

Subsidized Prediction Markets for Risk Averse Traders*

Stanko Dimitrov¹, Rahul Sami², and Marina Epelman¹

¹ Department of Industrial and Operations Engineering

² School of Information

University of Michigan, Ann Arbor, MI 48109

{sdimitro, rsami, mepelman}@umich.edu

Abstract. In this paper we study the design and characterization of prediction markets in the presence of traders with unknown risk-aversion. We formulate a series of desirable properties for any “market-like” forecasting mechanism. We present a randomized mechanism that satisfies all these properties while guaranteeing that it is myopically optimal for each trader to trade honestly, regardless of her degree of risk aversion. We observe, however, that the mechanism has an undesirable side effect: the traders’ expected reward, normalized against the inherent value of their private information, decreases exponentially with the number of traders. We prove that this is unavoidable: any mechanism that is myopically strategyproof for traders of all risk types, while also satisfying other natural properties of “market-like” mechanisms, must sometimes result in a player getting an exponentially small normalized expected reward.

1 Introduction and Related Work

Prediction markets are markets designed and deployed to aggregate information about future events by having agents with private beliefs trade in these markets. One market format that is gaining in popularity is the *market scoring rule* [8]. A market scoring rule is a market mechanism with an automated market maker that guarantees liquidity, effectively subsidizing the market to incentivize trade.

Hanson [8] has shown that, for risk-neutral agents who are myopic, it is optimal for each to reveal their true beliefs on the traded event. This results leaves two questions. The first question, partially addressed by Chen *et al.* [6], is: *What happens when agents take into account future payoffs?*

In this paper, we tackle the second question: *What happens when agents are not risk-neutral?* In practice, most people are better modeled as being risk-averse in their decision making. Therefore, we model traders as expected-utility maximizers with an arbitrary weakly monotone and concave utility function. Current prediction market mechanisms, like the Market Scoring Rule or the Dynamic Pari-mutuel Market [11], do not always give appropriate incentives to risk-averse traders. For example, a sufficiently risk-averse informed trader, who knows that an event will occur with 80% probability

* This work was supported by the National Science Foundation under award CCF-0728768. The proofs have been omitted from this short paper; the extended version can be found at <http://www.umich.edu/~rsami/papers/RiskAverseMarkets.pdf> [7]

even though it is currently priced at 50%, may not want to push up the price in a market because of the 20% chance of making a loss. This suggests that current subsidized prediction markets may converge to a non-truthful price in a sequential equilibrium.

If traders have known risk aversion, the scoring rules could be adjusted, retaining the original incentive properties. In this paper, we focus on the setting where traders have *unknown* risk aversion, and study whether it is possible to modify the market mechanism to guarantee myopic honesty and preserve other desirable properties. We first list a set of properties that any prediction market-like mechanism must satisfy: (1) myopic strategyproofness; (2) sequential trade, giving traders the opportunity to update beliefs; (3) a variant of *sybilproofness*, capturing the idea that trading under multiple identities does not yield any direct advantage; and (4) the expected subsidy should be bounded.

We propose one mechanism that satisfies all of these properties in the presence of traders with unknown risk-averse preferences. The key building block of our result is a sweepstakes technique, developed by Allen [2]. Unfortunately, the proposed mechanism reduces the expected reward exponentially with the number of agents.

We then establish that exponentially decreasing rewards are unavoidable for any mechanism satisfying all the properties listed above. To exclude trivial examples of decreasing rewards, we normalize all rewards by a measure of the intrinsic informativeness of a trader's private information. We show that exponential decrease in the *normalized* expected reward is necessary for any mechanism that satisfies the properties we propose in the presence of arbitrarily risk-averse agents.

1.1 Related Work

In this section we discuss some of the previous work in prediction markets and scoring rules. Hanson [8] introduced the concept of a market scoring rule, a form of subsidized prediction market, and proved a myopic strategyproofness property for risk-neutral traders, as well as a bound on the total subsidy. Pennock [11] introduced another mechanism, the dynamic pari-mutuel market, for a subsidized prediction market. Both these mechanisms introduce some of the properties in section 2. However, both mechanisms assume risk neutrality of the traders, which is not assumed in this paper. Lambert *et al.* [9] introduce a class of self-financed wagering mechanism along with the properties such mechanisms must satisfy. The authors assume risk neutral traders and an absence of subsidy in the mechanism. Chen and Pennock [5], and later generalized by Agrawal *et al.* [1], consider a risk-averse market maker in a subsidized market and show that the market maker has bounded subsidy in most forms of risk aversion. However, unlike our paper, the incentive consequence of risk-averse *traders* is not addressed.

