Timely Information from Prediction Markets

Grant Schoenebeck*
University of Michigan, Ann Arbor
schoeneb@umich.edu

Chenkai Yu
Tsinghua University, Beijing
yck17@mails.tsinghua.edu.cn

Fang-Yi Yu†
Harvard University
fangyiyu@seas.harvard.edu

ABSTRACT

Prediction markets are powerful tools to elicit and aggregate beliefs from strategic agents. However, in current prediction markets, agents may exhaust the social welfare by competing to be the first to update the market. We initiate the study of the trade-off between how quickly information is aggregated by the market, and how much this information costs. We design markets to aggregate timely information from strategic agents to maximize social welfare. To this end, the market must incentivize agents to invest the correct amount of effort to acquire information: quickly enough to be useful, but not faster (and more expensively) than necessary. The market also must ensure that agents report their information truthfully and on time. We consider two settings: in the first, information is only valuable before a deadline; in the second, the value of information decreases as time passes. We use both theorems and simulations to demonstrate the mechanisms.

KEYWORDS

Prediction Markets; Social Welfare; Timely Reports; Costly Information

ACM Reference Format:
Grant Schoenebeck, Chenkai Yu, and Fang-Yi Yu. 2021. Timely Information from Prediction Markets. In *Proc.of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), Online, May 3–7, 2021*, IFAAMAS, 9 pages.

1 INTRODUCTION

Eliciting information about uncertain events is crucial for informed decision making. Information is often acquired by individual agents. To achieve collective intelligence, the key problems are to elicit and aggregate *timely* and truthful information from dispersed agents.

Prediction markets (PMs) allow agents to bet on the occurrence of future events: the outcome of a presidential election, the winner of a football game, etc. Market prices reflect society's aggregated estimate of the outcome. However, prediction markets tend to only pay the *first* agent bringing information to the market. For example, in the market of a tennis match, if one player wins a set, likely the price of the market will shift dramatically. Because sportscasts are usually delayed by a few seconds, agents with real-time information (say with a confederate attending the match) reap the rewards by trading just seconds before others. This provides little to no societal value as the price would be updated seconds later anyhow. This practice is widespread. People, called "courtsiders", are paid to attend sports events and send back real-time information [10, 13]. The US Open ejected 20 spectators for courtsiding and banned them from future events [24]. While this information is useful to those profiting from it, it is hardly more than a waste of money and time for our society.

The importance of speed in trading on information is underscored by the infamous $300 million, 827 mile fiber optic cable from New York City to Chicago. The cable reduced the round-trip latency to 13.1 milliseconds [1] by offering a more direct route, bypassing Philadelphia, than the previous 1000 mile cable took with round trip latency 14.5 milliseconds.

We want to design systems that work well for society rather than promoting speed that is not needed but is merely a byproduct of the market design. Increasingly, businesses are seen not merely as profit-maximizing, but as responsible and responsive to their various stakeholders [25].

Maximizing welfare rather than profit has various economic motivations as well. To gain market share and maximize long-term revenue, a company may also want to benefit other companies that it makes deals with. For example, many sponsored search auctions maximize welfare instead of revenue [22]. Since profit is a lower bound of social welfare (when agents' utility is non-negative), higher welfare potentially leads to high profit. Welfare-maximizing auctions and profit-maximizing auctions are shown to be very close in terms of both welfare and efficiency [2, 7]. In our settings, though our mechanism maximizes social welfare, its profit is still high, illustrated in Figure 7.

Apart from social welfare concerns, another potential challenge with prediction markets is that agents may want to delay reporting their information to increase their rewards [4, 9, 20]. We show that in the settings we study, this is still a problem for traditional prediction markets, and we resolve this problem in our mechanisms.

1.1 Our Contribution

Motivated by the above concerns, we answer the following question: How to aggregate *timely* and truthful information to maximize social welfare?

We formulate the process as a principal-agent problem. The principal first suggests a contract which maps agents' reports to rewards. Then each agent decides his hidden actions (how much effort and how to report) strategically to maximize his utility. The principal knows neither agents' actions nor the relation between agents' actions and quality of information. This makes our problem different from those in standard contract theory. To resolve this, we design two new markets to maximize social welfare under the following two settings.

*Grant Schoenebeck is pleased to acknowledge the support of the National Science Foundation under grants NSF 1618187 and 2007256.
†Fang-Yi Yu is pleased to acknowledge the support of the National Science Foundation under grants NSF 1618187, NSF 2007256, CCF-1718549 and IIS-2007887.

Proc.of the 20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2021), U. Endriss, A. Nowé, F. Dignum, A. Lomuscio (eds.), May 3–7, 2021, Online.

In the single batch setting, a principal needs to decide by a particular deadline. To maximize social welfare, the principal needs to incentivize agents to invest the correct cost. Agents' costs are hidden, and they can misreport. For this setting, we propose the Fair Prediction Market (FPM, Mechanism 1), in which the expected reward for every truthful agent is the same.

In the sequential setting, the value of information decreases as time passes. Besides dealing with agents' hidden costs and misreporting behaviors, the principal also needs to encourage timely reports. We propose the Marginal Value Prediction (MVP) Market (Mechanism 2), in which every agent is paid by his contribution to the value of information.

Compared to the traditional prediction market, our mechanisms have more desirable properties, as shown in Table 1.

Table 1: Comparing the traditional prediction market (PM) and our mechanisms.

	PM	FPM	MVP Market
Timing	sequential	single-batch	sequential
Truthfulness	✓	✓	✓
Timeliness		N.A.	✓
Social Optimality		✓	✓

1.2 Related Works

One line of works studies when agents should report their signals in prediction markets [4, 9, 15, 20]. They find that whether information will be aggregated quickly depends on agents' information structure. Agents will delay reporting if their information is "complementary", and rush to report if it is "substitutional". Earlier in the finance literature, [14, 19, 21] analyze how private information is disseminated into real financial markets using different models. The market behavior depends on the numbers of insiders, noise traders, and market makers.

Chakraborty and Das [8] find that when agents are risk-averse, the market scoring rule acts as an opinion pool. Agents' risk aversion avoids the issue that agents will always pull the market price toward their own belief without ever reaching a consensus. We do not have this issue because, in our model, agents believe that the signals of each other are useful and are willing to do Bayesian updates.

For costly information, if the effort level is binary, it's well-known that we can scale the reward to compensate for the cost of effort. This approach encourages agents to invest effort and increases the liquidity of the prediction [23]. In our paper, we consider a more complicated setting: The effort level is continuous, and the relation between agents' actions and the quality of information is unknown. Azar and Micali [3] uses contract theory to delegate computation where acquiring data is costly, but assume the principal has some ability to verify the data. Moreover, our goal is to maximize social welfare, not just collect accurate decisions.

