

Information Elicitation for Decision Making*

Yiling Chen
School of Engineering and
Applied Sciences
Harvard University
Cambridge, MA 02138 USA
yiling@eecs.harvard.edu

Ian A. Kash
Center for Research on
Computation and Society
Harvard University
Cambridge, MA 02138 USA
kash@seas.harvard.edu

ABSTRACT

Proper scoring rules, particularly when used as the basis for a prediction market, are powerful tools for eliciting and aggregating beliefs about events such as the likely outcome of an election or sporting event. Such scoring rules incentivize a single agent to reveal her true beliefs about the event. Othman and Sandholm [16] introduced the idea of a decision rule to examine these problems in contexts where the information being elicited is conditional on some decision alternatives. For example, “What is the probability having ten million viewers if we choose to air new television show X? What if we choose Y?” Since only one show can actually air in a slot, only the results under the chosen alternative can ever be observed. Othman and Sandholm developed proper scoring rules (and thus decision markets) for a single, deterministic decision rule: always select the the action with the greatest probability of success. In this work we significantly generalize their results, developing scoring rules for other deterministic decision rules, randomized decision rules, and situations where there may be more than two outcomes (e.g. less than a million viewers, more than one but less than ten, or more than ten million).

Categories and Subject Descriptors

J.4 [Computer Applications]: Social and Behavioral Sciences—*Economics*

General Terms

Economics, Theory

Keywords

Information Elicitation, Proper Scoring Rules, Decision Rules, Decision Markets, Mechanism Design, Market Design, Prediction Markets

*This paper is based upon work supported by NSF under Grant No. CCF-0953516. Any opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the NSF.

Cite as: Information Elicitation for Decision Making, Yiling Chen and Ian A. Kash, *Proc. of 10th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2011)*, Tumer, Yolum, Sonenberg and Stone (eds.), May, 2–6, 2011, Taipei, Taiwan, pp. XXX-XXX. Copyright © 2011, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

1. INTRODUCTION

Eliciting information about uncertain events is crucial for informed decision making. Information is often possessed by individual agents. Truthfully eliciting such information, resolving conflicting beliefs, and aggregating the dispersed information are key problems for achieving collective intelligence in multi-agent systems. Proper scoring rules have been proposed to incentivize an expert to honestly report her probability assessment for an uncertain event [2, 7, 21, 22, 18, 23, 3, 6, 11]. Prediction markets, betting intermediaries where participants wager on the outcome of some event of interest, can output equilibrium market prices representing the consensus probability assessment based on the pooled information of market participants [1, 24, 17].

However, most work, with the exception of Othman and Sandholm [16] and Dimitrov and Sami [4], separates the information elicitation problem from the decision problem. Information elicitation and aggregation are studied without considering how the resultant information will be used. Participants of mechanisms are assumed to only care about their rewards within the mechanism. Incentives for the isolated elicitation problem may not motivate desirable behavior of participants when the elicited information is used for decision making, because participants may have vested interests in some decision outcome and may misreport their information hoping to achieve the decision outcome.

In this paper, we study the information elicitation problem situated in a decision making process. We consider a setting where a decision maker seeks to elicit information about the consequences of various actions that he could take and choose an action based on the elicited information. For example, a company needs to choose among three marketing campaigns and wants to elicit the probability that each campaign will result in a sales goal being reached. This setting was first investigated by Othman and Sandholm [16] when the decision maker uses a deterministic decision rule and there are two possible outcomes given any action. They derived a proper scoring rule for a particular decision rule. *Our main contribution is a theorem that significantly generalizes their result by characterizing all (strictly) proper scoring rules for all decision rules.*

The rest of the paper is organized as follows. In Section 2, we cover the relevant background and related work on the use of scoring rules for information elicitation. Then in Section 3, we provide a formal model of the problem of eliciting information from a single expert for the purpose of making a decision. We state and prove our main result, a characterization theorem for (strictly) proper scoring rules in the

decision making setting, in Section 4. Using this characterization, we provide a sufficient condition for the existence of a strictly proper scoring rule and a sufficient condition for when no such rule exists in Section 5. For situations where strictly proper scoring rules are not possible, we provide a derivation of “quasi-strictly proper” scoring rules in Section 6. We discuss issues that arise when eliciting information from multiple experts and some potential solutions in Section 7. We conclude in Section 8.

2. BACKGROUND AND RELATED WORK

Proper scoring rules have been designed to incentivize a risk-neutral expert to truthfully report her probability assessment for an uncertain event [2, 7, 21, 22, 18, 23, 14, 13, 20, 15]. Let v be a discrete random variable that has m mutually exclusive and exhaustive outcomes, $\mathcal{O} = \{o_1, \dots, o_m\}$. A scoring rule assigns score $s_o(\vec{p})$ to an expert who reports a probability assessment \vec{p} when outcome o happens. A scoring rule is *regular* if $s_o(\vec{p})$ is real valued for all o , except possibly $s_o(\vec{p}) = -\infty$ if $p_o = 0$. A regular scoring rule is (*strictly*) *proper* if a risk-neutral expert (strictly) maximizes her expected score by reporting truthfully. That is, \vec{p} is an optimal solution to $\max_{\vec{r} \in \Delta(\mathcal{O})} \sum_{o=1}^m p_o s_o(\vec{r})$ for any proper scoring rule s , and is the unique optimal solution if s is strictly proper. $\Delta(\mathcal{O})$ is the probability simplex over \mathcal{O} . For example, logarithmic scoring rule $s_o(\vec{p}) = a_o + b \log p_o$ and quadratic scoring rule $s_o(\vec{p}) = a_o + b(2p_o - \sum_i p_i^2)$, where $b > 0$ and a_o are arbitrary parameters, are strictly proper scoring rules.

Proper scoring rules are closely related to convex functions. In fact, the following characterization theorem of Gneiting and Raftery [6], which is credited to McCarthy [12] and Savage [18], gives the precise relationship between convex functions and proper scoring rules. In the theorem, \cdot denotes the vector inner product.