Several prediction market mechanisms are extensions of proper scoring rules. The notion of scoring rules was introduced by Brier [3], in the form of the quadratic scoring rule (which is proper), to measure the accuracy of weather forecasters. *Proper* scoring rules provide a way to reward forecasters such that honest reports are made. Most of the early work on scoring rules assumed that forecasters were risk neutral.

There has been some research on addressing risk aversion in scoring rules. Winkler and Murphy [12] showed that, if forecasters have a *known* risk type, scoring rules can be transformed to recapture the honest reporting property. One approach to handling

forecasters with unknown risk type, as proposed by Chen *et al.* [4] and Offerman *et al.* [10], is to figure out every participant's risk type by asking them a series of questions, and then calibrate their future reports using this data. This mechanism may work for a prediction market mechanism if the group of traders can be pre-screened. However, this may not be the case, and ideally we would like to have an "online" mechanism that can handle traders regardless of their risk type without any calibration. Allen [2] proposed one such "online" scoring rule for forecasters with arbitrary risk type. Our mechanism is based on Allen's result, and we discuss this idea in section 3.

2 Model, Notation, and Definitions

In this section and section 4 we consider a class of mechanisms defined by a set of properties satisfied by most subsidized prediction markets described in literature. We do not claim that the outlined properties are sufficient to completely characterize the space of prediction market mechanisms; rather, they identify a class of broad market-like mechanisms. We first describe our basic model of the information and interaction setting in which the mechanism operates, and then list the properties that the mechanisms we study must satisfy.

We consider a class of mechanisms designed to aggregate information from a set of *agents* (or *traders*) in order to forecast the outcome of a future event ω . Each agent i receives a private information signal, s_i , relevant to the outcome of the event; we assume s_i is binary, as is ω .

A central feature of the market-like mechanisms we consider is that the agents express a predicted probability of the event in the mechanism, through a sequence of public trades or *reports*. Other agents can update their beliefs based on the observed history of reports. We use $r_k \in [0, 1]$ to denote the k th report made in the market, and let $\mu_k = (r_1, \dots, r_{k-1})$ denote the history up to the start of the k th trade. r_k can thus depend on μ_k as well as any private information available to the trader making the report. We let n denote the total number of trades in the market.

We assume the identity of the agents making the reports cannot be verified, and the total number of agents participating is unknown. As each agent's signal is static, there is no need for any agent to trade more than once. Therefore, we will treat each report as if were made by separate traders and is natural for a market setting. However, an agent may masquerade as multiple agents, which is a consideration of sybilproofness.

Once the true outcome of the event is realized, $\omega = 1$ if the event occurs and $\omega = 0$ otherwise, the mechanism determines the reward for every agent. The reward for agent i , $\rho(r_i, \mu_i, n, \omega)$, is a function of the agent's report, market state, the total number of agents participating in the mechanism, and the event outcome. We allow the mechanism to randomize the distribution of the rewards, and we propose one such mechanism in section 3. We assume that the reward does not depend on the value of any reports made in the future. This is a nontrivial technical assumption that enables us to simplify the analysis of agents' myopic strategies, as agents can make decisions based on their current beliefs about the outcome, without forming beliefs about future agents' signals and strategies. This assumption is satisfied by most securities markets as well as market scoring rule markets, but not necessarily true for pari-mutuel markets.

Every agent i values the distributed reward according to her value function $V_i(\cdot)$, where $V_i(\cdot)$ is a weakly monotone increasing concave function. We make the normalizing assumption that $V_i(0) = 0$. In order to make her report, an agent maximizes her expected reward, with respect to her true belief p_i , over the outcome of the event and any randomization of the mechanism over the rewards, written as $E_{\omega \sim p_i} V_i(\rho(r, \mu_i, n, \omega))$. Though there may be other sources of uncertainty in the mechanism, we do not consider them in our model.

We identify the properties mechanisms should satisfy by examining literature in prediction market design and other wagering mechanisms. Hanson [8], in introducing the market scoring rule, had a subsidized prediction market be *myopically strategy proof* and have *bounded market subsidy*, both defined below. The same properties also hold in the dynamic pari-mutuel market introduced by Pennock [11]. As both of the mechanisms were subsidized, both had *guaranteed liquidity* by having a market maker that is always willing to trade with an agent. Prediction markets provide *anonymity*, i.e., the reward given due to a report is independent of who made the report. Finally, prediction markets are *sybilproof*, meaning that an agent reporting once with some information is no better off reporting twice in the market with the exact same information. Though anonymity and sybilproofness were not explicitly stated by Hanson or Pennock, they still hold in their proposed mechanisms and were explicitly defined by Lambert *et al.* [9]. We use a relaxed version of sybilproofness by requiring agents to be no better off reporting twice, as opposed to having the same payoff as presented by Lambert *et al.*. Using the notation established above, we formally define the desired properties:

