Decision Markets with Good Incentives

Citation

Chen, Yiling, Ian Kash, Mike Ruberry, and Victor Shnayder. 2011. "Decision Markets with Good Incentives." Lecture Notes in Computer Science 7090: 72-83.

Published Version

doi:10.1007/978-3-642-25510-6_7

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:11693950>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Decision Markets With Good Incentives [★]

Yiling Chen, Ian Kash, Mike Ruberry and Victor Shnayder

Harvard University

Abstract. Decision markets both predict and decide the future. They allow experts to predict the effects of each of a set of possible actions, and after reviewing these predictions a decision maker selects an action to perform. When the future is independent of the market, strictly proper scoring rules myopically incentivize experts to predict consistent with their beliefs, but this is not generally true when a decision is to be made. When deciding, only predictions for the chosen action can be evaluated for their accuracy since the other predictions become counterfactuals. This limitation can make some actions more valuable than others for an expert, incentivizing the expert to mislead the decision maker. We construct and characterize decision markets that are – like prediction markets using strictly proper scoring rules – myopic incentive compatible. These markets require the decision maker always risk taking every available action, and reducing this risk increases the decision maker’s worst-case loss. We also show a correspondence between strictly proper decision markets and strictly proper sets of prediction markets, creating a formal connection between the incentives of prediction and decision markets.

1 Introduction

To make an informed decision a decision maker must understand the likely consequences of their actions. Hanson proposed a “decision market” to directly predict these consequences [11]. His proposal consists of a set of conditional prediction markets, one for each possible action. After the markets close the decision maker could evaluate each action’s predicted effect on the set of possible outcomes, and choose the most preferred action. Conditional markets for actions not taken are voided.

Consider, for example, a project manager deciding between two developers, A and B. The manager prefers to hire the candidate more likely to complete a project on time, so it runs two conditional prediction markets. One determines the likelihood A will finish on time, conditional on A being hired, and the latter does the same for B. If the project manager has access to knowledgeable experts and these markets reflect the experts’ information then the manager can make an informed hiring decision.

[★] This material is based upon work supported by NSF Grant No. CCF-0915016. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors alone.

Using a prediction market to make a decision is natural and previous work has demonstrated they can produce accurate forecasts [1,2,18,3,8]. However, while a prediction market using a strictly proper scoring rule is myopic incentive compatible, Hanson’s proposed decision market is not. That is, in a prediction market an expert maximizes its score for a prediction by predicting consistent with its beliefs, but the same is not true when a decision is to be made [16,6].

We return to our hypothetical project manager and the design of its two prediction markets. The manager would like to reward experts for improving either market’s accuracy, but only one market’s condition will ever be realized since only one developer will be hired. The other market’s predictions will become unscored counterfactuals. If an expert has improved one market’s prediction more than the other’s, it has an incentive to convince the project manager to hire the associated developer regardless of how poor an employee that developer may be!

More concretely, if the markets currently predict developer A has a 60% and developer B a 80% chance of finishing the project on time, and an expert believes the correct likelihoods are 70% and 80%, respectively, truthful reporting can only improve the market’s accuracy for developer A. If developer B is hired this expert will receive a score of zero, but if A is hired they expect to score for a 10% improvement. Instead of being honest, then, the expert can pretend B is incompetent, lowering the market’s prediction for the likelihood B will finish on time to less than 70%, cause A to be hired instead, and enjoy the profits.

We address this manipulation and construct and characterize decision markets that are myopic incentive compatible, like prediction markets. Instead of a scoring rule, these markets use a decision scoring rule that can account for the likelihood actions are taken when scoring predictions. When a decision maker risks taking an action at random, these decision scoring rules allow the scores of unlikely actions to be amplified while the scores of likely actions are comparatively reduced, making risk neutral experts indifferent to their affect on the decision. We show this risk of taking an action at random is a requirement for myopic incentive compatible decision markets, and reducing this risk increases the decision maker’s worst-case loss. We also show that, for risk-neutral experts, every myopic incentive compatible decision market describes a game equivalent to that described by a myopic incentive compatible set of prediction markets, creating a formal connection between decision and prediction markets.

