# Eliciting Forecasts from Self-interested Experts: Scoring Rules for Decision Makers

Craig Boutilier
Department of Computer Science
University of Toronto, Toronto, CANADA
cebly@cs.toronto.edu

## ABSTRACT

Scoring rules for eliciting expert predictions of random variables are usually developed assuming that experts derive utility only from the quality of their predictions. We study more realistic settings in which (a) the principal is a decision maker who takes a decision based on the expert's prediction; and (b) the expert has an inherent *interest* in the decision. Not surprisingly, in such situations, the expert usually has an incentive to misreport her forecast to influence the choice of the decision maker. We develop a general model for this setting and introduce the concept of a *compensation rule*. When combined with the expert's inherent utility for decisions, a compensation rule induces a *net scoring rule* that behaves like a traditional scoring rule. Assuming full knowledge of expert utility, we provide a complete characterization of all (strictly) proper compensation rules. We then analyze the case when the expert's utility function is not fully known to the decision maker. We show bounds on: (a) expert incentive to misreport; (b) the degree to which an expert will misreport; and (c) decision maker loss in utility due to such uncertainty. These bounds depend in natural ways on the degree of uncertainty, the local degree of convexity of net scoring function, and properties of the decision maker's utility function. Finally, we briefly discuss the use of compensation rules in prediction markets.

## Categories and Subject Descriptors

I.2 [**Artificial Intelligence**]: Multiagent Systems

## General Terms

Algorithms, Economics, Theory

## Keywords

prediction markets, mechanism design, incentives, decision theory

## 1. INTRODUCTION

Eliciting predictions of uncertain events from knowledgeable *experts* is a fundamental problem in statistics, economics, operations research, artificial intelligence and a variety of other areas [18, 4]. Increasingly, robust mechanisms for prediction are being developed, proposed and/or applied in real-world domains ranging from elections and sporting events, to events of public interest (e.g., disease spread or terrorist action), to corporate decision making. Indeed, the very idea of crowd-sourcing and information (or prediction) markets is predicated on the existence of practical mechanisms for information elicitation and aggregation.

Prediction mechanisms must provide an expert agent with incentives to reveal a forecast they believe to be accurate. Many forms of "outcome-based" *scoring rules*, either individual or market-based, provide experts with incentives to: (a) provide sincere forecasts; (b) invest effort to improve the accuracy of their personal forecasts; and (c) participate in the mechanism if they believe they can improve the quality of the principal's forecast. However, with just a few exceptions (see, e.g., [14, 12, 17, 15, 2, 6, 8, 5]), most work fails to account for the ultimate use to which the forecast will be put. Furthermore, even these models assume that the experts who provide their forecasts derive no utility from the final forecast, or how it will be used, except insofar as they will be rewarded by the prediction mechanism itself.

In many settings, this assumption is patently false: the principal is will often exploit the elicited forecast in order to make a *decision* [10, 12, 15, 2, 5]. In corporate prediction markets, the principal may base strategic business decisions on internal predictions of uncertain events. In a hiring committee, the estimated probability of candidates accepting offers influences the order in which (and whether) offers are made. Providing appropriate incentives in the form of scoring rules is often difficult in such cases [15, 2, 6]. However, just as critically, experts often have *their own interests* in specific decisions. For example, in corporate settings, an expert from a certain division may have an incentive to misreport demand for specific products, thus influencing R&D decisions that favor her unit. In a hiring committee, a member may misreport the odds that a candidate will accept a competing position in order to bias the "offer strategy" in a way that favors her preferred candidate.

In this work, we develop a formal model of scoring rules that incentivize truthful forecasts even when experts have an interest in the decisions taken by the principal. Naturally such rules must compensate experts for "sacrificing their own interests." One might respond by ignoring forecasts from experts with such conflicts. Unfortunately, decision makers often *must rely on the advice and predictions of experts who have some stake in their decisions*. This is especially true in organizations (e.g., corporate R&D decisions, faculty hiring, etc.), when advice from employees or group members is solicited; but it also holds in any case where the principal is *uncertain* about an expert's true interests. Our model and analysis, rather than ignoring the issue and hoping for the best, provides insight into how best to reward experts for their forecasts. Other work has studied both decision making and incentive issues in prediction markets [14, 10, 15, 2, 8, 7, 17], but only rarely addresses the natural question of expert self-interest in decisions [12, 5].

Our basic building block is a scoring rule for a single expert who knows the principal's *policy*—i.e., mapping from forecasts to decisions—and where the principal knows the expert's utility for decisions. We show that the scoring rule must compensate the expert in a simple, intuitive way based on her utility function: we use a *compensation function* that induces a proper scoring rule. We provide a complete characterization of proper compensation functions,

**Appears in:** *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2012)*, Conitzer, Winikoff, Padgham, and van der Hoek (eds.), 4-8 June 2012, Valencia, Spain.

as well those which, in addition, satisfy *participation constraints*.

We then analyze *expert uncertainty* in the principal's policy, and *principal uncertainty* in the expert's utility. First, we note that the expert need not know the principal's policy prior to providing her forecast as long as she can verify which decision has been taken after the fact. Second, we observe that, in general, the principal cannot ensure truthful reporting without full knowledge of the expert's utility function. However, principals will almost always have *some* partial knowledge of expert utility. We show that *bounds* on this uncertainty give rise bounds on each of the following: (i) the expert's incentive to misreport; (ii) the deviation of the expert's misreported forecast from her true beliefs; and (iii) the loss in utility the principal will realize due to this uncertainty. The first two bounds rely on the notion of *strong convexity* of the net scoring function. The third uses natural properties of the principal's utility function. We show that these bounds can be significantly tightened using *local* strong convexity, requiring only sufficient (and differential) convexity near the decision boundaries of the principal's policy. We conclude by briefly discussing a market scoring rule (MSR) based on our one-shot compensation rule.

## 2. BACKGROUND: SCORING RULES

We first review scoring rules and prediction markets (see [18, 4] for more details). We assume an agent—the *principal*—must assess the distribution of a discrete random variable $\mathcal{X}$ with domain $X = \{x_1, \ldots, x_m\}$. Let $\Delta(X)$ denote the set of distributions over $X$, where $\mathbf{p} \in \Delta(X)$ is a nonnegative vector $\langle p_1, \ldots, p_m \rangle$ s.t. $\sum_i p_i = 1$. The principle can engage one or more experts to provide a forecast $\mathbf{p} \in \Delta(X)$. We first assume a single expert $E$. Consider a running example: the chief strategy officer (CSO) of a company asks a division head to estimate demand for a new product prior to committing R&D efforts to that product.

We assume $E$ has beliefs $\mathbf{p}$ about $\mathcal{X}$. To incentivize $E$ to report $\mathbf{p}$ faithfully (and devote reasonable effort to developing *accurate* beliefs), a variety of *scoring rules* have been proposed [16, 13, 9]. A *scoring rule* is a function $S : \Delta(X) \times X \to \mathbb{R}$ that provides a score (or payoff) $S(\mathbf{r}, x_i)$ to $E$ if she reports forecast $\mathbf{r}$ and the realized outcome of $\mathcal{X}$ is $x_i$, essentially rewarding $E$ for her predictive "performance" [13]. If $E$ has beliefs $\mathbf{p}$ and reports $\mathbf{r}$, her expected score is $S(\mathbf{r}, \mathbf{p}) = \sum_i S(\mathbf{r}, x_i) p_i$. We say $S$ is a *proper scoring rule* iff a truthful report is optimal for $E$:

$$S(\mathbf{p}, \mathbf{p}) \geq S(\mathbf{r}, \mathbf{p}), \quad \forall \mathbf{p}, \mathbf{r} \in \Delta(X) \tag{1}$$

We say $S$ is *strictly proper* if inequality (1) is strict for $\mathbf{r} \neq \mathbf{p}$. A popular strictly proper scoring rule is the log scoring rule, where $S(\mathbf{p}, x_i) = a \log p_i + b_i$ (for arbitrary constants $a > 0$ and $b_i$) [13, 16]. In what follows, we restrict attention to *regular* scoring rules in which payment $S(\mathbf{r}, x_i)$ is bounded whenever $r_i > 0$.