THEOREM 2.1 (GNEITING AND RAFTERY THEOREM 2).
A regular scoring rule is (strictly) proper if and only if

$$s_o(\vec{p}) = G(\vec{p}) - G'(\vec{p}) \cdot \vec{p} + G'_o(\vec{p}),$$

where $G : \Delta(\mathcal{O}) \rightarrow \mathbb{R}$ is a (strictly) convex function and $G'(\vec{p})$ is a subgradient of G at the point \vec{p} and $G'_o(\vec{p})$ is the o -th element of $G'(\vec{p})$.

Theorem 2.1 indicates that a regular scoring rule is (strictly) proper if and only if its expected score function $G(\vec{p}) = \sum_o p_o s_o(\vec{p})$ is (strictly) convex on $\Delta(\mathcal{O})$, and the vector with elements $s_o(\vec{p})$ is a subgradient of G at the point \vec{p} .

Hanson [9, 10] shows how a proper scoring rule designed to elicit information from a single expert can be turned into a mechanism for prediction markets, which aggregate the information of multiple experts. Such a mechanism is called a *market scoring rule* and is essentially a shared proper scoring rule. Hanson [8] also promotes the idea of *decision markets*. A decision market is a prediction market for conditional events. A decision maker who needs to decide among some actions can operate a conditional prediction market for each action. The conditional market elicits information on outcomes of some event of interest conditioned on the corresponding action being taken (e.g. probability that stock price of a company increases conditioned on A is hired as the CEO). The decision maker can decide on what action to take based on the elicited conditional probability distributions. Our work focuses on the incentive problem of eliciting

conditional information from a single expert using a scoring rule, but we also discuss the implications for using a market scoring rule for decision markets. Furthermore, while Hanson proposes the idea of a decision market, he does not provide any analysis or techniques showing how one could be implemented to correctly encourage participants to reveal their information. As we discuss in Section 7, this is a difficult problem.

The closest work to ours is that of Othman and Sandholm [16]. They pair a scoring rule for eliciting conditional probability distributions over two outcomes with a deterministic decision rule. This differs from the standard information elicitation problem using proper scoring rules, because only one action will be taken and used to determine the score of an expert, but the selected action depends on the reported conditional probability distributions. Othman and Sandholm show that for deterministic decision rules, to have a “quasi-strictly proper” scoring rule it is necessary that the decision rule only change its decision when probabilities are equal. A natural version of this is the MAX rule: decide on the action with the highest reported probability for the more desirable outcome. They construct a quasi-strictly proper scoring rule for MAX and then extend their results to decision markets. However, they show that it is impossible to achieve properness in decision markets using the MAX decision rule. As an open problem, they pose the question of how proper scoring rules can be derived for randomized decision rules. Our main theorem answers this question with a characterization of all (strictly) proper scoring rules for all decision rules. Thus, we extend their results to deterministic rules other than MAX, randomized decision rules, and situations with more than two outcomes.

Three other papers have considered the problem of information elicitation in other settings where the outcome is not independent of the predictions of experts. Shi, Conitzer, and Guo [19] examine settings where participants in a prediction market may also have an ability to influence the outcome. For example, participants in a market to predict terrorist attacks may be able to carry out acts of terrorism and employees of a company participating in a prediction market about when a new product will launch may have the ability to delay the launch. They show how to derive scoring rules that do not incentivize the participants to take these “undesirable” actions. However, unlike our work, the information elicited is not explicitly used for decision making. Dimitrov and Sami [4] examine incentive problems when there are two prediction markets for different but related events. This might cause a trader to report sub-optimally in one market to mislead a trader in another market. The first trader can also participate in the second market and profit from correcting the second trader. This is a situation where the payoff from the first market does not depend on the decision the second trader makes; instead, the first trader profits directly from the decision. Our work considers the opposite case: experts do not care what decision is made, except that the outcome (and thus their payoffs) depends on it. Gerding et al. [5] consider a model where experts need to be incentivized to make costly observations of the quality of service providers. They consider a number of approaches, including scoring rules, and all face tradeoffs between encouraging experts to invest effort and getting accurate reports. One of their approaches, basing scores on peer predictions, could potentially be helpful with resolving the issues with decision

markets discussed in Section 7.

3. OUR MODEL

Before we introduce our model, we note some notational conventions we use throughout the paper. We typically denote a matrix with a capital letter and an entry of a matrix P as $P_{i,o}$. We denote a vector as \vec{p} , except that when considering a row of a matrix as a vector we denote it P_i . We use the Frobenius inner product, $P : Q = \sum_{i,o} P_{i,o} Q_{i,o}$, for matrices.

Our model is essentially that of Othman and Sandholm [16], but adapted to allow randomized decision rules and more than two outcomes.

A decision maker needs to choose an action from a set $\mathcal{A} = \{1, \dots, n\}$. Each action may affect the probability of achieving each possible outcome from a set $\mathcal{O} = \{o_1, \dots, o_m\}$. Othman and Sandholm [16] considered the case of two outcomes, in which case we use \top to denote the more desirable outcome and \perp to denote the less.

The decision maker asks an expert to report a set of conditional probability distributions, denoted by a $n \times m$ matrix P , where $P_{i,o}$ is the probability of outcome o conditional on the decision maker taking action i . We use P_i to denote the i -th row of P , that is, the probability distribution over \mathcal{O} conditional on action i being taken. $P_i \in \Delta(\mathcal{O})$ for all actions i , where $\Delta(\mathcal{O})$ is the probability simplex over \mathcal{O} . We use \mathcal{P} to denote the space of P . In general, not every decision need potentially lead to each outcome. For example, we could model a decision maker that cares about which decision is made by having a disjoint set outcomes for each decision.

Based on the expert's report, the decision maker makes a decision using a decision rule $D : \mathcal{P} \rightarrow \Delta(\mathcal{A})$. $D_i(P)$ is the probability the decision maker assigns to action i given report P . In the special case of a deterministic decision rule, $D : \mathcal{P} \rightarrow \mathcal{A}$. The decision rule D is known by the expert.

