

Decision Scoring Rules

Caspar Oesterheld^[0000-0003-4222-7855] and
Vincent Conitzer^[0000-0003-1899-7884]

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

Abstract. We consider a setting in which a principal faces a decision and asks an external expert for a recommendation as well as a probabilistic prediction about what outcomes might occur if the recommendation were implemented. The principal then follows the recommendation and observes an outcome. Finally, the principal pays the expert based on the prediction and the outcome, according to some decision scoring rule. In this paper, we ask the question: What does the class of proper decision scoring rules look like, i.e., what scoring rules incentivize the expert to honestly reveal both the action he believes to be best for the principal and the prediction for that action? We first show that in addition to an honest recommendation, proper decision scoring rules can only incentivize the expert to reveal the expected utility of taking the recommended action. The principal cannot strictly incentivize honest reports on other aspects of the conditional distribution over outcomes without setting poor incentives on the recommendation itself. We then characterize proper decision scoring rules as ones which give or sell the expert shares in the principal's project. Each share pays, e.g., \$1 per unit of utility obtained by the principal. Owning these shares makes the expert want to maximize the principal's utility by giving the best-possible recommendation. Furthermore, if shares are offered at a continuum of prices, this makes the expert reveal the value of a share and therefore the expected utility of the principal conditional on following the recommendation.

1 Introduction

Consider a firm that is about to make a major strategic decision. It wishes to maximize the expected value of the firm. It hires an expert to consult on the decision. The expert is strictly better informed than the firm, but it is commonly understood that the outcome conditional on the chosen course of action is uncertain even for the expert. The firm can commit to a compensation package for the expert; compensation can be conditional both on the expert's predictions and on what happens (e.g., in terms of the value of the firm) after a decision is made. (The compensation cannot depend on what would have happened if another action had been chosen.) The firm cannot or does not want to commit to an arbitrary mapping from expert reports to actions: once the report is made, the firm will always choose the action that maximizes expected value, conditional on that report. What compensation schemes will incentivize the expert

to report truthfully? One straightforward solution is to give the expert a fixed share of the firm at the outset. Are there other schemes that also reward accurate predictions?

Our approach to formalizing and answering these questions is inspired by existing work on eliciting honest predictions about an event that the firm or *principal* cannot influence. In the single-expert case, such elicitation mechanisms are known as *proper scoring rules* [Brier, 1950; Good, 1952, Section 8; McCarthy, 1956; Savage, 1971; Gneiting and Raftery, 2007]. Formally, a scoring rule for prediction s takes as input a probability distribution \hat{P} reported by the expert, as well as the actual outcome ω , and assigns a *score* or *reward* $s(\hat{P}, \omega)$. A scoring rule s is proper if the expert maximizes his¹ expected score by reporting as \hat{P} his true beliefs about how likely different outcomes are. The class of proper scoring rules has been completely characterized in prior work [e.g., Gneiting and Raftery, 2007, Section 2]. This characterization also provides a foundation for the design of proper scoring rules that are *optimal* with respect to a specific objective (and potentially under additional constraints) [Osband, 1989; Neyman et al., 2020; D. Hartline et al., 2020], and as such can be viewed as analogous to characterizations of incentive-compatible choice functions in mechanism design, including characterizations based on cycle monotonicity [first derived by Rochet, 1987; and utilized in an optimization context by, e.g., Lavi and Swamy, 2007] or weak monotonicity [first derived by Bikhchandani et al., 2006; and utilized by, e.g., Ashlagi et al., 2012].

In this paper, we derive a similar characterization of what we call *proper decision scoring rules* – scoring rules that incentivize the expert to honestly report the best available action and an honest prediction about the outcome given that action. We show that proper decision scoring rules cannot give the expert *strict* incentives to report any properties of the outcome distribution under the recommended action, other than its expected utility (Section 3). For example, there is no proper decision scoring rule that strictly incentivizes the expert to honestly reveal the variance of the distribution for the recommended action. Intuitively, rewarding the expert for getting anything other than the mean of the distribution right will strictly incentivize him to recommend actions whose outcome is easy to predict as opposed to actions with high expected utility for the principal. Hence, the expert’s reward can depend only on the reported expected utility for the recommended action, and the realized utility. We then give a complete characterization of proper decision scoring rules (Section 4). In the case of a company maximizing its value, the mechanisms can be interpreted as offering the expert to buy, at varying prices, shares in the company (Section 5.1). The price schedule does not depend on the action chosen. Thus, given the chosen action, the expert is incentivized to buy shares up to the point where the price of a share exceeds the expected value of the share, thereby revealing the principal’s expected utility. Moreover, once the expert has some positive share in

¹ We use “he” for experts and “she” for the principal, i.e., the firm or person setting up the mechanism.

the principal’s utility, he will be (strictly) incentivized to recommend an optimal action.

2 Setting

We consider a setting in which a *principal* faces a choice from some finite set of at least two *actions* A . After taking any such action, the principal observes an outcome from some finite set Ω , to which she assigns a utility according to some utility function $u: \Omega \rightarrow \mathbb{R}$.

In addition to the principal, there is an expert who holds beliefs described by some vector of conditional probability distributions $P \in \Delta(\Omega)^A$, which specifies for each action $a \in A$ and outcome $\omega \in \Omega$ the probability $P(\omega \mid a)$ that outcome ω will be obtained if action a is taken by the principal.²

The principal may ask the expert to recommend some action \hat{a} – with the intention of getting the expert to report one that maximizes $\mathbb{E}_P[u(O) \mid \hat{a}]$, where O is the outcome distributed according to $P(\cdot \mid \hat{a})$ – and to also make a prediction $\hat{P}_{\hat{a}} \in \Delta(\Omega)$ about what outcome action \hat{a} will give rise to. The principal then always follows that recommendation. The principal could also ask the expert to report on what would happen if she took suboptimal actions $a \neq \hat{a}$. However, it makes little sense for the principal to use these reports for rewarding the expert. After all, these other predictions are never tested. Hence, if she gave him different (expected) rewards depending on whether he reports \hat{P}_a or \hat{P}'_a , he would prefer reporting one of them over the other regardless of which represents his beliefs. If we reward based on reports about unrecommended actions $(P(\cdot \mid a))_{a \neq \hat{a}}$, we would therefore strictly incentivize misreporting. For a formal statement and proof of this point, see Othman and Sandholm (2010, Theorems 1 and 4) or Chen, Kash, et al. (2014, Theorem 4.1). Since we are concerned with properly incentivizing the expert, we will therefore only consider the report $\hat{P}_{\hat{a}} \in \Delta(\Omega)$ about the recommended action \hat{a} .

