

Timing is Money: The Impact of Arrival Order in Beta-Bernoulli Prediction Markets

Blake Martin
Carnegie Mellon University
Pittsburgh, Pennsylvania, USA
blakemar@andrew.cmu.edu

Mithun Chakraborty, Sindhu Kutty
University of Michigan
Ann Arbor, Michigan, USA
{dcsmc,skutty}@umich.edu

ABSTRACT

Prediction markets are incentive-based mechanisms for eliciting and combining the diffused, private beliefs of traders about a future uncertain event such as a political election. Typically prediction markets maintain point estimates of forecast variables; however, *exponential family prediction markets* define a class of cost function-based market-making algorithms that maintain a complete, collective *belief distribution* over the underlying generative process of the event of interest (e.g. the probability density of the winner’s vote-share). We focus on concretizing a special case of this abstract framework, the algorithmic market maker being based on the beta distribution. We set up a multi-agent simulation of the market ecosystem to experimentally investigate the interaction of this microstructure with a *heterogeneous trading population*. We design a Bayesian trader model with explicit characterization of this heterogeneity with respect to two independent attributes: how rich a trader’s private information is and how much wealth they initially have at their disposal. We gauge the interplay of the above attributes with the *arrival order* of traders, particularly in terms of the net profit accrued by different trader types. Our results strongly suggest that early arrival can dominate both wealth and informativeness as a factor in determining trader compensation under a variety of experimental conditions.

CCS CONCEPTS

• Applied computing → Electronic commerce; • Computing methodologies → Model development and analysis.

KEYWORDS

Prediction markets, Exponential family distributions, Bayesian beliefs, Informativeness, Budget

ACM Reference Format:

Blake Martin and Mithun Chakraborty, Sindhu Kutty. 2021. Timing is Money: The Impact of Arrival Order in Beta-Bernoulli Prediction Markets. In *2nd ACM International Conference on AI in Finance (ICAIF’21)*, November 3–5, 2021, Virtual Event, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3490354.3494406>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
ICAIF’21, November 3–5, 2021, Virtual Event, USA

© 2021 Association for Computing Machinery.
ACM ISBN 978-1-4503-9148-1/21/11...\$15.00
<https://doi.org/10.1145/3490354.3494406>

1 INTRODUCTION

High-quality forecasts of uncertain events have always been necessary for decision-making and planning purposes in many spheres of human life. One way to produce such forecasts is to solicit and combine personal beliefs (or signals) about the event of interest from a potentially diverse population by polls, surveys, etc.; but this approach can be messy and prone to misinterpretation and manipulation. On the other hand, financial markets, where commodity prices emerge from the collective activity of a diverse trading population, have long been known to double as aggregators of similarly diffused information — their advantage lies in their built-in financial incentives and timely responses. *Prediction markets* [3] such as PredictIt (<https://www.predictit.org/>) and Iowa Electronic Markets (<https://iemweb.biz.uiowa.edu/>) represent a particular type of financial market that are designed with the express purpose of eliciting the personal beliefs of traders about an uncertain future event and reflecting these beliefs in the observable market “state”. Such a market offers trade in a carefully designed bundle of contracts, or *securities*, whose monetary value at market termination is tied to the realization of the forecast event. For instance, a market designed to predict the outcome of the U.S. presidential election may issue a single security that is finally worth \$1 if a Democrat wins the race and \$0 otherwise; the price of this security over the lifetime of the market can be interpreted as an aggregated probability that a Democrat will win the race. There is considerable empirical evidence of such markets being at least as effective as alternatives such as polls [4, 10, 17, 30]. There has also been a long line of theoretical research on prediction markets (see, e.g., [8, 16, 25]), especially after Robin Hanson introduced *market scoring rules* (MSR), a family of algorithms that act as automated market makers for prediction markets [18, 19]. An MSR under mild assumptions can be implemented in terms of a convenient and interpretable convex *cost function* which specifies the cost of buying/proceeds from selling shares in the market securities, given the total number of shares sold thus far [1, 7, 8]. This single function fully determines the *market microstructure*, hence plays a key role in traders’ decisions and the market’s aggregation characteristics (in particular, the *price function*, i.e. the first-order derivative of the cost function, represents the market’s probabilistic belief over the outcome space).

We seek to contribute to ongoing research on both the design of prediction markets for particular forecast objectives as well as the evaluation of their performance in the presence of various trader models. Let us explain using the prediction of a random binary event (e.g. a two-candidate presidential election as above) as a running example. The uncertainty in the realized outcome (winner) can then be modeled as a Bernoulli distribution with an unknown single parameter $p \in [0, 1]$ (the probability of a Democrat winning).

To maintain a point estimate of this p , one can use the classic *Logarithmic Market Scoring Rule* (LMSR), a member of the MSR family and the de facto standard algorithm for real-world prediction markets. An LMSR-based prediction market is incentive-compatible for myopic, risk-neutral traders with no budget constraints [19]; i.e., the trade that maximizes such a trader’s expected profit from the market also drives the market price to what this trader believes p to be, based on her private information. A plethora of work [21, 24, 25, 27] has established the impressive *convergence* or *aggregation* properties of prediction markets (MSRs in particular), i.e. how well the market state (e.g. price) tracks the belief that would be obtained by combining the private information of all participating traders under various assumptions on trader behavior.

But uncertainty about the forecast variable(s) can exist on multiple levels. In our running example, not only is the election winner uncertain until market closing but so is the underlying process represented by the Bernoulli parameter p . Is it feasible to design a market that elicits and expresses uncertainty at such “deeper” levels while retaining (some of the) desirable properties of MSRs? Abernethy et al. [1] laid the theoretical foundation for a family of market designs that can capture uncertainty when it is expressed in terms of a probability distribution belonging to the well-known *exponential family* [29]. We can model our unknown Bernoulli parameter as a continuous random variable with support on $[0, 1]$ following a beta distribution which belongs to the exponential family and is fully described by two independent parameters (see Section 2.1 for details); following Abernethy et al. [1], we can hence implement a specific exponential family market where the securities correspond to the above defining parameters of a beta distribution. Each trader effectively “reports” a parametric encoding of her *belief distribution* over p rather than a point probability p , and is compensated based on her report as well as the revealed “true” value of p (e.g. the vote-share of the winner could serve as a proxy for p). What do the financial (e.g. trader compensation) and informational (e.g. belief aggregation) characteristics of the ecosystem induced by this more complicated market design look like? These questions call for a careful consideration of the attributes and circumstances of traders, especially those that were not taken into account for designing the market.

