

Information Aggregation in Exponential Family Markets

JACOB ABERNETHY, University of Michigan, Ann Arbor
SINDHU KUTTY, University of Michigan, Ann Arbor
SÉBASTIEN LAHAIE, Microsoft Research, New York City
RAHUL SAMI, University of Michigan, Ann Arbor

We consider the design of prediction market mechanisms known as automated market makers. We show that we can design these mechanisms via the mold of *exponential family distributions*, a popular and well-studied probability distribution template used in statistics. We give a full development of this relationship and explore a range of benefits. We draw connections between the information aggregation of market prices and the belief aggregation of learning agents that rely on exponential family distributions. We develop a natural analysis of the market behavior as well as the price equilibrium under the assumption that the traders exhibit risk aversion according to exponential utility. We also consider similar aspects under alternative models, such as budget-constrained traders.

Categories and Subject Descriptors: J.4 [Social and Behavioral Sciences]: Economics; I.2.6 [Artificial Intelligence]: Learning

Additional Key Words and Phrases: logarithmic score; exponential family; maximum entropy; risk aversion; budget constraints

1. INTRODUCTION

Prediction markets are aggregation mechanisms that allow market prices to be interpreted as predictive probabilities on an event. Each trader in the market is assumed to have some private information that he uses to make a prediction on the outcome of the event. Traders are allowed to report their beliefs by buying and selling securities whose ultimate payoff depends on the future outcome. This will affect the state of the market, thus updating the predictive probabilities for the event. Further, since the trades are done sequentially, the trader may observe all past trades in the market and update his private beliefs based on this information. In this sense the market prices, which are in effect the prices at which the marginal trader is willing to buy or sell the available securities, can be interpreted as an aggregate “consensus probability forecast” of the event in question.

Much of the work on prediction market design has focused heavily on structural properties of the mechanism: incentive compatibility, the market maker loss, the available liquidity, the fluctuations of the prices as a function of the trading volume, to name a few. Absent from much of the literature is a corresponding *semantics* of the market behavior or the observed prices. That is, how can we interpret the equilibrium market state when we have a number of traders with diverse beliefs on the state of the world?

Rahul Sami’s current affiliation: Google, Inc. This work was done while the author was at the University of Michigan. Author’s addresses: J. Abernethy and S. Kutty, Department of Electrical Engineering and Computer Science, University of Michigan, {skutty,jabernet}@umich.edu. S. Lahaie, Microsoft Research, New York City, slahaie@microsoft.com. Rahul Sami, Google Inc., rahulsami@gmail.com.

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.

EC’14, June 8–12, 2014, Stanford, CA, USA.

ACM 978-1-4503-2565-3/14/06 ...\$15.00.

Copyright is held by the owner/author(s). Publication rights licensed to ACM.

<http://dx.doi.org/10.1145/2600057.2602896>

How is the market an aggregation mechanism? Do price changes relate to our usual Bayesian notion of information incorporation via posterior updating?

In the present work we show that a number of classical statistical tools can be leveraged to design a prediction market framework in the mold of *exponential family distributions*; we show that this statistical framework leads to a number of attractive properties and interpretations. Common concepts in statistics—including *entropy maximization*, *log loss*, and *Bayesian inference*—relate to natural aspects of our class of mechanisms. In particular, the central objects in our market framework can be interpreted via concepts used to define exponential families:

- the market’s *payoff function* corresponds to the *sufficient statistics* of the distribution;
- the vector of *outstanding shares* in the market corresponds to the *natural parameter* of the distribution;
- the market prices correspond to *mean parameters*;
- the market’s cost function corresponds to the distribution’s *log-partition* function.

We begin in Section 2 with a discussion of scoring rules based on exponential family distributions, and we show how the framework leads to a variety of scoring rules for continuous outcome spaces. We turn our attention to market design in Section 3 and give a full description of our proposed mechanisms. In addition to showing the syntactic relationship between exponential families and prediction markets, we explore a number of rich semantic implications as well. In particular, we show that our formulation allows us to analyze the evolution of the market under various trader models:

- Trader behavior varies depending on how they assimilate information; for example, should we consider our agents as Bayesians or frequentists. In Section 4 we consider traders that use a conjugate prior to update their beliefs, and we study how their trades would affect the market state.
- In Section 5 we consider *risk-averse* agents that optimize their bets according to exponential utility. In this case we can characterize precisely how a single trader interacts with the market, as well as the equilibrium reached given multiple traders; the latter result is achieved via a potential game argument. The eventual market state is a weighted combination of traders’ beliefs and the initial state; the weights are proportional to risk aversion parameters.
- In Section 6 we consider *budget-limited traders* who are constrained in how much they influence the market. We analyze the market under these circumstances; we are able to show that traders with good information can expect to profit and their influence over the market state increases over time whereas malicious traders have limited impact on the market.

A note about play-money markets: Even in prediction markets set up with play-money, risk aversion is plausible since it captures the fact that a trader high on the leaderboard may not want to risk her standing with large late trades. Further, in such markets especially, budget constraints are key since they help alleviate the problem that traders can easily manipulate prices with unreasonable trades. Our results are therefore particularly relevant in markets with play-money.

Related Work. The notion of an exponential family distribution is fundamental to this paper. For comprehensive introductions to these distributions, see [Barndorff-Nielsen 1978; Wainwright and Jordan 2008]. Exponential families are intimately tied to the notions of log loss and entropy, but can be generalized to other types of convex losses and information as shown by Grünwald and Dawid [2004], who also make a connection to scoring rules.

Scoring rules are a measure of prediction accuracy, and we are concerned here with scoring rules for statistic expectations, typically over infinite outcome spaces. Such

rules have been characterized by Savage [1971]; see also [Abernethy and Frongillo 2012; Gneiting and Raftery 2007; Lambert et al. 2008]. Our rules are of course special cases of this characterization, but it appears the range of elegant scoring rules that arise from exponential families has not been appreciated. Indeed, Gneiting and Raftery [2007] observe that specific instances of scoring rules for continuous outcomes are lacking, and survey various possibilities.

In a seminal paper, Hanson [2003] showed how to form a prediction market based on a sequentially-shared scoring rule, and specifically proposed the logarithmic market scoring rule (LMSR) based on log loss for finite outcome spaces [Hanson 2007]. The markets we introduce are direct generalizations of the LMSR to continuous outcomes, but take the form of cost-function based markets as introduced by Chen and Pennock [2007]. Gao et al. [2009] and Chen et al. [2013] also consider extending various market makers to infinite outcome spaces.

Prediction markets are known to perform well in practice [Pennock and Sami 2007; Pennock et al. 2001]. However, a sound theory for interpreting trader behavior and market prices is an ongoing area of study [Wolfers and Zitzewitz 2006]. At one extreme, agents are assumed myopic and risk-neutral, implying they move the market state to their belief [Chen and Vaughan 2010]. At the other extreme, agents are strategic and the market fully incorporates all information [Ostrovsky 2012]. The aggregation properties of market mechanisms have also been explored from a machine learning perspective in [Storkey 2011]. More closely related to our work, Frongillo and Reid [2013] also touch upon the statistical interpretations of prediction market costs and payoffs from the viewpoint of exponential families.

One of the aims of the present work is to understand the behavior of cost-function prediction markets with risk-averse agents, a subject which has received only limited attention; the only work we are aware of is [Frongillo et al. 2012, Sec. 3.1]. However, risk aversion is a fundamental component of mathematical finance and portfolio optimization, and there are close connections between the notion of a cost function and that of a convex risk measure [Föllmer and Schied 2002; Föllmer and Knispel 2011]. Indeed, they arise from the same axioms as noted by Othman and Sandholm [2011]. We see the potential to draw more on the mathematical finance literature to take into account risk aversion, as prediction markets can be viewed simply as single-period financial markets [Föllmer and Schied 2004, Part I].