Proper scoring rules can be fully characterized in terms of convex cost functions [13, 16]; here we review the formulation of Gneiting and Raftery [9]. Let $G : \Delta(X) \to \mathbb{R}$ be any convex function over distributions—we refer to $G$ as a *cost function*. We denote by $G^* : \Delta(X) \to \mathbb{R}^m$ some *subgradient* of $G$, satisfying

$$G(\mathbf{q}) \geq G(\mathbf{p}) + G^*(\mathbf{p}) \cdot (\mathbf{q} - \mathbf{p})$$

for all $\mathbf{p}, \mathbf{q} \in \Delta(X)$.[1] Such cost functions and associated subgradients can be used to derive any proper scoring rule.

THEOREM 1. *[13, 16, 9] A regular scoring rule $S$ is proper iff*

$$S(\mathbf{p}, x_i) = G(\mathbf{p}) - G^*(\mathbf{p}) \cdot (\mathbf{p}) + G_i^*(\mathbf{p}) \tag{2}$$

*for some convex $G$ and subgradient $G^*$. $S$ is strictly proper iff $G$ is strictly convex.*

[1]If $G$ is differentiable at $\mathbf{p}$ then the subgradient at that point is unique, namely, the gradient $\nabla G(\mathbf{p})$.

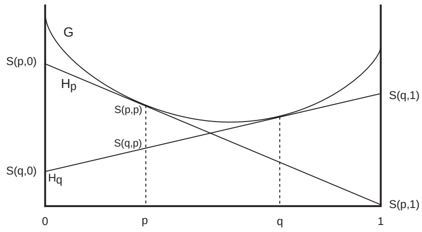


**Figure 1: Illustration of a proper scoring rule. If the expert reports $p$, her expected score (relative to her true beliefs) is given by the subtangent hyperplane $H_p$. For any report $q$ different from $p$, the expected score $S(q,p) = H_q \cdot p$ must be less than the expected score $S(p,p) = H_p \cdot p$ of truthful reporting.**

Intuitively, Eq. 2 defines a hyperplane

$$H_{\mathbf{p}} = \langle S(\mathbf{p}, x_1), \ldots, S(\mathbf{p}, x_m) \rangle,$$

for each point $\mathbf{p}$, that is subtangent to $G$ at $\mathbf{p}$. This defines a linear function, for any fixed report $\mathbf{p}$, giving the expected score of that report given beliefs $\mathbf{q}$: $S(\mathbf{p}, \mathbf{q}) = H_{\mathbf{p}} \cdot \mathbf{q}$. An illustration is given in Fig. 1 for a simple one-dimensional (two-outcome) scenario.

Several prediction market mechanisms allow the principal to extract information from multiple experts [18, 4]. *Market scoring rules (MSRs)* [11] allow experts to (sequentially) change the forecasted $\mathbf{p}$ using any proper $S$. An expert can change a forecast $\mathbf{p}'$ to $\mathbf{p}$ if she is willing to pay according to $S(\mathbf{p}', \cdot)$ and receive payment $S(\mathbf{p}, \cdot)$. If her true beliefs $\mathbf{p}$ differ from $\mathbf{p}'$ and the rule is strictly proper, then she has incentive to participate and report truthfully. Under certain conditions, MSRs can be interpreted as automated market makers [3]. Since each expert pays the amount due to the previous expert for her prediction, the net payment of the principal is the score associated with the final prediction.

Some prior work has studied incentives when prediction markets are used for decision making. Hanson [10] introduced the term *decision markets* to refer to the broad notion of prediction markets where experts offer forecasts for events conditional on some policy being adopted or a decision being taken. Osband [14] describes an interesting model for optimally incentivizing experts to "put effort" into deriving their forecasts. Othman and Sandholm [15] provide the first explicit, formal treatment of a principal who makes decisions based on expert forecasts. They address several key difficulties that arise due to the conditional nature of forecasts, but assume that the experts themselves have no direct interest in the decision that is taken. Chen and Kash [2, 6] extend this model to a wider class of informational settings and decision policies. Dimitrov and Sami [8] consider strategic behavior across multiple markets, where an expert may misreport her beliefs in one market to manipulate prices (and gain advantage) in another. Similarly, Conitzer [7] explores strategic aspects of prediction markets via connections to mechanism design, but again assumes that expert utility is derived solely from the mechanism's payoff. Shi et al. [17] consider experts that, once they report their forecasts, can take action to alter the probabilities of the outcomes in question. Unlike our model, they do not consider expert utility apart from the payoff offered by the mechanism (though, as in our model, the principal's utility function dictates the value of an expert report).

Hanson and Oprea [12] explicitly consider a single expert who has an interest in the final forecast of a prediction market and show that attempts to manipulate can in fact increase market accuracy (by incentivizing others to participate). Chen *et al.* [5] consider perfect Bayesian equilibria in a two-stage game with two participants predicting a boolean variable, one of whom has an interest in the final forecast. Both models, like ours, are motivated by expert interest in the principal's decision, but there are many important distinctions. Unlike our framework, neither approach explicitly models the principal's policy, utility, or loss due to manipulation. Both represent expert utility for the principal's decision indirectly through (very restrictive) payoff functions over the final forecast (either quadratic payoffs [12] or increasing payoffs [5]). Finally, neither model attempts to incentivize truthful reports by the manipulator.

# 3. SELF-INTERESTED EXPERTS

We now consider an expert who has a direct interest in the decision induced by her forecast. We devise a class of scoring rules that incentive self-interested agents to report their true beliefs. Such models are especially relevant in settings where expert opinions are sought from members internal to an organization. Rather than rejecting the forecasts of such experts, our model quantifies the impact of this self-interest and admits rules to circumvent it.

## 3.1 Model Formulation

A principal, or *decision maker (DM)*, elicits a forecast of $\mathcal{X}$ from expert $E$, and makes a decision based on this forecast. Let $D = \{d_1, \ldots, d_n\}$ be the set of possible decisions, and $u_{ij}$ be DM's utility should he take decision $d_i$ with $x_j$ ($j \leq m$) being the realization of $\mathcal{X}$. Letting $\mathbf{u}_i = \langle u_{i1}, \cdots, u_{im} \rangle$, the expected utility of decision $d_i$ given distribution $\mathbf{p}$ is $U_i(\mathbf{p}) = \mathbf{u}_i \cdot \mathbf{p}$. For any beliefs $\mathbf{p}$, DM takes the decision that maximizes expected utility, giving DM the *utility function* $U(\mathbf{p}) = \max_{i \leq n} U_i(\mathbf{p})$. Since each $U_i$ is a linear function of $\mathbf{p}$, $U$ is piecewise linear and convex (PWLC). Furthermore, each $d_i$ is optimal in a (possibly empty) convex region of belief space $D_i = \{\mathbf{p} : \mathbf{u}_i \cdot \mathbf{p} \geq \mathbf{u}_j \cdot \mathbf{p}, \forall j\}$. We assume DM acts optimally and that he has a policy $\pi : \Delta(X) \rightarrow D$ that selects some optimal decision $\pi(\mathbf{p})$ for any expert forecast $\mathbf{p}$. In what follows, we take $D_i = \pi^{-1}(d_i)$. We denote by $D_{ij}$ the (possibly empty) *boundary* between $D_i$ and $D_j$. Notice that for any $\mathbf{p} \in D_{ij}$ we must have $U_i(\mathbf{p}) = U_j(\mathbf{p})$. In our running example, the CSO is given a forecast probability $p$ of *high product demand* from the division head: he will authorize R&D if $p$ is above some threshold $\tau$, and abandon development if $p$ falls below $\tau$ (here $\tau$ is the indifference probability: $U_{\text{develop}}(\tau) = U_{\text{abandon}}(\tau)$).

