

Prediction Markets With Budget Constraints and Set Theoretic Beliefs

Nikhil R. Devanur, Microsoft Research
 Miro Dudik, Microsoft Research
 Zhiyi Huang, Univ. of Pennsylvania
 David Pennock, Microsoft Research

We consider two different scenarios in prediction markets where an agent cannot simply move the market price to his current belief. The first one is when the agent has a budget constraint which restricts the trades he could make. Extending recent work of [Fortnow and Sami 2012] on this problem, we give a rich geometric characterization of agent behavior in such cases. We show that the agent moves the price to the Bregman projection of his belief onto the budget set, the set of prices his budget constraint allows him to move to. Further we show that for the case of complete markets, an agent with budget B is equivalent two agents with budgets $B/2$ trading one after the other. In other words a sequence of agents with the same belief is equivalent to a single agent with the same belief and with budget equal to the sum of the budgets of all the agents. This is an indication that prediction markets can still aggregate information when agents have budget constraints if there are many agents with the same belief.

The second scenario we consider is when an agent only has limited information, in the sense that he only knows that the probability distribution belongs to a certain (convex) set. The optimal trade for the agent is the one that maximizes the worst case profit where the profit is taken over all distributions in his belief set. In this case we once again give a geometric characterization of the final price as a Bregman projection of the current price onto the belief set. We also relate this model to other models in the literature.

1. INTRODUCTION

A prediction market is a central clearinghouse for people with differing opinions about the likelihood of an event—say Barack Obama to win the U.S. Presidential election—to trade monetary stakes in the outcome with one another. At equilibrium, the price to buy a contract paying \$1 if Obama wins reflects a consensus of sorts on the probability of the event. At that price, and given the wagers already placed, no agent is willing to push the price further up or down. Prediction markets have a good track record of forecast accuracy in many domains [Goel et al. 2010; Wolfers and Zitzewitz 2006].

When agents are constrained in how much they can trade only by risk aversion, prediction market prices can be interpreted as a weighted average of traders' beliefs [Wolfers and Zitzewitz 2006; Beygelzimer et al. 2012], a natural reflection of the “wisdom of the crowd”. However, when agents are budget constrained, discontinuities and idiosyncratic results can arise [Manski 2006; Eisenberg and Gale 1959] that call into question whether the equilibrium price can be trusted to reflect any kind of useful aggregation.

In this paper, we examine more closely what happens in a prediction market when agents are budget constrained, and show that results can still be meaningful. We also look at the implications when traders have imprecise beliefs, meaning they know only

Permission to make digital or hard copies of part or all 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 show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2013 ACM 0000-0000/2013/06-ART39 \$15.00

DOI: <http://dx.doi.org/10.1145/0000000.0000000>

that the probabilities of events fall into some convex set, but are not sure of, or not willing to commit to, a single point belief.

We consider prediction markets with an automated market maker [Abernethy et al. 2011; Hanson 2007]. In such markets the market maker always offers a trade. All agents trade (buy or sell) with the market maker. This is in contrast to a two-sided market where the market only matches agents with a mutual agreement to trade. The interesting market makers are the ones that are (myopically) incentive compatible in the sense that the best strategy for an agent is to move the market price to equal his own belief. Such markets have been extensively studied and are fairly well understood.

In this paper we investigate scenarios where incentive compatibility breaks—that is, an agent cannot move the market price to exactly equal his belief—for two different reasons: budget constraints and imprecise beliefs.

The first scenario in which an agent may not be able to move the market price to his own belief is when he has a budget constraint. The meaning of a budget constraint is not immediately clear in this context. We use the definition of a *natural budget constraint* introduced by [Fortnow and Sami 2012]. Any trade an agent makes could potentially incur a loss. If an agent buys a security, then his loss would be the price he paid to buy the security in case it does not pay off. If the agent sells a security and if the event actually happens then he would have to pay out money as promised by the security. His loss would be the pay out minus the price he got from selling the security. In such cases, the market could ask the agent to post a collateral equal to make sure he has the ability to pay for the security. In general an agent could trade a bundle of securities, selling some and buying others. The worst case loss of such a trade is the maximum loss the agent could incur, taken over all possible outcomes. The natural budget constraint, for a certain budget, says that the agent is only allowed those trades for which the worst case loss is no greater than the budget. An alternate interpretation of this constraint for securities of binary events is obtained by considering selling a security for an event as equivalent to buying a security for the complement of the event. Then the agent only buys securities and in that case the natural budget constraint says that the total cost of the securities the agent buys cannot exceed his budget.

[Fortnow and Sami 2012] considered the following natural question: given that an agent cannot move the price to his own belief, is it possible that he moves the price in a straight line towards his belief? It seems quite natural that the movement of prices should be in the direction of the belief. Somewhat surprisingly they showed that for any *scoring rule*¹ there always exist situations in which the best strategy of the agent is not to move the price along the straight line towards his belief.

In this paper we shed more light on this seemingly surprising result by considering the *geometry* of the budget constraints. In particular our results are as follows.

- We give a simpler, geometric proof of the impossibility result of [Fortnow and Sami 2012].
- The budget set is the set of feasible prices that the agent can move the market price to, under the budget constraint. We show that the budget set is an intersection of *Bregman balls*² with centers at points corresponding to payoffs for all possible outcomes. The radii are determined by the current price.

¹ A scoring rule is a more general form of incentivizing agents to elicit information about probabilistic events. It specifies a payoff for an agent making a certain prediction as a function of the observed outcome.

²These are balls where the distances are measured by Bregman divergences. Bregman divergences generalize the notion of a distance to asymmetric functions. They are well studied and play a crucial role in convex optimization, machine learning, etc.

- We give a precise characterization of how an agent with a budget constraint moves the market price. The agent moves the market price to the Bregman projection³ of his belief onto the budget set. This is very intuitive, the agent does try to get as close to his belief as possible, only as measured by the Bregman divergence. A particularly interesting case is the *quadratic* scoring rule, for which the Bregman divergence is the square of the Euclidean distance. In this case the Bregman projection is just the Euclidean projection. The reason the new price is not directly in the direction of the belief is due to the shape of the budget set. Another message of this characterization is that the current price is only important to determine the budget set. It shows why the impossibility result of [Fortnow and Sami 2012] is not that surprising on hindsight. The projection may be locally insensitive to the current price. If the current price changes a little bit such that the tight constraint in the budget set does not change, then the projection also does not change.
- We consider another natural question that arises due to budget constraints. Suppose that there is a sequence of agents with the same belief but with small budgets, who move the market price one after the other. Is the eventual price the same as what a single agent, with the same belief and with a budget equal to the sum of the budgets of all the previous agents, would move the price to? We answer this question in the affirmative, for the case of complete markets, i.e., there is a security for every possible outcome of the market and the market elicits the complete probability distribution. This conclusion is heartening: it says that prediction markets can aggregate information even in the presence of budget constraints as long as there are sufficiently many agents with the same belief. We call this the *associativity of budgets* property.

