

Combinatorial Prediction Markets for Fusing Information from Distributed Experts and Models

Kathryn Blackmond Laskey
SEOR Department
George Mason University
Fairfax, VA, U.S.A.
klaskey@gmu.edu

Robin Hanson
Economics Department
George Mason University
Fairfax, VA, U.S.A.
rhanson@gmu.edu

Charles Twardy
C4I Center
George Mason University
Fairfax, VA, U.S.A.
ctwardy@gmu.edu

Abstract—Markets are a medium for information exchange between buyers and sellers. Prediction markets exploit the information transmission property of markets to improve forecasts of future events. Participants in a prediction market buy and sell assets that pay off if the underlying event occurs. Prices in a prediction market can be interpreted as consensus probabilities for the underlying events. Prediction markets thus provide a promising method to fuse information from a collection of human forecasters and/or computer forecasting algorithms. This paper describes the use of prediction markets as an information fusion mechanism, introduces combinatorial prediction markets as a mechanism for expressing and exploiting dependencies among base events, presents algorithms for performing basic market computations in combinatorial prediction markets, and introduces the SciCast combinatorial prediction market, a combinatorial prediction market for science and technology forecasting available at <http://scicast.org>.

Keywords—Prediction markets, Business forecasting, Combining expert judgment, Bayesian networks

I. INTRODUCTION

Markets are the dominant mechanism by which modern society allocates goods and services. Markets are also a medium for information exchange. In an efficient market, prices transmit information signals in a distributed fashion among market participants, encouraging expansion or contraction of production in response to demand, rationing of scarce resources, and elimination of surplus. Futures markets allow buyers and sellers to enter into contracts at the present time for exchanges that will occur at a specified future time. Prices in a futures market carry information about the anticipated future value of goods and services. These prices reflect participants' knowledge about factors that are likely to affect prices.

Improved accuracy in predictions can provide major competitive advantage to businesses and investors. Prediction markets exploit the information transmission property of

markets to improve predictions. A prediction market is a market in event futures formed for the purpose of predicting the underlying events [1-2]. Participants in a prediction market buy and sell assets that pay off if the underlying event occurs. Prices of assets can be interpreted as probabilities for the underlying events. In an efficient prediction market, prices aggregate the information of market predictions to make a more informed prediction than the individual participants could make on their own. Research has shown that aggregating judgments of a collection of forecasters improves accuracy of predictions [3-4], and that prediction markets are among the most accurate information aggregation methods [1]. Prediction markets can also be used as a mechanism for combining machine forecasters, often out-performing top ensemble forecasting methods [5]. The market mechanism provides a way to aggregate specialized algorithms, each participating only on problem instances for which it performs well. There is strong evidence for the advantages of combining human and machine forecasts [6]. Prediction markets provide a promising method for such combination.

This paper describes the use of prediction markets as an information fusion mechanism, introduces combinatorial prediction markets as a mechanism for expressing and exploiting dependencies among base events, presents algorithms for performing basic market computations in combinatorial prediction markets, and introduces the SciCast combinatorial prediction market, a combinatorial prediction market for science and technology forecasting available at <http://scicast.org>.

II. COMBINATORIAL PREDICTION MARKETS

Prediction markets allow forecasters to trade assets contingent on the occurrence of events of interest, yielding probabilistic forecasts in the form of event prices. For example, if an asset pays \$1 contingent on occurrence of event E , the price p of the asset in cents can be interpreted as the market probability of E . A forecaster whose probability for E is higher than the market price p expects to gain by purchasing the asset. The market mechanism tends to increase the assets of accurate forecasters and drive the assets of the least accurate forecasters toward zero. Those with greater forecasting ability will tend to have more assets to invest, and so will have more influence on market prices. In this way, the market naturally increases the influence of good forecasters. The market also encourages traders to self-select to forecast on topics about which they

This research was supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior National Business Center contract number D11PC20062. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoI/NBC, or the U.S. Government.

know the most. At any time, the market probability p of an event E can be interpreted as an aggregation of the collective information of market participants about E .

