Optimizing Bayesian Information Revelation Strategy in Prediction Markets: the Alice Bob Alice Case

Yuqing Kong[*1] and Grant Schoenebeck[†2]

1   University of Michigan, Ann Arbor, USA
    yuqkong@umich.edu
2   University of Michigan, Ann Arbor, USA
    schoeneb@umich.edu

─── Abstract ───

Prediction markets provide a unique and compelling way to sell and aggregate information, yet a good understanding of optimal strategies for agents participating in such markets remains elusive. To model this complex setting, prior work proposes a three stages game called the Alice Bob Alice (A-B-A) game – Alice participates in the market first, then Bob joins, and then Alice has a chance to participate again. While prior work has made progress in classifying the optimal strategy for certain interesting edge cases, it remained an open question to calculate Alice's best strategy in the A-B-A game for a general information structure.

In this paper, we analyze the A-B-A game for a general information structure and (1) show a "revelation-principle" style result: it is enough for Alice to use her private signal space as her announced signal space, that is, Alice cannot gain more by revealing her information more "finely"; (2) provide a FPTAS to compute the optimal information revelation strategy with additive error when Alice's information is a signal from a constant-sized set; (3) show that sometimes it is better for Alice to reveal partial information in the first stage even if Alice's information is a single binary bit.

1   Introduction

Prediction markets aggregate information from diverse resources in a compelling manner. However, despite their powerful function in real-life deployments, large holes remain in the theory of prediction markets. For example, a basic information revelation question – how should an agent reveal her information in a prediction market to maximize her expected payoff – is still not fully answered. To model the complex setting of prediction market and deal with the information revelation question, Chen et al. [4, 3], Chen and Waggoner [5] propose and study a three stages game called the Alice Bob Alice (A-B-A) game – Alice participates in the market first, then Bob joins, and then Alice has an opportunity to participate again. They also define two special information structures – "substitutes" and "compliments" – and

show that when traders' information are substitutes (compliments), Alice should reveal her information as soon (late) as possible. However, apart from those extreme cases, it remained an open question to calculate Alice's optimal information revelation strategy in the A-B-A game for a general information structure which is the main focus of the current paper.

Computing optimal information revelation strategy is also a key problem in the Bayesian persuasion literature ([11, 7, 2]). Bayesian persuasion, proposed by Kamenica and Gentzkow [11], models a situation where a informed sender (partially) reveals her information to persuade an uninformed receiver to adopt an action. In the Bayesian persuasion model, the informed sender first commits to an information revelation strategy and then announces a signal based upon the committed strategy. Borrowing this idea of commitment from Bayesian persuasion, we consider the A-B-A game where Alice makes a commitment to her information revelation strategy before the game. We call the strategy in the A-B-A game with commitment a *Bayesian* information revelation strategy. Like in the Bayesian persuasion case, this power of commitment makes sense if we expect the game to be repeated [11]. People cannot lie to others in the market and make money forever. After many rounds of market activities, people will identify rules to effectively translate the leaked information.

Our Results

The current paper analyzes A-B-A for the general information structures and

1) proves a general revelation principle style theorem (Section 5.2.2) for the A-B-A game by showing it is enough for Alice to use her private signal space as her announced signal space, that is, Alice cannot gain more by revealing her information more "finely".

2) leverages the intuition of the aforementioned result to give a fully polynomial time approximation scheme (FPTAS) to compute the optimal information revelation strategy with additive error when Alice's information is a signal from a constant-sized set (Section 5.2.2);

3) shows that sometimes it is better for Alice to reveal partial information in the first stage, even when Alice's information is a binary bit (Section 4, Appendix B).

Before our result, it was not known how to compute (regardless of time complexity) an optimal strategy for Alice even when her signal was binary, or that such a strategy existed for a general information structure.

1.1   Related Work

Bayesian Persuasion (BP) Model

Conceptually, the Bayesian Persuasion model is different than the A-B-A with commitment. In the A-B-A game, both Alice and Bob sell information to the market and Bob is informed as well and is actually a competitor of Alice. In contrast, in BP, the sender sells the information to the uninformed receiver and has different kind of utility function with the receiver. Technically, in BP, the goal function is a linear function of the revelation strategy while the A-B-A is much more complicated.

Despite those differences, Kamenica and Gentzkow [11] also show a "revelation principle" style statement which is similar with the statement in the current paper – it is enough for the sender to draw her announced signal from the receiver's action space. That is, the sender cannot gain more by making her announced signal space more complicated. More formally, if the receiver's action space is A_r and the sender's private information space is X_s, the sender can obtain the optimal utility by just optimizing over the space of all "simple" strategies

$M : X_s \times A_r \mapsto [0, 1]$ such that $M(x, a)$ represents the probability the sender who has private information x and announces action a regardless of other "complicated" strategies. However, the proof of our A-B-A revelation principle is much more complicated than that in the BP case. In the A-B-A game, Bob's possible actions (best responses) are infinite, while in the BP case they are finite. Moreover, Alice's utility depends non-linearly on Bob's action which means even though Bob's action space has a good structure, any simple proof is unlikely to work. To prove the A-B-A revelation principle, we use a totally different method involving linear programming.

Information Revelation Problem

As mentioned in the introduction, Chen et al. [4, 3], Chen and Waggoner [5] propose and study the A-B-A game. When Alice's information and Bob's information are independent with each other, their information is defined as "compliments". On the other hand, when Alice's information and Bob's information are conditionally independent with each other (conditioning on the event they want to forecast), their information is defined as "substitutes". Chen et al. [4, 3], Chen and Waggoner [5] show that Alice should reveal her information as late (early) as possible when their information is "compliments" ("substitutes"). In those extreme cases, Alice cannot obtain better utility by partially revealing her information which is not true in the general case which is studied in the current paper.

Azar et al. [1] consider a model where the market price is a reverse Gaussian random walk and the expert who has a less noisy signal should decide a time to announce her signal. However, in their model, the expert only has a chance to participate in the market once while in the A-B-A game, Alice has multiple chances to participate the market and can partially reveal her information at first to obtain better utility. Moreover, the information structure considered in Azar et al. [1] is limited by their assumptions while the current paper considers the general information structure.

2 Preliminaries

2.1 Prediction Markets

In this section, we introduce the market scoring rule (MSR) model [9, 10]. We first introduce the main technical tools in the MSR model – proper scoring rules [13], which are used to measure the *score* (accuracy) of the forecast.

Proper Scoring Rules [8, 13]

A scoring rule which we denote $PS : \Sigma \times \Delta_\Sigma \mapsto \mathbb{R}$ takes in a signal $\sigma \in \Sigma$ and a distribution over signals $\mathbf{p} \in \Delta_\Sigma$ and outputs a real number. A scoring rule is *proper* if, whenever the first input is drawn from a distribution $\mathbf{p}$, the expectation of PS is maximized if the second input is $\mathbf{p}$. That is, $\mathbf{p} \in \arg\max_{\mathbf{p}'} \mathbb{E}_{\sigma \sim \mathbf{p}}[PS(\sigma, \mathbf{p}')]$. A scoring rule is called *strictly proper* if $\mathbf{p}$ uniquely maximizes $\mathbb{E}_{\sigma \sim \mathbf{p}}[PS(\sigma, \mathbf{p}')]$. We will assume throughout that the scoring rules we use are strictly proper. By slightly abusing notation, we can extend a scoring rule to be $PS : \Delta_\Sigma \times \Delta_\Sigma \mapsto \mathbb{R}$ by simply taking $PS(\mathbf{p}, \mathbf{p}') = \mathbb{E}_{\sigma \leftarrow \mathbf{p}}(\sigma, \mathbf{p}')$. Any proper scoring rule is linear in the first term.

Fix an outcome space Σ for a signal σ. Let $\mathbf{q} \in \Delta_\Sigma$ be a reported distribution.

■ **Figure 1** An example of nice convex function $G : [0,1] \mapsto \mathbb{R}$, $G(0) = G(1) = 0$.

▶ **Example 1.** (Logarithmic Scoring Rule) A *logarithmic scoring rule* maps a signal σ and reported distribution $\mathbf{q}$ to a payoff as follows:

$$LSR(\sigma, \mathbf{q}) = \log(\mathbf{q}(\sigma)).$$

Let the signal σ be drawn from some random process with distribution $\mathbf{p} \in \Delta_\Sigma$. Then the expected payoff of the logarithmic scoring rule is

$$\mathbb{E}_{\sigma \leftarrow \mathbf{p}}[LSR(\sigma, \mathbf{q})] = LSR(\mathbf{p}, \mathbf{q}) = \sum_\sigma \mathbf{p}(\sigma) \log \mathbf{q}(\sigma)$$

This value will be maximized if and only if $\mathbf{q} = \mathbf{p}$

▶ **Example 2.** (Quadratic Scoring Rule / Brier scoring rule) A *quadratic scoring rule* which maps a signal σ and reported distribution $\mathbf{q}$ to a payoff as follows:

$$QSR(\sigma, \mathbf{q}) = 2\mathbf{q}(\sigma) - \sum_{\sigma'} \mathbf{q}(\sigma')^2 - 1.$$

Let the signal σ be drawn from some random process with distribution $\mathbf{p} \in \Delta_\Sigma$. Then the expected payoff of the logarithmic scoring rule is

$$\mathbb{E}_{\sigma \leftarrow \mathbf{p}}[QSR(\sigma, \mathbf{q})] = QSR(\mathbf{p}, \mathbf{q}) = 2\langle \mathbf{p}, \mathbf{q} \rangle - \langle \mathbf{q}, \mathbf{q} \rangle - 1$$

This value will be maximized if and only if $\mathbf{q} = \mathbf{p}$.