2. GENERALIZED LOG SCORING RULES

We consider a measurable space consisting of a set of (mutually exclusive, exhaustive) outcomes \mathcal{X} together with a σ -algebra \mathcal{F} . An agent or expert has a *belief* over potential outcomes taking the form of a probability measure absolutely continuous with respect to a base measure ν .¹ Throughout we represent the belief as the corresponding density p with respect to ν . Let \mathcal{P} denote the set of all such probability densities.

We are interested in eliciting information about the agent's belief, in particular expectation information. Let $\phi : \mathcal{X} \rightarrow \mathbf{R}^d$ be a vector-valued random variable or *statistic*, where d is finite. The aim is to elicit $\mu = \mathbf{E}_p[\phi(x)]$ where x is the random outcome. A *scoring rule* is a device for this purpose. Let

$$\mathcal{M} = \{ \mu \in \mathbf{R}^d : \mathbf{E}_p[\phi(x)] = \mu, \text{ for some } p \in \mathcal{P} \}$$

be the set of realizable statistic expectations. Unless otherwise mentioned, all integrals including expectations are taken with respect to base measure ν , which is therefore

¹Recall that a measure P is absolutely continuous with respect to ν if $P(A) = 0$ for every $A \in \mathcal{F}$ for which $\nu(A) = 0$. In essence the base measure ν restricts the support of P . In our examples ν will typically be a restriction of the Lebesgue measure for continuous outcomes or the counting measure for discrete outcomes.

implicit in many of the formulas. A scoring rule $S : \mathcal{M} \times \mathcal{X} \rightarrow \mathbf{R} \cup \{-\infty\}$ pays the agent $S(\hat{\mu}, x)$ according to how well its report $\hat{\mu} \in \mathcal{M}$ agrees with the eventual outcome $x \in \mathcal{X}$. The following definition is due to Lambert et al. [2008].

Definition 2.1. A scoring rule S is *proper* for statistic ϕ if for each $\mu \in \mathcal{M}$ and $p \in \mathcal{P}$ with expected statistic μ , we have for all $\hat{\mu} \neq \mu$

$$\mathbf{E}_p[S(\mu, x)] \geq \mathbf{E}_p[S(\hat{\mu}, x)]. \quad (1)$$

Given a proper scoring rule S any affine transformation $\tilde{S}(\mu, x) = aS(\mu, x) + b(x)$ of the rule, with $a > 0$ and b an arbitrary real-valued function of the outcomes, again yields a proper scoring rule termed *equivalent* [Dawid 1998; Gneiting and Raftery 2007]. Throughout we will implicitly apply such affine transformations to obtain the clearest version of the scoring rule. We will also focus on scoring rules where inequality (1) is strict to avoid trivial cases such as constant scoring rules.

Classically, scoring rules take in the entire density p rather than just some statistic, and incentive compatibility must hold over all of \mathcal{P} . When the outcome space is large or infinite, it is not feasible to directly communicate p , so the definition allows for summary information of the belief.

Note that Definition 2.1 places only mild information requirements on the part of the agent to ensure truthful reporting. Because condition (1) holds for all p consistent with expectation μ , it is enough for the agent to simply know the latter and not the complete density to be properly incentivized. However, the agent must also agree with the support of the density as implicitly defined by base measure ν .

2.1. Proper Scoring from Maximum Entropy

Our starting point for designing proper scoring rules is the classic logarithmic scoring rule for eliciting probabilities in the case of finite outcomes. This rule is simply $S(p, x) = \log p(x)$, namely we take the log likelihood of the reported density at the eventual outcome. To generalize the rule to expected statistics rather than full densities, we consider a subset of densities $\mathcal{D} \subseteq \mathcal{P}$. If \mathcal{D} and \mathcal{M} are bijective, then we say that \mathcal{M} parametrizes \mathcal{D} and write $p(\cdot; \mu)$ for the density mapping to μ . Given such a family parametrized by the relevant statistics, the generalized log scoring rule is then

$$S(\mu, x) = \log p(x; \mu). \quad (2)$$

Even though the log score is only applied to densities from \mathcal{D} , according to Definition 2.1 it must work over all densities in \mathcal{P} . It turns out this is possible if \mathcal{D} is chosen appropriately, drawing on a well-known duality between maximum likelihood and maximum entropy [Grünwald and Dawid 2004].

Exponential Families. We let $p(x; \mu)$ be the maximum entropy distribution with expected statistic μ . Specifically, it is the solution to the following program:²

$$\min_{p \in \mathcal{P}} F(p) \quad \text{s.t.} \quad \mathbf{E}_p[\phi(x)] = \mu, \quad (3)$$

where the objective function is the negative entropy of the distribution, namely

$$F(p) = \int_{x \in \mathcal{X}} p(x) \log p(x) d\nu(x).$$

²We assume that the minimum is finite and achieved for all $\mu \in \mathcal{M}$. Some care is needed to ensure this holds for specific statistics and outcome spaces. For example, taking outcomes to be the real numbers, there is no maximum entropy distribution with a given mean μ (one can take densities tending towards the uniform distribution over the reals), but there is always a solution if we constrain both the mean and variance.

Note that the explicit set of constraints in (3) are linear, whereas the objective is convex. We let $G : \mathcal{M} \rightarrow \mathbf{R}$ be the optimal value function of (3), meaning $G(\mu)$ is the negative entropy of the maximum entropy distribution with expected statistics μ .

It is well-known that solutions to (3) are *exponential family* distributions; see, for instance [Wainwright and Jordan 2008]. These are distributions whose densities with respect to ν take the form

$$p(x; \theta) = \exp(\langle \theta, \phi(x) \rangle - T(\theta)). \quad (4)$$

The density is stated here in terms of its *natural* parametrization $\theta \in \mathbf{R}^d$, where θ arises as the Lagrange multiplier associated with the linear constraints in (3). The term $T(\theta)$ essentially arises as the multiplier for the normalization constraint (the density must integrate to 1), and so ensures that (4) is normalized:

$$T(\theta) = \log \int_{\mathcal{X}} \exp(\langle \theta, \phi(x) \rangle) d\nu(x). \quad (5)$$

The function T is known as the *log-partition* or *cumulant* function corresponding to the exponential family. Its domain is $\Theta = \{\theta \in \mathbf{R}^d : T(\theta) < +\infty\}$, called the natural parameter space, and we assume throughout that it is nonempty.³ The exponential family is *regular* if Θ is open—almost all exponential families of interest, and all those we consider in this work, are regular. The family is *minimal* if there is no $\alpha \in \Theta$ such that $\langle \alpha, \phi(x) \rangle$ is a constant over \mathcal{X} (ν -almost everywhere); minimality is a property of the associated statistic ϕ , usually called the *sufficient statistic* in the literature.

Some commonly studied distributions that are in fact exponential families include the binomial, beta, Poisson, exponential, and normal distributions; the sufficient statistics and parametrizations for the exponential and normal will be covered in Section 2.2. The following proposition collects the relevant results on regular exponential families; proofs may be found in Wainwright and Jordan [2008, Prop. 3.1–3.2, Thm. 3.3–3.4] and see also Banerjee et al. [2005a, Lem. 1, Thm. 2]. A convex function T is of *Legendre type* if it is proper, closed, strictly convex and differentiable on the interior of its domain, and $\lim_{\theta \rightarrow \bar{\theta}} \|\nabla T(\theta)\| = +\infty$ when $\bar{\theta}$ lies on the boundary of the domain.

PROPOSITION 2.1. *Consider a regular exponential family with minimal sufficient statistic. The following properties hold:*

- (1) *T and G are of Legendre type, and $T = G^*$ (equivalently $G = T^*$).*
- (2) *The gradient map ∇T is one-to-one and onto the interior of \mathcal{M} . Its inverse is ∇G which is one-to-one and onto the interior of Θ .*
- (3) *The exponential family distribution with natural parameter $\theta \in \Theta$ has expected statistic $\mu = \mathbf{E}_p[\phi(x)] = \nabla T(\theta)$.*
- (4) *The maximum entropy distribution with expected statistic μ is the exponential family distribution with natural parameter $\theta = \nabla G(\mu)$.*

In the above T^* denotes the convex conjugate of T , which here can be evaluated as $T^*(\mu) = \sup_{\theta \in \Theta} \langle \theta, \mu \rangle - T(\theta)$. Similarly, $G^*(\theta) = \sup_{\mu \in \mathcal{M}} \langle \theta, \mu \rangle - G(\mu)$.