In this work, we model traders as *Bayesian entities* that each privately observe a small number of Bernoulli trials with the same parameter p that determines the “true” outcome; they then compute their posterior belief distribution based on their private observations and the current market state, and trade so as to maximize their expected net compensation with respect to this belief. The *size* of this private sample can be viewed as a measure of the quality of the trader’s information source or, from the perspective of the market, the reliability of her report. We call this property the *informativeness* of a trader. Intuitively, it would be unfair if a trader of lower informativeness sustained a higher profit in expectation from the market than one of higher informativeness. Moreover, some or all traders may have a *budget* constraint that limits their trading ability, and hence their *effective* informativeness as perceived by the market. These trader attributes may interact in complex and non-intuitive ways with the microstructure: in cost function-based market making, the traders arrive sequentially and directly interact with the algorithmic market maker only (a pure-dealer model),

their profit being decided by the *marginal* value — rather than the *absolute* value — of the information they inject into the market, *relative* to the current market state induced by all previous trades. If a highly informative trader shows up after a large amount of information has already been incorporated into the market, can they end up booking a smaller profit than a less informative trader arriving earlier? How serious is this *sequence effect*? This is particularly important when traders cannot fully control their arrival time owing to market regulations or circumstances extraneous to the market. For instance, the market may employ tie-breaking to resolve simultaneous arrival or a trader may be willing or able to trade only after her source delivers the private observations on which she bases her belief and hence her trading decision.

1.1 Our Contributions

Our broad motivation is to develop a detailed understanding of *price formation* and *trader compensation* for an *exponential family prediction market* with a *heterogeneous* trader population.

- We design the *beta-Bernoulli prediction market* ecosystem where a central automated market maker maintains and updates a beta belief distribution about the parameter of a Bernoulli event (Section 3.1); we adopt a simulation-based approach to systematically investigate the properties of the market maker in the presence of trading agents described below. To our knowledge, this is the first experimental assessment of any exponential family market.
- We formulate an agent-based model of trading under Bayesian belief updates (Section 3.2): salient features of an agent include her *informativeness*, measured by the number of privately observed draws from the true Bernoulli distribution, and her *budget*, i.e. an upper bound on the *debt* (negative wealth) she can be in when the market closes.
- We conduct and report two sets of experiments: in the first, traders with unlimited budget are divided into two types based on their asymmetric informativeness; in the other, traders of the same informativeness are divided into two types depending on whether or not they are budget-limited. In each set, we construct different trader sequences based on types (e.g. randomized, interleaved, etc.). We measure separately the interaction of these two features with a trader’s arrival time in influencing market convergence and trader compensation. In the second set, we simulate several *rounds* of forecasting, each having its own market and ending in the revelation of a different Bernoulli event but involving the same agents; we track the evolution of traders’ budgets over these rounds.

A major insight from our experimental results is that the arrival time of a trader can outweigh both informativeness and budget in terms of effect on trader compensation (or trading power). This is not a counter-intuitive finding but the value of our study lies in the in systematically *measuring this effect* and identifying conditions under which it is prominent, laying the groundwork for further experimental (and theoretical) analysis of different exponential family prediction markets.

1.2 Further Related Work

Our work is based on the exponential family market framework due to Abernethy et al. [1] (see [15] for follow-up work on elicitation for exponential family distributions in a non-market context).

This paper does provide some analytical results on convergence for some trader models but traders are assumed to be homogeneous and with infinite budgets. Unlike papers on convergence / aggregation in market equilibrium that abstract away from microstructure ([5, 26, 28] and references therein), we study convergence under a cost function-based microstructure; moreover, our market is designed to capture uncertainty at a finer granularity compared to predecessors in this vein [21, 24, 27]. There has been both theoretical [11] and empirical [27] work on budget restrictions for traders in cost function-based trading but we operationalize budgets differently (Section 3.2.2). In our second set of experiments, we consider a sequence of forecast rounds and study the growth of budget over these rounds, reminiscent of Beygelzimer et al. [5]; but we admit a microstructure and distinct characterizations of budget and informativeness. Our work also differs from the strand on *information structure* of potentially manipulative traders, i.e. correlation among traders' private signals, and its impact on market equilibria ([2, 9, 22] and references therein) — we focus on a quantification of a myopic trader's individual informativeness. Recent work on prediction markets over interval securities [14] allows traders to express their confidence by trading on an interval rather than a point estimate; however, our market is set up to elicit the entire (parametrized) posterior distribution, given that it belongs to a well-defined family. Some other recent papers on prediction markets that deserve mention are: Dudík et al. [13] who look at tradeoffs among forecast error components in prediction markets rather than the interplay of trader attributes; Laskey et al. [23] who deal with *combinatorial* prediction markets that elicit fine-grained information of a fundamentally different nature.

2 PRELIMINARIES

2.1 The Exponential Family of Distributions

The exponential family [29] is a set of distributions whose probability density function over a d -dimensional real-valued random variable $x \in \mathbb{R}^d$ can be expressed in the form

$$p(x|\eta) = \exp\{\eta \cdot \phi(x) - C(\eta)\}. \quad (1)$$

Here $\eta \in \mathbb{R}^d$ is the vector of *natural parameters* of the distribution, $\phi(x) \in \mathbb{R}^d$ is that of *sufficient statistics*, and $C(\eta)$ is called the *log-partition function*. Exponential family distributions are the solution to the constrained optimization problem of finding the maximum entropy distribution with a given value of expected sufficient statistics. Many commonly used distributions like the Gaussian, Poisson, and importantly for this paper, beta and Bernoulli distributions, are all members of this family.

The Bernoulli distribution is defined over a binary outcome space $X \in \{0, 1\}$ and is fully described by a single parameter $p = \Pr[X = 1] = \mathbb{E}[X] \in [0, 1]$. In situations where this parameter itself is unknown and we wish to maintain and update our uncertainty about its value, we can model this uncertainty as a distribution over the parameter space $[0, 1]$. A natural choice in this case is the *beta distribution* which is the *conjugate prior* of the Bernoulli data distribution [12] and is itself a member of the exponential family. The probability density function of the beta distribution is specified

in terms of two independent parameters $\alpha, \beta \in \mathbb{R}^+$:

$$\text{BETA}(p; \alpha, \beta) = \frac{p^{\alpha-1}(1-p)^{\beta-1}}{B(\alpha, \beta)}, \quad p \in [0, 1],$$

where $B(\alpha, \beta) = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}$; $\Gamma(\cdot)$ denotes the well-known *gamma function* [20]. The expectation of the distribution is given by $\mathbb{E}[p] = \frac{\alpha}{\alpha+\beta}$. For $\alpha = \beta = 1$, the distribution is identical to the uniform distribution over the interval $[0, 1]$. Suppose a Bernoulli random variable is parameterized by $p \in [0, 1]$. Given a prior belief distribution $\text{BETA}(p; \alpha, \beta)$ over p and m observations $\{x_1, \dots, x_m\}$ drawn from the true Bernoulli distribution, the Bayesian posterior is also a beta distribution with updated parameter values

$$\alpha^{(\text{posterior})} = \alpha + \sum_{i=1}^m x_i; \quad \beta^{(\text{posterior})} = \beta + m - \sum_{i=1}^m x_i. \quad (2)$$

2.2 Exponential Family Markets

Exponential family markets [1] use the structure of an exponential family distribution to construct a cost function-based prediction market with interesting theoretical elicitation and aggregation properties. Classical LMSR is a special case of this market where the cost function is based on the log-partition function of the Bernoulli probability mass function.

