



Computing Equilibria of Prediction Markets via Persuasion

Jerry Anunrojwong^{1,2(✉)}, Yiling Chen³, Bo Waggoner⁴, and Haifeng Xu^{3,5}

¹ Massachusetts Institute of Technology, Cambridge, USA
jerryanunroj@gmail.com

² Chulalongkorn University, Bangkok, Thailand

³ Harvard University, Cambridge, USA

⁴ University of Colorado Boulder, Boulder, USA

⁵ University of Virginia, Charlottesville, USA

Abstract. We study the computation of equilibria in prediction markets in perhaps the most fundamental special case with two players and three trading opportunities. To do so, we show equivalence of prediction market equilibria with those of a simpler signaling game with commitment introduced by Kong and Schoenebeck [18]. We then extend their results by giving computationally efficient algorithms for additional parameter regimes. Our approach leverages a new connection between prediction markets and Bayesian persuasion, which also reveals interesting conceptual insights.

1 Introduction

Prediction markets allow participants to buy and sell financial contracts whose payoff is contingent on the outcome of a future event. The market aggregates these decisions, which reveal beliefs about the event, into a collective prediction. Researchers study their game-theoretic properties to understand how these markets function in practice as well as how to better design them to encourage information elicitation and aggregation.

The widely-studied *scoring-rule based markets (SRM)* [13] utilize *proper scoring rules* $R(\mathbf{p}, e)$, which assign a score to each prediction \mathbf{p} on any given outcome e of the event. Each participant $t = 1, \dots, T$ arrives and updates the market prediction from \mathbf{p}^{t-1} to \mathbf{p}^t , and receives a payoff of her improvement in score, $R(\mathbf{p}^t, e) - R(\mathbf{p}^{t-1}, e)$, after the event outcome e is revealed.

Despite the apparent simplicity of this game, its equilibria have been challenging to describe. We have two primary motivations for doing so. First, prediction markets are popular in practice, and understanding the properties of their equilibria may be helpful in determining how to design such markets. Second, the SRM is a very simple but apparently deep extensive-form signaling game. Understanding it may lead to general insights regarding value of information and connections to other signaling settings. Therefore, this paper seeks algorithms and characterizations that further our understanding of these games.

The Alice-Bob-Alice (ABA) Game and Prior Work. Historically, equilibria of markets have proven difficult to describe even in the special but perhaps the most fundamental “Alice-Bob-Alice” (ABA) case. Here there are only two players and three trading opportunities. Alice observes a private signal from a set \mathcal{A} while Bob receives a private signal from a set \mathcal{B} . They can be correlated with each other and with the (random) event being predicted, which has outcomes drawn from a set \mathcal{E} . Alice, participating at $t = 1$, can choose to predict truthfully, withhold information, or even bluff and make a knowingly false prediction. This might mislead Bob into a poor prediction at $t = 2$, leaving Alice the opportunity to improve the market score significantly at $t = 3$.

A sequence of works [5, 6, 9, 11] focused on the popular log scoring rule and found conditions under which Alice fully reveals all information in stage 1 as well as cases where she reveals no information. Chen and Waggoner [7] generalized these results to a characterization of pairs (players’ signals, scoring rule) under which the first player is always truthful (termed *informational substitutes*) or withholds all information (*informational complements*). All of the results mentioned so far extend to general prediction markets with any number of players, yet solving the Alice-Bob-Alice case was often the key step.

However, one major open problem left in [7] is the computational tractability of determining whether players’ signals satisfy the substitutes condition, complements condition, or neither. The aforementioned papers also leave open what happens in the “neither” case, i.e. when Alice uses some nontrivial strategy in the first stage. To our knowledge, Kong and Schoenebeck [18] are the first to address these questions. It introduced a signaling game, the *Alice-Bob-Alice game with commitment*, that simplifies some aspects of prediction markets from an analysis perspective. Payoffs are defined as in the Alice-Bob-Alice SRM above. But instead of directly making a prediction in round 1, Alice reports according to some signaling scheme conditioned on her private information. Bob observes Alice’s signal and Alice is assigned \mathbf{p}^1 = the posterior event distribution conditioned on this signal. Crucially, Alice must commit to this signaling scheme and it is known to Bob in advance, so she cannot bluff or mislead him by deviating to another signal or prediction. For this game, [18] gave a fully polynomial-time approximation scheme (FPTAS) for computing an optimal signaling scheme of Alice when the number of possible realizations of Alice’s private information, $|\mathcal{A}|$, is constant, and the scoring rule satisfies a rather strong separability and smoothness condition.

