

Prediction Mechanisms That Do Not Incentivize Undesirable Actions^{*}

Peng Shi, Vincent Conitzer, and Mingyu Guo

Department of Computer Science, Duke University,
Durham, NC 27708, USA
ps46@duke.edu, {conitzer, mingyu}@cs.duke.edu

Abstract. A potential downside of prediction markets is that they may incentivize agents to take undesirable actions in the real world. For example, a prediction market for whether a terrorist attack will happen may incentivize terrorism, and an in-house prediction market for whether a product will be successfully released may incentivize sabotage. In this paper, we study *principal-aligned* prediction mechanisms—mechanisms that do not incentivize undesirable actions. We characterize all principal-aligned proper scoring rules, and we show an “overpayment” result, which roughly states that with n agents, any prediction mechanism that is principal-aligned will, in the worst case, require the principal to pay $\Theta(n)$ times as much as a mechanism that is not. We extend our model to allow uncertainties about the principal’s utility and restrictions on agents’ actions, showing a richer characterization and a similar “overpayment” result.

Key words: Prediction Markets, Proper Scoring Rules, Mechanism Design

1 Introduction

Prediction markets reward agents for accurately assessing the probability of a future event.¹ Typically, agents buy or sell securities according to their beliefs, and they are rewarded based on the outcome that materializes. Empirical studies suggest that prediction markets make very accurate predictions, sometimes beating the best experts and polls [1, 15]. Currently, online markets such as NewsFutures and Intrade elicit public predictions about a wide variety of topics, and many technology companies, including HP [12], Google, Microsoft, and Yahoo!, use in-house prediction markets to elicit employees’ predictions on future products.²

A potential downside of prediction markets is that they may incentivize agents to take undesirable actions in the real world, if those actions affect the probability of the event. For example, the idea of organizing markets to predict terrorist activity, which was once seriously considered by the U.S. Department of Defense, has been dismissed in part based on the consideration that terrorists may stand to profit from such markets.³ As another example, consider a software company that runs an in-house prediction market to assess whether a product will be released on

^{*} This work is supported by NSF IIS-0812113, the Sloan Foundation, and a Yahoo! Faculty Research Grant. We thank the anonymous reviewers for helpful comments.

¹ For literature reviews on prediction markets, see [11, 14, 15].

² NewsFutures and InkingMarkets both provide services to help companies run in-house prediction markets, thus making running these markets accessible to non-technology-based companies.

³ In July 2003, reports surfaced about a DARPA project to use prediction markets to guide policy decisions, and a possible topic was terrorist attacks. This ignited a political uproar and the proposal was quickly dropped. The arguments against it include creating incentives for a person to perform a violent act, as well as the distasteful thought of any person benefiting from such an attack. A survey of the proposal, the debate, and the aftermath can be found in [9].

time. The company may be concerned that the market provides an incentive to an employee to sabotage a timely release, if the employee predicts a late release.

On the other hand, not all real-world actions are necessarily undesirable. A terrorism prediction market may also incentivize an agent to *prevent* a terrorist attack, if she is predicting that no such attack will take place. Similarly, the software company’s in-house prediction market may incentivize an employee to work extra hard to finish her component of the product in time, if she is predicting that the product will be released on time.

The question that we study in this paper is the following: is it possible to design prediction mechanisms that do not incentivize undesirable actions? Here, an action is *undesirable* if it reduces the expected utility of the principal (center, organizer) of the prediction mechanism (e.g., the Department of Defense or the software company in the above examples). We call such mechanisms *principal-aligned* since they, in some sense, align the agents’ incentives with those of the principal.⁴

The rest of this paper is organized as follows. In Section 2, we study proper scoring rules, which incentivize a single agent to truthfully report her subjective probabilities about an event. After reviewing proper scoring rules and a known characterization theorem, we give a complete characterization of principal-aligned proper scoring rules. In Section 3, we consider settings with n agents. We show a negative “overpayment” result that indicates (roughly stated) that a principal-aligned prediction mechanism will, in the worst case, require the principal to pay n times as much as a mechanism that is not principal-aligned.

In Section 4, we extend our model to allow uncertainties about the principal’s utility, as well as restrictions on how agents’ actions can affect the underlying probabilities. We only want to disincentivize actions that are plausible under these restrictions and are definitely undesirable. We show how this provides richer structure to the class of principal-aligned proper scoring rules, and show that given sufficient uncertainty, any proper scoring rule can be transformed into a principal-aligned proper scoring rule by adding constant bonuses. We also show a similar overpayment result under the extended model.

2 Principal-Aligned Proper Scoring Rules

2.1 Review of Proper Scoring Rules

Let $\Omega = \{1, 2, \dots, m\}$ be the outcome space, with m possible outcomes. Let $\mathbb{P} = \{\mathbf{p} \in \mathbb{R}^m : 0 < p_i < 1, \sum_{i=1}^m p_i = 1\}$ be the set of probability distributions over the outcomes.⁵ We define the standard basis \mathbf{O}_i ($i = 1, 2, \dots, m$) as follows: \mathbf{O}_i is the vector in which the i -th element equals 1 and all other elements equal 0. Note that while \mathbf{O}_i is not in \mathbb{P} , all vectors in \mathbb{P} are in the span of the \mathbf{O}_i ’s.

⁴ We note that it is possible for the prediction mechanisms in this paper to incentivize *desirable* actions. However, in this paper we will not have an explicit model for the agent’s costs for taking desirable actions (such as putting in extra effort so that a product is released on time), and as a result we will not be able to solve for the agent’s optimal action. All we can say is that the agent’s optimal action will be no less desirable than it would have been without the existence of the prediction mechanism. Moreover, whatever action the agent decides to take, she will be incentivized to report the true distribution conditional on that action.

We also assume that the prediction itself does not have secondary effects on the agent’s utility (an example of this would be the case where the agent’s manager observes the prediction of an early release date, and as a result will punish the agent if the product is not released early).

⁵ We assume that the probability of any outcome is positive. This assumption helps us handle peculiar cases with discontinuities at the boundary, and makes the ensuing math more elegant. We can handle the edge cases by taking limits.

Definition 1 A scoring rule is a function $S : \mathbb{P} \times \Omega \rightarrow \mathbb{R}$. For each report $\mathbf{r} \in \mathbb{P}$ (on the underlying distribution) and each outcome $i \in \Omega$, it specifies a payment $S(\mathbf{r}, i)$. The expected payment \tilde{S} under the scoring rule S depends on both the report \mathbf{r} and the true probability distribution \mathbf{p} over the outcomes. \tilde{S} can be written as $\tilde{S}(\mathbf{r}, \mathbf{p}) = \sum_{i=1}^m S(\mathbf{r}, i)p_i$.

