

Implementing Securities Based Decision Markets with Stochastic Decision Rules

Doctoral Consortium

Wenlong Wang
 Massey University
 Auckland, New Zealand
 W.Wang1@massey.ac.nz

ABSTRACT

Incentivised decision markets are mechanisms that allow selecting one action among a set of actions based on properly incentivised forecasts about the actions’ consequences. Existing research on decision markets is based on scoring rules. Because scoring rules based decision markets involve two-side liabilities that can be difficult to track, we here study more convenient securities based decision markets. We present a decision market setting that prices securities using a cost function derived from a scoring rule. In such a decision market setting, traders will have the same expected utility as forecasters measured by the corresponding scoring rule. Moreover, we identify differences between scoring rules based decision markets and securities based decision markets in terms of actual payoffs. Lastly, we describe an insurance mechanism that can shift risk from the market creator to a risk-neutral third party.

KEYWORDS

prediction markets; decision markets; securities based implementation

ACM Reference Format:

Wenlong Wang. 2020. Implementing Securities Based Decision Markets with Stochastic Decision Rules. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020*, IFAAMAS, 3 pages.

1 INTRODUCTION

1.1 Problem Description

Prediction markets are popular tools for aggregating distributed information into often highly accurate forecasts. Participants in prediction markets trade contracts with payoffs tied to the outcome of future events. The pricing of these contracts reflects aggregated information about the probabilities associated with the possible outcomes. A frequently used contract type is Arrow-Debreu securities that pay \$1 when a particular outcome is realised and otherwise pay \$0. If such security is traded at \$0.30, this can be interpreted as a forecast for that outcome to occur at 30% chance. Potential caveats with the interpretation of prices in prediction markets as probabilities have been discussed in the literature [11], but are not seen as critical for typical applications [9, 11].

In many practical prediction markets applications, such as recreational markets on political events, participants trade directly with

each other, and one participant’s gain is the other participant’s loss. Prediction markets can, however, also be designed to offer net benefits to the participants. Such **incentivised prediction markets** can be used by an agent who is willing to compensate the market participants for the information obtained from the market [3, 9, 10]. Incentivised prediction markets rely on market maker algorithms to trade with the participants, and on cost functions [7] to update prices based on past transactions. These cost functions are closely related to proper scoring rules such as the Brier (or quadratic) scoring rule and the logarithmic scoring rule [1, 8], which measure the accuracy of forecasts and allow rewarding a single expert based on the forecast and actual outcome. The market maker in an incentivised prediction market subsidises the entire market rather than single experts; its maximal loss is finite and its expected loss depends on how much the participants ‘improve’ on the information entailed by the initial market maker pricing [9].

Accurate forecasts, as obtained from prediction markets, can be of tremendous value for decision makers. Commercial companies, for instance, can benefit substantially from accurate forecasts regarding the future demand for their products. However, many decision-making problems require conditional forecasts [6]. To decide, for instance, between alternative marketing campaigns, a company needs to understand how each of the alternatives will affect sales. In other words, it needs to predict, and choose between, “alternative futures”. Finding mechanisms that properly incentivise participants to provide their information for such conditional forecasts is non-trivial but can be achieved in **decision markets** [2, 4–6, 12].

Properly incentivised decision markets work in a stepwise process to select one among several mutually exclusive actions. First, forecasts about the expected future consequences of each action are elicited in a step analogous to incentivised prediction markets. Second, a decision rule is used to select an action based on the forecasted consequences. Once an action has been selected, and its consequences are revealed, payoffs are provided for the forecasts as elicited in the first step. Importantly, the decision rule in properly incentivised decision markets is stochastic, with each action being picked with a strictly positive probability [5, 6]. Payoffs can be adjusted to ensure that the participants’ expected payoffs in decision markets remain analogous to those made in properly incentivised prediction markets [5, 6], and that game-theoretical results on strategic interactions between participants in prediction markets [3] carry over.