In general, proper scoring rules can be constructed from convex functions. Given a bounded convex function $H : \Delta_\Sigma \mapsto \mathbb{R}$, we define $PS_H : \Sigma \times \Delta_\Sigma \mapsto \mathbb{R}$ such that

$$PS_H(\sigma, \mathbf{p}) = H(\mathbf{p}) - \langle H'(\mathbf{p}), \mathbf{p} \rangle + H'_\sigma(\mathbf{p})$$

where $\langle , \rangle$ denotes the inner product of two vectors and H'_σ denotes the partial derivative of H with respect to the σ^{th} entry.

▶ **Fact 3.** *[8] When $H : \Delta_\Sigma \mapsto \mathbb{R}$ is (strictly) convex, $PS_H : \Sigma \times \Delta_\Sigma \mapsto \mathbb{R}$ is (strictly) proper and $\forall \mathbf{p}$, $PS_H(\mathbf{p}, \mathbf{p}) = H(\mathbf{p})$.*

To control the convergence rate analysis in the future, we consider a special class of proper scoring rules.

▶ **Definition 4** (Nice convex functions)**.** We say a convex real function $G : [0,1] \mapsto \mathbb{R}$ is *nice* if (i) $G(0) = G(1) = 0$ and (ii) there exists a constant $\lambda > 0$ such that when ϵ is sufficiently small, $\max\{|G(\epsilon)|, |G(1 - \epsilon)|\} \leq \epsilon^\lambda$. We denote the set of all such nice functions as $\mathcal{G}$.

▶ **Definition 5** (PS^G, H_G, $\mathcal{PSG}$). Given a bounded strictly convex real function $G : [0, 1] \mapsto \mathbb{R}$, (Figure 1), we define $H_G : \Delta_\Sigma \mapsto \mathbb{R}$ as a function such that $H_G(\mathbf{p}) := \sum_\sigma G(\mathbf{p}(\sigma))$ for any $\mathbf{p} \in \Delta_\Sigma$. We define $PS^G(\sigma, \mathbf{p}) := PS_{H_G}(\sigma, \mathbf{p})$. We call such a proper scoring rule *good* if G is a nice convex function, and define the set of all nice proper scoring scoring rules as $\mathcal{PSG}$.

Now we explain the restrictions of nice convex functions $\mathcal{G}$. If we pick $G(x) = x \log x$ which is a nice convex function (Example 1), the proper scoring rule is the common used log scoring rule and $|PS^G(\mathbf{p}, \mathbf{p})| = -H_G(\mathbf{p}) = -\sum_\sigma \mathbf{p}(\sigma) \log \mathbf{p}(\sigma)$ which is the Shannon entropy of distribution $\mathbf{p}$ [6].

Note that entropy can be interpreted as the uncertainty of distribution $\mathbf{p}$. For example, when $\mathbf{p} = (0, 1, 0, 0, 0)$, there is no uncertainty, the entropy is 0. When we use other $G(x)$, we still want $|PS^G(\mathbf{p}, \mathbf{p})| = -H_G(\mathbf{p})$ to be interpreted as the uncertainty of distribution $\mathbf{p}$. Therefore, we put the restriction $G(0) = G(1) = 0$. Then if when there is no uncertainty, $\mathbf{p}$ must be an extreme point of Δ_Σ, $H_G(\mathbf{p}) = G(1) + G(0) + ... + G(0) = 0$.

The second restriction is needed when we analyze the convergence rate in the future. We hope $G(x)$ never changes too fast. Note that it's a weaker condition than lipschiz condition since $G(x) = x \log x$ does not satisfy the lipschiz condition but satisfy our restriction since $|G(\epsilon)| = \epsilon \log \epsilon \leq \epsilon^\lambda, \forall 0 < \lambda < 1$ and $|G(1 - \epsilon)| = (1 - \epsilon) \log(1 - \epsilon) \leq \log(1 - \epsilon) \leq \epsilon$.

Note this special class is still rich and several commonly used proper scoring rules, including the aforemetioned examples, belong to this class.

▶ **Remark.** Setting $G(x) = x \log x$, the nice proper scoring rule PS^G is the logarithmic scoring rule. Setting $G(x) = (x - \frac{1}{2})^2 - \frac{1}{4}$, the nice proper scoring rule PS^G is the quadratic scoring rule.

▶ **Remark.** Note that affine transformations preserve convexity. Without loss of generality, we assume for any $PS^G \in \mathcal{PSG}$, $|PS^G(\mathbf{p}, \mathbf{p})| = |H_G(\mathbf{p})| \leq 1, \forall \mathbf{p}$. This also means our results apply to the Brier Scoring rule, which is a shift of the quadradic scoring rule.

Market Scoring Rule Model

The theoretical prediction market model used in the current paper is the market scoring rule (MSR) model which is proposed by Hanson [9, 10]. In this model, market price corresponds to people's beliefs for the event. When a trader changes the market belief from p_1 to p_2, the automated market maker market scoring rule (MSR) will pay the trader the "accuracy" of forecast p_2 minus the "accuracy" of forecast p_1. The "accuracy" of the forecast is measured by the proper scoring rules [13]. We provide a formal definition in the below paragraph.

▶ **Definition 6** (Prediction market $PM(PS, X_E)$ [9, 10]). Let X_E be a random variable that people want to forecast. The market maker sets up an initial belief for X_E. Every agent can modify the market belief. When an agent changes the market belief from p_1 to p_2, her payment will be the score of belief p_2 minus the score of belief p_1, that is, $PS(X_E, p_2) - PS(X_E, p_1)$ after X_E is revealed.

2.2 Notation

For two random variables X, Y which are drawn from space $[n] \times [m]$, we define

$$\Pr[\mathbf{Y}|X = i] := (\Pr[Y = 1|X = i], \Pr[Y = 2|X = i], \cdots, \Pr[Y = m|X = i]).$$

For any function $f : [n] \mapsto \mathbb{R}$, $\mathbb{E}_X f(X) = \sum_i \Pr[X = i] f(i)$.

For a matrix M, we use $M_{i,\cdot}$ to denote the i^{th} row, $M_{\cdot,j}$ to denote the j^{th} column, and $M_{i,j} = M(i,j)$ to denote the entry at the i^{th} row and the j^{th} column. In particular, $M_{i,\cdot}$ is a row vector and $M_{\cdot,j}$ is a column vector.

A matrix M is a *transition matrix* if every entry of M is non-negative and the sum of all entries in each *row* is 1. Throughout this paper, we denote by $\mathcal{M}_{n \times m}$ the set of all transition matrices that have dimension $n \times m$, and denote $\mathcal{M}_{n \times *} = \bigcup_{m \geq n} \mathcal{M}_{n \times m}$.

Given a random variable X that has n possible outcomes, an transition matrix $M \in \mathcal{M}_{n \times m}$ defines a **transition probability** that transforms X to $M(X)$ such that $M(X)$ is a new random variable that has m possible outcomes where $\Pr[M(X) = j | X = i] = M_{i,j}$.

If the distribution of X is represented by an $1 \times n$ row vector $\mathbf{p}$, then the distribution over $M(X)$ is $\mathbf{p}M$ and $\Pr[M(x) = j] = \mathbf{p} \cdot M_{\cdot,j}$.

3   Alice Bob Alice Game with Commitment

We analyse the Alice Bob Alice game with commitment in this section, that is, Alice commits a signaling scheme before the game. The random event of interest is X_E. X_E is drawn from a signal space $\Sigma_E, |\Sigma_E| = n_E$. We use a proper scoring rule based prediction market $PM(PS, X_E)$ to pay Alice and Bob. Suppose Alice's private information is X_A and Bob's private information is X_B. X_A is drawn from a signal space $\Sigma_A, |\Sigma_A| = n_A$ and X_B is drawn from a signal space $\Sigma_B, |\Sigma_B| = n_B$.

We assume both Alice and Bob are rational.

▶ **Definition 7** (signaling scheme). Given that a signal space $\Sigma, |\Sigma| = m$, we define Alice's signaling scheme M as an $n_A \times m$ transition matrix such that $M_{x_A, \sigma} = M(x_A, \sigma)$ is the probability Alice announces signal $\sigma \in \Sigma$ given private information $X_A = x$.

A signaling scheme M defines a transition probability. We define X_σ as a random variable such that $X_\sigma := M(X_A)$, that is,

$$X_A \xrightarrow{M} X_\sigma.$$

Alice Bob Alice Game with Commitment (X_A, X_B, X_E, PS)

Stage 0 Alice commits her signaling scheme M.

Stage 1 Alice receives a signal $\sigma_A \in \Sigma_A$, implements her signaling scheme, and announces the result $\sigma \in \Sigma$. Alice changes the market belief for event X_E from the original prior forecast $\mathbf{p}_0 = \Pr[\mathbf{X_E}]$ to

$$\mathbf{p}_1 = \Pr[\mathbf{X_E} | X_\sigma = \sigma, M].$$

Stage 2 Bob changes the market belief to $\mathbf{p}_2$ (which is a function of $M, \mathbf{p}_1, X_B$ and Bob's strategy).