The rest of the paper is organized as follows. Section 1.1 describes previous work on prediction and decision markets. Section 2 provides a formal description of prediction markets in our notation, and Section 3 describes our decision market model. Section 4 presents our construction and characterization results. Section 5 extends these results, describing optimal behavior for a risk neutral decision maker and connecting prediction and decision markets. Finally, Section 6 discusses further research challenges and concludes.

1.1 Related work

Strictly proper scoring rules have long been understood to be able to truthfully elicit a single risk-neutral expert’s beliefs over the outcome of an uncertain event [14,17,10]. Hanson [12,13] showed these same rules could be used to myopically incentivize any number of experts to be honest in a prediction market, and described an extension of scoring rules – market scoring rules – that prevent the market maker’s worst-case loss from growing with the number of experts. Importantly, all strictly proper scoring rules require the market maker correctly observe the event’s outcome. We formally describe these rules in Section 2.

When making a decision, some outcomes are not observed, and strictly proper scoring rules do not generally myopically incentivize an expert to be truthful. Othman and Sandholm [16] first formalized this incentive problem. They considered a single expert predicting and a decision maker picking their most preferred action based on the expert’s predictions, and they showed the expert can be incentivized to honestly reveal the decision maker’s most preferred action. They describe this decision rule as MAX—simply picking the best action from what’s available. Chen and Kash [6] also considered a single expert but allowed both deterministic and stochastic decision rules. Given a decision rule they characterized all scoring rules incentivizing a single risk-neutral expert to predict truthfully.

But while strictly proper scoring rules can be used for a single expert and extend to prediction markets, these scoring and decision pairs do not have a such a natural extension. In a prediction market, an expert’s expected score for a prediction is fixed once the prediction is made, and the same is true when a single expert is informing a decision. In a decision market, however, a prediction’s score may not be fixed until the market closes, creating new strategic complexities. In fact, Othman and Sandholm showed that no scoring rule can myopically incentivize experts to predict honestly in a decision market using their MAX rule, and we extend this result in Section 4. We also describe myopic incentive compatible decision markets for the first time.

Recently, manipulation in the presence of outside incentives has been studied [9,5]. In this paper we do not consider outside payoffs. The decision maker’s choice of action may affect an expert’s utility, but not because of any inherent preferences over actions that expert may have.

2 Prediction Markets: Background and Notation

This section presents the standard market scoring rule model of a prediction market, first described by Hanson [12,13], and defines our notation.

Let $\mathcal{O}$ be a finite, mutually exclusive, and exhaustive set of outcomes. A *prediction market* is a sequential game played by any number of risk-neutral, expected-value-maximizing *experts* predicting the likelihood of these outcomes. The market opens at round zero with some initial prediction $p^0 \in \Delta(\mathcal{O})$, where $\Delta(\mathcal{O})$ is the set of probability distribution over outcomes. At each round after the market opens, an arbitrarily chosen expert makes a prediction $p \in \Delta(\mathcal{O})$,

and we let p^t be the prediction made in round t . The market closes at some round $\bar{t}$, after which an outcome o^* is observed and experts are scored for each prediction by a *scoring rule*,

$$s : \mathcal{O} \times \Delta(\mathcal{O}) \rightarrow \mathbb{R} \cup \{-\infty\},$$

where $\mathbb{R}$ is the set of real numbers. We write $s_o(p) \equiv s(o, p)$ as a shorthand, and an expert's payment for a prediction is the difference between the scores of its and the immediately preceding prediction. Letting $\mathcal{T}$ be the set of rounds when an expert made a prediction, its total payoff is

$$\sum_{t \in \mathcal{T}} s_{o^*}(p^t) - s_{o^*}(p^{t-1}).$$

Markets with this sequential difference payoff structure are described as *market scoring rule* markets.