A. Logarithmic Market Scoring Rule

Markets require a way to match buyers with sellers. The double auction approach common to many prediction markets, in which buyers list bid prices and sellers list ask prices for assets they wish to trade, can result in illiquid markets in which traders with information have no incentive to trade [7]. Hanson [8] suggested the use of automated market makers to address this problem. An automated market maker will offer to buy or sell any asset on any relevant event at a price that depends on the current market probability and the quantity of the asset to be bought or sold. Trades with the automated market maker change the market probability. Thus, there are two ways to view an interaction with a market maker: as buying or selling a contingent asset with consequent change in the market probability, or as editing the market probability with consequent change in one's assets contingent on whether the event occurs. The two views – trading shares or editing probabilities – are formally equivalent. For this reason, we use the terms edit and trade interchangeably.

The most natural and commonly used market scoring rule is the logarithmic market scoring rule, or LMSR. In an LMSR market, a user with current assets a who changes the market probability of event E from p to q will have assets

$$\begin{aligned} a + b \ln\left(\frac{q}{p}\right) & \quad \text{if } E \text{ occurs} \\ a + b \ln\left(\frac{1-q}{1-p}\right) & \quad \text{if } E \text{ does not occur.} \end{aligned} \tag{1}$$

Thus, a forecaster will gain if she increases the probability of an event that occurs or decreases the probability of an event that does not occur. Furthermore, a forecaster who disagrees with the current market forecast has incentive to trade, in that her expected assets increase if she moves the probability in the direction of her beliefs. Thus, a LMSR based market maker addresses the illiquidity problem in prediction markets. Forecasters are incentivized to jointly form consensus market probabilities that aggregate their collective information.

Liquidity does come at a cost -- the market maker must be willing to subsidize the market. However, the loss to the market maker is bounded by an amount proportional to the entry of the market distribution set by the market maker at the start of trading.

B. Combinatorial Prediction Markets

Questions in a prediction market may be related to each other. For example, consider the question: *When will the first reported collision occur between an autonomous car and a human driven vehicle on a public road?* A forecaster might judge that the answer depends on the answer to the question: *When will a fully autonomous (self-driving) car made by a major auto maker drive more than 60 miles in the United States?* In a combinatorial prediction market, forecasters can link these questions, giving different probabilities for the first

question conditional on different answers to the second question.

An attractive feature of a LMSR based combinatorial market is that conditional forecasts satisfy intuitive independence properties: a trader who increases $p(A|B)$ gains if B and A both occur; loses if B occurs but A does not; and neither gains nor loses if B does not occur [9].

Because the number of event combinations is exponential in the number of base events, computing market prices becomes intractable when there are very many base events. For this reason, factored representations such as Bayesian networks have been used as a tractable representation for the market joint probability distribution [10-13].

An advantage of the combined probability and asset management approach presented in [12] is the ability to reuse assets. Previous proposals for asset management in combinatorial markets required a forecaster to make an upfront payment for the right to purchase a contingent asset that pays off if a certain event happens. No trades are allowed if the forecaster runs out of assets. In a combinatorial market, however, this approach can tie up assets unnecessarily. Consider, for example, a forecaster who wishes to make a conditional trade on event A given mutually exclusive events B and C , where the current market probabilities are $p(A|B) = x$ and $p(A|C) = y$. This could be accomplished by setting up two conditional assets. The first conditional asset costs $\$x$, pays $\$1$ if A and B both occur, pays $\$0$ if B occurs but A does not, and refunds the purchase price $\$x$ if B does not occur. The second conditional asset costs $\$y$, and pays $\$1$ if A and C both occur, pays $\$0$ if C occurs but A does not, and refunds the purchase price $\$y$ if C does not occur. Because B and C are mutually exclusive, there is no risk of losing the entire amount, because at least one of the purchase prices will be refunded. If B occurs, the forecaster loses at most $\$x$ from the first asset and is refunded $\$y$ for the second asset, for a net loss of no more than $\$x$. If C occurs, the forecaster loses at most $\$y$ from the second asset and is refunded $\$x$ for the first asset, for a net loss of no more than $\$y$. If neither B nor C occurs, both purchase prices are refunded, and there is no loss. Therefore, the maximum loss is $\max\{x, y\}$. In a market with no asset reuse, the forecaster has to tie up the entire purchase price $\$x+y$ until the outcome of B or C is determined, even though no more than $\max\{x, y\}$ is being risked.