Our model of DM utility is slightly more restricted than that of [15, 2], who allow the utility of each decision to depend on a different random variable, and assume that a variable will be observed *only if* the corresponding decision is taken. We also focus on principals that maximize expected utility given $E$'s report (i.e., DM uses the *max decision rule* [15]), though we discuss stochastic DM policies in Sec. 4.2.

Suppose expert $E$ is asked by DM to provide a forecast of $\mathcal{X}$. Assume that $E$ knows DM's policy $\pi$—knowledge of DM's utility function is sufficient but not required, see Sec. 4.1—and that $E$ has her own utility function or *bias* $b$, where $b_{i,j}$ is $E$'s utility should DM take decision $d_i$ and $x_j$ is the realization of $\mathcal{X}$. Define $\mathbf{b}_i = \langle b_{i,1}, \cdots, b_{i,m} \rangle$; and let $E$'s expected utility for $d_i$ given $\mathbf{p}$ be $B_i(\mathbf{p}) = \mathbf{b}_i \cdot \mathbf{p}$. For example, the division head (expert) may see increased corporate influence if R& D is authorized, but see her power wane if the product fails to materialize.

As with DM, $E$'s *optimal utility function* $B^*$ is PWLC:

$$B^*(\mathbf{p}) = \max_i \mathbf{b}_i \cdot \mathbf{p}. \tag{3}$$

Denote by $D^*(\mathbf{p})$ the decision $d_i$ that maximizes Eq. 3, i.e., $E$'s preferred decision given beliefs $\mathbf{p}$ (see Fig. 2).[2]

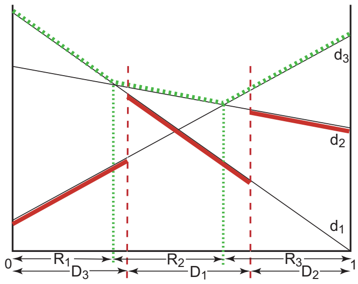


**Figure 2: Expert utility function. $E$'s utility for each decision $d_i$ is given by the corresponding hyperplane (here, thin black line). $E$'s "optimal" utility $B^*$ is the PWLC function shown by the dotted green line (i.e., the upper surface), with each $R_i$ denoting the regions of belief space where $D^*(\mathbf{p}) = d_i$. Regions $D_i$ represent DM's policy, where $\pi(\mathbf{p}) = d_i$ for $\mathbf{p} \in D_i$. The thick red lines denote the (discontinuous) utility $B^\pi$ that $E$ will receive from (truthfully) reporting her belief.**

Since DM is pursuing his own policy $\pi$, $E$'s actual utility for a specific report $\mathbf{r}$ under beliefs $\mathbf{p}$ is given by

$$B^\pi(\mathbf{r}, \mathbf{p}) = \mathbf{b}_{\pi(\mathbf{r})} \cdot \mathbf{p}; \tag{4}$$

that is, if she reports $\mathbf{r}$, DM will take decision $\pi(\mathbf{r}) = d_k$ for some $k$, and she will derive benefit $B_k(\mathbf{p})$. We refer to $B^\pi(\mathbf{r}, \mathbf{p})$ as $E$'s *inherent utility* for reporting $\mathbf{r}$. Similarly, $B(\mathbf{r}, x_i) = b_{i,\pi(\mathbf{r})}$ is $E$'s inherent utility for report $\mathbf{r}$ under realization $x_i$. This is the inherent benefit she derives from the decision she induces DM to take. This is illustrated in Fig. 2. Note that $E$'s utility for reports, given any fixed beliefs $\mathbf{p}$, is not generally continuous, with potential (jump) discontinuities at DM's decision boundaries.

Without some scoring rule, there is a clear incentive for $E$ to misreport her true beliefs to induce DM to take a decision that $E$ prefers, thereby causing DM to take a suboptimal decision. For instance, in Fig. 2, if $E$'s true beliefs $\mathbf{p}$ lie in $R_1$, her preferred decision is $d_1$; but truthful reporting will induce DM to take decision $d_3$. $E$ has greater inherent utility for reporting (any) $\mathbf{r} \in D_1$. Indeed, her gain from the misreport is $\mathbf{p} \cdot (\mathbf{b}_1 - \mathbf{b}_3)$. Equivalently, $E$ stands to lose $\mathbf{p} \cdot (\mathbf{b}_1 - \mathbf{b}_3)$ by reporting truthfully. Intuitively, a proper scoring rule would remove this incentive to misreport.

## 3.2 Compensation Rules

If DM knows $E$ utility function, he could reason about $E$'s incentive to misreport and revise his decision policy accordingly. Of course, this would naturally lead to a Bayesian game requiring analysis of its Bayes-Nash equilibria, and generally leaving DM with uncertainty about $E$'s true beliefs.[3] Instead, we wish to derive a scoring rule that DM can use to incentivize $E$ to report truthfully.

Unsurprisingly, such rules must compensate $E$ for the utility she foregoes by reporting her beliefs truthfully rather than influencing DM to act in a way that furthers $E$'s own interests. A *compensation function* $C : \Delta(X) \times X \rightarrow \mathbb{R}$ maps reports and outcomes into payoffs, like a standard scoring rule. However, $C$ does not fully determine $E$'s utility for a report; we must also account for the inherent utility $E$ derives from the decision she brings about. A compensation function $C$ induces a *net scoring function*:

$$S(\mathbf{p}, x_i) = C(\mathbf{p}, x_i) + B^\pi(\mathbf{p}, x_i) \tag{5}$$

[2]We assume that $D$ includes no dominated decisions; i.e., for any $d \in D$, we have $\pi^{-1}(d) \neq \emptyset$. If not, the subset of $D$ with only nondominated decisions should be used, since DM never needs to compensate $E$ for a decision he would never take.

[3]See Dimitrov and Sami [8] and Conitzer [7] for such a game-theoretic treatment of prediction markets (without decisions).

$E$'s expected net score for report $\mathbf{r}$ under beliefs $\mathbf{p}$ is $S(\mathbf{r},\mathbf{p}) = C(\mathbf{r},\mathbf{p}) + B^\pi(\mathbf{r},\mathbf{p})$, where $C(\mathbf{r},\mathbf{p}) = \sum_i p_i C(\mathbf{r}, x_i)$ is $E$'s expected compensation.