Budish et al. [6] investigate the continuous limit order book market and the high-frequency trading arms race. They show examples that such arms races induce rents that harm the liquidity of the market. They propose a frequent batch market which discretizes the time to mitigate the necessity to be first, and show that the above-mentioned example does not hold in their new markets.

1.3 Outline.

In Section 2, we provide some basic notations, assumptions, and definitions; frame the problem we want to solve; and describe how prediction markets work. In Section 3, we show how prediction markets may fail to collect timely reports. In Section 4, we propose Mechanism 1 for the single-batch setting and show that it is truthful and maximizes the social welfare. In Section 5, we propose Mechanism 2 for the sequential setting and show that it is truthful, timely and maximizes social welfare. At last, we present some concrete examples to compare our mechanism to prediction markets in Section 6.

2 PRELIMINARIES

There is a ***principal*** and a set of ***agents*** $\mathcal{N} = \{1, \dots, n\}$.[1] Let $\mathcal{Y}$ be the outcome space, $y \in \mathcal{Y}$ be the true outcome, and Y be the random variable for the outcome. For each agent $i \in \mathcal{N}$, let X_i be the private information of agent i, and $\mathcal{X}_i$ be the set of possible values of X_i where $\mathcal{X}_i$ is finite. The principal wants to collect information from agents to better predict y, and her utility depends on the value of information.[2]

2.1 Information Structure

We assume agents' signals $X_1, \dots, X_n$ are i.i.d. conditioning on the outcome Y. Every agent knows the joint distribution $\Pr[X_k, Y]$. The principal knows the prior of the outcome $\Pr[Y]$, but she may not know $\Pr[X_k \mid Y]$.

Example 2.1. Consider a binary outcome space $\mathcal{Y} = \{0, 1\}$, and binary signal spaces $\mathcal{X}_i = \{0, 1\}$ for all $i \in \mathcal{N}$. The prior of the outcome is given by $\Pr[Y = 1] = \alpha$. There is a noise level $\beta \in [0, 1/2)$ such that each signal is an independent noisy observation of the outcome $\Pr[X_i = Y] = 1 - \beta$ for all $i \in \mathcal{N}$ and $y \in \{0, 1\}$. Suppose the principal wants to predict Y and her utility is 1 (-1) if correct (incorrect). Let $\alpha = 1/2$. If she only knows the prior of the outcome, her expected utility is 0. However, if all agents collectively provide their prediction $\Pr[Y \mid X_1, \dots, X_n]$, the principal's utility is greater than 0. For example, if $n = 1$, her expected utility is $1 - 2\beta > 0$.

2.2 The Mechanism Design Problem

Acquiring signals is costly for agents. Hence, agents may not bother to invest effort and may misreport their information. The principal needs to incentivize agents to invest some effort to acquire their signals, without observing how much effort they actually invest. Formally, the mechanism has three stages:

(1) The principal publishes a contract which maps reports and the outcome to payments. Agents accept (or refuse).

[1] Throughout the paper, we use 'she' and 'he' for the principal and agents respectively.
[2] In reality, it could be other people who value the information in the market, use it to do something outside the market, and then get utility based on the quality of the information. Without loss of generality, we simply aggregate all such utility into that of the principal. [11, 16]

(2) Each agent chooses a hidden effort level and submits a report to the principal.
(3) The true outcome (that agents are guessing at) is revealed. The principal rewards each agent according to the contract.

Each agent's utility is his reward from the principal minus his effort level. Agents are rational and maximize their expected utility based on their beliefs over future events. The principal's utility is the value of information minus the rewards given to the agents. The ***social welfare*** is the total utility of all agents and the principal.

We want to design mechanisms whose resulting social welfare equals that in the *centralized* setting, where every agent's action is controlled by the principal. In particular, we hope that the agents invest the same amount of effort and report just the same as in the centralized setting.

Throughout the paper, we are only looking for symmetric equilibrium. This is reasonable since agents' signals are i.i.d. and they have no a priori means to coordinate. However, it is important to know that a non-symmetric strategy profile can be better than a symmetric one in terms of social welfare. For instance, as shown in Figure 1, social welfare may decrease as the number of agents increases, a phenomenon shared across a wide range of economic models: the tragedy of commons, the game of chicken, etc. In such cases, if we break the symmetry and only allow a restricted number of agents to participate, the social welfare will increase.

At each time during the whole process, for each agent, given the information he has and the strategies of other agents, he has a belief over all uncertainty — including the randomness of the world and other agents' private information. His immediate strategy should maximize his expected utility given his current information. If a strategy profile satisfies the above property, it is called a ***perfect Bayesian equilibrium***.[3]

We say a mechanism is ***individually rational*** if every agent has non-negative expected utility in every Nash equilibrium.

2.3 Prediction Markets with Scoring Rules

Scoring rules have a very long history [5, 12, 16, 17]. Market scoring rules were introduced by Hanson [18] to study prediction markets. A *scoring rule* for an outcome y is a function $S : \Delta_{\mathcal{Y}} \times \mathcal{Y} \to \mathbb{R}$,[4] so that $S(p, y)$ is the score assigned to a prediction p when the true outcome realized is $Y = y$. Formally,

Definition 2.2 (Proper Scoring Rule). $S : \Delta_{\mathcal{Y}} \times \mathcal{Y} \to \mathbb{R}$ is called a ***proper scoring rule*** if for any $b, p \in \Delta_{\mathcal{Y}}$,

$$\mathop{\mathbb{E}}_{Y\sim b}[S(b, Y)] \geq \mathop{\mathbb{E}}_{Y\sim b}[S(p, Y)].$$

A proper scoring rule is *strict* if the equality holds only if $b = p$. In other words, reporting one's belief results in a higher score than reporting other distributions. Such a report is said to be *truthful*.

[3]Here is a more rigorous definition. In an extensive-form game, for each game state h, let $\pi^\sigma(h)$ denote its reach probability according to strategy profile σ. We simply use $\pi(h)$ if σ is clear from the context. For each information set I, let $\pi(I) = \sum_{h\in I} \pi(h)$. Let $u^\sigma(h)$ denote h's expected utility according to σ. For each information set I with $\pi(I) > 0$, let $u(I) = \mathbb{E}_{h\in I}[u(h)] = \sum_{h\in I} u(h)\pi(h)/\pi(I)$. For a strategy profile σ, let $\sigma_{I\to s}$ denote the same strategy profile except the strategy at information set I is changed to s. A strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ is a ***perfect Bayesian equilibrium*** if for every information set I with $\pi(I) > 0$, for every strategy s at I, $u^\sigma(I) \geq u^{\sigma_{I\to s}}(I)$.