In a more complex scenario with many conditional trades, a considerable amount of assets could be unnecessarily tied up. It is straightforward in the above example to correct the problem and adjust the price of the assets, but pricing combinatorial assets is in general a difficult problem. The method of [12] and its more computationally efficient variant [13] support asset reuse. These methods would tie up no more than $\max\{x, y\}$ of the forecaster's assets to support this pair of trades, and in general would arrive at correct prices for many complex combinatorial trades.

C. Market Maker Functions

A market maker in a combinatorial prediction market needs to be able to perform the following functions:

1. Query to find the conditional probability $p(T | A)$ of target event T given conditioning events A .
2. Query to find the expected value $\sum_x p_x a_x^u$ of the user u 's assets, where the expectation is taken over all joint states of the combinatorial space of events in the market, and p_x and a_x^u are, respectively, the market probability and u 's assets in state x .
3. Check that allowing a trade by user u will not make u 's assets a_x^u negative in any state x .
4. Implement an edit by u to change $p(T | A)$ to $q(T | A)$.
5. If an event E occurs, set the probability of E to 1, pay off all users who made contingent trades involving E , and update the probabilities of all events depending on E . This operation is called *resolving* the event E .
6. Add new events to the market.

Sun et al. [12] described a method to perform these tasks when the consensus joint probability distribution can be represented in factored form as a junction tree and each edit is allowed to involve only variables that all belong to the same clique of the junction tree. This approach uses parallel junction trees for the joint distribution and for each user's assets. Task 1 is performed using the standard junction tree algorithm. Edits (Task 4) are represented as soft evidence in the junction tree. Sun et al. showed how to adapt the junction tree algorithm to compute expected assets (Task 2) and edit limits (Task 3). Resolving events (Task 5) involves conditioning on the corresponding random variables and then removing them from the junction tree. Adding new events (Task 6) and new links also involve changing the junction tree.

The parallel junction tree method of [12] maintains a separate junction tree for each participant with the same structure as the junction tree representing the market joint distribution. In the typical case in which most users trade sparsely relative to the full joint space, using the same global junction tree for all users is very inefficient for both storage and computation. The Dynamic Asset Cluster (DAC) asset management method [13] achieves much greater efficiency by

maintaining a separate trade-based asset structure for each user. A user's asset model consists of a set of asset blocks, where each asset block compactly represents a subset of trades, and the entire set represents all the trades the user has made. Expected assets can be calculated directly from the asset blocks and the probability junction tree. Edit limits and minimum assets are found by performing graph transformations to transform the set of asset blocks into an asset junction tree, and propagating minimum assets in the user-specific trade-based junction tree.

III. DAGGRE: GEOPOLITICAL FORECASTING

The DAGGRE geopolitical forecasting market [14] was a combinatorial prediction market that used the method of [12] for probability and asset management. Participants were recruited from email solicitations, articles on blogs and newspapers, and personal recruiting at professional events. Small financial incentives were given for participation, but to program restrictions precluded compensating participants for accuracy. Over the 20 months the DAGGRE market was active, more than 3000 forecasters participated, with an average of about 150 forecasts per week. Of the 400 questions on the market, about 200 were shared among five teams in an IARPA-funded forecasting tournament. All five teams reported forecasts daily to IARPA, and forecasts were scored against ground truth by averaging the daily Brier score [15] over the period of time the question was active. Although early DAGGRE results were unreliable due to software issues, from February 2012 through May 2013, the DAGGRE market accuracy was about 38% greater than the baseline system. The DAGGRE prediction market closed in June of 2013 and reopened in November 2013 with a shift in focus to science and technology forecasting.