We adapt the definition of *proper* scoring rules to the case of compensation rules, recognizing that compensation is in full control of DM, while the net score is not:

DEFINITION 2. *A compensation function $C$ is* proper *iff the expected net score function $S$ satisfies $S(\mathbf{p},\mathbf{p}) \geq S(\mathbf{q},\mathbf{p})$ for all $\mathbf{p},\mathbf{q} \in \Delta(X)$. $C$ is* strictly proper *if the inequality is strict.*

One natural way to structure the compensation function is to use $C$ to compensate $E$ for the loss in inherent utility incurred by reporting her true beliefs $\mathbf{p}$ (relative to her best report), thus removing incentive for $E$ to misreport. This gives rise to a very specific compensation function $C_b$ that accounts for this loss:

$$C_b(\mathbf{p}, x_i) = b_{i,D^*(\mathbf{p})} - b_{i,\pi(\mathbf{p})}. \tag{6}$$

$C_b(\mathbf{p}, x_i)$ is simply the difference between $E$'s *realized* utility for her optimal decision (relative to her report $\mathbf{p}$) and the actual decision she induced. $C_b$ gives rise to the specific net scoring function:

$$\begin{aligned} S_b(\mathbf{p}, x_i) &= C_b(\mathbf{p}, x_i) + B^\pi(\mathbf{p}, x_i) & (7)\\ &= (b_{i,D^*(\mathbf{p})} - b_{i,\pi(\mathbf{p})}) + b_{i,\pi(\mathbf{p})}) = b_{i,D^*(\mathbf{p})} & (8)\end{aligned}$$

Since $E$'s expected net score under beliefs $\mathbf{p}$ is identical to her expected utility for the optimal decision $D^*(\mathbf{p})$, truthful reporting results, showing $C_b$ to be a proper compensation rule.

$C_b$ is just one straightforward mechanism for proper scoring with self-interested experts. We can generalize the approach to provide a complete characterization of all proper (and strictly proper) compensation functions. We derived $C_b$ by compensating $E$ for her *loss* due to truthful reporting. This is more "generous" than necessary: we need only *remove the potential gain* from misreporting. The key element of $C_b$ is not the "compensation term" $b_{i,D^*(\mathbf{p})}$, but the penalty term $-b_{i,\pi(\mathbf{p})}$, which prevents $E$ from benefiting by changing DM's decision. Any such gain is subtracted from her compensation via the penalty term $-b_{i,\pi(\mathbf{p})}$. We require only that the positive compensation term is convex: it need bear no connection to $E$'s actual utility function to incentivize truthfulness. Indeed, we can fully characterize the space of proper and strictly proper compensation functions:

THEOREM 3. *A compensation rule $C$ is proper for $E$ iff*

$$C(\mathbf{p}, x_i) = G(\mathbf{p}) - G^*(\mathbf{p}) \cdot \mathbf{p} + G^*_i(\mathbf{p}) - b_{i,\pi(\mathbf{p})} \tag{9}$$

*for some convex function $G$, and subgradient $G^*$ of $G$. $C$ is strictly proper iff $G$ is strictly convex.*[4]

An illustration of a cost function $G(\mathbf{p})$ that gives rise to a proper compensation function is shown in Fig. 3(a).

The fact that the specific rule $C_b$ is proper follows directly by observing that the net score $S_b$ can be derived from Eq. 2 by letting $G(\mathbf{p}) = B^*(\mathbf{p}) = \max_{i \leq n} B_i(\mathbf{p})$ be $E$'s optimal utility function (which is PWLC, hence convex), and using the subgradient $G^*(\mathbf{p})$ given by the hyperplane corresponding to the optimal decision $D^*(\mathbf{p})$ at that point.[5] Of course, $C_b$ is not *strictly* proper, since it is induced by a non-strictly convex cost function $G = B^*$. In particular, for any region $R(d)$ of belief space where a single decision $d$ is optimal for $E$, every report $\mathbf{p} \in R(d)$ has the same expected net score, hence there is no "positive" incentive for truthtelling.

[4]All proofs are available in working paper *arXiv:1106.2489*.

[5]At interior points of $E$'s decision regions, the hyperplane is the unique subgradient. At $E$'s decision boundaries, an arbitrary subgradient can be used.

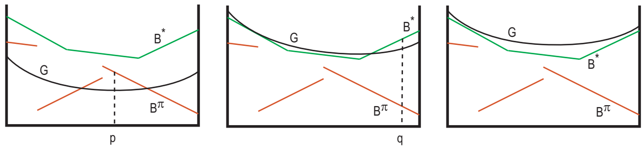


**Figure 3: Illustration of cost functions $G$ for strictly proper compensation rules. $E$'s optimal utility $B^*$ is the PWLC function shown in green, and $E$'s inherent utility $B^\pi$ is the discontinuous function in red. The net scoring function $G(\mathbf{p}) = S(\mathbf{p},\mathbf{p})$, the convex curve, induces an expected compensation function $C(\mathbf{p},\mathbf{p})$ by subtracting $B^\pi$. (a) A strictly proper rule that violates weak participation at point $p$. (b) A rule that satisfies weak participation but violates strong participation at point $q$. (c) A rule that satisfies strong participation.**

The characterization of Thm. 3 ensures truthful reporting, but may not provide incentives for participation. Indeed, the expert may be *forced to pay DM in expectation* for certain beliefs. Specifically, if $G(\mathbf{p}) < B^\pi(\mathbf{p})$, $E$'s expected compensation $C(\mathbf{p},\mathbf{p})$ is negative. Unless the DM can "force" $E$ to participate, this will cause $E$ to avoid providing a forecast if her beliefs are $\mathbf{p}$ (e.g., see point $p$ is Fig. 3(a)). In general, we'd like to provide $E$ with non-negative expected compensation. We can do this by insisting the compensation rule weakly incentives participation:

DEFINITION 4. *A compensation function $C$ satisfies* weak participation *iff for any beliefs $\mathbf{p}$, $E$'s expected compensation for truthful reporting $C(\mathbf{p},\mathbf{p})$ is non-negative.*

(Fig. 3(b) illustrates a cost function $G$ that induces a compensation rule $C$ satisfying weak participation.)

THEOREM 5. *A proper compensation rule $C$ satisfies weak participation iff it meets the conditions of Thm. 3 and $G(\mathbf{p}) \geq B^\pi(\mathbf{p})$ for all $\mathbf{p} \in \Delta(X)$.*

While desirable, weak participation does not ensure participation in general. Consider a compensation function defined with a convex cost function $G(\mathbf{p})$. If $E$ participates, she maximizes her net payoff by reporting her true beliefs $\mathbf{p}$. But suppose $G(\mathbf{p}) < B^*(\mathbf{p})$. While $E$ may not be certain how DM will act without her input (e.g., she may not know DM's "default beliefs" precisely), she may nevertheless have beliefs about DM's default policy. And, if $E$ believes DM will take decision $D^*(\mathbf{p})$ if she provides no forecast, then she is better off taking the expected payoff $B^*(\mathbf{p})$ based on her inherent utility and not participating (which has a lower payoff of $G(\mathbf{p})$). (See point $q$ in Fig. 3(b).) To prevent this, we say $C$ *strongly incentivizes* participation if, no matter what $E$ believes about DM's default policy (i.e., his action given no reporting), she will not sacrifice expected utility by participating in the mechanism.

DEFINITION 6. *A compensation function $C$ satisfies* strong participation *iff, for any decision $d_i \in D$, for any beliefs $\mathbf{p}$, $E$'s net score for truthful reporting is no less than $B_i(\mathbf{p})$.*

Strong participation means that $E$ has no incentive to abstain from participation (and need not "take her chances" that DM will make a decision she likes). This definition is equivalent to requiring that $E$'s expected utility for truthful reporting, as a function of $\mathbf{p}$, is at least as great as her optimal utility function, i.e., $S(\mathbf{p},\mathbf{p}) \geq B^*(\mathbf{p})$ for all $\mathbf{p} \in \Delta(X)$. Fig. 3(c) illustrates such a compensation rule.