To encourage the expert to make an accurate prediction, the decision maker rewards her using a scoring rule $S : \mathcal{A} \times \mathcal{O} \times \mathcal{P} \rightarrow \mathbb{R} \cup \{-\infty\}$. For notational convenience, we use $S_{i,o}(P)$ to represent $S(i, o, P)$, the score for the report P when action i is taken and outcome o happens. Note that we allow the expert's reward to depend on the decision made, a feature not necessary for the deterministic decision rules considered in Othman and Sandholm's model. We assume the expert is risk neutral and only cares about her reward according to the scoring rule. Specifically, she does not care what decision is made, other than to the extent that it affects her expected score.

We now define regular, proper, strictly proper, and quasi-strictly proper scoring rules for a decision rule.

DEFINITION 3.1. *A scoring rule S is regular for decision rule D if $S_{i,o}(P) \in \mathbb{R}$ unless $P_{i,o} = 0$.*

The definition is analogous to that of regular scoring rules in Section 2. An expert may get a score of $-\infty$ only if an event occurred to which she assigned probability 0. We consider only regular scoring rules for a decision rule in this paper because if this condition is not met an expert can get $-\infty$ in expectation, making the scoring rule unappealing.

Let $V(P, Q)$ denote the expected score of an expert who believes that the true conditional probabilities are P but reports the probabilities Q , i.e. $V(P, Q) = \sum_{i,o} D_i(Q) P_{i,o} S_{i,o}(Q)$.

We define (strictly) proper scoring rules for a decision rule as follows, which is a direct generalization of (strictly) proper scoring rules in Section 2.

DEFINITION 3.2. *A regular scoring rule S is proper for a decision rule D if*

$$V(P, P) \geq V(P, Q)$$

for all P and all $Q \neq P$. It is strictly proper for the decision rule if the inequality is strict.

Othman and Sandholm showed that no deterministic decision rule has a strictly proper scoring rule, but showed one that satisfies a slightly weaker condition. Intuitively, the decision maker does not care if the expert is not strictly incentivized to tell the truth about the probabilities for actions he does not take as long as he learns the true conditional probabilities for the action he takes. We formally define the notion for randomized decision rules below.

DEFINITION 3.3. *A regular scoring rule S is quasi-strictly proper for a decision rule D if it is proper (i.e.*

$$V(P, P) \geq V(P, Q)$$

for all P and all $Q \neq P$), and

$$V(P, P) > V(P, Q)$$

for all P and Q such that $P_k \neq Q_k$ for some $k \in \sigma_Q$, where $\sigma_Q = \{i | D_i(Q) > 0\}$ is the support of $D(Q)$ ¹.

A quasi-strictly proper scoring rule for a decision rule ensures that an expert is strictly incentivized to truthfully report her conditional probability distributions for actions that will be taken with positive probabilities by the decision maker, although she may lie about her conditional probability distributions for actions that won't be taken without changing her expected score.

4. CHARACTERIZING PROPER SCORING RULES FOR DECISION RULES

In this section, we state and prove our main theorem, a characterization of all regular (strictly) proper scoring rules for arbitrary (randomized) decision rules. We show that any scoring rule of a particular form is (strictly) proper for the corresponding decision rule and that every regular (strictly) proper scoring rule for a decision rule is of this form. This form, similar to the one used by Gneiting and Raftery in Theorem 2.1, relies on the (strict) convexity of a function G , which can be thought of as the expected truthful score function $V(P, P) = \sum_{i,o} D_i P_{i,o} S_{i,o}(P)$. Our theorem can be interpreted as saying that a scoring rule is (strictly) proper for D if and only if $G(P) = V(P, P)$ is (strictly) convex and satisfies some additional conditions. G need not be differentiable in general (for example with a deterministic decision rule), so rather than using the gradient of G , the theorem uses the notion of a subgradient. At a point where G is differentiable, the gradient is the unique subgradient.

¹Othman and Sandholm give a different definition of quasi-strict properness, which does not account for the possibility that a decision rule may be effectively "tied." For example the scoring rule they give for the MAX decision rule violates their definition when $P_{1,\top} = 0.5$, $P_{2,\top} = 0.5$, $Q_{1,\top} = 0.4$, and $Q_{2,\top} = 0.5$, but satisfies our definition.

The resulting theorem is quite powerful. For an arbitrary decision rule and scoring rule it provides a simple test to determine whether the scoring rule is proper for the decision rule. For an arbitrary decision rule, it gives a method of constructing proper scoring rules. Additionally, generalizations of many of Othman and Sandholm's results [16] characterizing properties of proper scoring rules for deterministic decision rules to situations where the decision rule does not have full support (for example that there are no strictly proper scoring rules and that all proper scoring rules satisfy an independence of irrelevant alternatives condition), are simple corollaries of our theorem.

THEOREM 4.1. *A regular scoring rule is (strictly) proper for a decision rule D if and only if*

$$S_{i,o}(P) = \begin{cases} G(P) - G'(P) : P + \frac{G'_{i,o}(P)}{D_i(P)} & D_i(P) > 0 \\ \Pi_{i,o}(P) & D_i(P) = 0 \end{cases} \quad (1)$$

where $G : \mathcal{P} \rightarrow \mathbb{R} \cup \{-\infty\}$ is a (strictly) convex function, $G'(P)$ is a subgradient of G at the point P with $G'_{i,o}(P) = 0$ when $D_i(P) = 0$, and $\Pi_{i,o} : \mathcal{P} \rightarrow \mathbb{R} \cup \{-\infty\}$ is an arbitrary function that can take a value of $-\infty$ only when $P_{i,o} = 0$.

PROOF. Consider a regular scoring rule S satisfying (1). We first show that it must be (strictly) proper. Let $\sigma_P = \{i \mid D_i(P) > 0\}$. We have,

$$\begin{aligned} V(P, P) &= \sum_{i,o} D_i(P) P_{i,o} S_{i,o}(P) \\ &= \sum_{i \in \sigma_P, o} D_i(P) P_{i,o} \left(G(P) - G'(P) : P + \frac{G'_{i,o}(P)}{D_i(P)} \right) \\ &= G(P) - G'(P) : P + \sum_{i \in \sigma_P, o} G'_{i,o}(P) P_{i,o} \\ &= G(P) - G'(P) : P + G'(P) : P = G(P). \end{aligned}$$

The fourth equality relies on the condition that $G'_{i,o}(P) = 0$ when $D_i(P) = 0$. Because G is convex and G' is a subgradient, for $Q \neq P$

$$\begin{aligned} V(P, Q) &= \sum_{i,o} D_i(Q) P_{i,o} S_{i,o}(Q) \\ &= \sum_{i \in \sigma_Q, o} D_i(Q) P_{i,o} \left(G(Q) - G'(Q) : Q + \frac{G'_{i,o}(Q)}{D_i(Q)} \right) \\ &= G(Q) + (P - Q) : G'(Q) \leq G(P) = V(P, P). \end{aligned}$$

This gives us that S is a proper scoring rule for D . The inequality is strict if G is strictly convex, in which case S is strictly proper for D .