Proper Log Scoring. We are now in a position to analyze the log scoring rule under exponential family distributions. From our discussion so far, we have that an exponential family density can be parametrized either by the natural parameter θ , or by the mean parameter μ , and that the two are related by the invertible gradient map $\mu = \nabla T(\theta)$. We will write $p(x; \theta)$ or $p(x; \mu)$ given the parametrization used.

³This may impose additional conditions on the choice of base measure ν in certain special cases. For instance, if the density over \mathbf{R} is known to have unit variance and statistic $\phi(x) = x$, then one uses the base measure $\nu(x) = e^{-x^2/2}$ leading to a normal distribution with fixed variance [Barndorff-Nielsen 1978].

The following observation is crucial. Let $\tilde{p} \in \mathcal{P}$ be any arbitrary density with expected statistic μ , let $p(\cdot; \mu)$ be the exponential family with the same expected statistic, and let $\hat{\mu} \in \mathcal{M}$ be an alternative report. Then from (4) we have

$$\mathbf{E}_{\tilde{p}}[\log p(x; \hat{\mu})] = \mathbf{E}_{p(\cdot; \mu)}[\log p(x; \hat{\mu})] = \langle \hat{\theta}, \mu \rangle - T(\hat{\theta}), \quad (6)$$

where $\hat{\theta} = \nabla G(\hat{\mu})$ is the natural parameter for the exponential family with statistic $\hat{\mu}$. Note that the expected log score does not depend on the entire maximum entropy distribution but just its expected statistic μ . This is how we can achieve proper scoring according to Definition 2.1.

THEOREM 2.2. *Consider the logarithmic scoring rule $S(\mu, x) = \log p(x; \mu)$ defined over a set of densities \mathcal{D} parametrized by \mathcal{M} . The scoring rule is proper if and only if \mathcal{D} is the exponential family with statistic ϕ , where $\mu = \mathbf{E}_{p(x; \mu)}[\phi(x)]$.*

PROOF. Let $\mu, \hat{\mu} \in \mathcal{M}$ be the agent's true belief and an alternative report, and let $p \in \mathcal{P}$ be a density consistent with μ . Let $\theta = \nabla G(\mu)$ and $\hat{\theta} = \nabla G(\hat{\mu})$, and note that $\mu = \nabla T(\theta)$. We have

$$\begin{aligned} \mathbf{E}_p[\log p(x; \mu)] - \mathbf{E}_p[\log p(x; \hat{\mu})] &= \langle \theta, \mu \rangle - T(\theta) - \langle \hat{\theta}, \mu \rangle + T(\hat{\theta}) \\ &= T(\hat{\theta}) - T(\theta) - \langle \hat{\theta} - \theta, \mu \rangle \\ &= T(\hat{\theta}) - T(\theta) - \langle \hat{\theta} - \theta, \nabla T(\theta) \rangle. \end{aligned} \quad (7)$$

The latter is positive by the strict convexity of T , which shows that the log score is proper. For the converse, assume the defined log score is proper. By the Savage characterization of proper scoring rules for expectations (see Abernethy and Frongillo [2012, Thm. 11]), we must have $S(\mu, x) = G(\mu) - \langle \nabla G(\mu), \mu - \phi(x) \rangle$ for some strictly convex function G . Let $T = G^*$, so that $\nabla G = \nabla T^{-1}$, and let $\theta = \nabla G(\mu)$. Then the above can be written as

$$\begin{aligned} \log p(x; \mu) &= G(\mu) - \langle \nabla G(\mu), \mu - \phi(x) \rangle \\ &= \langle \theta, \mu \rangle - T(\theta) - \langle \theta, \mu - \phi(x) \rangle = \langle \theta, \phi(x) \rangle - T(\theta), \end{aligned}$$

which shows that $p(x; \mu)$ takes the form of an exponential family. \square

It has been understood since [McCarthy 1956; Savage 1971] that scoring rules for densities and expectations can be characterized in terms of an underlying convex function. The generalized log scoring rule of Theorem 2.2 obtains by applying the modern characterization of Abernethy and Frongillo [2012, Thm. 11] to the optimal value function G (negative maximum entropy). Banerjee et al. [2005b, Thm. 6] prove a similar result by showing a bijection between exponential families and Bregman divergences.

The notion of a Bregman divergence in fact provides further intuition for the result. Note that (7) is the definition of the Bregman divergence with respect to strictly convex function T , written D_T . Therefore we have

$$\mathbf{E}_p[\log p(x; \mu)] - \mathbf{E}_p[\log p(x; \hat{\mu})] = D_T(\hat{\theta}, \theta) = D_G(\mu, \hat{\mu}),$$

where the last equality is a well-known identity relating the Bregman divergences of T and $T^* = G$. The equation states that the agent's regret from misreporting its mean parameter does not depend on the full density p , only the mean μ .

2.2. Examples: Moments over the Real Line

Theorem 2.2 leads to a straightforward procedure for constructing score rules for expectations. Define the relevant statistic, and consider the maximum entropy (equivalently, exponential family) distribution consistent with the agent's reported mean μ .

The scoring rule compensates the agent according to the log likelihood of the eventual outcome according to this distribution. The interpretation is that the agent is only providing partial information about the underlying density, so the principal first infers a full density according to the principle of maximum entropy, and then scores the agent using the standard log score. We stress that the exponential families framework is a tool in the construction of scoring rules at this point; we are not actually modeling any agent's belief via an exponential family.

An advantage of this generalization of the log score is that, for many domains (multi-dimensional included) and expectations of interest, it leads to novel closed-form scoring rules. By examining the log densities of various exponential families, we can for instance obtain scoring rules for several different combinations of the arithmetic, geometric, and harmonic means, as well as higher order moments. The following examples illustrate the construction. Here, we have used μ and σ for the true mean and variance and $\hat{\mu}$ and $\hat{\sigma}$ for the elicited mean and variance respectively (and $\hat{\Sigma}$ for the covariance matrix of the multivariate normal).

Example 2.3. As a base measure we take the counting measure over the finite set $\{1, 2, \dots, d\}$, and we consider the statistic $\phi(i) \in \{0, 1\}^d$ that maps outcome i to the unit vector with component i equal to 1. The expectation of $\phi(i)$ is simply the entire probability distribution p , which is a multinomial distribution (an instance of an exponential family). We recover the standard log scoring rule $S(\hat{p}, i) = \log \hat{p}(i)$.

Example 2.4. As base measure we take the Lebesgue restricted to $[0, +\infty)$, and we consider the statistic $\phi(x) = x$ so that we are simply eliciting the mean. The maximum entropy distribution with a given mean μ is the exponential distribution, and taking its log density gives the scoring rule

$$S(\hat{\mu}, x) = -\frac{x}{\hat{\mu}} - \log \hat{\mu}. \quad (8)$$

We stress that although this rule is derived from the exponential distribution, Theorem 2.2 implies that it elicits the mean of any distribution supported on the non-negative reals (e.g., Pareto, lognormal). Indeed, it is easy to see that the expected score (8) depends only on the mean of the agent's belief because it is linear in x . As a generalization of this example, the maximum entropy distribution for the k -th moment $\phi(x) = x^k$ with respect to the same base measure is the Weibull distribution. Taking its log density leads to the following equivalent scoring rule:

$$S(\hat{\mu}, x) = (k-1) \log x - k \log \hat{\mu} - \Gamma \left(1 + \frac{1}{k} \right)^k \left(\frac{x}{\hat{\mu}} \right)^k, \quad (9)$$

where Γ denotes the gamma function (the extension of the factorial to the reals). We have not found either scoring rule (8) or (9) in the literature.