THEOREM 7. *Proper compensation rule $C$ satisfies strong participation iff it meets the conditions of Thm. 3 and $G(\mathbf{p}) \geq B^*(\mathbf{p})$ for all $\mathbf{p} \in \Delta(X)$.*

OBSERVATION 8. *Compensation rule $C_b$ is the unique minimal (non-strictly) proper rule satisfying strong participation. That is, no compensation rule offers lower compensation for any report without violating strong participation.*

In general, if we insist on strong participation, DM must provide potential compensation up to the level of $E$'s maximum utility *gap*:

$$g(B) = \max_{i \le m, j, k \le n} b_{i,k} - b_{i,j}.$$

However, this degree of compensation is needed only if DM and $E$ have "directly conflicting" interests (i.e., DM takes a decision whose realized utility is as far from optimal as possible from $E$'s perspective). In such cases, one would expect $E$'s utility to be significantly less than DM's. If not, this compensation is not worthwhile for DM. Conversely, if $E$'s interests are well aligned with those of DM, the total compensation required will be small. The most extreme case of well-aligned utility is one where functions $\pi$ and $D^*$ coincide, i.e., $\pi(\mathbf{p}) = D^*(\mathbf{p})$ for all beliefs $\mathbf{p}$, in which case, no compensation is required. Specifically, compensation function $C_b(\mathbf{p}) = 0$ for all $\mathbf{p}$; and while $C_b$ is not strictly proper, the only misreports that $E$ will contemplate (i.e., that do not reduce her net score) are those that cannot change DM's decision (i.e., cannot impact DM's utility). As a consequence, DM should elicit forecasts from an expert who either (a) has well-aligned interests in the decisions being contemplated; (b) has interest whose magnitude is small (hence requires modest compensation) relative to DM's own utility; or (c) can be "forced" to make a prediction (possibly at negative *net* cost). Fortunately, these conditions often obtain in many settings, especially organizational or corporate settings. Employee incentives are usually reasonably well-aligned with those of corporate decision makers; and when external consultants are used, while their interests are not aligned with those of the principal, their stake in specific decisions is usually minimal.

# 4. POLICY AND UTILITY UNCERTAINTY

We now relax two key assumptions from Section 3.1: that $E$ knows DM's policy, and that DM knows $E$'s utility.

## 4.1 Policy Uncertainty

We first consider the case where DM does not want to disclose his policy to $E$. For example, suppose DM wanted to forego a truthful compensation rule $C$ and simply rely on a proper scoring rule of the usual form that ignores the $E$'s inherent utility. Thm. 3 shows that DM cannot prevent misreporting in general if he ignores $E$'s inherent utility; hence he can suffer a loss in his own utility. However, by refusing to disclose his policy $\pi$, DM could reduce the incentive for $E$ to misreport. Without accurate knowledge of $\pi$, $E$ would be forced to rely on uncertain beliefs about $\pi$ to determine the utility of a misreport, generally lowering her incentive. However, this will not remove the misreporting incentive completely. For instance, referring to Fig. 2, suppose DM does not disclose $\pi$. If $E$ believes *with sufficient probability* that the decision boundary between $d_3$ and $d_1$ is located at the point indicated, she will misreport any forecast $\mathbf{p}$ in region $D_3$ sufficiently close to that boundary should DM use a scoring rule rather than a compensation rule. As such, refusing to disclose his policy can be used by DM to reduce, but not eliminate, the incentive to misreport.[6]

Our analysis in the previous section assumed that $E$ used her knowledge of $\pi$ to determine the report that maximizes her net score. However, DM does not need to disclose $\pi$ to make good use of a compensation rule. He can specify a compensation rule *implicitly* by announcing his net scoring function $S(\mathbf{p}, x_i)$ (or the cost function $G$ and subgradient $G^*$) and promising to deduct $B_d \cdot \mathbf{p}$ from this score for whatever decision $d$ he ultimately takes. *$E$ need not know* in advance what decision will be taken to be incentivized to offer a truthful forecast. Nor does $E$ ever need to know what decisions *would have been taken* had she reported differently. Thus the only information $E$ needs to learn about $\pi$ is the value of $\pi(\mathbf{p})$ at her reported forecast $\mathbf{p}$; and even this need not be revealed until after the decision is taken (and its outcome realized).[7]

## 4.2 Uncertainty in Expert Utility

We now consider the more interesting issues that arise when DM is uncertain about the parameters $\mathbf{b}$ of $E$'s utility function. If the DM has a distribution over $\mathbf{b}$, one obvious technique is to specify a proper compensation rule using the expectation of $\mathbf{b}$. This may work reasonably well in practice, depending on the nature of the distribution; but it follows immediately from Thm. 3 that this approach will not induce truthful reporting in general.

Rather than probabilistic beliefs, we suppose that DM has *constraints* on $\mathbf{b}$ that define a bounded feasible region $\mathcal{B} \subseteq \mathbb{R}^{mn}$ in which $E$'s utility parameters must lie. We will confine our analysis to a simple, but natural class of constraints, specifically, upper and lower bounds on each utility parameter; i.e., assume DM has upper and lower bounds $b_{i,j}\Uparrow$ and $b_{i,j}\Downarrow$, respectively, on each $b_{i,j}$. This induces a hyper-rectangular feasible region $\mathcal{B}$. If $\mathcal{B}$ is a more general region (e.g., a polytope defined by more general linear constraints), our analysis below can be applied to the tightest "bounding box" of the feasible region.[8] Again by Thm. 3, DM cannot define a proper compensation rule in general: without certain knowledge of $E$'s utility, any proposed "deduction" of inherent utility from $E$'s compensation could mistaken, leading to an incentive to misreport. However, we show this incentive is bounded.

Requiring that DM have some information about $E$'s utility for DM's decisions may, at first glance, seem like too stringent a requirement. However, in many contexts, including organizational settings like those discussed above, it would, in fact, be highly unusual for this *not* to be the case. For instance, it would be unheard of for a CSO not to have some rough, albeit imprecise, idea of the benefit a division head would derive from undertaking R&D for a new product. More to the point, ignoring $E$'s potential biases makes it impossible for DM to have any confidence in her forecast. Our analysis sheds light on how much, and what type of, effort DM should invest in assessing $E$'s biases.

Under conditions of utility uncertainty, it is natural for DM to restrict his attention to "consistent" compensation rules:

DEFINITION 9. *Let $\mathcal{B}$ be the set of feasible expert utility functions. A compensation rule is* consistent *with $\mathcal{B}$ iff it has the form, for some (strictly) convex $G$ and $\tilde{\mathbf{b}} \in \mathcal{B}$:*

$$C(\mathbf{p}, x_i) = G(\mathbf{p}) - G^*(\mathbf{p}) \cdot \mathbf{p} + G^*_i(\mathbf{p}) - \tilde{b}_{i,\pi(\mathbf{p})}. \qquad (10)$$

Notice that consistent compensation rules are naturally linear: intuitively, we select a single consistent estimate of each parameter $\tilde{b_{i,j}} \in [b_{i,j}\Downarrow, b_{i,j}\Uparrow]$, treat $E$ as if this were her true (linear) utility function, and define $C$ using this estimate. We say DM is *$\delta$-certain* of $E$'s utility iff $b_{i,j}\Uparrow - b_{i,j}\Downarrow \le \delta$ for all $i, j$. Then we can bound the incentive for $E$ to misreport as follows:

[6]A similar argument shows that a stochastic policy can be used to reduce misreporting incentive, e.g., the *soft max* policy that sees DM take decision $d_i$ with probability proportional to $e^{\lambda u_i(\mathbf{p})}$. .

[7]Some mechanism to *verify* the decision *post hoc* may be needed.

[8]General linear constraints on $E$'s parameters could be could be inferred, for example, from observed behavior.

THEOREM 10. *If DM is $\delta$-certain of E's utility, then E's incentive to misreport under any consistent compensation rule is bounded by $2\delta$. That is, $S(\mathbf{r},\mathbf{p}) - S(\mathbf{p},\mathbf{p}) \leq 2\delta$.*

We can limit the misreporting incentive further by using a *uniform* compensation rule.

