

On Risk Measures, Market Making, and Exponential Families

JACOB D. ABERNETHY

University of Michigan

and

RAFAEL M. FRONGILLO

Harvard University

and

SINDHU KUTTY

University of Michigan

In this note we elaborate on an emerging connection between three areas of research: (a) the concept of a *risk measure* developed within financial mathematics for reasoning about risk attitudes of agents under uncertainty, (b) the design of *automated market makers* for prediction markets, and (c) the family of probability distributions known as *exponential families*.

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

General Terms: Economics, Theory

Additional Key Words and Phrases: exponential family; entropic risk measure; exponential utility

1. INTRODUCING RISK MEASURES

Imagine that a farmer must decide between several crops to plant for the upcoming growing season. The cost of each crop is different and the yield of each crop depends differently on weather conditions; one may be better suited for cold temperatures, another for heavy rainfall, and yet another for drought. The farmer's profits for the harvest will thus be determined not only by the cost of the seeds and planting, but also by the suitability of the chosen crops for the actual weather conditions during the season. Given that the weather is uncertain, how should the farmer choose the crops to maximize her profit? Generally speaking, we can model this type of problem by specifying a set Ω of future states of the world (in this case, the weather during the season), and considering the agent's *position* $X : \Omega \rightarrow \mathbb{R}$, which specifies the monetary payoff to the agent (in this case, the farmer's profit) in each such state $\omega \in \Omega$. Now the agent's actions each induce a different position, and the problem reduces to measuring the quality of such positions and choosing the best.

There are, of course, many ways to evaluate a financial position, including von Neumann-Morgenstern expected utility theory. In this note, we focus on the concept of a (*convex*) *risk measure*, which was introduced by the academic finance community [Artzner et al. 1999; Delbaen 2002; Föllmer and Schied 2004] and has appeared more recently in computer science [Othman and Sandholm 2011; Hu and

Authors' addresses: jabernet@umich.edu, raf@cs.berkeley.edu, skutty@umich.edu

Storkey 2014]. Here, an agent chooses the position which minimizes her risk measure. In general, risk measures are convex functions of positions which satisfy certain axioms, such as monotonicity ($X > X' \Rightarrow \rho(X) \leq \rho(X')$) and cash invariance ($\rho(X + c\mathbf{1}) = \rho(X) - c$, where $\mathbf{1}$ denotes the sure payoff of \$1). A major focus of this note will be the *entropic risk measure*; for any probability measure p over Ω , entropic risk is given by:

$$\rho_p(X) := \log \int_{\Omega} \exp(-X) dp \quad (1)$$

The entropic risk measure is related to one popular measure of utility of wealth m , $u(m) = -\exp(-m)$, commonly known as the *exponential utility function*. An agent holding some belief distribution p on Ω who maximizes expected exponential utility (under p) is identical to an agent who minimizes the entropic risk ρ_p . That is, for any two positions X, X' , we have that $\mathbb{E}_p[u(X)] \geq \mathbb{E}_p[u(X')] \iff \rho_p(X) \leq \rho_p(X')$. More generally, we can always construct a risk measure from any concave utility function $u(\cdot)$ and belief distribution p . As noted by [Föllmer and Schied 2004], we may define a risk measure $\rho_{u,p}(X) := \inf \{m : \mathbb{E}_{\omega \sim p}[u(X(\omega) + m)] \geq u_0\}$. That is, $\rho_{u,p}(X)$ is the least amount of money the agent needs to maintain expected utility above some default threshold u_0 , while holding position X .

It is typical to restrict the space of positions with respect to a payoff function $\phi : \Omega \rightarrow \mathbb{R}^d$, such that each position X under consideration can be written $X(\omega) = r^\top \phi(\omega)$ for some $r \in \mathbb{R}^d$. Given a fixed ϕ , we will often abuse notation and write $\rho_p(r)$ in place of $\rho_p(X)$ for the X defined above, and consider r to be the “compact” position. As we will see, this compact form lends itself well to the prediction market setting, where the component $\phi_i(\cdot)$ corresponds to the payout amount for the i th outcome-contingent contract sold in the market. In addition, for any $r, q \in \mathbb{R}^d$ it is convenient to define $\rho_p(r|q) := \rho_p(r + q) - \rho_p(q)$ which may be interpreted as the *relative risk* of r given a current position q .