The second scenario we consider is when agents have *set theoretic beliefs* instead of point beliefs. What we mean by this is that the agent may not know the exact probability distribution. Instead, we assume that the agent only knows that the probability distribution belongs to a certain set. For example, suppose that there are two binary events, A and B. The agent may only know that the probability of event A is greater than the probability of event B, or that the probability of either A or B is greater than 50%. We model the behavior of such an agent by saying that he only trades when he is guaranteed a positive expected payoff for all probability distributions in his belief set. In particular the optimal trade for the agent is the one that maximizes his worst case expected payoff where the worst case is taken over all probability distributions in his belief set.

Our main result for set theoretic beliefs is once again a geometric characterization of the optimal trade for an agent. We show that the optimal trade for an agent is to move the market price to the Bregman projection of the current price onto his belief set. This characterization in terms of the Bregman projection is not as intuitive as in the case of budget constraints.⁴ A priori, it is not even clear that there is always a trade that gives a strictly positive worst case payoff whenever the current price is not in the belief set. It turns out that for any hyperplane that separates the current price from the belief set, the trade that corresponds to the normal of this hyperplane always has a positive worst case payoff. This implies that the final price should end up on

³ The Bregman projection of a point onto a set is the point in the set that is closest to the given point when measured in terms of a Bregman divergence.

⁴ Bregman projections are actually of two types, where the point to be projected could be the first argument or the second. These could be different due to the asymmetry of a Bregman divergence. The Bregman projection in the budget constraint case is when the point is the first argument and the Bregman projection in the set theoretic belief case is when the point is the second argument.

the boundary of the belief set, since we can always find a trade with a positive worst case payoff otherwise. This does not say why this final point should be the Bregman projection. We still don't have an intuitive explanation for this phenomenon. This is reflected in the difference in the proof of this result as compared to the one for the budget constraint. While the proof for the budget constraint is straight forward, the proof of this result uses strong convex programming duality. The proof idea is the same as the one used by [Grunwald and Dawid 2004] who used it to show the equivalence between finding an entropy maximizing distribution and a distribution that minimizes a particular worst case loss. We also need an assumption that the belief set is convex and satisfies Slater's condition, which is a mild condition under which strong duality holds. This result is once again a theoretical justification for the effectiveness of prediction markets. It shows that prediction markets can aggregate information even when the agents are not completely sure of the probability distribution.

Our set theoretic belief model gives an interesting alternate perspective on the problem of eliciting properties of probability distributions considered by [Lambert and Shoham 2009]. There, a property is modeled as a partition of the set of all probability distributions. They consider proper scoring rules to elicit the set in the partition that contains the agent's belief. They showed that a proper scoring rule exists if and only if all the sets in the partition form a Voronoi diagram. We give an alternative perspective to this problem via our set theoretic belief model and compare the advantages and disadvantages of our model with theirs in Section 5.

Our set theoretic belief model is also related to [Bhattacharjee and Goel 2006, 2007] who consider the problem of ranking a set of n binary events by decreasing order of probability. They give a mechanism where agents with limited information, such as that event A is more likely than event B can be incentivized to participate in the mechanism. The mechanism aggregates information from all participants and produces a ranking. We note that their mechanism is actually equivalent to a prediction market, the agents with limited information fit exactly in our model of set theoretic beliefs.

Other related work

There is a rich literature on scoring rules and prediction markets. Some of the much studied scoring rules are the quadratic scoring rule of [Brier 1950] and the logarithmic market scoring rule (LMSR) of [Hanson 2007]. We consider convex cost function based prediction markets, which is without loss of generality. Even though they did not phrase it this way, [Gneiting and Raftery 2007] showed that any proper scoring rule is equivalent to a convex cost function based prediction market. Cost function based prediction markets have been considered earlier by [Hanson 2003; Chen and Pennock 2007]. [Chen and Vaughan 2010] showed an interesting equivalence between complete cost function based prediction markets and no-regret learning algorithms. [Abernethy et al. 2011] used this connection to design market makers for more general prediction markets. They also showed that under certain assumptions the choice of a convex cost function based prediction market is without loss of generality.

[Fortnow and Sami 2012] introduced the notion of a natural budget constraint for scoring rules and considered the question of whether there exists a scoring rule such that the forecast of the agent is always on the straight line joining the market price and the agent's belief. They showed that no such scoring rule exists and also showed that it is not always possible to infer the agent's belief from his forecast, even when his budget is public knowledge. We shed more light on their impossibility results and give a richer characterization of the agent behavior in the presence of budget constraint.

2. PRELIMINARIES

Consider a probability space with a finite set of outcomes Ω . A *security* is a financial instrument whose payoff depends on the realization of an outcome in Ω . In other words, the payoff of a security is simply a random variable of the probability space. A security can be traded before the realization is observed with the intention that the price of a security serves as a prediction for the expected payoff of the outcome. A prediction market consists of one or more securities, which are represented by a vector of random variables, denoted by $\phi : \Omega \rightarrow \mathbb{R}^n$. We will refer to ϕ simply as *a security*, keeping in mind that it is really a set of n securities.

An *automated market maker* always offers to trade a security, for the right price. In fact the price is the current prediction of the market maker for the expectation of ϕ . A cost function based market maker is based on a differentiable convex *cost function*, $C : \mathbb{R}^n \rightarrow \mathbb{R}$. The domain of C is the vector of the *number of outstanding shares*⁵ for the security ϕ , which we denote by $q \in \mathbb{R}^n$. We also refer to q as the *state* of the market.

The *instantaneous price* of the security is simply the gradient of C at q :

$$p(q) := \nabla C(q) .$$

The price of a security changes continuously as more of the security is traded, so it is useful to consider the cost of trading a given quantity of the security. Buying $\delta \in \mathbb{R}^n$ units of the security (where a negative value corresponds to selling) costs

$$\int_q^{q+\delta} p(\bar{q}) d\bar{q} = C(q + \delta) - C(q) .$$

When the outcome ω is realized, the δ units of the security pays off an amount of $\delta \cdot \phi(\omega)$. Thus, the realized utility of a trader whose trade δ moved the market state from q to $q' = q + \delta$ is

$$U(q', \omega; q) := (q' - q) \cdot \phi(\omega) - C(q') + C(q) .$$

Let \mathcal{M} be the convex hull of the payoff vectors:

$$\mathcal{M} := \text{conv}\{\phi(\omega) : \omega \in \Omega\} .$$

It is easy to see that \mathcal{M} contains exactly the vectors $\mu \in \mathbb{R}^n$ which can be realized as expected payoffs $\mathbb{E}[\phi]$ for some probability distribution over Ω . For a trader who believes that $\mathbb{E}[\phi] = \mu$, the expected utility takes form

$$U(q', \mu; q) := \mathbb{E}[U(q', \omega; q)] = (q' - q) \cdot \mu - C(q') + C(q) .$$

A key property satisfied by expected utility is *path independence*: for any $q, \bar{q}, q' \in \mathbb{R}^n$

$$U(q', \mu; \bar{q}) + U(\bar{q}, \mu; q) = U(q', \mu; q) ,$$

i.e., risk-neutral traders have no incentive to split their trades. For a risk-neutral trader, $q' \in \mathbb{R}^n$ is an optimal action if and only if

$$\mu = \nabla C(q') = p(q')$$

(this follows from the first-order optimality conditions). In other words, the trader is incentivized to move the market to the prices corresponding to his belief as long as such prices exist. Therefore, we sometimes refer to the price reached by a trader's action as the *prediction*.