Definition 2 A scoring rule $S : \mathbb{P} \times \Omega \rightarrow \mathbb{R}$ is (weakly) proper if $\forall \mathbf{p}, \mathbf{r} \in \mathbb{P}, \tilde{S}(\mathbf{p}, \mathbf{p}) \geq \tilde{S}(\mathbf{r}, \mathbf{p})$. It is strictly proper if equality occurs if and only if $\mathbf{r} = \mathbf{p}$.

Definition 3 Given convex function $G : \mathbb{P} \rightarrow \mathbb{R}$, a subgradient⁶ is a vector function $G^* : \mathbb{P} \rightarrow \mathbb{R}^m$ such that $\forall \mathbf{x}, \mathbf{y} \in \mathbb{P}, G(\mathbf{y}) \geq G(\mathbf{x}) + G^*(\mathbf{x}) \cdot (\mathbf{y} - \mathbf{x})$. Given convex function $G : \mathbb{P} \rightarrow \mathbb{R}$, a subtangent hyperplane at $\mathbf{x} \in \mathbb{P}$ is a linear function $H_{\mathbf{x}} : \mathbb{P} \rightarrow \mathbb{R}$ such that $H_{\mathbf{x}}(\mathbf{x}) = G(\mathbf{x})$, and $\forall \mathbf{y} \in \mathbb{P}, H_{\mathbf{x}}(\mathbf{y}) \leq G(\mathbf{y})$.

The following characterization of proper scoring rules was discovered first by Savage [13], but the version shown here is due to Gneiting and Raftery [5]. The intuition behind the characterization is illustrated in Figure 1.

Theorem 1 (Gneiting & Raftery). Given a convex function $G : \mathbb{P} \rightarrow \mathbb{R}$ and a subgradient G^* , setting $H_{\mathbf{r}}(\mathbf{p}) = G(\mathbf{r}) + G^*(\mathbf{r}) \cdot (\mathbf{p} - \mathbf{r})$ defines a family of subtangent hyperplanes such that $H_{\mathbf{r}}$ is subtangent at \mathbf{r} . Setting $S(\mathbf{r}, i) = H_{\mathbf{r}}(\mathbf{O}_i)$ defines a proper scoring rule. i.e.,

$$S(\mathbf{r}, i) = G(\mathbf{r}) - G^*(\mathbf{r}) \cdot \mathbf{r} + G^*(\mathbf{r}) \cdot \mathbf{O}_i = G(\mathbf{r}) - G^*(\mathbf{r}) \cdot \mathbf{r} + G_i^*(\mathbf{r})$$

Conversely, any proper scoring rule can be written in terms of a subgradient of some convex function G in the above fashion. We call G the cost function for the rule.




Fig. 1. Geometric Intuition for Theorem 1: Every proper scoring rule corresponds to a convex cost function G . Suppose an agent reports \mathbf{r} ; then, the rule’s payment for each outcome is the corresponding intercept of the subtangent hyperplane at \mathbf{r} . Hence, for a given report, her possible expected payments are represented by the corresponding subtangent hyperplane, which is always below the upper envelope G . When the agent reports the true distribution \mathbf{p} , she attains the upper envelope G .

Any proper scoring rule S corresponds to a unique convex cost function G where $G(\mathbf{p}) = \tilde{S}(\mathbf{p}, \mathbf{p})$ is the maximum expected payment the agent can obtain if the true probability is \mathbf{p} . Conversely, any convex function corresponds to some proper scoring rule.⁷

⁶ G^* always exists. If G is differentiable at $\mathbf{x} \in \mathbb{P}$, then the subgradient at \mathbf{x} is the gradient: $G^*(\mathbf{x}) = (\nabla G)(\mathbf{x})$. Otherwise, there may be many choices of $G^*(\mathbf{x})$.

⁷ If G is not differentiable, then many families of subtangent hyperplanes can be specified, each of which corresponds to a proper scoring rule.

Currently, the cost function G is defined only in the open set \mathbb{P} ; we define G on the boundary of \mathbb{P} by taking limits. Since G is already continuous in \mathbb{P} ,⁸ this makes G continuous everywhere.

2.2 Principal-Aligned Proper Scoring Rules

We now develop the concept of “principal-aligned” proper scoring rules. Suppose that agents can take actions in the real world to change the underlying distribution of the actual outcome. A principal-aligned rule disincentivizes any action that harms the principal in expectation.

Formally, let $\mathbf{u} \in \mathbb{R}^m$ be a vector whose i th component is the principal’s utility for outcome i . Note that given the true distribution $\mathbf{p} \in \mathbb{P}$, the principal’s expected utility is $\mathbf{p} \cdot \mathbf{u}$.

Definition 4 A proper scoring rule S is aligned with vector \mathbf{u} if the cost function G satisfies: $\forall \mathbf{p}_1, \mathbf{p}_2 \in \mathbb{P}$, if $(\mathbf{p}_2 - \mathbf{p}_1) \cdot \mathbf{u} > 0$, then $G(\mathbf{p}_2) \geq G(\mathbf{p}_1)$. We call S strictly aligned if the inequality is always strict.

Since we have $\tilde{S}(\mathbf{p}_2, \mathbf{p}_2) = G(\mathbf{p}_2)$ and $\tilde{S}(\mathbf{r}, \mathbf{p}_1) \leq G(\mathbf{p}_1)$, this definition says that if the true probability is \mathbf{p}_2 , then the agent prefers reporting \mathbf{p}_2 over changing the true probability to \mathbf{p}_1 and reporting some \mathbf{r} .

We call \mathbf{u} uniform if $\mathbf{u} = \alpha \mathbf{1}$ for some α . Note that because $\mathbf{1} \perp (\mathbf{p}_2 - \mathbf{p}_1) \forall \mathbf{p}_1, \mathbf{p}_2 \in \mathbb{P}$, the above definition does not say anything when \mathbf{u} is uniform (that is, when the principal is indifferent among all outcomes).

Definition 5 A convex function $G : \mathbb{P} \rightarrow \mathbb{R}$ is non-decreasing with respect to direction $\mathbf{u} \in \mathbb{R}^m$ if $\forall \mathbf{x}_1, \mathbf{x}_2 \in \mathbb{P}$, $(\mathbf{x}_2 - \mathbf{x}_1) \cdot \mathbf{u} > 0$ implies $G(\mathbf{x}_2) \geq G(\mathbf{x}_1)$. It is strictly increasing if the above inequality is strict.

Lemma 1. If convex function $G : \mathbb{P} \rightarrow \mathbb{R}$ is non-decreasing with respect to non-uniform direction \mathbf{u} , then $G(\mathbf{x}) = g(\mathbf{x} \cdot \mathbf{u})$ for some single-variable non-decreasing convex function g . The statement remains true when non-decreasing is replaced by strictly increasing.