Our Results. Our first result establishes a formal connection between ABA game with and without commitment. We prove that Alice’s optimal commitment in the ABA game is also (up to negligible ϵ) part of an equilibrium in the corresponding prediction market (without commitment). This shows, perhaps surprisingly, that any equilibrium that can be achieved when Alice is forced to commit to a signaling scheme can also be achieved in a market without commitment or explicit signaling. In other words, finding equilibria in prediction markets reduces to a pure signaling problem.

Given this result, we focus our attention on designing algorithms for the ABA game with commitment. Here, we extend the results of [18] to several other cases, although we do not solve the Alice-Bob-Alice game in full generality. Our results are built upon an interesting connection between Alice’s signaling problem and *Bayesian persuasion* [16, 17]—in some sense, Alice’s signaling scheme in round 1 is “persuading” Bob to make certain reports. We formalize this connection by proving that Alice’s signaling problem reduces to Bayesian persuasion of a *privately informed* receiver, but with a persuasion objective that is specific to prediction markets. As a direct application of this connection, we exhibit an efficient and exact algorithm for Alice’s optimal signaling in the case $|\mathcal{B}| = O(1)$ but under the assumption that the expected scoring function is piece-wise linear with polynomially many pieces. Though this restriction appears restrictive, we hope this result may serve as a stepping stone to future work. Next, we leverage techniques from algorithmic persuasion to design an FPTAS for the case $|\mathcal{A}| = O(1)$ under a natural smoothness assumption on the scoring function. This result strictly generalizes—and interestingly, also much simplifies—the main result of Kong and Schoenebeck [18]. Finally, to show the generality of our technique, we use a similar idea to design an FPTAS for the case that both $|\mathcal{B}|, |\mathcal{E}| = O(1)$.

2 Preliminaries

2.1 Signals and Probabilities

A *signal* is a random variable, denoted by a capital letter, taking values in an *outcome space* written in calligraphics. In particular, there are four signals of interest in this paper: E , A , B , and S . The signal E is a future event we would like to predict having a finite set of outcomes \mathcal{E} . The goal of a prediction market is to elicit forecasts about E in the form of probability distributions in $\Delta(\mathcal{E})$, the probability simplex over \mathcal{E} . For an outcome $e \in \mathcal{E}$, we write $\Pr[e]$ as shorthand for $\Pr[E = e]$, and so on for the other signals.

In this paper, there will always be two players, Alice and Bob. Alice observes a signal A with finite outcome space \mathcal{A} , while Bob observes B in the finite space \mathcal{B} . There is a prior distribution $\mu(e, a, b)$ on the joint realizations of $e \in \mathcal{E}$, $a \in \mathcal{A}$, and $b \in \mathcal{B}$. The prior distribution is common knowledge to Alice and Bob. Alice will be choosing to send a signal S in space \mathcal{S} . A *signaling scheme* is represented as a function $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ where $\pi(s, a) = \Pr[S = s, A = a]$ such that π satisfies $\sum_{s \in \mathcal{S}} \pi(s, a) = \sum_{e, b} \mu(e, a, b)$ for all $a \in \mathcal{A}$.

2.2 Prediction Market Model

Proper Scoring Rules. A scoring rule is a function $R : \Delta(\mathcal{E}) \times \mathcal{E} \rightarrow \mathbb{R} \cup \{-\infty\}$ that assigns a score $R(\mathbf{w}, e)$ to the prediction \mathbf{w} when the event E of our interest is realized to e . We write $R(\mathbf{w}'; \mathbf{w}) = \mathbb{E}_{E \sim \mathbf{w}} R(\mathbf{w}', E)$ for the expected score of prediction \mathbf{w}' when E is drawn from \mathbf{w} . It is *strictly proper* if for all $\mathbf{w} \neq \mathbf{w}'$, $R(\mathbf{w}'; \mathbf{w}) < R(\mathbf{w}; \mathbf{w})$. That is, for any belief \mathbf{w} , one uniquely maximizes expected score by reporting \mathbf{w} . We rely on the following characterization.

Proposition 1 ([12, 21, 23]). *For every strictly proper scoring rule R , there exists a strictly convex function $G : \Delta(\mathcal{E}) \rightarrow \mathbb{R}$ such that $R(\mathbf{w}; \mathbf{w}) = G(\mathbf{w})$. Conversely, from every strictly convex G , one can construct a strictly proper scoring rule R such that $G(\mathbf{w}) = R(\mathbf{w}; \mathbf{w})$.*

Example 1. The log scoring rule is defined as $R(\mathbf{w}, e) = \log w_e$, i.e. the logarithm of the probability assigned to e . Its “expected score function” is $G(\mathbf{w}) = \sum_e w_e \log w_e = -H(\mathbf{w})$, the negative of Shannon entropy. The quadratic scoring rule is $R(\mathbf{w}, e) = 2w_e - \|\mathbf{w}\|_2^2$. Its expected score function is $G(\mathbf{w}) = \|\mathbf{w}\|_2^2$. Both are strictly proper.