Let $x \in \mathbb{R}^d$ be a random variable of interest. A market maker that intends to elicit and aggregate beliefs of the expectation of a vector of statistics $\phi(x)$, chooses an appropriate maximum-entropy distribution based on these statistics — let this be the exponential family distribution $p(x|\eta)$ as defined in Equation (1) with sufficient statistics $\phi(x)$. The market maker defines the microstructure based on this distribution, including: *securities* (the commodities traded in the market), a *cost function* (that determines the cost of purchasing these securities) and a *payoff function* (i.e. cash per share that each security liquidates to at market termination). Corresponding to each dimension of the sufficient statistics vector, the market maker issues a security or contract that pays off the value of the corresponding statistic $\phi(\cdot)$ for the realized outcome x^* . In other words, if a trader holds $\delta \in \mathbb{R}^d$ shares, where each dimension δ^i corresponds to the trader's holdings of the i^{th} security, then the trader receives a payment of $\delta \cdot \phi(x^*)$ from the market maker from these holdings on market closing. After every trade, the vector of all traders' share holdings of all securities, called the *outstanding shares*, determines the market "state" — this corresponds to the natural parameter η for exponential family markets [1]. The market maker prices securities based on the log-partition function which serves as the cost function in this market — it charges the trader $C(\eta + \delta) - C(\eta)$ to purchase δ shares at market state η . Hence, the traders' net profit/compensation from these holdings on market closing is $\delta \cdot \phi(x) - (C(\eta + \delta) - C(\eta))$. The instantaneous *price vector* (the cost of purchasing an infinitesimal amount of shares) at market state η is given by $dC/d\eta$ which is also equal to the expectation of the sufficient statistics vector (equivalently, the payoff vector) with the respect to the distribution induced by η , $\mathbb{E}_\eta[\phi(x)]$.

3 MARKET AND TRADER MODELS

A traditional LMSR market maker's goal is to update its belief over the (usually discrete) outcome space of some uncertain real-world

event; e.g. a Bernoulli random variable $X \in \{0, 1\}$ where 1 (resp. 0) stands for the win of the Democratic (resp. Republican) candidate in a U.S. presidential race (before it is called). The market price at any time corresponds to the expected value of this random variable (the probability p of a Democrat winning in the above example).

In our model, the market institution has a different goal: instead of learning the expected value (a *point estimate*) of X , we are interested in learning a belief *distribution* over the parameters of the probability distribution of X . Thus, our random variable of interest has support $[0, 1]$; we model our uncertainty over this variable via a *beta distribution* which is the conjugate prior for the Bernoulli data distribution [12]. Since the beta distribution is fully described by parameters $\alpha, \beta > 0$, an update to the belief distribution is equivalent to an update in the two-dimensional (α, β) parameter space.

Another feature of our model is that, at market closing, the “true” Bernoulli success probability p_{true} is made public. While the revelation of the outcome is not the same as the revelation of the specifics of the outcome-generating process, there are scenarios where we can observe a proxy for the process parameters; the vote-share of the winning candidate in a binary election can serve as such a proxy (see, e.g. Chakraborty et al. [6] for a vote-share prediction market).

3.1 Beta Market Maker

We now specify the design of the beta distribution-based market-making algorithm for the above prediction task in accordance with the principles laid down in [1]. The goal of the algorithmic market maker is to learn an aggregate belief distribution over $p_{true} \in [0, 1]$:

- **Securities and payoffs.** The market offers trade in two securities whose payoffs are the sufficient statistics computed at the value $p_{true} \in [0, 1]$ revealed¹ at market termination, i.e. $\ln p_{true}$ and $\ln(1 - p_{true})$.
- **Outstanding shares and market belief.** The market maintains and continuously displays (for the trading population) the current outstanding shares $\eta = [\eta_1, \eta_2]^T$. This corresponds to the natural parameters of (and hence fully describes) the market’s current belief $\text{BETA}(p; \alpha, \beta)$ as follows: $\eta_1 = \alpha - 1$, $\eta_2 = \beta - 1$.
- **Cost function.** This corresponds to the beta log-partition function [29] and is also known to the traders:

$$C(\eta) = \ln \Gamma(\eta_1 + 1) + \ln \Gamma(\eta_2 + 1) - \ln \Gamma(\eta_1 + \eta_2 + 2).$$

In particular, the outstanding share vector readily gives us the market’s current instantaneous prices of the two securities (or their expected payoffs) as: $\mathbb{E}_\eta[\ln p] = \psi(\alpha) - \psi(\alpha + \beta)$ and $\mathbb{E}_\eta[\ln(1 - p)] = \psi(\beta) - \psi(\alpha + \beta)$, where $\psi(y) = \frac{d\Gamma(y)/dy}{\Gamma(y)}$ is the *digamma function*. We may also infer the expected value of the unknown Bernoulli parameter as $\mathbb{E}_\eta[p] = \frac{\eta_1 + 1}{\eta_1 + \eta_2 + 2}$.

3.2 Bayesian Trader Model

In our model, traders are *myopic*, i.e. on each market entry, they act as if it is their last chance to interact with the market before market closing, and *risk-neutral*, i.e. they trade to optimize their expectation of their net compensation from the market (perhaps subject to budget constraints, see Section 3.2.2). Before trading, the

trader formulates her own belief about p as follows: before each of her arrivals, a trader privately observes a sample of m datapoints drawn independently from the true Bernoulli distribution: $x_i \sim_{i.i.d.} \text{BERNOULLI}(p_{true}) \quad \forall i \in \{1, \dots, m\}$. The trader is *Bayesian*: she uses the market state η to define her Bayesian prior belief $\text{BETA}(\eta_1 + 1, \eta_2 + 1)$; using her sample, she constructs her Bayesian posterior belief which is a beta distribution with updated parameters as in Equations (2). Due to the incentive-compatibility of exponential family markets [1], the trader’s optimal trade updates the market’s state to η^{new} as follows: $\eta_1^{new} = \eta_1 + \sum_{i=1}^m x_i$ and $\eta_2^{new} = \eta_2 + m - \sum_{i=1}^m x_i$. One immediate observation is that η_1, η_2 increase (in general, weakly) with each trade under this model; in other words, traders express their beliefs through purchases only, eliminating the need for (short-)selling.