Scoring rules are *regular* when only predictions assigning zero likelihood to the observed event are scored negative infinity, and *proper* when a risk-neutral expert's expected score for a prediction is maximized when predicting consistent with its belief. Formally, a rule is proper if for all beliefs $q \in \Delta(\mathcal{O})$ over the likelihood of outcomes and predictions p

$$\sum_{o \in \mathcal{O}} q_o s_o(q) \geq \sum_{o \in \mathcal{O}} q_o s_o(p),$$

where q_o is the believed likelihood of outcome o . A rule is *strictly proper* when the inequality is strict unless $q = p$, uniquely maximizing an expert's score when they predict consistent with their beliefs. An example of a strictly proper scoring rule is $s_o(p) = a_o + b \log p_o$ with $a_o \in \mathbb{R}$ and $b > 0$. When experts uniquely maximize their score for a prediction by predicting consistent with their beliefs we describe the mechanism as *myopically incentive compatible*.

In aggregate, experts receive a payoff of $\sum_{t=1}^{\bar{t}} s_{o^*}(p^t) - s_{o^*}(p^{t-1}) = s_{o^*}(p^{\bar{t}}) - s_{o^*}(p^0)$, so the market institution's worst-case loss is

$$\max_{o^*, p^{\bar{t}}} s_{o^*}(p^0) - s_{o^*}(p^{\bar{t}}).$$

Note that the market institution's payment is bounded and independent of the number of experts. In practice, a market institution's budget must be at least their worst-case loss.

Ideally, the final prediction is an accurate consensus of experts' beliefs. Bayesian experts, for example, update their beliefs as they observe other's predictions. However, while a market using a strictly proper scoring rule is myopic incentive compatible, it is not incentive compatible in general. An expert participating in multiple rounds may provide a prediction inconsistent with its belief, with the intention to mislead other experts and later capitalize on their mistakes [4]. Despite such possible manipulations by forward-looking Bayesian experts, previous work has shown that under certain conditions prediction markets that are

myopic incentive compatible can fully aggregate information in finite rounds [4] or in the limit [15]. In this paper, however, we do not restrict experts to be Bayesian but allow arbitrary – or free – beliefs, as is typical when working with scoring rules.

3 Decision Market Model

A prediction market is a special case of a decision market. Both use the same sequential market structure, but a decision market uses a decision rule to pick from a set of actions before the outcome is observed, and which action is chosen may affect the likelihood an outcome occurs. Unlike previously proposed models of decision markets, we score experts using a decision scoring rule instead of a standard scoring rule. This more general function is necessary to recreate the myopic incentive compatibility of a prediction market for the broadest possible class of decision markets.

Let $\mathcal{A}$ be a finite set of actions, and $\mathcal{O}$ a set of outcomes as before. Without loss of generality and for notational convenience we assume the outcomes for every action are the same. As in a prediction market, a decision market opens with an initial prediction, but instead of a single probability distribution it is a set of conditional distributions, one for each action, denoted $P^0 \in \Delta(\mathcal{O})^{|\mathcal{A}|}$. Experts still report sequentially, and we let P^t be the prediction made in round t , P_a^t that prediction's distribution over outcomes given action a is chosen, and $P_{a,o}^t$ be that conditional distribution's likelihood for outcome o .

After the market closes, the decision maker selects an action using a *decision rule*

$$D : \Delta(\mathcal{O})^{|\mathcal{A}|} \rightarrow \Delta(\mathcal{A}),$$

applied to the final report $P^{\bar{t}}$, drawing an action a^* from $\mathcal{A}$ according to the distribution $D(P^{\bar{t}})$. We say that a decision rule has *full support* if it only maps to distributions with full support. As a shorthand we write d for a distribution over actions and d_a the likelihood action a is drawn from the set.

Once the action is selected, an outcome o^* is revealed, and experts are scored for each prediction by a *decision scoring rule*

$$S : \mathcal{A} \times \mathcal{O} \times \Delta(\mathcal{A}) \times \Delta(\mathcal{O})^{|\mathcal{A}|} \rightarrow \mathbb{R} \cup \{-\infty\},$$

written $S_{a,o}(d, P) \equiv S(a, o, d, P)$. Paralleling scoring rules, we describe decision scoring rules as *regular* when only predictions assigning zero likelihood to the observed event are scored negative infinity.