Others have considered principals who take suboptimal actions with some (small) probability [Chen, Kash, et al., 2014; cf. Zermelo, 2011; Zermelo, 2012]. This can help incentivize the expert to report honestly. For instance, Chen, Kash, et al. (2014) show that taking an action a with a small positive probability is enough to make the expert honestly report his belief $P(\cdot \mid a)$. The principal may

² Throughout this paper, we assume that the expert could hold *any* beliefs $P \in \Delta(\Omega)^A$. An alternative setting would be one in which the expert has exclusive access to some private piece of information e from some set H . Each such piece of evidence gives rise to a posterior $P(\cdot \mid \cdot, e) \in \Delta(\Omega)^A$ and the principal knows both the possible pieces of evidence and how they map to posteriors. In some cases, this gives rise to ways of scoring that are not available here. For instance, the outcome ω might inform the principal directly about which e the expert observed [cf. Boutilier, 2012, Section 3.1; Carroll, 2019]. Extending our characterizations to such settings appears nontrivial. In contrast, in our setup, the principal is not required to know the evidence structure, making the scoring rules under consideration more generally applicable.

thus learn from the expert about the effects of all actions, not just the one she ends up taking. However, randomizing also has a number of disadvantages. First and obviously, the principal prefers taking the best action all of the time over taking the best action *almost* all of the time. Second, the principal, who will finally make the decision, may be unable to credibly commit to taking, say, the worst available action if the dice demand it. The expert may not trust that the principal will really go through with a promise (or, rather, threat) of choosing a suboptimal action some of the time (cf. Chen, Kash, et al., 2014, Section 5). Third, if the probability of suboptimal actions being chosen is small, then rewards or punishments based on the outcomes of these events must be scaled up in inverse proportion to that probability to generate proper incentives (Chen, Kash, et al., 2014, Theorem 4.2). While in theory this poses no problem, in practice there are limits to rewards and punishments, due to budgets, limited liability, and other constraints. Operating within these constraints will thus require suboptimal actions to be chosen significantly more often. Fourth, taking each action with positive probability is only an option when the set of available actions is at most countable. Fifth and finally, in the real world an action may produce different effects depending on the probability with which it is taken. For instance, if the decision mechanism is transparent, actions may be less effective when they are chosen with low probability as a result of randomization over suboptimal actions, because they will then catch employees by surprise.

In our setting, the principal has no way of directly verifying the information she receives. We also assume that the expert has no intrinsic interests in the principal’s endeavors.³ To nevertheless incentivize the expert to report his private information honestly, the principal may therefore use a *decision scoring rule* (DSR) $s: \Delta(\Omega) \times \Omega \rightarrow \mathbb{R}$, which maps a report and an outcome observed after taking \hat{a} onto a reward for the expert. This reward could be financial, but it could also be given in some social currency, e.g., a number of points listed on some website. As we have noted earlier, we do not let the score depend on predictions about what would happen if an action other than \hat{a} were to be taken. Furthermore, we do not let the score depend on what action is recommended – other than through the outcome obtained after implementing \hat{a} . It is easy to see why that would set poor incentives for some beliefs. We do not give a formal proof here to avoid the introduction of alternative formalisms. However, such a proof could easily be conducted as part of the proof of Lemma 1. Incidentally, because the DSR does not take as input the recommended action, affecting the principal’s outcome by making a good recommendation is indistinguishable from affecting the principal’s outcome in other ways. Our work, therefore, generalizes beyond this pure recommendation setting. It also means that any proper DSR (as defined and characterized in this paper) not only incentivizes the expert to

³ Of course, if the principal knows what intrinsic preferences the expert has over outcomes, the principal can exactly compensate these preferences with payments. Hence, our characterization would still characterize the *net* scoring rules that the principal can induce. If the principal does not know the expert’s preferences, however, propriety of the net scoring rule can in general not be ensured.

make good recommendations a^* but also to take unobservable actions in the principal’s favor whenever he can⁴ (cf. the notion of “principal alignment” in prediction markets, as discussed by Shi et al., 2009).

Ideally, the principal sets up s such that the expert is incentivized to recommend an action \hat{a} from $\text{opt}(P) = \arg \max_{a \in A} \mathbb{E}_{O \sim P} [u(O) \mid a]$, the set of optimal actions, and further to report $\hat{P}_{\hat{a}} = P(\cdot \mid \hat{a})$ honestly. The most basic form of this requirement is (non-strict) propriety: among the reports that give the expert the highest expected score should always be one that consists of an optimal action and an honest prediction. Formally, we can define this as follows.

Definition 1. *We say that a DSR s is proper if for all beliefs $P(\cdot \mid \cdot) \in \Delta(\Omega)^A$ and all possible recommendations $\hat{a} \in A$ and predictions $\hat{P}_{\hat{a}} \in \Delta(\Omega)$ we have*

$$\mathbb{E}_{O \sim P} [s(\hat{P}_{\hat{a}}, O) \mid \hat{a}] \leq \mathbb{E}_{O \sim P} [s(P(\cdot \mid a^*), O) \mid a^*] \quad (1)$$

for some $a^* \in \text{opt}(P)$.

We limit our attention to designing proper DSRs. However, while this propriety implies that the expert has no bad incentives, it does not require that the expert has any good incentives. For example, any constant s is (non-strictly) proper. We might therefore be interested in the structure of *strictly* proper DSRs, i.e., ones where inequality 1 is strict unless \hat{a} is optimal and $\hat{P}_{\hat{a}} = P(\cdot \mid \hat{a})$ is reported honestly. As we will see (Lemma 2), no DSR is strictly proper in this sense. We will therefore define partially strict versions of propriety.