P1: Myopically Strategyproof: If an agent making trade i has true belief p_i , and trades only once in a market, her reported belief will be her true belief. Mathematically,

$$\forall n, i \in \{1..n\}, \mu_i p_i = \operatorname{argmax}_{r \in [0,1]} E_{\omega \sim p_i} V_i(\rho(r, \mu_i, n, \omega)). \quad (1)$$

Further, we also require $\max_{r \in [0,1]} E_{\omega \sim p_i} V_i(\rho(r, \mu_i, n, \omega)) \geq 0$ so that myopic strategyproofness includes a standard individual rationality condition.

P2: Sybilproofness: An agent is no worse off reporting once honestly than making any two consecutive reports $r^{(1)}, r^{(2)}$ with the same information. Mathematically,

$$\forall n, i \in \{1..n\}, \mu_i E_{\omega \sim p_i} V_i(\rho(p_i, \mu_i, n, \omega)) \geq E_{\omega \sim p_i} V_i(\rho(r^{(1)}, \mu_i, n+1, \omega) + \rho(r^{(2)}, \mu_{i+1}, n+1, \omega)). \quad (2)$$

P3: Bounded Subsidy: The expected subsidy the market maker needs to invest into the market is bounded by a value β :

$$\forall n, i \in \{1..n\}, \mu_i, r_i \sum_{\text{all players } i} E_{\omega} \rho(r_i, \mu_i, n, \omega) < \beta$$

To summarize, we define the class of market-like mechanisms to be all mechanisms that are *anonymous*, *guarantee liquidity*, *myopically strategy proof*, *sybilproof*, and have *bound market subsidy*.

Before we introduce our results, we must introduce the concepts of information structure, report informativeness, and normalized expected reward.

Information Structure: We define an *information structure* to consist of a set of possible signal realizations for each trader, and the posterior probability of events given a subset of signal realizations (equivalently, the joint probability of signal realizations and the true outcomes).

Informativeness: For a given information structure, we define *informativeness* of an agent k , given a history μ_k , as the expected reduction in forecasting error, as measured by the reduction in entropy of the event, after conditioning on k 's signal.

Normalized Expected Reward: The informativeness and the reward of each agent may deviate. Therefore, in order to compare the reward an agent receives from a report, we define the *normalized expected reward* as the ratio of the expected reward to the informativeness of the report given the history up to that point.

3 Proposed Mechanism

In this section we review the work presented by Allen [2] and then present one mechanism that satisfies the properties outlined in section 2.

Allen shows that an agent with a monotone value function $V(\cdot)$ with unknown risk preference will set her report, \hat{p} , to her true belief, p , on an outcome ω if she participates in a sweepstakes. According to the sweepstakes, the agent will receive a reward of 1 with probability $q(\hat{p}) = 1 - (1 - \hat{p})^2$ if the event occurs and probability $\hat{q}(\hat{p}) = 1 - \hat{p}^2$ if the event does not occur. Allen's result follows from the fact that the expected value is linear in probabilities and the value function is monotonically increasing.

Now consider the following serial sweepstakes that is a derivative of Allen's result:

1. An agent observes the previous agents' reports, and plays an individual sweepstake as defined by Allen with sweepstake functions described above.
2. The outcome of the event is observed.
3. If there are n reports in the mechanism, then each player reporting \hat{p} wins a reward of 1 with probability $q(\hat{p}) = \frac{1}{4^n}(1 - (1 - \hat{p})^2)$ if the event occurs and $\hat{q}(\hat{p}) = \frac{1}{4^n}(1 - \hat{p}^2)$ if the event does not occur.

The mechanism above possesses all of the properties outlined in section 2; however, the mechanism distributes rewards that are exponentially decreasing with the number of agents. Moreover, if all the reports are equally informative, then the normalized expected reward also decreases exponentially with the number of reports.

4 Impossibility Result

Theorem 1. *If an anonymous, guaranteed liquidity mechanism satisfies properties P1–P3, then, there is a family of information structures $I^{(n)}$, each parameterized by a number n of agents, such that, even if all agents perform perfect Bayesian updating according to the structure $I^{(n)}$ and report their posteriors honestly, the minimum normalized expected reward of an agent must decrease exponentially with n .*

We start by showing that if agents with arbitrary risk-averse preferences are to participate in a mechanism in our class, all rewards must be non-negative. We then observe that the informativeness of a report is a constant multiple of the square of the differences between the posteriors after every report, so long as the posteriors are bounded in $[0.5 - c, 0.5 + c]$, for $c \leq 0.2$. We then show that for any two sequential reports made under two different posterior beliefs, the expected reward from the reports under the first posterior is a constant multiple of the expected reward under the second posterior, so long as the two posterior beliefs are bounded within $[0.5 - c, 0.5 + c]$, $c \leq 0.2$.