Letting $\mathcal{T}$ again be rounds where an expert made a prediction, its total payoff is

$$\sum_{t \in \mathcal{T}} S_{a^*, o^*}(d, P^t) - S_{a^*, o^*}(d, P^{t-1}),$$

so the market institution's worst-case loss is

$$\max_{P^{\bar{t}}, a \in \mathcal{A}, o \in \mathcal{O}} S_{a,o}(d, P^0) - S_{a,o}(d, P^{\bar{t}}), \quad (1)$$

where $\bar{\mathcal{A}}$ is the support of $D(P^i)$. Previous work on decision markets used a similar model, but with a *conditional scoring rule*

$$s_c : \mathcal{A} \times \mathcal{O} \times \Delta(\mathcal{O})^{|\mathcal{A}|} \rightarrow \mathbb{R} \cup \{-\infty\},$$

instead of a decision scoring rule.

As we show in the next section, however, considering the likelihood an action is selected is necessary to create the same myopic incentive compatibility as in prediction markets.

4 Decision Market Incentives

In a prediction market, a strictly proper scoring rule uniquely maximizes an expert's score for a prediction when they predict consistent with their beliefs. The same is not always true in a decision market. While both markets can reward improvements over a prior prediction, a decision market only observes and scores the improvement in the prediction for the chosen action. Since this action is a function of the market's final prediction, experts may have an incentive to change this prediction (either directly or by manipulating other experts) to create a distribution over actions more likely to score their greatest improvement.

In this section we extend the myopic incentives of prediction markets to decision markets, demonstrating how to construct myopic incentive compatible decision markets, and characterizing some of their properties. While myopic incentive compatibility does not guarantee that an expert who participates in multiple rounds will predict consistent with its beliefs in *every* round, previous work has shown that myopic incentives are sufficient to aggregate experts' private information at perfect Bayesian equilibria under certain conditions [4,15].

4.1 Myopic Incentive Compatibility

We first provide a formal treatment of myopic incentive compatibility for decision markets. Recall, for a prediction market, myopic incentive compatibility requires an expert always maximize their score for a prediction when they predict consistent with their beliefs. Assume a decision market uses decision rule D and decision scoring rule S . Then an expert with beliefs Q over the conditional outcomes who expects that d will be the final distribution over actions has an expected score for a prediction P of

$$\sum_{a \in \mathcal{A}} d_a \sum_{o \in \mathcal{O}} Q_{a,o} S_{a,o}(d, P).$$

And a myopic incentive compatible decision market must account not only for an expert's prediction, but also the likelihood each action is taken.

Definition 1. A decision market (D, S) with a regular decision scoring rule S is proper if

$$\sum_{a \in \mathcal{A}} d_a \sum_{o \in \mathcal{O}} Q_{a,o} S_{a,o}(d, Q) \geq \sum_{a \in \mathcal{A}} d'_a \sum_{o \in \mathcal{O}} Q_{a,o} S_{a,o}(d', P),$$

for all beliefs Q , distributions d and d' in the codomain of D and predictions P . The market is strictly proper if the inequality is strict unless $P = Q$.

If the market is strictly proper we also describe it as myopic incentive compatible, analogous to our treatment of prediction markets. When a decision market is not strictly proper there exist final predictions and beliefs such that experts maximize their score for a prediction by misrepresenting their beliefs.

4.2 A Simple Construction For Strictly Proper Decision Markets

Given a decision rule with full support, a simple construction can extend any strictly proper scoring rule into a decision scoring rule that makes the resulting decision market strictly proper, too.

Theorem 1. *Let D be a decision rule with full support. Then there exists a decision scoring rule S such that (D, S) is strictly proper.*

Proof. The proof is by construction. Let s be any strictly proper scoring rule. Construct

$$S_{a,o}(d, P) = \frac{1}{d_a} s_o(P_a). \quad (2)$$

(D, S) is strictly proper: an expert's expected score for a prediction is

$$\sum_{a \in \mathcal{A}} d_a \sum_{o \in \mathcal{O}} Q_{a,o} \frac{1}{d_a} s_o(P_a) = \sum_{a \in \mathcal{A}} \sum_{o \in \mathcal{O}} Q_{a,o} s_o(P_a),$$

the sum of the expected scores of the same prediction in a set of prediction markets, one for each action, using a strictly proper scoring rule. Since predicting consistent with beliefs maximizes the expected score of the expert in each market, it maximizes the sum of the expected scores. $\square$

This constructive result positively answers Chen and Kash's open question whether it is possible to construct decision markets with good incentives [6].