Definition 2. *We say that s is right-action proper if it is proper and for all beliefs $P(\cdot \mid \cdot) \in \Delta(\Omega)^A$ and all possible recommendations $\hat{a} \in A$ and predictions $\hat{P}_{\hat{a}} \in \Delta(\Omega)$,*

$$\mathbb{E}_P [s(\hat{P}_{\hat{a}}, O) \mid \hat{a}] = \max_{a^* \in \text{opt}(P)} \mathbb{E}_P [s(P(\cdot \mid a^*), O) \mid a^*] \quad (2)$$

implies $\hat{a} \in \text{opt}(P)$. We call s strictly proper w.r.t. the mean if eq. 2 implies

$$\mathbb{E}_{O \sim P(\cdot \mid \hat{a})} [u(O)] = \mathbb{E}_{O \sim \hat{P}_{\hat{a}}} [u(O)]. \quad (3)$$

Right-action propriety should be a main goal. Fortunately, such scoring rules do indeed exist. One class of such rules is especially easy to identify. For any $c_1, c_2 \in \mathbb{R}$ with $c_1 > 0$, we can use

$$s(\hat{P}_{\hat{a}}, \omega) = c_1 u(\omega) + c_2. \quad (4)$$

⁴ From our definition of propriety it will immediately follow that this is the case when the expert knows which interventions he will make at the time of submitting his prediction. Much more surprisingly, our characterization will show that he will be incentivized to take action in the principal’s favor, even if that renders his earlier prediction inaccurate. In other words, if an expert makes a recommendation and prediction today, and then on the next day is to his surprise and unobservable to the principal given an opportunity to increase the principal’s expected utility beyond what he predicted earlier, he will gladly seize that opportunity.

If we imagine that the principal is some company whose utility is the company’s overall value, then this corresponds to giving the expert some share in the company [Chen, Kash, et al., 2014, Section 5; cf. Johnstone et al., 2011], which is of course a common approach in principal-agent problems. It is much less obvious what the entire classes of proper and right-action proper DSRs look like, and on what other aspects of the report \hat{P}_a they can set strict incentives. As we will see, the only further form of strictness that proper DSRs can achieve is strict propriety w.r.t. the mean. This is why we only define these two forms of partially strict propriety here.

While right-action propriety appears most important, setting incentives on the predictions for what will happen after taking the recommended action can be useful for a variety of reasons. For example, consider again a principal owning a firm. By eliciting predictions, she may hope to inform auxiliary decisions. For instance, the principal may wish to know the expected value of the firm to decide at which prices she would be willing to sell some shares, whether to buy a luxury apartment with a view of Central Park or a modest flat in Brooklyn, etc. Similarly, if the recommended project is risky (if the variance of the utility is high according to the reported probability distribution), the principal may wish to hedge against the uncertainty and hold off on other major decisions that require financial security (acquiring another company, buying said apartment, starting a family, etc.).⁵ Another reason to reward accurate predictions can be motivated by an alternative interpretation of scoring rules themselves. Instead of using scoring rules to set incentives on an expert’s future recommendations and predictions, we could also use them [in line with the name and the original intention of, e.g., Brier, 1950] to evaluate experts based on their past record. While making good recommendations is paramount, we would all else equal regard an expert as more competent (and more likely to be helpful in the future) if he can make accurate predictions about what outcomes his recommendations give rise to.

If we are willing to drop the demand of getting honest recommendations, then the characterization of scoring rules for prediction (see, e.g., Gneiting and Raftery, 2007, Section 2) tells us which DSRs are strictly proper w.r.t. the reported probability distribution (see Chen, Kash, et al., 2014, Sections 3-4). It is informative to work through an example of such a scoring rule for prediction and why it strictly incentivizes giving a suboptimal recommendation. Consider the quadratic scoring rule (originally proposed by and sometimes named after Brier, 1950): $s(\hat{P}, \omega) = 2\hat{P}(\omega) - \sum_{\omega' \in \Omega} \hat{P}(\omega')^2$. In a context in which no action needs to be selected and an expert must report only a probability distribution, it is well known that the expert is best off reporting the distribution truthfully.

⁵ In principle, the principal could, of course, let the expert make auxiliary decisions on her behalf. However, this may often be impractical. For example, some of the auxiliary decisions may require knowledge (say, about the NYC housing market or about how the principal would like different apartments) that the given expert or any external expert lacks. Further, many auxiliary decisions may not lead to an observable, quantifiable type of utility. This is especially true for private decisions like buying an apartment.

Hence, even if we allow the expert to choose the principal’s action and thereby the random variable he is being scored on, it will elicit honest predictions for that random variable. However, s is not a proper *decision* scoring rule. It generally incentivizes the expert to recommend an action a that makes the outcome easiest to predict. For instance, suppose that $\Omega = \{\omega_1, \dots, \omega_m\}$, that the optimal action a^* leads to the uniform distribution $O_{a^*} = \frac{1}{m} * \omega_1 + \dots + \frac{1}{m} * \omega_m$, while a' leads to $O_{a'} = 1 * \omega_1$ deterministically. Then the expert will (assuming $m > 1$) always prefer recommending the suboptimal a' , since

$$\mathbb{E}[s(P_{a'}, O_{a'})] = 1 > \frac{1}{m} = \frac{2}{m} - \sum_{\omega' \in \Omega} \frac{1}{m^2} = \mathbb{E}[s(P_{a^*}, O_{a^*})]. \quad (5)$$

3 Only means matter

We have argued [and, as noted, it has been formally proven in earlier work: Othman and Sandholm, 2010, Theorems 1 and 4; Chen, Kash, et al., 2014, Theorem 4.1] that proper DSRs cannot strictly incentivize the expert to honestly report on what would happen if a non-recommended action were taken. Next, we prove that if a DSR is to be proper, it can only strictly incentivize the expert to be honest about the optimal (recommended) action and the expected utility of that action. That is, no proper DSRs are strictly proper w.r.t. any other aspect of the report. For example, proper DSRs cannot strictly incentivize the expert to honestly reveal the variance of the utility given that the recommended action is taken (aside from the special case of $|\Omega| = 2$, in which the mean reveals the entire distribution).