2. MARKET MAKING IN PREDICTION MARKETS

Risk measures provide a surprisingly natural object to design a *prediction market* via an *automated market maker*. Prediction markets facilitate aggregation of information via financial incentives, and market designers typically aim to yield accurate predictions of uncertain future events. Goods in these markets correspond to securities with payoffs contingent on some future outcome. The goal is that the prices of these securities should reflect a useful aggregate of information from market participants.

Much attention has been given to the design of automated market makers which facilitate the market by offering to trade with any party at a given price. The task of the market maker is to adjust these prices according to demand. Various formulations of automated market makers have been proposed, such as *market scoring rules* [Hanson 2003], and the *constant-utility market maker* [Chen and Pennock 2007]; one that has received considerable attention is the *cost function market maker* [Abernethy et al. 2013; Chen and Pennock 2007]. In this framework, the market maker possesses a cost function $C : \mathbb{R}^d \rightarrow \mathbb{R}$ and a current “liability” $q \in \mathbb{R}^d$; a trader purchasing a bundle of securities $r \in \mathbb{R}^d$ pays $C(q+r) - C(q)$ to the market maker, who then updates the liability to $q + r$.

In [Abernethy et al. 2013], the cost function C is required to satisfy certain axioms (e.g. no arbitrage, expressiveness), which turn out to be essentially the same axioms as those for risk measures mentioned above. Thus, we may equivalently think of the market maker as possessing a compact position q and risk measure $\rho(q) := C(-q)$. Then a trader wishing to purchase bundle r must pay the market maker $\rho(q-r) - \rho(q) = \rho(-r|q)$.¹ It is easy to check that this transaction leaves the market maker’s risk unchanged regardless of r , via the cash-invariance principle. Finally, it is interesting to note that the constant-utility market maker of [Chen and Pennock 2007], when viewed as a risk measure, is the same as $\rho_{u,p}$ from above.

3. MARKET SEMANTICS OF EXPONENTIAL FAMILY DISTRIBUTIONS

We will now switch gears to talk about a popular family of probability distributions that turn out to be naturally connected with the entropic risk measure. Given access to empirical averages of some statistics of data, a natural question to ask is if we can find a distribution whose expected statistics match these observations. *Exponential family distributions* arise as the unique distribution which produce the desired statistics while maximizing Shannon entropy.

Nearly all of the popular probability distributions utilized in the literature can be expressed as an exponential family, including the *Gaussian*, the *multinomial*, the *Poisson*, etc. Let us consider, for example, the Gaussian distribution on a real-valued variable x . If instead of the typical parameters of mean μ and variance σ^2 , we use the *natural* parameters $\theta = (\frac{\mu}{\sigma^2}, \frac{-1}{2\sigma^2})$ and we define the vector function $\phi(x) := (x, x^2)$, then we see that the probability density of the Gaussian can be rewritten as $p_\theta(x) \propto \exp(\theta^\top \phi(x))$.

The probability density of all exponential families have a similar form, which we now describe. Given a space Ω and any function $\phi : \Omega \rightarrow \mathbb{R}^d$, we can define a probability density for every parameter vector θ in some feasible set Θ as

$$p_\theta(\omega) := \exp(\theta^\top \phi(\omega) - A(\theta)) \quad \text{where} \quad A(\theta) := \log \int_\Omega \exp(\theta^\top \phi(\omega)) d\nu. \quad (2)$$

The function $A(\cdot)$ is the normalization factor and is commonly known as the *log partition function*. It may not be lost on the reader that the definition of $A(\cdot)$ is conspicuously similar to the entropic risk defined in (1). Indeed, recent work [Abernethy et al. 2014; Frongillo 2013] has explored an alternative semantic interpretation of the exponential family framework: one can design a market maker by using the log partition function $A(\cdot)$ as a cost function (equivalently, risk measure). That is, we can imagine a market maker selling d reference securities which pay out according to function $\phi(\omega) \in \mathbb{R}^d$ upon observing the outcome ω . The market maker can interpret its position q as a vector of natural parameters θ , so that when traders request to purchase a share bundle $r \in \mathbb{R}^d$, the market maker charges the trader $A(\theta + r) - A(\theta)$ and, for outcome ω , pays the trader $r^\top \phi(\omega)$.

