

Decision Rules and Decision Markets

Abraham Othman
Computer Science Department
Carnegie Mellon University
aothman@cs.cmu.edu

Tuomas Sandholm
Computer Science Department
Carnegie Mellon University
sandholm@cs.cmu.edu

ABSTRACT

We explore settings where a principal must make a decision about which action to take to achieve a desired outcome. The principal elicits the probability of achieving the outcome by following each action from a self-interested (but decision-agnostic) expert. We prove results about the relation between the principal's decision rule and the scoring rules that determine the expert's payoffs. For the most natural decision rule (where the principal takes the action with highest success probability), we prove that no symmetric scoring rule, nor any of Winkler's asymmetric scoring rules, have desirable incentive properties. We characterize the set of differentiable scoring rules with desirable incentive properties and construct one. We then consider decision markets, where the role of a single expert is replaced by multiple agents that interact by trading with an automated market maker. We prove a surprising impossibility for this setting: an agent can always benefit from exaggerating the success probability of a suboptimal action. To mitigate this, we construct automated market makers that minimize manipulability. Finally, we consider two alternative decision market designs. We prove the first, in which all outcomes live in the same probability universe, has poor incentive properties. The second, in which the experts trade in the probability of the outcome occurring unconditionally, exhibits a new kind of no-trade phenomenon.

Categories and Subject Descriptors

J.4 [Social and Behavioral Sciences]: Economics

General Terms

Economics, Theory

Keywords

Decision Rules, Decision Markets, Mechanism Design, Market Design, Prediction Markets, Elicitation

1. INTRODUCTION

Prediction markets have served as a reliable tool for estimating the winners of political elections and sports games [Berg et al., 2001]. However, due to legal restrictions severely limiting their use, the latest wave of prediction markets have focused not on the general, but on the specific. Instead of creating large-scale markets

Cite as: Decision Rules and Decision Markets, Abraham Othman and Tuomas Sandholm, *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 625–632

Copyright © 2010, International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

on publicly verifiable events, these services bill themselves as accumulators of organizational knowledge on specific internal (e.g., corporate) events. For instance, the prediction market startup *Inkling Markets* suggests that companies use their service “to uncover and quantify risk in your organization”. The recently launched *Crowdcast* lists questions like “When will your product really ship?” and “How much will it really cost?” on their homepage.

We contend that the current corporate prediction markets do not actually capture the problems most prominent for their clients. The chief challenge facing businesses is not whether the decisions they have made will prove successful, but rather what actions to take to assure the best chance of success.

Our work approaches this issue directly. In our model, a principal has to choose an action from a set of possibilities (e.g., “Hire additional sales staff”, “Double R&D funding”) in order to maximize the probability of achieving a desirable outcome (e.g., “Become the top-grossing company in our market space”, “Sell more than a million widgets”, or “Achieve profitability by the end of the year”). The principal elicits from an expert, for each action, the probability of achieving the desired outcome. Based on those probabilities, the principal then (deterministically) chooses an action according to a decision rule. Upon success or failure in achieving the desired outcome, the principal rewards the expert according to a pre-determined scoring rule. Three functions define the space: the decision rule, the payoff for success, and the payoff for failure.

We first prove general properties of the way these functions relate. We then provide an in-depth study of the most natural deterministic decision rule, the max decision rule, which selects the action that has the highest reported success probability. For the single-agent setting, we show that no symmetric scoring rule, nor the asymmetric ones from the literature, give the agent the right incentives. We construct an asymmetric scoring rule that does, and provide a characterization of a space of such rules.

Along the same lines as the original construction of market scoring rules [Hanson, 2003, 2007, Pennock and Sami, 2007], we continue by attempting to expand our scoring rule system into an automated market maker for our multiagent setting, to create *decision markets*. Surprisingly, we show that *every* market of this kind suffers from a peculiar type of manipulability, where an agent benefits from exaggerating the success probability of a suboptimal action. We show that this kind of manipulability applies even under an infinite stream of self-interested agents visiting the market—even if each agent's beliefs about the probabilities are exactly accurate. We design families of asymptotically-optimal pricing rules for decision markets that minimize manipulability.

Finally, we study two alternate decision market designs. We show the first one suffers from significant manipulability of a different kind, and for the second we show a new kind of no-trade

result. We conclude with a discussion of future research directions.

2. ELICITING FROM A SINGLE AGENT

We are ultimately interested in the study of decision markets—prediction markets used for decision support. These markets are multiagent tools for aggregating information in the form of prices which correspond to probabilities. Decision markets therefore have two components: a multiagent component, and a probability elicitation component. As a building block towards creating decision markets, in this section we study the elicitation of truthful probabilities from a single agent.

If we relax the need to elicit probabilities, the problem can be solved straightforwardly. For instance, it is trivial to incentivize an agent to truthfully report which action has the highest probability of achieving a desired outcome. This can be done by simply giving the agent a lump sum payment if the outcome is reached, which aligns the incentives of the agent and the principal. However, as we discuss in this section, expanding the query to include probabilities makes the elicitation problem much more complex.

We first describe the basic framework of the elicitation process. Then we explore the space of all deterministic decision rules, and finally characterize the space of scoring rules that should be used with the max decision rule.

2.1 Setting

In our setting, a principal cares about achieving beneficial outcome o , for example, unemployment under 8% by the beginning of next year. He needs to choose one of n actions, $D = \{d_1, \dots, d_n\}$, to try to achieve the outcome. He asks an expert to tell him the probability of achieving the outcome o under each of the alternative actions. Based on the agent's report, the principal makes his decision. The agent is risk neutral and does not care which action is taken or whether the outcome is achieved, except to the extent that it affects her compensation.