Example 2.5. As a base measure we take the Lebesgue over the real numbers \mathbf{R} . We are interested in eliciting the mean μ and variance σ^2 , so as a statistic we take $\phi(x) = (x, x^2)$ for which $\mathbf{E}_p[\phi(x)] = (\mu, \mu^2 + \sigma^2)$. The max entropy distribution for a given mean and variance is the Gaussian, whose log density gives the scoring rule

$$S((\hat{\mu}, \hat{\sigma}^2), x) = -\frac{(x - \hat{\mu})^2}{\hat{\sigma}^2} - \log \hat{\sigma}^2. \quad (10)$$

Again, we stress that this scoring rule elicits the mean and variance of any density over \mathbf{R} , not just those of a normal distribution. The construction easily generalizes to a multi-dimensional outcome space by taking the log density of the multivariate normal:

$$S((\hat{\mu}, \hat{\Sigma}), x) = -(x - \hat{\mu})' \hat{\Sigma}^{-1} (x - \hat{\mu}) - \log |\hat{\Sigma}|. \quad (11)$$

Here the statistics being elicited are the mean vector μ and the covariance matrix Σ . These scoring rules have been studied by Dawid and Sebastiani [1999] as rules that only depend on the mean and variance of the reported density. They note that these rules are weakly proper (because they do not distinguish between densities with the same first and second moments), but do not make the point that knowledge of the full density is not necessary on the part of the agent.

In the above, Example 2.5 illustrates an important point about parametrizations of the elicited expectations. The variance σ^2 cannot be written as $\mathbf{E}[\phi(x)]$ for any ϕ , because the mean μ enters the definition of σ^2 but is not available when ϕ is defined (indeed it is elicited in tandem with the variance).⁴ Instead one must use the first two *uncentered* moments $\mathbf{E}[x]$ and $\mathbf{E}[x^2]$. These are in bijection with μ and σ^2 , so the resulting scoring rule can be re-written in terms of the latter. Therefore, it is possible to elicit not just expectations but also bijective transformations of expectations.

3. EXPONENTIAL FAMILY MARKETS

In a single-agent setting, a scoring rule is used to *elicit* the agent’s belief. In a multi-agent setting, a prediction market can be used to *aggregate* the beliefs of the agents. In his seminal paper Hanson [2003] introduced the idea of a market scoring rule, which inherits the appealing elicitation and aggregation properties of both in order to perform well in thin or thick markets. In this section, we adapt the generalized log scoring rule to a market scoring rule which leads to markets with simple closed-form cost functions for many statistics of interest.

3.1. Prediction Market

In a prediction market an agent’s expected belief μ is elicited indirectly through the purchase and sale of contingent claim securities. Under this approach, each component i of the statistic ϕ is interpreted as the payoff function of a security; that is, a single share of security i pays off $\phi_i(x)$ when $x \in \mathcal{X}$ occurs. Thus if the portfolio of shares held by the agent is $\delta \in \mathbf{R}^d$, where entry δ_i corresponds to the number of shares of security i , then the payoff to the agent when x occurs is the inner product $\langle \delta, \phi(x) \rangle$.

To be concrete, recall from Example 2.3 that in the classic finite-outcome case the statistic has a component for each outcome x such that $\phi_x(x') = 1$ if $x' = x$ and 0 otherwise. Therefore the corresponding security pays 1 dollar if outcome x occurs. (These are known as Arrow-Debreu securities.) In Example 2.4 the one-dimensional statistic is $\phi(x) = x$, corresponding to a security whose payoff is linear in the outcome $x \in \mathbf{R}_+$. (This amounts to a futures contract.)

The standard way to implement a prediction market in the literature, due to Chen and Pennock [2007], is via a centralized market maker. The market maker maintains a convex, differentiable cost function $C : \mathbf{R}^d \rightarrow (-\infty, +\infty]$, where $C(\theta)$ records the revenue collected when the vector of outstanding shares is θ . The cost to an agent of purchasing portfolio δ under a market state of θ is $C(\theta + \delta) - C(\theta)$, and therefore the instantaneous prices of the securities are given by the gradient $\nabla C(\theta)$.

A risk-neutral agent will choose to acquire shares up to the point where, for each share, expected payoff equals marginal price. Formally, if the agent acquires portfolio δ , moving the market state vector to $\theta' = \theta + \delta$, then we must have

$$\mathbf{E}_p[\phi(x)] = \nabla C(\theta'). \tag{12}$$

⁴This is an intuitive but far from formal explanation for the fact that the dimension of the message space, or *elicitation complexity*, for eliciting the variance is at least 2 [Lambert et al. 2008].

In this way, by its choice of δ , the agent reveals that its expected belief is $\mu = \nabla C(\theta')$. We stress that this observation relies on the assumptions that 1) the agent is risk-neutral, 2) the agent does not incorporate the market's information into its own beliefs, and 3) the agent is not budget constrained. We will examine relaxations of each assumption in later sections.

3.2. Information-Theoretic Interpretation

In the remainder of this paper we focus on the following cost function, which arises from the “generalized” logarithmic market scoring rule (LMSR):

$$C(\theta) = \log \int_{x \in \mathcal{X}} \exp\langle \theta, \phi(x) \rangle d\nu(x). \quad (13)$$

This is of course exactly the log-partition function (5) for the exponential family with sufficient statistic ϕ , and we recover the classic LMSR using outcome indicator vectors as statistics. Because an agent would never select a portfolio with infinite cost, the effective domain (i.e., the possible vectors of outstanding shares) of C is $\Theta = \{\theta \in \mathbf{R}^d : C(\theta) < +\infty\}$, which gives an economic interpretation to the natural parameter space of an exponential family.

The correspondence between the cost function (13) and the log-partition function (5) suggests the following interpretation. The market maker maintains an exponential family distribution over the state space \mathcal{X} parametrized by share vectors that lie in Θ . When an agent buys shares, it moves the distribution's natural parameter so that the market prices matches its beliefs, or in other words the market's mean parametrization matches the agent's expectation.

There is a well-known duality between scoring rules and cost-function based markets [Abernethy et al. 2013; Hanson 2003]. In our context, recall from (6) that

$$\mathbf{E}_{\tilde{p}}[\log p(x; \hat{\mu})] = \langle \hat{\theta}, \mu \rangle - T(\hat{\theta})$$

where \tilde{p} is the agent's belief and $\hat{\mu}$ the agent's report. The expected log score from reporting $\hat{\mu}$ is exactly the same as the expected payoff from buying portfolio of shares $\hat{\theta} = \nabla C(\hat{\mu})$ (assuming an initial market state of 0), as $\langle \hat{\theta}, \mu \rangle$ is the expected revenue and $T(\hat{\theta})$ is the cost. As in Section 2 this reasoning relies on the assumption of risk-neutrality, not on any specific form for the agent's belief.

The agent's expected profit from moving the share vector from θ to θ' is

$$\begin{aligned} & \langle \theta' - \theta, \mu \rangle - C(\theta') + C(\theta) \\ &= C(\theta) - C(\theta') - \langle \theta - \theta', \nabla C(\theta) \rangle \\ &= D_C(\theta, \theta') = D_{C^*}(\mu', \mu), \end{aligned}$$

recalling (7). Now Banerjee et al. [2005b] have observed (among others) that the Kullback-Leibler divergence between two exponential family distributions is the Bregman divergence, with respect to the log-partition function, between their natural parameters. The agent's expected profit is therefore the KL divergence between the market's implied expectation and the exponential family corresponding to the agent's expectation, a well-known property from the classical LMSR [Hanson 2007].

3.3. Examples: Real Line and the Sphere

Let us now revisit our scoring rule examples from Section 2 in the context of prediction markets. The relevant entities now are the payoff function, the effective domain of shares, and the cost function.