Stage 3 Alice changes the market belief to $\mathbf{p}_3$ (which is a function of $M, \mathbf{p}_1, \mathbf{p}_2, X_A$ and Alice's strategy).

Payment Both Alice and Bob are paid according to proper scoring rule based prediction market $PM(PS, X_E)$. Suppose the initial market belief is $\mathbf{p}_0 = \Pr[X_E | X_\sigma]$. Alice and Bob's payments are

$$\mu_A = (PS(X_E, \mathbf{p}_1) - PS(X_E, \mathbf{p}_0)) + (PS(X_E, \mathbf{p}_3) - PS(X_E, \mathbf{p}_2))$$
$$\mu_B = PS(X_E, \mathbf{p}_2) - PS(X_E, \mathbf{p}_1)$$

correspondingly.

Fix the joint distribution over random variables X_A, X_B, X_E. Consider an A-B-A game with commitment $(X_A, X_B, X_E, PS \in \mathcal{PSG})$. We assume Alice and Bob will optimally respond in stage 2 and stage 3. Actually we will see since the market uses strictly proper scoring rule, both Alice and Bob's optimal responses in stage 2 and stage 3 are unique (Claim 11, 12). In this case, both Alice and Bob's expected payments can be seen as a function of Alice's signaling scheme M. We define Alice's expected payment as $\mu_A^*(M)$; and Bob's expected payment as $\mu_B^*(M)$.

We define $\mu^*(M) := \mu_A^*(M) + \mu_B^*(M)$. We also define $\mu^* := \sup_M \mu^*(M)$, $\mu_B^\dagger := \inf_M \mu_B^*(M)$ (note that it's an infimum here), and $\mu_A^* := \mu^* - \mu_B^\dagger$. Note that $\mu_A^*(M) \leq \mu_A^*$.

▶ **Definition 8** (Optimizing Signaling Scheme Problem). Consider an A-B-A game with commitment $(X_A, X_B, X_E, PS \in \mathcal{PSG})$. An optimizing signaling scheme problem is the problem of constructing Alice's optimal signaling scheme M^* such that $\mu_A^*(M^*) = \mu_A^*$ if exists or a series of signaling schemes $\{M^*(\epsilon)\}_\epsilon$ such that

$$\mu_A^*(M^*(\epsilon)) \xrightarrow{\epsilon \to 0} \mu_A^*.$$

4 Summary of the Main Results

▶ **Theorem 9** (Optimizing Signaling Scheme). *Given the joint distribution over random variables X_A, X_B, X_E, consider an A-B-A game with commitment $(X_A, X_B, X_E, PS^G \in \mathcal{PSG})$. When Alice's private signal is from a constant-sized set, that is, n_A is a constant integer T, for all sufficiently small $0 < \epsilon < 1$, there exists an $O\left((LP(\frac{1}{\epsilon}+1)^T) + n_B n_E(\frac{1}{\epsilon}+1)^T) + n_B^2 n_E\right)$ time algorithm that constructs the signaling scheme $M^*(\epsilon) \in \mathcal{M}_{n_A \times n_A}$ such that*

$$\mu_A^*(M^*(\epsilon)) \geq \mu_A^* - \Theta(|\epsilon| + n_E |G(\epsilon)| + n_E |G(1-\epsilon)|)$$

where $LP(k)$ is the time complexity of linear programming with k variables.

Moreover, when Alice commits to signaling scheme $M^*(\epsilon)$, the optimal responses of Bob and Alice in stage 2 and stage 3 are

$$\mathbf{p}_2^* = \Pr[\mathbf{X_E}|X_\sigma, X_B] \qquad \mathbf{p}_3^* = \Pr[\mathbf{X_E}|X_A, X_B]$$

respectively where $X_A \xrightarrow{M^*(\epsilon)} X_\sigma$.

▶ **Corollary 10.** *Given the joint distribution over random variables X_A, X_B, X_E, consider an A-B-A game with commitment $(X_A, X_B, X_E, PS^G \in \mathcal{PSG})$. When Alice's private signal is from a constant-sized set, that is, n_A is a constant integer T, there is a FPTAS for Alice to optimize her signaling scheme with additive error.*

We defer the full proof to the end of Section 5.

Proof Sketch of Theorem 9

We will first give a game theoretic analysis for the optimal strategy for Alice and Bob in stage 2 and stage 3. The definition of the proper scoring rules implies that in stage 2 and stage 3, Alice and Bob should honestly report their best forecast for event X_E at that stage. The best forecast should be the posterior probability of X_E conditioning on all possible information they have at that time.

Step 1 Minimizing Bob's optimal expected payment Fixing the joint distribution over Alice and Bob's private information and the event, Bob's optimal expected payment $\mu_B^*(M)$ is a function of Alice's signal scheme $M \in \mathcal{M}^{n \times *}$. To calculate the signal scheme $M^\dagger$ for Alice to minimize Bob's optimal expected payment, we will prove that

Figure 2 The optimal signal scheme when X_A is a binary random variable given different joint distributions over X_A, X_B, X_E (see Appendix for the numerical values of the three joint distributions). Left: $M^* = \begin{bmatrix} x & 1-x \\ x & 1-x \end{bmatrix}, \forall x$ Alice's optimal strategy is revealing no information. Middle: $M^* = \begin{bmatrix} 0.18 & 0.82 \\ 1 & 0 \end{bmatrix}$ or $\begin{bmatrix} 0.82 & 0.18 \\ 0 & 1 \end{bmatrix}$ Alice's optimal strategy is revealing partial information. Right: $M* = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ or $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ Alice's optimal strategy is revealing full information.

Revelation principle it is sufficient to optimize over $M \in \mathcal{M}^{n \times n}$, we prove this by showing a "decomposibility" property of Bob's optimal expected payment $\mu_B^*(M)$;

Continousness when $M \approx M'$, $\mu_B^*(M) \approx \mu_B^*(M')$.

Then we use a linear programming based algorithm to approximate the signal scheme $M^\dagger$ for Alice to minimize Bob's optimal expected payment.

Step 2 Maximizing the total expected payment We will perturb $M^\dagger$ to $M^* \approx M^\dagger$ such that adopting the signaling scheme M^* guarantees that the sum of Alice and Bob's optimal expected payment obtains the upper-bound.

After finishing the above two steps, we obtain M^* which gives Alice the expected payment which is close to optima since M^* both minimizes Bob's optimal expected payment and maximizes the sum of Alice and Bob's expected payment.

$$\begin{aligned} \mu_A^*(M^*) =& \mu^*(M^*) - \mu_B^*(M^*) \\ =& \mu^* - \mu_B^*(M^*) \\ \approx& \mu^* - \mu_B^*(M^\dagger) \qquad\qquad (M^* \approx M^\dagger, \text{ continousness of } \mu_B^*(M)) \\ =& \mu^* - \mu_B^\dagger \\ =& \mu_A^* \end{aligned}$$

Experimental Results

We show that sometimes it is better for Alice to reveal partial information in the first stage even if Alice's information is a single binary bit by providing the optimal signal scheme of Alice in the A-B-A game in three scenarios (Figure 2). We give the numerical values of the three scenarios in appendix.

5 Optimizing Signaling Scheme

5.1 Game Theoretic Analysis of A-B-A

This section shows that in stage 2 and stage 3, Alice and Bob should honestly report their best forecast for event X_E at that stage. The best forecast is the posterior probability of X_E conditioning on all possible information they have at that time. Moreover, the sum of the expected payments of Alice and Bob obtains its upper-bound if Alice learns the exact value of X_B from Bob's behavior and plays the strategy in stage 3 that fully aggregates the information.

▶ **Claim 11.** *Given the joint distribution over random variables X_A, X_B, X_E, consider an A-B-A game with commitment $(X_A, X_B, X_E, PS^G \in \mathcal{PSG})$. Given that Alice commits signaling scheme M, $X_A \xrightarrow{M} X_\sigma$, Bob's optimal action is changing the market belief to $\mathbf{p}_2^* = \Pr[\mathbf{X_E}|X_\sigma, X_B]$ in stage 2 and his optimal expected payment is*

$$\mu_B^*(M) = \mathbb{E}_{X_\sigma, X_B}[H_G(\Pr[\mathbf{X_E}|X_\sigma, X_B]) - H_G(\Pr[\mathbf{X_E}|X_\sigma])].$$

Recall that $-H_G(\mathbf{p})$ can be interpreted as the uncertainty / "entropy" of distribution $\mathbf{p}$. The uncertainty of event X_E decreases with more information. Therefore, Bob's expected payment can be interpreted as the contribution of Bob's private information X_B to decrease the uncertainty if the event X_E given the existence of the partial information X_σ of Alice.