Definition 1. A (deterministic) *decision rule* is a function $\mathbb{D} : (0, 1)^n \mapsto D$ that maps an expert report to an action.

To avoid degeneracy, we assume that the decision rule maps a positive measure of space to each possible action.

Definition 2. A *scoring rule* is a pair of functions $f, g : (0, 1)^n \mapsto \mathfrak{R}$, such that an expert reporting $\mathbf{p} \in (0, 1)^n$ receives a payoff of $f(\mathbf{p})$ if the outcome is achieved and $g(\mathbf{p})$ if it is not.

Definition 3. A scoring rule is *continuous* if both $f(\cdot)$ and $g(\cdot)$ are continuous.

Definition 4. A scoring rule/decision rule pair is *strictly proper* if it strictly incentivizes truthful report of the expert's belief. That is, let the expert have belief \mathbf{p} , and let $u(\mathbf{p})$ represent the expert's expected utility from a report of \mathbf{p} . Then:

$$u(\mathbf{p}) > u(\mathbf{p}')$$

for all $\mathbf{p}' \neq \mathbf{p}$.

2.2 Results

We proceed to prove several restrictions on how scoring rules and deterministic decision rules interact.

THEOREM 1. *There does not exist any strictly proper scoring rule/decision rule pair.*

PROOF. Suppose there exists a deterministic decision rule \mathbb{D} and a strictly proper scoring rule for it. Now consider two belief vectors \mathbf{p}_1 and \mathbf{p}_2 such that they report the same probability p for the chosen action. Consider first an expert with true belief \mathbf{p}_1 . Since the scoring rule is strictly proper, we have

$$pf(\mathbf{p}_1) + (1-p)g(\mathbf{p}_1) > pf(\mathbf{p}_2) + (1-p)g(\mathbf{p}_2)$$

Now consider an expert with true belief \mathbf{p}_2 . Since the scoring rule is strictly proper, we have

$$pf(\mathbf{p}_2) + (1-p)g(\mathbf{p}_2) > pf(\mathbf{p}_1) + (1-p)g(\mathbf{p}_1)$$

These two equations cannot hold simultaneously, generating a contradiction. ■

We now relax the requirement of strict incentives to report the truthful estimates for all actions. We will only require that the incentive to report the truth is strict for the action chosen, and that there is no strict incentive to lie about the other probabilities.

Definition 5. Let the agent have true belief $\mathbf{p} = (p_1, \dots, p_n)$. Let d_i be the action taken by the deterministic decision rule \mathbb{D} if the agent reports truthfully. Let $u(\mathbf{p})$ represent the agent's expected utility from reporting \mathbf{p} given the decision rule/scoring rule pair. If

$$u(\mathbf{p}) \geq u(\mathbf{p}')$$

for all $\mathbf{p}' \neq \mathbf{p}$, and

$$u(\mathbf{p}) > u(\mathbf{p}^i)$$

where \mathbf{p}^i represents any report for which the i th component does not equal p_i (that is, the agent does not report $\mathbb{P}(o|d_i)$ truthfully). Then that scoring rule/decision rule pair is *quasi-strictly proper*.

THEOREM 2. *Let \mathbf{x} be a belief vector on the boundary of the decision rule \mathbb{D} between reports that map to action d_1 and those that map to d_2 . If there exists a quasi-strictly proper scoring rule for \mathbb{D} , then it must induce, for an agent with private belief \mathbf{x} , the same expected utility for d_1 or d_2 being taken.*

PROOF. Without loss of generality, assume $\mathbb{D}(\mathbf{x}) = d_1$. Since \mathbf{x} is on the boundary of \mathbb{D} , we have that for all $\epsilon > 0$ there exists a vector \mathbf{y} such that $\mathbb{D}(\mathbf{y}) = d_2$ and

$$|\mathbf{x} - \mathbf{y}| < \epsilon$$

Denote the expected utility of an expert with private belief \mathbf{p} reporting \mathbf{q} by $u_{\mathbf{p}}(\mathbf{q})$. Now suppose for contradiction that for all ϵ ,

$$u_{\mathbf{x}}(\mathbf{x}) > u_{\mathbf{x}}(\mathbf{y}),$$

i.e., the agent has different expected utilities for action d_1 or d_2 being taken. Since the rule is quasi-strictly proper, for an agent with true belief \mathbf{y} we must have

$$u_{\mathbf{y}}(\mathbf{y}) \geq u_{\mathbf{y}}(\mathbf{x})$$

We will show that there exists an ϵ that induces a \mathbf{y} such that

$$u_{\mathbf{y}}(\mathbf{x}) > u_{\mathbf{y}}(\mathbf{y})$$

Because the rule is quasi-strictly proper, the payoff from any report must be finite. This implies that because $|\mathbf{x} - \mathbf{y}| < \epsilon$, there exists some constant M such that for all \mathbf{z}

$$|u_{\mathbf{x}}(\mathbf{z}) - u_{\mathbf{y}}(\mathbf{z})| < M\epsilon$$

So we have some $\delta > 0$ such that for all $\epsilon > 0$

$$u_{\mathbf{y}}(\mathbf{x}) > u_{\mathbf{y}}(\mathbf{x}) - \delta > u_{\mathbf{y}}(\mathbf{y}) - 2M\epsilon$$

Now choose ϵ small enough so that

$$\delta > 2M\epsilon$$

At that ϵ we have

$$u_{\mathbf{y}}(\mathbf{x}) > u_{\mathbf{y}}(\mathbf{y})$$

This contradicts the rule being quasi-strictly proper. ■

This result holds for any payoff rule, including those that are discontinuous. If we restrict the payoff functions to be continuous, we can prove an additional restriction on the types of decision rules that admit quasi-strictly proper scoring rules. Specifically, the decision rule can only switch between actions when their reported success probabilities are equal.