Fig. 2. Intuition behind Lemma 1. Consider the hyperplane normal to \mathbf{u} through \mathbf{x} . On one side of this hyperplane is \mathbb{P}_x , and for any \mathbf{y} in the region \mathbb{P}_x , $G(\mathbf{x}) \geq G(\mathbf{y})$. By continuity of G , $G(\mathbf{x})$ must be constant along the hyperplane (the boundary of \mathbb{P}_x).

⁸ A well-known fact about convex functions is that they are everywhere continuous in the interior of their domain.

Proof. For each $\mathbf{x} \in \mathbb{P}$, define the set $\mathbb{P}_{\mathbf{x}} = \{\mathbf{y} | \mathbf{y} \in \mathbb{P}, (\mathbf{x} - \mathbf{y}) \cdot \mathbf{u} > 0\}$. $\mathbb{P}_{\mathbf{x}}$ is non-empty because \mathbf{u} is not uniform.⁹ Moreover, $\mathbb{P}_{\mathbf{x}}$ is open because it is the intersection of open sets \mathbb{P} and $\{\mathbf{y} | (\mathbf{x} - \mathbf{y}) \cdot \mathbf{u} > 0\}$.

Since G is non-decreasing with respect to \mathbf{u} , $G(\mathbf{x}) \geq \sup_{\mathbf{y} \in \mathbb{P}_{\mathbf{x}}} G(\mathbf{y})$. But \mathbf{x} lies in the closure $\overline{\mathbb{P}_{\mathbf{x}}}$, so by continuity of G , $G(\mathbf{x}) \leq \sup_{\mathbf{y} \in \mathbb{P}_{\mathbf{x}}} G(\mathbf{y})$. This means that for all $\mathbf{x} \in \mathbb{P}$, $G(\mathbf{x}) = \sup_{\mathbf{y} \in \mathbb{P}_{\mathbf{x}}} G(\mathbf{y})$.

Note now that $\forall \mathbf{x}_1, \mathbf{x}_2 \in \mathbb{P}$, whenever $\mathbf{x}_1 \cdot \mathbf{u} = \mathbf{x}_2 \cdot \mathbf{u}$, we have $\mathbb{P}_{\mathbf{x}_1} \equiv \mathbb{P}_{\mathbf{x}_2}$, which implies that $G(\mathbf{x}_1) = G(\mathbf{x}_2)$. Hence, $G(\mathbf{x}) = g(\mathbf{x} \cdot \mathbf{u})$ for some single-variable function g . Moreover, g must be convex because G is convex, and g must be non-decreasing because G is non-decreasing w.r.t. \mathbf{u} . The above proof still holds if non-decreasing is replaced by strictly increasing. Figure 2 illustrates the intuition behind this lemma.

Combining Definition 4 and Lemma 1, we get the following characterization of principal-aligned proper scoring rules.

Theorem 2. *Given any non-uniform principal utility vector $\mathbf{u} \in \mathbb{R}^m$, every proper scoring rule S aligned with \mathbf{u} corresponds to a cost function of the form $G(\mathbf{p}) = g(\mathbf{p} \cdot \mathbf{u})$, where g is a single-variable non-decreasing convex function. (By Theorem 1, this implies $S(\mathbf{r}, i) = G(\mathbf{r}) - G^*(\mathbf{r}) \cdot \mathbf{r} + G_i^*(\mathbf{r})$,¹⁰ where G^* is a subgradient of G .) Conversely, all such rules are aligned with \mathbf{u} .*

The above statement remains true when aligned is replaced by strictly aligned, and non-decreasing is replaced by strictly increasing.

Note that $\mathbf{p} \cdot \mathbf{u}$ is the principal’s expected utility. The above theorem implies that for a fixed report, the agent’s expected reward is a non-decreasing function of the principal’s expected utility. Hence the interests of the agent and the principal are aligned.

3 Principal-Aligned Prediction Mechanisms

Now, suppose that we want to elicit predictions from n agents in a principal-aligned way. One possible method is to use a principal-aligned proper scoring rule for each agent, and allow agents to see previous agents’ reports. However, to incentivize agents to respond thoughtfully, each proper scoring rule requires some subsidy to implement, and hence this method requires $\Theta(n)$ subsidy. On the other hand, if we do not require principal-alignment, we can use a *market scoring rule* and implement this with $\Theta(1)$ subsidy.¹¹ In fact, we will show that this $\Theta(n)$ gap always exists between the “cheapest” principal-aligned prediction mechanism and the “cheapest” non-principal-aligned mechanism. We formalize this in the next section.

3.1 One-round Prediction Mechanisms

We first make a technical note: As shown in [2, 4], in all current implementations of prediction markets, agents may try to deceive others by giving false signals, to their own later profit. This

⁹ This is because $\exists \mathbf{y} \in \mathbb{P}$ such that $(\mathbf{x} - \mathbf{y}) \cdot \mathbf{u} \neq 0$, in which case either \mathbf{y} or $2\mathbf{x} - \mathbf{y}$ is in $\mathbb{P}_{\mathbf{x}}$.

¹⁰ When g is differentiable, $S(\mathbf{r}, i) = g(\mathbf{r} \cdot \mathbf{u}) + (u_i - \mathbf{r} \cdot \mathbf{u})g'(\mathbf{r} \cdot \mathbf{u})$.

¹¹ Under a market scoring rule, each agent is paid according to a proper scoring rule, but must pay the proper-scoring rule payment of the previous agent (*i.e.*, the i th agent’s expected payment is $\tilde{S}(\mathbf{r}_i, \mathbf{p}) - \tilde{S}(\mathbf{r}_{i-1}, \mathbf{p})$). As a result, from the perspective of the principal, almost all the payments cancel out, and the total amount that the principal must pay ($\sum_i \tilde{S}(\mathbf{r}_i, \mathbf{p}) - \tilde{S}(\mathbf{r}_{i-1}, \mathbf{p}) = \tilde{S}(\mathbf{r}_n, \mathbf{p}) - \tilde{S}(\mathbf{r}_0, \mathbf{p})$) depends only on the last prediction and not on the number of agents. See [6] for further explanations of market scoring rules.

makes analyzing incentives in a prediction market difficult without strong assumptions such as myopic agents.¹² Since our aim is to show a negative result about the minimum subsidy required in principal-aligned mechanisms, rather than to resolve such strategic issues, we focus on a *one-round* model in which agents can participate at most once, hence ruling out such strategic play. A negative result in this restricted model will carry over to general multi-agent prediction mechanisms.¹³

Definition 6 *We define a one-round prediction mechanism as follows. There is an event with a finite set of disjoint outcomes Ω , and n agents are asked to give a probability for each outcome. There is some rule which decides the order in which agents report their predictions. Define E_j as the set of agents whose reports agent j cannot influence. When agent j reports prediction \mathbf{r}_j and event $i \in \Omega$ occurs, she receives a payment of $f_j(\{\mathbf{r}_k | k \in E_j\}, \mathbf{r}_j, i)$.*