[4]$\Delta_{\mathcal{Y}}$ is the set of all probability distributions over $\mathcal{Y}$.

Prediction markets with a scoring rule proceed as follows: A public belief $p \in \Delta_{\mathcal{Y}}$ is maintained in the market. Initially, $p = p_0$. Agents can change the market belief arbitrarily, resulting in a belief sequence $p_0, p_1, p_2, \ldots$. After the outcome is revealed, for every $k \geq 1$, the agent who changes p_{k-1} to p_k is rewarded (or charged, if negative) by

$$S(p_k, y) - S(p_{k-1}, y),$$

where S is a proper scoring rule.

In prediction markets, if each agent is only allowed to report once (and they believe other agents are rational), then the reports will be truthful. In particular, each agent will report his posterior distribution using the market information as the prior and his private information as the evidence.

Proposition 2.3. *If signals are independent conditioned on the outcome, then agents can perform a Bayesian update knowing only the current market belief (and the likelihood of their own signals). In particular, they need not know the history of updates.*

3 LIMITATIONS OF PREDICTION MARKETS

In this section, we show by examples how the original form of prediction markets may perform undesirably: agents may 1) invest too much effort, which decreases the social welfare, and 2) intentionally delay their reports.

3.1 Inflated Effort and Poor Welfare

Current theories of prediction markets do not consider agents' effort to discover signals: How much effort should an agent invest? For some easy, accurate information, it would be a waste of resources if everyone invests a lot.

We model the cost of information by an ***access function*** $F : \mathbb{R}_{\geq 0} \to [0, 1]$ that maps from an effort level (the cost that an agent spends) to the probability of getting signals. Each agent i decides his effort level c_i and then obtains his signal X_i with probability $F(c_i)$. We assume the access function is the same for every agent as common knowledge but unknown to the principal.

Suppose agents' signals are structured as in Example 2.1 with $\alpha = 1/2$ and $\beta = 0$, i.e., all X_i's are identical to the outcome and are exact substitutes to each other. Also, only the first agent who changes the belief will get one unit of reward, and others will get zero. In this case, every agent wants to be the first. If multiple agents get signals, we assume each of them receives the reward with equal probability since they are symmetric.

Suppose the principal's value of information is described by a proper scoring rule S. She needs to choose a proper scoring rule $\tilde{S}$ for the prediction market. Unfortunately, she is unable to maximize the social welfare because she doesn't know the access function F and the information structure $\Pr[X \mid Y]$ (characterized by β in this case). We will see in this section that, if she chooses $\tilde{S} = S$ for prediction markets, the social welfare could be very poor. This is in stark contrast to the mechanisms we propose later in this paper, where the social welfare is maximized when $\tilde{S} = S$, without the need to know anything about F or $\Pr[X \mid Y]$. Our results are summarized in Figure 1.

We consider two examples of access functions and compute the social welfare of the market as the number of agents increases. When every agent's effort is c, social welfare is the total value minus

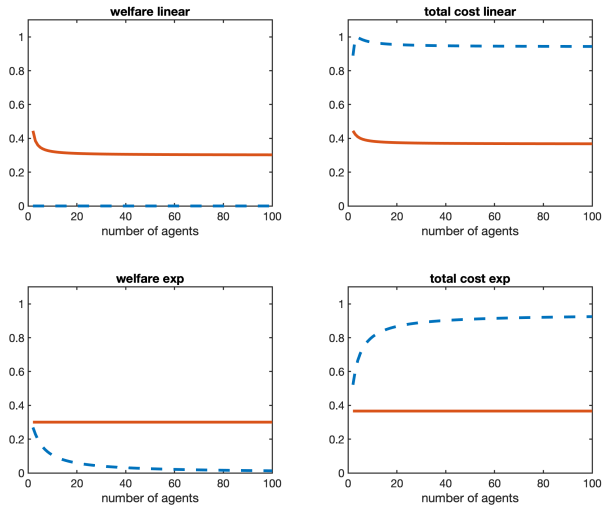

Figure 1: Red solid: Optimal (Centralized). Blue dashed: Prediction markets. Top: Linear access function $F(c) = 3c$ with $c \in [0, 1/3]$. Bottom: Exponential access function $F(c) = 1-e^{-3c}$ with $c \geq 0$. In both (linear and exponential) cases, the total costs (right column) are high in the strategic setting. In the linear case, agents spend all potential social value and have zero social welfare. In the exponential case, the total cost increases as the number of agents increases.

cost. The value is 1 when at least one agent derives his signal, and the total cost is cn, so in expectation,

$$W^{(n)}(c) = \left(1 - (1 - F(c))^n\right) - cn. \tag{1}$$

Proposition 3.1 (Linear Access). *Given $n \geq 2$, $\lambda > 1$ and a linear access function is $F(c) = \lambda c$ for $c \in [0, 1/\lambda]$, the optimal social welfare with n agents is $W^{(n)}(c_{\text{opt}}) = 1 - \lambda^{-n/(n-1)} - n\left(\lambda^{-1} - \lambda^{-n/(n-1)}\right) > 0$ in the centralized setting, but the social welfare $W^{(n)}(c_{\text{self}})$ is zero in the strategic setting.*

Proposition 3.2 (Exponential Access). *Given $n \geq 2$, $\lambda > 1$ and an exponential access function $F(c) = 1 - e^{-\lambda c}$ for $c \geq 0$. As $n \to \infty$, the optimal social welfare with n agents is $W^{(n)}(c_{\text{opt}}) = 1 - (1 + \ln \lambda)/\lambda > 0$, but the social welfare in the strategic setting is $W^{(n)}(c_{\text{self}}) = O(1/n)$.*

3.2 Delayed Report

Unfortunately, prediction markets do not guarantee timely reports, i.e., one may wait for others to report. This is undesirable when the value of information decays quickly. Suppose agents' signals are structured as in Example 2.1 with $\alpha = 0.02$, $\beta = 0.2$. Consider the quadratic scoring rule, where we let $G(p) = \|p\|_2^2$ and $S(p, y) = G(p) + \langle \nabla G(p), \delta_y - p \rangle = 2p(y) - \|p\|_2^2$. Contrary to the intuition, the marginal value of a report does not monotonically decrease. As shown in Figure 2, the largest increase in the scoring rule is due to the third report, not the first or the second. Thus, an agent who believes he is likely to be the first will wait before reporting.