To prove that result, we need a simple lemma. From the definition of propriety it follows that if one action is better than another, the expert must weakly prefer recommending the better action, even if, say, the worse action makes the outcome more predictable. But perhaps if two actions’ expected utilities are the same, a proper DSR could induce the expert to strictly prefer recommending one of the two (say, the one with an easier-to-predict distribution over outcomes)? It turns out that this is not the case. That is, we show that under honest predicting the expected scores for two different recommendations are the same whenever the expected utilities of the two recommendations are the same.

Lemma 1. *Let s be a proper DSR and $P_a, P_{a'} \in \Delta(\Omega)$. Then, if*

$$\min_{\omega \in \Omega} u(\omega) < \mathbb{E}_{O \sim P_a} [u(O)] = \mathbb{E}_{O \sim P_{a'}} [u(O)] < \max_{\omega \in \Omega} u(\omega) \quad (6)$$

it must be the case that $\mathbb{E}_{O \sim P_a} [s(P_a, O)] = \mathbb{E}_{O \sim P_{a'}} [s(P_{a'}, O)]$.

Proof. Let $\omega_L = \arg \min_{\omega \in \Omega} u(\omega)$ and $\omega_H = \arg \max_{\omega \in \Omega} u(\omega)$ (with ties broken arbitrarily). Let $L = u(\omega_L)$ and $H = u(\omega_H)$. Then, for $p \in (0, 1)$, let $R_p = p * \omega_H + (1 - p) * \omega_L$ and Q_p be the distribution of that random variable. Note first that because s is proper, $\mathbb{E}[s(Q_p, R_p)]$ is non-decreasing in p . We claim that

$\mathbb{E}[s(Q_p, R_p)]$ is also continuous in p . From this, we will directly derive the claim of the lemma.

Such continuity properties have often been proven in the literature. For example, Frongillo and Kash (2014, p. 1f.) note – translated to our setting – that for proper s , $\mathbb{E}[s(Q_p, R_p)] = \max_{p' \in [0,1]} \mathbb{E}[s(Q_{p'}, R_p)]$. So $\mathbb{E}[s(Q_p, R_p)]$ as a function of p is the pointwise maximum of a set of functions in p . These individual functions are (by the definition of expectation) affine in p . It can be shown that the pointwise maximum of a set of affine functions is convex. Finally, a convex function defined on an open interval is continuous on that interval. For completeness, we also give a more elementary (longer) real analysis-style proof of continuity in Appendix A.

We now use continuity to show that $\mathbb{E}_{O \sim P_a}[s(P_a, O)]$ must be the same for all P_a with the same non-degenerate mean μ (i.e., $L < \mu < H$). For s to be proper, it must be the case that $\mathbb{E}_{O \sim P_a}[s(P_a, O)] \geq \mathbb{E}[s(Q_p, R_p)]$ for all p with $\mu > pH + (1-p)L$ and $\mathbb{E}_{O \sim P_a}[s(P_a, O)] \leq \mathbb{E}[s(Q_p, R_p)]$ for all p with $\mu < pH + (1-p)L$. But because $\mathbb{E}[s(Q_p, R_p)]$ is continuous, this implies $\mathbb{E}_{O \sim P_a}[s(P_a, O)] = \mathbb{E}[s(Q_p, R_p)]$ for the p s.t. $\mu = pH + (1-p)L$. So for any P_a with mean μ , this uniquely fixes the expected score under honestly reporting P to the same $\mathbb{E}[s(Q_p, R_p)]$. \square

It is worth noting that the proof is based on the lack of “space” in the set \mathbb{R} of possible scores. We could imagine experts who maximize a lexicographic score. Then our result only shows that the lexically highest value of the scores – under honest reporting – of two equally good recommendations must be the same. But the lexically lower values could be given according to some scoring rule for prediction (such as the quadratic scoring rule) and thus make the expert prefer one of two recommendations with equal expected utility for the expert.

We have now shown that the expected utility of an action uniquely determines the expected reward the expert gets for recommending that action and honestly predicting the outcome given that action. Next, we show – perhaps more surprisingly – that in a sense, the mean is the only piece of information the principal can elicit from the expert. That is, as long as the expert honestly reports the expected utility of the recommendation, he can almost arbitrarily mis-predict the outcome to a proper DSR without affecting his expected score.

Lemma 2. *Let s be a proper DSR and $P_a, \hat{P}_a \in \Delta(\Omega)$. Then if*

$$\min_{\omega \in \Omega} u(\omega) < \mathbb{E}_{P_a}[u(O)] = \mathbb{E}_{\hat{P}_a}[u(O)] < \max_{\omega \in \Omega} u(\omega) \quad (7)$$

and $\text{supp}(P_a) \subseteq \text{supp}(\hat{P}_a)$, it must be the case that

$$\mathbb{E}_{P_a}[s(P_a, O)] = \mathbb{E}_{P_a}[s(\hat{P}_a, O)].$$

Proof. If $\mathbb{E}_{P_a} [u(O)] = \mu = \mathbb{E}_{\hat{P}_a} [u(O)]$ and $\text{supp}(P_a) \subseteq \text{supp}(\hat{P}_a)$, there is a P'_a and a $p \in (0, 1]$ s.t. $\hat{P}_a = pP_a + (1-p)P'_a$ and $\mathbb{E}_{P'_a} [u(O)] = \mu$. Then

$$\mathbb{E}_{\hat{P}_a} [s(\hat{P}_a, O)] = p\mathbb{E}_{P_a} [s(\hat{P}_a, O)] + (1-p)\mathbb{E}_{P'_a} [s(\hat{P}_a, O)] \quad (8)$$

$$\stackrel{s \text{ is proper}}{\leq} p\mathbb{E}_{P_a} [s(\hat{P}_a, O)] + (1-p)\mathbb{E}_{P'_a} [s(P'_a, O)] \quad (9)$$

$$\stackrel{s \text{ is proper}}{\leq} p\mathbb{E}_{P_a} [s(P_a, O)] + (1-p)\mathbb{E}_{P'_a} [s(P'_a, O)] \quad (10)$$

$$\stackrel{\text{Lemma 1}}{=} \mathbb{E}_{\hat{P}_a} [s(\hat{P}_a, O)]. \quad (11)$$