3.2.1 Trader Informativeness. We define the informativeness of a trader simply as her private sample size m as above. Since these datapoints are drawn from the true distribution, larger the sample size, the higher the quality of the posterior. Interestingly, since the market displays η at all times and it is updated in the above Bayesian manner, the informativeness of the latest trader can be inferred readily as $m = \eta_1^{new} + \eta_2^{new} - (\eta_1 + \eta_2)$.

3.2.2 Trader Budget. Until now, we have implicitly assumed that every trader has an unlimited budget for trading and hence is free to inject all her information into the market state. But, a trader may be constrained by a finite budget $B > 0$ to limit her worst-case exposure in the market. More formally, a budget-limited trader with privately observed sequence of Bernoulli datapoints x_1, x_2, \dots, x_m will trade as follows. For a current market state of η , she chooses to trade on the posterior induced by the leading sub-sequence x_1, \dots, x_k , for the largest value of $k \in \{1, 2, \dots, m\}$ such that

$$\max_p \left\{ C \left(\eta + \left[\sum_{i=1}^k x_i, k - \sum_{i=1}^k x_i \right]^T \right) - C(\eta) - \left[\sum_{i=1}^k x_i, k - \sum_{i=1}^k x_i \right] \cdot [\ln p, \ln(1 - p)] \right\} \leq B.$$

In our simulations, we maximize over $p \in [\epsilon, 1 - \epsilon]$ for some small $\epsilon > 0$ to exclude the degenerate instances corresponding to sure events. If the trader returns, she retains the last $m - k$ datapoints $x_{k+1}, x_{k+2}, \dots, x_m$ that she failed to incorporate and collects the remaining k datapoints for the new trade afresh.

4 EXPERIMENTS

4.1 Experimental Setup

We conduct two sets of experiments, SET1 and SET2, on a beta-Bernoulli market ecosystem (Section 3). The starting market state is set at $\eta = [0, 0]^T$ (i.e. no outstanding shares) or, equivalently, $\alpha = \beta = 1$ (i.e. a uniform prior distribution over the unknown Bernoulli parameter).

4.1.1 SET1: Informativeness and Trader Ordering. In SET1, we have two types of (unlimited-budget) traders with different levels of informativeness. We aim to compare the *average* difference in compensation between trader types across multiple instances of the ‘same’ market i.e., with the same initial market state, revealed value

¹For our simulations, we ensure $p_{true} \in (0, 1)$ to prevent degenerate cases.

of p_{true} and trader ordering (described below). We report the average compensations for the two trader types over different revealed values of p_{true} and different orderings.

HI (high informativeness) traders have a private sample size $m = 5$, and LI (low informativeness) traders have $m = 1$. In each market instance, there are $n = 500$ traders of each type, i.e. $2n = 1000$ traders in total. The revelation at market closure, i.e. after all $2n$ traders have completed trading, has three possibilities $p_{true} \in \{0.25, 0.5, 0.75\}$, determining each trader's compensation for that instance. We consider four different *orderings* of the trader population with respect to types:

HI First: All n HI traders arrive first, then all n LI traders.

LI First: All n LI traders arrive first, then all n HI traders.

Interleaved: LI and HI traders arrive alternately, starting with LI.

Random: A random permutation of all $2n$ traders is chosen at the start of the market and they trade in sequence.

We measure the per-trader average net compensation of HI traders and the same quantity for LI traders over 10^4 market instances for a fixed value of p_{true} and fixed trader ordering (the instances differing in each trader's observed random datapoints).

4.1.2 SET2: Budgets and Trader Ordering. In SET2, we have two types of traders with the same level of informativeness: those that are budget-limited (LB) and those with unlimited budget (UB). We track the growth in budget of LB traders over multiple *market rounds* defined as follows. In each round, a new market starts at $\eta = [0, 0]^T$; each trader is *multi-shot* and can re-enter the market multiple times in a single round until the market in that round terminates according to a *stop criterion* (see below). At the start of each round, the Bernoulli parameter p_{true} is drawn from the uniform distribution $U[0.05, 0.95]$; it governs all traders' private samples and is revealed when the round terminates, determining all payoffs. We track the overall compensation characteristics of a trader population over a sequence of $N_{max} = 25$ rounds. We repeat these experiments for traders with different starting levels of limited budget. As in Beygelzimer et al. [5], the fixed population of traders can be viewed as a panel of experts responding to a series of different binary forecast questions and having their *weights* (akin to trading power/budgets) readjusted by the response aggregator according to their performance history.

In each simulation over N_{max} rounds, we have $n = 5$ traders of either type in the population. On each entry (opportunity to trade) in any market round, every trader has a private sample of the same size $m = 5$ at her disposal. Each UB trader uses all her m (fresh) datapoints and the current market state to determine her posterior belief, and hence her trade, every time she enters. Thus, a UB trader acts as a multi-shot version of an HI trader from SET1. Each LB trader is endowed with an initial budget $B_0 \in (0, \infty)$ at the start of round $N = 1$. In our experiments, $B_0 \in \{2, 5, 10\}$. LB traders possess potentially different budgets B_{N-1} at the start of each round N of the market. At her i^{th} entry in round N , an LB trader makes her budget-constrained trading choice as described in Section 3.2.2 with her "residual" budget $B = B_{N-1}^i$; $B_{N-1}^1 = B_{N-1}$ and B_{N-1}^i for every $i > 1$ is obtained from B_{N-1} by properly accounting for the cost function-based payments and share holdings (if any) due to all of the trader's first $i - 1$ entries in this round. At the end of round N , an LB trader's net compensation from this round is added to

B_{N-1} to produce her budget B_N for the start of the next round. We define trader orderings as follows:

UB First: All n UB traders come first, then all n LB traders.

LB First: All n LB traders come first, then all n UB traders.

Interleaved: UB and LB traders alternate, starting with LB.

Random: A random permutation of all $2n$ traders is chosen at the start of every round.

For the first three orderings, we have an arbitrary internal order within the LB (similarly, UB) sub-population at the start that is maintained over rounds in each simulation; for all orderings, re-entries are handled in a round-robin fashion over individual traders. Each simulation is characterized by a sequence of N_{max} realized values of p_{true} sampled independently over rounds as well as random draws for private samples of each trader. Results are averaged over 10^3 simulations for each trader ordering and initial budget.

Stop criterion: We terminate a market round either when *each* trader reaches an upper bound k_{max} on the number of opportunities to trade or when the prices converge (in the following sense), whichever comes sooner. We split each round into successive blocks of 10 trader entries each; if η^{old} and η^{new} denote respectively the outstanding shares vector before and after such a block, we track the Euclidean distance between the corresponding market prices $\left\| \mathbb{E}_{\eta^{new}} [[\ln p, \ln(1-p)]^T] - \mathbb{E}_{\eta^{old}} [[\ln p, \ln(1-p)]^T] \right\|_2$; if this difference is at most a small $\varepsilon > 0$, we deem the market prices to have converged. In our experiments, we choose $k_{max} = 60$ and $\varepsilon = 0.01$.