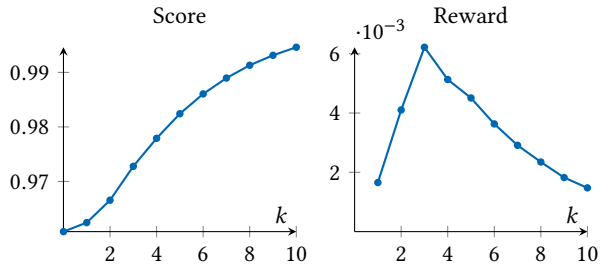

Figure 2: Left: The expected score after k reports. Right: The expected reward for the k-th report, which is the difference between two consecutive values in the left plot. Note that it's not decreasing!

4 SINGLE BATCH MODEL

We want to design new mechanisms to deal with the above issues. In particular, we assume the market belief is used by someone whose utility is the "quality" of the market belief minus the rewards she gives to the agents. Then, we maximize social welfare and encourage truthful and timely reports. We first ignore the time factor and consider a simple case, where each agent i:

(1) chooses to invest $c_i \in \mathbb{R}$ much effort,
(2) gets a signal X_i with probability $F(c_i)$, and reports b_i to the mechanism,
(3) receives some reward r_i from the mechanism.

We want to design a reward function (contract) from agents' reports to rewards in order to maximize the social welfare. If p is the aggregated belief from agents' reports, the value of information is represented as a (strictly) proper scoring rule $S(p, y)$, where y is the outcome. The principal's utility U, each agent i's utility u_i, and the social welfare W, are given by: $U = \mathbb{E}[S(p, Y)] - \sum_i \mathbb{E}[r_i]$, $u_i = \mathbb{E}[r_i] - c_i$, $W = \mathbb{E}[S(p, Y)] - \sum_i c_i$. Here the expectation is taken over all randomness (ex-ante), i.e., agents compute it based on the information in Stage 1. Note that given the information valuation S and the information structure, assuming agents are truthful, the ex-ante social welfare only depends on the agents' effort $\boldsymbol{c} = (c_1, \ldots, c_n)$. We call $\boldsymbol{c}^*$ an *optimal effort profile* if

$$\boldsymbol{c}^* \in \arg\max_{\boldsymbol{c}} \left(\mathop{\mathbb{E}}_{p,Y}[S(p, Y)] - \sum_i c_i \right), \tag{2}$$

where p is the Bayesian posterior of the outcome given all signals.

4.1 Mechanism and Theorem

To incentivize agents to invest the optimal effort, we want to design a mechanism that, given agents' reports, outputs an aggregated belief $p \in \Delta_{\mathcal{Y}}$ and a reward for each agent. This task is challenging for two reasons:

- The agents' efforts $c_1, \ldots, c_n$ and reports $b_1, \ldots, b_n$ are decided by each agent individually.
- The joint distribution $\Pr[X_k, Y]$ and the effort function $F(c)$ are not known to the principal, and thus naïvely eliciting agents' signals X_k does not work.

Our mechanism is shown in Mechanism 1. Each agent k is asked to report $b_{k,y} = \Pr[X_k \mid Y = y]$ for each $y \in \mathcal{Y}$ (one of

them can be omitted). If the mechanism knows $\frac{\Pr[Y=y \mid X_1,\ldots,X_{k-1}]}{\Pr[Y\neq y \mid X_1,\ldots,X_{k-1}]}$, then it is easy to updated it to $\frac{\Pr[Y=y \mid X_1,\ldots,X_k]}{\Pr[Y\neq y \mid X_1,\ldots,X_k]}$ given $b_{k,y}$ because $\frac{\Pr[Y=y \mid X_1,\ldots,X_k]}{\Pr[Y\neq y \mid X_1,\ldots,X_k]} = \frac{\Pr[Y=y \mid X_1,\ldots,X_{k-1}]}{\Pr[Y\neq y \mid X_1,\ldots,X_{k-1}]}\frac{b_{k,y}}{1-b_{k,y}}$, which follows by applying Bayes to both the numerator and denominator of the first two fractions, and then using the fact that X_k is conditionally independent of $X_1,\ldots,X_{k-1}$. As a result, we can update $\Pr[Y=y \mid X_1,\ldots,X_{k-1}]$ to $\Pr[Y=y \mid X_1,\ldots,X_k]$ because we can first compute $\frac{\Pr[Y=y \mid X_1,\ldots,X_{k-1}]}{\Pr[Y\neq y \mid X_1,\ldots,X_{k-1}]}$, use this and $b_{k,y}$ to compute $\frac{\Pr[Y=y \mid X_1,\ldots,X_k]}{\Pr[Y\neq y \mid X_1,\ldots,X_k]}$, and then transform this back to $\Pr[Y=y \mid X_1,\ldots,X_k]$. We succinctly denote the above process as $p_{k,y} = Update(p_{k-1,y}, b_{k,y})$ where

$$Update(p_{k,y}, b_{k,y}) = \frac{p_{k,y}b_{k,y}}{(1-p_{k,y})(1-b_{k,y}) + p_{k,y}b_{k,y}} \tag{3}$$

and $p_{k,y} = \Pr[Y=y \mid X_1,\ldots,X_k]$.

Mechanism 1: Fair Prediction Market

Input: a report profile $(b_1,\ldots,b_n)$ where $b_k = (b_{k,1},\ldots,b_{k,d-1})$ is the information provided by agent k, describing what the agent k's (claimed) values of $\Pr[X_k \mid Y=y]$ for each y. For those who do not obtain a signal, we assume their $b_{k,y} = 1/2$ for $y = 1,\ldots,d-1$, where $d = |\mathcal{Y}|$.

Output: the reward r_k for each agent k, and the aggregated belief p_n

for $k = 1$ **to** n **do**
  Let $(\pi_1,\ldots,\pi_n)$ be a random permutation with $\pi_n = k$
  $p_0 \leftarrow \Pr[Y]$
  for $j = 1$ **to** n **do**
    for $y = 1$ **to** $d-1$ **do** // This loop goes through every $y \in \mathcal{Y}$ and update the corresponding entry of p according to the information provided by agent π_j.
      $p_{j,y} \leftarrow Update(p_{j-1,y}, b_{\pi_j,y})$ as defined in (3)
    $p_{j,d} \leftarrow 1 - \sum_{y=1}^{d-1} p_{j,y}$
  $r_k \leftarrow S(p_n, y^*) - S(p_{n-1}, y^*)$ // y^* is the true outcome

A report is said to be *truthful* if it results in a Bayesian update on the market belief, where the prior is the previous market belief, the posterior is the new market belief, and the evidence is the reporter's signal. We pay each agent by his improvement on the market belief as if he were the last one to update. This mechanism is "fair" in the sense that if everyone is truthful then everyone receives the same expected reward.