Proof of Claim 11. Bob only has one chance to participate. Thus, he will always report his truthful forecast which is conditioning on his own information and the information Alice conveys to him. Therefore, $\Pr[\mathbf{X_E}|X_\sigma, X_B]$ is Bob's optimal action and

$$\begin{aligned}
&\mathbb{E}_{X_E, X_\sigma, X_B} PS^G(X_E, \mathbf{p}_2) - PS^G(X_E, \mathbf{p}_1) \\
=&\mathbb{E}_{X_\sigma, X_B} \mathbb{E}_{X_E|X_\sigma, X_B}[PS^G(X_E, \mathbf{p}_2) - PS^G(X_E, \Pr[\mathbf{X_E}|X_\sigma])] && (\mathbf{p}_1 = \Pr[\mathbf{X_E}|X_\sigma]) \\
\leq&\mathbb{E}_{X_E, X_\sigma, X_B}[PS^G(X_E, \Pr[\mathbf{X_E}|X_\sigma, X_B]) - PS^G(X_E, \Pr[\mathbf{X_E}|X_\sigma])] \\
&&& \text{(definition of proper scoring rule)} \\
=&\mathbb{E}_{X_E, X_\sigma, X_B} PS^G(X_E, \Pr[\mathbf{X_E}|X_\sigma, X_B]) \\
&- \mathbb{E}_{X_E, X_\sigma} PS^G(X_E, \Pr[\mathbf{X_E}|X_\sigma]) \\
&&& \text{(The second part is independent of Bob's information } X_B) \\
=&\mathbb{E}_{X_\sigma, X_B}[H_G(\Pr[\mathbf{X_E}|X_\sigma, X_B]) - H_G(\Pr[\mathbf{X_E}|X_\sigma])] && \text{(Fact 3)}
\end{aligned}$$

◀

▶ **Claim 12.** *Given the joint distribution over random variables X_A, X_B, X_E, consider an A-B-A game with commitment $(X_A, X_B, X_E, PS^G \in \mathcal{PSG})$. When Alice learns the exact value of X_B from Bob's behavior, Alice's optimal action is changing the market belief to $\mathbf{p}_3^* = \Pr[\mathbf{X_E}|X_A, X_B]$ in stage 3 and the optimal sum of expected payment is*

$$\mu^* = \mathbb{E}_{X_E, X_A, X_B}[H_G(\Pr[\mathbf{X_E}|X_A, X_B]) - H_G(\Pr[\mathbf{X_E}])].$$

The optimal total expected payment can be interpreted as the contribution of Alice and Bob's private information to decrease the uncertainty of the event X_E and can be obtained when in stage 3, Alice learns all information.

Proof of Claim 12. According to the definition of the proper scoring rules, the optimal $\mathbf{p}_3$ in stage 3 is the forecast that conditions on all information X_A, X_B. When Alice learns the

exact value of X_B from Bob's behavior, she can play this optimal strategy in stage 3 which makes the sum of expected payments of Alice and Bob optimal.

$$\mathbb{E}_{X_E, X_A, X_B} PS^G(X_E, \mathbf{p}_3) - PS^G(X_E, \mathbf{p}_0)$$
$$=\mathbb{E}_{X_A, X_B} \mathbb{E}_{X_E | X_A, X_B} PS^G(X_E, \mathbf{p}_3) - \mathbb{E}_{X_E} PS^G(X_E, \mathbf{p}_0)$$
$$\leq \mathbb{E}_{X_E, X_A, X_B} [PS^G(X_E, \Pr[\mathbf{X_E} | X_A, X_B]) - PS^G(X_E, \mathbf{p}_0)]$$

(definition of proper scoring rule)

$$=\mathbb{E}_{X_E, X_A, X_B} PS^G(X_E, \Pr[\mathbf{X_E} | X_A, X_B]) - \mathbb{E}_{X_E} PS^G(X_E, Pr[\mathbf{X_E}]) \qquad (\mathbf{p}_0 = Pr[\mathbf{X_E}])$$
$$=\mathbb{E}_{X_A, X_B} [H_G(\Pr[\mathbf{X_E} | X_A, X_B]) - H_G(\Pr[\mathbf{X_E}])] \qquad \text{(Fact 3)}$$

◀

5.2 Minimizing Bob's Expected Payment

5.2.1 Decomposability and Continuousness of Bob's Expected Payment

This section will show two important properties required of $\mu_B^*(M)$ for the optimization step.

▶ **Definition 13** (Decomposability). Recall that $\mathcal{M}_{n \times *}$ is the set of all transition matrices which have n rows. A function $F : \mathcal{M}_{n \times *} \mapsto \mathbb{R}$ is *decomposable* if there exists a function $f : \mathbb{R}^n \mapsto \mathbb{R}$ satisfying $f(\lambda \mathbf{v}) = \lambda f(\mathbf{v})$ for any $\lambda \in \mathbb{R}^+, \mathbf{v} \in \mathbb{R}^n$ such that for any

$$M = [M_{.,1} \quad M_{.,2} \quad \cdots \quad M_{.,m}] \in \mathcal{M},$$

we have

$$F(M) = \sum_{j=1}^{m} f(M_{.,j}).$$

Given the joint distribution over random variables X_A, X_B, X_E, consider an A-B-A game with commitment $(X_A, X_B, X_E, PS^G \in \mathcal{PSR})$. Recall that we define Bob's optimal expected payment as $\mu_B^*(M)$ and random variables X_A, X_B, and X_E have n_A, n_B, and n_E possible outcomes respectively. The signal space of Alice is $\Sigma, |\Sigma| = m$.

▶ **Lemma 14** (Decomposability). $\mu_B^*(M)$ *is a decomposable function of* $M \in \mathcal{M}_{n \times *}$.

▶ **Lemma 15** (Continuousness). *For every* $M \in \mathcal{M}_{n_A \times m}$, *if* $\max_{i,j} |M_{i,j} - M'_{i,j}| \leq \epsilon$, *then*

$$|\mu_B^*(M') - \mu_B^*(M))| \leq \Theta(n_A m(n_E |G(\epsilon)| + n_E |G(1 - \epsilon)| + \epsilon)).$$

We defer the full proofs to the appendix.

Proof sketch

Recall that

$$\mu_B^*(M) = \mathbb{E}_{X_\sigma, X_B} [H_G(\Pr[\mathbf{X_E} | X_\sigma, X_B]) - H_G(\Pr[\mathbf{X_E} | X_\sigma])].$$

For the decomposability,

$$\mu_B^*(M) = \mathbb{E}_{X_\sigma, X_B} [H_G(\Pr[\mathbf{X_E} | X_\sigma, X_B]) - H_G(\Pr[\mathbf{X_E} | X_\sigma])] \qquad \text{(Claim 11)}$$
$$= \sum_\sigma \Pr[X_\sigma = \sigma] \mathbb{E}_{X_B | X_\sigma = \sigma} [H_G(\Pr[\mathbf{X_E} | X_\sigma = \sigma, X_B]) - H_G(\Pr[\mathbf{X_E} | X_\sigma = \sigma])]$$

We define $\phi_j(M) := \Pr[X_\sigma = j]$, $\psi_j(M) := \mathbb{E}_{X_B | X_\sigma = \sigma} H_G(\Pr[\mathbf{X_E} | X_\sigma = \sigma, X_B]) - H_G(\Pr[\mathbf{X_E} | X_\sigma = \sigma])$.

We will show

(i) $\phi_j(M)$ is a linear function of $M_{\cdot,j}$

(ii) $\psi_j(M)$ only depends on the "shape" of $M_{\cdot,j} - \frac{M_{\cdot,j}}{S(M_{\cdot,j})}\mathbf{1}$ – where $S(M_{\cdot,j})$ is the sum of vector $M_{\cdot,j}$

Combining both (i) and (ii), we can see $\phi_j(M)\psi_j(M) = \Phi(M_{\cdot,j})\Psi(M_{\cdot,j})$ only depends on $M_{\cdot,j}$ and moreover, for any $\lambda \in \mathbb{R}^+$,

$$\Phi(\lambda M_{\cdot,j})\Psi(\frac{\lambda M_{\cdot,j}}{S(\lambda M_{\cdot,j})}) = \lambda \Phi(M_{\cdot,j})\Psi(\frac{M_{\cdot,j}}{S(M_{\cdot,j})})$$

that is, it preserves the scalar multiplication of $M_{\cdot,j}$. Therefore, $\mu_B^*(M) = \sum_j \phi_j(M)\psi_j(M)$ is a decomposable function of M. We defer the proofs of (i) and (ii) to the appendix.

The proof for continousness is a little bit tricky. When we perturb M a little bit, it's possible that $\Pr[\mathbf{X_E}|M(X_A)])$ changes a lot. Consider an extreme case where the prior of X_A is a uniform distribution and we pick transition probability M such that $\Pr[M(X_A) = j] = 0.000001$. We add $\epsilon = 0.01$ to M_{ij} to obtain M'. In this case, $M'_{ij} >> M'_{kj}, k \neq i$. Thus, conditioning on $M'(X_A) = j$ the probability $X_A = i$ is close to 1. Therefore, $\Pr[\mathbf{X_E}|M'(X_A) = j] \approx \Pr[\mathbf{X_E}|X_A = i]$. However, since we can still freely determine the "shape" of $M_{\cdot,j}$, we can make $\Pr[\mathbf{X_E}|M(X_A) = j]$ far away from $\Pr[\mathbf{X_E}|X_A = i] \approx \Pr[\mathbf{X_E}|M'(X_A) = j]$ even if $M \approx M'$. Fortunately, this bad case only happens when $\Pr_M[X_\sigma = j]$ is very small. We will show that the product $\Pr[X_\sigma = j]H_G(\Pr[\mathbf{X_E}|M(X_A) = j])$ is robust with respect to M.

The key property needed in the proof of contiousness is the convexity of function $G(x) : [0,1] \mapsto \mathbb{R}$. For a convex function $G(x)$, its derivative is a monotone function which implies that the absolute value of the derivative is maximized at the endpoints. That is why the values of $|G(\epsilon)|$ and $|G(1-\epsilon)|$ dominate the convergence rate.