4.2 Results

4.2.1 SET1. We will first look at the evolution and convergence of the market and then at measured trader compensation characteristics. We omit results for $p_{true} = 0.75$ as they mirror those for $p_{true} = 0.25$ due to symmetry. Since our market maintains a complete beta probability density over the random variable p after every trade (instead of a point estimate), we can assess its aggregation quality by how well it concentrates probability around p_{true} . Figure 1 depicts the market's (averaged) belief after 100 trades, 500 trades (mid-point of a market's lifetime) and 1000 trades (market termination) for each of our trader orderings. We see the (averaged) posterior belief *converging in expectation* to p_{true} as early as 100 trades regardless of trader ordering! For instance, even for **LI First** which has the slowest rate of information injection by design, the average (α, β) parameters after 100 trades for $p_{true} = 0.25$ are (26.01, 75.99) with respective standard errors (0.04, 0.04), and the (symmetrical bell-shaped) beta distribution corresponding to the above average parameters has an expectation of 0.2550. Over subsequent trades, the belief becomes progressively more peaked at p_{true} . After 500 trades, the (averaged) beta distribution for **LI First** concentrates over 99% of its probability in the interval 0.2509 ± 0.05 ; the corresponding intervals are slightly tighter for the other orderings. The speeds of both convergence in expectation and concentration around the mean are understandably the highest for **HI First**, the lowest for **LI First**, and intermediate and comparable for **Interleaved** and **Random**. After 1000 trades, the distributions for the same p_{true} are virtually indistinguishable for all orderings, as expected from our design. For $p_{true} = 0.5$, the expectation of the market's starting belief $U[0, 1]$ is 0.5 which already equals p_{true} ;

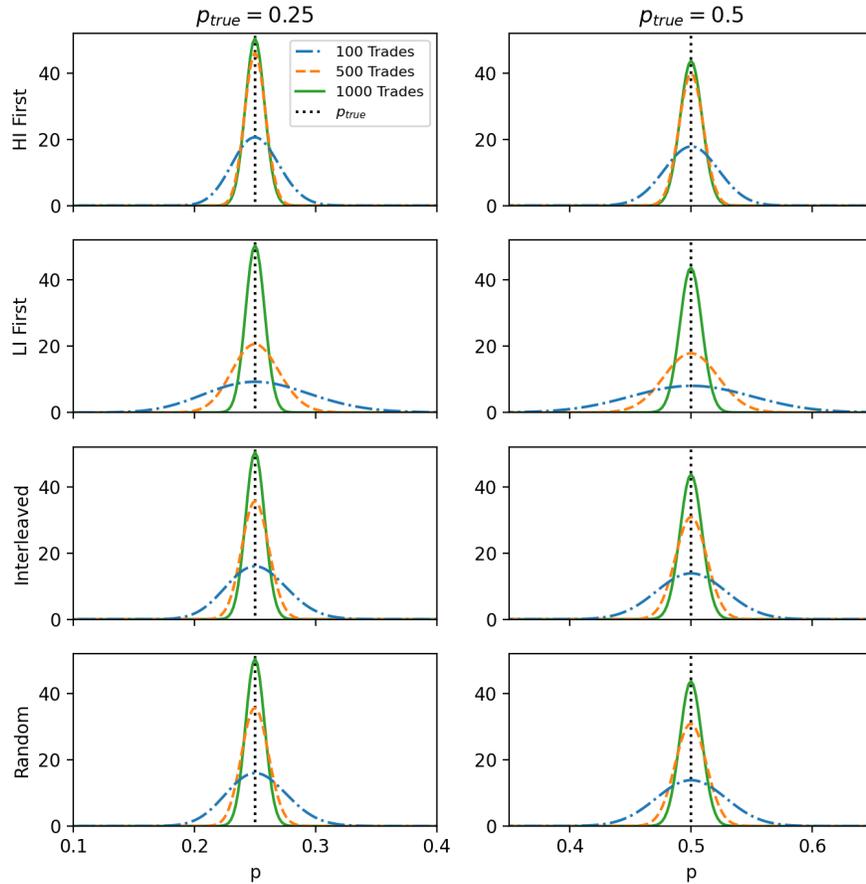


Figure 1: Posterior distribution (probability density) induced by outstanding shares averaged over 10^4 simulations after 100, 500, and 1000 trades, for $p_{true} \in \{0.25, 0.5\}$ and all 4 trader orderings in SET1. The support is $[0, 1]$ but we exclude the thin tails for visual clarity.

however, the market still performs the aggregation function of progressively reducing uncertainty about the (correct) mean (although the peak at market closing is not as sharp as that for $p_{true} = 0.25$).

p_{true}	Ordering	$10^3 \times (\text{Avg HI})$	$10^3 \times (\text{Avg LI})$
0.25	HI First	6.65 ± 0.03	0.19 ± 0.02
	LI First	1.77 ± 0.04	5.05 ± 0.03
	Interleaved	5.55 ± 0.05	1.31 ± 0.04
	Random	5.60 ± 0.05	1.22 ± 0.04
0.50	HI First	6.39 ± 0.03	0.17 ± 0.02
	LI First	1.77 ± 0.04	4.79 ± 0.03
	Interleaved	5.49 ± 0.04	1.04 ± 0.04
	Random	5.48 ± 0.04	1.07 ± 0.04

Table 1: Net compensations of HI and LI traders, averaged over all traders of respective types for each simulation, then averaged over 10^4 simulations (with 95% confidence bounds), for $p_{true} \in \{0.25, 0.5, 0.75\}$ and all 4 trader orderings in SET1.

Table 1 shows the expected per-trader average compensation for each trader type (HI, LI) estimated by averaging over 10^4 market

instances for each of 2×4 experimental conditions. It is noteworthy that the numbers are comparable for the (systematically) **Interleaved** and (uniformly) **Random** orderings. But the salient takeaway is that, for every p_{true} considered, HI traders get less compensation on average than LI traders when LI traders arrive first; interleaving, even with a leading LI trader, does not have this effect. Figure 1 offers a qualitative explanation of this phenomenon: even with the low information injection rate of **LI First**, the market gets “saturated” with information long before (~ 100 trades) all LI traders have completed their trades (500 trades); subsequent trades (by HI traders) produce minimal marginal improvement to aggregation quality, hence these traders tend to be compensated less (regardless of the absolute informational content of their trade). The significantly large difference in compensation is still surprising!