⁵We allow trading fractions of a security. Negative values correspond to short-selling.

There may be several actions yielding the same expected utility. We make a standard assumption that our market is *expressive* in the sense that

$$\mathcal{M} \subseteq \text{cl } \text{im } p$$

where cl denotes the closure in the standard topology and $\text{im } p := \{p(q) : q \in \mathbb{R}^n\}$ is the image of the price map. In this case, risk-neutral traders can always express beliefs $\mu \in \mathcal{M}$ (potentially in the limit of infinitely many trades, if $\mu \in \mathcal{M} \setminus \text{im } p$).

Example 2.1. The first example of a cost function, which is expressive for arbitrary payoffs ϕ , is the quadratic cost function defined by $C(q) = \frac{1}{2}\|q\|^2$. In this case, $p(q) = q$, and $U(q', \mu; q) = \frac{1}{2}\|q - \mu\|^2 - \frac{1}{2}\|q' - \mu\|^2$. It is clear that the expected utility is maximized when $p(q') = q' = \mu$.

Example 2.2. Our second example is Hanson's logarithmic market-scoring rule (LMSR), which is expressive for the *complete market* defined by the indicator payoff function $\phi : \Omega \rightarrow \mathbb{R}^\Omega$

$$\phi_{\omega'}(\omega) = \begin{cases} 1 & \text{if } \omega = \omega' \\ 0 & \text{otherwise.} \end{cases}$$

In this case \mathcal{M} is the simplex in \mathbb{R}^Ω and beliefs μ are in one-to-one correspondence with probability distributions over Ω . The LMSR cost function is

$$C(q) = \ln \left(\sum_{\omega \in \Omega} e^{q_\omega} \right)$$

with the price

$$p_\omega(q) = \frac{e^{q_\omega}}{\sum_{\omega' \in \Omega} e^{q_{\omega'}}} = e^{q_\omega - C(q)} .$$

For $\mu \in \mathcal{M}$, the expected utility function takes form

$$U(q', \mu; q) = \sum_{\omega} \mu_{\omega} \left(\ln p_{\omega}(q') - \ln p_{\omega}(q) \right) = \text{KL}(\mu \| p(q)) - \text{KL}(\mu \| p(q')) ,$$

where $\text{KL}(\mu \| \nu) = \sum_{\omega} \mu_{\omega} \ln(\mu_{\omega} / \nu_{\omega})$ is the KL-divergence. KL-divergence is not symmetric, but it is non-negative, and zero only if the arguments are equal. Thus, the expected utility is clearly maximized if and only if $\mu = p(q')$.

The above two examples illustrate that the expected utility can be written as the difference of two terms measuring the distance between the belief and the market state. This is true not just for the two costs listed above but for arbitrary convex differentiable costs C . The distance measure above is the mixed *Bregman divergence*.⁶ To define the Bregman divergence formally, first let $C^* : \mathbb{R}^n \rightarrow \mathbb{R}^*$, where $\mathbb{R}^* = \mathbb{R} \cup \{\infty\}$, be the convex conjugate of C :

$$C^*(\nu) := \sup_{q' \in \mathbb{R}^n} [q' \cdot \nu - C(q')] .$$

Since C^* is a supremum of linear functions, it is convex lower-semicontinuous. Up to a constant, it characterizes the maximum achievable utility on an outcome ω as

$$\sup_{q' \in \mathbb{R}^n} U(q', \omega; q) = C^*(\phi(\omega)) + [C(q) - q \cdot \phi(\omega)] .$$

⁶Our notion of Bregman divergence is more general than typically assumed in the literature.

The term in the brackets is always finite, but C^* might be positive infinite. We make a standard assumption that $C^*(\phi(\omega)) < \infty$ for all $\omega \in \Omega$, i.e., that the maximum achievable utility is bounded by a finite constant. Since the maximum utility of a trader is also the maximum loss of the market maker, this assumption means that the loss of the market maker is bounded by a finite constant. By convexity, this implies that $C^*(\mu) < \infty$ for all $\mu \in \mathcal{M}$. The Bregman divergence derived from C is the function $D : \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}^*$ measuring the maximum expected utility under belief μ

$$D(q, \mu) := C(q) + C^*(\mu) - \mu \cdot q = \sup_{q' \in \mathbb{R}^n} U(q', \mu; q) .$$

From the convexity of C and C^* and the definition of C^* , it is clear that:

- D is convex and lower-semicontinuous in each argument separately
- D is non-negative
- D is zero if and only if $p(q) = \nabla C(q) = \mu$

By the bounded loss assumption, Bregman divergence is finite on $\mu \in \mathcal{M}$. For $\mu \in \mathcal{M}$, we can write

$$U(q', \mu; q) = D(q, \mu) - D(q', \mu) . \quad (1)$$

Thus, maximizing the expected utility is the same as minimizing the Bregman divergence between the state q' and the belief μ .

For the quadratic cost, we have $C^*(\nu) = \frac{1}{2}\|\nu\|^2$ and $D(q, \nu) = \frac{1}{2}\|q - \nu\|^2$. For LMSR, we have $C^*(\nu) = -\sum_{\omega \in \Omega} \nu_\omega \ln \nu_\omega$, with the usual convention $0 \ln 0 = 0$, and $D(q, \nu) = \text{KL}(\nu \| p(q))$.

We will slightly abuse notation and frequently identify ω with $\phi(\omega)$. For instance, we will write $D(q, \omega)$ instead of the more verbose $D(q, \phi(\omega))$. This is a natural identification since $\phi(\omega)$ corresponds to the belief realized by the distribution that puts all the mass on the outcome ω .

3. BUDGET CONSTRAINTS

TODO-MD: we need better pictures demonstrating various constructions in the paper. At the minimum, in Figure 3, replace the third figure by a figure demonstrating the construction of L with the segments ℓ_0, ℓ_1 and maybe the outline of the initial convex hull $\text{conv}(X, \mu)$.