Because the expression at the beginning is the same as the expression in the end, the \leq -inequalities in the middle must be equalities. Therefore, because $p > 0$, it must be the case that $\mathbb{E}_{P_a} [s(P_a, O)] = \mathbb{E}_{P_a} [s(\hat{P}_a, O)]$. \square

Lemma 2 implies that proper DSRs cannot be strictly proper w.r.t. anything but the mean (and the recommended action). Thus, we will henceforth only consider DSRs $s(\hat{\mu}, \omega)$, which take only the reported mean as input. Note that not all proper scoring rules can be expressed as a scoring rule that depends only on the mean. For one, we could punish the expert if the support of the reported probability distribution does not contain the observed outcome. Of course, unless $\min/\max_{\omega \in \Omega} u(\omega)$ will occur with certainty, the expert has no reason not to report full support. Furthermore, we could let the submitted probability distribution determine the scoring rule in ways that do not affect the *expected* score.⁶ Since none of these dependencies on details of the submitted probability distribution seem helpful, we will ignore them.

Next we argue that in a proper DSR, $s(\hat{\mu}, \omega)$ can only depend on $u(\omega)$ (and $\hat{\mu}$, of course), i.e., on the utility of the obtained outcome rather than the outcome itself.

Lemma 3. *Let $s: \mathbb{R} \times \Omega \rightarrow \mathbb{R}$ be a proper DSR; $\omega_1, \omega_2 \in \Omega$ be two outcomes with $u(\omega_1) = u(\omega_2)$; and $\hat{\mu} \in \mathbb{R}$ be a non-extreme report, i.e., a report with $\min_{\omega \in \Omega} u(\omega) < \hat{\mu} < \max_{\omega \in \Omega} u(\omega)$. Then $s(\hat{\mu}, \omega_1) = s(\hat{\mu}, \omega_2)$.*

Proof. Choose a $p \in (0, 1]$ and $\omega_3 \in \Omega$ s.t. the two random variables $Y_1 = p * \omega_1 + (1-p) \cdot \omega_3$ and $Y_2 = p * \omega_2 + (1-p) \cdot \omega_3$ both have an expected utility of $\hat{\mu}$. Then:

$$\begin{aligned} ps(\hat{\mu}, \omega_1) + (1-p)s(\hat{\mu}, \omega_3) &= \mathbb{E}[s(\hat{\mu}, Y_1)] \\ &\stackrel{\text{Lemma 1}}{=} \mathbb{E}[s(\hat{\mu}, Y_2)] \\ &= ps(\hat{\mu}, \omega_2) + (1-p)s(\hat{\mu}, \omega_3). \end{aligned}$$

Because p is positive, it follows $s(\hat{\mu}, \omega_1) = s(\hat{\mu}, \omega_2)$ as claimed. \square

⁶ For example, in the characterization given in Theorem 1, we can let the expert modify $f(\hat{\mu})$ at a set of values of $\hat{\mu}$ with measure zero as long as f remains positive and non-decreasing. This means that we can let the expert choose the value of $f(\hat{\mu})$ exactly at the jump points.

Note that if $\hat{\mu} = \min/\max_{\omega \in \Omega} u(\omega)$, then the result does not hold when $u(\omega_1) = u(\omega_2) \neq \hat{\mu}$. This is because the extreme reports $\hat{\mu} = \min/\max_{\omega \in \Omega} u(\omega)$ mean that the expert predicts such ω_1, ω_2 never to occur. There is some arbitrariness in how we punish such predictions and in particular, we could punish different ω_1, ω_2 with the same utility differently (though we cannot think of a reason why it would be helpful to do so). As with the non-full-support reports excluded in Lemma 2, we could carry this case through the rest of this paper and in our characterization provide separate rules for how scores are to be assigned in the case that $\hat{\mu} = \min/\max_{\omega \in \Omega} u(\omega)$. Such rules would not be too difficult to provide. However, in most of the real-world settings we have in mind (such as predicting a company’s profit), it appears unlikely that a sensible expert would ever provide such an extreme report. We will therefore ignore this degenerate case for the rest of this paper. Nonetheless, depending on one’s area of interest it is helpful to keep in mind that in the case of these extreme predictions, the principal has some additional degrees of freedom in how to punish the expert’s incorrect prediction.

We thus limit attention to DSRs that can depend only on the *utility* of the obtained outcome. So from now on, we will consider scoring rules $s: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ that map a reported mean $\hat{\mu}$ and the obtained utility y onto a score $s(\hat{\mu}, y)$.

4 Characterization

Now that we have shown that we can limit our attention to scoring rules s that map a reported expected utility and an observed utility onto a score, we can finally characterize proper decision scoring rules. The change in inputs to s also allows us to consider scoring rules independently of any utility function and outcome set, which in turn lets us ignore the degenerate cases of the reported mean μ being the lowest-possible and highest-possible utility. With this, we can characterize proper DSRs as follows.

Theorem 1. *A DSR $s: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ is proper if and only if*

$$s(\hat{\mu}, y) = f(\hat{\mu})(y - \hat{\mu}) + \int_0^{\hat{\mu}} f(x)dx + C \quad (12)$$

for some non-negative, non-decreasing f and constant $C \in \mathbb{R}$. If this condition is satisfied, then furthermore, s is right-action proper if $f > 0$ and strictly proper w.r.t. the mean if and only if f is strictly increasing.

While this is the mathematically simplest description, Eq. 12 is not very intuitive. We will give some interpretations and alternative statements of Theorem 1 in the next section. In particular, we give the interpretation of proper DSRs as selling shares in Section 5.1. The reader may wish to read these first to gain some intuition to help better understand the following proof.