4.3 Strictly Proper Decision Markets Have Full Support

The construction in Theorem 1 requires a decision rule have full support, and makes experts' expected scores independent of future reports while their actual scores vary inversely with the likelihood an action is chosen. Surprisingly, every strictly proper decision market with a differentiable decision scoring rule has these properties. We prove the necessity of full support before characterizing all strictly proper decision market with differentiable decision scoring rules.

Theorem 2. *Let D be a decision rule. A decision scoring rule S that makes (D, S) strictly proper exists if and only if D has full support.*

Proof. First we prove that if a decision rule D does not have full support, there is no decision scoring rule S such that (D, S) is strictly proper. We proceed by contradiction. Let D be a decision rule without full support, choose a final report P so that $d = D(P)$ has $d_k = 0$ for some $k \in \mathcal{A}$, and let S be a decision scoring rule such that (D, S) is strictly proper. Let $Q, Q' \in \Delta(\mathcal{O})^{|\mathcal{A}|}$ be two beliefs differing only on action k : for all $a \neq k$ and all o , $Q_{a,o} = Q'_{a,o}$; $\exists o$ $Q_{k,o} \neq Q'_{k,o}$. Consider the expected utility of an expert with each of these beliefs reporting truthfully, while the final report remains P . One of these utilities must be weakly greater than the other. Without loss of generality, let

$$\sum_{a \in \mathcal{A}} \sum_{o \in \mathcal{O}} d_a Q_{a,o} S_{a,o}(d, Q) \geq \sum_{a \in \mathcal{A}} \sum_{o \in \mathcal{O}} d_a Q'_{a,o} S_{a,o}(d, Q'), \quad (3)$$

Because Q and Q' only differ on action k , and $d_k = 0$,

$$\sum_{a \in \mathcal{A}} \sum_{o \in \mathcal{O}} d_a Q_{a,o} S_{a,o}(d, Q) = \sum_{a \in \mathcal{A}} \sum_{o \in \mathcal{O}} d_a Q'_{a,o} S_{a,o}(d, Q), \quad (4)$$

Combining lines (3) and (4) contradicts strict properness with respect to Q' .

The other direction, which shows how to construct a strictly proper decision market for any decision rule with full support, follows by the construction in the proof of Theorem 1. $\square$

This result extends Othman and Sandholm's impossibility result for deterministic decision markets [16] to the more general class of decision markets without full support. The theorem does not apply to non-strictly proper decision markets, however; for example, all constant decision scoring rules are proper for all decision rules.

4.4 Decision Markets With Good Incentives

While Theorem 1 provided a simple construction to create a strictly proper decision market, we now characterize all strictly proper decision and decision scoring rule pairs. The proof of Theorem 3 parallels similar characterizations of proper scoring rules given by Gneiting and Raftery [10] and of strictly proper pairs for a single expert [6], and appears in the full version of the paper¹.

Theorem 3. *A decision market (D, S) , where S is regular and D has full support, is (strictly) proper if*

$$S_{a,o}(d, P) = \frac{1}{d_a |\mathcal{A}|} (G(P) - G'(P) : P + |\mathcal{A}| G'_{a,o}(P)) \quad (5)$$

where $G : \Delta(\mathcal{O})^{|\mathcal{A}|} \rightarrow \mathbb{R}$ is a (strictly) convex function, $G'(P)$ is a subgradient of G at P and $:$ denotes the Frobenius inner product. Conversely, if S is differentiable in P and (D, S) is (strictly) proper, then S can be written in the form of (5) for some (strictly) convex G .

¹ Available from the authors' personal webpages.

Like the construction of Theorem 1, the characterization shown in Theorem 3 requires an expert’s expected score to be independent of the final report, and that the realized score vary inversely with the likelihood an action is taken. It does, however, allow more complicated constructions than the normalized strictly proper scoring rules used in Theorem 1. For example, given a decision rule D with full support, defining

$$S_{a,o}(d, P) = \frac{1}{d_a |\mathcal{A}|} (2|\mathcal{A}| P_{a,o} - \sum_{i,j} P_{i,j}^2), \quad (6)$$

makes (D, S) a strictly proper decision market.

