*, abs/1510.04991, 2015. URL: http://arxiv.org/abs/1510.04991.