Automated Prediction Market. In this paper we focus on the popular *automated scoring-rule market (SRM)* framework of [13]. The market is parameterized by a finite set of event outcomes \mathcal{E} , a strictly proper scoring rule R , and an initial prediction $\mathbf{p}^0 \in \Delta(\mathcal{E})$. The participants arrive in a fixed, predefined order. Each round $t = 1, \dots, T$, the arriving participant observes the previous prediction \mathbf{p}^{t-1} and replaces it with a prediction \mathbf{p}^t . At the end, the event outcome $E = e$ is observed and the arriving participant at time t is paid

$$R(\mathbf{p}^t, e) - R(\mathbf{p}^{t-1}, e). \quad (1)$$

One of the key properties this payoff rule inherits from R is “one-step” truthfulness:

Fact 1. *If every player arrives only once, then it is a strictly dominant strategy to set \mathbf{p}^t to the player’s true posterior belief conditioned on all information they have observed.*

This follows immediately because R is a proper scoring rule and the second term in (1) is not under the player’s control.

However, if players participate multiple times, it might be beneficial to withhold information (or possibly even bluff). This motivates study of the *Alice-Bob-Alice (ABA) market*, a prediction market with two players and three rounds where Alice participates in rounds 1 and 3 while Bob participates in round 2. Despite its apparent simplicity, this special case captures many of the challenges of general markets and has been studied in e.g. [5, 11, 18].

Equilibrium in Markets. In the prediction market game, a strategy for Alice consists of a pair of possibly-randomized functions σ_1, σ_3 defining her predictions at rounds 1 and 3. We have $\sigma_1 : \mathcal{A} \rightarrow \Delta(\mathcal{E})$, i.e. Alice plays $\mathbf{p}^1 = \sigma_1(A)$. Next, $\sigma_3 : \mathcal{A} \times \Delta(\mathcal{E}) \times \Delta(\mathcal{E}) \rightarrow \Delta(\mathcal{E})$, where Alice at round 3 plays $\mathbf{p}^3 = \sigma_3(A, \mathbf{p}^1, \mathbf{p}^2)$. Similarly, a strategy for Bob is a possibly-randomized function $\sigma_2 : \mathcal{B} \times \Delta(\mathcal{E}) \rightarrow \Delta(\mathcal{E})$ where he plays $\mathbf{p}^2 = \sigma_2(B, \mathbf{p}^1)$.

For $t \in \{1, 2, 3\}$, define the expected net score for the prediction at round t to be

$$u_t((\sigma_1, \sigma_3), \sigma_2) = \mathbb{E}_{A, B, E, \sigma_1, \sigma_2, \sigma_3} [R(\mathbf{p}^t, E) - R(\mathbf{p}^{t-1}, E)].$$

Alice’s total expected utility is $u_A((\sigma_1, \sigma_3), \sigma_2) := u_1 + u_3$. Similarly, Bob’s expected utility is $u_B((\sigma_1, \sigma_3), \sigma_2) := u_2$.

A set of strategies $((\sigma_1, \sigma_3), \sigma_2)$ are a *Bayes-Nash equilibrium (BNE)* if each is a best response to the other, i.e. for all (σ'_1, σ'_3) , $u_A((\sigma'_1, \sigma'_3), \sigma_2) \leq u_A((\sigma_1, \sigma_3), \sigma_2)$, and similarly for all σ'_2 , $u_B((\sigma_1, \sigma_3), \sigma'_2) \leq u_B((\sigma_1, \sigma_3), \sigma_2)$.

In extensive-form games such as prediction markets, BNE can include “non-credible” threats. For example perhaps in BNE, Bob may threaten to reveal no information in the second round if Alice deviates from the equilibrium strategy. This is not credible because, if Alice were to actually deviate, Bob’s best response would still be to predict truthfully according to his beliefs. Therefore, in this paper we focus on *perfect Bayesian equilibrium (PBE)*. Informally, a BNE $((\sigma_1, \sigma_3), \sigma_2)$ is a PBE if, off the equilibrium path, these strategies still best-respond according to some beliefs that are consistent with Bayesian updating on the player’s own signal and some information about their opponent’s signal. See the full version for a formal definition.