This characterization of the exponential family distribution with market semantics gives rise to a number of nice interpretations:

- (1) Given the market maker’s position θ , the “market prices” of the d securities announced by the market maker are identical to the *mean parameters* of p_θ .²

¹Note the change of sign, as risk measures deal with gains whereas cost functions deal with losses.

²The mean parameters of a distribution p are defined as $\mu_\theta := \mathbb{E}_{\omega \sim p}[\phi(\omega)]$.

- (2) We can view the market as simply updating its belief according to new information (i.e. trade), or we can alternatively view the market maker as simply maintaining constant risk according to the entropic risk measure.
- (3) We can imagine a trader in this market, with some initial belief p_θ , who trades to minimize the entropic risk measure $\rho_{p_\theta}(\cdot)$; as mentioned, this is equivalent to the trader maximizing the expected exponential utility. When the trader invests in a set of shares r , this will clearly affect future investment decisions. But the effect on the trader has two potential interpretations: (i) the agent updates the risk measure to $\rho_{p_\theta}(\cdot|r)$ or (ii) the agent replaces the original belief p_θ with an updated belief $p_{\theta+r}$. Indeed, the risk measure $\rho_{p_\theta}(\cdot|r)$ is *identical* to $\rho_{p_{\theta+r}}(\cdot)$, suggesting that one's portfolio and one's belief parameters are interchangeable quantities within this market framework.
- (4) We can imagine a trader who knows the true distribution p^* , and that $p^* = p_{\hat{\theta}}$ is a member of the exponential family. If the market maker's position is currently θ , then the trader has the potential to earn expected profit in the amount of $D_A(\theta, \hat{\theta}) = \text{KL}(p_{\hat{\theta}}||p_\theta)$, the Kullback-Leibler divergence between $p_{\hat{\theta}}$ and p_θ .³ The trader achieves this by purchasing $\hat{\theta} - \theta$ shares.

It is worth noting that many of the above properties hold more generally for other distribution families, in particular the class of *generalized exponential families*.⁴

4. INFORMATION AGGREGATION IN PREDICTION MARKETS

A central thread of research in the prediction market literature seeks to understand how the market aggregates the information of its participants. Results in this vein depend heavily on the equilibrium concept used, as well as how trader behavior is modeled. Many existing results show natural aggregation properties of the market prices, or equivalently, mean parameters [Wolfers and Zitzewitz 2006; Othman and Sandholm 2010; Frongillo et al. 2012]. Here we will present aggregation results which operate in the *share space*, or equivalently, in the natural parameters.

Consider a market maker with cost function based on the log partition function (2) as described above. Assuming that traders in this market wish to maximize expected utility with respect to their beliefs, we seek to characterize the market equilibrium, which we define to be the final market state after which no trader wishes to continue trading. It was shown by [Abernethy et al. 2014] that, when each trader i has exponential utility with risk tolerance parameter b_i and exponential-family belief parameters $\hat{\theta}_i$, the market equilibrium becomes

$$\theta_{\text{final}} = \theta_{\text{init}} + \sum_{j=1}^n \delta_j = \frac{\theta_{\text{init}} + \sum_{i=1}^n b_i \hat{\theta}_i}{1 + \sum_{i=1}^n b_i}, \quad (3)$$

where δ_i are trader i 's security purchases and θ_{init} is the initial market state. In other words, the equilibrium state is a risk-tolerance-weighted average of the natural parameters of the traders and the market maker, with the more risk tolerant traders taking on proportionally more of the final position.

³The notation $D_f(x, y)$ refers to the *Bregman divergence* defined as $f(x) - f(y) - \nabla f(y)^\top (x - y)$.