Does the **LI first** ordering *always* result in a higher (expected per-trader average) compensation for LI traders than for HI traders, regardless of the number of traders $2n$ and the value of p_{true} ? To answer this question, we derived analytic expressions for the expected per-trader average compensations of LI traders (denoted by \widehat{R}_n^{LI}) and HI traders (denoted by \widehat{R}_n^{HI}) in an **LI first** sequence as a function of n (the number of traders of each type), m (HI

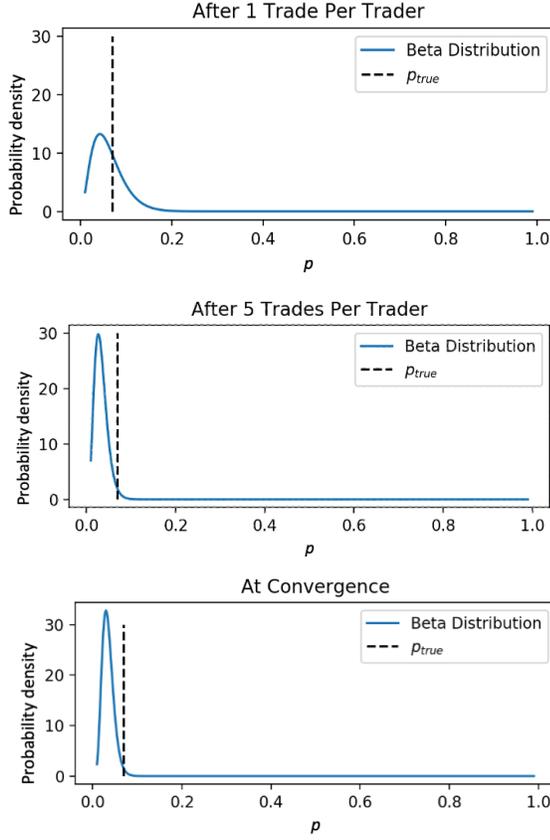


Figure 2: Evolution of market belief (top to bottom) in round 25 of the worst case ($\Delta_f \approx 0.766$) for LB First ordering and $B_0 = 10$.

informativeness), and p_{true} . This expectation is with respect to the joint distribution of all traders' (independent) private samples. We found by direct calculation that, for $m = 5$ and each $p_{true} \in [0.05, 0.10, 0.15, \dots, 0.95]$, there exists a value of $n \in \{1, 2, \dots\}$, denoted by n^* , such that $\widehat{R}_n^{HI} \geq \widehat{R}_n^{LI}$ for every $n < n^*$ and $\widehat{R}_n^{HI} < \widehat{R}_n^{LI}$ for every $n \geq n^*$. Moreover, n^* increases steadily with p_{true} up to $p_{true} = 0.5$ and then decreases symmetrically, with $n^* = 22$ for $p_{true} = 0.5$ (peak) and $n^* = 15$ for $p_{true} \in \{0.25, 0.75\}$. We can thus conclude that there exist small values of n for which HI traders can expect to sustain a higher compensation on average than LI traders even if the former appear later in the sequence; however, for large enough values of n (e.g. $n = 500$ in our experiments), early arrival always dominates informativeness in this regard.

4.2.2 Trader Order and Budgets. As for SET1, we start with qualitative results on aggregation. We measure the aggregation efficacy of a market instance (each round in each simulation) by Δ_f , the Euclidean distance between the final expected sufficient statistics of the market $[\mathbb{E}_\eta[\ln p], \mathbb{E}_\eta[\ln(1-p)]]^T$ and the corresponding vector induced by the revealed Bernoulli parameter p_{true} , i.e. $[\ln(p_{true}), \ln(1-p_{true})]^T$. Values of Δ_f (lower is better) that we

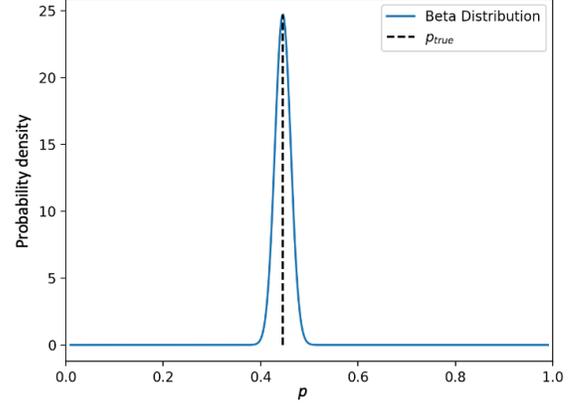


Figure 3: Final market belief in round 25 of the best case ($\Delta_f \approx 0.001$) for LB First ordering and $B_0 = 10$.

compute over all instances in our experiment suggest that the market is a good aggregator on average. For example, the average value of Δ_f for round 25 (the round in which LB traders are least constrained by budget) over all 10^3 simulations for **LB First** and $B_0 = 10$ is 0.099; for reference, the Euclidean distance between the above sufficient statistics vectors evaluated at the extreme values of the support of p_{true} , i.e. $\{0.05, 0.95\}$, is 4.16; that from either extreme to $p = 0.5$ is 2.39.

However, we will now discuss outliers revealed by the distribution of Δ_f across instances (deferring more detailed statistics on Δ_f to the full version). We found that, for $B_0 = 10$ and the **LB First** ordering, the highest and lowest values of Δ_f in round 25 across all 10^3 simulations are approximately 0.766 and 0.001 respectively; we name the corresponding market instances the *worst case* and the *best case* respectively. Figure 2 shows select snapshots of the market's evolving belief in the worst case. The market continually shifts its probability density in the correct direction but ends up concentrating the probability around an incorrect mean p^* that is close to but more extreme than p_{true} ($p^* = p_{true} - 0.035$). Contrast this with Figure 3 which depicts the final market belief in the best case where p_{true} is close to 0.5. This suggests that the market under our budget constraints performs as a better aggregator for less certain events, i.e. when p_{true} is safely distant from the extremes of the support $[0.05, 0.95]$.

We now come to the evolution of the budget of LB traders as a result of augmentation of the initial budget for each new round with compensation earned in the previous round. As seen in Figure 4, LB traders are able to grow their budget over successive rounds, as expected. The earlier the LB traders arrive, the higher is the rate of this growth; the growth (on average) appears convex in general and virtually linear for higher values of B_0 considered. One observation worth mentioning is that **Interleaved** significantly dominates **Random** in terms of growth rate for all values of B_0 considered. Perhaps more interestingly, the rate of growth displays a non-monotonic dependence on the size of the initial budget — with $B_0 = 2, 5, 10$, the budget after round 25 under **LB First** ordering is approximately 1.85, 2, and 1.6 times its initial value (on average).