IV. SCICAST: FORECASTING SCIENCE AND TECHNOLOGY INNOVATION

SciCast, launched in November 2013, is the largest and most advanced science and technology prediction market in the world. SciCast is a community-driven initiative that allows scientists, technologists, and technology watchers around the globe to forecast science and technology trends, and to contribute forecasting questions. SciCast participants can move

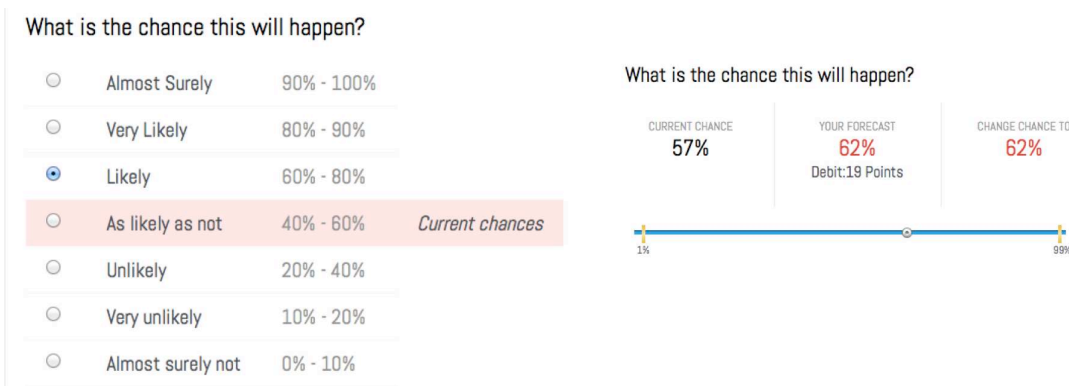


Figure 1. SciCast Interface for Editing Probabilities

1 If the answer to one of these related questions is:

Your Assumption	Question
<p>TIP: Forecast different scenarios using these assumptions. If an answer occurs, what will happen in the question asked below?</p> <p>False</p> <p>+ Add another assumption from our existing questions OR suggest a new related question</p> <p>START OVER</p>	<p>Will one of the top 5 insurance companies in California offer a liability insurance package for driverless vehicles cheaper than the equivalent package for a human driven car by an under 25 year-old male driver by July 1st, 2018?</p>

1 If the answer to one of these related questions is:

Your Assumption	Question
<p>TIP: Forecast different scenarios using these assumptions. If an answer occurs, what will happen in the question asked below?</p> <p>True</p> <p>+ Add another assumption from our existing questions OR suggest a new related question</p> <p>START OVER</p>	<p>Will one of the top 5 insurance companies in California offer a liability insurance package for driverless vehicles cheaper than the equivalent package for a human driven car by an under 25 year-old male driver by July 1st, 2018?</p>

2 Then: When will a fully autonomous (self-driving) car made by a major auto maker drive more than 60 miles in the United States?

POSSIBLE ANSWERS AND CURRENT CHANCE	MINIMUM AVAILABLE POINTS
Before January 1, 2015	2% 4,894
Between January 1, 2015 and June 30, 2015	14% YOUR PREVIOUS FORECASTS
Between July 1, 2015 and December 31, 2015	21%
Between January 1, 2016 and June 30, 2016	19%
Between July 1, 2016 and December 31, 2016	12%
Will not occur before January 1, 2017	32%

2 Then: When will a fully autonomous (self-driving) car made by a major auto maker drive more than 60 miles in the United States?