⁴Introduced by [Grünwald and Dawid 2004], these families are maximum-entropy distributions for entropy functions other than Shannon entropy. For details, see [Frongillo 2013, Chap. 4.3].

This result is quite natural, but appears to depend on the synergy between exponential families and exponential utility. Surprisingly, it is shown in [Barrieu and Karoui 2007; Frongillo and Reid 2014] that this result extends to risk-tolerance families of arbitrary risk measures: if the market maker is risk-constant with risk measure ρ , and each trader i seeks to minimize risk measure $\rho_i(X) = b_i\rho(X/b_i)$, then the equilibrium state is again the weighted average given by eq. (3).

While we have characterized the equilibrium state of these markets, it remains to understand how to reach it. In particular, does a more dynamic model of trader activity converge to this equilibrium? We consider a very simple dynamic: at each time step, a trader is selected at random, who computes the optimal trade δ_t given her current position in the market and the current market state. Then as long as every trader has a nonzero probability of being selected in each round, the unique fixed point of this dynamic is again the market equilibrium state (3). Moreover, we can bound the *rate* of this convergence: the optimality gap at time t , as measured by the sum of risks, is $O(1/t)$. These convergence results continue to hold even beyond the risk-tolerance setting, for *any* choice of risk measures ρ_i for the traders.

REFERENCES

- ABERNETHY, J., CHEN, Y., AND VAUGHAN, J. W. 2013. Efficient market making via convex optimization, and a connection to online learning. *ACM Transactions on Economics and Computation* 1, 2, 12.
- ABERNETHY, J., KUTTY, S., LAHAIE, S., AND SAMI, R. 2014. Information aggregation in exponential family markets. In *Proceedings of the fifteenth ACM conference on Economics and computation*. ACM, 395–412.
- ARTZNER, P., DELBAEN, F., EBER, J.-M., AND HEATH, D. 1999. Coherent measures of risk. *Mathematical finance* 9, 3, 203–228.
- BARRIEU, P. AND KAROUI, N. E. 2007. Pricing, hedging and optimally designing derivatives via minimization of risk measures. *arXiv:0708.0948 [math]*. arXiv: 0708.0948.
- CHEN, Y. AND PENNOCK, D. 2007. A utility framework for bounded-loss market makers. In *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence*. 49–56.
- DELBAEN, F. 2002. Coherent risk measures on general probability spaces. In *Advances in finance and stochastics*. Springer, 1–37.
- FRONGILLO, R. 2013. Eliciting private information from selfish agents.
- FRONGILLO, R., DELLA PENNA, N., AND REID, M. 2012. Interpreting prediction markets: a stochastic approach. In *Advances in Neural Information Processing Systems 25*. 3275–3283.
- FRONGILLO, R. M. AND REID, M. D. 2014. Risk dynamics in trade networks. *arXiv:1410.0413*.
- FÖLLMER, H. AND SCHIED, A. 2004. *Stochastic Finance: An Introduction in Discrete Time*. Walter de Gruyter.
- GRÜNWARD, P. AND DAWID, A. 2004. Game theory, maximum entropy, minimum discrepancy and robust Bayesian decision theory. *The Annals of Statistics* 32, 4, 1367–1433.
- HANSON, R. 2003. Combinatorial information market design. *Information Systems Frontiers* 5, 1, 107–119.
- HU, J. AND STORKEY, A. 2014. Multi-period trading prediction markets with connections to machine learning. *arXiv preprint arXiv:1403.0648*.
- OTHMAN, A. AND SANDHOLM, T. 2010. When do markets with simple agents fail? In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*. 865–872.
- OTHMAN, A. AND SANDHOLM, T. 2011. Liquidity-sensitive automated market makers via homogeneous risk measures. In *Proceedings of the 7th International Conference on Internet and Network Economics*. WINE’11. Springer-Verlag, Berlin, Heidelberg, 314–325.
- WOLFERS, J. AND ZITZEWITZ, E. 2006. Interpreting prediction market prices as probabilities. Tech. rep., National Bureau of Economic Research.