DEFINITION 11. *A consistent compensation rule is* uniform *if each parameter is estimated by $\tilde{b}_{i,\pi(\mathbf{p})} = \lambda b_{i,j}\!\Downarrow + (1-\lambda) b_{i,j}\!\Uparrow$ for some fixed $\lambda \in [0,1]$.*

For example, if DM uses the lower bound (or midpoint, or upper bound, etc.) of each parameter interval uniformly, we call the compensation rule uniform.

COROLLARY 12. *If DM is $\delta$-certain of E's utility, then E's incentive to misreport under any uniform compensation rule is bounded by $\delta$. That is, $S(\mathbf{r},\mathbf{p}) - S(\mathbf{p},\mathbf{p}) \leq \delta$.*

While bounding the *incentive* to misreport is useful, it is more important to understand the impact such misreporting can have on DM. Fortunately, this too can be bounded. The (strict) convexity of $G$ means that the greatest incentive to misreport occurs at the decision boundaries of DM's policy $\pi$ in Thm. 10. Since, by definition, DM is *indifferent* between the adjacent decisions at any decision boundary, misreports in a bounded region around decision boundaries have limited impact on DM's utility; the *amount* by which $E$ will misreport is bounded using the "degree of convexity" of the cost function $G$, which in turn bounds DM's utility loss.

DEFINITION 13. *Let G be a convex cost function with subgradient $G^*$. We say $G$ is* robust relative to $G^*$ with factor $m > 0$ *iff, for all $\mathbf{p},\mathbf{q} \in \Delta(X)$:*[9]

$$G(\mathbf{q}) \geq G(\mathbf{p}) + G^*(\mathbf{p}) \cdot (\mathbf{q} - \mathbf{p}) + m||\mathbf{q}-\mathbf{p}||_2 \qquad (11)$$

It is not hard to see that $m$-robustness of the pair $G, G^*$ imposes a minimum "penalty" on any expert misreport, as a function of its distance from her true beliefs:

OBSERVATION 14. *Let $C$ be a proper compensation rule based on an $m$-robust cost function $G$ and subgradient $G^*$. Let $S$ be the induced net scoring function. Then*

$$S(\mathbf{p},\mathbf{p}) - S(\mathbf{q},\mathbf{p}) \geq m||\mathbf{q}-\mathbf{p}||_2.$$

Together with Thm. 10, this gives a bound on the degree to which an expert will misreport when an uncertain DM uses a consistent compensation rule.

COROLLARY 15. *Let DM be $\delta$-certain of E's utility and use a consistent compensation rule based on an $m$-robust cost function and subgradient. Let $\mathbf{p}$ be E's true beliefs. Then the report $\mathbf{q}$ that maximizes E's net score satisfies $||\mathbf{q}-\mathbf{p}||_2 \leq \frac{2\delta}{m}$. If the compensation rule is uniform, then $||\mathbf{q}-\mathbf{p}||_2 \leq \frac{\delta}{m}$.*

In other words, $E$'s utility-maximizing report must be within a bounded distance of her true beliefs if DM uses an $m$-robust cost function to define the compensation rule.

The notion of $m$-robustness is a slight variant of the notion of *strong convexity* [1] in which we use the specific subgradient $G^*$ to measure the "degree of convexity." In the specific case of twice differentiable cost function $G$, we say $G$ is *strongly convex with factor $m$* iff $\nabla^2 G(\mathbf{p}) \succeq mI$ for all $\mathbf{p} \in \Delta(X)$; i.e., if the matrix $\nabla^2 G(\mathbf{p}) - mI$ is positive definite [1]. $m$-convexity is a sufficient condition for the robustness we seek.

[9] The definition of $m$-robustness can be recast using any reasonable metric, e.g., $L_1$-norm or KL-divergence; but the $L_2$-norm is most convenient below when we relate robustness to strong convexity.

COROLLARY 16. *Let DM be $\delta$-certain of E's utility and use a consistent compensation rule based on an $m$-convex, twice differentiable cost function $G$. Let $\mathbf{p}$ be E's true beliefs. Then the report $\mathbf{q}$ that maximizes E's net score satisfies $||\mathbf{q}-\mathbf{p}||_2 \leq \sqrt{\frac{4\delta}{m}}$. If the compensation rule is uniform, then $||\mathbf{q}-\mathbf{p}||_2 \leq \sqrt{\frac{2\delta}{m}}$.*

Robustness (or strong convexity) allows us to globally bound the maximum degree to which $E$ will misreport. This allows us to give a simple, global bound on the loss in DM utility that results from his uncertainty about the expert's utility function. Recall that DM's utility function $U_i$ for any decision $d_i$ is linear, hence has a constant gradient $\nabla U_i$. (We abuse notation and simply write $\nabla U_i$ for $\nabla U_i(\mathbf{p})$.) The function $U_i - U_j$ is also linear, given by parameter vector $(\mathbf{u}_i - \mathbf{u}_j)$. Let $\mathbf{e}_k$ denote the $n$-dimensional unit vector with a 1 in component $k$ and zeros elsewhere.

THEOREM 17. *Let DM be $\delta$-certain of E's utility and use a consistent compensation rule based on an $m$-robust cost function and subgradient. Assume $E$ reports to maximize her net score. Then DM's loss in utility relative to a truthful report by $E$ is at most $\max_k[\mathbf{e}_k^T \max_{i,j} \nabla(U_i - U_j)]\sqrt{n}\frac{2\delta}{m}$. If the compensation rule is uniform, then the bound is tightened by a factor of two.*

The same proof applies to strongly convex cost functions, with $\sqrt{n}\frac{2\delta}{m}$ replacing the term $\sqrt{n\frac{4\delta}{m}}$ in the bound above.

The results above all rely on the *global* robustness or *global* strong convexity of the cost function $G$. Designing a specific cost function (and if not differentiable, choosing its subgradients) can be challenging if we try to ensure uniform $m$-robustness or $m$-convexity across the entire probability space $\Delta(X)$. But recall that $E$ can only impact DM's utility if her misreport causes DM to change his decision. This means that the cost function *need only induce strong penalties for misreporting near decision boundaries*. Furthermore, the strength of these penalties should be related to the rate at which DM's utility is negatively impacted. For example, suppose $\mathbf{p}$ lies on the decision boundary between region $D_i$ and $D_j$. If $|\nabla(U_i - U_j)|$ is large, then a misreport in the region around $\mathbf{p}$ will cause a greater loss in utility than if $|\nabla(U_i - U_j)|$ is small. This suggests that the cost function should be more strongly convex (or more robust) near decision boundaries whose corresponding decisions differ significantly in utility, and can be less strong when the decisions are "similar." See Fig. 4 for an illustration of this point. Furthermore, the cost function need only be robust or strongly convex in a *local region* around these decision boundaries. In particular, suppose $G$ is $m$-robust in some local region around the decision boundary between $D_i$ and $D_j$. The degree of robustness bounds the maximum deviation from truth that $E$ will contemplate. If the region of $m$-robustness includes these maximal deviations, that will be sufficient to bound DM's utility loss for any true beliefs $E$ has in that region. Outside of these regions, no misreport by $E$ will cause DM to change his decisions (relative to a truthful report).