Definition 7 *A one-round prediction mechanism is truthful if regardless of others' reports, an agent's expected utility is maximized when she reports her true subjective probability vector \mathbf{p} . The mechanism satisfies voluntary participation if reporting the true \mathbf{p} always yields non-negative expected payment. The mechanism is feasible if it is truthful and satisfies voluntary participation.*

Almost all current prediction mechanisms, when we add the restriction that each agent participates at most once, fit into this framework.¹⁴

Lemma 2. *A one-round prediction mechanism is feasible if and only if for each agent j , the payment function f_j , holding fixed the reports $\mathbf{r}_{-j} = \{\mathbf{r}_k | k \neq j\}$, is a proper scoring rule S with a non-negative cost function. We call S the corresponding proper scoring rule in this situation.*

Definition 8 *A one-round prediction mechanism is aligned with principal utility vector \mathbf{u} if for every agent j and every combination of other agents' reports ($\mathbf{r}_{-j} = \{\mathbf{r}_k | k \neq j\}$), the corresponding proper scoring rule is aligned with \mathbf{u} .*

¹² [3] studies creating prediction markets that are fully incentive-compatible, without assuming myopic agents. However, the mechanisms proposed are quite different from typical prediction markets and it is not yet clear that they can be made to work in practice.

¹³ This is because in any general prediction mechanism the agents may choose to act sequentially in a one-round fashion. Hence, relaxing the one-round requirement can only increase the worst-case subsidy required. The notions of "one-round" and "subsidy required" are made precise later.

¹⁴ For example, a market scoring rule can be viewed as a one-round prediction mechanism with a sequential order of reporting. When agent j reports, she sees the reports of agents $k < j$ and her payment for outcome i is $f_j(\{\mathbf{r}_k | k < j\}, \mathbf{r}_j, i) = S(\mathbf{r}_j, i) - S(\mathbf{r}_{j-1}, i)$, where $S(\cdot, i)$ is a proper scoring rule.

Both Hanson's Market Maker [7, 6] and Pennock's Dynamic Parimutuel Market [11, 10], when agents participate at most once, are one-round prediction mechanisms. Agents act sequentially; when agent j acts, she buys shares or bets on outcomes, which in either model corresponds to some report of probabilities \mathbf{r}_j . When she acts she also observes the current state of the market, which is a function of the previous agents' reports $\{\mathbf{r}_k | k < j\}$. In Hanson's Market Maker, the payment function follows a market scoring rule, and thus does not depend on later agents' reports. In a Dynamic Parimutuel Market the situation is more complicated, but if we make Pennock's assumption in [10] that the current implicit payoff is the same as the final payoff, then the payment function also only depends on previous reports.

As another example, the weighted-score mechanism for self-financed wagering considered in [8] is a feasible one-round prediction mechanism in which agents report simultaneously. The payment function for j is the difference between a proper scoring rule based on j 's report and a weighted average of the other agents' proper scoring rules. (This allows the mechanism to be self-financed.) This fits into our model, because in this framework, no agent can influence another's report, so we can allow each agent's payment to depend on every other agent's report.

3.2 Minimum Subsidy for Principal-Aligned Prediction Mechanisms

We now formalize the notion that any principal-aligned prediction mechanism requires $\Theta(n)$ subsidy to implement. Here, the subsidy is the minimum amount the principal must have to be solvent in expectation, no matter what the true probability and agents' reports are.

Definition 9 *A one-round prediction mechanism requires subsidy M if $\forall \epsilon > 0$, for some true probability \mathbf{p} and some reports $\mathbf{r} = \{\mathbf{r}_j\}$, the total expected payment*

$$\sum_{i,j} p_i f_j(\{\mathbf{r}_k | k \in E_j\}, \mathbf{r}_j, i) \geq M - \epsilon$$

Before deriving our minimum subsidy result, we introduce a notion of minimum incentives. In order to elicit useful predictions, we cannot simply offer the trivial scoring rule $S \equiv 0$. We assume that each agent will give a thoughtful report only if she can gain $c > 0$ by reporting accurately. More precisely, to be meaningful, a prediction mechanism cannot allow agents to always obtain within c of the optimal expected payment by giving a constant report.

Definition 10 *We say a proper scoring rule provides incentive c if the difference between the greatest attainable expected payment and the greatest expected payment the agent can guarantee with some constant report r is at least c .¹⁵ A one-round prediction mechanism guarantees incentive c if for each agent j and each combination of others' reports \mathbf{r}_{-j} , the corresponding proper scoring rule provides incentive c .*

Lemma 3. *Given a non-uniform principal utility vector \mathbf{u} , let \mathbf{O}_{\max} correspond to an optimal outcome for the principal (i.e., it maximizes $\mathbf{O}_i \cdot \mathbf{u}$ among all i). If a proper scoring rule aligned with \mathbf{u} provides incentive c , and its cost function $G(\mathbf{p})$ is non-negative, then $G(\mathbf{O}_{\max}) \geq c$.*

Proof. Suppose on the contrary that $G(\mathbf{O}_{\max}) < c$. Let \mathbf{O}_{\min} correspond to the principal's worst outcome (i.e., it minimizes $\mathbf{O}_i \cdot \mathbf{u}$ among all i), and let $d = \|\mathbf{O}_{\max} - \mathbf{O}_{\min}\|$. Since \mathbf{u} is not uniform, $d > 0$. For $0 < \epsilon < d$, define $r = \frac{\epsilon}{d}\mathbf{O}_{\max} + \frac{d-\epsilon}{d}\mathbf{O}_{\min}$. Let G^* be the subgradient of G corresponding to S (in the sense of Theorem 1).

First note that $\tilde{S}(\mathbf{r}, \mathbf{O}_{\max}) \geq 0$. This is because by the definition of subgradients and by $\mathbf{u} \cdot (\mathbf{O}_{\min} - \mathbf{r}) < 0$, we have $G^*(\mathbf{r}) \cdot (\mathbf{O}_{\min} - \mathbf{r}) \leq G(\mathbf{O}_{\min}) - G(\mathbf{r}) \leq 0$. Using Theorem 1 and the fact that $(\mathbf{O}_{\min} - \mathbf{r})$ is collinear with $(\mathbf{r} - \mathbf{O}_{\max})$, we have $\tilde{S}(\mathbf{r}, \mathbf{r}) - \tilde{S}(\mathbf{r}, \mathbf{O}_{\max}) = G^*(\mathbf{r}) \cdot (\mathbf{r} - \mathbf{O}_{\max}) \leq 0$, which implies that $\tilde{S}(\mathbf{r}, \mathbf{O}_{\max}) \geq G(\mathbf{r}) \geq 0$.