2.3 ABA Game with Commitment

Although prediction market equilibria generally capture relative value of information, there are several technical complications. First, in principle it could be that a prediction of Alice’s does not reveal her signal for the coincidental reason that two signals give the same posterior belief. For example, in the case where both players receive a uniformly random bit and $E = A \oplus B$ (the XOR), Alice’s posterior on E is uniformly random regardless of which signal she receives. Second is the question of *commitment*. It might be that equilibria of prediction markets do not completely reflect the relative value of information and idealized signaling schemes because Alice is unable to commit to such a scheme.

This motivates us to study the more mathematically clean *ABA game with commitment*. Introduced in [18], this “game” can be phrased as a single-player decision problem, fully specified by $\{G, \mu\}$ where: convex function $G : \Delta(\mathcal{E}) \rightarrow \mathbb{R} \cup \{-\infty\}$ is chosen by the designer; μ is the prior on (A, B, E) . Alice makes the only decision in the game by selecting a signaling scheme $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$. This signaling scheme is announced to Bob. Nature draws $(A, B, E) \sim \mu$ and draws $S \sim \pi(\cdot | A)$. Bob observes the signal S , updates to a posterior $\mathbf{p}_{S,B}$, and receives utility $R(\mathbf{p}_{S,B}, E) - R(\mathbf{p}_S, E)$. Then Alice receives utility $R(\mathbf{p}_S, E) - R(\mathbf{p}, E) + R(\mathbf{p}_{A,B}, E) - R(\mathbf{p}_{S,B}, E)$ in total. Crucially, this payoff structure makes the game constant-sum since for each $A = a, B = b, E = e$, the sum of Alice’s and Bob’s utilities equals $R(\mathbf{p}_{a,b}, e) - R(\mathbf{p}, e)$, which is fixed.¹

The interpretation of these payoffs is that Alice comes to the prediction market, announces signal S , and predicts the posterior conditioned on S . Then, Bob arrives, sees S , announces B , and predicts the posterior conditioned on both S and B (via Bayesian update). Finally, Alice arrives, announces A , and

¹ This is a slight departure from the formalization of the game in [18]. There, Alice did not automatically observe Bob’s signal, causing complications in the case where Bob’s report $\mathbf{p}_{S,B}$ could be the same for two different outcomes $b, b' \in \mathcal{B}$.

predicts the posterior given both A and B . In other words, as phrased by [7, 14], Alice receives the *marginal value* of signal S over the prior; then Bob receives the marginal value of B over S ; and finally, Alice receives the marginal value of A over S, B .

2.4 Bayesian Persuasion

The ABA game turns out to be relevant to the Bayesian persuasion model. A persuasion game is played between a *sender* and a *receiver*. The receiver is faced with selecting an action i from $[k] = \{1, \dots, k\}$. Both the sender and receiver utility depend on the receiver’s action as well as a state of nature e supported on \mathcal{E} . Formally, the sender and receiver payoff function are $v(i, e)$ and $u(i, e)$ where $i \in [k]$ and $e \in \mathcal{E}$.

Particularly relevant to this work is the model of *Bayesian persuasion with a privately informed receiver*, first studied by Kolotilin *et al.* [17]. Here, the sender and receiver each observe a private signal regarding the state of nature E , which may be correlated with each other. Let $A \in \mathcal{A}$ and $B \in \mathcal{B}$ denote the (random) signal observed by the sender and receiver, respectively. The joint distribution of A, B, E is public knowledge and denoted as $\mu(e, a, b)$. The Bayesian persuasion model studies how the sender can maximize her expected utility by *committing* to a signaling scheme $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ to strategically influence the receiver’s belief about e and consequently his optimal action.² Here, again, \mathcal{S} is the set of signal outcomes. In Sect. 4, we will formalize the connection to prediction markets, which involves Alice “persuading” Bob to make certain reports but with a particular form of sender objectives specific to prediction markets.

3 Equivalence with and Without Commitment

In this section, we show that Alice’s optimal signaling scheme in the ABA game with commitment yields an approximate PBE in the Alice-Bob-Alice prediction market (without commitment). Thus, we can next focus on solving the ABA game with commitment. In this section, to simplify technicalities, we assume that the proper scoring rule R has a *differentiable* convex expected score function G .

First, we formalize the sense in which Alice uses a signaling scheme even in a prediction market. This perspective has appeared in prior works on equilibria of markets, though a precise result may not have been stated. Informally, it says that in *any* equilibrium, Alice’s equilibrium strategy can be written as reporting the posterior conditioned on a signal she draws from a private scheme. Recall from Fact 1 that, because Bob only participates once and the market uses a strictly proper scoring rule, his unique best response is always to report truthfully according to his information and beliefs.