We can summarize this as follows:

DEFINITION 18. *$G$ is* locally robust relative to $G^*$ *in the $\varepsilon$-neighborhood around $\mathbf{p}$ with factor $m > 0$ iff, for all $\mathbf{q} \in \Delta(X)$ s.t. $||\mathbf{q}-\mathbf{p}||_2 \leq \varepsilon$:*

$$G(\mathbf{q}) \geq G(\mathbf{p}) + G^*(\mathbf{p}) \cdot (\mathbf{q} - \mathbf{p}) + m||\mathbf{q}-\mathbf{p}||_2 \qquad (12)$$

*Local strong convexity is defined similarly.*

Now suppose DM wishes to bound his loss due to misreporting by $E$ by some factor $\sigma > 0$. This can be accomplished using a locally robust cost function:

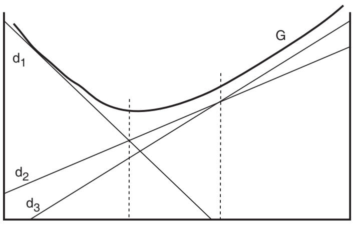


**Figure 4: A locally strongly convex cost function $G$. Here $G$ has is more strongly convex in the neighborhood of decision boundary $D_{12}$ than the boundary $D_{23}$. This means an expert willing to sacrifice compensation (e.g., to gain inherent utility due to DM uncertainty) can offer a report that deviates more from her true beliefs in the neighborhood around $D_{23}$ and than it can in the neighborhood around $D_{12}$ for the same loss in compensation. However, since $d_2$ and $d_3$ are more similar than $d_1$ and $d_2$ (w.r.t. DM utility), i.e., the gradient $|\nabla(U_2 - U_3)|$ is less than $|\nabla(U_1 - U_2)|$, DM will lose less utility "per unit" of misreport in the neighborhood of $D_{23}$. Note: the cost function $G$ is drawn above DM's utility function $U$ for illustration only—in general, it will lie below $U$.**

THEOREM 19. *Let DM be $\delta$-certain of E's utility and fix $\sigma > 0$. For any pair of decisions $d_i, d_j$ with non-empty decision boundary $D_{ij}$, define*

$$m_{ij} = \frac{\max_k(\mathbf{e}_k^T \nabla[U_i - U_j])\sqrt{n}2\delta}{\sigma}; \quad \varepsilon_{ij} = \frac{\sigma}{\max_k(\mathbf{e}_k^T \nabla[U_i - U_j])\sqrt{n}}.$$

*Let G be a convex cost function with subgradient $G^*$ such that, for all $i, j$ and any $\mathbf{p} \in D_{ij}$, (a) G is locally robust with factor $m_{ij}$ in the $\varepsilon_{ij}$-neighborhood around $\mathbf{p}$; (b) no other decision boundary lies within the $\varepsilon_{ij}$-neighborhood around $\mathbf{p}$. Let DM use a consistent compensation rule based on $G, G^*$. Assume E reports to maximize her net score. Then DM's loss in utility relative to a truthful report by E is at most $\sigma$. If the compensation rule is uniform, the result holds with $m_{ij}$ and $\varepsilon_{ij}$ decreased by a factor of two.*

This result can be generalized to the case where the degree of robustness around one decision boundary is relaxed sufficiently so that the neighborhood within which $E$ can profitably misreport crosses more than one decision boundary (i.e., when another decision boundary overlaps the $\varepsilon_{ij}$-neighborhood around $D_{ij}$). Utility loss will increase, but it can be bounded using the maximum gradient $\nabla(U_i - U_j)$ over decisions that can be swapped. The result can also be adapted to locally strongly convex cost functions.

These results quantify the "cost" to the decision maker of his imprecise knowledge of the expert's utility function, i.e., his worst-case expected utility relative to what he could have achieved if he had full knowledge of $E$'s utility (i.e., with truthful reporting by $E$). This analysis, however, does more than merely bound the risk facing a principal who solicits forecasts from self-interested experts. It also suggests: (a) ways in which the principal might expend effort to refine his knowledge of expert utility or bias; and (b) procedures for optimizing compensation rules when dealing with such experts. Regarding the first issue, while it goes beyond the scope of this paper, a more fine-grained analysis in the style of that used here—based on DM uncertainty regarding $E$'s *specific* utility parameters—can be used to justify and focus DM efforts when assessing $E$'s utility. On issue (b), the characterization of DM loss using local robustness or convexity has operational significance in the design of compensation rules. Indeed, it suggests an procedure for designing cost function $G$ (and induced compensation rule $C$) so as to minimize DM utility loss. Intuitively, $G$ should optimize two conflicting objectives: minimizing the bound $\sigma$ on utility loss (requiring an *increase* the degree of convexity at decision boundaries); and minimizing expected compensation $c$ (requiring a *decrease* in convexity). We believe specific classes of spline functions should prove useful for addressing this tradeoff.

Finally, note that if we relax the constraint that DM choose the decision $d_i$ with maximum expected utility, we can exploit local robustness to induce truthful forecasts. Suppose DM uses the softmax decision policy (see footnote 6): this stochastic policy makes $E$'s utility $B^\pi(\mathbf{r}, \mathbf{p})$ continuous in her report $\mathbf{r}$. An similar analysis similar using local convexity shows that DM induces truthtelling if the degree of convexity compensates for the gradient of $B^\pi$ at decision boundaries (since policy randomness removes the discontinuities in $B^\pi$). Of course, this comes at a cost: the DM is committed to taking suboptimal actions with some probability, leading to interesting tradeoffs between "acting optimally" but risking misleading reports vs. "acting suboptimally" given a truthful report.

# 5. MARKET SCORING RULES

We provide a brief sketch how to exploit compensation functions when DM aggregates the forecasts of multiple experts. One natural means of doing so is to use a *market scoring rule (MSR)* [11] that sequentially applies a scoring rule based on how an expert alters the prior forecast (see Sec. 2). An MSR based on a scoring rule $S$ has the $k$th expert pay the $k$–1st expert for her forecast according to $S$, and have the principal pay only final expert for her forecast using $S$. Thus, the principal's total payment is bounded by the maximal payment to a single expert [11]. When experts are self-interested, however, difficulties emerge; e.g., Shi et al. [17] show that experts who can alter the outcome distribution after making a forecast *each* require compensation to prevent them from manipulating the distribution to the detriment to the principal. A related form of *subsidy* arises in our decision setting.

Following [17], we assume a collection of $n$ experts, each of whom can provide alter the forecast $\mathbf{p}$ exactly once.[10] An "obvious" MSR in our model would simply adopt a proper compensation rule, and have each expert pay the either the *compensation* or the *net score* due to the expert who provided the incumbent forecast, and receive her payment from the next expert. If we use compensation, we run into strategic issues. With a proper compensation rule, an expert $k$ reports truthfully based on her *net score (total utility)*, consisting of both compensation and the inherent utility of the decision she induces. In a market setting, $k$'s proposed decision may be *changed* by the next expert's forecast. This (depending on her beliefs about other expert opinions) may incentivize $k$ to misreport in order to *maximize her compensation* rather than her net score. Overcoming such strategic issues seems challenging.