Now consider a regular proper scoring rule S for D . We will show that it must be of the form of (1). Define $G(P) = V(P, P) = \sup_Q V(P, Q)$ and $G'_{i,o}(P) = D_i(P) S_{i,o}(P)$. Each $V(P, Q)$ is a convex function of P . G is a point-wise supremum of a set of convex functions and hence is convex itself. If $D_i(P) = 0$, $G'_{i,o}(P) = 0$ by definition. For $Q \neq P$, we have

$$\begin{aligned} G(P) + (Q - P) : G'(P) &= \sum_{i,o} D_i(P) P_{i,o} S_{i,o}(P) + \sum_{i,o} (Q_{i,o} - P_{i,o}) D_i(P) S_{i,o}(P) \\ &= \sum_{i,o} D_i(P) Q_{i,o} S_{i,o}(P) = V(Q, P) \leq V(Q, Q) = G(Q). \end{aligned}$$

The inequality is due to the properness of S . It is strict if S is strictly proper. Thus, G' is a subgradient of the convex function G , and G is strictly convex if S is strictly proper. Finally, for any P and j such that $D_j(P) > 0$,

$$\begin{aligned} G(P) - G'(P) : P + \frac{G'_{j,o'}(P)}{D_j(P)} &= \sum_{i,o} D_i(P) P_{i,o} S_{i,o}(P) - \sum_{i,o} D_i(P) S_{i,o}(P) P_{i,o} + \frac{G'_{j,o'}(P)}{D_j(P)} \\ &= \frac{G'_{j,o'}(P)}{D_j(P)} = \frac{D_j(P) S_{j,o'}(P)}{D_j(P)} = S_{j,o'}(P). \end{aligned}$$

So, S is of the proper form. \square

In Theorem 4.1, $\Pi_{i,o}$ allows arbitrary scores to be assigned when $D_i(P) = 0$, subject to the constraint that the scoring rule is regular. As the decision rule never takes action i , the score that would be assigned if it did is essentially arbitrary and does not affect the expected score of the expert. The theorem has another condition with no parallel in Theorem 2.1 by Gneiting and Raftery: the requirement that $G'_{i,o}(P) = 0$ if $D_i(P) = 0$. This can be read as requiring the expert's expected score to be independent of her reports about the probabilities for actions that will not be taken. For proper scoring rules for a decision rule D , this condition is satisfied by the trivial convex function $G(P) = 0$, which pays the expert nothing no matter what she reports, so she is weakly indifferent to truthful reporting. However, this condition may not be satisfied by any strictly convex function, resulting in a non-existence of strictly proper scoring rules for D , an issue to which we return in Section 5.

We conclude this section with a number of examples of proper scoring rules that can be derived using Theorem 4.1.

- For the two-outcome case, taking $G(P) = \max_i P_{i,\top}^2$ for the MAX decision rule gives the proper scoring rule derived by Othman and Sandholm [16].
- More generally, with more than two outcomes the decision maker may have some utility $u(o)$ for each outcome and want to use the deterministic decision rule that selects the action i that maximizes expected utility $U_i(P) = \sum_o u(o) P_{i,o}$. In this case, he can use $G(P) = \max_i U_i(P)^2$, which gives the proper scoring rule $S_{i,o} = 2U_i(P)u(o) - U_i(P)^2$. Note that this rule is not strictly proper, but we will see in Section 6 that it is quasi-strictly proper.
- For the two outcome case with randomized decision rule $D_i(P) = P_{i,\top} / \sum_j P_{j,\top}$, taking $G(P) = \sum_i P_{i,\top}^2$ gives us the strictly proper scoring rule $S_{i,\top} = \sum_j 2P_{j,\top} - P_{j,\top}^2$ and $S_{i,\perp} = -\sum_j P_{j,\top}^2$, which is reminiscent of the quadratic scoring rule.

5. STRICT PROPERNESS

In addition to characterizing all proper scoring rules for a particular decision rule, Theorem 4.1 characterizes the strictly proper scoring rules as well. However, as Othman and Sandholm [16] observed for the case of deterministic rules, some decision rules may not have any strictly proper scoring rules. More generally, we would like to know whether, given a decision rule D , there exists a strictly convex G satisfying the requirements of Theorem 4.1, and thus a strictly

proper scoring rule. In this section, we give sufficient conditions for both the existence and non-existence of strictly proper scoring rules.

For a strictly proper scoring rule to exist, we need to find a strictly convex function G that satisfies the condition from Theorem 4.1 that $G'_{i,o}(P) = 0$ whenever $D_i(P) = 0$. When $D(P)$ always has full support (i.e. $D_i(P)$ is never 0) this is trivially satisfied by any strictly convex function. This gives us a sufficient condition for a decision rule to have a strictly proper scoring rule.

COROLLARY 5.1. *If a decision rule D always has full support ($D_i(P) > 0$ for all i and P) then it has a strictly proper scoring rule.*

PROOF. Any strictly convex function G , for example $G(P) = \sum_{i,o} P_{i,o}^2$, satisfies the requirements of Theorem 4.1 and thus yields a strictly proper scoring rule. \square

On the other hand, Othman and Sandholm show that no deterministic decision rule has a strictly proper scoring rule. The following corollary establishes a larger class of decision rules for which this is the case, namely those for which the probability distribution over actions chosen by the decision rule does not have full support and there is a case where the probabilities outside the support can be changed without changing the probability distribution over actions.