Proposition 4.1. *Agent k makes a correct Bayesian update iff $b_{k,y} = \frac{\Pr[X_k \mid Y=y]}{\Pr[X_k \mid Y\neq y]}$ for all $y \in \mathcal{Y}$.*

Note that if $p_{k,y}$ is known for $|\mathcal{Y}|-1$ different y's, then the last one follows directly since their sum is 1, so each agent needs to report only $|\mathcal{Y}|-1$ values, the same number as in the original prediction market, where each agent reports a probability distribution.

Theorem 4.2. *Assume $\frac{\mathrm{d}^2F(c)}{\mathrm{d}c^2} < 0$.[5] Mechanism 1 is individually rational, and there exists a strict perfect Bayesian equilibrium $\boldsymbol{\sigma}$ in which the expected social welfare is maximized (over all symmetric strategy profiles), and $\boldsymbol{\sigma}$ satisfies the following properties:*

Effort Optimality *The effort profile $\mathbf{c}$ is optimal (as in (2)).*

Truthfulness *Each agent makes a Bayesian update on the market belief.*

4.2 Proof Sketch

In order to proof Theorem 4.2, we first show some lemmas.

Lemma 4.3 (Truthfulness). *Every report b_k will be truthful, assuming other reports are truthful. Any deviation will result in a strictly worse expected reward.*

Lemma 4.4 (Effort Optimality). *Assume $\frac{\mathrm{d}^2F(c)}{\mathrm{d}c^2} < 0$. Agents are incentivized to invest the "right" amount of effort (2) that maximizes the expected social welfare, assuming all reports are truthful.*

Lemma 4.5. *The expected score $\mathbb{E}[S(p,Y)]$ of distribution p only depends on the number of previous updates but not who have made updates, assuming all reports are truthful.*

Proof of Theorem 4.2. Each agent i makes two decisions, and his strategy can be written as $\sigma_i = (c_i, b_i)$, where c_i is the effort he invests at the beginning and b_i is his report. By Lemma 4.4, he will not deviate from c_i. By Lemma 4.3, he will not deviate from b_i.

The proof for individual rationality is simple: If in a Nash equilibrium, agent i gets negative expected utility, then he can deviate to $c_i = 0$ and get zero utility. This means he is not in a Nash equilibrium. Thus in every Nash equilibrium, every agent has non-negative expected utility. □

Below is the proofs of the lemmas. Lemma 4.4 is perhaps the most interesting among the three.

Proof of Lemma 4.3. Assuming other reports are truthful, for avery agent k, we have $p_{n-1} = \Pr[Y \mid \text{all signals except agent } k\text{'s}]$. By the property of strictly proper scoring rule, his best strategy is to make $p_n = \Pr[Y \mid \text{all signals}]$ for every possible p_{n-1}. Any deviation will lower his reward. This is achievable due to Proposition 2.3. □

Proof of Lemma 4.4. Since agents are symmetric to each other, we look for a symmetric equilibrium, where every agent i invests the same amount of effort $c_i = c$. Let $v_k = \mathbb{E}[S(p_k,Y) - S(p_0,Y)]$, the expected increase of the score after k updates.

The expected social welfare is given by:

$$W = \mathbb{E}[S(p_k,Y)] - nc = \mathbb{E}[v_k] + \mathbb{E}[S(p_0,Y)] - nc.$$

Setting the derivative to be zero, we have

$$\begin{aligned} 0 = \frac{\mathrm{d}W}{\mathrm{d}c} &= \frac{\mathrm{d}}{\mathrm{d}c}\mathbb{E}[v_k] - n \\ &= \frac{\mathrm{d}F(c)}{\mathrm{d}c}\frac{\mathrm{d}}{\mathrm{d}F(c)}\sum_{k=0}^{n}\binom{n}{k}F(c)^k(1-F(c))^{n-k}v_k - n \\ &= \frac{\mathrm{d}F(c)}{\mathrm{d}c}\sum_{k=0}^{n}\binom{n}{k}\Big(kF(c)^{k-1}(1-F(c))^{n-k} \end{aligned}$$

[5] decreasing marginal benefit, a very common assumption in economics.

$$- (n-k)F(c)^k(1-F(c))^{n-k-1}\Big)v_k - n$$

$$= \frac{\mathrm{d}F(c)}{\mathrm{d}c} n\Bigg(\underbrace{\sum_{k=1}^{n}\binom{n-1}{k-1}F(c)^{k-1}(1-F(c))^{n-k}v_k}_{\text{substitute } k \text{ by } k+1} - \sum_{k=0}^{n-1}\binom{n-1}{k}F(c)^k(1-F(c))^{n-k-1}v_k\Bigg) - n$$

$$= n\frac{\mathrm{d}F(c)}{\mathrm{d}c}\sum_{k=0}^{n-1}\binom{n-1}{k}F(c)^k(1-F(c))^{n-k-1}(v_{k+1}-v_k) - n. \quad (4)$$

On the other hand, the expected utility of agent i is

$$\begin{aligned} u_i &= \Pr[\text{get signal}]\,\mathbb{E}[\text{reward} \mid \text{agent } i \text{ reports}] - \text{cost} \\ &= F(c_i)\mathop{\mathbb{E}}_{k,X_1,\ldots,X_k,Y}[S(p_k,Y) - S(p_{k-1},Y) \mid \text{agent } i \text{ reports}] - c_i \\ &= F(c_i)\mathop{\mathbb{E}}_{k}[v_k - v_{k-1} \mid \text{agent } i \text{ reports}] - c_i \\ &= F(c_i)\sum_{k=0}^{n-1}\binom{n-1}{k}F(c)^k(1-F(c))^{n-k-1}(v_{k+1}-v_k) - c_i, \end{aligned}$$

where c is the amount of effort invested by other agents. Because $F(c_i)$ is concave in c_i, u_i is also concave in c_i. We also know that $F(c_i)$ is upper bounded (by 1), so $\lim_{c_i\to\infty}\mathrm{d}F(c_i)/\mathrm{d}c_i = 0$ and thus $\lim_{c_i\to\infty}\mathrm{d}u_i/\mathrm{d}c_i = -1$. Therefore, there is a unique c_i that maximizes the utility, which is either (a) $c_i = 0$ or (b) the point with zero derivative. In case (b), we have

$$0 = \frac{\mathrm{d}u_i}{\mathrm{d}c_i} = \frac{\mathrm{d}F(c_i)}{\mathrm{d}c_i}\sum_{k=0}^{n-1}\binom{n-1}{k}F(c)^k(1-F(c))^{n-k-1}(v_{k+1}-v_k) - 1. \quad (5)$$