POSSIBLE ANSWERS AND CURRENT CHANCE	MINIMUM AVAILABLE POINTS
Before January 1, 2015	3% 4,894
Between January 1, 2015 and June 30, 2015	23% YOUR PREVIOUS FORECASTS
Between July 1, 2015 and December 31, 2015	32%
Between January 1, 2016 and June 30, 2016	24%
Between July 1, 2016 and December 31, 2016	7%
Will not occur before January 1, 2017	11%

Figure 2: SciCast allows users to make conditional forecasts

probabilities of a question up or down in accordance with their beliefs (see Figure 1). When the answer becomes known, participants win or lose points according to (1), gaining or losing depending on whether they move the probability in the right direction. In a friendly competition, participants vie for position on the leaderboard and engage in animated discussions of future technological trends. Participants can also propose questions to be included in the market.

SciCast users can make predictions on hundreds of questions on science and technology innovation. Questions can be binary (true/false), multiple-choice (a set of possible answers), or scaled continuous (with a continuous range of

values for which forecasters predict an expected value). Participants can enter forecasts either by using a survey-style interface to select a probability range (lefthand side of Figure 1) or by moving a slider to change the probability to a new value (righthand side of Figure 1). The former method moves the probability a small amount in the direction of the user's expressed belief. The latter method allows a direct edit of the probability and shows the user how many points are put at risk to make the change.

Once questions have been linked, forecasters can make conditional forecasts, as shown in Figure 2. The question, *When will a fully autonomous (self-driving) car made by a*

Trends

Trend graph data is updated nightly.

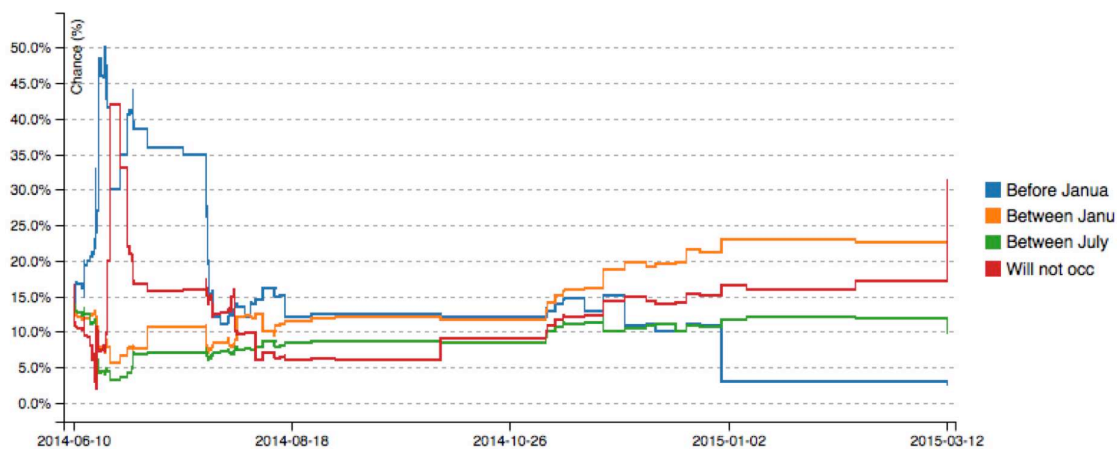


Figure 3: SciCast Users can View How Probabilities Have Changed over Time

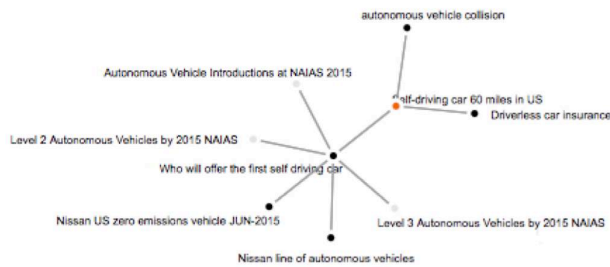


Figure 4: Linked Questions on Driverless Cars

major auto maker drive more than 60 miles in the United States? has been linked to the question, *Will one of the top 5 insurance companies in California offer a liability insurance package for driverless vehicles cheaper than the equivalent package for a human driven car by an under 25 year-old male driver by July 1st, 2016?* Therefore, a user who wants to make a forecast on the former question will be offered the opportunity to condition the forecast on the answer to the latter question. Figure 2 shows screenshots from SciCast depicting the market conditional distributions for the former question at given the two answers to the latter question. We see that

availability of liability insurance moves the market distribution toward earlier dates of introduction of driverless cars.