COROLLARY 5.2. *If there exist $P \neq Q$ such that*

1. $D(P) = D(Q)$,
 2. $\sigma_P \subset \mathcal{A}$, and
 3. $P_{i,o} = Q_{i,o}$ for all $i \in \sigma_P$ and all o ,
- then D does not have a strictly proper scoring rule.

PROOF. Consider such a P and Q and a proper scoring rule S . By Theorem 4.1, $G'_{i,o}(P) = G'_{i,o}(Q) = 0$ for all $i \notin \sigma_P$, so $G(P) = G(Q)$ and

$$\begin{aligned} V(P, Q) &= \sum_{i,o} P_{i,o} D_i(Q) S_{i,o}(Q) \\ &= \sum_{i \in \sigma_Q, o} P_{i,o} D_i(Q) (G(Q) + G'(Q) : Q + G'_{i,o}(Q)) \\ &= \sum_{i \in \sigma_P, o} P_{i,o} D_i(P) (G(P) + G'(P) : P + G'_{i,o}(P)) \\ &= V(P, P). \end{aligned}$$

Thus S is not strictly proper. \square

While Corollary 5.2 shows that a subset of decision rules that do not have full support do not have a strictly proper scoring rule, the following open problem remains.

OPEN PROBLEM 1. *Characterize when decision rules that do not have full support and also do not satisfy the additional conditions of Corollary 5.2 have a strictly proper scoring rule.*

6. QUASI-STRICT PROPERNESS

We saw in Section 5 that, while we can always construct a proper scoring rule, many decision rules do not have any strictly proper scoring rules. The mere existence of a proper scoring rule is unsatisfying; the scoring rule that gives the expert a score of 0 no matter what the decision and outcome

is proper for every decision rule but gives the expert no particular incentive to reveal her beliefs. Strictly proper scoring rules fix this problem by ensuring that truthful reporting is uniquely optimal. While not quite as satisfying, a quasi-strictly proper scoring rule provides the weaker promise that, no matter what optimal report the expert makes, she reported her true beliefs over the outcome space *for the actions the decision maker might take*. In this section, we give a derivation of quasi-strictly proper scoring rules for a class of decision rules.

To build intuition about how quasi-strictly proper scoring rules can be derived, we first examine a set of sufficient conditions for a scoring rule to be quasi-strictly proper for the MAX decision rule.

LEMMA 6.1 (OTHMAN AND SANDHOLM [16]). *Let f and g be functions such that*

1. f and g are twice differentiable on $(0, 1)$,
2. $h(p) = pf(p) + (1-p)g(p)$ is strictly increasing on $[0, 1]$,
3. $pf'(p) + (1-p)g'(p) = 0$ for all $p \in [0, 1]$, and
4. $pf''(p) + (1-p)g''(p) < 0$ for all $p \in [0, 1]$.²

Then $S_{i,\top}(P) = f(P_{i,\top})$ and $S_{i,\perp}(P) = g(P_{i,\top})$ is quasi-strictly proper for the MAX decision rule.

A subset of these conditions also suffices to prove that a function h used in the construction is strictly convex.

LEMMA 6.2. *Let f and g be functions such that*

1. f and g are twice differentiable on $(0, 1)$,
2. $pf'(p) + (1-p)g'(p) = 0$ for all $p \in [0, 1]$,
3. $pf''(p) + (1-p)g''(p) < 0$ for all $p \in [0, 1]$.

Then $h(p) = pf(p) + (1-p)g(p)$ is strictly convex on $[0, 1]$.

PROOF. $h'(p) = f(p) - g(p) + pf'(p) + (1-p)g'(p)$ and $h''(p) = 2(f'(p) - g'(p)) + pf''(p) + (1-p)g''(p)$. Because $pf'(p) + (1-p)g'(p) = 0$, we have $h'(p) = f(p) - g(p)$ and $h''(p) = f''(p) - g''(p)$. Combining our two equations for h'' gives $f''(p) - g''(p) + pf''(p) + (1-p)g''(p) = 0$, or $f''(p) > g''(p)$. Thus $h''(p) > 0$ and h is strictly convex. \square

This is not a coincidence; we now show how such strictly convex functions can be used to construct quasi-strictly proper scoring rules for a large class of decision rules. In the simple case of deterministic decision rules, the members of this class share the feature that the desirability of each action can be computed as a strictly convex function of the conditional probabilities reported for that action and the decision rule simply takes the maximum of these desirabilities. For example, the MAX decision rule for two outcomes can be expressed as $D(P) \in \operatorname{argmax}_i P_{i,\top}^2$, where $h(P_i) = P_{i,\top}^2$ is a strictly convex function of P_i . More generally the decision rule may randomize over several actions and may a priori exclude some actions or combinations of actions from consideration. Thus, our construction proceeds by selecting a subset of the power set of actions, associating a strictly convex function with each, and showing that every corresponding decision rule has a quasi-strictly proper scoring rule.

²Othman and Sandholm do not explicitly state this condition, but it is implicit from the proof of their Theorem 7.

LEMMA 6.3. Let $\beta \subseteq 2^{\mathcal{A}} - \{\emptyset\}$, and for each $b \in \beta$, $G^b : \Delta(\mathcal{O})^{|\mathcal{b}|} \rightarrow \mathbb{R}$ be a strictly convex function, $D(P)$ have support $\text{argmax}_{b \in \beta} G^b(P_b)$ for all P , and P_b be the submatrix of P consisting of those rows whose action is in b . Then, the scoring rule from Theorem 4.1 with $G(P) = \max_{b \in \beta} G^b(P_b)$ is quasi-strictly proper for D .

PROOF. Let P and Q be given and let b be the support of $D(Q)$ (the decisions made with positive probability). Then by Theorem 4.1 and the strict convexity of G^b ,

$$\begin{aligned} V(P, Q) &= G(Q) + (P - Q) : G'(Q) \\ &= G^b(Q_b) + (P_b - Q_b) : (G^b)'(Q_b) \\ &\leq G^b(P_b) \leq \max_{c \in \beta} G^c(P_c) = V(P, P). \end{aligned}$$

If $P_b \neq Q_b$ the first inequality is strict. If $b \notin \text{argmax}_{c \in \beta} G^c(P_c)$ the second is strict. Thus the scoring rule is quasi-strictly proper for D . \square

In the statement of Lemma 6.3, β is the set of possible supports the decision rule considers. For each such support b , G^b is a strictly convex function that determines how “good” that support is given the probabilities ($G^b(P_b)$). Our construction applies to any decision rule that always has a support that is “best” according to the various G^b ($\text{argmax}_{b \in \beta} G^b(P_b)$). There are many such rules, as this condition restricts only the support, not the actual decision probabilities. Each has a different quasi-strictly proper scoring rule, but they can all be derived from the same convex function $G(P) = \max_{b \in \beta} G^b(P_b)$.