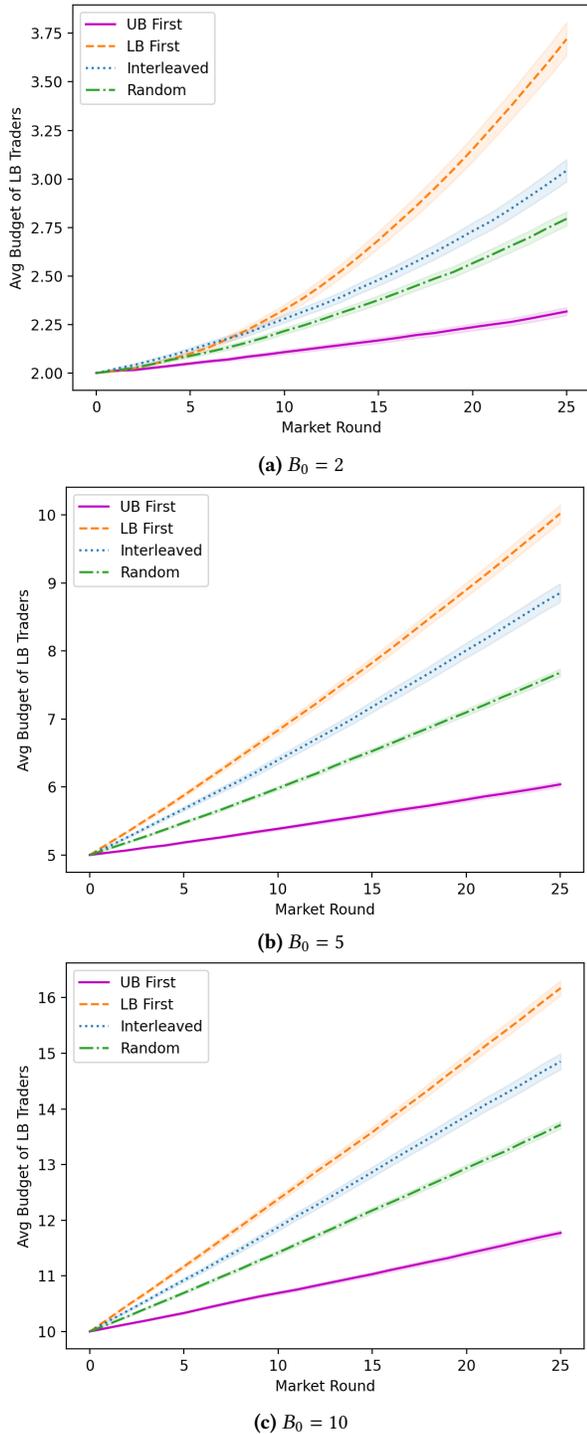


Figure 4: Change in budget of LB traders with an initial budget of B_0 over rounds 1 – 25, averaged over 10^3 simulations for $p_{true} \sim U[0.05, 0.95]$ and all 4 trader orderings.

However, averaging over all LB traders obscures the impact of the *internal order* among them within each round on the enhancement

of their trading power from one round to the next. To gauge this impact, we measured the number of datapoints that each LB trader is able to incorporate into her trading decision (owing to the constraint in Section 3.2.2) over successive rounds. In the left and right panels of Figure 5, we show this number, averaged over all 10^3 simulations, for the first and last LB traders (in terms of the round-robin sequence of market entries) respectively, all with an initial budget of 10, for three representative market rounds under **LB First** and **UB First**; results for other experimental conditions are qualitatively similar and are omitted. Note that the average number of market entries of a trader over all rounds in all simulations is in the range $[10.9, 11.8]$ over all traders, the number of re-entries being almost always determined by the ϵ -based and not the k_{max} -based stop criterion (Section 4.2.2) – we cut off the horizontal axis at 10.

The takeaways from Figure 5 are as follows. The number of datapoints used per entry (weakly) decreases with re-entries within a round; this is understandable since the residual budget of each trader, and hence the amount of information she is able to trade on, (weakly) decreases until market closure at the end of the round. Of course, all LB traders are able to trade on more information in later rounds, in particular round 25, than in the initial rounds due to their increased start-of-the-round budget. However, while the first LB trader can use all her $m = 5$ datapoints at (or possibly before) round 10 under **LB First**, the average number of datapoints that the last LB trader manages to use remains virtually fixed at around 4.5 which is strictly smaller than $m = 5$ – this situation appears comparable to the less favorable **UB First** ordering.

5 CONCLUSION AND FUTURE WORK

We presented a study on the effects of arrival time on trader compensation under a new market ecosystem designed to capture the aggregate belief of a trader population about a random outcome-generating process. While trader belief updates follow a previously studied Bayesian process, we modeled the difference in trader *informativeness* in a novel way and studied its interplay with other design parameters. When traders have differing levels of informativeness, we systemically gauged the extent to which earlier arrival can lead to *larger* average compensation for traders with *lower* informativeness. When traders have varying levels of *initial budget*, we showed that those who trade earlier become *effectively budget-unlimited* within a few market rounds whereas those who arrive later are *always constrained* by their budgets, even after several rounds. Thus our findings strongly indicate that any systematic study of inherent trader attributes in prediction markets must appropriately control for the confounding effects of trader ordering.

One direction we are actively pursuing is to establish the theoretical underpinnings of all our numerical results (similar to the last paragraph of Section 4.2.1). Further experimental assessment of design and modeling issues is also on our agenda. For example, budget-limited traders could strategize over arrival time or could choose other ways to incorporate their available information into the market – perhaps by choosing the ‘best’ subset of datapoints that their budget allows, or by solving a constrained optimization problem as in Devanur et al. [11]. We believe that our framework lays the foundation for further investigations along these lines.

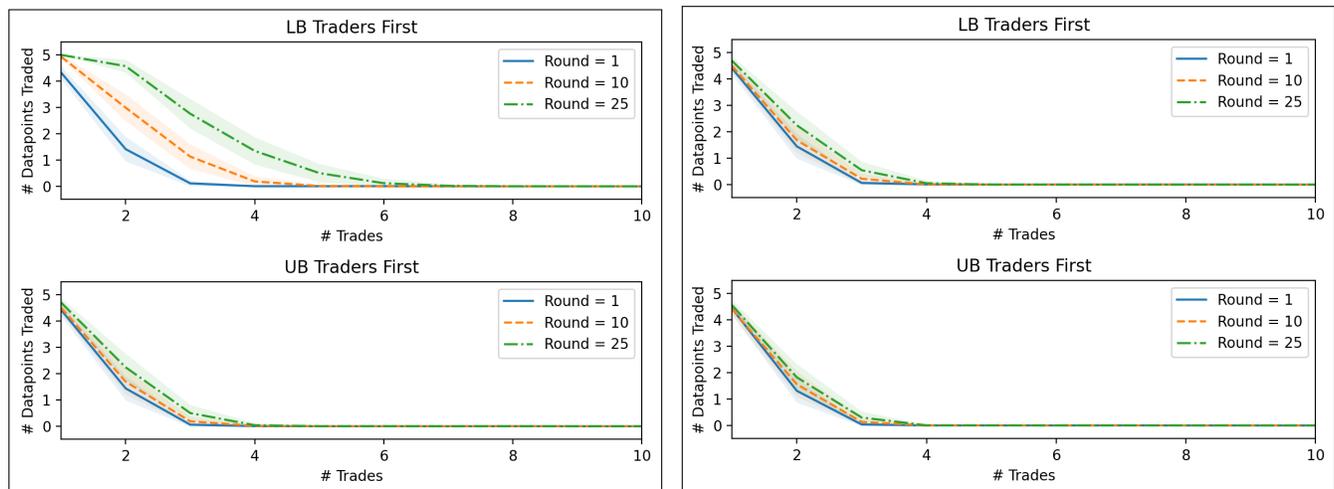


Figure 5: Average number of datapoints used to make a budget-constrained trading decision by the first (left) and last (right) trader in the sub-sequence of multi-shot LB traders within each of rounds 1, 10, and 25 over 10^3 simulations ($p_{true} \sim \text{UNIFORM}[0.05, 0.95]$) with $B_0 = 10$ for two different trader orderings.