Because the scoring rule provides incentive c and because of convexity, either

$$\begin{aligned} G(\mathbf{O}_{\max}) - \tilde{S}(\mathbf{r}, \mathbf{O}_{\max}) &\geq c \\ \text{or } G(\mathbf{O}_{\min}) - \tilde{S}(\mathbf{r}, \mathbf{O}_{\min}) &\geq c \end{aligned}$$

However, the first inequality cannot hold, because $G(\mathbf{O}_{\max}) < c$ by assumption and $\tilde{S}(\mathbf{r}, \mathbf{O}_{\max}) \geq 0$.

Moreover, the second inequality cannot hold for sufficiently small ϵ . This is because

$$c - 0 > G(\mathbf{O}_{\max}) - G(\mathbf{r}) \geq G^*(\mathbf{r}) \cdot (\mathbf{O}_{\max} - \mathbf{r}) = G^*(\mathbf{r}) \cdot (\mathbf{r} - \mathbf{O}_{\min}) \left(\frac{d - \epsilon}{\epsilon}\right)$$

¹⁵ In mathematical language, this states that \forall dummy report \mathbf{r} , \exists distribution \mathbf{p} such that reporting truthfully instead of reporting the dummy value \mathbf{r} pays off by at least c : $\tilde{S}(\mathbf{p}, \mathbf{p}) - \tilde{S}(\mathbf{r}, \mathbf{p}) \geq c$.

So as $\epsilon \rightarrow 0$, we need $G^*(\mathbf{r}) \cdot (\mathbf{r} - \mathbf{O}_{min}) \rightarrow 0$. This along with the identity $\tilde{S}(\mathbf{r}, \mathbf{O}_{min}) - \tilde{S}(\mathbf{r}, \mathbf{r}) = G^*(\mathbf{r}) \cdot (\mathbf{O}_{min} - \mathbf{r})$ implies that as $\epsilon \rightarrow 0$, $\tilde{S}(\mathbf{r}, \mathbf{O}_{min}) \rightarrow G(\mathbf{r})$. Because G is continuous and bounded below by 0, we have that as $\epsilon \rightarrow 0$, $G(\mathbf{r}) \rightarrow G(\mathbf{O}_{min}) \geq 0$, so $\tilde{S}(\mathbf{r}, \mathbf{O}_{min}) \rightarrow G(\mathbf{O}_{min})$. Hence, for sufficiently small ϵ , the second inequality also fails. Contradiction.

Therefore, $G(\mathbf{O}_{max}) \geq c$.




Fig. 3. Illustration of the proof of Lemma 3. As shown, the geometric requirement imposed by the rule being principal-aligned, the requirement of providing incentive c , and the requirement $G(\mathbf{r}) \geq 0$, force $G(\mathbf{O}_{max}) \geq c$.

Theorem 3. *Let \mathbf{u} be a non-uniform principal utility vector and let n be the number of agents. Any feasible one-round prediction mechanism that guarantees incentive c and is aligned with \mathbf{u} requires subsidy cn .*

Proof. For each agent j and each combination of others' reports, the corresponding proper scoring rule must have a non-negative cost function, must be aligned with \mathbf{u} , and must provide incentive c . By Lemma 3, $G(\mathbf{O}_{max}) \geq c$. Hence, if all agents report some \mathbf{r} that is arbitrarily close to \mathbf{O}_{max} , we get, by the continuity of convex function G , that the total expected payment can be arbitrarily close to $nG(\mathbf{O}_{max}) \geq cn$.

Remark 1. Suppose we sacrifice principal-alignment; then, we can implement the market scoring rule, based on, say, the quadratic scoring rule with cost function $G(\mathbf{p}) = c \frac{m}{m-1} (\sum_i p_i^2 - \frac{1}{m})$. This is a one-round prediction mechanism guaranteeing incentive c and requiring only subsidy c , which is the minimum possible.¹⁶ This yields the following implication of Theorem 3: *Suppose we want a prediction mechanism that guarantees incentive c ; the cheapest principal-aligned mechanism requires n times as much subsidy as the cheapest non-principal-aligned mechanism in the worst case.*

In practice, one way to effectively provide this $\Theta(n)$ subsidy is to tap into pre-established means of payments. For example, suppose a company wishes to host an in-house prediction market on whether a product will succeed, without providing incentives to sabotage the product. Suppose outcome 2 denotes success, 1 denotes failure. Consider a cost function $G = g(\mathbf{O}_2 \cdot \mathbf{p})$, where g is an arbitrary non-decreasing convex function. Let $c = g(1) - g(0) - g'(0)$. The company can implement the corresponding market scoring rule using this cost function, allowing each employee to participate at most once, and promise a bonus of c if the product succeeds to each employee. (Promising a bonus for good performance is a commonplace practice anyway.) This produces a principal-aligned prediction mechanism requiring a subsidy of $cn + O(c)$, which by Theorem 3 is asymptotically the best possible.

¹⁶ By definition, guaranteeing incentive c requires subsidy c .

4 Extensions: Uncertain Utilities and Restricted Actions

In practice, there may be uncertainty about the principal’s utility vector \mathbf{u} , and there may be restrictions on the change in underlying probabilities that the agents’ actions can bring about. In this section, we show that adding these features to our model provides a richer characterization. Moreover, if the agents can perform actions that are certainly undesirable, then principal-aligned prediction mechanisms still require $\Theta(n)$ subsidy. We formalize these concepts via the following definitions.

Definition 11 An action model is a function $A : \mathbb{P} \rightarrow 2^{\mathbb{P}}$ such that $\forall \mathbf{p} \in \mathbb{P}$, $A(\mathbf{p})$ is a convex set satisfying $\mathbf{p} \in A(\mathbf{p})$.¹⁷ Intuitively, if \mathbf{p} is the initial underlying probability vector, then for any $\mathbf{p}' \in A(\mathbf{p})$, there is some action that the agent can perform to change the probability vector to \mathbf{p}' (and the agent is not able to change it to any probability vector outside of $A(\mathbf{p})$).

Definition 12 Let T be a set of possible utility vectors for the principal and let A be an action model. A proper scoring rule is aligned with T under A if $\forall \mathbf{p} \in \mathbb{P}$ and $\forall \mathbf{p}' \in A(\mathbf{p})$ such that $\mathbf{u} \cdot \mathbf{p}' < \mathbf{u} \cdot \mathbf{p} \forall \mathbf{u} \in T$,¹⁸ the cost function $G(\mathbf{p}) \geq G(\mathbf{p}')$. A one-round prediction mechanism is aligned with T under A if the corresponding proper scoring rule is always aligned with T . In all following references to T , we assume that T is not always uniform: $T \setminus \{\alpha \mathbf{1}\} \neq \emptyset$. (The principal is not definitely indifferent among all outcomes.)




Fig. 4. Example of uncertain utilities and restricted actions: citizens are asked about 3 possible outcomes: no terrorist attack, a successful attack, and a failed attack. The range of the principal’s possible utilities, and the corresponding utility set T projected onto \mathbb{P} , are shown. Under this action model, citizens can greatly increase the probability of an attack (by committing their own attack), but have much less influence over the success probability of an attack. (This depends heavily on security and other factors that citizens presumably can affect less.)