Lemma 6.3 allows us to derive quasi-strictly proper scoring rules for deterministic decision rules with two outcomes using $\beta = \mathcal{A}$ and $G^b = h$. For example, we can derive quasi-strictly proper scoring rules for MAX (e.g. $h(P_i) = P_{i,\top}^2$, mentioned previously, gives Othman and Sandholm’s rule [16]), MIN (e.g. $h(P_i) = P_{i,\perp}^2$), and even strange rules like “probability farthest from 0.5” ($h(P_i) = (P_{i,\top} - 0.5)^2$). We can take $\beta \subset \mathcal{A}$ to allow for decisions rules that only allow certain actions (e.g. “choose whichever of actions 1 and 3 is more likely to succeed”). With more than two outcomes it allows rules like the expected utility maximization rule from Section 4. We can also apply this construction to the randomized case. For example, we saw in Section 4 a construction of a scoring rule that is strictly proper for the decision rule $D_i(P) = P_{i,\top} / \sum_j P_{j,\top}$. Lemma 6.3 tells us that a version of this rule that disregards some actions and uses an appropriately modified scoring rule is quasi-strictly proper. In particular, if $\alpha \subset \mathcal{A}$ is the set of actions considered then Lemma 6.3 can be applied with $\beta = \{\alpha\}$ and $G^\alpha(P) = \sum_{i \in \alpha} P_{i,\top}^2$.

The proof of Lemma 6.3 actually proves something stronger than quasi-strict properness. In particular, it shows that, unless the support of $D(Q)$ is a maximizer of $\max_{c \in \beta} G^c(P_c)$, $V(P, P) > V(P, Q)$. Thus, not only does the expert have a strict incentive to report the true probabilities for the actions the decision maker ends up randomizing over, she also has a strict incentive to ensure this set is one that the decision rule considers “optimal.”

One interesting observation about these scoring rules in the deterministic case is that they can all be viewed as strictly proper scoring rules when outcomes are exogenous. For example taking $h(P_i) = P_{i,\top}^2$ and $D(P) = i$ gives the scoring rule $s_\top(P_i) = 2P_{i,\top} - P_{i,\top}^2$ and $s_\perp(P_i) = -P_{i,\top}^2$,

which is a variant of the well known quadratic scoring rule (which is strictly proper).

In fact, we can show that this is generally true for quasi-strictly proper scoring rules derived according to Lemma 6.3 with a deterministic decision rule. When D is a deterministic decision rule, its support given P must be a singleton action. Hence, β in Lemma 6.3 equals \mathcal{A} . For each element of β (i.e. each action i), we set $G^i(P_i) = h(P_i)$, where h is strictly convex. Thus, the decision rule will take action k where $k \in \text{argmax}_{i \in \mathcal{A}} h(P_i)$. We assume that the decision rule breaks ties arbitrarily when there are more than one actions that have the same highest value of $h(P_i)$. We have $G(P) = \max_{i \in \mathcal{A}} h(P_i) = h(P_k)$ and can derive a quasi-strictly proper scoring rule $S_{i,o}(P)$ according to expression (1) in Theorem 4.1. Clearly, for the chosen action k , $S_{k,o}(P)$ only depends on P_k . We would like to consider whether $S_{k,o}(P)$ can be viewed as a strictly proper scoring rule of P_k assuming that action k is always chosen no matter how P_k changes. Let $\vec{q} = \text{argmin}_{\vec{p}} h(\vec{p})$. \vec{q} is unique because h is strictly convex. We construct a scoring rule $s_o(\vec{p}) = S_{k,o}(Q^{\vec{p}})$ where $Q^{\vec{p}}$ is a probability matrix where $Q_k^{\vec{p}} = \vec{p}$ and $Q_j^{\vec{p}} = \vec{q}$ for all $j \neq k$.

COROLLARY 6.1. $s_o(\vec{p})$ constructed above is strictly proper.

PROOF. For all \vec{p} , $D(Q^{\vec{p}}) = k$, so $s_o(\vec{p}) = S_{k,o}(Q^{\vec{p}}) = h(\vec{p}) - h'(\vec{p}) \cdot \vec{p} + h'(\vec{q}) \cdot \vec{p}$. By Theorem 2.1 and the strict convexity of h , s is strictly proper. \square

Thus, for deterministic decision rules, the quasi-strictly proper scoring rules derived according to Lemma 6.3 are strictly proper given a chosen action. As we will see in Section 7.1, this is a potentially useful property if one of these rules is used as a basis for a decision market.

7. DISCUSSION: DECISION MARKETS

Scoring rules are useful in their own right as a tool to elicit information from a single expert, but collectively a group of experts may provide better information. In this section, we discuss some challenges and observations on using decision markets to elicit information from multiple experts.

For a standard market scoring rule, using a strictly proper scoring rule s , the market maintains a probability distribution over outcomes \vec{p} . At any time, a trader can change this to \vec{q} , and in doing so accepts the following bet: if outcome o occurs then the market pays her $s_o(\vec{q}) - s_o(\vec{p})$ (which may be negative). A trader who only participates once maximizes her expected payoff by changing the market probability to match her true beliefs. We can use the same approach for decision markets. A decision market maintains a market probability matrix P . Any trader can change this to Q , accepting the bet that if action i is taken and outcome o occurs then she receives $S_{i,o}(Q) - S_{i,o}(P)$. At the close of the market, there is some final probability matrix F , and the decision is made according to $D(F)$.