Theorem 3 also illustrates that our expansion of the payment rule in decision markets from scoring rules to decision scoring rules is necessary to obtain myopic incentive compatibility, because scoring rules do not allow a dependence on d . Scoring rules function properly in the special case of a prediction market because for any constant decision rule a strictly proper scoring rule is sufficient to create myopic incentive compatibility.

5 Extensions

In this section we discuss how a decision maker can approximate a deterministic rule, and what the optimal decision rule for a risk-neutral decision maker would be. We also demonstrate a correspondence between any strictly proper decision market and a set of strictly proper prediction markets, suggesting a framework for applying previous prediction market results to decision markets.

5.1 Approximating Deterministic Decisions

Deterministic decision rules, like MAX, are natural. Unfortunately, no strictly proper decision market can use a deterministic decision rule. It is possible, however, to approximate deterministic decision rules with stochastic ones, but better approximations of a deterministic decision rule increase the decision maker’s worst case loss.

Corollary 1. *Every strictly proper decision market (D, S) where $\inf_{P \in \Delta(\mathcal{O})^{|\mathcal{A}|}} D_a(P) = 0$ for some action a has unbounded worst-case loss.*

We omit the proof as it follows directly from the inverse relationship between scores and the likelihood of actions required by Theorem 3.

5.2 Expected Utility Maximizing Decision Rules

A natural question is how a decision maker should design a strictly proper decision market to maximize their expected utility. A decision maker’s utility is the payoff they receive for the observed outcome minus the cost of paying experts. Since the expected payment to experts is independent of the decision rule used,

an expected utility maximizing decision rule maximizes the likelihood the most preferred action is taken, subject to the decision maker’s budget constraint. We call this decision rule APPROX-MAX, and since picking a decision scoring rule is analogous to picking a scoring rule for a prediction market, we take it as given when defining the decision rule.

Given a budget b and decision scoring rule S , and the final reports $P^{\bar{t}}$, we compute a minimal feasible probability for each action a ,

$$p_a = \max_o \frac{S_{a,o}([1]^{|a|}, P^{\bar{t}}) - S_{a,o}([1]^{|a|}, P^0)}{b},$$

where $[1]^{|a|}$ is a vector of ones. This expression computes the decision maker’s worst-case payment to experts for each action, unweighted, then divides that value by the budget to find a feasible inverse multiplier for the decision scoring rule, which is equal to the minimal feasible probability. If $\sum_{a \in \mathcal{A}} p_a > 1$ then no decision rule fits the decision maker’s budget, but otherwise a “probability surplus” of $1 - \sum_{a \in \mathcal{A}} p_a$ can be assigned arbitrarily. APPROX-MAX adds this surplus to the minimal feasible probability of the most preferred action to maximize the decision maker’s expected utility.

5.3 A Correspondence Between Decision Markets and Prediction Markets

Strictly proper decision markets constructed using the technique in Theorem 1 have an expected score for a prediction P , given beliefs Q , of

$$\sum_{a \in \mathcal{A}} \sum_{o \in \mathcal{O}} Q_{a,o} s_o(P_a),$$

where s is a strictly proper scoring rule. This is also equal to an expert’s expected score for a set of predictions in a set of independent prediction markets, one for each action in $\mathcal{A}$. This equivalence holds more generally: every strictly proper decision market has a corresponding set of prediction markets. Theorem 4 states this correspondence formally.

Theorem 4. *Every strictly proper decision market (D, S) , where S is differentiable, has a corresponding strictly proper set of prediction markets. This correspondence implies that when the previous prediction in both settings is the same, the expected score for a new prediction, given any beliefs, is also the same.*

Informally, this theorem implies risk-neutral experts are indifferent to participating in a strictly proper decision market or the corresponding set of strictly proper prediction markets. Their available predictions and expected scores for each prediction are the same in both settings. This correspondence suggests that results applying to sets of prediction markets may also apply directly to decision markets.