¹⁷ We require $\mathbf{p} \in A(\mathbf{p})$ because the agent can always choose to do nothing. Moreover, $A(\mathbf{p})$ is convex because if by committing action a , the agent can change the underlying probability to $\mathbf{q} \in A(\mathbf{p})$, and by committing a' , the agent can change it to \mathbf{q}' , then by committing a with probability λ and a' with probability $(1 - \lambda)$, the agent can change the underlying probability to $\lambda \mathbf{q} + (1 - \lambda) \mathbf{q}'$.

¹⁸ Here, we only guard against actions that harm the principal’s expected utility for *all* possible vectors $\mathbf{u} \in T$. Alternatively we might have guarded against actions that are bad for *any* setting of \mathbf{u} , but this case is not interesting because the proof of Lemma 1 implies that the only proper scoring rule aligned with two linearly independent utility vectors is the trivial rule with $G \equiv 1$.

Definition 13 Given a set T of possible utilities for the principal and an action model A for the agents, the corresponding bad direction function is $B : \mathbb{P} \rightarrow 2^{\mathbb{R}^m}$,

$$B(\mathbf{p}) = \{\mathbf{p}' - \mathbf{p} \mid \mathbf{p}' \in A(\mathbf{p})\} \cap \left(\bigcap_{\mathbf{u} \in T} \{\mathbf{x} \mid \mathbf{x} \cdot \mathbf{u} < 0\} \right)$$

We say that B is contained in a strict cone if $\exists \epsilon > 0$ and $\exists \mathbf{y}$ such that $\forall \mathbf{p} \in \mathbb{P}$ and $\forall \mathbf{x} \in B(\mathbf{p})$, $\mathbf{y} \cdot \mathbf{x} \geq \epsilon \|\mathbf{x}\| \|\mathbf{y}\|$.

Define the magnitude of the worst action as $d = \sup\{\|\mathbf{x}\| \mid \mathbf{x} \in B(\mathbf{p}), \mathbf{p} \in \mathbb{P}\}$. Note that $d = 0$ implies that agents cannot perform any action that is certainly bad.

Note that B is the intersection of various half-planes, and the set of differences of vectors in A , so $B(\mathbf{p})$ is always a convex set. Moreover, if T projected onto \mathbb{P} is $m - 1$ dimensional, then B is contained in a strict cone.

4.1 Principal-Aligned Proper Scoring Rules in the Extended Model

The following characterization follows immediately from Definition 13.

Theorem 4. Given a set T of possible utility vectors and an action model A , let the corresponding bad direction function be B . A proper scoring rule is aligned with T under A if and only if its cost function G satisfies $\forall \mathbf{p} \in \mathbb{P}$, and $\forall \mathbf{p}'$ with $\mathbf{p}' - \mathbf{p} \in B(\mathbf{p})$, $G(\mathbf{p}') \leq G(\mathbf{p})$.

Definition 14 A proper scoring rule is bounded if $S(\mathbf{r}, i)$ is bounded, that is, $\exists C$ s.t. $|S(\mathbf{r}, i)| \leq C \forall \mathbf{r} \in \mathbb{P}, i \in \Omega$. Equivalently (by Theorem 1), a proper scoring rule is bounded if and only if both the cost function $G(\mathbf{p})$ and the corresponding subgradient $G^*(\mathbf{p})$ are bounded.

The following theorem shows that under sufficient uncertainty, any proper scoring rule can be modified to be principal-aligned by adding constant bonuses.

Theorem 5. Suppose that the bad direction function B is contained in a strict cone; then, any bounded¹⁹ proper scoring rule S can be modified to be principal-aligned by adding constants $\{k_i\}$ so that $S'(\mathbf{r}, i) = S(\mathbf{r}, i) + k_i$. Moreover, $\|\mathbf{k}\| = \frac{M}{\epsilon}$, where M is an upper bound²⁰ on the norm of the subgradient $\|G^*(\mathbf{p})\|$, and ϵ is as in Definition 13.

Proof. Because M upper-bounds $\|G^*(\mathbf{p})\|$, we have that $\forall \mathbf{r}', \mathbf{r} \in \mathbb{P}$,

$$\|G(\mathbf{r}) - G(\mathbf{r}')\| \leq M \|\mathbf{r} - \mathbf{r}'\|$$

Let \mathbf{y} be such that $\mathbf{y} \cdot \mathbf{x} \geq \epsilon \|\mathbf{x}\| \|\mathbf{y}\| \forall \mathbf{x} \in B(\mathbf{r}) \forall \mathbf{r} \in \mathbb{P}$. Let $\mathbf{k} = -\frac{M}{\epsilon \|\mathbf{y}\|} \mathbf{y}$. Now if S' is associated with cost function G' , we have $\forall \mathbf{r}' - \mathbf{r} \in B(\mathbf{r})$,

$$\begin{aligned} G'(\mathbf{r}') - G'(\mathbf{r}) &= G(\mathbf{r}') - G(\mathbf{r}) + (\mathbf{r}' - \mathbf{r}) \cdot \mathbf{k} \\ &\leq G(\mathbf{r}') - G(\mathbf{r}) - \epsilon \frac{M}{\epsilon} \|\mathbf{r}' - \mathbf{r}\| \\ &\leq \|\mathbf{r}' - \mathbf{r}\| (M - M) \\ &= 0 \end{aligned}$$

which implies that S' is principal-aligned by Theorem 4.

¹⁹ In practice, we can make any unbounded proper scoring rule bounded by restricting reports away from the boundary.

²⁰ M exists by the assumption that the scoring rule S is bounded. See Definition 14.

4.2 Principal-Aligned Prediction Mechanisms in the Extended Model

We now generalize the minimum subsidy result in Section 3.2 to the context of uncertain utilities and restricted actions. First, we note that the overpayment result does not hold if we allow the proper scoring rule to not be strict: as in the example shown in Figure 5, sometimes we can treat all reports for which we have relative preferences²¹ to be the same, and implement a market scoring rule on the collapsed classes. In this case, the agents can obtain optimal payments while ignoring the directions for which the principal has preferences.




Fig. 5. Suppose that there are 3 outcomes and no restrictions on actions. T is such that the only bad direction is \mathbf{b} . Consider a market scoring rule on the relative probability of outcomes 1 and 2, guaranteeing incentive c . For every probability distribution $\mathbf{r} \in \mathbb{P}$ we follow the contour lines in the diagram and treat this as some \mathbf{r}' on the segment connecting \mathbf{O}_1 and \mathbf{O}_2 . This makes the cost function G (for the corresponding proper scoring rule) in every situation constant along the contour lines. The resultant mechanism is principal-aligned, guarantees incentive c , but requires only $O(c)$ subsidy.