Proof. “ \Leftarrow ” We first show that the scoring rules of the given structure are strictly proper w.r.t. the best action and strictly proper w.r.t. the mean if f is strictly increasing.

We first demonstrate that for whatever action \hat{a} the expert recommends, he is (strictly) incentivized to report that action's mean honestly. So let $U = u(O_{\hat{a}})$ be the random variable representing the utility resulting from choosing \hat{a} and let $\mu = \mathbb{E}[U]$. Now let $d > 0$. We show that the expert prefers reporting μ over reporting $\mu + d$:

$$\mathbb{E}[s(\mu + d, U)] = s(\mu + d, \mu) \quad (13)$$

$$= -df(\mu + d) + \int_0^{\mu+d} f(x)dx + C \quad (14)$$

$$= s(\mu, \mu) - df(\mu + d) + \int_{\mu}^{\mu+d} f(x)dx \quad (15)$$

$$\stackrel{f \text{ non-decr.}}{\leq} s(\mu, \mu) - df(\mu + d) + df(\mu + d) \quad (16)$$

$$= s(\mu, \mu) \quad (17)$$

$$= \mathbb{E}[s(\mu, U)] \quad (18)$$

For strictly increasing f , the inequality in the middle is strict. The same argument with a few flipped signs applies if we subtract (rather than add) d in the report.

It is left to show that it is optimal to recommend the best action. So let $U_{a^*} = u(O_{a^*})$ and $U_a = u(O_a)$ with $\mu := \mathbb{E}[U_a]$ and $\mathbb{E}[U_{a^*}] = \mu + d$ for some $d > 0$. Then the expert prefers recommending a^* with truthfully reported mean $\mu + d$ over recommending \hat{a} with truthfully reported mean μ :

$$\mathbb{E}[s(\mu + d, U_{a^*})] = s(\mu + d, \mu + d) \quad (19)$$

$$= \int_0^{\mu+d} f(x)dx + C \quad (20)$$

$$= s(\mu, \mu) + \int_{\mu}^{\mu+d} f(x)dx \quad (21)$$

$$\stackrel{f \text{ is non-negative}}{\geq} s(\mu, \mu) \quad (22)$$

$$= \mathbb{E}[s(\mu, U_a)]. \quad (23)$$

If f is strictly positive, then the inequality in the second-to-last line is strict and so s is right-action proper.

“ \Rightarrow ” Let s be a right-action proper DSR. We now show that s is of the form given in the theorem.

First, we show that $s(\hat{\mu}, y)$ is affine in y , i.e., that

$$s(\hat{\mu}, y) = f(\hat{\mu})y - g(\hat{\mu}) \quad (24)$$

for some functions f and g , with f non-negative and non-decreasing. Let $\hat{\mu} \in \mathbb{R}$ be a reported mean and let X be a random variable over \mathbb{R} with mean μ . Consider $Y = p * X + (1 - p) * x'$ and $Y' = p * \mu + (1 - p) * x'$, where p and x' are such

that both Y and Y' have mean $\hat{\mu}$. Then:

$$p\mathbb{E}[s(\hat{\mu}, X)] + (1-p)\mathbb{E}[s(\hat{\mu}, x')] = \mathbb{E}[s(\hat{\mu}, Y)] \quad (25)$$

$$\stackrel{\text{Lemma 1}}{=} \mathbb{E}[s(\hat{\mu}, Y')] \quad (26)$$

$$= ps(\hat{\mu}, x) + (1-p)\mathbb{E}[s(\hat{\mu}, x')]. \quad (27)$$

Hence, even if $\hat{\mu}$ is not the mean of X ($\hat{\mu} \neq \mu$), we have $\mathbb{E}[s(\hat{\mu}, X)] = s(\hat{\mu}, \mu)$ for all $\hat{\mu}, X$. This exactly characterizes $s(\hat{\mu}, \cdot)$ as being affine and therefore of the form in eq. 24. Further, notice that for s to be proper f has to be non-negative (otherwise, the expert would always be best off recommending the action with the lowest expected value) and for f to be right-action proper it has to be strictly positive.

It is left to show that f must be non-decreasing and that

$$g(\mu) = f(\mu)\mu - \int_0^\mu f(x)dx - C \quad (28)$$

for some $C \in \mathbb{R}$. For both of these, we will need a relationship between the rates at which f and g change. For $s(\hat{\mu}, \mu)$ to be maximal at $\hat{\mu} = \mu$, it has to be the case that for all $d > 0$

$$s(\mu + d, \mu) \leq s(\mu, \mu), \quad (29)$$

which – using eq. 24 – we can rewrite as

$$g(\mu + d) - g(\mu) \geq \mu \cdot (f(\mu + d) - f(\mu)). \quad (30)$$

Similarly, it has to be the case that for $d > 0$,

$$s(\mu, \mu + d) \leq s(\mu + d, \mu + d), \quad (31)$$

which we can rewrite as

$$g(\mu + d) - g(\mu) \leq (\mu + d) \cdot (f(\mu + d) - f(\mu)). \quad (32)$$

Note that all of these inequalities must be strict if s is to be strictly proper w.r.t. the mean.

We now show that f is non-decreasing. From ineq.s 30 and 32, it follows that for all positive d

$$\mu \cdot (f(\mu + d) - f(\mu)) \leq (\mu + d) \cdot (f(\mu + d) - f(\mu)), \quad (33)$$

which implies that $f(\mu + d) - f(\mu) \geq 0$ for all $d > 0$. If s is to be strictly proper w.r.t. the mean, then this inequality is strict.