5.2.2 Optimizing a Decomposable and Continuous Function

▶ **Definition 16.** $F : \mathcal{M}_{n \times *} \mapsto \mathbb{R}$ is $C(\epsilon, n)$-continuous if for all sufficiently small $0 < \epsilon < 1$, for every $M, M' \in \mathcal{M}_{n \times n}$, if $\max_{i,j}|M_{i,j} - M'_{i,j}| \le \epsilon$, then

$$|F(M') - F(M))| \le C(\epsilon, n).$$

▶ **Theorem 17.** *If $F : \mathcal{M}_{n \times *} \mapsto \mathbb{R}$ is decomposable, then*

$$\min_{M \in \mathcal{M}_{n \times *}} F(M) = \min_{M \in \mathcal{M}_{n \times n}} F(M) \qquad and \qquad \max_{M \in \mathcal{M}_{n \times *}} F(M) = \max_{M \in \mathcal{M}_{n \times n}} F(M).$$

Moreover, when F is $C(\epsilon, n)$-continuous, for all sufficiently small $0 < \epsilon < 1$, there exists an $O\left(LP(\frac{1}{\epsilon}+1)^n)\right)$ time algorithm that outputs $M^{-*}(\epsilon), M^{+*}(\epsilon) \in \mathcal{M}_{n \times n}$ such that

$$F(M^{-*}(\epsilon)) \le \min_{M \in \mathcal{M}_{n \times *}} F(M) + C(\epsilon, n) \qquad and \qquad F(M^{+*}(\epsilon)) \ge \max_{M \in \mathcal{M}_{n \times *}} F(M) - C(\epsilon, n)$$

where $LP(k)$ is the time complexity of linear programming with k variables.

Proof.

▶ **Claim 18.** *For any $M \in \mathcal{M}_{n \times *}$, there exists $M^-, M^+ \in \mathcal{M}_{n \times n}$ such that*

$$F(M^-) \le F(M) \le F(M^+).$$

---

[1] if $M_{\cdot,j} = (0, 0, ..., 0)^\top$, we define $\frac{M_{\cdot,j}}{S(M_{\cdot,j})}$ as $(0, 0, ..., 0)^\top$.

Part 1: Revelation Principle

The above claim directly implies the first result. It remains to show the claim. Let's construct the below linear program. Fix any $M \in \mathcal{M}_{n \times m}$,

$$\min_{\mathbf{x}} \sum_j x_j f(M_{\cdot, j}) \tag{1}$$

$$s.t. \sum_j x_j M_{\cdot, j} = [1, 1...1]_{1 \times n}^\top \tag{2}$$

$$x_j \geq 0, \forall j \tag{3}$$

Note that $\sum_j M_{\cdot, j} = [1, 1...1]_{1 \times n}^\top$ since M is a transition matrix. Thus, the above linear program must have a solution which implies that it must have a basic feasible solution (bfs) $\mathbf{x}^*$ [12]. Since $\mathbf{x}^*$ is a bfs, it must have at least $m - n$ zero entries. Therefore, there exists a size n subset $\{j_1, j_2, ..., j_n\}$ such that for any $j \notin \{j_1, j_2, ..., j_n\}$, $x_j^* = 0$.

Let $M^- = [x_{j_1}^* M_{j_1} \quad x_{j_2}^* M_{j_2} \quad ... \quad x_{j_n}^* M_{j_n}]$.
$\sum_k x_{j_k}^* M_{j_k} = \sum_j x_j^* M_{\cdot, j} = [1, 1...1]_{1 \times n}^\top$. Thus, M^- is a transition matrix. Moreover,

$$\begin{aligned}
F(M^-) &= F(\sum_k x_{j_k}^* M_{j_k}) \\
&= \sum_k x_{j_k}^* f(M_{j_k}) &\text{(Decomposability of } F) \\
&= \sum_j x_j^* f(M_{\cdot, j}) \leq \sum_j f(M_{\cdot, j}) = F(M) &\text{(For any } j \notin \{j_1, j_2, ..., j_n\}, x_j^* = 0)
\end{aligned}$$

Therefore, we finish our construction of M^-. The construction of M^+ is similar.

Part 2: LP Based Algorithm

The LP based algorithm for $\arg\min_{M \in \mathcal{M}_{n \times *}} F(M)$ is in Algorithm 1. Solving $\arg\max_{M \in \mathcal{M}_{n \times *}} F(M)$ is similar.

We can first enumerate all possible row vectors of $M \in \mathcal{M}_{n \times *}$ which are $(\frac{1}{\epsilon} + 1)^n$ in number. Then we can construct linear programming (1) with the all possible row vectors and solve the linear programming to obtain a basic feasible solution. ◀

5.3   Maximizing the Total Expected Payment

In this section, we will show that for any signaling scheme M, we can always perturb M a little bit to M' such that M' maximizes the optimal sum of expected payments of Alice and Bob.

The A-B-A game is a constant-sum game if Bob is required to reveal (announce) his full information to Alice, and Alice plays rationally in the final round. This is because the market belief will be changed from $\Pr[X_E]$ to $\Pr[X_E | X_A, X_B]$ regardless of Alice's signaling scheme. In this case, information is fully aggregated. Therefore when Bob is required to announce his full information, minimizing Bob's expected payment is equivalent to maximizing Alice's expected payment since ABA is a constant sum game.

However, when Bob is not required to announce his full information, minimizing Bob's expected payment is not equivalent to maximizing Alice's expected payment. For example, we consider the case where both X_A and X_B are i.i.d. binary bits (equal 1 with probability (w.p.) $\frac{1}{2}$, equal 0 w.p. $\frac{1}{2}$), and $X_E = X_A \oplus X_B$. To minimize Bob's expected payment, Alice

Algorithm 1 Minimizing Decomposable Function (MDF): $\arg\min_{M \in \mathcal{M}_{n \times *}} F(M)$

```
function MDF(f(·), n, ε)
N = 1/ε
enumerate all [ℓ₁/N, ℓ₂/N, ..., ℓₙ/N]
where ℓ₁, ℓ₂, ..., ℓₙ ∈ {0, 1, 2, ..., N} and denote them by {M·,ⱼ}ⱼ₌₁^(N+1)ⁿ

solve the following linear program and return a BFS x*


        minₓ ∑ⱼ xⱼ f(M·,ⱼ)
s.t.    ∑ⱼ xⱼ M·,ⱼ = [1, 1...1]ᵀ₁ₓₙ
        xⱼ ≥ 0, ∀j

pick a size n subset {j₁, j₂, ..., jₙ} s.t. for any j ∉ {j₁, j₂, ..., jₙ}, x*ⱼ = 0

return [x*ⱼ₁ M·,ⱼ₁  x*ⱼ₂ M·,ⱼ₂  ...  x*ⱼₙ M·,ⱼₙ]
```

should reveal no information in the first stage. However, in this case, rational Bob will not change the market belief and thus leaks no information of Bob's signal. In this case, the total payment for Alice and Bob is 0. On the other hand, if Alice commits a signaling scheme where with some small probability ϵ she announces X_A, and with probability $1 - \epsilon$ she flips a coin and announces the result. This signaling scheme entices Bob to move the market thus Alice can identify Bob's full information to guarantee almost optimal expected payment of Alice. Actually via a similar idea, for any signaling scheme M, we will construct signaling schemes $\{M(\epsilon)\}_\epsilon \approx M$ such that for all sufficiently small $\epsilon > 0$,

$$\mu^*(M(\epsilon)) = \mu^*.$$

Recall that n_A is the size of Alice's private information space, n_B is the size of Bob's private information space, n_E is the size the event X_E's outcome space and m is the size of Alice's announced signal space.

▶ **Lemma 19** (Perturbing Signaling Scheme). *Given the joint distribution over random variables X_A, X_B, X_E, consider an A-B-A game with commitment $(X_A, X_B, X_E, PS^G) \in \mathcal{PSR}$. For any signaling scheme M, for all sufficiently small $0 < \epsilon < 1$, we can always use $O(mn_A n_B^2 n_E)$ time* [2] *to perturb M to $M(\epsilon)$ such that*

$$\mu^*(M(\epsilon)) = \mu^* \qquad and \qquad \max_{x_A, \sigma} |M(x_A, \sigma) - M(\epsilon)(x_A, \sigma)| \leq \Theta(mn_B^2 \epsilon).$$

We defer the construction of the perturbing signaling scheme and the proof of Lemma 19 to Appendix A.

Now we are ready to give a full proof for the main theorem – Theorem 9 and Corollary 10.

Proof of Theorem 9. We construct the optimal signaling scheme of Alice using two steps.

Step 1: Minimizing Bob's expected payment we first use $O\left(\left(\frac{1}{\epsilon}\right)^T n_B n_E\right)$ time to construct the linear programming in Algorithm 1. We then spend $O\left(L\left(\frac{1}{\epsilon}\right)^T\right)$ time to solve the linear programming to obtain signaling scheme $M^\dagger(\epsilon) \in \mathcal{M}^{T \times T}$. Note that $\mu_B^*(M^\dagger(\epsilon)) \geq \mu_B^\dagger - \Theta(\epsilon + n_E |G(\epsilon)| + n_E |G(1 - \epsilon)|)$ based on Lemma 15 and the fact that we enumerate an ϵ-net of all possible column vectors in Algorithm 1.