² Such a signaling scheme is also called an *experiment* by Kolotilin *et al.* [17]. We remark that their model is a special case of the general model we described here, with independent A, B and binary receiver actions.

Lemma 1. *In perfect Bayesian equilibrium of the Alice-Bob-Alice prediction market, without loss of generality, Alice’s strategy is to predict \mathbf{p}_S for some signaling scheme π and associated random signal S .*

Therefore, from here on we will describe Alice’s strategy in prediction markets as a signaling scheme π , keeping in mind that she does not publicly announce her signal and does not have to commit to the scheme.

Before we proceed, we will give some necessary definitions.

Definitions. First, let us define $V = \mathbb{E}_{A,B,E} R(\mathbf{p}_{A,B}, E) - R(\mathbf{p}, E)$ where \mathbf{p} is the prior. This is the difference in expected score between the prior and the posterior conditioned on both players’ signals (it can also be written $\mathbb{E}_{A,B} G(\mathbf{p}_{A,B}) - G(\mathbf{p})$). Next, let us define the notation $u_B(\pi'; \pi)$ as follows. In the prediction market game, suppose Alice draws from π while Bob believes she is drawing from π' . If \mathbf{p}^1 is in the support of π' given Bob’s signal B , then he does a Bayesian update to an incorrect (in general) posterior belief \mathbf{p}^2 and reports it. If \mathbf{p}^1 is not in the support of Alice’s π' strategy (“off the equilibrium path”), then Bob forms some belief over Alice’s signal and uses this to again form an incorrect posterior belief \mathbf{p}^2 . We define $u_B(\pi'; \pi)$ to be Bob’s expected utility in this case, for some off-path beliefs of Bob.

The core idea occurs in the following lemma, which shows that, under some conditions, Alice prefers to deviate to the optimal signaling scheme.

Lemma 2. *Suppose that, in the ABA game with commitment, π^* brings Alice higher utility than π . Then in the Alice-Bob-Alice prediction market, if Alice plays π and always learns Bob’s signal after his report, then Alice improves utility by deviating to π^* .*

To prove our main result, we also need the following continuity claim.

Lemma 3. *In the prediction market with differentiable G , fixing Bob’s strategy, Alice’s expected utility is continuous in π ; and similarly, fixing Alice’s strategy, Bob’s expected utility is continuous with respect to each of his reports at the second stage (i.e. outcomes of \mathbf{p}^2) as well as each of the probabilities he places on each report.*

These results allow us to prove the main result of this section.

Theorem 1. *Let π^* be the optimal signaling scheme for the ABA game with commitment, i.e. the minimizer of $u_B(\pi; \pi)$. Then for any ϵ , there is an ϵ -PBE of the Alice-Bob-Alice prediction market in which Alice plays within ϵ of π^* .*

4 ABA Game with Commitment Is Bayesian Persuasion

In this section, we formally establish the connection between the ABA game with commitment (denoted as **ABA-Commit**) and the Bayesian Persuasion (BP) game with a privately informed receiver (denoted as **BP-Private**). Besides revealing interesting conceptual insights, this connection also enables us to directly employ ideas from Bayesian persuasion to design an efficient algorithm for the ABA game when the size of Bob’s signal space is a constant and the expected score function G is k -piecewise linear.

4.1 Reducing ABA-Commit to BP-Private

We start by simplifying the equilibrium analysis of the ABA game with commitment. Since Bob has only one chance to participate in the ABA game, his optimal strategy is simply to reveal his original signal at $t = 2$ (assuming tie breaking in favor of more information) and Alice will also reveal all her information at $t = 3$. Therefore, the only non-trivial stage is Alice's optimal commitment at the first stage. Since the game is constant-sum, so maximizing Alice's utility is equivalent to minimizing Bob's utility. As a result, solving the ABA game with commitment boils down to compute *Alice's optimal commitment* (to a signaling scheme) at the first stage to minimize Bob's utility.

For convenience and clarity, we state the result for piecewise linear convex function G , however this connection holds for arbitrary convex G function (see remarks at the end of the theorem proof).