Equation (5) describes how an agent's decision c_i should best response to those of others c. In a symmetric equilibrium, $c_i = c$. Then, surprisingly, Equation (5) becomes equivalent to Equation (4). In case (a), both individually optimal and socially optimal solutions are $c_i = 0$. In other words, the distributed maximization of each agent's utility can result in the maximization of social welfare. □

Proof of Lemma 4.5. Since signals are identically distributed and the effort function is the same for every agent, an agent is indistinguishable from another. The lemma simply follows. □

5 SEQUENTIAL MODEL

In this section, we consider a setting that involves time. Its difference from the setting of the previous section is that signals are not received by agents immediately but will be eventually. Formally, before receiving signal X_i, each agent i suffers from a latency T_i, which is a random variable with c.d.f. $F_c(t)$. For instance, $F_c(t) = 1 - e^{-\lambda ct}$ means $T_i \sim \mathrm{Exp}(\lambda c_i)$. Here, $F_c(t)$ — a generalization of the access function $F(c)$ used in the previous section — depends on time. In summary, each agent $i \in \mathcal{N}$:

(1) chooses to invest $c_i \in \mathbb{R}$ effort,
(2) obtains a signal X_i at time $T_i \geq 0$ generated from c.d.f $F_{c_i}(\cdot)$, decides a time $s_i \geq 0$ and a report b_i to send to the mechanism at time s_i,
(3) receives some reward r_i from the mechanism.

We assume that as long as the agents invest non-zero effort, they always obtain their signals before the true outcome being revealed because this is far in the future. The value of the information (market belief) evolves over time. Let $p(t) \in \Delta_{\mathcal{Y}}$ denote the market belief at time t. The value of a belief history $\{p(t)\}_{t>0}$ is defined as: $V = \int_{t>0} S(p(t), y)h(t)\mathrm{d}t$, where y is the outcome, S is a strictly proper scoring rule characterizing the quality of the market belief, and ***time value function*** h is a function characterizing how the value of information diminishes through time. For instance, $h(t) = \eta e^{-\eta t}$ means that the value of information decays exponentially. This would be appropriate if the principal needs to make decision at a random time $\tau \geq 0$ generated from an exponential distribution with parameter η.

We want to design a mechanism that takes agents' online reports as inputs, maintains a real-time market belief, and finally outputs the reward given to each agent. We also want this mechanism to be truthful, *timely*, and social-welfare-maximizing. We say a mechanism is timely if every agent reports immediately after he gets a signal. The principal's utility U, each agent i's utility u_i, and the social welfare W, are given by: $U = \mathbb{E}[V] - \sum_i \mathbb{E}[r_i]$, $u_i = \mathbb{E}[r_i] - c_i$, $W = \mathbb{E}[V] - \sum_i c_i$. Given proper scoring rule S and the information structure, assuming the agents aggregate their information in a *truthful* and *timely* manner, the expected social welfare only depends on the agents' effort $\boldsymbol{c} = (c_1, \ldots, c_n)$. In this section, we call $\boldsymbol{c}^*$ an *optimal effort profile* if

$$\boldsymbol{c}^* \in \arg\max_{\boldsymbol{c}}\Big(\mathop{\mathbb{E}}_{p,Y}[V] - \sum_i c_i\Big). \quad (6)$$

5.1 Mechanism and Theorem

Besides the challenges involved in Section 4 — including the hidden effort, the unknown information structure and effort function, and the potential manipulation of agents' reports — we also need to deal with another complexity: An agent can choose any time to report, not necessarily just at the time he receives his signal, and even before it (i.e., $s_i < t_i$). We restrict our focus to mechanisms where each agent can report only once. Note that this assumption exists in the previous literature, e.g., in the traditional prediction market, there would be no truthfulness guarantee without this assumption. For this sequential setting, we propose Mechanism 2, which updates a market belief using agents' reports one by one. Agents report their information in the same structure as in Section 4. The mechanism also computes counterfactual market beliefs with one of the reports skipped. In particular, the counterfactual belief for agent i's absence is what the market belief would be if agent i does not report. The reward for each agent depends on both the actual and counterfactual market beliefs.

Theorem 5.1. *Assume $h(t) > 0$ for all $t > 0$, $\int_{t>0} h(t)\mathrm{d}t < \infty$, $\frac{\mathrm{d}^2F_c(t)}{\mathrm{d}c^2} < 0$,[6] and $F_c(t)$ is the c.d.f. of a non-negative random variable for $c \geq 0$. Mechanism 2 is individually rational, and there exists a strict perfect Bayesian equilibrium that is socially optimal (over all symmetric strategy profiles) and satisfies the following properties:*

Effort Optimality *Every agent invests the "right" amount of effort as* (6).

[6] decreasing marginal benefit

Mechanism 2: Marginal Value Prediction (MVP) Market

Input: online report $b_j = (b_{j,1}, \ldots, b_{j,d-1})$ from each agent j
Output: real-time market belief p_k, and the reward r_i for each agent i

1 $p_0 \leftarrow \Pr[Y]$
2 **for** $i = 1$ **to** n **do**
3 $\quad \tilde{p}^i_0 \leftarrow \Pr[Y]$ `// $\tilde{p}^i$ is the counterfactual market belief for agent $i$'s absence`
4 **for** $k = 1$ **to** n **do**
5 $\quad$ wait until receive a report from an agent j and denote it by b_j
6 $\quad t_j \leftarrow$ current time
7 $\quad$ **for** $y = 1$ **to** $d-1$ **do** `// Update the market belief with agent $j$'s report`
8 $\quad\quad p_{k,y} \leftarrow Update(p_{k-1,y}, b_{j,y})$ as defined in (3)
9 $\quad p_{k,d} \leftarrow 1 - \sum_{y=1}^{d-1} p_{k,y}$
10 $\quad$ **for** $i = 1$ **to** n **do**
11 $\quad\quad$ **if** $i = j$ **then**
12 $\quad\quad\quad \tilde{p}^i_k \leftarrow \tilde{p}^i_{k-1}$
13 $\quad\quad$ **else** `// Update the counterfactual belief for $i$'s absence with $j$'s report`
14 $\quad\quad\quad$ **for** $y = 1$ **to** $d-1$ **do**
15 $\quad\quad\quad\quad \tilde{p}^i_{k,y} \leftarrow Update(\tilde{p}^i_{k-1,y}, b_{j,y})$ as defined in (3)
16 $\quad\quad\quad \tilde{p}^i_{k,d} \leftarrow 1 - \sum_{y=1}^{d-1} \tilde{p}^i_{k,y}$
`// After the true outcome $y^*$ reveals`
17 Let $k(t) = \sum_{j=1}^{n} \mathbb{I}[t_j < t]$, $p(t) = p_{k(t)}$, and $\tilde{p}^i(t) = \tilde{p}^i_{k(t)}$.
18 **for** $i = 1$ **to** n **do**
19 $\quad r_i \leftarrow \int_{t>0} (S(p(t), y^*) - S(\tilde{p}^i(t), y^*)) h(t) \mathrm{d}t$