THEOREM 3. *Let decision rule \mathbb{D} have a boundary between some two actions d_1 and d_2 where $\mathbb{P}(o|d_1) \neq \mathbb{P}(o|d_2)$. Then there exists no continuous quasi-strictly proper scoring rule for \mathbb{D} .*

PROOF. For shorthand, denote the boundary point ($\mathbb{P}(o|d_1)$, $\mathbb{P}(o|d_2)$) by (p, q) . We know from the proof of Theorem 2 that we must have, for an agent with belief (p, q) ,

$$pf(p, q) + (1 - p)g(p, q) = qf(p, q) + (1 - q)g(p, q)$$

Rearranging,

$$f(p, q)(p - q) - g(p, q)(p - q) = 0$$

If $p \neq q$, this holds only if

$$f(p, q) = g(p, q)$$

This means the expert receives the same payoff regardless of whether the outcome is achieved or not. Obviously, no rule of this kind can be quasi-strictly proper. This yields a contradiction. ■

We will now prove another property of quasi-strictly proper scoring rules. It will help us further narrow the space of scoring rules to consider.

Definition 6. In an *independent of irrelevant alternatives (IIA)* scoring rule, the expert can arbitrarily alter her report of the probabilities of the actions that are not taken without altering her expected utility. Formally, let \mathbf{p}_1 and \mathbf{p}_2 represent two reports such that $\mathbb{D}(\mathbf{p}_1) = \mathbb{D}(\mathbf{p}_2)$, and the reported probability of the chosen action d_i is equal (i.e., $p_1^i = p_2^i$). Letting $u(\mathbf{p})$ represent the agent's expected utility from reporting \mathbf{p} , $u(\mathbf{p}_1) = u(\mathbf{p}_2)$.

THEOREM 4. *For any decision rule, any non-IIA scoring rule incentivizes an agent with some beliefs to misreport her beliefs.*

PROOF. Consider a non-IIA scoring rule $(f'(\mathbf{p}), g'(\mathbf{p}))$. There exist two vectors \mathbf{p}_1 and \mathbf{p}_2 where the reported probability of the chosen action is equal and the expert's payoff from reporting \mathbf{p}_1 is not equal to that from reporting \mathbf{p}_2 . Then either an expert with belief \mathbf{p}_1 is incentivized to report \mathbf{p}_2 , or an expert with belief \mathbf{p}_2 is incentivized to report \mathbf{p}_1 . ■

This result allows us, without loss of generality, to represent scoring rules not as general functions of the belief vector, but as scalar functions only of the probability of success of the action taken (i.e., $f(p^{\mathbb{D}(\mathbf{p})})$ and $g(p^{\mathbb{D}(\mathbf{p})})$ instead of $f(\mathbf{p})$ and $g(\mathbf{p})$).

2.3 The max decision rule

We will now study the most natural decision rule that satisfies the condition of Theorem 2 (that the rule switches actions at equal probabilities), the max decision rule.

Definition 7. The *max decision rule* selects the action that has highest reported probability of achieving the outcome: $\arg \max_i \mathbf{p}$. Ties are broken in some fixed way.

The rest of this section will discuss creation of quasi-strictly proper scoring rules for the max decision rule. First, we show that symmetric scoring rules are not quasi-strictly proper for the max decision rule. Then we construct an asymmetric scoring rule that is quasi-strictly proper.

2.3.1 Inadequacy of symmetric scoring rules

As Winkler [1994] discusses, most of the literature on scoring rules has studied symmetric scoring rules only. These include the log rule

$$\begin{aligned} f(p) &= \log p \\ g(p) &= \log 1 - p \end{aligned}$$

and the quadratic rule

$$\begin{aligned} f(p) &= 1 - (1 - p)^2 \\ g(p) &= 1 - p^2 \end{aligned}$$

A symmetric scoring rule is defined as follows.

Definition 8. A *symmetric* scoring rule sets

$$f(1 - p) = g(p)$$

Symmetric scoring rules are intuitively desirable because they offer a pleasing equivalence between the positive and negative sides of an event occurring. Consider asking an expert whether it will rain tomorrow or not. Using a symmetric rule, an expert who reports p to the question "Will it rain tomorrow?" achieves the exact same payoff, regardless of the realized outcome, as an expert who reports $1 - p$ to the question "Will it *not* rain tomorrow?". This fits well with our intuitive understanding of how a scoring rule works—in each case the expert has supplied identical information to the principal, and is therefore compensated identically.

However, as the following example shows, symmetric scoring rules are not quasi-strictly proper for the max decision rule, because they can induce an expert to manipulate her report. The following example shows this for the most famous scoring rule.

EXAMPLE 1. *The principal can take one of two actions, d_1 or d_2 , in an effort to achieve beneficial outcome o . The expert agent has true belief:*

$$\mathbb{P}(o|d_1) = .5 \text{ and } \mathbb{P}(o|d_2) = .25$$

The interested party uses the max decision rule, and uses the log scoring rule, which is strictly proper for traditional elicitation:

$$f(x) = 1 + \log(x) \quad g(x) = 1 + \log(1 - x)$$

If the agent reports truthfully, action d_1 is taken. The agent receives expected utility

$$.5f(x) + .5g(x) = 1 + \log(.5) \approx .7$$

If the agent instead reports $\mathbb{P}(o|d_1) < .25 = \mathbb{P}(o|d_2)$, then action d_2 is taken (because it has the higher reported probability). The agent receives expected utility

$$.25f(x) + .75g(x) = 1 + (.25 \log(.25) + .75 \log(.75)) \approx .756$$

So the agent achieves higher expected utility by manipulating her response and causing the principal to take the inferior action.