SciCast provides a graphical view of the evolution of probability distributions over time, as depicted in Figure 3. Users can update their forecasts at any time, to take account of new information such as news reports about driverless cars. The continually updated and reshaped information helps both the public and private sectors better monitor developments in a variety of industries. SciCast is a real-time indicator of what participants think is going to happen in the future. This information can provide valuable inputs to businesses and investors making decisions about products or investments.

SciCast offers users the opportunity to make conditional forecasts given the outcomes of events linked to a given question. Users may also add new links if they believe the probability distribution for a question should depend on the answer to another question. Clicking the link “+ Add another assumption from our existing questions,” in Figure 2 brings up a screen from which the user can browse existing questions and add one as an assumption for a forecast. In order to add a link, users must make a sufficiently large edit, i.e., the link will not be added if the probability distributions given different answers to the conditioning question do not differ by a sufficient amount.




Figure 5: The SciCast Question Network

SciCast users can also view all questions that are linked to a given question. Figure 4 shows all the questions linked to the driverless car question; Figure 5 shows the entire SciCast question linkage graph as of May, 2015.

SciCast has over 11,000 registered users, who have made over 120,000 forecasts on a total of 1188 questions. On average, there have been over 600 questions active on any given day, and forecasters have averaged about 240 forecasts per day. The formal evaluation against an external baseline ended with the transition from geopolitical to science and technology forecasting, and was replaced with a comparison against our own internal benchmark consisting of an unweighted average of forecasts made using the survey-style interface. SciCast has performed 18% better than this internal benchmark. The DAGGRE benchmark was also an unweighted average of forecasts, but unlike the forecasters providing forecasts for the DAGGRE benchmark, the SciCast forecasters providing survey-style edits could see the market estimates before making their forecasts. A preliminary accuracy study conducted during the summer of 2014 found strong effects of monetary incentives on participation, with large, statistically significant spikes in participation when monetary rewards were given for participation and/or accuracy. No statistically significant effect of incentives on accuracy was found. A larger follow-on study conducted from November 2014 through March of 2015 involved more questions and had larger financial incentives. This study found statistically significant gains in both activity and accuracy when monetary incentives for accuracy were offered.

V. A PLATFORM FOR RESEARCH

In addition to making forecasts, participating in discussions, and authoring questions, SciCast users can access forecast data from the SciCast data mart for their own research.

SciCast partner organizations include the American Association for the Advancement of Science (AAAS), the Institute of Electrical and Electronics Engineers (IEEE), the Information Systems Audit and Control Association, Inc. (ISACA), the American Medical Informatics Association (AMIA), and TechCast. By posing questions to SciCast, partner organizations are able to bring the SciCast technology to bear on questions of interest to their organizations.

The SciCast research team is conducting several research studies to help us learn how to use the wisdom of crowds to improve forecasts. These include studies comparing SciCast with quantitative forecasting methods, studies to assess the benefits of linking questions, studies examining the effect of including computerized traders in the SciCast prediction market, and studies of the effects of offering incentives for accuracy and participation. We are collaborating with researchers at DRDC Canada, and the TechCast project at George Washington University. The SciCast team is led by George Mason University and includes Inkling Markets, Gold Brand Software, Tuuyi, and KaDSci.

VI. CONCLUSION

SciCast aims to increase the accuracy of science and technology forecasts by tapping the collective intelligence of thousands of users around the world. Users spend points to

make forecasts, and gain more than they spend if they turn out to be right. Better forecasters gain more points and therefore more influence, improving system accuracy. SciCast allows forecasters to link questions and make conditional edits: almond yield can depend on honeybee collapse, or photovoltaic price-performance on the progress of multijunction arrays. At any time, SciCast provides a real-time assessment of the probability of any of its questions, given the information currently available to its participants.