Trading in prediction markets needs an investment of capital. It is possible that an agent loses money on the trade, in particular $U(q', \omega; q)$ could be negative for some ω . One restriction on how an agent trades could be that he is unable to sustain a big loss, due to a budget constraint. In this section we study how such a restriction affects the behavior of an agent trading in the market. We consider the notion of *natural budget constraint* defined by [Fortnow and Sami 2012] which states that the loss of the agent is at most his budget, for all $\omega \in \Omega$. Formally, given a starting market state q and a budget of $B \geq 0$, the set of market states that satisfy the natural budget constraint is

$$\mathcal{B}(q, B) := \{q' : U(q', \omega; q) \geq -B \text{ for all } \omega \in \Omega\} .$$

We also refer to this set as the budget set. Next lemma shows that $\mathcal{B}(q, B)$ can be viewed as an intersection of D -balls centered at ω for $\omega \in \Omega$. Specifically, let $\mathcal{B}_\omega(r)$ denote the ball of radius r centered around ω :

$$\mathcal{B}_\omega(r) := \{q' : D(q', \omega) \leq r\} .$$

LEMMA 3.1. $\mathcal{B}(q, B) = \bigcap_{\omega \in \Omega} \mathcal{B}_\omega(D(q, \omega) + B)$.




Fig. 1. \circ —current state, \times —belief, \diamond —projected belief. (top-left) three circles bounding the allowed final states for budget 0.1 and projections of two different beliefs; (top-right) two different budget levels (0.1 and 0.03); (bottom) sequence of projections for varying budgets.

PROOF. The proof is immediate from Eq. (1). Note that $q' \in \mathcal{B}(q, B)$ if and only if $U(q', \omega; q) \geq -B$ for all ω . For a single ω this is equivalent to $D(q', \omega) \leq D(q, \omega) + B$, i.e., $q' \in \mathcal{B}_\omega(D(q, \omega) + B)$. Combining across all ω gives the intersection. \square

By convexity and lower-semicontinuity of D , Bregman balls $\mathcal{B}_\omega(r)$ are convex and closed. Therefore the budget set is also convex and closed. It is also non-empty since $q \in \mathcal{B}(q, B)$. The next question is, given such a budget constraint, what is the optimal action of an agent? Again, Eq. (1) implies that a risk-neutral agent will want to move the market to the point that is closest to his belief μ , while staying in the budget set:

$$\operatorname{argmax}_{q' \in \mathcal{B}(q, B)} U(q', \mu; q) = \operatorname{argmin}_{q' \in \mathcal{B}(q, B)} D(q', \mu) .$$

In general, there may be multiple q' optimizing this objective, or, as in the case of a belief $\mu \in \mathcal{M} \setminus \operatorname{im} p$ (and an infinite budget), the argmax set might be empty and the optimum may be only achievable in a limit. However, as the next theorem shows, there exists a unique price $\hat{\nu}$ that is reached in any of the cases. (TODO?: can we show that for finite B there always exists some (non-unique) action \hat{q} solving the problem? Maybe we can show that if the maximum were only obtained in a sequence, i.e., at a point $\hat{\nu} \notin \operatorname{im} p$, then an infinite budget is required relative to some outcome ω .)

THEOREM 3.2. *For any $q \in \mathbb{R}^n$, $B \geq 0$ and $\mu \in \mathcal{M}$, there exists a unique $\hat{\nu}$ such that any sequence $\{q_k\}_{k=1}^\infty$ maximizing*

$$\max_{q' \in \mathcal{B}(q, B)} U(q', \mu; q) \tag{2}$$

satisfies $p(q_k) \rightarrow \hat{\nu}$.

PROOF. TODO-MD: The general idea is to take the Fenchel dual of Eq. (2), show that the dual has a unique solution $\hat{\nu}$ and that for the maximizing sequence, $D(q_k, \hat{\nu}) \rightarrow 0$. \square

In the next section, we analyze Eq. (2) in more detail. In the remainder of this section, we revisit the question considered by Fortnow and Sami [2012]. They asked whether there is a market scoring rule such that the optimal states under the budget constraint induce prices that lie on the straight line connecting the initial price and the true expectation, and answered in the negative. We show a simple geometric proof of this result for cost function based markets.

TODO: I started modifying the proof below, but then realized that it is not completely straightforward. There are various hidden continuity assumptions, e.g., $p(q)$ is implicitly assumed to be continuous and invertible around the initial state—invertibility is assumed when we say that it is possible to maintain ν'_0 at the same distance from ω_0 as ν_0 . Also, obviously, the statement should require at least two dimensional \mathcal{M} .

LEMMA 3.3. *For all C , there exist q , B and μ such that an optimal action \hat{q} yields market prices $p(\hat{q})$ which are not on the line joining $p(q)$ and μ .*

PROOF. Suppose for now that we can choose q , B and μ such that an optimal state \hat{q} is on the boundary of the ball $\mathcal{B}_{\omega_0}(D(q, \omega_0) + B)$ for some *unique* ω_0 , i.e.,

$$\begin{aligned} D(\hat{q}, \omega_0) &= D(q, \omega_0) + B \\ D(\hat{q}, \omega) &< D(q, \omega) + B \text{ for all } \omega \neq \omega_0. \end{aligned} \tag{3}$$

Further, $p(q) \neq \phi(\omega_0)$. OLD PROOF CONTINUES with $\nu_0 := p(q)$ and $\mu^* := p(\hat{q})$: Suppose that ν_0, μ^* and μ are all on a straight line. (If not, then we are done already.) We claim that there exists ν'_0 not also on the same straight line such that condition Eq. (3) also holds when ν_0 is replaced by ν'_0 . The existence of such a ν'_0 follows from the continuity of D^* : if we choose ν'_0 close enough to ν_0 then the strict inequalities should continue to hold. Given that the strict inequalities hold, one can always choose ν'_0 so that the divergence from $\phi(\omega_0)$ does not change (since $\nu_0 \neq \phi(\omega_0)$) and that it is not on the given straight line. This leaves the budget constraint essentially unchanged, \mathcal{B}_{ω_0} which is the only tight constraint is unchanged while the change in the other constraints is too small to make them tight. Therefore μ^* is still the projection of μ onto the budget set and hence the optimal prediction, with ν'_0 as the starting point and budget B . Since ν'_0 is not on the straight line containing μ and μ^* , it is the counter example as needed.

We still need to justify the original choice of ν_0 , B and μ . Essentially we need to find a point on the boundary of the budget set where there is only one tight constraint. We give one particular construction here. We start with the choice of μ^* being any point such that the largest divergence to one of the $\phi(\omega)$'s is unique, i.e. $\arg \max_{\omega} \{D^*(\phi(\omega); \mu^*)\}$ is unique, and say is equal to $\{\omega_0\}$. Notice that condition Eq. (3) holds if we choose $\nu_0 = \phi(\omega_0)$ and $B = D^*(\phi(\omega_0); \mu^*)$. Again by the continuity of D^* , if we choose ν_0 to be close enough to $\phi(\omega_0)$ and set $B = D^*(\phi(\omega_0); \mu^*) - D^*(\phi(\omega_0); \nu_0)$ then the inequalities in Eq. (3) will continue to hold. Finally we need to pick μ such that its projection is μ^* . This is easily accomplished by picking μ to be just outside the boundary near μ^* . \square

4. BUDGET ADDITIVITY

In this section, we consider another natural question that arises due to budget constraints. Suppose that there is a sequence of agents with the same belief but with

small budgets, who move the market price one after the other. Is the eventual price the same as what a single agent, with the same belief and with a budget equal to the sum of the budgets of all the previous agents, would move the price to? We give affirmative answer to this question, and refer to this property as *budget additivity*.