We will now expand the manipulation shown in this example into a proof that covers every symmetric scoring rule, not just the log rule.

Let us denote the agent's truthful expected utility by $\mathbb{T}(p) \equiv pf(p) + (1-p)g(p)$. This is the agent's expected utility when her maximum probability among her believed probabilities is p and she reports (all the probabilities) truthfully.

THEOREM 5. *If a scoring rule is quasi-strictly proper for the max decision rule, then the truthful expected utility function $\mathbb{T}(p)$ is strictly increasing in p .*

PROOF. Suppose the function is not strictly increasing, so there exists $p_l < p_u$ such that

$$\mathbb{T}(p_l) \geq \mathbb{T}(p_u)$$

Then an expert with true belief (p_l, p_u) will not be disincentivized to instead report $(p_l, p_l - \epsilon)$. ■

(It is interesting that even though the scoring rule motivates an agent to report truthfully, the agent may wish she had a higher value to report.)

Using the theorem above, we can now prove a general impossibility about symmetric scoring rules.

THEOREM 6. *No symmetric scoring rule is quasi-strictly proper for the max decision rule.*

PROOF. We will show that no symmetric scoring rule admits a monotonically increasing truthful expected utility function, and thus cannot be quasi-strictly proper. Suppose there exists a symmetric quasi-strictly proper scoring rule that admits a monotonically increasing truthful expected utility function, and consider $p > .5$. Since the truthful expected utility function is strictly increasing, it must be the case that

$$\begin{aligned} \mathbb{T}(p) &> \mathbb{T}(1-p) \\ pf(p) + (1-p)g(p) &> (1-p)f(1-p) + pg(1-p) \\ pf(p) + (1-p)g(p) &> (1-p)g(p) + pf(p) \\ 0 &> 0 \end{aligned}$$

where the intermediate step is derived from the definition of a symmetric scoring rule. ■

2.3.2 Quasi-strictly proper scoring rules for the max decision rule

Since we cannot use a symmetric scoring rule for the max decision rule, we turn our attention to asymmetric rules. Unfortunately, the best-known family of asymmetric scoring rules, those devised by Winkler [1994], are not quasi-strictly proper and thus cannot be used for our setting. As Figure 2 of Winkler [1994] illustrates, the truthful expected utility function $\mathbb{T}(p)$ for any of Winkler's rules reaches a minimum in $(0, 1)$, and is thus not strictly increasing. By Theorem 5, this means none of Winkler's asymmetric rules are not quasi-strictly proper.

We proceed to completely characterize the set of differentiable quasi-strictly proper scoring rules, and show that the set is non-empty by constructing a quasi-strictly proper scoring rule for the max decision rule. The following result will guide our characterization.

THEOREM 7. *If an IIA, differentiable scoring rule has a strictly increasing truthful expected utility function $\mathbb{T}(p)$ and it strictly incentivizes an agent to truthfully report the probability of the maximum element of her belief, then it is quasi-strictly proper for the max decision rule.*

PROOF. Suppose for the purpose of contradiction that $\mathbb{T}(p)$ is strictly increasing and the scoring rule strictly incentivizes the agent

to report truthfully the probability of her maximum belief element, but the scoring rule is not quasi-strictly proper.

Because the scoring rule is not quasi-strictly proper, yet it strictly incentivizes the report of the probability of the maximum belief element, it must be the case that the agent is motivated to misreport a non-maximum element of her belief. Since the scoring rule is IIA and the principal is using the max decision rule, it must also be the case that the agent would manipulate her report of a non-maximum belief element to increase it beyond the probability of her maximum belief element.

Let the optimal action have probability p , the non-optimal action have true probability $p_l < p$, and the manipulated report be $p_u > p$. Because the agent wants to manipulate her report, we must have

$$pf(p) + (1-p)g(p) < p_l f(p_u) + (1-p_l)g(p_u)$$

Now consider instead if the agent had reported p_u as her maximum probability. She would have expected utility

$$pf(p_u) + (1-p)g(p_u) < pf(p) + (1-p)g(p)$$

Simplifying the above lines yields

$$p(f(p_u) - g(p_u)) < p_l(f(p_u) - g(p_u))$$

And because $p > p_l$, this implies $f(p_u) < g(p_u)$.

We will show that the supposition implies that $f(p) > g(p)$ for all $p \in (0, 1)$, yielding a contradiction.

Taking the derivative of $\mathbb{T}(p)$ with respect to p for a strictly increasing function yields

$$\mathbb{T}'(p) = f(p) + pf'(p) - g(p) + (1-p)g'(p) > 0$$

Now consider that the scoring rule strictly incentivizes truthful report of the selected action. That is, for all $p \in (0, 1)$,

$$\arg \max_{x \in (0,1)} pf(x) + (1-p)g(x) = p$$

The first order condition of this equation is

$$pf'(p) + (1-p)g'(p) = 0$$

By incorporating this into the equation for $\mathbb{T}'(p)$, we see that it simplifies to

$$f(p) > g(p)$$

In particular, take $p = p_u$ to get $f(p_u) > g(p_u)$, a contradiction. ■

The above proof tells us also how to characterize the set of quasi-strictly proper scoring rules. In particular, when combined with the encouragement of strictly truthful reporting of the maximum probability, the set of differentiable scoring rules with strictly increasing truthful utility functions are those for which

$$f(p) > g(p)$$

Furthermore, the set of differentiable rules which strictly encourage truthful reporting of the maximum probability are those for which

$$\frac{f'(p)}{g'(p)} = \frac{p-1}{p}$$

The differentiable scoring rules that satisfy these two equations are exactly those that are quasi-strictly proper for the max decision rule.