In practice it seems desirable to always strictly incentivize the agents to report as close to the true probability as possible. One way to formalize this is the notion of *uniform incentivization*.

Definition 15 Suppose that $f : [0, \infty) \rightarrow \mathbb{R}$ is convex,²² with $f(0) = 0$, $f(x) > 0 \forall x > 0$. A proper scoring rule provides uniform incentives according to f if $\forall \mathbf{r}, \mathbf{p} \in \mathbb{P}$,

$$\tilde{S}(\mathbf{p}, \mathbf{p}) - \tilde{S}(\mathbf{r}, \mathbf{p}) \geq f(\|\mathbf{p} - \mathbf{r}\|)$$

A one-round prediction mechanism guarantees uniform incentives according to f if for every agent j and every combination of others' reports, the corresponding proper scoring rule provides uniform incentives according to f .

Intuitively, f can be thought of as a measure of incentivization locally at $\mathbf{p} \in \mathbb{P}$, and providing uniform incentives according to f guarantees a certain level of incentivization at all $\mathbf{p} \in \mathbb{P}$. For a one-round prediction mechanism, guaranteeing uniform incentives according to some f corresponds to maintaining a minimal standard of incentivization for all agents in all situations. Conventional scoring rules such as the quadratic, the logarithmic, and the spherical scoring rules all guarantee uniform incentives according to some f .

Theorem 6. Suppose that a feasible one-round prediction mechanism with n agents

²¹ By “relative preference” between reports \mathbf{r}' and \mathbf{r} , we mean that either $\mathbf{r}' - \mathbf{r} \in B(\mathbf{r})$ or $\mathbf{r} - \mathbf{r}' \in B(\mathbf{r}')$.

²² The requirement that f is convex is natural since cost function G is convex.

- is principal aligned with set T of possible utility vectors, under action model A (assume that the magnitude of the worst action $d > 0$.²³);
- guarantees uniform incentives according to some f .

Then the mechanism requires subsidy $n \cdot f(\frac{d}{2})$ (which is $\Theta(n)$).

Proof. It is straightforward to show that the requirements that f is convex, $f(0) = 0$, and $f(x) > 0 \forall x > 0$, imply that f is continuous and non-decreasing.

For a given j and a given combination of others' reports, let the corresponding proper scoring rule be S with cost function G . By the definition of d (Definition 13), $\forall \epsilon > 0, \epsilon < d, \exists \mathbf{x}, \mathbf{y} \in \mathbb{P}$, such that $\mathbf{y} - \mathbf{x} \in B(\mathbf{x})$ and $\|\mathbf{y} - \mathbf{x}\| = d - \epsilon$. Let $\mathbf{z} = \frac{\mathbf{x} + \mathbf{y}}{2}$. By uniform incentivization according to f ,

$$\begin{aligned}\tilde{S}(\mathbf{x}, \mathbf{x}) - \tilde{S}(\mathbf{z}, \mathbf{x}) &\geq f\left(\frac{d}{2} - \frac{\epsilon}{2}\right) \\ \tilde{S}(\mathbf{y}, \mathbf{y}) - \tilde{S}(\mathbf{z}, \mathbf{y}) &\geq f\left(\frac{d}{2} - \frac{\epsilon}{2}\right)\end{aligned}$$

Let G^* be the subgradient corresponding to the scoring rule S (in the sense of Theorem 1); we have

$$\begin{aligned}\tilde{S}(\mathbf{z}, \mathbf{x}) - G(\mathbf{z}) &= G^*(\mathbf{z}) \cdot (\mathbf{x} - \mathbf{z}) \\ \tilde{S}(\mathbf{z}, \mathbf{y}) - G(\mathbf{z}) &= G^*(\mathbf{z}) \cdot (\mathbf{y} - \mathbf{z})\end{aligned}$$

Adding the four expressions and dividing by 2, we get

$$\frac{1}{2}(G(\mathbf{x}) + G(\mathbf{y})) = \frac{1}{2}(\tilde{S}(\mathbf{x}, \mathbf{x}) + \tilde{S}(\mathbf{y}, \mathbf{y})) \geq G(\mathbf{z}) + f\left(\frac{d}{2} - \frac{\epsilon}{2}\right)$$

Now, by voluntary participation, $G(\mathbf{z}) \geq 0$. Since $\mathbf{y} - \mathbf{x} \in B(\mathbf{x})$, $G(\mathbf{x}) \geq G(\mathbf{y})$. Hence, $G(\mathbf{x}) \geq f(\frac{d}{2} - \frac{\epsilon}{2})$. If all agents report \mathbf{x} and the true probability vector is \mathbf{x} , then the principal is forced to pay $nf(\frac{d}{2} - \frac{\epsilon}{2})$. Taking the limit as $\epsilon \rightarrow 0$ and using monotonicity and continuity of f , we get the desired result.

In other words, even when the principal's utilities are uncertain and when agents' actions might be limited, suppose that some surely undesirable action exists. Then, under the requirements of feasibility and uniform incentivization, *the cheapest principal-aligned mechanism requires $\Theta(n)$ times as much subsidy as the cheapest non-principal-aligned mechanism.*

One practical implication of our $\Theta(n)$ subsidy results is that to run a useful prediction market without incentivizing undesirable actions, it may be impractical to let agents join freely. This is because agents may join just to get the subsidy, without providing any useful additional information. However, it may be practical for an organization to run a principal-aligned in-house prediction market, because in this context the number of agents is naturally limited.

Suppose that sufficient uncertainties in utilities exist, Theorem 5 implies that organizations with in-house prediction markets can take away bad incentives by providing constant bonuses to align agents' incentives with those of the principal.

References

1. Joyce Berg, Robert Forsythe, Forrest Nelson, and Thomas Rietz. *Results from a Dozen Years of Election Futures Markets Research*, volume 1 of *Handbook of Experimental Economics Results*, chapter 80, pages 742–751. Elsevier, 2008.
2. Yiling Chen, Daniel M. Reeves, David M. Pennock, Robin D. Hanson, Lance Fortnow, and Rica Gonen. Bluffing and strategic reticence in prediction markets. In *WINE '07: Workshop on Internet and Network Economics*, 2007.

²³ *i.e.* it is possible for agents to perform a certainly undesirable action.