4.1. Optimality conditions

The optimization problem under the budget constraints can be written as

$$\begin{aligned} \max_{q \in \mathbb{R}^n} & U(q, \mu; q_0) \\ \text{s.t.} & U(q, \omega; q_0) \geq -B \quad \forall \omega \in \Omega . \end{aligned} \quad (4)$$

The set of solutions of Eq. (4) will be denoted $\hat{Q}(B; q_0)$. The belief μ is assumed to be fixed throughout this section, so we suppress the dependence on μ . In Theorem 3.2, we showed that all the elements $\hat{Q}(B; q_0)$ must yield the same prices. The following theorem gives a more detailed structure:

LEMMA 4.1 (KKT LEMMA). *Let $q_0 \in \mathbb{R}^n$. Then $q \in \hat{Q}(B; q_0)$ for some $B \in \mathbb{R}$ if and only if there exists $X \subseteq \Omega$ such that*

- (a) $U(q, x; q_0) = U(q, x'; q_0)$ for all $x, x' \in X$
- (b) $U(q, \omega; q_0) \geq U(q, x; q_0)$ for all $x \in X$ and $\omega \in \Omega$
- (c) $p(q) \in \text{conv}(X, \mu)$

PROOF. We begin by forming a Lagrangian of Eq. (4), with non-negative multipliers $\lambda = (\lambda_\omega)_{\omega \in \Omega}$:

$$L(q, \lambda) = U(q, \mu; q_0) + \sum_{\omega} \lambda_{\omega} (U(q, \omega; q_0) + B) .$$

By differentiability and concavity of the objective and constraints, KKT conditions are both necessary and sufficient for optimality. KKT conditions state that q and λ solve the above problem if and only if the following hold:

- *primal feasibility*: $U(q, \omega; q_0) \geq -B$ for all $\omega \in \Omega$
- *dual feasibility*: $\lambda \geq 0$
- *first-order optimality*: $\nabla_1 L(q, \lambda) = 0$
- *complementary slackness*: $\lambda_{\omega} (U(q, \omega; q_0) + B) = 0$ for all $\omega \in \Omega$

We first show that KKT conditions imply (a)–(c). Assume that KKT conditions hold. Set X to be the set of outcomes with tight constraints, i.e., $X = \{x \in \Omega : U(q, x; q_0) = -B\}$. For this X , the conditions (a) and (b) hold by primal feasibility. We prove (c) by analyzing first-order optimality. First note that:

$$\nabla_1 U(q, \nu; q_0) = \nu - \nabla C(q) = \nu - p(q) .$$

Thus, first-order optimality is equivalent to

$$\begin{aligned} \nabla_1 U(q, \mu; q_0) + \sum_{\omega} \lambda_{\omega} \nabla_1 U(q, \omega; q_0) &= 0 \\ \mu - p(q) + \sum_{\omega} \lambda_{\omega} (\omega - p(q)) &= 0 \\ p(q) &= \frac{\mu + \sum_{\omega} \lambda_{\omega} \omega}{1 + \sum_{\omega} \lambda_{\omega}} . \end{aligned}$$

By complementary slackness, $\lambda_{\omega} = 0$ for $\omega \in \Omega \setminus X$, so this shows (c).

Now assume that (a)–(c) hold. If $X = \emptyset$, then KKT conditions hold for $B = \max_{\omega} [-U(q, \omega; q_0)]$ and $\lambda = 0$. If $X \neq \emptyset$, then KKT conditions hold for $B = -U(q, x; q_0)$ (where $x \in X$), and λ_{ω} representing $p(q)$ as a convex combination of X and μ . \square

4.2. Perpendiculars

Before we state our main result, we need to define the notion of a Bregman perpendicular to an affine space. For the quadratic cost, this notion will exactly coincide with the usual Euclidean perpendicular.

Let A be an affine space in \mathbb{R}^n and $q \in \mathbb{R}^n$ be a market state such that $p(q) \notin A$. Let a be an arbitrary point in A and define parallel affine spaces to A as

$$A_{\lambda} := A + \lambda(p(q) - a)$$

for $\lambda \in \mathbb{R}$. Finally, let

$$\nu_{\lambda} := \operatorname{argmin}_{\nu \in A_{\lambda}} D(q, \nu) .$$

By strict convexity of D in the second argument, ν_{λ} is either a singleton or an empty set. Let Λ be the set of λ 's where ν_{λ} is not an empty set and $\nu_{\lambda} \in \operatorname{im} p$ (i.e., ν_{λ} can be realized by some market state). Thus, for $\lambda \in \Lambda$, C^* and hence D are subdifferentiable at ν_{λ} , and by first-order optimality,

$$(\partial C^*(\nu_{\lambda}) - q) \cap A^{\perp} = (\partial C^*(\nu_{\lambda}) - q) \cap A_{\lambda}^{\perp} = \partial_2 D(q, \nu_{\lambda}) \cap A_{\lambda}^{\perp} \neq \emptyset$$

where ∂ denotes the subgradient and $(\cdot)^{\perp}$ the orthogonal complement.

A *Bregman perpendicular* to A through q is a map $\gamma : \lambda \mapsto (\nu_{\lambda}, q_{\lambda})$ defined over $\lambda \in \Lambda$, where q_{λ} is some point from $\partial C^*(\nu_{\lambda}) \cap (q + A^{\perp})$; for $\lambda = 1$ we specifically choose $q_{\lambda} = q$. We stress that in our definition points ν_{λ} are uniquely determined, but q_{λ} are not necessarily unique (except for q at $\lambda = 1$). For any pair of points $q_{\lambda}, q_{\lambda'}$ and any $a, a' \in A$, we have $(q_{\lambda'} - q_{\lambda}) \cdot (a' - a) = 0$, hence the name perpendicular.

Note that γ has a sense of direction as determined by λ , increasing from A towards q . We will say that q_{λ} precedes $q_{\lambda'}$ if $\lambda < \lambda'$ (and similarly for ν_{λ} and $\nu_{\lambda'}$). We will use the shorthand $[\gamma]_{\nu}$ to refer to the map $\lambda \mapsto \nu_{\lambda}$ and the shorthand $[\gamma]_q$ to refer to the map $\lambda \mapsto q_{\lambda}$.

TODO-MD: need to show the continuity of $[\gamma]_{\nu}$. This is true over $\operatorname{ri} \operatorname{dom} C^*$. Can we show it more broadly over $\operatorname{dom} \partial C^*$?

TODO: work out perpendiculars for LMSR and quadratic.