One scoring rule that satisfies these two conditions is

$$\begin{aligned} f(p) &= -(p-1)^2 + 1 \\ g(p) &= -p^2 \end{aligned}$$




Figure 1: An illustration of an asymmetric quasi-strictly proper scoring rule for the max decision rule. The upper line is the expert’s payoff if the objective is achieved, and the lower line is his payoff if it is not achieved.

By inspection, it is easy to see that this scoring rule meets the first-order condition of encouraging truthful report. Moreover, as is evident from Figure 1, $f(p) > g(p)$ for $p \in (0, 1)$, implying that the truthful expected utility function, $\mathbb{T}(p)$, is strictly increasing, and therefore the scoring rule is quasi-strictly proper for the max decision rule.

3. DECISION MARKETS

In this section, we discuss replacing our single expert with a pool of experts who produce a probability estimate through the mediation of an automated market maker. A prediction market is used to elicit estimates of the probability of achieving a desirable outcome through each possible action. Rather than acting on the report of an expert agent, the principal sets a deadline, examines the market prices at that deadline, uses them as probability estimates $\mathbb{P}(o|d_i)$ and takes the action d_i as if $\mathbb{P}(o|d_i)$ had been the probability report from the expert, as in the previous section.

The equating of prices and probabilities is standard and sound for risk-neutral agents because the contract in the market pays out a dollar if outcome o is achieved. Therefore, the expected payoff to an agent for the contract corresponding to the action taken is $\mathbb{P}(o|d_i)$, which is also the price that the agent would be willing to pay for it.

After the principal makes his decision d , the market continues, in that all of the trades placed on $\mathbb{P}(o|d)$ remain in effect, but the contract is renamed $\mathbb{P}(o)$. All contracts in the other markets are canceled, because $\mathbb{P}(o|d_i)$ has no meaning if $d \neq d_i$ has been selected. We might be concerned that expanding out to multiple different markets, one for each $\mathbb{P}(o|d_i)$, might leave agents hopelessly budget constrained. However, a trader who moves some amount of funds into the market can trade as if he had moved that money into the market for *every* contract, because only one of the markets will persist into the future.

3.1 Setting and prior work

The idea to use markets as tools in decision analysis stems from the work of Berg and Rietz [2003], who used conditional markets relating to the choice of vice-presidential nominees in the 1996 vice-presidential Iowa Electronic Markets. Wolfers and Zitzewitz [2006] also explore conditional prediction markets in an election context with an eye towards teasing out causations from correla-

tions. The difference between our work and these prior efforts is that those markets served as cameras, to capture the state of the outside world, rather than as engines, which shape the decisions made in the world. The markets we describe here tell the principal how to make his decision.

The conversion from scoring rules into automated market makers has been the topic of a flurry of recent research. It stems largely from the work of Hanson [2003, 2007]. A good overview of the automated market maker concept is given by Pennock and Sami [2007]. We provide a brief explanation of how these market makers work before discussing the impacts of their use.

Definition 9. A pricing rule is a continuously differentiable, monotonically increasing function $q(p) : (0, 1) \mapsto \mathfrak{R}$. If

$$\lim_{p \rightarrow 0} q(p) = -\infty$$

then that pricing rule is *unbounded from below*.

Since pricing rules are monotonically increasing they are invertible and we will sometimes refer to them as a function $p(q)$ which maps quantities into prices.

Agents interact in the marketplace by spending their money to obtain shares which will have value at expiry. Say an agent seeks to move the market from q aggregate demand to $q' > q$ aggregate demand (he is betting for the event in question to occur). The agent pays

$$\int_q^{q'} p(q) dq$$

and receives a payoff of $q' - q$ dollars if the event in question occurs. Now consider an agent wishing to bet against the event, moving the aggregate demand from q to $q'' < q$. That agent pays:

$$\int_{q''}^q 1 - p(q) dq$$

and receives a payoff of $q - q''$ if the event in question does not occur. Note that an agent who moves the market from price p to price p' and then back to price p has no expected benefit and no net cost—his costs are precisely equal to the payoff he is guaranteed to receive regardless of whether the event occurs or not.

Because every interaction with a pricing rule corresponds exactly to the payoffs of some strictly proper scoring rule, interactions with the automated market maker are trivially truthful for a myopic agent who holds a private belief (which she does not change based on others’ past and future bets).

The impossibility results of this section rely on examining the behavior of agents at the end of the trading period, where there are no strategic impacts from the arrival of other agents to the market. While there is no guarantee that a market mechanism with desirable incentive properties for its final participant will work well in practice, it is almost certainly true that any market mechanism with poor incentive properties at its close will exhibit pathologies in practice. (Furthermore, good incentive properties for the last agent can imply good incentive properties for *every* agent—for instance, Chen et al. [2007] demonstrated a plausible information aggregation setting in which agents are not incentivized to act strategically built on the fact that the last agent acting does not have these incentives.)

We will focus on decision markets operating under the max decision rule because it is the most natural decision rule to use. However, the model we describe in this section could be used to study any decision rule.

3.2 A general impossibility result for the max decision rule

The direct approach to creating an automated market maker with good incentive properties would be to convert the strictly-proper scoring rule for the max decision rule we constructed into a pricing rule. However, as we show next, not only that rule, but every pricing rule, suffers from manipulability where agents are motivated to misprice actions and force the principal to take a non-optimal action.