6 Conclusion

We extended the myopic incentive compatibility of prediction markets to decision markets. We proved that this extension requires the decision maker use a decision rule with full support, and showed how to construct a strictly proper decision market for any such decision rule, answering an open question posed by Chen and Kash [6]. We characterized the set of myopic incentive compatible decision markets, and show that it is possible to approximate any deterministic decision rule with a stochastic decision rule, although better approximations cause higher worst-case loss for the decision maker. We also showed a correspondence between strictly proper decision and sets of prediction markets, suggesting a unifying technique to apply results to both types of markets.

There remain many interesting research questions involving decision markets. Requiring decision makers commit to a randomized decision rule poses an important practical challenge. Returning to our example from the introduction, the project manager must be willing to risk hiring a slower developer for the privilege of running a myopic incentive compatible decision market. This is simply not credible behavior—managers prefer to hire faster developers. Designing a more credible mechanism is likely to be a prerequisite for the deployment of decision markets in practice.

Another practical concern is extending our decision market results to a cost function framework. Instead of requiring experts provide an entire probability distribution, cost function based prediction markets allow traders to buy and sell contracts associated with particular outcomes [7]. The price of each contract is expected to represent the likelihood that outcome occurs. These interfaces are similar to that provided by stock markets, and there is an equivalence between scoring rule and cost function markets. The same equivalence holds for decision markets, but our decision scoring rules require contracts with variable payouts or large upfront costs, both undesirable features. Designing a more natural contract structure for a decision may be of considerable practical value.

References

1. J. E. Berg, R. Forsythe, F. D. Nelson, and T. A. Rietz. Results from a dozen years of election futures markets research. In C. A. Plott and V. Smith, editors, *Handbook of Experimental Economic Results*. 2001.
2. J. E. Berg and T. A. Rietz. Prediction markets as decision support systems. *Information Systems Frontier*, 5:79–93, 2003.
3. K.-Y. Chen and C. R. Plott. Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem. Working paper No. 1131, California Institute of Technology, Division of the Humanities and Social Sciences, 2002.
4. Y. Chen, S. Dimitrov, R. Sami, D. M. Reeves, D. M. Pennock, R. D. Hanson, L. Fortnow, and R. Gonen. Gaming prediction markets: Equilibrium strategies with a market maker. *Algorithmica*, 58(4):930–969, 2010.

5. Y. Chen, X. A. Gao, R. Goldstein, and I. A. Kash. Market manipulation with outside incentives. In *AAAI '11: Proceedings of the 25th Conference on Artificial Intelligence*, 2011.
6. Y. Chen and I. A. Kash. Information elicitation for decision making. In *AAMAS '11: Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems*, 2011.
7. Y. Chen and D. M. Pennock. A utility framework for bounded-loss market makers. In *UAI '07: Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence*, pages 49–56, 2007.
8. S. Debnath, D. M. Pennock, C. L. Giles, and S. Lawrence. Information incorporation in online in-game sports betting markets. In *EC '03: Proceedings of the 4th ACM conference on Electronic commerce*, pages 258–259, New York, NY, USA, 2003. ACM.
9. S. Dimitrov and R. Sami. Composition of markets with conflicting incentives. In *EC '10: Proceedings of the 11th ACM conference on Electronic commerce*, pages 53–62, New York, NY, USA, 2010. ACM.
10. T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
11. R. Hanson. Decision markets. *IEEE Intelligent Systems*, 14(3):16–19, 1999.
12. R. D. Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003.
13. R. D. Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):1–15, 2007.
14. J. McCarthy. Measures of the value of information. *PNAS: Proceedings of the National Academy of Sciences of the United States of America*, 42(9):654–655, 1956.
15. M. Ostrovsky. Information aggregation in dynamic markets with strategic traders. In *EC '09: Proceedings of the tenth ACM conference on Electronic commerce*, page 253, New York, NY, USA, 2009. ACM.
16. A. Othman and T. Sandholm. Decision rules and decision markets. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 625–632, 2010.
17. L. J. Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.
18. J. Wolfers and E. Zitzewitz. Prediction markets. *Journal of Economic Perspective*, 18(2):107–126, 2004.