3. Vincent Conitzer. Prediction markets, mechanism design, and cooperative game theory. In *UAI '09: Conference on Uncertainty in Artificial Intelligence*, 2009.
4. Stanko Dimitrov and Rahul Sami. Non-myopic strategies in prediction markets. In *EC '08: Proceedings of the 9th ACM conference on Electronic Commerce*, pages 200–209, New York, NY, USA, 2008. ACM.
5. Tilmann Gneiting and Adrian E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, March 2007.
6. Robin Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003.
7. Robin Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):1–15, 2007.
8. Nicolas S. Lambert, John Langford, Jennifer Wortman, Yiling Chen, Daniel Reeves, Yoav Shoham, and David M. Pennock. Self-financed wagering mechanisms for forecasting. In *EC '08: Proceedings of the 9th ACM conference on Electronic Commerce*, pages 170–179, New York, NY, USA, 2008. ACM.
9. Robert Looney. DARPA’s policy analysis market for intelligence: Outside the box or off the wall. *Strategic Insights*, 2(9), Sep 2003.
10. David M. Pennock. A dynamic pari-mutuel market for hedging, wagering, and information aggregation. In *EC '04: Proceedings of the 5th ACM conference on Electronic Commerce*, pages 170–179, New York, NY, USA, 2004. ACM.
11. David M. Pennock and Rahul Sami. Computational aspects of prediction markets. In *Algorithmic Game Theory*, chapter 26, pages 651–674. Cambridge University Press, 2007.
12. Charles R. Plott and Kay-Yut Chen. Information aggregation mechanisms: Concept, design and implementation for a sales forecasting problem. Working Paper 1131, California Institute of Technology, Division of the Humanities and Social Sciences, March 2002.
13. Leonard J. Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, 66(336):783–801, 1971.
14. J. Wolfers and E. Zitzewitz. Prediction markets. *The Journal of Economic Perspectives*, 18(2):107–126, 2004.
15. Justin Wolfers and Eric Zitzewitz. Prediction markets in theory and practice. Research Papers 1927, Stanford University, Graduate School of Business, February 2006.

A Omitted Proofs

Proof. [of Theorem 1] Suppose that $H_{\mathbf{r}}$ is a subtangent hyperplane (recall Definition 3) and setting $S(\mathbf{r}, i) = H_{\mathbf{r}}(\mathbf{O}_i)$. We have $\forall \mathbf{p} \in \mathbb{P}$,

$$\tilde{S}(\mathbf{r}, \mathbf{p}) = \sum_{i=1}^m H_{\mathbf{r}}(\mathbf{O}_i)p_i = H_{\mathbf{r}}(\mathbf{p})$$

(Since $H_{\mathbf{r}}$ is a linear function constrained a $m - 1$ dimensional space, it is determined by a weighted average of the m intercepts we specify.) Therefore, because $H_{\mathbf{r}}$ lies below $G(\mathbf{p})$,

$$\tilde{S}(\mathbf{r}, \mathbf{p}) = H_{\mathbf{r}}(\mathbf{p}) \leq G(\mathbf{p}) = H_{\mathbf{p}}(\mathbf{p}) = \tilde{S}(\mathbf{p}, \mathbf{p})$$

Hence, S is a proper scoring rule. If $H_{\mathbf{r}}$ is strictly subtangent, then the above inequality is strict $\forall \mathbf{p} \neq \mathbf{r}$ and S is strictly proper.

Conversely, suppose S is a proper scoring rule, define $G(\mathbf{p}) = \tilde{S}(\mathbf{p}, \mathbf{p})$. Note that $G(\mathbf{p}) = \sup_{\mathbf{r} \in \mathbb{P}} \tilde{S}(\mathbf{r}, \mathbf{p})$. For fixed \mathbf{r} , $\tilde{S}(\mathbf{r}, \mathbf{p}) = \sum_i S(\mathbf{r}, i)p_i$ is a linear and therefore convex function of $\mathbf{p} \in \mathbb{P}$. Because the supremum of a class of convex functions is convex, G is convex. Now define $H_{\mathbf{r}}(\mathbf{p}) = \tilde{S}(\mathbf{r}, \mathbf{p})$. $H_{\mathbf{r}}$ is a linear function in \mathbf{p} , and $H_{\mathbf{r}}(\mathbf{r}) = G(\mathbf{r})$. Moreover, $\forall \mathbf{p} \in \mathbb{P}$,

$$H_{\mathbf{r}}(\mathbf{p}) = \tilde{S}(\mathbf{r}, \mathbf{p}) \leq \tilde{S}(\mathbf{p}, \mathbf{p}) = G(\mathbf{p})$$

Hence, $H_{\mathbf{r}}$ is a subtangent hyperplane to G . (It is strictly subtangent if S is strict.)

B Relaxing Uniform Incentivization When Uncertainties In Utility are Independent

Given additional structure the uniform incentivization requirement in Theorem 6 can be relaxed. Here we prove a stronger result when uncertainties in the principal’s utilities are independent and there are no restrictions in how agents can change the underlying probabilities.

Definition 16 *Uncertainties in the principal’s utilities are independent if the utility u_i for outcome i is in some range $u_i \in [l_i, h_i]$, and all combinations across outcomes are possible. In other words, the set of possible utility vectors is the Cartesian product*

$$T = \times_{i \in \Omega} [l_i, h_i]$$

Our weaker notion of incentivization is *providing incentive for outcome i* , which stipulates that the agent cannot simply report a dummy value for the probability of event i .

Definition 17 *A proper scoring rule provides incentive c for outcome $i \in \Omega$ if fixing the relative probabilities of other outcomes, the difference, between the greatest expected payment and the greatest the agent can guarantee by reporting some constant p_i^* for i , is at least c . (Here the constant p_i^* can depend on the relative probabilities of other outcomes.) A one-round prediction mechanism guarantees incentive c for outcome i if for every agent j and every combination of others’ reports, the corresponding proper scoring rule provides incentive c for outcome i .*

Since the action model is trivial ($A(\mathbf{p}) \equiv \mathbb{P}$), the bad direction function (Definition 13) is $B(\mathbf{p}) \equiv B = \bigcap_{\mathbf{u} \in T} \{\mathbf{x} | \mathbf{x} \cdot \mathbf{u} < 0, \mathbf{x} \cdot \mathbf{1} = 0\}$. We call B the *bad direction set*.

Lemma 4. *Suppose the bad direction set B is non-empty, then there exists outcomes i, j s.t. $l_i > h_j$. i.e., we are certain that outcome i is better than j for the principal. We call such outcomes i or j distinguishable.*

Note that if outcome i is not distinguishable, then under some setting of $\mathbf{u} \in T$, it is the worst outcome, and under some other setting, it is the best: i.e., we have no idea whatsoever whether i is good or bad. Suppose that we only provide incentives for outcomes that are not distinguishable, then we can run a market scoring rule on those outcomes and implement this with constant subsidy. However, by ignoring distinguishable outcomes, we have discarded the “principal-alignment” facet of the problem, as the principal no longer has clear preferences on any outcome.

The stronger overpayment theorem is as follows.

Theorem 7. *Suppose that a feasible one-round prediction mechanism for n agents*

- *is aligned to the set T of possible principal’s utilities, in which uncertainties are independent (assume that the bad-direction set B is non-empty);*
- *guarantees incentive c for some distinguishable outcome i .*

Then the mechanism requires subsidy cn .