Example 3.1. We consider a finite outcome space $\mathcal{X} = \{1, 2, \dots, d\}$ and set up a market with an Arrow-Debreu security ϕ_i for every outcome i : the security pays off 1 if i occurs and 0 otherwise. The log-partition function recovers the classical LMSR cost function: $C(\theta) = \log \sum_{i \in \mathcal{X}} e^{\theta_i}$. The effective domain is $\Theta = \mathbf{R}^d$, meaning that a security can always be bought or sold short.

Example 3.2. We consider outcomes over the positive reals \mathbf{R}_+ and set up a market for the expected outcome, consisting of a single security that pays off $\phi(x) = x$. The log partition function of the exponential distribution leads to the following cost function:

$$C(\theta) = -\log(-\theta).$$

The effective domain is $\Theta = \{\theta \in \mathbf{R} : \theta < 0\}$. This means the market must start with a negative number of outstanding shares for the security, and the number of shares must stay negative. The market maker need not explicitly enforce this, because by the Legendre property of C the cost tends to $+\infty$ as the outstanding shares approach the boundary, which is straightforward to see in this example.

Example 3.3. We consider outcomes over the real line \mathbf{R} and set up a market with securities corresponding to the first two uncentered moments (i.e, agents are betting on the return and volatility). The securities are defined by the payoffs $\phi(x) = (x, x^2)$. The log partition function of the normal distribution, under its natural parametrization, leads to the following cost function:

$$C(\theta) = -\frac{\theta_1^2}{4\theta_2} - \frac{1}{2} \log(-2\theta_2).$$

The effective domain is $\Theta = \{(\theta_1, \theta_2) \in \mathbf{R}^2 : \theta_2 < 0\}$. Again, we have here an instance where it is not possible for the number of outstanding shares of the second security to exceed 0. However, an arbitrary amount of the securities can be sold short, which corresponds to increasing the variance of the market's estimate.

Example 3.4. As another example let the outcome space be the d -dimensional unit sphere. This setting was considered by Abernethy et al. [2013] who provide a cost function implicitly defined through a variational characterization. The maximum entropy approach leads to another alternative. We have a security for each of the d dimensions, and security i simply pays off $\phi_i(x) = x_i$, where $x \in \mathbf{R}^d$ is the unit-norm outcome. The maximum entropy distribution over the sphere with such sufficient statistics is the von Mises-Fisher distribution. The log partition function corresponds to

$$C(\theta) = I_{\frac{d}{2}-1}(\|\theta\|) - \left(\frac{d}{2} - 1\right) \log \|\theta\|,$$

where I_r refers to the modified Bessel function of first kind and order r ; see Banerjee et al. [2005b] for an explanation of these quantities. The effective domain of θ is the positive orthant in \mathbf{R}^d . The mean parametrization of the von Mises-Fisher distribution gives a generalized log scoring rule for the expected outcome components, but it is unwieldy and involves several special functions.

4. BAYESIAN TRADERS WITH LINEAR UTILITY

In the standard model of cost-function based prediction markets, a sequence of myopic, risk-neutral agents arrive and trade in the market [Chen and Pennock 2007; Chen and Vaughan 2010]. As we saw in Section 3.1, such a trader moves the prices to its own expectation μ . However, this means that the market does not perform meaningful aggregation of agents' beliefs, as the final prices are simply the final agent's expectation.

In this section we examine the aggregation behavior of the market when agents are Bayesian and take into account the current market state when forming their beliefs. This requires more structure to their beliefs. For this section and the remainder of the paper, we will assume that agents have *exponential family beliefs*.

The exponential families framework is well-suited to reasoning about Bayesian updates. As before let the data distribution be given by $p(x; \theta) = \exp(\langle \theta, \phi(x) \rangle - T(\theta))$ where T is the log partition function and ϕ are the sufficient statistics. Instead of direct beliefs about the data distribution, the agent maintains a conjugate prior over the parameters θ . Every exponential family admits a conjugate prior of the form

$$p(\theta; n\nu, n) = \exp(\langle n\nu, \theta \rangle + nT(\theta) - \psi(\nu, n)).$$

Note that this is also an exponential family with natural parameter $(n\nu, n)$ where $\nu \in \mathbf{R}^d$ and n is a positive integer. The sufficient statistic maps θ to $(\theta, T(\theta))$, and the log partition function ψ is defined as the normalizer as usual. For a complete treatment of exponential families conjugate priors, see for instance [Barndorff-Nielsen 1978]. Now Diaconis and Ylvisaker [1979, Thm. 2] and Jewell [1974] have shown that

$$\mathbf{E}_{\theta \sim (n\nu, n)} \mathbf{E}_{x \sim \theta} [\phi(x)] = \nu, \quad (14)$$

meaning that $\nu = n\nu/n$ is the posterior mean. Thus, it is helpful to think of the prior as being based on a ‘phantom’ sample of size n and mean ν . Suppose now that the agent observes an empirical sample with mean $\hat{\mu}$ and size m . By a standard derivation [see Diaconis and Ylvisaker 1979], the posterior conjugate prior parameters become $n\nu \leftarrow n\nu + m\hat{\mu}$ and $n \leftarrow n + m$, and the posterior expectation (14) evaluates to

$$\frac{n\nu + m\hat{\mu}}{n + m}. \quad (15)$$

Thus the posterior mean is a convex combination of the prior and empirical means, and their relative weights depend on the phantom and empirical sample sizes.

Consider Bayesian agents maintaining an exponential family conjugate prior over the data model’s natural parameters (equivalently, the expected security payoffs). Each agent has access to a private sample of the data of size m with mean statistic $\hat{\mu}$. If n agents have arrived before to trade, then the current market prices μ correspond to the phantom sample, and the phantom sample size is nm . After forming the posterior (15) with these substitutions, the (risk-neutral) agent purchases shares δ to move the current market share vector to $\nabla C(\theta + \delta) = \frac{n\nu + \mu}{n+1}$. As a result, the final market prices under this behavior are a simple average of the agent’s mean parameters and the initial market prices. We note that to facilitate such belief updating, the market should post the number of trades since initialization.

5. RISK-AVERSE TRADERS WITH EXPONENTIAL UTILITY

In this section we relax the standard assumption that agents in the market are risk-neutral. We show that with sufficient extra structure to the agents’ beliefs and utilities, the market performs a clean aggregation of the agents’ beliefs via a simple weighted average. Assume that the agent has an exponential utility function for wealth w :

$$U_a(w) = -\frac{1}{a} \exp(-aw). \quad (16)$$

Here a controls the risk aversion: the agent’s aversion grows as a increases, and as a tends to 0 we approach linear utility (risk-neutrality). Specifically, a is the Arrow-Pratt coefficient of absolute risk aversion, and exponential utilities of the form (16) are the unique utilities that exhibit constant absolute risk aversion [Varian 1992, Chap. 11].

If wealth is distributed according to a probability measure P , then the *certainty equivalent* of a random amount of wealth is defined as

$$CE(w) = U_a^{-1}(\mathbf{E}_P[U_a(w)]).$$

Suppose that the agent's belief over outcomes is a density p with respect to base measure ν . There is a close relationship between the log-partition function and the certainty equivalent under exponential utility [see Ben-Tal and Teboulle 2007].

LEMMA 5.1. *The certainty equivalent of the agent's expected profit, with exponential utility, when acquiring shares δ under a market state of θ is*

$$\log a - T^p(-a\delta) - aC(\theta + \delta) + aC(\theta), \quad (17)$$

where T^p is the log partition function (5) with a base measure of $p \, d\nu$. Furthermore, if the agent's belief is an exponential family with natural parameter $\hat{\theta}$, we have

$$T^p(\delta) = T(\hat{\theta} + \delta) - T(\hat{\theta}),$$

where T is the usual log partition function with base measure ν .