However, as Othman and Sandholm [16] observed, traders’ incentives are not as perfectly aligned in a decision market, even if S is a proper scoring rule for D . A trader’s payoff relying on $D(F)$ points to two key issues. First, in order to determine her expected utility for a report, a trader needs to know what F will be, which is not determined until the market closes. One way to resolve this issue is to follow Othman and Sandholm and consider the last trader in the market, whose report is F (or equivalently assume that

traders are myopic and all assume they are the last trader). Second, in a standard market scoring rule, a trader's expected payment to the market institution given beliefs \vec{q} is $\vec{q} \cdot s(\vec{p})$, which is independent of her report. Thus, she chooses her report \vec{r} to maximize $\vec{q} \cdot s(\vec{r})$, and strict properness of s is sufficient. However, a myopic trader who reports R in a decision market makes an expected payment of $\sum_{i,o} D_i(R) Q_{i,o} S_{i,o}(P)$ given beliefs Q . This is *not* independent of R . Thus, although properness of S means that Q maximizes $\sum_{i,o} D_i(R) Q_{i,o} S_{i,o}(R)$, unlike the simple market scoring rule case, it does not follow that Q maximizes $\sum_{i,o} D_i(R) Q_{i,o} (S_{i,o}(R) - S_{i,o}(P))$.

Othman and Sandholm give the following example for the two-outcome, two-action case under the MAX decision rule with scoring rule $S_{i,\top}(P) = 2P_{i,\top} - P_{i,\top}^2$, $S_{i,\perp}(P) = -P_{i,\top}^2$ and show that all scoring rules for MAX have similar manipulations. Suppose a trader believes the true probabilities are $(Q_{1,\top}, Q_{2,\top}) = (0.8, 0.75)$, but the current market probabilities are $(0.8, 0.3)$. If the trader reports her true belief, her net expected payment is 0, but if she reports $(0.8, 0.81)$ her expected payment is $\bar{0}.15$. In essence, she only gets paid for correcting the value of $P_{2,\top}$ if she convinces the decision maker to choose decision 2. To make matters worse, any later trader with similar beliefs is weakly indifferent to correcting the market, so the market may get stuck at these wrong probabilities. Furthermore, this manipulation is "safe;" if action 1 is chosen in the end the trader's payment is 0 and if action 2 is chosen her expected payment is positive.

Clearly this is not a desirable outcome. Othman and Sandholm propose to address this problem choosing a scoring rule that minimizes, but does not eliminate, the incentive for a trader to perform such a manipulation. In the remainder of this section we consider several other approaches.

7.1 Faith in Markets

Suppose that, rather than being myopic, a trader believes the market will "get it right" in the end. That is, if her beliefs are Q , she believes that, regardless of her report F will eventually equal Q . Then she believes that the portion of her payment to the market institution based on the current market probabilities is $\sum_{i,o} D_i(Q) Q_{i,o} S_{i,o}(P)$, which is independent of her report. Thus she wants to optimize $\sum_{i,o} D_i(Q) Q_{i,o} S_{i,o}(R)$ by selecting R . In this case, $S_{i,o}(R)$ need not to be proper for D . In fact, we can replace $S_{i,o}(R)$ with any standard proper scoring rule $s(R_i)$ and traders are incentivized to report their true beliefs. Thus, if traders believe in the market, the decision maker can simply use a standard proper scoring rule!

Of course, as Othman and Sandholm's example shows, traders may have good reason to believe that the market will not get it right. In particular, traders near the close of the market may have an incentive to distort the probabilities. Luckily, for deterministic decision rules, Corollary 6.1 shows that the quasi-strictly proper scoring rules we derive are in fact proper scoring rules in this sense! Thus by using such a rule we can simultaneously provide myopic traders an incentive for truthful reporting in many, though not all, situations and provide traders with faith in the market an incentive for truthful reporting.

7.2 Differing Beliefs

We saw that one potential way around Othman and Sand-

holm's example is if some traders have a different belief about the final market prediction F . Another possibility is some traders have different beliefs about the true probability matrix. For example, consider a trader arriving with beliefs $(0.79, 0.74)$ and market probabilities $(0.8, 0.81)$. Depending on her beliefs about F , the trader has an incentive to change at least one of the probabilities to match her beliefs, so the market will no longer be "stuck" at $(0.8, 0.81)$.

7.3 Randomized Decision Rules

The negative example involves a deterministic decision rule. Potentially, randomized rules could have better incentive properties. However, simply adding randomness is not a panacea, as the following example shows.

Suppose $\mathcal{A} = \{1, 2\}$ and $\mathcal{O} = \{\top, \perp\}$. In section 4, we saw that the scoring rule $S_{i,\top} = \sum_j 2P_{j,\top} - P_{j,\top}^2$ and $S_{i,\perp} = \sum_j -P_{j,\top}^2$ is strictly proper for the decision rule $D_i(P) = P_{i,\top}/(P_{1,\top} + P_{2,\top})$. Suppose the current market probabilities are $(0.8, 0.7)$ and a myopic trader arrives with the belief $(0.8, 0.7)$. Ideally, we would like her to not make a prediction as the current market probabilities match her belief. However, it turns out to be optimal for her to report $(0.75, 0.75)$. More generally, we have the following lemma.

LEMMA 7.1. *Suppose $\mathcal{A} = \{1, 2\}$, $\mathcal{O} = \{\top, \perp\}$, $D_i(P) = P_{i,\top}/(P_{1,\top} + P_{2,\top})$, and $S_{i,\top} = \sum_j 2P_{j,\top} - P_{j,\top}^2$ and $S_{i,\perp} = \sum_j -P_{j,\top}^2$. Suppose the current market probabilities are P and a myopic trader arrives with belief P . The trader's optimal report is $Q_{i,\top} = (P_{1,\top} + P_{2,\top})/2$. Furthermore, suppose the current market probabilities are Q when the trader arrives. Then her optimal report is still Q .*

The proof is straightforward calculus and is therefore omitted. In many ways this example is worse than Othman and Sandholm's. Even if the market reaches the correct probabilities, the next trader to arrive will change them to values that give the decision maker no useful information. Furthermore, these uninformative values are stable. While this feature makes this particular randomized decision rule and scoring rule pair a poor choice for use in a decision market, the more general question is open.