Theorem 2. *For any ABA-commit instance $\{G, \mu\}$ where G is k -piecewise linear and μ is the prior over (A, B, E) , there is a BP-private instance such that Alice's optimal commitment is the same as the sender's optimal commitment in the BP-private instance, which is described as follows: (1) the instance has the same joint prior μ over the sender signal A , receiver signal B and event E ; (2) The receiver utility function $U_G(i, e)$ is uniquely determined by G with action set $[k] = \{1, 2, \dots, k\}$; (3) The sender utility as a function of any signaling scheme $\pi : \mathcal{S} \times \mathcal{A} \rightarrow [0, 1]$ is given by*

$$\text{Sender Obj} = \mathbb{E} \max_s \max_{i \in [k]} \sum_{e \in E} [U_G(i, e) \cdot \Pr(e|s)] - \mathbb{E} \max_{s,b} \max_{i \in [k]} \sum_{e \in E} [U_G(i, e) \cdot \Pr(e|s, b)]. \quad (2)$$

4.2 A Direct Application of the Reduction

As a direct application of the reduction in Sect. 4.1, we now show how to use this connection to compute Alice's optimal commitment when $|\mathcal{B}|$ is constant and the expected score function G is k -piecewise linear. Our algorithm is polynomial in k but exponential in the constant $|\mathcal{B}|$, as described in the following theorem.

Theorem 3. *When G is k -piecewise linear, there exists a $\text{poly}(k^{|\mathcal{B}|}, |\mathcal{A}|, |\mathcal{E}|)$ -time algorithm that computes Alice's optimal signaling scheme to commit to.*

In the introduction, we discussed the connection between ABA-commit with informational substitutes and complements. Two signals are strong substitutes if the optimal signaling scheme is to always reveal all information, and two signals are strong complements if the optimal signaling scheme is to always reveal no information. We can use the algorithm in this section to compute the signaling scheme exactly. Therefore, the following corollary is immediate.

Corollary 1. *If G is k -piecewise linear, then there exists a $\text{poly}(k^{|\mathcal{B}|}, |\mathcal{A}|, |\mathcal{E}|)$ -time algorithm that tests whether two signals A and B are strong substitutes, complements, or neither.*

5 FPTAS for Different Parameter Regimes

In this section, we develop Fully Polynomial Time Approximation Schemes (FPTAS) for the ABA game with commitment for different parameter regimes. These results cover a wider range of settings, and in particular, strictly generalize the main result of Kong and Schoenebeck [18]. Moreover, our algorithm is much simpler than that in [18] and is inspired by ideas that have also been used in the previous literature of algorithmic Bayesian persuasion.

While we do not use the explicit correspondence with the Bayesian persuasion instance developed in Sect. 4 here, we use key analytical techniques from the persuasion literature. Namely, the signaling scheme can be equivalently viewed as a distribution of posteriors and the only constraint on that distribution is the *Bayes-plausibility* constraint: the expectation of the posteriors equal the prior. We then show that under a Lipschitz-like constraint on G , a small perturbation of the posterior leads to a small perturbation of Alice's payoff. We can therefore discretize the space of posteriors within ϵ precision and show that there exists an approximately optimal signaling mechanism whose induced posteriors lie only on those grid points. When the total number of grid points are polynomially bounded, we obtain efficient algorithms. This idea has been employed in algorithmic persuasion (e.g., [4, 8]).

We start by defining the continuity condition we need on the expected score function G .

Definition 1 (Local Hölder Continuity). *A function $G : \mathbb{R}^n \rightarrow \mathbb{R}$ is (α, β) -locally Hölder continuous if there exists $\alpha > 0, \beta \in (0, 1]$ and some $c \in (0, 1)$ such that $|G(\mathbf{x}) - G(\mathbf{y})| \leq \alpha |\mathbf{x} - \mathbf{y}|^\beta$ for any \mathbf{x}, \mathbf{y} such that $|\mathbf{x} - \mathbf{y}| \leq c$.*

Note that local Hölder continuity is a natural and weak continuity assumption, which holds for almost any reasonable scoring rule. In particular, it is weaker than the standard Hölder continuity, which requires the above condition to hold for any \mathbf{x}, \mathbf{y} , not only those with $|\mathbf{x} - \mathbf{y}| \leq c$. Hölder continuity is then weaker than the Lipschitz continuity which corresponds to the case of $\beta = 1$. Moreover, we will see later that α does not have to be an absolute constant; only that α is polynomial-sized is enough for an FPTAS.

To obtain an FPTAS for the case with constant $|\mathcal{A}|$, Kong and Schoenebeck [18] defined another notion of continuity of G , which they call *niceness* condition formally described as follows. It turns out that *niceness* condition is a stronger requirement than the local Hölder continuity. So any function satisfying their condition also satisfies ours, including quadratic and log scoring rules.