SciCast topics include: agriculture, biology and medicine, chemistry, computational sciences, energy, engineered technologies, global change, information systems, mathematics, physics, science and technology business, social sciences, space sciences and transportation. Scientists, statisticians, entrepreneurs, policymakers, technical traders and futurists of all stripes are encouraged to improve our forecasts, link questions, and pose new questions. SciCast is accessible at <https://scicast.org>.

REFERENCES

- [1] Y. Chen and D. M. Pennock, "Designing Markets for Prediction." *AI Magazine* 31, no. 4, 2011, pp. 42–52.
- [2] K. Arrow, R. Forsythe, M. Gorham, R. Hahn, R. Hanson, J. O. Ledyard, S. Levmore, R. Litan, P. Milgrom, F. D. Nelson, G. R. Neumann, M. Ottaviani, T. C. Schelling, R. J. Shiller, V. L. Smith, E. Snowberg, C. R. Sunstein, P. C. Tetlock, P. E. Tetlock, H. R. Varian, J. Wolfers, and E. Zitzewitz, "The promise of prediction markets," *Science*, vol. 320, no. 5878, 2008, pp. 877–878.
- [3] Solomonoff, R. J., "Complexity-based induction systems: Comparisons and convergence theorems." *IEEE Transactions on Information Theory* vol. 24, 1978, pp. 422–432.
- [4] James Surowiecki. "The Wisdom of Crowds," 2005, Anchor.
- [5] Barbu, A. and Lay, N. (2011). "An introduction to artificial prediction markets for classification." *Journal of Machine Learning Research* vol. 13, 2012, pp. 2177-2204.
- [6] Nagar, Y. and Malone, T. (2011). Making business predictions by combining human and machine intelligence for making predictions. In *Proceedings of the Thirty Second International Conference on Information Systems (ICIS 2011)*, Shanghai. ICIS. Expanded from NIPS 2010.
- [7] Y. Chen, S. Dimitrov, R. Sami, D. Reeves, D. Pennock, R. Hanson, L. Fortnow, and R. Gonen, "Gaming prediction markets: Equilibrium strategies with a market maker," *Algorithmica*, vol. 58, no. 4, pp. 930–969, 2010. [Online]. Available: <http://dx.doi.org/10.1007/s00453-009-9323-2>
- [8] R. Hanson, "Combinatorial information market design," *Information Systems Frontiers*, vol. 5, no. 1, pp. 107–119, 2003.
- [9] R. Hanson, "Logarithmic market scoring rules for modular combinatorial information aggregation," *Journal of Prediction Markets*, vol. 1, no. 1, pp. 3–15, 2007.
- [10] Y. Chen, S. Goel, and D. M. Pennock, "Pricing combinatorial markets for tournaments," in *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC-2008)*, 2008, pp. 305–314.
- [11] D. Pennock and L. Xia, "Price updating in combinatorial prediction markets with Bayesian networks," in *Proceedings of the Twenty-Seventh Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-11)*. Corvallis, Oregon: AUAI Press, 2011, pp. 581–588.
- [12] W. Sun, R. Hanson, K. B. Laskey, and C. Twardy, "Probability and asset updating using Bayesian networks for combinatorial prediction markets," in *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence (UAI-12)*, Catalina, CA, 2012.

- [13] W. Sun, S. Matsumoto, R. Hanson, K. B. Laskey, C. Twardy, and B. Goldfeder “Trade-based asset models for combinatorial prediction markets,” in Proceedings of the UAI Bayesian Modeling Applications Workshop, 2014.
- [14] W. A. Powell, R. Hanson, K. B. Laskey, and C. Twardy, “Combinatorial prediction markets: An experimental study,” in Scalable Uncertainty Management, ser. Lecture Notes in Computer Science, W. Liu, V. Subrahmanian, and J. Wijsen, Eds. Berlin Heidelberg: Springer, 2013, vol. 8078, pp. 283–296.