4.3. Main result

To state our main result, we need to introduce a few definitions and make some regularity assumptions.

Let $\operatorname{aff}(X)$ denote the *affine hull* of the set X (i.e., the smallest affine set including X) and $\operatorname{ri} X$ the *relative interior* of X (i.e., interior relative to the affine hull). Since Ω is finite, the realizable set $\mathcal{M} = \operatorname{conv} \Omega$ is a polytope. Its boundary can therefore be decomposed into *faces*. More precisely, we will say that $X \subseteq \Omega$, $X \neq \emptyset$, forms a *face* of \mathcal{M} if there exists a hyperplane H intersecting \mathcal{M} in $\operatorname{conv} X$, such that all of $\mathcal{M} \setminus \operatorname{conv} X$ lies on one side of H .

With the definitions at hand, we make the following regularity assumptions on the cost function, initial state, belief, and the payoff function:

ASSUMPTION 1 (INFORMATIVE PRICES). $D(q, \nu)$ is only a function of $p(q)$ and ν .

ASSUMPTION 2 (REALIZABLE INITIALIZATION). $p(q_0) \in \mathcal{M}$.

ASSUMPTION 3 (NON-DEGENERATE BELIEF). $\mu \in \operatorname{ri} \mathcal{M}$.

ASSUMPTION 4 (ACUTE ANGLES). *For any face $X \subseteq \Omega$ and any q such that $p(q) \in \mathcal{M}$, let γ be a perpendicular to $\text{aff}(X)$ passing through q . Let (ν_λ, q_λ) and $(\nu_{\lambda'}, q_{\lambda'})$ be points on γ such that $\lambda' > \lambda$ and $\nu_\lambda, \nu_{\lambda'} \in \mathcal{M}$. Then for any $\omega \in \Omega$ and $x \in X$:*

$$(q_{\lambda'} - q_\lambda) \cdot (\omega - x) \geq 0 \quad .$$

The assumptions of realizable initialization and non-degenerate belief are introduced to simplify our analysis and exclude degenerate special cases. The informative prices assumption is more restrictive. However, it is satisfied by many market scoring rules, including the quadratic cost and LMSR. (TODO: we should list a wider class of costs that satisfy this assumption. For instance, costs derived from the traditional Bregman divergences. Also, can we either remove this assumption or show that it is necessary?) Finally, acute angles seem to be a necessary condition and we will show that without acute angles, budget additivity need not hold. (TODO: prove characterization when the acute angles hold for quadratic and work out counterexamples showing that the assumption is necessary. Can we also show something stronger, e.g., that the budget additivity holds if and only if acute angles hold?)

Recall that $\hat{Q}(B; q_0)$ denotes the set of solution to Eq. (4). If $Q_0 \subseteq \mathbb{R}^n$, we will also use the shorthand $\hat{Q}(B; Q_0) := \bigcup_{q \in Q_0} \hat{Q}(B; q)$. Taking advantage of the informative prices assumption, we also introduce the notation $\hat{\nu}(B; \nu_0)$ for the unique prices achieved by solving Eq. (4) starting from any $q_0 \in p^{-1}(\nu_0)$. We have that $\hat{Q}(B; q_0) = p^{-1}(\hat{\nu}(B; p(q_0)))$. We are ready to state our main result.

THEOREM 4.2 (BUDGET ADDITIVITY). *Assuming informative prices, realizable initialization, non-degenerate belief and acute angles, for all $B, B' \geq 0$,*

$$\hat{Q}(B + B'; q_0) = \hat{Q}(B'; \hat{Q}(B; q_0)) \quad ,$$

or, equivalently,

$$\hat{\nu}(B + B'; \nu_0) = \hat{\nu}(B'; \hat{\nu}(B; \nu_0)) \quad ,$$

where $\nu_0 = p(q_0)$.

PROOF. If $\nu_0 = \mu$ then the statement trivially holds, so in the remainder we assume $\nu_0 \neq \mu$. The plan is to exhibit a sequence of complementary pairs $(\nu_0, q_0), (\nu_1, q_1), \dots, (\nu_k, q_k)$, $\nu_k = \mu$, connected by directed curves called segments $\ell_0, \dots, \ell_{k-1}$. Our intention will be to show that the budget associativity holds along the union of these segments $L = \bigcup_{i=0}^{k-1} \ell_i$, and that $[L]_\nu$ contains the solutions $\{\hat{\nu}(B; \nu_0) : B \geq 0\}$.

To proceed, we associate each segment with an active set $X_i \subseteq \Omega$ such that $[\ell_i]_\nu \subseteq \text{conv}(X_i, \mu)$, and prove that for any pair of points $q, q' \in [\ell_i]_q$ where q precedes q' , the KKT lemma holds, i.e.,

- (1) $U(q', x; q) = U(q', x'; q)$ for all $x, x' \in X_i$
- (2) $U(q', \omega; q) \geq U(q', x; q)$ for all $x \in X_i$

We refer to this statement as the Segment Lemma (proved in the next section).

The key point of our construction will be the monotonicity of the sequence X_i , such that $\Omega \supseteq X_0 \supseteq X_1 \supseteq \dots \supseteq X_{k-1} \neq \emptyset$. The proof of the theorem then follows by using path independence of the utility function with the Segment Lemma. We first show that for any $(q, \nu) \in \ell_i$ and $(q', \nu') \in \ell_j$ where $i < j$, we have $q' \in \hat{Q}(B; q)$ for some budget B . By the monotonicity of sequence X_0, X_1, \dots, X_{k-1} , and the Segment Lemma we can

derive that for all $x \in X_j$ and $\omega \in \Omega$

$$\begin{aligned} U(q', \omega; q) &= U(q', \omega; q_j) + U(q_j, \omega; q_{j-1}) \cdots + U(q_{i+1}, \omega; q) \\ &\geq U(q', x; q_j) + U(q_j, x; q_{j-1}) \cdots + U(q_{i+1}, x; q) \\ &= U(q', x; q) . \end{aligned}$$

And analogously, for all $x, x' \in X_j$,

$$U(q', x; q) = U(q', x'; q) .$$

Thus, by the KKT lemma, $q' \in \hat{Q}(B; q)$, or equivalently

$$\nu' = \hat{\nu}(B; \nu) \quad (5)$$

where $B = -U(q', x; q)$ for $x \in X_j$ (all $x \in X_j$ yield the same value B).