Finally, it is left to show that g is structured as described above. By telescoping, for any $n \in \mathbb{N}_{>0}$ and any $\hat{\mu} \in \mathbb{R}$ we can write:

$$g(\hat{\mu}) = g(0) + \sum_{i=1}^n g\left(\frac{i\hat{\mu}}{n}\right) - g\left(\frac{(i-1)\hat{\mu}}{n}\right). \quad (34)$$

Since relative to any f , g can only be unique up to a constant, we will write C instead of $g(0)$. From equations 30 and 32, it follows that

$$\sum_{i=1}^n \frac{(i-1)\hat{\mu}}{n} \left(f\left(\frac{i\hat{\mu}}{n}\right) - f\left(\frac{(i-1)\hat{\mu}}{n}\right) \right) \quad (35)$$

$$\leq g(\hat{\mu}) - C \quad (36)$$

$$\leq \sum_{i=1}^n \frac{i\hat{\mu}}{n} \left(f\left(\frac{i\hat{\mu}}{n}\right) - f\left(\frac{(i-1)\hat{\mu}}{n}\right) \right) \quad (37)$$

for all $n \in \mathbb{N}_{>0}$.

We would now like to find g by taking the limit w.r.t. $n \rightarrow \infty$ of the two series. To do so, we will rewrite the two sums to interpret them as the (right and left) Riemann sums of some function.⁷ It is

$$\sum_{i=1}^n \frac{i\hat{\mu}}{n} \left(f\left(\frac{i\hat{\mu}}{n}\right) - f\left(\frac{(i-1)\hat{\mu}}{n}\right) \right) \quad (38)$$

$$= \sum_{i=1}^n \frac{i\hat{\mu}}{n} f\left(\frac{i\hat{\mu}}{n}\right) - \frac{(i-1)\hat{\mu}}{n} f\left(\frac{(i-1)\hat{\mu}}{n}\right) - \sum_{i=1}^n \frac{\hat{\mu}}{n} f\left(\frac{(i-1)\hat{\mu}}{n}\right) \quad (39)$$

$$= \hat{\mu}f(\hat{\mu}) - \sum_{i=1}^n \frac{\hat{\mu}}{n} f\left(\frac{(i-1)\hat{\mu}}{n}\right). \quad (40)$$

The last step is due to telescoping of the left-hand sum. Analogously,

$$\sum_{i=1}^n \frac{(i-1)\hat{\mu}}{n} \left(f\left(\frac{i\hat{\mu}}{n}\right) - f\left(\frac{(i-1)\hat{\mu}}{n}\right) \right) \quad (41)$$

$$= \sum_{i=1}^n \frac{i\hat{\mu}}{n} f\left(\frac{i\hat{\mu}}{n}\right) - \frac{(i-1)\hat{\mu}}{n} f\left(\frac{(i-1)\hat{\mu}}{n}\right) - \sum_{i=1}^n \frac{\hat{\mu}}{n} f\left(\frac{i\hat{\mu}}{n}\right) \quad (42)$$

$$= \hat{\mu}f(\hat{\mu}) - \sum_{i=1}^n \frac{\hat{\mu}}{n} f\left(\frac{i\hat{\mu}}{n}\right). \quad (43)$$

First note that the subtrahends are the left and right Riemann sums of f on $[0, \hat{\mu}]$. Because f is non-decreasing on \mathbb{R} , it is integrable (e.g. Rudin, 1976, Theorem 6.9). That is, both the left and right Riemann sum converge to the

⁷ In fact, we could immediately interpret them as Riemann sums of the function f^{-1} for the partition $(f(\frac{i\hat{\mu}}{n}))_{i=1, \dots, n}$. This works out but leads to a number of technical issues that are cumbersome to deal with: if f is discontinuous, then $(f(\frac{i\hat{\mu}}{n}))_{i=1, \dots, n}$ might not get arbitrarily fine, and f^{-1} could be empty somewhere between $f(0)$ and $f(\hat{\mu})$; and if f is constant on some interval, then f^{-1} contains more than one element. We can avoid these issues by first re-writing the above sum. This re-writing corresponds directly to a well-known, very intuitive formula for the integral of the inverse (e.g., Key, 1994, Theorem 1).

integral:

$$\sum_{i=1}^n \frac{\hat{\mu}}{n} f\left(\frac{i\hat{\mu}}{n}\right) \xrightarrow{n \rightarrow \infty} \int_0^{\hat{\mu}} f(x) dx \xleftarrow{n \rightarrow \infty} \sum_{i=1}^n \frac{\hat{\mu}}{n} f\left(\frac{(i-1)\hat{\mu}}{n}\right). \quad (44)$$

So for $n \rightarrow \infty$, the lower and upper bound on $g(\hat{\mu})$ converge to the same value. Hence, $g(\hat{\mu})$ must be that value, i.e.

$$g(\hat{\mu}) = C + \hat{\mu}f(\hat{\mu}) - \int_0^{\hat{\mu}} f(x) dx. \quad (45)$$

From this, eq. 12 follows as claimed. \square

As an example, we can construct the simplest possible right-action proper DSR $s(\hat{\mu}, y) = c_1 y + c_2$ for $c_1 > 0$ (see eq. 4) from using the constant function $f = c_1$ and $C = c_2$. Assuming that the true and reported mean must always be non-negative, the second-simplest example (which up to a factor of 1/2 is also given by Chen and Kash, 2011, end of Section 4) arises from $f(\hat{\mu}) = \hat{\mu}$, which gives

$$s(\hat{\mu}, y) = \hat{\mu}y - \frac{1}{2}\hat{\mu}^2. \quad (46)$$

Because s is positive affine in y , it is again easy to see that the expert wants to recommend the best action. A straightforward analysis shows that s is also strictly proper w.r.t. the mean of the optimal action.