REFERENCES

- [1] Jacob Abernethy, Sindhu Kutty, Sébastien Lahaie, and Rahul Sami. 2014. Information aggregation in exponential family markets. In *Proceedings of the 15th ACM conference on Economics and Computation*. 395–412.
- [2] Jerry Anunrojwong, Yiling Chen, Bo Waggoner, and Haifeng Xu. 2019. Computing Equilibria of Prediction Markets via Persuasion. In *International Conference on Web and Internet Economics*. Springer, 45–56.
- [3] Kenneth J Arrow, Robert Forsythe, Michael Gorham, Robert W Hahn, Robin Hanson, Daniel Kahneman, John O Ledyard, Saul Levmore, Robert Litan, Paul R Milgrom, Forrest D. Nelson, George R. Neumann, Marco Ottaviani, Charles R. Plott, Thomas C. Schelling, Robert J. Shiller, Vernon L. Smith, Erik Snowberg, Shyam Sunder, Cass R. Sunstein, Paul C. Tetlock, Philip E. Tetlock, Hal R. Varian, Justin Wolfers, and Eric Zitzewitz. 2007. Statement on Prediction Markets. *AEI-Brookings Joint Center Related Publication* 7 (2007).
- [4] J. Berg, R. Forsythe, F. Nelson, and T. Rietz. 2008. Results from a dozen years of election futures markets research. *Handbook of Experimental Economics Results* 1 (2008), 742–751.
- [5] Alina Beygelzimer, John Langford, and David M Pennock. 2012. Learning performance of prediction markets with Kelly bettors. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 3*. 1317–1318.
- [6] M. Chakraborty, S. Das, A. Lavoie, M. Magdon-Ismail, and Y. Naamad. 2013. Instructor Rating Markets. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*. 159–165.
- [7] Yiling Chen and David M. Pennock. 2007. A utility framework for bounded-loss market makers. In *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence (UAI)*. 49–56.
- [8] Yiling Chen and Jennifer Wortman Vaughan. 2010. A new understanding of prediction markets via no-regret learning. In *Proceedings of the 11th ACM Conference on Electronic Commerce*. 189–198.
- [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] B. Cowgill and E. Zitzewitz. 2015. Corporate prediction markets: Evidence from Google, Ford, and Firm X. *Review of Economic Studies* 82, 4 (2015), 1309–1341.
- [11] Nikhil Devanur, Miroslav Dudík, Zhiyi Huang, and David M Pennock. 2015. Budget constraints in prediction markets. In *Proceedings of the Thirty-First Conference on Uncertainty in Artificial Intelligence*. 238–247.
- [12] Persi Diaconis and Donald Ylvisaker. 1979. Conjugate Priors for Exponential Families. *The Annals of Statistics* 7, 2 (1979), 269–281.
- [13] Miroslav Dudík, Sébastien Lahaie, Ryan Rogers, and Jennifer Wortman Vaughan. 2017. A decomposition of forecast error in prediction markets. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*. 4374–4383.
- [14] Miroslav Dudík, Xintong Wang, David M Pennock, and David M Rothschild. 2021. Log-time Prediction Markets for Interval Securities. *20th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (2021).
- [15] Rafael Frongillo, Yiling Chen, and Ian Kash. 2015. Elicitation for aggregation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 29.
- [16] Rafael M Frongillo, Nicolas D Penna, and Mark D Reid. 2012. Interpreting prediction markets: a stochastic approach. In *Advances in Neural Information Processing Systems*. 3266–3274.
- [17] Andreas Graefe and J Scott Armstrong. 2011. Comparing face-to-face meetings, nominal groups, Delphi and prediction markets on an estimation task. *International Journal of Forecasting* 27, 1 (2011), 183–195.
- [18] R. Hanson. 2003. Combinatorial information market design. *Information Systems Frontiers* 5, 1 (2003), 107–119.
- [19] R. Hanson. 2007. Logarithmic Market Scoring Rules for Modular Combinatorial Information Aggregation. *Journal of Prediction Markets* 1, 1 (2007), 3–15.
- [20] Peter D Hoff. 2009. *A first course in Bayesian statistical methods*. Vol. 580. Springer.
- [21] Krishnamurthy Iyer, Ramesh Johari, and Ciamac C Moallemi. 2014. Information Aggregation and Allocative Efficiency in Smooth Markets. *Management Science* 60, 10 (2014), 2509–2524.
- [22] 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.
- [23] Kathryn Blackmond Laskey, Wei Sun, Robin Hanson, Charles Twardy, Shou Matsumoto, and Brandon Goldfeder. 2018. Graphical model market maker for combinatorial prediction markets. *Journal of Artificial Intelligence Research* 63 (2018), 421–460.
- [24] Michael Ostrovsky. 2012. Information aggregation in dynamic markets with strategic traders. *Econometrica* 80, 6 (2012), 2595–2647.
- [25] D. Pennock and R. Sami. 2007. Computational Aspects of Prediction Markets. In *Algorithmic Game Theory*, N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani (Eds.). Cambridge University Press.
- [26] David M Pennock. 1999. *Aggregating probabilistic beliefs: Market mechanisms and graphical representations*. Ph.D. Dissertation. University of Michigan.
- [27] Rajiv Sethi and Jennifer Wortman Vaughan. 2016. Belief aggregation with automated market makers. *Computational Economics* 48, 1 (2016), 155–178.
- [28] Amos J Storkey, Zhanxing Zhu, and Jinli Hu. 2015. Aggregation Under Bias: Rényi Divergence Aggregation and Its Implementation via Machine Learning Markets. In *Machine Learning and Knowledge Discovery in Databases*. Springer, 560–574.
- [29] Martin J. Wainwright and Michael I. Jordan. 2008. Graphical Models, Exponential Families, and Variational Inference. *Foundations and Trends in Machine Learning* 1, 1–2 (2008), 1–305.
- [30] Justin Wolfers and Eric Zitzewitz. 2004. Prediction Markets. *Journal of Economic Perspectives* 18, 2 (2004), 107–126.