PROOF. Explicitly, the certainty equivalent of the profit is

$$\begin{aligned} & CE(\langle \delta, \phi(x) \rangle - [C(\theta + \delta) - C(\theta)]) \\ &= -\log \int_{\mathcal{X}} \frac{1}{a} \exp(\langle -a\delta, \phi(x) \rangle + a[C(\theta + \delta) - C(\theta)]) p(x) d\nu(x) \\ &= \log a - a[C(\theta + \delta) - C(\theta)] - \log \int_{\mathcal{X}} \exp\langle -a\delta, \phi(x) \rangle p(x) d\nu(x) \\ &= \log a - a[C(\theta + \delta) - C(\theta)] - T^p(-a\delta). \end{aligned}$$

For the second part of the result, we have

$$\begin{aligned} T^p(\delta) &= \log \int_{\mathcal{X}} \exp\langle \delta, \phi(x) \rangle p(x; \hat{\theta}) d\nu(x) \\ &= \log \int_{\mathcal{X}} \exp(\langle \delta + \hat{\theta}, \phi(x) \rangle - T(\hat{\theta})) d\nu(x) \\ &= T(\hat{\theta} + \delta) - T(\hat{\theta}) + \log \int_{\mathcal{X}} \exp(\langle \delta + \hat{\theta}, \phi(x) \rangle - T(\hat{\theta} + \delta)) d\nu(x) \\ &= T(\hat{\theta} + \delta) - T(\hat{\theta}) + \log \int_{\mathcal{X}} p(x; \hat{\theta} + \delta) d\nu(x) = T(\hat{\theta} + \delta) - T(\hat{\theta}), \end{aligned}$$

where the last line follows from the fact that density $p(x; \hat{\theta} + \delta)$ integrates to 1. \square

Recall that for the generalized LMSR, the cost function C is exactly the log partition function T . We are therefore lead to the following understanding of a risk-averse agent's behavior in such a market.

THEOREM 5.2. *Suppose an agent has exponential utility with coefficient a and exponential family beliefs with natural parameter $\hat{\theta}$. In the generalized LMSR market with current market state θ , the agent's optimal trade δ moves the state vector to*

$$\theta + \delta = \frac{1}{1+a} \hat{\theta} + \frac{a}{1+a} \theta. \quad (18)$$

PROOF. The agent's optimal trade maximizes its expected utility, or equivalently the certainty equivalent. From Lemma 5.1 and that $T = C$, the agent maximizes

$$\log a - T(\hat{\theta} - a\delta) + T(\hat{\theta}) - aT(\theta + \delta) + aT(\theta).$$

This objective is strictly concave, from the strict convexity of T . The optimum is therefore characterized by the first-order condition $\nabla T(\hat{\theta} - a\delta) = \nabla T(\theta + \delta)$. As the gradient map ∇T is one-to-one, this is solved by equating the arguments, which leads to $\delta = (\hat{\theta} - \theta)/(1 + a)$ and (18). \square

We stress that the aggregation in Theorem 5.2 arises directly from the agents' choice of trades. The market maker does not need to know the risk aversion parameters. Note that as a tends to 0 we approach risk neutrality and the agent moves the share vector all the way to its private estimate $\hat{\theta}$. As a grows larger (the agent grows more risk averse) the agent makes smaller trades to reduce its exposure, and the final state stays closer to the current state θ . Update (18) implies that, under the conditions of the theorem, a market that receives a sequence of myopic traders aggregates their natural parameters in the form of an exponentially weighted moving average. The final market estimates (i.e., prices) are obtained by applying ∇T to this average.

Liquidity Adjustment. In practice the centralized market maker allows itself some control over the *liquidity* in the market, which captures how responsive prices are to trades. To adjust liquidity we consider the parametrized cost $C_\lambda(\theta) = \frac{1}{\lambda}C(\lambda\theta)$. Here λ is construed as the *inverse liquidity*, or price responsiveness. A larger setting of λ means fewer shares need to be bought to reach the same prices.⁵

In the context of the generalized LMSR we write T rather than C where T is the log partition function, with liquidity-adjusted version T_λ . Let $\hat{\mu}$ be the agent's mean belief with corresponding natural parameter $\hat{\theta} = \nabla T^{-1}(\hat{\mu})$. Recall that a risk-neutral agent moves the share vector so that the prices match its mean parameter. Therefore, define the *target shares* as $\tilde{\theta} = \nabla T_\lambda^{-1}(\hat{\mu})$. The target shares $\tilde{\theta}$ and natural parameter $\hat{\theta}$ are related by $\nabla T(\hat{\theta}) = \nabla T_\lambda(\tilde{\theta}) = \hat{\mu}$. In addition it is straightforward to check that $\nabla T_\lambda(\tilde{\theta}) = \nabla T(\lambda\tilde{\theta})$ so we have

$$\tilde{\theta} = \hat{\theta}/\lambda. \quad (19)$$

Higher price responsiveness means fewer shares must be bought to make the market prices match the agent's expectation, so the natural parameter is scaled down accordingly. With a liquidity adjustment the analysis of Theorem 5.2 can be extended and yields the following result, where as before θ and μ are the market's outstanding shares and prices respectively.

COROLLARY 5.3. *Under the conditions of Theorem 5.2 and an inverse liquidity of λ , the agent's optimal trade δ moves the state vector to*

$$\theta + \delta = \frac{\lambda}{\lambda + a}\tilde{\theta} + \frac{a}{\lambda + a}\theta = \frac{\lambda}{\lambda + a}\nabla T_\lambda^{-1}(\hat{\mu}) + \frac{a}{\lambda + a}\nabla T_\lambda^{-1}(\mu). \quad (20)$$

According to (20), as λ grows large the agent moves the market state closer to the *target shares*, rather than its true natural parameter. Note that the target shares themselves depend on λ by (19), but the update can be directly written in terms of the agent's beliefs as in the right-hand side of (20).

5.1. Repeated Trading and the Effective Belief

In previous sections we have analyzed trader behavior assuming it is his first entry into the market. We now pose the question: how will a trader reason about a possible

⁵The liquidity adjustment to the cost function takes the same form as the risk-aversion adjustment to the exponential utility in (16). In convex analysis, this transformation is known as the perspective function [Hiriart-Urruty and Lemaréchal 2000, p. 90].

future investment when the trader holds an existing portfolio? In the context of a trader possessing an exponential family belief together with exponential utility, we show that we can explicitly analyze how an agent incorporates an existing portfolio. The key conclusion is that a trader will reason about a future investment simply as though he had updated his belief and had no prior investment.

Suppose an exponential utility agent has exponential family belief parametrized by natural parameter $\hat{\theta}$. Based on this belief, let δ_1 be the vector of shares the agent has purchased on first entry in the market. On a subsequent entry into this market with market state θ' , his optimal purchase δ_2^* is given by the solution of

$$\arg \max_{\delta_2} \mathbb{E}_{x \sim p(x; \hat{\theta})} U [\langle \delta_1 + \delta_2, \phi(x) \rangle - C(\delta_1 + \theta) + C(\theta) - C(\delta_2 + \theta') + C(\theta')].$$

Then if $\theta'' = \hat{\theta} - a\delta_1$ is the effective belief, the trader's optimal purchase is given by $\delta_2 = (\theta'' - \theta')/(1 + a)$, moving the share vector to $\theta' + \delta_2 = \frac{1}{1+a}\theta'' + \frac{a}{1+a}\theta'$, which is a convex combination of the effective belief and the current market state.

THEOREM 5.4. *Suppose an exponential utility maximizing trader with utility parameter a who has belief $\hat{\theta}$ makes a purchase δ in a market. On subsequently re-entering the market, he will behave identically to an exponential utility maximizing trader with belief $\hat{\theta} - a\delta$ and no prior exposure in the market.*

Theorem 5.4 implies that financial exposure can be equivalently understood as changing the privately held beliefs.