Definition 2 (Niceness Condition [18]). *A function $G : \Delta_n \rightarrow \mathbb{R}$ is λ -nice if there exists a function $g : [0, 1] \rightarrow \mathbb{R}$ such that $G(\mathbf{x}) = \sum_i g(x_i)$ for every $\mathbf{x} \in \Delta_n$, $g(0) = g(1) = 0$, g is convex, and there exists a constant $\lambda \in (0, 1)$ such that for sufficiently small ϵ , $\max(|g(\epsilon)|, |g(1 - \epsilon)|) \leq \epsilon^\lambda$.*

Proposition 2. *Any function that is λ -nice for some $\lambda \leq 1$ is $(n^{1-\lambda}, \lambda)$ -locally Hölder continuous.³*

The niceness condition is a relatively strong requirement, especially as requires the expected score function G to be separable in all arguments $G(x) = \sum_i g(x_i)$. It happens to hold for log and quadratic scoring rules, but it is certainly not a property we generally expect to hold; the spherical scoring rule has $G(x) = (\sum_i x_i^2)^{1/2}$ which is not separable.

5.1 Constant Number of Alice’s Signal Outcomes

We now consider the setting of [18] with constant size of Alice’s signal space, i.e., $d \equiv |\mathcal{A}|$ is a constant. Kong and Schoenebeck [18] prove that when G satisfies the niceness condition, there is an FPTAS for this case. Here we exhibit another FPTAS for this setting based on the aforementioned idea from persuasion but under the (weaker) assumption of local Hölder continuity. This thus strictly generalizes the result in [18].

Let $\Delta_d \equiv \Delta(\mathcal{A})$ denote the set of all possible distributions over signal realizations of A . Let $\mathbf{p} \in \Delta_d$ denote a generic posterior distribution over Alice’s signal space. Throughout we always use $\|\mathbf{z}\| = \sum_i |z_i|$ to denote the l_1 norm of a vector \mathbf{z} . For a function f , denote by $\{f(e)\}_{e \in \mathcal{E}}$ a vector of dimension $|\mathcal{E}|$ whose entries are $f(e)$ for $e \in \mathcal{E}$. We prove the following theorem.

Theorem 4. *Assume that $|\mathcal{A}|$ is a constant, and the G function is (α, β) -locally Hölder continuous for some $\alpha, \beta > 0$ and bounded within $[-L, L]$ for some L . Then there exists a $\text{poly}(|\mathcal{B}|, |\mathcal{E}|, 1/\delta, L)$ -time algorithm that computes Alice’s δ -optimal signaling scheme.*

5.2 Constant Number of Event Outcomes and Bob’s Signal Outcomes

Next we exhibit an FPTAS for another parameter regime: both $n_E \equiv |\mathcal{E}|$ and $n_B \equiv |\mathcal{B}|$ are constant. The proof uses the same technique as in the previous section. The key idea is that Alice’s signaling scheme can be viewed equivalently as a distribution over posterior distributions $\mathbf{v} \in \Delta(\mathcal{E} \times \mathcal{B})$ jointly over the event and the Bob’s private signal, and that this distribution captures *all* of the information needed. Compared to Theorem 3, this result does not require k -piecewise linearity of G but requires that $|\mathcal{E}|$ is a constant. Moreover, this result is an FPTAS whereas Theorem 3 gives an exact algorithm.

Theorem 5. *Assume that $|\mathcal{E}|$ and $|\mathcal{B}|$ are constants, and the G function is (α, β) -locally Hölder continuous for some $\alpha, \beta > 0$ and bounded within $[-L, L]$ for some L . Then there exists a $\text{poly}(|\mathcal{A}|, 1/\delta, L)$ -time algorithm that computes Alice’s δ -optimal signaling scheme.*

³ Note that if $\lambda > 1$ in the λ -nice condition, or if $\beta > 1$ in the (α, β) -local Hölder continuity condition, then G is identically zero so we are not interested in those trivial cases.

6 Conclusion and Directions

In this work, we took steps toward better understanding of equilibria of prediction markets, identifying informational substitutes and complements, and connections between these problems and other signaling games including Bayesian persuasion.

While these results extend the work of [18] in several ways – connecting Alice’s optimal commitment to the original prediction market game, generalizing results for the case of fixed $|\mathcal{A}|$, and new algorithms for other cases – much open work still remains. A first direction is to give efficient algorithms with fewer assumptions, e.g. if $|\mathcal{B}|$ is bounded but we have fewer restrictions on G . A second direction is to prove intractability results, which do not yet exist for this game, although the problem appears quite challenging. It would also be interesting to understand whether the problem of testing whether signals are informational substitutes is tractable or not, and whether computing Alice’s optimal signaling scheme is algorithmically easier than testing substitutes. Finally, one can ask how these results extend to larger prediction market games.