The theorem proof follows by inductively building a family of information structures starting with the structure of two agents both making symmetric reports. In the base structure $I^{(2)}$, two agents start with a common prior of 0.5 on the event and receive a binary signal. If the agents report honestly, each agent will change the posterior report by $\pm \frac{c}{2}$. This results in information structure $I^{(2)}$ with bounded posteriors between $[0.5 - c, 0.5 + c]$. We now consider a sybil attack in this setting, where one of the agents is making two reports under the same priors. To make this consideration we construct $I^{(3)}$. This construction is dependent on the mechanism: If the expected reward of the first agent is larger than 4 times the expected reward of the second, we construct a structure $I^{(3')}$, otherwise $I^{(3'')}$.

In either case, we split one of the reports into two, such that the histories up to the split report are consistent with $I^{(2)}$. By the sybilproofness property, we know that the sum of the expected rewards of the split reports is no larger than the original. Even accounting for the reduction in informativeness, we show that there is a report with normalized expected reward of at most γ times that of the split report in $I^{(2)}$, where $\gamma = \frac{0.5+c}{0.5-c} \cdot \frac{80}{81} < 1$ for a suitable c .

From $I^{(3)}$ we construct $I^{(4)}$ in a similar manner. Iteratively applying this procedure we show that there exists at least one report in $I^{(n)}$ with normalized expected reward that is exponentially smaller than one of the two reports in $I^{(2)}$.

5 Conclusion

In this paper we present one mechanism that satisfies properties in section 2 that allows agents with arbitrary risk-averse value function to participate. However, this mechanism requires that the normalized expected reward exponentially decrease with the number of agents. We then show that as long as the risk aversion structure of the agents is not known, for any mechanism in the class of interest that allows agents with arbitrary risk-averse value functions to participate, the normalized expected reward must decrease exponentially with the number of agents.

References

- [1] Agrawal, S., Delage, E., Peters, M., Wang, Z., Ye, Y.: A unified framework for dynamic parimutuel information market design. In: EC 2009: Proceedings of the tenth ACM conference on Electronic commerce, pp. 255–264. ACM, New York (2009)
- [2] Allen, F.: Discovering personal probabilities when utility functions are unknown. *Management Science* 33(4), 542–544 (1987)

- [3] Brier, G.W.: Verification of forecasts expressed in terms of probability. *Monthly Weather Review* 78(1), 1–3 (1950)
- [4] Chen, K.-Y., Fine, L.R., Huberman, B.A.: Forecasting uncertain events with small groups. In: *Proceedings of the Third ACM Conference on Electronic Commerce (EC 2001)*, Tampa, Florida, pp. 58–64 (2001)
- [5] Chen, Y., Pennock, D.M.: A utility framework for bounded-loss market makers. In: *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence (UAI 2007)*, pp. 49–56 (2007)
- [6] Chen, Y., Dimitrov, S., Sami, R., Reeves, D., Pennock, D., Hanson, R., Fortnow, L., Gonen, R.: Gaming prediction markets: Equilibrium strategies with a market maker. *Algorithmica*
- [7] Dimitrov, S., Sami, R., Epelman, M.: Subsidized prediction markets for risk averse traders (2009) (manuscript), <http://www.umich.edu/~rsami/papers/RiskAverseMarkets.pdf>
- [8] Hanson, R.: Combinatorial information market design. *Information Systems Frontiers* 5(1), 107–119 (2003)
- [9] Lambert, N.S., Langford, J., Wortman, J., Chen, Y., Reeves, D., Shoham, Y., Pennock, D.M.: Self-financed wagering mechanisms for forecasting. In: *EC 2008: Proceedings of the 9th ACM conference on Electronic commerce*, pp. 170–179. ACM, New York (2008)
- [10] Offerman, T., Sonnemans, J., van de Kuilen, G., Wakker., P.P.: A truth-serum for non-bayesians: Correcting proper scoring rules for risk attitudes. Working paper (June 2008), <http://www1.fee.uva.nl/creed/pdf/files/propscr30jun08.pdf>
- [11] Pennock, D.M.: A dynamic pari-mutuel market for hedging, wagering, and information aggregation. In: *EC 2004: Proceedings of the 5th ACM conference on Electronic commerce*, pp. 170–179. ACM Press, New York (2004)
- [12] Winkler, R.L., Murphy, A.H.: Nonlinear utility and the probability score. *Journal of Applied Meteorology* 9, 143–148 (1970)