5.2. Equilibrium Market State for Exponential Utility Agents

We have shown that every exponential-utility maximizing trader picks the share vector δ so that the eventual market state can be represented as a convex combination of the current market state and the natural parameter of his (exponential family) belief distribution. In this section we will compute the equilibrium state in an exponential family market with multiple such traders.

We draw on a well-known result from game theory regarding the class of *potential games*. We say a function $f(\vec{x})$ is at a *local optimum* if changing any coordinate of \vec{x} does not increase the value of f .

THEOREM 5.5 (MONDERER AND SHAPLEY [1996]). *Let $U_i(\vec{\delta})$ be the utility function of the i^{th} trader given the strategy profile $\vec{\delta} = (\delta_1, \dots, \delta_i, \dots, \delta_n)$. Assume there exists a potential function $\Phi(\vec{\delta})$ such that, for any $\vec{\delta}$ and δ'_i ,*

$$U_i(\vec{\delta}) - U_i(\vec{\delta}_{-i}, \delta'_i) = \Phi(\vec{\delta}) - \Phi(\vec{\delta}_{-i}, \delta'_i).$$

Then $\vec{\delta}$ is a pure-strategy Nash equilibrium if and only if $\vec{\delta}$ is a local optimum for $\Phi(\cdot)$.

In the exponential family market, the cost function C is identical to the log partition function T defined in (5). Let $\vec{\delta}$ be the matrix of share vectors purchased by every trader in the market at equilibrium. Let θ be the initial market state, $\hat{\theta}_i$ the natural parameter of trader i 's belief distribution, and a_i his risk aversion parameter.

Define a potential function as $\Phi(\vec{\delta}) = T(\theta + \sum_i \delta_i) + \sum_i \frac{1}{a_i} T(\hat{\theta}_i - a_i \delta_i)$. Rather than working directly with the utilities of every trader, we will work with the log of their utility values.⁶ Now the log-utility of trader i is

$$U_i(\vec{\delta}) = -T(\theta + \sum_j \delta_j) + T(\theta + \sum_{j \neq i} \delta_j) - \frac{1}{a_i} T(\hat{\theta}_i - a_i \delta_i) + \frac{1}{a_i} T(\hat{\theta}_i).$$

⁶Note: the potential function analysis still applies for any monotonically increasing transformation of the traders' utility functions.

We can now apply Theorem 5.5, hence the equilibrium state is obtained by jointly maximizing $\Phi(\vec{\delta})$ for each δ_i :

$$\nabla_{\delta_i} \Phi(\vec{\delta}) = \nabla T \left(\theta + \sum_{j=1}^n \delta_j \right) - \nabla T(\hat{\theta}_i - a_i \delta_i) = 0.$$

This leads to the following expression for the final market state.

$$\theta + \sum_{j=1}^n \delta_j = \frac{\theta + \sum_{i=1}^n \left(\frac{\hat{\theta}_i}{a_i} \right)}{1 + \sum_{i=1}^n \frac{1}{a_i}}$$

We see that the equilibrium state is a convex combination of the initial market state and all agent beliefs, with the latter weighted according to risk tolerance.

We finally note that we have not considered the dynamics by which such an equilibrium would be reached, nor the rate of convergence etc., yet we think such questions provide fruitful directions for future research.

6. BUDGET-LIMITED AGGREGATION

In this section, we consider the evolution of the market state when traders are budget-limited. We assume that the traders trade in multiple instances of the market. That is, after the initial setup and trades, the outcome is revealed and the traders are paid off. Then, the market runs again with the same set of traders and the outcome drawn from a (possibly different) exponential family distribution. The goal of this section is to analyze how the budgets of traders evolve under these circumstances. As before, the market price is interpreted as a probability density over the outcome space and the share vector as the natural parameter of an exponential family distribution. Consistent with the connections drawn in Section 2 and throughout, we measure the error in prediction using the standard log loss.

We show that traders with faulty information can only impose a limited amount of additional loss to the market's prediction. Further, since informative traders experience an expected increase in budget, they will eventually be unconstrained and allowed to carry out unrestricted trades. Taken together, this means that while the market suffers limited damage from ill-informed traders, it is also able to make use of all the information from informative traders in the long run.

Budget-limited trades. Let α be the budget of a trader in the market. Suppose that with infinite budget, the trader would have moved the market state from θ to $\hat{\theta}$, where $\hat{\theta}$ represents his true belief. Now suppose further that $\alpha < C(\hat{\theta}) - C(\theta)$; that is, the trader's budget disallows purchasing enough shares to move the market state to his belief. In this case, we want to budget-limit the trader's influence on the market state.

Let the current market state be given by θ and let the final market state be $\theta' = \lambda \hat{\theta} + (1 - \lambda)\theta$ where $\lambda = \min \left(1, \frac{\alpha}{C(\hat{\theta}) - C(\theta)} \right)$. The cost to the trader to move the market state from θ to θ' is at most his budget α and is called his *budget-limited trade*.

Note that while an uninformative trader is unconstrained in his belief, the informative trader's belief corresponds to the parameter of the true distribution.

Limited Damage. We will now quantify the error in prediction that the market maker might have to endure as a result of ill-informed entities entering the market. We assume that these entities trade in multiple instances of the market; thus the exposure of the market maker is over several rounds. The log loss function for θ shares held is defined as $L(\theta, x) = -\log p(x; \theta) = C(\theta) - \langle \theta, \phi(x) \rangle$.

LEMMA 6.1. *The loss induced on the market by an uninformative trader is bounded by his initial budget.*

PROOF. First consider the change in budget of a trader i over multiple rounds of the prediction market. Let his budget at rounds t and $t-1$ be α_i^t and α_i^{t-1} respectively. The change in budget for trader i moving the market state from θ to θ' with outcome x^t is

$$\alpha_i^t - \alpha_i^{t-1} = C(\theta) - C(\theta') - (\theta - \theta')^T \phi(x^t) = L(\theta, x^t) - L(\theta', x^t) = \Delta_i^t$$

Here Δ_i^t is called the myopic impact of a trader i in round t . Thus, the myopic impact captures incremental gain in prediction due to the trader in a round and is equal to the change in his budget in that round.

Since the market evolves so that the budget of any trader never falls below zero, the total myopic impact in T rounds caused due to trader i is

$$\Delta_i := \sum_{t=1}^T \Delta_i^t = \sum_{t=1}^T (\alpha_i^t - \alpha_i^{t-1}) = \alpha_i^T - \alpha_i^0 \geq -\alpha_i^0$$

□

An interesting aspect of Lemma 6.1 is that the log loss can be quantified in the same units as the traders' budgets.

Budget of Informative Traders. We now characterize the expected change in budget for an informative trader.

LEMMA 6.2. *Let θ be the current market state. Suppose that an informative trader with belief distribution parametrized by $\hat{\theta}$ moves the market state to the budget-limited state $\theta' = \lambda\hat{\theta} + (1-\lambda)\theta$. Then, the expectation (over the trader's belief) of the trader's profit is strictly positive whenever his budget is positive and his belief differs from the previous market position θ .*

PROOF. Let C be the log partition function T of the belief distribution. The payoff is given by the sufficient statistics $\phi(x)$. Then the trader's expected net payoff is given by

$$\begin{aligned} \mathbb{E}_{x \sim P_{\hat{\theta}}}[C(\theta) - C(\theta') - (\theta - \theta')^T \phi(x)] &= T(\theta) - \theta^T \nabla T(\hat{\theta}) - (T(\theta') - \theta'^T \nabla T(\hat{\theta})) \\ &= D_T(\theta, \hat{\theta}) - D_T(\theta', \hat{\theta}) \geq \lambda D_T(\theta, \hat{\theta}) \geq 0 \end{aligned}$$

where $D_T(\cdot, \cdot)$ is the Bregman divergence based on T . The second to last inequality holds since $D_T(\theta', \hat{\theta})$ is convex in θ' and we have:

$$D_T(\theta', \hat{\theta}) = D_T(\lambda\hat{\theta} + (1-\lambda)\theta, \hat{\theta}) \leq \lambda D_T(\hat{\theta}, \hat{\theta}) + (1-\lambda)D_T(\theta, \hat{\theta}) = (1-\lambda)D_T(\theta, \hat{\theta})$$

A trader who adjusts the market state may expect profit of at least $\lambda D_T(\theta, \hat{\theta})$. □

We note one important aspect of Lemma 6.2: the expectation is taken with respect to each trader's belief at the time of trade, rather than with respect to the true distribution. This is needed because we have made no assumptions about the optimality of the traders' belief updating procedure. If we assume that the traders' belief formation is optimal, then this growth result will extend to the true distribution as well.