Truthfulness *Each agent makes a Bayesian update on the market belief.*
Timeliness *For all $i \in \mathcal{N}$, $s_i = t_i$.*

5.2 Intuition and Proof Sketch

Our core idea is to pay each agent by the actual value of information minus the counterfactual value of information as if he had not *updated the market belief*. Let $\tilde{V}^i$ be the counterfactual value w.r.t. agent i. The reward (in Line 19 of Mechanism 2) is given by

$$V - \tilde{V}^i = \int_{t>0} (S(p(t), y) - S(\tilde{p}^i(t), y)) h(t) \mathrm{d}t,$$

where $\tilde{p}^i(t)$ is what the market belief would be at time t if agent i had not changed anything in the market. Figure 3 gives an intuition for $\mathbb{E}[V - \tilde{V}^i]$. Note that we are not talking about a counterfactual value for the case as if an agent had not *participated the game*. The number of agents is still n, and other agents do the same.

Lemma 5.2 (Effort Optimality). *Agents are incentivized to invest the "right" amount of effort that maximizes the expected social welfare, assuming all updates are timely and truthful.*

Lemma 5.3 (Truthfulness). *No matter what time an agent makes his update, a truthful update is better than a non-truthful one, assuming all other updates are truthful.*

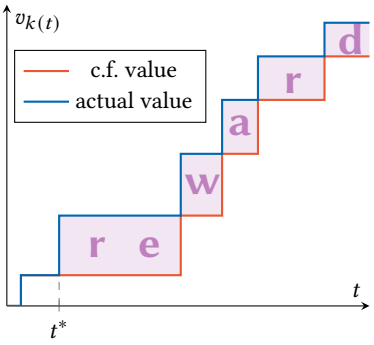

Figure 3: *Ex-ante* total reward received by an agent (taken expectation over *all* agents' signals) for a fixed time sequence of signal discovery. $v_k = \mathbb{E}[S(p_k, y) - S(p_0, y)]$ is the expected increase of score due to the first k reports, $k(t)$ is the number of report up to time t, and t^* is the time of report. A late report (larger t^*) reduces the reward. A non-truthful report shifts the *actual curve* downward, also reducing reward.

Lemma 5.4 (Timeliness). *Every agent is incentivized to update the market belief as soon as he gets his signal, assuming all updates are truthful.*

Proof of Theorem 5.1. Each agent i decides how much effort to make at the beginning. By Lemma 5.2, he will not deviate at this decision. Then, at each time, he decides whether to report and if so, what to report. By Lemma 5.4, he will report if he has a signal. By Lemma 5.3, he will report truthfully. □

The proofs of the lemmas are similar to those in Section 4 and are postponed to the supplementary materials.

5.3 Connection to VCG

There are both similarities and differences between our mechanisms and the VCG mechanism. They are similar because both of them have a payoff function that can be interpreted as an actual term minus a counterfactual term. Also, in our mechanisms, the utility function of an agent is — to some extent — aligned with the social welfare as a function of his action, as in VCG.

However, a straightforward application of VCG fails. In VCG, we need to compute the utility of each agent, which is impossible here, because the amount of effort each agent invests is never revealed. In our mechanisms, the alignment of the agent's utility and social welfare is achieved implicitly without the principal computing them. In addition, VCG deals with a single-stage game, while our mechanisms deal with multi-stage games. This is to say that the signals must be discovered before they can be (truthfully) revealed. Finally, VCG guarantees the DSIC (dominant-strategy incentive-compatible) property, which is not the case in our setting, where agents respond to others when choosing the effort level.

6 SIMULATIONS

The equilibrium depends on a variety of parameters. In this section, we analyze how it is affected by ease, noise, and substitutability of the information. We compare our socially optimal mechanism

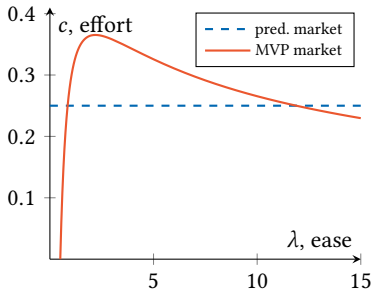

Figure 4: How ease affects agent's effort.

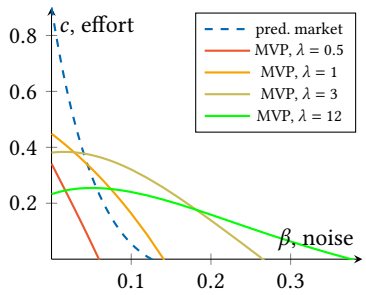

Figure 5: How noise affects agent's effort.

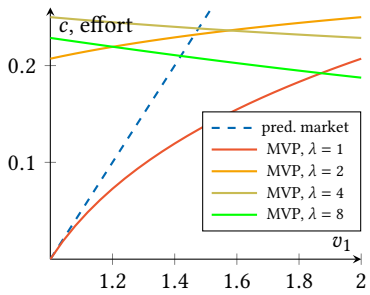

Figure 6: How substitutability v_1/v_2 (proportional to v_1 since v_2 is fixed) affects agent's effort.

(Mechanism 2) with the traditional prediction market.[7] We assume the value of information decays exponentially: $h(t) = \eta e^{-\eta t}$ with parameter $\eta = 1$. The latency of signal discovery is also exponentially distributed: $F_c(t) = 1 - e^{-\lambda c t}$. Here, λ can be viewed as the ease of collecting the information, as the larger λ is, the shorter latency the agent suffers. There are $n = 2$ agents unless otherwise stated. To simplify the calculation, we let $v_k = \mathbb{E}[S(p_k, y) - S(p_0, y)]$. It is an intermediate variable that depends on the information structure $\Pr[X \mid Y]$ (or β).

6.1 Ease

Let $v_1 = 2$ and $v_2 = 3$, defined above. In our mechanism, the amount of effort in equilibrium automatically adapts to the ease of collecting information, even though the mechanism does not know anything about the parameters! Figure 4 is a visualization. In prediction markets, agents invest too much for very easy ($\lambda \to \infty$) information, thus making the updates unnecessarily quick. For information that takes a long time to discover, agents still invest equally much effort in prediction markets, while in our mechanism, they do not invest anything because we have the $\eta e^{-\eta t}$ term in the value of information — the information value decays so quickly that the gain in information value is overwhelmed by the amount of effort invested.