---

[2] Note that the running time is independent with ϵ

Step 2: Maximizing the total expected payment: Lemma 19 shows for *any* signaling scheme M, there exists an $O(n_B^2 n_E)$ algorithm that constructs $M(\epsilon')$ such that

$$\max_{x_A, \sigma} |M(x_A, \sigma) - M(\epsilon')(x_A, \sigma)| \leq \Theta(n_B^2 \epsilon')$$

and $\mu^*(M(\epsilon')) = \mu^*$. Since the running time of the perturbing method is independent with ϵ', we can pick sufficiently small ϵ' such that $n_B^2 \epsilon' \leq \epsilon$. Thus, we can still use $O(n_B^2 n_E)$ time to perturb $M^\dagger(\epsilon)$ to $M^*(\epsilon)$ such that

$$\max_{x_A, \sigma} |M^\dagger(\epsilon)(x_A, \sigma) - M^*(\epsilon)(x_A, \sigma)| \leq \epsilon$$

and $\mu^*(M^*(\epsilon)) = \mu^*$.
We can see

$$
\begin{aligned}
\mu_A^*(M^*(\epsilon)) &= \mu^*(M^*(\epsilon)) - \mu_B^*(M^*(\epsilon)) \\
&= \mu^* - \mu_B^*(M^*(\epsilon)) \\
&\geq \mu^* - \mu_B^*(M^\dagger(\epsilon)) - \Theta(n_A^2(\epsilon + n_E|G(\epsilon)| + n_E|G(1 - \epsilon)|)) \quad \text{(Lemma 15)} \\
&\geq \mu^* - \mu_B^\dagger - \Theta(n_A^2(\epsilon + n_E|G(\epsilon)| + n_E|G(1 - \epsilon)|)) \quad \text{(step 1)} \\
&= \mu^* - \mu_B^\dagger - \Theta(\epsilon + n_E|G(\epsilon)| + n_E|G(1 - \epsilon)|) \\
&\qquad\qquad\qquad\qquad\qquad (n_A = T \text{ is a constant integer}) \\
&= \mu_A^* - \Theta(\epsilon + n_E|G(\epsilon)| + n_E|G(1 - \epsilon)|)
\end{aligned}
$$

When Alice commits to signaling scheme $M^*(\epsilon)$, according to Claim 11 and Claim 12,

$$\mathbf{p}_2^* = \Pr[\mathbf{X_E}|X_\sigma, X_B] \qquad \mathbf{p}_3^* = \Pr[\mathbf{X_E}|X_A, X_B]$$

where $X_A \xrightarrow{M^*(\epsilon)} X_\sigma$.   ◀

Proof of Corollary 10. $\tau = \Theta(|\epsilon| + n_E|G(\epsilon)| + n_E|G(1 - \epsilon)|) \leq \Theta(\epsilon^{\min\{\lambda, 1\}} n_E)$ implies that there exists a positive constant C such that when τ is sufficiently small, $\frac{1}{\epsilon} \leq C(\frac{n_E}{\tau})^{\frac{1}{\min\{1, \lambda\}}}$. Theorem 9 implies that we have a $poly(\frac{1}{\tau}, n_B, n_E)$ complexity algorithm to have $\pm\tau$ approximation.   ◀

6 Discussion

Our FPTAS result depends on the assumption that Alice's signal space is constant size. In fact, the complexity of our algorithm depends exponentially on the size of Alice's signal space. Note that before our results even the case where Alice's signal is binary was an open question. In practice, to apply our results one might reduce the signal space size by merging similar signals to hopefully obtain a good approximation.

A potential future direction is proving a hardness result for optimal information revelation problem in the ABA case when Alice's signal is arbitrarily large, or finding a better dependence on the size of Alice's signal space.

References

1   Yossi Azar, Amir Ban, and Yishay Mansour. When should an expert make a prediction? In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pages 125–142. ACM, 2016.

2   Yakov Babichenko and Siddharth Barman. Computational aspects of private bayesian persuasion. *arXiv preprint arXiv:1603.01444*, 2016.

3   Yiling Chen, Stanko Dimitrov, Rahul Sami, Daniel M Reeves, David M Pennock, Robin D Hanson, Lance Fortnow, and Rica Gonen. Gaming prediction markets: Equilibrium strategies with a market maker. *Algorithmica*, 58(4):930–969, 2010.

4   Yiling Chen, Daniel M Reeves, David M Pennock, Robin D Hanson, Lance Fortnow, and Rica Gonen. Bluffing and strategic reticence in prediction markets. In *International Workshop on Web and Internet Economics*, pages 70–81. Springer, 2007.

5   Yiling Chen and Bo Waggoner. Informational substitutes. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages 239–247. IEEE, 2016.

6   Thomas M. Cover and Joy A. Thomas. *Elements of information theory (2. ed.)*. Wiley, 2006.

7   Shaddin Dughmi and Haifeng Xu. Algorithmic bayesian persuasion. *arXiv preprint arXiv:1503.05988*, 2015.

8   Tilmann Gneiting and Adrian E Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.

9   Robin Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003.

10 Robin Hanson. Logarithmic markets coring rules for modular combinatorial information aggregation. *The Journal of Prediction Markets*, 1(1):3–15, 2012.

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

12 David G Luenberger. *Introduction to linear and nonlinear programming*, volume 28. Addison-Wesley Reading, MA, 1973.

13 Robert L Winkler. Scoring rules and the evaluation of probability assessors. *Journal of the American Statistical Association*, 64(327):1073–1078, 1969.

A   Proofs of Lemma 14, 15, 19

Proof of Lemma 14. We start to prove (i) and (ii).

Proof of (i)

Recall that if the distribution of X is represented by a row vector $\mathbf{p}$, then the distribution over $M(X)$ is $\mathbf{p}M$ and $\Pr[M(x) = j] = \mathbf{p} \cdot M_{\cdot,j}$.

Thus $\Pr[X_\sigma = j] = \Pr[M(X_A) = j] = \Pr[\mathbf{X_A}] \cdot M_{\cdot,j}$ is a linear function of the vector $M_{\cdot,j}$.

Proof of (ii)

For $\psi_j(M)$, since

$$
\begin{aligned}
\Pr[X_B = x_B, X_E = x_E | X_\sigma = j] &= \frac{\Pr[X_B = x_B, X_E = x_E, X_\sigma = j]}{\Pr[X_\sigma = j]} \\
&= \frac{\Pr[X_B = x_B, X_E = x_E, \mathbf{X_A}] \cdot M_{\cdot,j}}{\Pr[\mathbf{X_A}] \cdot M_{\cdot,j}} \\
&= \frac{\Pr[X_B = x_B, X_E = x_E, \mathbf{X_A}] \cdot \frac{M_{\cdot,j}}{S(M_{\cdot,j})}}{\Pr[\mathbf{X_A}] \cdot \frac{M_{\cdot,j}}{S(M_{\cdot,j})}}
\end{aligned}
$$

only depends on $\frac{M_{\cdot,j}}{S(M_{\cdot,j})}$. Since $\psi_j(M)$ only depends on the joint distribution over (X_B, X_E) conditioning on $X_\sigma = j$, thus, $\psi_j(M)$ only depends on $\frac{M_{\cdot,j}}{S(M_{\cdot,j})}$ as well.   ◀

Proof of Lemma 15. It is sufficient to show that if $M' = M$ except $M'_{ij} = M_{ij} + \epsilon$, $M'_{ik} = M_{ik} - \epsilon$,

$$
|\mu_B^*(M') - \mu_B^*(M)| \le \Theta(n_E|G(\epsilon)| + n_E|G(1-\epsilon)| + \epsilon).
$$

Note that when $G(x)$ is convex, for sufficiently small constant ϵ, $n_E|G(\epsilon)| + n_E|G(1-\epsilon)| + \epsilon$ is an increasing function of ϵ. Therefore, if $M' = M$ except $M'_{ij} = M_{ij} + \epsilon'$, $M'_{ik} = M_{ik} - \epsilon'$ where $\epsilon' \le \epsilon$, we still have

$$
|\mu_B^*(M') - \mu_B^*(M)| \le \Theta(n_E|G(\epsilon)| + n_E|G(1-\epsilon)| + \epsilon).
$$

The proof of Lemma 14 shows that $\mu_B^*(M)$ can be decomposed as m parts and the only part that relates to $M_{\cdot,j}$ is

$$
\Pr[M(X_A) = j]\mathbb{E}_{X_B|M(X_A)=j}[H_G(\Pr[\mathbf{X_E}|M(X_A) = j, X_B]) - H_G(\Pr[\mathbf{X_E}|M(X_A) = j])]
$$

We define $\gamma(M_{\cdot,j}) := \Pr[M(X_A) = j]H_G(\Pr[\mathbf{X_E}|M(X_A) = j])$ and would like to the below claim.

▶ **Claim 20.** *When $M'_{\cdot,j} = M_{\cdot,j}$ except $M'_{ij} = M_{ij} + \epsilon$,*

$$
|\gamma(M_{\cdot,j}) - \gamma(M'_{\cdot,j})| \le \Theta(n_E|G(\epsilon)| + n_E|G(1-\epsilon)| + \epsilon).
$$

The above claim is valid for every possible joint distribution over X_E, X_A. The set of all possible joint distributions over X_E, X_A equals the set of all possible joint distributions over X_E, X_A conditioning X_B for any X_B. Therefore, the above claim also implies that part $\Pr[M(X_A) = j]\mathbb{E}_{X_B|M(X_A)=j}H_G(\Pr[\mathbf{X_E}|M(X_A) = j, X_B])$ also only fluctuates at most $\Theta(n_E|G(\epsilon)| + n_E|G(1-\epsilon)| + \epsilon)$ when we perturb M_{ij} at most ϵ. We have similar analysis for $M_{\cdot,k}$, therefore, Lemma 15 follows.