To finish the proof of the theorem, note that $[L]_\nu$ is connected and contains $\nu_k = \mu$, so Eq. (5) implies that $[L]_\nu$ consists exactly of the points $\hat{\nu}(B; q_0)$ for $B \geq 0$. Moreover, if $\nu = \hat{\nu}(B; \nu_0) \in [\ell_j]_\nu$ and $\nu' = \hat{\nu}(B + B'; \nu_0) \in [\ell_{j'}]_\nu$ then some pairs of the form (ν, q) and (ν', q') must lie on L , and for any $x \in X_{j'} \subseteq X_j \subseteq X_0$,

$$U(q', x; q) = U(q', x; q_0) - U(q, x; q_0) = -(B + B') + B = -B' .$$

Hence $\nu' = \hat{\nu}(B'; \nu)$ and the theorem follows, provided that we can show the segment lemma. \square

4.4. Segment Lemma

Consider the initial state q_0 and its price ν_0 . As in the proof of the theorem, assume that $\nu_0 \neq \mu$. Let X be the smallest face such that $\nu_0 \in \text{conv}(X, \mu)$. Let γ be a perpendicular to $\text{aff}(X)$ going through q_0 . The curve $[\gamma]_\nu$ passes through ν_0 and eventually reaches the boundary of $\text{conv}(X, \mu)$ at some point ν_1 . The portion of γ going from ν_0 to ν_1 is our first segment ℓ . It has the same sense of direction as γ .

LEMMA 4.3 (SEGMENT LEMMA). *Let $q, q' \in [\ell]_q$ such that q precedes q' . Then:*

- (1) $U(q', x; q) = U(q', x'; q)$ for all $x, x' \in X$
- (2) $U(q', \omega; q) \geq U(q', x; q)$ for all $x \in X$ and $\omega \in \Omega$

PROOF. For the first part, note that since $\ell \subseteq \gamma$ and γ is perpendicular to $\text{aff}(X)$, we have

$$U(q', x'; q) - U(q', x; q) = (q' - q) \cdot (x' - x) = 0 .$$

For the second part, we directly appeal to the acute angles assumption:

$$U(q', \omega; q) - U(q', x; q) = (q' - q) \cdot (\omega - x) \geq 0 .$$

\square

The above construction gives us the first segment ℓ_0 and the corresponding set X_0 . The segment terminates at the point (ν_1, q_1) . There are two possibilities:

- (1) $\nu_1 = \mu$; in this case we are done;
- (2) ν_1 lies on a lower-dimensional face of $\text{conv}(X_0, \mu)$; in this case, we can use the above construction again, starting with q_1 , and obtaining a new set $X_1 \subseteq X_0$ and a new segment ℓ_1 ; and iterate

Note that the above process eventually ends, because with each iteration, the dimension of X decreases.

5. PREDICTION MARKETS WITH SET THEORETIC BELIEFS

The usual assumption in the study of prediction markets is that the agent knows (or has a belief about) the expectation of the security. In this section, we relax this assumption to allow for set theoretic beliefs, by which we mean that all that the agent knows is that the expectation lies in some subset of \mathbb{R}^n , called the *belief set*. For example, suppose that we have two securities for the indicator variables for two events A and B in the probability space. It is possible that an agent knows that event A is more likely to happen than event B , without knowing the exact joint probability distribution of A and B . This can be represented by the set $\mu_A \geq \mu_B$ intersected with the subset of all possible pairs of probabilities, which is $[0, 1]^2$.

We consider the incentives for such an agent to trade in the prediction market. Our model of the agent is that he will trade in the market if he is guaranteed to profit no matter what the true expectation is, as long as it is in his belief set. Taking this further, the agent's prediction will be the one that maximizes the worst-case payoff, where the worst case is taken over all expectations in his belief set. Given a belief set $\mathbb{B} \subseteq \mathcal{M}$ and an initial state q , the agent aims to maximize

$$\max_{q' \in \mathbb{R}^n} \min_{\mu \in \mathbb{B}} U(q', \mu; q) . \quad (6)$$

As with point beliefs, the maximum may only be achievable in the limit, but it will always coincide with a unique target price, as the next theorem shows. (TODO?: can we also prove the converse? i.e., any sequence of actions $\{q_k\}_{k=1}^\infty$ with $p(q_k) \rightarrow \hat{\mu}$ maximizes the expected utility?)

THEOREM 5.1. *Suppose that an agent has a non-empty belief set $\mathbb{B} \subseteq \mathcal{M}$ that is closed and convex. Then for all q , any sequence of actions maximizing Eq. (6) moves prices to the Bregman projection of q onto \mathbb{B} , which is*

$$\hat{\mu} := \operatorname{argmin}_{\mu \in \mathbb{B}} D(q, \mu) .$$

PROOF. Rewrite Eq. (6) as follows

$$\max_{q' \in \mathbb{R}^n} \min_{\mu \in \mathbb{B}} U(q', \mu; q) = \min_{\mu \in \mathbb{B}} \max_{q' \in \mathbb{R}^n} U(q', \mu; q) \quad (7)$$

$$\begin{aligned} &= \min_{\mu \in \mathbb{B}} \max_{q' \in \mathbb{R}^n} \{q' \cdot \mu - C(q') - q \cdot \mu + C(q)\} \\ &= \min_{\mu \in \mathbb{B}} \{C^*(\mu) - q \cdot \mu + C(q)\} \end{aligned} \quad (8)$$

$$= \min_{\mu \in \mathbb{B}} D(q, \mu) = D(q, \hat{\mu}) . \quad (9)$$

In Eq. (7) we are able to switch max and min by strong duality, because utility is linear in μ , concave in q' , and \mathbb{B} is compact and convex. (TODO-MD: need a reference here, e.g., we could use some Fenchel duality result from Rockafellar). In Eq. (8) we simply use the definition of C^* . To complete the proof, let $\{q_k\}_{k=1}^\infty$ be a maximizing sequence in Eq. (6) and $\mu_k \in \mathbb{B}$ the corresponding worst-case beliefs, i.e., $\mu_k := \operatorname{argmin}_{\mu \in \mathbb{B}} U(q_k, \mu; q)$. We need to show that $p(q_k) \rightarrow \hat{\mu}$.

From Eqs. (7)–(9), we know that $U(q_k, \mu_k; q) \rightarrow D(q, \hat{\mu})$. Also, from the optimality of μ_k relative to q_k , and the optimality of $\hat{\mu}$ on the right-hand side of Eq. (7), we obtain

$$U(q_k, \mu_k; q) \leq U(q_k, \hat{\mu}; q) \leq D(q, \hat{\mu}) .$$

Thus, $U(q_k, \hat{\mu}; q) \rightarrow D(q, \hat{\mu})$. By Eq. (1) and the non-negativity of D , we then must have

$$D(q_k, \hat{\mu}) \rightarrow 0 .$$

This means that $p(q_k) \rightarrow \hat{\mu}$. (TODO-MD: The last step needs a reference. It should be standard for Bregman divergences, but our definition is a bit more general than the most. We should be able to prove this by writing $D(q_k, \hat{\mu}) = C^*(\hat{\mu}) - C^*(p(q_k)) - (\hat{\mu} - p(q_k)) \cdot q_k$, and using strict convexity of C^* : Proceed by contradiction. If $p(q_k) \not\rightarrow \hat{\mu}$, then one of the orthants with the origin at $\hat{\mu}$ must contain infinitely many $p(q_k)$ separated from $\hat{\mu}$ by a hyperplane H that cuts across all rays in the orthant. The “Taylor series” lower bound on C^* implied by D gets worse along the rays in the orthant, so it suffices to consider the lower bounds of points in the intersection of the orthant and the hyperplane H . This intersection, call it K , is compact. For each $p(q_k)$ in the orthant, pick the corresponding point ν_k in K and calculate the Taylor-series gap. Note that the gap of ν_k converges to zero, because the gap for $p(q_k)$ converges to zero. Let ν be a limit point of the sequence (exists by compactness). Consider $C^*(\hat{\mu}/2 + \nu/2)$. Since the Taylor series gap converges to zero, we must have $C^*(\hat{\mu}/2 + \nu/2) \geq C^*(\hat{\mu})/2 + C^*(\nu)/2$, contradicting strict convexity of C^* .) \square