6.2 Noise

Consider the scenario described in Section 3.2, with $\alpha = 0.1$. Recall that $\beta = \Pr[X \neq Y]$ is the probability each signal differs from the true outcome. It can also be regarded as the noise of information. It turns out that in prediction markets, agents invest too much not only for very easy information but also for very accurate ($\beta \to 0$) information, as shown in Figure 5. When the signals become weak enough, agents no longer invest anything in either prediction markets (because a late reporting is encouraged) or our mechanism (because the gain in value is too little).

6.3 Substitutability

Recall that v_k is the expected increase of the score due to the first k updates. We fix $v_2 = 2$, and see what happens when v_1 changes. Here, v_1/v_2 can be considered as the substitutability of information.

[7] Following Section 3.1, we study the case where $\tilde{S} = S$.

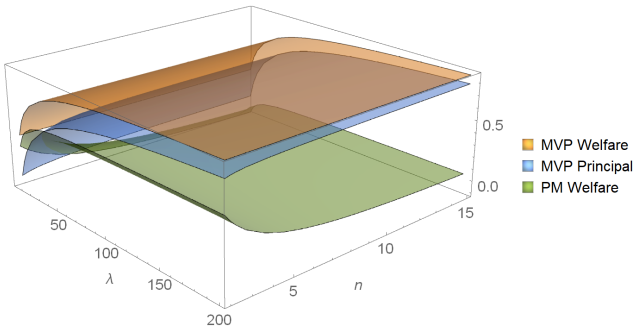

Figure 7: A comparison of the social welfare of the MVP market, the principal's utility of MVP market, and the social welfare of the traditional predictions markets. When there are many agents (n is large) and the signals do not take much time to obtain (λ is large), the MVP market has both high social welfare and high utility for the principal.

As shown in Figure 6, in prediction markets, when value is more concentrated in the first report (higher substitutability), agents invest more effort to get a signal quickly. Surprisingly, this is not always the case in our mechanism. When signals are very easy (quick to obtain), agents tend to invest *less* when value is more concentrated in the first signal. This is because even though the first report brings high value, after the second report, the marginal value of the first report (against other signals) becomes much less. Moreover, the easier to obtain a signal, the stronger this effect.

6.4 Social Welfare and Principal's Utility

As we argue that agents invest too much effort for easy and accurate information in the traditional prediction market, one may wonder how bad the effect of such behavior could be on social welfare. We find that the social welfare approaches to 0 as the number of agents grows large, while in our proposed mechanism, both social welfare and the principal's utility are high, as shown in Figure 7. In this example, we assume $v_0 = 0$ and $v_k = 1$ for each $k \geq 1$ (or equivalently, $\beta = 0$). This is similar to what we show in Section 3.1.

REFERENCES

[1] Jerry Adler. 2012. Raging Bulls: How Wall Street Got Addicted to Light-Speed Trading. *wired* 20 (August 2012). Issue 10. Accessed June, 26, 2020 at https://www.wired.com/2012/08/ff_wallstreet_trading/.

[2] Gagan Aggarwal, Gagan Goel, and Aranyak Mehta. 2009. Efficiency of (revenue-) optimal mechanisms. In *Proceedings of the 10th ACM conference on Electronic commerce*. 235–242.

[3] Pablo D Azar and Silvio Micali. 2018. Computational principal–agent problems. *Theoretical Economics* 13, 2 (2018), 553–578.

[4] Yossi Azar, Amir Ban, and Yishay Mansour. 2016. When Should an Expert Make a Prediction?. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. ACM, 125–142.

[5] Glenn W Brier and Roger A Allen. 1951. Verification of weather forecasts. In *Compendium of meteorology*. Springer, 841–848.

[6] Eric Budish, Peter Cramton, and John Shim. 2015. The high-frequency trading arms race: Frequent batch auctions as a market design response. *The Quarterly Journal of Economics* 130, 4 (2015), 1547–1621.

[7] Jeremy Bulow and Paul Klemperer. 1994. *Auctions vs. negotiations.* Technical Report. National Bureau of Economic Research.

[8] Mithun Chakraborty and Sanmay Das. 2015. Market scoring rules act as opinion pools for risk-averse agents. In *Advances in Neural Information Processing Systems*. 2359–2367.

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

[10] S Cox. 2015. Why tennis courtsiding was my dream job. *BBC News Mag* (2015).

[11] A Philip Dawid and Steffen L Lauritzen. 2005. The geometry of decision theory. In *Proceedings of the Second International Symposium on Information Geometry and its Applications*. University of Tokio, 22–28.

[12] Bruno De Finetti. 1937. La prévision: ses lois logiques, ses sources subjectives. In *Annales de l'institut Henri Poincaré*, Vol. 7. 1–68.

[13] Craig Dickson. 2015. 'Courtsiding' in Sport: Cheating, Sharp Practice or Merely Irritating? *LawInSport e-journal, March* (2015).

[14] F Douglas Foster and S Viswanathan. 1996. Strategic trading when agents forecast the forecasts of others. *The Journal of Finance* 51, 4 (1996), 1437–1478.

[15] Xi Alice Gao, Jie Zhang, and Yiling Chen. 2013. What you jointly know determines how you act: Strategic interactions in prediction markets. In *Proceedings of the fourteenth ACM conference on Electronic commerce*. 489–506.

[16] Tilmann Gneiting and Adrian E Raftery. 2007. Strictly proper scoring rules, prediction, and estimation. *J. Amer. Statist. Assoc.* 102, 477 (2007), 359–378.

[17] Irving John Good. 1992. Rational decisions. In *Breakthroughs in statistics*. Springer, 365–377.

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

[19] Craig W Holden and Avanidhar Subrahmanyam. 1992. Long-lived private information and imperfect competition. *The Journal of Finance* 47, 1 (1992), 247–270.

[20] Yuqing Kong and Grant Schoenebeck. 2018. Optimizing Bayesian information revelation strategy in prediction markets: the Alice Bob Alice case. In *9th Innovations in Theoretical Computer Science Conference (ITCS 2018)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik.

[21] Albert S Kyle. 1985. Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society* (1985), 1315–1335.

[22] Noam Nisan. 2009. *Revenue vs. Efficiency in Auctions.* https://agtb.wordpress.com/2009/06/09/revenue-vs-efficiency-in-auctions/

[23] Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V Vazirani. 2007. *Algorithmic game theory*. Cambridge university press.

[24] Ben Rothenberg. 2017. Man Barred From Open Last Year Is Arrested After Returning. *New Yolk Times On-line* (2017). accessed November 15, 2019.

[25] Business Round Table. 2020. *Our Commitment: Statement on the Purpose of a Corporation.* https://opportunity.businessroundtable.org/ourcommitment/