Alternatively, each expert might pay the net score due her predecessor. Unfortunately, an arbitrary proper compensation rule may not pay expert $k$ enough score to "cover her costs" (e.g., if $k$–1's inherent utility is much higher than $k$'s). However, if we set aside issues associated with incentive for participation for the moment, the usual MSR approach can be adapted as follows: we fix a *single (strictly) convex cost function G for all experts*, and define the compensation rule $C^k$ for expert $k$ using $G$ in the usual way:

$$C^k(\mathbf{p}, x_i) = G(\mathbf{p}) - G^*(\mathbf{p}) \cdot \mathbf{p} + G_i^*(\mathbf{p}) - b_{i,\pi(\mathbf{p})}^k,$$

where $\mathbf{b}^k$ is $k$'s utility function (bias). If $G$ satisfies strong participation for all experts (i.e., if $G(x_i) \geq B^*(x_i)$ for all $i$), then any

[10]This means we need not explicitly reason about how experts update their beliefs given the forecasts of others.

expert $k$ whose beliefs $\mathbf{p}[k]$ differ from the forecast $\mathbf{p}[k–1]$ provided by $k$–1 will have an expected net score (given $\mathbf{p}[k]$) greater than her expected payment to $k$–1 and will maximize her utility by providing a truthful forecast. In particular, denote $k$'s expected payment to $k$–1 by $\rho(k, k–1)$; then we have:

$$\begin{aligned}\rho(k, k–1) &= (H_{\mathbf{p}[k–1]} - \mathbf{b}^{k–1}_{\pi(\mathbf{p}[k–1])}) \cdot \mathbf{p}[k] \; + \; \mathbf{b}^{k–1}_{\pi(\mathbf{p}[k–1])} \cdot \mathbf{p}[k] \\ &= H_{\mathbf{p}[k–1]} \cdot \mathbf{p}[k] \\ &\leq H_{\mathbf{p}[k]} \cdot \mathbf{p}[k].\end{aligned}$$

Hence $k$'s expected payment $\rho(k, k–1)$ is less than her expected net utility, leaving her with a (positive) net gain of $(H_{\mathbf{p}[k]} - H_{\mathbf{p}[k–1]}) \cdot \mathbf{p}[k]$. However, this gain may be smaller than the inherent utility she derives from the decision induced by $k$–1, namely, $\mathbf{b}^{k}_{\pi(\mathbf{p}[k–1])} \cdot \mathbf{p}[k]$. Hence this scheme may not incentivize participation. In cases where DM can force participation, such a scheme can be used; but in general, the self-subsidizing nature of standard MSRs *cannot* be exploited with self-interested experts.

To incentivize participation, DM can subsidize these payments. In the most extreme case, DM simply pays each displaced expert her net utility, which removes any incentives to misreport, but at potentially high cost. In certain circumstances, we can reduce the DM subsidy to the market by having him pay only the inherent utility $b^{k–1}_{i,\pi(\mathbf{p}[k–1])}$ (given realized outcome $x_i$) of the displaced expert $k$–1, and requiring the displacing expert $k$ to pay the compensation $H_{i,\mathbf{p}[k–1]}$. Under certain conditions on the relative utility of different experts for different decisions, this is sufficient to induce participation; that is, $k$'s net gain for participating exceeds her inherent utility for the incumbent decision.

For instance, suppose all experts have the same utility function $\mathbf{b}$ (e.g., consider experts in the same division of a company who are asked to predict the outcome of some event, and have different estimates, but have aligned interests in other respects). In this case, $k$'s net gain for reporting her true beliefs is:

$$\begin{aligned}&(H_{\mathbf{p}[k]} - (H_{\mathbf{p}[k–1]} - \mathbf{b}_{\pi(\mathbf{p}[k–1])})) \cdot \mathbf{p}[k] \\ &\quad = (H_{\mathbf{p}[k]} - H_{\mathbf{p}[k–1]}) \cdot \mathbf{p}[k] + \mathbf{b}_{\pi(\mathbf{p}[k–1])} \cdot \mathbf{p}[k] \\ &\quad \geq \mathbf{b}_{\pi(\mathbf{p}[k–1])} \cdot \mathbf{p}[k].\end{aligned}$$

Hence $k$'s expected net gain is at least as great as her inherent expected utility for the decision induced by $k$–1, and strictly greater if her beliefs differ from those of $k$–1. Thus participation is assured.

Indeed, the argument holds even if the utility functions are not identical: we require only that $k$'s expected utility for the decision she displaces is less than the expected utility (given $k$'s beliefs) to be offered to her predecessor $k$–1. A sufficient condition for this is that $\mathbf{b}^k \leq \mathbf{b}^{k–1}$ (pointwise). This suggests that if the DM can elicit predictions of the experts in a particular order, he should do so by eliciting forecasts of those with the greatest utility first. Even with identical expert utility functions, there seems to be no escape from the requirement that DM subsidize the market at a level that grows linearly with the number of agents (as in [17]). Further development of MSRs in this setting should prove to be quite interesting.

# 6. CONCLUDING REMARKS

We have presented a model for the analysis of the incentives facing experts who have a vested interest in the decision taken by the principal, defining *compensation rules* that are necessary and sufficient to induce truthful forecasts and ensure participation. The analysis allows for uncertainty in the knowledge of both parties, exploiting various forms of robustness or convexity. We also provided some initial steps toward MSRs based on compensation rules. Rather than rejecting self-interested experts outright, our model allows the principal to assess the risks of using such experts, and design compensation to mitigates these risks.

Of course, our model and analysis are just first steps toward a comprehensive treatment of self-interested experts. Many interesting directions remain, including: the development of effective procedures for the design of cost functions that minimize utility loss and compensation when expert utility is unknown; the design and analysis of more refined market-scoring rules; finer-grained modeling of DM utility loss to guide DM's elicitation or assessment effort of expert utility/interests; and analyzing the tradeoffs when DM accepts restrictions on his possible decisions (potentially acting suboptimally) to reduce expert misreporting.

**Acknowledgements:** Thanks to Yiling Chen, Vince Conitzer, Ian Kash, Tuomas Sandholm, Jenn Wortman Vaughan and the anonymous reviewers for helpful discussions and suggestions. This work was supported by NSERC.

# 7. REFERENCES

[1] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.

[2] Y. Chen and I. Kash. Information elicitation for decision making. *AAMAS-11*, pp.175–182, Taipei, 2011.

[3] Y. Chen and D. Pennock. A utility framework for bounded-loss market makers. *UAI-07*, pp.49–56, 2007.

[4] Y. Chen and D. Pennock. Designing markets for prediction. *AI Magazine*, 31(4):42–52, 2011.

[5] Y. Chen, X. A. Gao, R. Goldstein, and I. A. Kash. Market manipulation with outside incentives. *AAAI-11*, pp.614–619, San Francisco, 2011.

[6] Y. Chen, I. A. Kash, M. Ruberry, and V. Shnayde. Decision markets with good incentives. *WINE-11*, pp.72–83, Singapore, 2011.

[7] V. Conitzer. Prediction markets, mechanism design, and cooperative game theory. *UAI-09*, pp.101–108, 2009.

[8] S. Dimitrov and R. Sami. Non-myopic strategies in prediction markets. *ACM EC'08*, pp.200–209, 2008.

[9] T. Gneiting and A. Raftery. Strictly proper scoring rules, prediction, and estimation. *J. Amer. Stat. Assoc.*, 102:359–378, 2007.

[10] R. Hanson. Decision markets. *IEEE Intelligent Systems*, 14:16–19, 1997.

[11] R. Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *J. Prediction Markets*, 1:3–15, 2007.

[12] R. Hanson and R. Oprea. A manipulator can aid prediction market accuracy. *Economica*, 76(302):304–314, 2009.

[13] J. McCarthy. Measures of the value of information. *PNAS*, 42(9):654–655, 1956.

[14] K. Osband. Optimal forecasting incentives. *J. Pol. Econ.*, 97(5):1091–1112, 1989.

[15] A. Othman and T. Sandholm. Decision rules and decision markets. *AAMAS-10*, pp.625–632, Toronto, 2010.

[16] L. J. Savage. Elicitation of personal probabilities and expectations. *J. Amer. Stat. Assoc.*, 66:783–801, 1971.

[17] P. Shi, V. Conitzer, M. Guo. Prediction mechanisms that do not incentivize undesirable actions. *WINE-09*, pp.89–100, 2009.

[18] J. Wolfers and E. Zitzewitz. Prediction markets. *J. Econ. Perspectives*, 18(2):107–126, 2004.