The objective of this work is to identify and resolve potential obstacles to the implementation of a decision market platform where market creators can create decision markets to harness the wisdom

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand. © 2020 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

of the crowd for their decision making problems. We particularly focus on a securities based implementation of decision markets.

2 IN PROGRESS AND FUTURE WORK

2.1 Securities based Decision Markets with Stochastic Decision Rules

It is well established that the incentives provided by market scoring rules are equivalent to scoring rules that are typically used to incentivised single experts. Market scoring rules can be seen as scoring rules that are applied sequentially to experts, and are updated after each report. The updating is done such that experts improve previous reports and the market creator only pays the difference in the score between initial forecast and the final forecast [9, 10].

For implementation, securities trading is often seen as advantageous [7]. This is because it reduces the complexity of liabilities of forecasters and also is more natural for forecasters with experience in recreational prediction markets or financial markets. Liabilities in a scoring rules based framework are complex to track because they are two-sided. The decision maker is liable to pay all forecasters who improve the forecasts, and forecasters are liable to pay the market creator if they worsen the forecasts. Such a two-sided liability is not preferred in practical implementations.

In our setting, the market creator has a finite set of actions $\mathcal{A} = \{\alpha_1, \alpha_2, \dots, \alpha_m\}$ to choose from. For each action α_j , there is a conditional market which has a set of possible outcomes Ω_j . Both action set \mathcal{A} and outcome sets Ω_j are collectively exhaustive and mutually exclusive. We denote the outcome with index i for action α_j as ω_i^j and corresponding securities as q_i^j . Let $C(\vec{q}_j)$ be the cost function of a conditional market in a securities based decision market, which computes the total amount spent for purchasing outstanding securities \vec{q}_j . Let $\vec{\phi} = (\phi_1, \phi_2, \dots, \phi_m)$ be the decision rule distribution and we have $\sum_{i=1}^m \phi_i = 1$. Assume the market being resolved by paying off $\$1/\phi_j$ per share for the outstanding securities q_i^j which represents the observed outcome ω_i^j and $\$0$ for any other securities (including securities in conditional markets which corresponding action is not selected).

Assume a trader in our securities based decision market changes the outstanding securities distribution from \vec{q}_k to $^* \vec{q}_k$ for all α_k . Then the profit she gains from such a trade is denoted as \hat{S}_i^j and given by

$$\hat{S}_i^j = \frac{1}{\phi_j} (*q_i^j - q_i^j) - \sum_{k=1}^m (C(*\vec{q}_k) - C(\vec{q}_k)) \quad (1)$$

where α_j is the selected action and the observed outcome is ω_i^j .

In our work, we prove that a trader in above securities based market setting have the same expected utility as a forecaster in a scoring rules based framework from which the cost function has been derived. In terms of actual payoff, a comparison between the two types of decision market implementation is made. Furthermore, we compare worst-case loss between the two types of implementation. The market creator for securities based decision markets appears to be exposed to a bound-less worst-case loss because the loss has an additional dependence on the quantities of outstanding securities. We argue this is an inevitable trade-off for implementing a decision market in practice.

2.2 Worst-case Loss and Insurance

In prediction markets with market scoring rules, the worst-case loss for the market creator has been shown to be finite and is known ex-ante. In decision markets, this is no longer the case. In this work, we study worst-case loss in securities based decision market implementations and present a preliminary approach that allows the market creator to use insurance to limit worst-case losses.

The worst-case losses are even worse in the securities based decision markets as the expected utility of the traders is depending on the quantity of outstanding securities as well as the decision rule. Assume a risk-neutral insurer provide an insurance contract that charges I and pays I/ϕ with probability ϕ . Theoretically, the market creator can insure her decision market with a carefully selected I to hedge against the risk that she is exposed. Note that the cost function $C(\vec{q}_j)$ represents the amount spent for purchasing securities \vec{q}_j in conditional market corresponded to action α_j . The market creator can use this amount to purchase the insurance by setting $I_j = C(\vec{q}_j)$ for every action α_j . Then the worst-case loss no longer depends on the quantity of outstanding securities.