It remains to show the claim.

Proof of Claim 20. Without loss of generality we assume $\epsilon \ge 0$, otherwise we can exchange $M_{\cdot,j}$ and $M'_{\cdot,j}$. In the proof, we will fix ϵ as a small constant and figure out the worst case of the joint distribution over X_A, X_E and original scheme M such that $|\gamma(M_{\cdot,j}) - \gamma(M'_{\cdot,j})|$ is maximized.

Recall that $\gamma(M_{\cdot,j}) = \Pr[M(X_A) = j]H_G(\Pr[\mathbf{X_E}|M(X_A) = j])$.

Part 1

We will first figure out the explicit relationship between ϵ, $\gamma(M'_{\cdot,j})$ and $\gamma(M_{\cdot,j})$.

$$\Pr[M(X_A) = j] = \sum_{x_A} \Pr[X_A = x_A] M_{x_A,j} = \Pr[X_A = i] M_{ij} + K(-i)$$

Thus, $\Pr[M'(X_A) = j] = \Pr[M(X_A) = j] + \epsilon \Pr[X_A = i]$.

$$
\begin{aligned}
\mathbf{q} := \Pr[\mathbf{X_E}|M(X_A) = j] =& \frac{\Pr[\mathbf{X_E}, M(X_A) = j]}{\Pr[M(X_A) = j]} \\
=& \frac{\sum_{x_A} \Pr[X_A = x_A] M_{x_A,j} \Pr[\mathbf{X_E}|X_A = x_A]}{\Pr[M(X_A) = j]} \\
=& \frac{\Pr[X_A = i] M_{ij} \Pr[\mathbf{X_E}|X_A = i] + K(-i)\mathbf{q}_{-i}}{\Pr[M(X_A) = j]} \\
& (\mathbf{q}_{-i} \text{ is a distribution over } X_E \text{ that is independent with } M_{ij}) \\
=& \frac{\Pr[X_A = i] M_{ij}\mathbf{q}_i + K(-i)\mathbf{q}_{-i}}{\Pr[M(X_A) = j]} \qquad (\mathbf{q}_i := \Pr[\mathbf{X_E}|X_A = i])
\end{aligned}
$$

Thus, $\Pr[\mathbf{X_E}|M'(X_A) = j] = \mathbf{q}_i \frac{\epsilon \Pr[X_A=i]}{\Pr[M(X_A)=j]+\epsilon \Pr[X_A=i]} + \mathbf{q}\frac{\Pr[M(X_A)=j]}{\Pr[M(X_A)=j]+\epsilon \Pr[X_A=i]}$.
$\Pr[\mathbf{X_E}|M'(X_A) = j]$ is a convex combination of the original forecast for X_E – $\mathbf{q} = \Pr[\mathbf{X_E}|M(X_A) = j]$ and the posterior forecast for X_E conditioning on $X_A = i$.
We define $\tau(\epsilon) := \frac{\epsilon \Pr[X_A=i]}{\Pr[M(X_A)=j]+\epsilon \Pr[X_A=i]}$. Note that

$$\tau^* := \frac{\epsilon \Pr[X_A = i]}{1 + \epsilon \Pr[X_A = i]} \le \tau(\epsilon) \le 1 \tag{4}$$

$\tau^* \le \frac{\epsilon}{1+\epsilon}$.
Then we have $\Pr[\mathbf{X_E}|M'(X_A) = j] = \tau(\epsilon)\mathbf{q}_i + (1 - \tau(\epsilon))\mathbf{q}$. By replacing $\Pr[M(X_A) = j]$ and $\Pr[M'(X_A) = j]$ by $\tau(\epsilon)$ and $\Pr[X_A = i]$, we obtain

$$\gamma(M'_{\cdot,j}) = \epsilon \Pr[X_A = i]\frac{1}{\tau(\epsilon)} H_G(\tau(\epsilon)\mathbf{q}_i + (1 - \tau(\epsilon))\mathbf{q})$$

$$\gamma(M_{\cdot,j}) = \epsilon \Pr[X_A = i]\frac{1 - \tau(\epsilon)}{\tau(\epsilon)} H_G(\mathbf{q})$$

Part 2

We start to calculate the worst $\tau(\epsilon)$, $\mathbf{q}$ and $\mathbf{q}_i$ that maximizes the gap between $\gamma(M_{\cdot,j})$ and $\gamma(M'_{\cdot,j})$. We first tune $\tau(\epsilon)$ and then tune $\mathbf{q}$ and $\mathbf{q}_i$.

Part 2: Tuning $\tau(\epsilon)$

In this part, we focus on $\tau(\epsilon)$ and omit other variables. We define $g(\tau) := H_G(\tau\mathbf{q}_i + (1-\tau)\mathbf{q})$. Note that $g(\tau)$ is a convex function since H_G is a convex function and $|g(\tau)|$ is also bounded by 1 (Remark 2.1).

Then we have

$$\gamma(M'_{\cdot,j}) = \epsilon \Pr[X_A = i]\frac{1}{\tau(\epsilon)}g(\tau(\epsilon)) \qquad \gamma(M_{\cdot,j}) = \epsilon \Pr[X_A = i]\frac{1 - \tau(\epsilon)}{\tau(\epsilon)}g(0)$$

$$|\gamma(M'_{.,j}) - \gamma(M_{.,j})| = |\epsilon \Pr[X_A = i](1 - \tau(\epsilon))\frac{1}{\tau(\epsilon)}[g(\tau(\epsilon)) - g(0)] + \epsilon \Pr[X_A = i]g(\tau(\epsilon))|$$

$$\leq |\epsilon \Pr[X_A = i]\frac{g(\tau(\epsilon)) - g(0)}{\tau(\epsilon) - 0} + |\epsilon \Pr[X_A = i]|$$

$$(1 - \tau(\epsilon) \leq 1, \; g(x) \text{ is bounded by 1 since } |H_G(\mathbf{p})| \text{ is bounded by 1 (Remark 2.1).})$$

Recall that

$$\tau^* := \frac{\epsilon \Pr[X_A = i]}{1 + \epsilon \Pr[X_A = i]} \leq \tau(\epsilon) \leq 1 \tag{5}$$

Note that $\frac{g(\tau) - g(0)}{\tau - 0}$ is an increasing function of $0 \leq \tau \leq 1$ when g is a convex function, thus, $|\frac{g(\tau) - g(0)}{\tau - 0}|$ is maximized in the endpoints,

$$|\gamma(M'_{.,j}) - \gamma(M_{.,j})| \leq \epsilon \Pr[X_A = i]|\frac{g(\tau) - g(0)}{\tau - 0}| + \epsilon$$

$$\leq \epsilon \Pr[X_A = i] \max\{|\frac{g(\tau^*) - g(0)}{\tau^* - 0}|, |\frac{g(1) - g(0)}{1 - 0}|\} + \epsilon$$

$$\leq 2|[g(\tau^*) - g(0)]| + 2\epsilon$$

$$(g(x) \text{ is bounded by 1 since } |H_G(\mathbf{p})| \text{ is bounded by 1 (Remark 2.1).})$$

Part 2: Tuning q and $\mathbf{q}_i$.

It remains to compute the upper-bound of $|g(\tau^*) - g(0)|$.

$$\max |g(\tau^*) - g(0)|$$

$$= \max_{\mathbf{q}_i, \mathbf{q}} |H_G(\tau^* \mathbf{q}_i + (1 - \tau^*)\mathbf{q}) - H_G(\mathbf{q})|$$

$$\leq \max_{\mathbf{q}_i, \mathbf{q}} \sum_{\sigma \in \Sigma_E} |G(\tau^* \mathbf{q}_i(\sigma) + (1 - \tau^*)\mathbf{q}(\sigma)) - G(\mathbf{q}(\sigma))|$$

Consider $h(x, y) := |G(\tau^* x + (1 - \tau^*)y) - G(y)|$.

$$|G(\tau^* x + (1 - \tau^*)y) - G(y)| = |G(\tau^*(x - y) + y) - G(y)|$$

Fix $\tau^*(x - y)$. Note that $\tau^*(x - y)$ can be less than the any small constant by picking sufficiently small constant ϵ. Since $G(x)$ (Figure 1) is a convex function, $G'(x)$ is a monotone function. Thus, $|G'(x)|$ is maximized in end points. $h(x, y) = |G(\tau^*(x - y) + y) - G(y)|$ is maximized if y is one of the end points, that is, $y = 0, 1$.