Given a continuous density the probability a trader will form exactly the same beliefs as the current market position is 0, and thus, each trader will have positive expected profit on almost all sequences of observed samples and beliefs. This result suggests that, eventually, every informative trader will have the ability to influence the market state in accordance with his beliefs, without being budget limited.

Notice that Lemma 6.2 only required that the market state to which the trader moves be representable as a convex combination of the current market state and his belief. This means that the result holds for exponential utility traders aiming to maximize their utility by Theorem 5.2. In this case, the trader who moves the market state can expect his profit to be positive and at least $\frac{1}{a}D_T(\theta, \hat{\theta})$ where a is the exponential utility parameter. When the cost function is adjusted to C_λ with an inverse liquidity parameter λ (Section 5), the trader receives expected payoff of at least $\frac{1}{a}D_T(\lambda\theta, \hat{\theta})$.

ACKNOWLEDGMENTS

We would like to sincerely thank Rafael Frongillo and the anonymous referees for valuable feedback.

REFERENCES

- Jacob Abernethy, Yiling Chen, and Jennifer Wortman Vaughan. 2013. Efficient market making via convex optimization, and a connection to online learning. *ACM Transactions on Economics and Computation* 1, 2 (May 2013).
- Jacob Abernethy and Rafael M. Frongillo. 2012. A characterization of scoring rules for linear properties. In *Proceedings of the 25th Annual Conference on Learning Theory*.
- Arindam Banerjee, Inderjit S. Dhillon, and Joydeep Ghosh. 2005a. Clustering with Bregman divergences. *Journal of Machine Learning Research* 6 (2005), 1705–1749.
- Arindam Banerjee, Inderjit S. Dhillon, Joydeep Ghosh, and Suvrit Sra. 2005b. Clustering on the unit hypersphere using von Mises-Fisher distributions. *Journal of Machine Learning Research* 6 (2005), 1345–1382.
- Ole Barndorff-Nielsen. 1978. *Information and Exponential Families in Statistical Theory*. Wiley Publishers.
- Aharon Ben-Tal and Marc Teboulle. 2007. An old-new concept of convex risk measures: The optimized certainty equivalent. *Mathematical Finance* 17, 3 (2007), 449–476.
- Yiling Chen and David M. Pennock. 2007. A utility framework for bounded-loss market makers. In *In Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence (UAI)*. 49–56.
- Y. Chen, Mike Ruberry, and J. Wortman Vaughan. 2013. Cost function market makers for measurable spaces. In *Proceedings of the 14th ACM Conference on Electronic Commerce*.
- Y. Chen and J.W. Vaughan. 2010. A new understanding of prediction markets via no-regret learning. In *Proceedings of the 11th ACM Conference on Electronic Commerce (EC)*. 189–198.
- A. P. Dawid. 1998. *Coherent measures of discrepancy, uncertainty and dependence with applications to Bayesian predictive experimental design*. Technical Report 139. University College London, Dept. of Statistical Science.
- A Philip Dawid and Paola Sebastiani. 1999. Coherent dispersion criteria for optimal experimental design. *Annals of Statistics* 27 (1999), 65–81.
- Persi Diaconis and Donald Ylvisaker. 1979. Conjugate priors for exponential families. *Annals of Statistics* 7, 2 (1979), 269–281.
- Hans Föllmer and Thomas Knispel. 2011. Entropic risk measures: Coherence vs. convexity, model ambiguity and robust large deviations. *Stochastics and Dynamics* 11 (2011), 333–351.
- Hans Föllmer and Alexander Schied. 2002. Convex measures of risk and trading constraints. *Finance and Stochastics* 6, 4 (2002), 429–447.
- Hans Föllmer and Alexander Schied. 2004. *Stochastic Finance: An Introduction in Discrete Time*. Walter de Gruyter.
- Rafael M. Frongillo, Nicolás Della Penna, and Mark D. Reid. 2012. Interpreting pre-

- diction markets: a stochastic approach. In *Proceedings of Neural Information Processing Systems (NIPS)*.
- Rafael M. Frongillo and Mark D. Reid. 2013. Convex foundations for generalized MaxEnt models. In *Proceedings of the 33rd International Workshop on Bayesian Inference and Maximum Entropy Methods in Science and Engineering (MaxEnt)*.
- Xi Gao, Yiling Chen, and David M. Pennock. 2009. Betting on the real line. In *Proceedings of the 5th International Workshop on Internet and Network Economics (WINE)*. 553–560.
- T. Gneiting and A.E. Raftery. 2007. Strictly proper scoring rules, prediction, and estimation. *J. Amer. Statist. Assoc.* 102, 477 (2007), 359–378.
- Peter Grünwald and A. P. Dawid. 2004. Game theory, maximum entropy, minimum discrepancy and robust Bayesian decision theory. *Annals of Statistics* 32, 4 (2004), 1367–1433.
- R. Hanson. 2003. Combinatorial information market design. *Information Systems Frontiers* 5, 1 (2003), 105–119.
- R. Hanson. 2007. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets* 1, 1 (2007), 3–15.
- Jean-Baptiste Hiriart-Urruty and Claude Lemaréchal. 2000. *Fundamentals of Convex Analysis*. Springer.
- William S Jewell. 1974. Credible means are exact Bayesian for exponential families. *Astin Bulletin* 8, 1 (1974), 77–90.
- N. S Lambert, D. M Pennock, and Y. Shoham. 2008. Eliciting properties of probability distributions. In *Proceedings of the 9th ACM Conference on Electronic Commerce*. 129–138.
- John McCarthy. 1956. Measures of the value of information. *Proceedings of the National Academy of Sciences (PNAS)* 42, 9 (1956), 654.
- Dov Monderer and Lloyd S Shapley. 1996. Potential games. *Games and Economic Behavior* 14, 1 (1996), 124–143.
- Michael Ostrovsky. 2012. Information aggregation in dynamic markets with strategic traders. *Econometrica* 80, 6 (November 2012), 2595–2647.
- Abraham Othman and Tuomas Sandholm. 2011. Liquidity-sensitive automated market makers via homogeneous risk measures. In *Internet and Network Economics*. Springer, 314–325.
- David M Pennock, Steve Lawrence, C Lee Giles, Finn Arup Nielsen, and others. 2001. The real power of artificial markets. *Science* 291, 5506 (2001), 987–988.
- David M. Pennock and Rahul Sami. 2007. Computational aspects of prediction markets. In *Algorithmic Game Theory*, Noam Nisan, Tim Roughgarden, Eva Tardos, and Vijay V. Vazirani (Eds.). Cambridge University Press, Chapter 26.
- Leonard J. Savage. 1971. Elicitation of personal probabilities and expectations. *J. Amer. Statist. Assoc.* 66 (1971), 783–801.
- Amos J. Storkey. 2011. Machine learning markets. In *Proceedings of AI and Statistics (AISTATS)*. 716–724.
- Hal R. Varian. 1992. *Microeconomic Analysis*. W. W. Norton and Company.
- Martin J. Wainwright and Michael I. Jordan. 2008. Graphical models, exponential families, and variational inference. *Foundations and Trends in Machine Learning* 1 (January 2008), 1–305. Issue 1-2.
- Justin Wolfers and Eric Zitzewitz. 2006. *Interpreting prediction market prices as probabilities*. Technical Report. National Bureau of Economic Research.