3 CONCLUSION

Previous studies of decision markets have not focused on security based implementations. To address this gap, we present a decision market setting that allows traders to earn an identical expected utility as forecasters who are measured by the corresponding scoring rule. Furthermore, we compare the actual payoffs between the two types of market and find that the worst-case loss of securities based market creator depends on outstanding securities. Lastly, we characterise a potential insurance solution to the bound-less worst-case loss problem which will be addressed in future.

ACKNOWLEDGMENTS

This work is supported by the Marsden fund (grant MAU1710).

REFERENCES

- [1] J. Eric Bickel. 2007. Some Comparisons among Quadratic, Spherical, and Logarithmic Scoring Rules. *Decision Analysis* 4, 2 (6 2007), 49–65. <https://doi.org/10.1287/deca.1070.0089>
- [2] Craig E. Boutilier. 2012. Eliciting forecasts from self-interested experts: Scoring rules for decision makers. In *11th International Conference on Autonomous Agents and Multiagent Systems 2012, AAMAS 2012: Innovative Applications Track*, Vol. 2. International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS), Department of Computer Science, University of Toronto, Toronto, Canada, 1008–1015.
- [3] Yiling Chen, Stanko Dimitrov, Rahul Sami, Daniel M. Reeves, David M. Pennock, Robin D. Hanson, Lance Fortnow, and Rica Gonen. 2010. Gaming prediction markets: Equilibrium strategies with a market maker. *Algorithmica (New York)* 58, 4 (2010), 930–969. <https://doi.org/10.1007/s00453-009-9323-2>
- [4] Yiling Chen, Ian Kash, Mike Ruberry, and Victor Shnayder. 2011. Decision markets with good incentives. (2011), 72–83 pages. <https://doi.org/10.1007/978-3-642-25510-6-7>
- [5] Yiling Chen and Ian A. Kash. 2011. Information elicitation for decision making. In *10th International Conference on Autonomous Agents and Multiagent Systems 2011, AAMAS 2011*, Vol. 1. International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS), School of Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138, United States, 161–168.
- [6] Yiling Chen, Ian A. Kash, Michael Ruberry, and Victor Shnayder. 2014. Eliciting predictions and recommendations for decision making. *ACM Transactions on Economics and Computation* 2, 2 (6 2014), 1–27. <https://doi.org/10.1145/2556271>
- [7] Yiling Chen and David M. Pennock. 2007. A utility framework for bounded-loss market makers. In *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence, UAI 2007*. Conference on Uncertainty in Artificial Intelligence, Yahoo Research, 45 W. 18th St., New York, NY, United States, 49–56.

- [8] Tilmann Gneiting and Adrian E Raftery. 2007. Strictly Proper Scoring Rules, Prediction, and Estimation. *J. Amer. Statist. Assoc.* 102, 477 (3 2007), 359–378. <https://doi.org/10.1198/016214506000001437>
- [9] Robin D. Hanson. 2003. Combinatorial information market design. *Information Systems Frontiers* 5, 1 (2003), 107–119. <https://doi.org/10.1023/A:1022058209073>
- [10] Robin D. Hanson. 2007. Logarithmic market scoring rules for modular combinatorial information aggregation. *George Mason University* 1, 1 (2007), 3–15. <https://doi.org/10.5750/jpm.v1i1.417>
- [11] C F Manski. 2006. Interpreting the predictions of prediction markets. *Economics Letters* 91, 3 (2006), 425–429. <https://doi.org/10.1016/j.econlet.2006.01.004>
- [12] Abraham Othman and Tuomas Sandholm. 2010. Decision rules and decision markets. In *Proceedings of the International Joint Conference on Autonomous Agents and Multiagent Systems, AAMAS*, Vol. 1. International Foundation for Autonomous Agents and Multiagent Systems (IFAAMAS), 625–632.