EXAMPLE 2. Consider an agent that arrives at the end of the trading period to a market with two possible actions, d_1 and d_2 . That agent has the knowledge that the probability of achieving the outcome by following each action is:

$$\mathbb{P}(o|d_1) = .8 \text{ and } \mathbb{P}(o|d_2) = .75$$

(Furthermore, say that these subjective probabilities are actually the correct objective probabilities—that is, our counterexample does not stem from the agent having inaccurate beliefs.) For the purposes of this example, the principal is employing the linear pricing rule $q(p) = p$. Suppose that the agent arrives to the market and sees the prices $(.8, .2)$. By bringing the market prices to the agent's true belief, $(.8, .75)$, the agent has zero expected utility. This is because the principal will pursue d_1 , which the agent holds no stake in. If the agent instead raises the price to $(.8, .8 + \epsilon)$, then she has positive expected utility: even though moving the market price beyond $.75$ incurs negative expected utility, this loss is more than countered by the expected gains from the agent's purchases at low prices (from $.2$ to $.75$). Figure provides a graphical intuition for the agent's expected utility. She gains positive utility $.5(.55^2)$ from raising the price to $.75$, but loses $.5(.05 + \epsilon)^2$ from raising the price from $.75$ to $.8 + \epsilon$. Her net utility is $\approx .15 > 0$.




Figure 2: A graphical depiction of a manipulating agent raising the probability of achieving the outcome if an action is taken from $.2$ to $.8$ given a true belief of $.75$. The agent makes positive expected utility for $\mathbb{P}(o|d_i) < .75$ and negative expected utility for $.75 < \mathbb{P}(o|d_i) \leq .8$. The areas expressed in the figure correspond to the amount of expected utility for the manipulator.

THEOREM 8. All pricing rules $q(p)$ are susceptible to a form of manipulation in which an agent raises the price of a non-optimal action higher than the best action.

PROOF. Let $\mathbb{P}(o|d_1) = p^*$ and $\mathbb{P}(p|d_1) = p^* + m/2$, where

$$z = \int_{q(\epsilon)}^{q(p^*)} p^* - p(x) dx$$

and m solves

$$\int_{q(p^*)}^{q(p^* + m/2)} q(p^* + m/2) - q(x) dx \leq z/2$$

Then an agent facing initial prices $(\epsilon, p^* + m/2)$ will have greater positive utility for moving prices to cause d_1 to be chosen rather than d_2 . ■

This type of manipulation can occur even in the face of an infinite stream of agents with accurate beliefs.

EXAMPLE 3. Imagine a second omniscient agent arriving in the market in our example after the participation of the first agent. That agent would not be motivated to change the prices he sees, and would leave the market at $(.8, .8 + \epsilon)$. By repeating this argument, one can see that even an infinite stream of agents would not correct the market prices, leading the principal to select an action less likely to generate the beneficial outcome.

An alternative manipulation for the agent in the first example would be to bring the price of d_2 to its true value and then artificially lower the price of d_1 below it. However, this manipulation would be rectified by the next agent and is not resilient in the way artificially raising d_2 beyond d_1 is.

3.3 Families of optimal pricing rules

Even though Theorem 8 shows all markets of this type can be manipulated, we note that the possibility of manipulation relies on an agent balancing the positive expected utility from bringing a price to its true value against the negative expected utility of raising it above its true value. In this section, we derive two families of pricing rules that minimize a risk-neutral principal's loss from manipulation, which is equivalent to minimizing the amount that an agent would be incentivized to distort the price.

Let the true maximum probability of achieving the principal's objective be p , and let the last mover in the market be an agent with correct beliefs. We do not assume anything about the behavior of the other market participants: market prices faced by the agent we are studying are arbitrary. Let U be the principal's utility of achieving the beneficial outcome, and let 0 be the principal's utility of not achieving the outcome, so a risk-neutral principal has expected utility pU from implementing the best action.

Definition 10. A pricing rule has an approximation ratio $k > 1$ if the expected utility of the principal after the participation of the manipulating agent is never less than pU/k .

Definition 11. A pricing rule has an additive approximation $k > 0$ if the expected utility of the principal after the participation of the manipulating agent is never less than $U(p - k)$.

THEOREM 9. If a pricing rule is unbounded from below, then it has an unbounded worst-case approximation ratio.

PROOF. A pricing rule unbounded from below admits a manipulating agent to derive an unbounded amount of positive expected utility by raising the price of an action with a price near 0. Formally, let there be two possible actions, d_1 and d_2 . Let the true probabilities (and our agent's beliefs) be $\mathbb{P}(o|d_1) = \frac{1}{2}$ and $\mathbb{P}(o|d_2) = \frac{1}{2k}$. Let the market prices before our agent participates be $\frac{1}{2}$ and ϵ . Because the pricing rule is unbounded from below,

$$\lim_{q' \rightarrow -\infty} \int_{q'}^{q(\frac{1}{2k})} 1/2k - p(q) dq = \infty$$

Since the pricing rule $q(p)$ has pre-image $(0, 1)$,

$$\int_{q(\frac{1}{2k})}^{\frac{1}{2}} p(q) - \frac{1}{2k} dq < \infty$$

Therefore, our agent will have positive expected utility from pushing the price of d_2 above $\frac{1}{2}$, causing action d_2 to be chosen, and d_2 is k times worse than d_1 for the principal. This holds for any k . Thus the pricing rule does not have finite approximation ratio. ■

This result means that the most commonly used pricing rule in practice, the logarithmic market scoring rule, has bad worst-case incentive properties, because it is unbounded from below [Hanson, 2007, Pennock and Sami, 2007]. We now proceed to derive families of optimal pricing rules that are bounded from below.