When $y = 0, 1$, $h(x, y) = |G(\tau^*(x - y) + y) - G(y)| = |G(\tau^*(x - y) + y)|$. Note that $|G(x)|$ and $|G(1 - x)|$ are increasing functions when x is sufficiently small when $G(x)$ is convex. Thus, to maximize $h(x, y)$, we should pick $|x - y| = 1$ to maximize $|\tau^*(x - y)|$. Therefore, $|G(\tau^*(x - y) + y)| \leq \max\{|G(\tau^*)|, |G(1 - \tau^*)|\} \leq |G(\tau^*)| + |G(1 - \tau^*)|$.

$$\max |g(\tau^*) - g(0)|$$

$$\leq \max_{\mathbf{q}_i, \mathbf{q}} \sum_{\sigma \in \Sigma_E} |G(\tau^* \mathbf{q}_i(\sigma) + (1 - \tau^*)\mathbf{q}(\sigma)) - G(\mathbf{q}(\sigma))|$$

$$\leq n_E \max_{x, y} h(x, y)$$

$$\leq n_E(|G(\tau^*)| + |G(1 - \tau^*)|)$$

Recall that $\tau^* := \frac{\epsilon \Pr[X_A = i]}{1 + \epsilon \Pr[X_A = i]} \le \epsilon,$

$$
\begin{aligned}
|\gamma(M'_{\cdot,j}) - \gamma(M_{\cdot,j})| &\le 2|g(\tau^*) - g(0)| + 2\epsilon \\
&\le 2n_E(|G(\tau^*)| + |G(1 - \tau^*)|) + 2\epsilon \\
&\le 2n_E(|G(\epsilon)| + |G(1 - \epsilon)|) + 2\epsilon \\
&\text{(when } x \text{ is small, both } |G(x)| \text{ and } |G(1 - x)| \text{ are increasing functions.)}
\end{aligned}
$$

◄

◄

Proof of Lemma 19.

▶ **Definition 21** (Bad pair). We say a pair $(x_B, x'_B), x_B, x'_B \in \Sigma_B, x_B \ne x'_B$ is bad for a event E if

$$\Pr[X_E = \cdot | E, X_B = x_B] = \Pr[X_E = \cdot | E, X_B = x'_B].$$

Intuitively, a bad pair cannot be distinguished conditioning on event E.

According to Claim 12, to guarantee the total expected payment being maximal, we should guarantee the information is fully aggregated, that is, Alice should identify all "meaningful" information from Bob. If there exists a pair $(x_B, x'_B), x_B, x'_B \in \Sigma_B, x_B \ne x'_B$ such that

$$\Pr[X_E = \cdot | X_A = x_A, X_B = x_B] = \Pr[X_E = \cdot | X_A = x_A, X_B = x'_B],$$

then distinguishing the event $X_B = x_B$ and the event $X_B = x'_B$ is not meaningful for Alice since she can think of x_B, x'_B as one signal without losing any information. Therefore, without loss of generality, *we assume every pair (x_B, x'_B) is good for at least one $X_A = x_A$.*

We start our construction for $M(\epsilon)$. Adopting $M(\epsilon)$ helps Alice identify all meaningful information of Bob means for every announced signal σ, there does not exist any *bad* pair $(X_B = x_B, X_B = x'_B)$ conditioning on Alice announced $X_\sigma^\epsilon = \sigma$ where $X_\sigma^\epsilon := M(\epsilon)(X_A)$.

We define $X_\sigma := M(X_A)$. We will show by changing each entry of M at most ϵ at most, we can fix a bad pair and will not produce more bad pairs. In the end, we fix all bad pairs and construct $M(\epsilon)$ that satisfies (i) and (ii).

For a bad pair (x_B, x'_B) conditioning on that event that Alice announces $X_\sigma = \sigma$ in stage 1. We know

$$
\begin{aligned}
1 &= \frac{\Pr[X_E = x_E | X_B = x_B, X_\sigma = \sigma]}{\Pr[X_E = x_E | X_B = x'_B, X_\sigma = \sigma]} \\
&= \frac{\Pr[X_E = x_E, X_B = x_B, X_\sigma = \sigma] \Pr[X_B = x'_B, X_\sigma = \sigma]}{\Pr[X_E = x_E, X_B = x'_B, X_\sigma = \sigma] \Pr[X_B = x_B, X_\sigma = \sigma]} \qquad (6) \\
&= \frac{\sum_{X_A = x_A} \Pr[X_E = x_E, X_B = x_B, X_A = x_A] M(x_A, \sigma)}{\sum_{X_A = x_A} \Pr[X_E = x_E, X_B = x'_B, X_A = x_A] M(x_A, \sigma)} \\
&\quad \times \frac{\sum_{X_A = x_A} \Pr[X_B = x'_B, X_A = x_A] M(x_A, \sigma)}{\sum_{X_A = x_A} \Pr[X_B = x_B, X_A = x_A] M(x_A, \sigma)} \qquad (7)
\end{aligned}
$$

for all x_E.

Note that we have assumed that for every pair is good for at least one $X_A = x_A$. Suppose (x_B, x'_B) is good for $X_A = x_A$. If we add an ϵ in entry $M(x_A, \sigma)$ (if $M(x_A, \sigma) > 1 - \epsilon$ we can subtract an ϵ and the analysis is similar. We can pick sufficiently small ϵ to guarantee either $0 \le M(x_A, \sigma) + \epsilon \le 1$ or $0 \le M(x_A, \sigma) - \epsilon \le 1$) and tune other $M(x_A, \sigma'), \sigma' \ne \sigma$ arbitrarily

■ **Table 1** Experiment inputs

	Left input		Middle input		Right input	
$\Pr[X_A = 1, \mathbf{X_B}, \mathbf{X_E}]$	0.0900	0.0900	0.2209	0.0947	0.0100	0.0100
	0.0900	0.1800	0.0947	0.0199	0.0100	0.0200
$\Pr[X_A = 2, \mathbf{X_B}, \mathbf{X_E}]$	0.0900	0.0900	0.0947	0.0406	0.0100	0.0100
	0.0900	0.2800	0.0406	0.3942	0.0100	0.9200

such that $M(x_A, \cdot)$ remains to be a valid distribution over Σ, and denote the new signaling scheme as M' and define $X'_\sigma := M'(X_A)$, we will know

$$\frac{\Pr[X_E = x_E | X_B = x_B, X'_\sigma = \sigma]}{\Pr[X_E = x_E | X_B = x'_B, X'_\sigma = \sigma]}$$

$$= \frac{(\Pr[X_E = x_E, X_B = x_B, X_\sigma = \sigma] + \epsilon \Pr[X_E = x_E, X_B = x_B, X_A = x_A])}{(\Pr[X_E = x_E, X_B = x'_B, X_\sigma = \sigma] + \epsilon \Pr[X_E = x_E, X_B = x'_B, X_A = x_A])} \tag{8}$$

$$\cdot \frac{(\Pr[X_B = x'_B, X_\sigma = \sigma] + \epsilon \Pr[X_B = x'_B, X_A = x_A])}{(\Pr[X_B = x_B, X_\sigma = \sigma] + \epsilon \Pr[X_B = x_B, X_A = x_A])} \qquad \text{(according to formula (6))}$$

Since (x_B, x'_B) is bad for $X_\sigma = \sigma$ but good for $X_A = x_A$, there exists x_E such that

$$\frac{\Pr[X_E = x_E, X_B = x_B, X_\sigma = \sigma] \Pr[X_B = x'_B, X_\sigma = \sigma]}{\Pr[X_E = x_E, X_B = x'_B, X_\sigma = \sigma] \Pr[X_B = x_B, X_\sigma = \sigma]} = 1$$

$$\frac{\Pr[X_E = x_E, X_B = x_B, X_A = x_A] \Pr[X_B = x'_B, X_A = x_A]}{\Pr[X_E = x_E, X_B = x'_B, X_A = x_A] \Pr[X_B = x_B, X_A = x_A]} \neq 1$$

Thus there exists constant $\lambda \neq 0$ and μ such that the difference between the denominator and numerator of formula (8) can be represented as

$$\lambda \epsilon^2 + \mu \epsilon$$

Since $\lambda \neq 0$, we can always pick a sufficiently small $\epsilon_0 > 0$ such that $\lambda \epsilon^2 + \mu \epsilon \neq 0$ for all $0 < \epsilon < \epsilon_0$. For good pair (y_B, y'_B) for signal σ',

$$\frac{\Pr[X_E = x_E, X_B = y_B, X_\sigma = \sigma'] \Pr[X_B = y'_B, X_\sigma = \sigma']}{\Pr[X_E = x_E, X_B = y'_B, X_\sigma = \sigma'] \Pr[X_B = y_B, X_\sigma = \sigma']} \neq 1$$

the difference between the denominator and numerator of good pair's formula will change from $c \neq 0$ to $c + \Theta(\epsilon)$ which can be non-zero as well when ϵ is sufficiently small. Thus, we won't produce more bad pairs.

The time depends on the number of bad pairs we need to fix since we fix them one by one. The number of bad pairs is at most $O(mn_B^2)$. Thus we need $O(mn_A n_B^2 n_E)$ time to finish the construction and $\max_{x_A, \sigma} |M(x_A, \sigma) - M(\epsilon)(x_A, \sigma)| \leq \Theta(mn_B^2 \epsilon)$. ◀

B   Experiment Inputs

The proper scoring rule we use in the experiment is the logarithmic scoring rule. Now we give the input joint distributions over X_A, X_B, X_E for the three examples. All of X_A, X_B, X_E are binary random variables in this case. Therefore, we can show the input joint distribution via two 2×2 matrices. $U_1 = \Pr[X_A = 1, \mathbf{X_B}, \mathbf{X_E}]$ is a matrix where $U_1(i, j) = \Pr[X_A = 1, X_B = i, X_E = j]$ and $U_2 = \Pr[X_A = 2, \mathbf{X_B}, \mathbf{X_E}]$ is a matrix where $U_2(i, j) = \Pr[X_A = 2, X_B = i, X_E = j]$.