[Lambert and Shoham 2009] considered the elicitation problem where the goal of the market maker is to elicit a weaker information structure instead of an exact probability distribution. This is modeled by considering a partition of the space of all probability distributions and the goal is to elicit which set in the partition contains the actual probability distribution. [Lambert and Shoham 2009] then considered scoring rules where the agent is asked to report a set in the partition and is compensated based on his report and the eventual outcome. The motivation for considering such a scoring rule instead of simply eliciting the exact probability distribution and inferring the answer from that, is that the dimension of the complete market needed to elicit the entire probability distribution may be prohibitively huge. Reporting a set in the partition would be much easier for the agent when the number of sets in the partition is much smaller than the number of possible outcomes. The conclusion of [Lambert and Shoham 2009] is that strictly proper scoring rules exist if and only if the partition forms a (weighted) Voronoi diagram.

We note that quite often for interesting partitions one does not need to run a complete market to elicit the information about which partition contains the belief of the agent. There exist securities, with a dimension smaller than the entire outcome space, that are sufficient to elicit the information needed. For example, suppose that the outcome space consists of n binary events, i.e., each outcome corresponds to a binary string of size n with the i^{th} bit indicating whether the i^{th} event occurred or not. If the goal is to elicit the most likely event then it is sufficient to have n securities, one for each event, instead of 2^n securities, one for each outcome. In spite of this, there is another reason why eliciting the probabilities would be infeasible, that the agent himself might not know the exact probabilities. In the above example, an agent might only know which is the most likely event, or he might know that the most likely event is one out of two or that event i is more likely than event j . The set theoretic belief model shows that it is still possible to aggregate information from such agents.

The advantage of the set theoretic belief model over the model in [Lambert and Shoham 2009] is that this allows different agents to have different information structures while the scoring rule remains the same. It works both ways, if an agent has a weaker information than what is needed then he can still contribute, while if an agent has a stronger information than what is needed then he has the opportunity to give a stronger prediction. An example of the latter is as follows, consider once again the example of n binary events earlier where the goal is to elicit the most likely event. If an agent knows not only which is the most likely event but also that the probability of this event is at least 10% more than the next highest probability, then that would be reflected in his prediction. Even though we are not interested in this margin, it leads

to a greater confidence in our prediction of the most likely event. Also we only need the sets in the partition to be convex, which is a weaker requirement than them being Voronoi diagrams as needed by [Lambert and Shoham 2009].

6. CONCLUSIONS

In this paper, we studied two different scenarios in prediction markets where an agent cannot simply move the market price to his current belief: (1) when the agent has a budget constraint that restricts the trades he could make; or (2) when the agent has limited information about the underlying probability distribution, in the sense that he only knows that the probability distribution lies in a convex set.

For the budget constrained case, we presented a geometric characterization of the user's behavior: there is a convex budget set within which the user can move the market state, and the optimal strategy of the user is a Bregman projection of his belief onto the budget set. Further, we extended this characterization to show that complete prediction markets have *budget associativity*, i.e., a sequence of agents with the same belief is equivalent to a single agent with the same belief and with budget equal to the sum of the budgets of all the agents.

For the set theoretic belief case, we assume that the user aims to maximize the worst case profit where the profit is taken over all distributions in his belief set. Again, we characterized the user's behavior geometrically by showing the user's optimal strategy is a Bregman projection of the current price vector onto the belief set.

REFERENCES

- ABERNETHY, J., CHEN, Y., AND VAUGHAN, J. W. 2011. An optimization-based framework for automated market-making. In *ACM Conference on Electronic Commerce*. 297–306.
- BEYGEZIMER, A., LANGFORD, J., AND PENNOCK, D. 2012. Learning performance of prediction markets with kelly bettors. *CoRR abs/1201.6655*.
- BHATTACHARJEE, R. AND GOEL, A. 2006. Incentive based ranking mechanisms. In *EC Workshop, Economics of Networked Systems*.
- BHATTACHARJEE, R. AND GOEL, A. 2007. Algorithms and incentives for robust ranking. In *SODA*. 425–433.
- BRIER, G. 1950. Verification of forecasts expressed in terms of probability. *Monthly Weather Review* 78, 13.
- CHEN, Y. AND PENNOCK, D. M. 2007. A utility framework for bounded-loss market makers. In *UAI*. 49–56.
- CHEN, Y. AND VAUGHAN, J. W. 2010. A new understanding of prediction markets via no-regret learning. In *ACM Conference on Electronic Commerce*. 189–198.
- EISENBERG, E. AND GALE, D. 1959. Consensus of subjective probabilities: The parimutuel method. *Annals of Mathematical Statistics* 30, 165–168.
- FORTNOW, L. AND SAMI, R. 2012. Multi-outcome and multidimensional market scoring rules. *CoRR abs/1202.1712*.
- GNEITING, T. AND RAFTERY, A. E. 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102, 477, 359–378.
- GOEL, S., REEVES, D. M., WATTS, D. J., AND PENNOCK, D. M. 2010. Prediction without markets. In *Proceedings of the 11th ACM conference on Electronic commerce*. EC '10. ACM, New York, NY, USA, 357–366.
- GRUNWALD, P. D. AND DAWID, A. P. 2004. Game theory, maximum entropy, minimum discrepancy and robust bayesian decision theory. *Ann. Stat.* 32, math.ST/0410076. IMS-AOS-AOS-231. 4, 1367–1433.

- HANSON, R. 2003. Combinatorial information market design. *Information Systems Frontiers* 5, 1, 107–119.
- HANSON, R. 2007. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets* 1, 1, 3–15.
- LAMBERT, N. S. AND SHOHAM, Y. 2009. Eliciting truthful answers to multiple-choice questions. In *ACM Conference on Electronic Commerce*. 109–118.
- MANSKI, C. F. June 2006. Interpreting the predictions of prediction markets. *Economics Letters* 91, 3, 425–429.
- WOLFERS, J. AND ZITZEWITZ, E.
- WOLFERS, J. AND ZITZEWITZ, E. 2006. Interpreting prediction market prices as probabilities. IZA Discussion Papers 2092, Institute for the Study of Labor (IZA). Apr.