3.3.1 Multiplicatively optimal

A pricing rule $q(p)$ with approximation ratio $k > 1$ satisfies, for all p ,

$$\int_0^p q(x) dx \leq \int_p^{kp} q(kp) - q(x) dx$$

Since the two sides are trivially equal at $p = 0$, the following relationship of their derivatives with respect to p must hold:

$$\begin{aligned} q(p) &\leq p(k-1)kq'(kp) + q(p) - q(kp) \\ q(kp) &\leq p(k-1)kq'(kp) \end{aligned}$$

Substituting $q(p) = p^n$ yields

$$k = 1 + 1/n,$$

which is asymptotically optimal for large n .

3.3.2 Additively optimal

A pricing rule $q(p)$ with additive error k satisfies, for all p ,

$$\int_0^p q(x) dx \leq \int_p^{p+k} q(p+k) - q(x) dx$$

Since we have equality between the two sides at $p = 0$, it is sufficient to consider the derivatives of both sides with respect to p . By the rule for taking derivatives under the integral sign, we can differentiate both sides with respect to p to get

$$q(p) \leq q(p) + kq'(p+k) - q(p+k)$$

which simplifies to

$$q(p+k) \leq kq'(p+k)$$

Taking equality and solving the resulting differential equation shows that the pricing rule

$$q(p) = e^{p/k}$$

has additive approximation k , which is optimal as k approaches 0.

3.4 Alternative market designs

We have shown that our market format, regardless of the specific pricing rule used, is susceptible to producing a final price prediction which induces the market maker to take a non-optimal course of action. Put another way, it is possible for the market to get stuck in a set of prices that do not produce the optimal action. Though in the last section we developed families of asymptotically worst-case optimal pricing rules for our market format in an effort to mitigate this shortcoming, one might wonder whether some alternative format might have better properties.

In this section, we explore two alternative market designs, which we dub *joint elicitation* and *direct elicitation*.

3.4.1 Joint elicitation: $\mathbb{P}(o \cap d_i)$

One might think the impossibility results of the previous section were a function of the different contracts existing in separate probability universes, and that by collapsing the state space into a single universe, one might avoid manipulation.

In our *joint elicitation market*, there exist contracts for actions joined with the probability of the outcome being achieved (i.e., $\mathbb{P}(o \cap d_i)$). There is also a contract for estimating the chosen action, $\mathbb{P}(d_i)$. These prices enable the principal to determine the probabilities pertinent to his decision by using Bayes' rule:

$$\mathbb{P}(o|d_i) = \frac{\mathbb{P}(o \cap d_i)}{\mathbb{P}(d_i)}$$

The fundamental difference between a joint estimation and a conditional estimation market is that all the contracts in the joint markets exist within the same probability universe. That is, in a conditional elicitation market attempting to measure $\mathbb{P}(o|d_i)$, the agent receives a null outcome (i.e., does not have to pay and does not get paid anything) if d_i is not chosen by the principal. On the other hand, in a joint market attempting to measure $\mathbb{P}(o \cap d_i)$, the agent is guaranteed to receive a payoff regardless of the action ultimately taken.

Unfortunately, joint elicitation has poor incentives for the agent who is last to act in the market. She can bring the price of the currently winning action to zero (that is, both $\mathbb{P}(d_i) = 0$ as well as $\mathbb{P}(o \cap d_i) = 0$), and select another, arbitrary action (say, d_j) to bring to a non-zero probability. The agent benefits from this manipulation since neither d_i nor $o \cap d_i$ occur.

3.4.2 Direct elicitation: $\mathbb{P}(o)$

We now consider another market design, which we coin *direct elicitation*, in which the principal makes a decision based on the price of the contract that pays out if the *outcome* actually occurs, $\mathbb{P}(o)$.

Formally, the principal has a decision rule $\mathbb{D} : (0, 1) \mapsto D$, which maps the price of the outcome to an action. This rule can be published along with the market or not; all that matters for the impossibility results of this section is that the agent has an actionable belief about the principal's decision rule.

The key difference between a market of this form and a traditional prediction market is whether the principal is impacted by market prices, or whether the market serves only as a reflection of events, that is, whether the market is an engine or a camera. We contend that many of the existing corporate prediction markets in practice are in fact engines while they were designed as cameras. For example, if a principal sees a high probability that a product will be delayed months, perhaps he will devote additional expenditures to seeing it arrive on time, or maybe he will cancel the product entirely. The important thing is that the principal adjusts his behavior in response to market price. This is almost certainly the case for internal corporate markets—or else why pay a startup to run a prediction market for you in the first place?

Now, imagine that the principal can exert high effort h or low effort l to achieving his desirable outcome o . The agent knows that the true probabilities are $\mathbb{P}(o|h) = 1$ and $\mathbb{P}(o|l) = 0$. Let the market price be .5 before the agent acts (and that there are no other agents that will act after her, or that every agent that will act afterwards is rational). Imagine that the principal exerts effort as follows. If he observes market price $\mathbb{P}(o) > .5$, he exerts low effort (perhaps because such effort is costly), and ultimately fails to

achieve his goal. Similarly, if he observes price $\mathbb{P}(o) < .5$, he will exert high effort and achieve his goal. Now, consider how a rational agent will interact with such a market. If she buys and drives the price above .5, she will regret buying. If she sells and drives the price below .5, she will regret selling. Thus, she will not interact with the market.

Though our contradiction has been couched in the language of adverse selection, there is no loss of generality here. The decision to exert high or low effort is in line with any other decision which could positively or negatively affect the probability of achieving the goal.