References

1. Abernethy, J., Chen, Y., Vaughan, J.W.: Efficient market making via convex optimization, and a connection to online learning. *ACM Trans. Econ. Comput.* **1**(2), 12 (2013)
2. Berg, J., Forsythe, R., Nelson, F., Rietz, T.: Results from a dozen years of election futures markets research. In: *Handbook of Experimental Economics Results* (2008)
3. Bergemann, D., Morris, S.: The comparison of information structures in games: Bayes correlated equilibrium and individual sufficiency. Technical report 2, May 2016
4. Bhaskar, U., Cheng, Y., Ko, Y.K., Swamy, C.: Hardness results for signaling in Bayesian zero-sum and network routing games. In: *Proceedings of the 2016 ACM Conference on Economics and Computation*, pp. 479–496. ACM (2016)
5. Chen, Y., et al.: Gaming prediction markets: equilibrium strategies with a market maker. *Algorithmica* **58**(4), 930–969 (2010)
6. Chen, Y., Reeves, D.M., Pennock, D.M., Hanson, R.D., Fortnow, L., Gonen, R.: Bluffing and strategic reticence in prediction markets. In: Deng, X., Graham, F.C. (eds.) *WINE 2007*. LNCS, vol. 4858, pp. 70–81. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-77105-0_10
7. Chen, Y., Waggoner, B.: Informational substitutes. In: *56th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2016* (2016)
8. Cheng, Y., Cheung, H.Y., Dughmi, S., Emamjomeh-Zadeh, E., Han, L., Teng, S.H.: Mixture selection, mechanism design, and signaling. In: *2015 IEEE 56th Annual Symposium on Foundations of Computer Science*, pp. 1426–1445. IEEE (2015)
9. Dimitrov, S., Sami, R.: Non-myopic strategies in prediction markets. In: *Proceedings of the 9th ACM Conference on Electronic Commerce, EC 2008*, pp. 200–209. ACM (2008)
10. Freeman, R., Pennock, D.M., Vaughan, J.W.: The double clinching auction for wagering. In: *Proceedings of the 18th Conference on Economics and Computation (EC)* (2017)

11. Gao, X.A., Zhang, J., Chen, Y.: What you jointly know determines how you act: strategic interactions in prediction markets. In: Proceedings of the 14th ACM Conference on Electronic Commerce, EC 2013, pp. 489–506. ACM (2013). <https://doi.org/10.1145/2482540.2482592>
12. Gneiting, T., Raftery, A.E.: Strictly proper scoring rules, prediction, and estimation. *J. Am. Stat. Assoc.* **102**(477), 359–378 (2007)
13. Hanson, R.: Combinatorial information market design. *Inf. Syst. Front.* **5**(1), 107–119 (2003)
14. Howard, R.A.: Information value theory. *IEEE Trans. Syst. Sci. Cybern.* **2**(1), 22–26 (1966)
15. Iyer, K., Johari, R., Moallemi, C.C.: Information aggregation and allocative efficiency in smooth markets. *Manag. Sci.* **60**(10), 2509–2524 (2014)
16. Kamenica, E., Gentzkow, M.: Bayesian persuasion. *Am. Econ. Rev.* **101**(6), 2590–2615 (2011)
17. Kolotilin, A., Mylovannov, T., Zapechelnuyuk, A., Li, M.: Persuasion of a privately informed receiver. *Econometrica* **85**(6), 1949–1964 (2017)
18. Kong, Y., Schoenebeck, G.: Optimizing Bayesian information revelation strategy in prediction markets: the Alice Bob Alice case. In: 9th Innovations in Theoretical Computer Science Conference, ITCS 2018 (2018)
19. Lambert, N.S., et al.: An axiomatic characterization of wagering mechanisms. *J. Econ. Theory* **156**, 389–416 (2014)
20. Lambert, N.S., et al.: Self-financed wagering mechanisms for forecasting. In: Proceedings of the 9th ACM Conference on Electronic Commerce, EC 2008, pp. 170–179. ACM (2008)
21. McCarthy, J.: Measures of the value of information. *Proc. Nat. Acad. Sci.* **42**(9), 654–655 (1956)
22. Ostrovsky, M.: Information aggregation in dynamic markets with strategic traders. *Econometrica* **80**(6), 2595–2647 (2012)
23. Savage, L.J.: Elicitation of personal probabilities and expectations. *J. Am. Stat. Assoc.* **66**(336), 783–801 (1971)
24. Tetlock, P.E., Gardner, D.: *Superforecasting: The Art and Science of Prediction*. Broadway Books, New York (2016)
25. Witkowski, J., Freeman, R., Vaughan, J.W., Pennock, D.M., Krause, A.: Incentive-compatible forecasting competitions. In: AAAI (2018)