OPEN PROBLEM 2. *Is there a randomized decision rule and corresponding scoring rule with good incentive properties for decision markets?*

7.4 Increasing Market Maker Loss

Another option is to consider a more drastic change to the design of the market. Othman and Sandholm's example shows that a decision market can get "stuck" with a prediction like $(0.8, 0.81)$ that no rational agent has an incentive to fix, because of the form of the expected payment of a myopic trader: $\sum_{i,o} D_i(R) Q_{i,o} (S_{i,o}(R) - S_{i,o}(P))$. Suppose we gave each trader only the side of the bet based on her prediction (i.e. $\sum_{i,o} D_i(R) Q_{i,o} (S_{i,o}(R))$ assuming she is myopic). Then if S is proper she would have an incentive to report her true probability rather than leaving the market at the current probability.

This approach loses the shared nature of market scoring rules and may create a large loss for the market maker. In particular, using a market scoring rule, if the initial prediction is P^0 and traders update this as P^1, \dots, P^f , the market maker's total payments to traders when action i is chosen

and outcome o occurs are $S_{i,o}(P^1) - S_{i,o}(P^0)$, $S_{i,o}(P^2) - S_{i,o}(P^1)$, \dots , $S_{i,o}(P^f) - S_{i,o}(P^{f-1})$, for a total of $S_{i,o}(P^f) - S_{i,o}(P^0)$. If traders' payments are changed to be based only on their own predictions, this property disappears.

While paying each trader based solely on her own prediction can be expensive, if this is done occasionally the loss may be acceptable. For example, once per hour or once per day the market maker could select a random trader to whom to make such an offer. The loss of the market maker is linear in the number of such offers made.

8. CONCLUSION

We examined the problem of information elicitation for decision making. One agent, a decision maker, wants to choose a distribution over a set of actions based on the probability distribution over outcomes for each action. Another agent, an expert, has a belief about these probabilities that the decision maker wants to elicit. Such elicitation is done through scoring rules. Othman and Sandholm [16] studied this problem for deterministic decision rules, with many of their results focusing on the MAX decision rule in particular. Our main result significantly generalized their results by providing a complete characterization of (strictly) proper scoring rules for arbitrary decision rules.

This characterization allowed us to give a sufficient condition for a decision rule to have a strictly proper scoring rule and a sufficient condition for no strictly proper scoring rule to exist. As these sufficient conditions do not cover all decision rules, an open problem remains. We also showed how our characterization allows us to derive quasi-strictly proper scoring rules in a number of cases where strictly proper scoring rules do not exist.

Finally, we discussed how the elicitation problem becomes more complicated when there are multiple experts, an observation also made by Othman and Sandholm [16]. A natural approach is to use a proper scoring rule for a decision rule to make a decision market, in the same way proper scoring rules are used to make prediction markets. However, this introduces two main problems. First, since only one decision is made in the end, an agent trading in the market has to base her decisions on beliefs about what the final market probabilities would be, a strategic problem with no parallel in prediction markets. Second, since no individual trader controls the final decision, scoring rules that encourage truthful revelation when they do have control no longer have the safe effect. We examined several ways this problem might be tackled in practice.

9. 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] G. W. Brier. Verification of forecasts expressed in terms of probability. *Monthly Weather Review*, 78(1):1–3, 1950.
- [3] J. Cervera and J. Munoz. Proper scoring rules for fractiles. *Bayesian Statistics*, 5:513–519, 1996.
- [4] 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, 2010.
- [5] E. Gerding, K. Larson, and N. Jennings. Eliciting expert advice in service-oriented computing. In *Proceedings of the 11th International Workshop on Agent-Mediated Electronic Commerce*, pages 29–42, 2009.
- [6] T. Gneiting and A. E. Raftery. Strictly proper scoring rules, prediction, and estimation. *J. of the American Statistical Association*, 102(477):359–378, March 2007.
- [7] I. J. Good. Rational decisions. *Journal of the Royal Statistical Society, Series B (Methodological)*, 14(1):107–114, 1952.
- [8] R. Hanson. Decision markets. *IEEE Intelligent Systems*, 14(3):16–19, 1999.
- [9] R. D. Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003.
- [10] R. D. Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):1–15, 2007.
- [11] N. Lambert, D. M. Pennock, and Y. Shoham. Eliciting properties of probability distributions. In *EC '08: Proceedings of the 9th ACM conference on Electronic commerce*, pages 129–138, 2008.
- [12] 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.
- [13] A. H. Murphy and R. L. Winkler. Probability forecasting in meteorology. *Journal of the American Statistical Association*, 79(387):489–500, 1984.
- [14] R. G. Nelson and D. A. Bessler. Subjective probabilities and scoring rules: Experimental evidence. *Journal of Agricultural Economics*, 71(2):363–369, 1989.
- [15] F. M. O'Carroll. Subjective probabilities and short-term economic forecast: An empirical investigation. *Applied Statistics*, 26(3):269–278, 1977.
- [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] D. M. Pennock and R. Sami. Computational aspects of prediction markets. In N. Nisan, T. Roughgarden, E. Tardos, and V. Vazirani, editors, *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [18] L. J. Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.
- [19] P. Shi, V. Conitzer, and M. Guo. Prediction mechanisms that do not incentivize undesirable actions. In *WINE '09: Proceedings of the 5th International Workshop on Internet and Network Economics*, pages 89–100, 2009.
- [20] D. J. Spiegelhalter. Probabilistic prediction in patient management and clinical trials. *Statistics in Medicine*, 5(5):421–433, 1986.
- [21] R. L. Winkler. The quantification of judgment: Some methodological suggestions. *Journal of the American Statistical Association*, 62(320):1105–1120, 1967.
- [22] R. L. Winkler. Scoring rules and the evaluation of probability assessors. *Journal of the American Statistical Association*, 64(327):1073–1078, 1969.
- [23] R. L. Winkler, J. Mu'oz, J. Cervera, J. Bernardo, G. Blattenberger, J. Kadane, D. Lindley, A. Murphy, R. Oliver, and D. R'sos-Insua. Scoring rules and the evaluation of probabilities. *TEST: An Official Journal of the Spanish Society of Statistics and Operations Research*, 5(1):1–60, June 1996.
- [24] J. Wolfers and E. Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, 18(2):107–126, 2004.