Furthermore, the no-trade result also holds if an automated market maker were not present. In a regular limit order market, a rational agent would be willing to sell at any price under .5 and buy at any price greater than .5. If the market uses conventional order clearing rules—that any time a bid price exceeds an ask price a trade takes place at the earlier-placed price, no other trader in the market would match the agent’s orders.

Our no-trade result is fundamentally different than the classic no-trade result of Milgrom and Stokey [1982]. In their setting, rational agents are unable to “agree to disagree”, so speculative trade cannot exist. Purely speculative trade is zero-sum. Since a rational agent knows that no other rational agent would offer a counterparty a trade with positive expected value, no trade occurs. In contrast, our no-trade result stems from the fact prices are endogenous (self-referential): the price affects the action, which affects the success probability, which affects the price.

The difference between the settings is perhaps most clear if we imagine an automated market maker interacting with a rational risk-neutral agent with correct beliefs. In the Milgrom and Stokey setting, the agent would adjust the market price to the true probability of the event occurring. In our setting, the agent does not interact with the automated market maker at all.

4. CONCLUSION

We initiated the study of decision rules and decision markets in settings where a principal needs to select an action and is advised by a self-interested but decision-agnostic expert about the success probabilities of alternative actions.

We began by investigating the properties of general deterministic decision rules in the context of eliciting from a single expert. We proved results about the relations between the principal’s decision rule and the rules that specify the expert’s payoff if the desired outcome is, and is not, achieved. For the most natural decision rule (where the principal takes the action with highest success probability), we showed that no symmetric scoring rule, nor any of Winkler’s asymmetric scoring rules, are quasi-strictly proper. We characterized the set of differentiable quasi-strictly proper scoring rules and constructed an asymmetric scoring rule that is quasi-strictly proper.

Moving to decision markets where multiple experts interact by trading, we showed a surprising impossibility for every automated market maker, where an agent is incentivized to artificially raise the price of a non-optimal action (again under the decision rule where the principal takes the action with highest success probability). To counter this impossibility, we constructed two families of asymptotically optimal pricing rules against this form of manipulation, one additively optimal and the other multiplicatively optimal. Finally, we considered two alternative market designs for decision markets. The first, in which all outcomes live in the same probability universe, has even worse incentives (for the final participant). The second, in which the experts trade on the probability of the outcome occurring unconditionally, exhibits a new kind of no-trade

result.

There are many interesting research directions in this new area. The first would be examining the space of discontinuous scoring rules, being aware that such rules would nevertheless have to obey Theorem 2. Another area of exploration would be combinatorial decision rules, in which the principal could take more than one of the available actions.

We focused on deterministic decision rules because they are natural. It is known that adding randomness can significantly increase the incentive-compatible feasible space of mechanisms in other settings—prominent examples include voting [Gibbard, 1977] and sybill-proof (false name resistant) voting [Wagman and Conitzer, 2008]. Future research should study randomized decision rules in our setting as well. Trivial randomized solutions—like shifting ϵ probability to each action—do not get around our impossibilities, like the nonexistence of strictly proper scoring rules (because the manipulations in our examples and proofs yield the manipulator benefits that are not infinitesimally small).

It would also be interesting to explore a larger characterization of the peculiar no-trade impossibility of Section 3.4.2. Does it apply in other economic interactions besides markets of this sort? The key seems to be that any action an agent takes has negative utility because of the effects of taking that action.

Another important direction concerns what happens when agents are interested in the action taken by the principal. For instance, an expert could advocate doubling a company’s advertising budget because she works in the marketing department. This type of setting is related to recent work by Shi et al. [2009], who study a setting where an agent can perform an action after a market runs such that the market can incentivize the agent to act counter to the principal’s goal. Whether or not the agents can take actions after the decision market, it would be desirable to use decision markets to balance the utilities of agents impacted by the principal’s decision with the principal’s desire to achieve his goal.

Acknowledgements

This work is supported by NSF IIS-0905390.

References

- J. Berg and T. Rietz. Prediction markets as decision support systems. *Information Systems Frontiers*, 5(1):79–93, 2003.
- J. Berg, R. Forsythe, F. Nelson, and T. Rietz. Results from a Dozen Years of Election Futures Markets Research. *Handbook of Experimental Economics Results*, 2001.
- Y. Chen, D. Reeves, D. Pennock, R. Hanson, L. Fortnow, and R. Gonen. Bluffing and Strategic Reticence in Prediction Markets. In *WINE*, 2007.
- A. Gibbard. Manipulation of schemes that mix voting with chance. *Econometrica*, 45:665–681, 1977.
- R. Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003.
- R. Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):1–15, 2007.
- P. Milgrom and N. Stokey. Information, Trade and Common Knowledge. *Journal of Economic Theory*, 26(1):17–27, 1982.
- D. Pennock and R. Sami. Computational Aspects of Prediction Markets. In *Algorithmic Game Theory*, chapter 26, pages 651–674. Cambridge University Press, 2007.
- P. Shi, V. Conitzer, and M. Guo. Prediction Mechanisms That Do Not Incentivize Undesirable Actions. In *WINE*, 2009.
- L. Wagman and V. Conitzer. Optimal false-name-proof voting rules with costly voting. In *AAAI*, pages 190–195, 2008.
- R. Winkler. Evaluating probabilities: Asymmetric scoring rules. *Management Science*, pages 1395–1405, 1994.
- J. Wolfers and E. Zitzewitz. Five Open Questions About Prediction Markets. Technical Report 1975, Institute for the Study of Labor (IZA), Feb. 2006.