5 Interpretation and alternative statements

5.1 Selling shares at different prices

We can interpret the proper scoring rules of Theorem 1 as ones where the principal sells $f(\hat{\mu})$ shares in the project, for an overall price of $f(\hat{\mu})\hat{\mu} - \int_0^{\hat{\mu}} f(x) dx$. Since this expression for the price is not very intuitive, let us re-write it a bit. For technical convenience, assume that f is strictly increasing and continuous and therefore invertible.⁸ By a well-known (and intuitive, see Figure 1) formula for the integral of the inverse (see, e.g., Key, 1994, Theorem 1), we have

$$f(\hat{\mu})\hat{\mu} - \int_0^{\hat{\mu}} f(x) dx = \int_{f(0)}^{f(\hat{\mu})} f^{-1}(z) dz. \quad (47)$$

Hence

$$s(\hat{\mu}, y) = f(\hat{\mu})y - \int_0^{f(\hat{\mu})} f^{-1}(z) dz + C' \quad (48)$$

for some constant $C' \in \mathbb{R}$. Now imagine that instead of reporting a mean $\hat{\mu}$ to s , the expert is offered shares in the project at various prices, with prices

⁸ As noted in the proof of Theorem 1, if f is not strictly increasing and not continuous, this analysis still works but only with a few cumbersome complications.

starting at 0. Then, the expert will start buying shares at the lowest prices and continue buying up to the price that is equal to the expected value of the project (thereby revealing that value). Let $q = f(\hat{\mu})$ denote the total number of shares bought by the agent if his reported expected value is $\hat{\mu}$. Then $f^{-1}(z)$ is the price of the z -th share (ordered by price). (Note that if f is strictly increasing, f^{-1} is, too.) Again, to act optimally, the expert stops buying shares when the cost of the marginal share is exactly the value of a single share, i.e., when $f^{-1}(q)$ ($= f^{-1}(f(\hat{\mu})) = \hat{\mu}$) is the expected utility of the principal. Hence, if the expert has bought a set of shares indicating that the value of such a share is $\hat{\mu}$, he will have paid a total of $\int_0^{f(\hat{\mu})} f^{-1}(z) dz$ for those shares; if the realized value of the project is y , those shares will be worth $f(\hat{\mu})y$; adding an arbitrary constant C' to the expert's reward, we obtain the formula in Equation 48.

As an example, we can rewrite the scoring rule resulting from $f(\hat{\mu}) = \hat{\mu}$ (see eq. 46) as $s(\mu, y) = \mu y - \int_0^\mu z dz$ to easily see why it is strictly proper w.r.t. the mean: This scoring rule corresponds to the case where we offer the same number of shares at every price above 0.

Figure 1 illustrates how f^{-1} serves as a pricing function. As noted earlier, it also shows why eq. 47 is true; further, it illustrates why (up to a constant) the expert's payoff under honest reporting is $\int_0^\mu f(x) dx$ – rotating the graph in Figure 1 counterclockwise by 90° shows that this is the area marked “expert's payoff”.

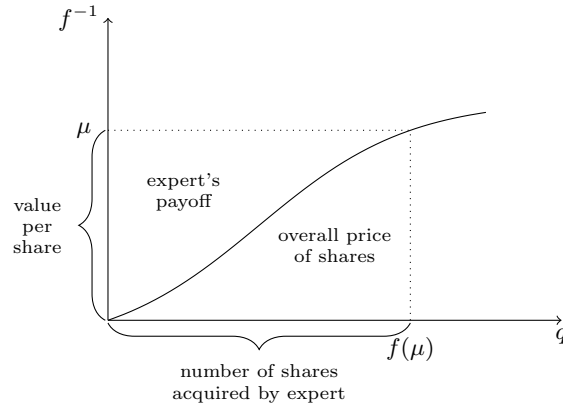


Fig. 1. An illustration of the main result (Theorem 1), one of its equivalent statements/interpretations (eq. 48) and their relationship to each other. The function graphed is f^{-1} . It is assumed that the expert reports μ honestly.

It is worth noting that scoring rules for eliciting mere predictions (see Gneiting and Raftery, 2007, for an overview and introduction) can be interpreted in a similar way. Roughly, to elicit the probability of some outcome ω , we can offer

the expert Arrow-Debreu securities on ω – assets which pay some fixed amount if ω occurs and are worthless otherwise – at different prices [cf. Savage, 1971; Schervish, 1989; Gneiting and Raftery, 2007, Section 3.2].

5.2 A characterization of differentiable scoring rules

In the proof of Theorem 1, we first show that $s(\hat{\mu}, y) = f(\hat{\mu})y - g(\hat{\mu})$ and then infer how f and g relate to each other for s to be maximal at $\hat{\mu} = y$. If s and hence f and g are differentiable in $\hat{\mu}$, it is immediately clear what to do: for any fixed μ , $\frac{d}{d\hat{\mu}}s(\hat{\mu}, \mu)$ has to be 0 at $\hat{\mu} = \mu$. This gives us the following corollary for differentiable scoring rules. The theorem in this form makes it easier to compare our result to that of Othman and Sandholm (2010, Section 2.3.2) (discussed in Section 6.1 of our paper), as well as to some results on direct elicitation of properties (see Section 6.3 and in particular Footnote 10).

Corollary 1. *A differentiable DSR $s: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ is proper if and only if there are differentiable f, g s.t. $s(\hat{\mu}, y) = f(\hat{\mu})y - g(\hat{\mu})$ with $g'(\hat{\mu}) = \hat{\mu}f'(\hat{\mu})$ for all $\hat{\mu} \in \mathbb{R}$; $f' \geq 0$; and $f \geq 0$. Furthermore, s is right-action proper if $f > 0$ and strictly proper w.r.t. to the reported mean of the optimal action if and only if $f' > 0$.*

5.3 Characterization in terms of convex functions and subgradients

Existing work on elicitation often uses the terminology of convex functions and their subgradients (see Section 6.2). Indeed, our result can be put in these terms as well.

Corollary 2. *A DSR $s: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ is proper if and only if there is a convex, non-decreasing h with a subgradient h' s.t. $s(\hat{\mu}, y) = h'(\hat{\mu})(y - \hat{\mu}) + h(\hat{\mu})$. Furthermore, s is right-action proper if $h' > 0$ and strictly proper w.r.t. to the reported mean of the optimal action if and only if h is strictly convex (i.e., if h' is strictly increasing).*

Proof. Follows directly from Theorem 1 and the equivalence of convex functions and integrals over subderivatives, see e.g., Theorem 24.2 and Corollary 24.2.1 of Rockafellar (1970). \square

6 Related work

6.1 Othman and Sandholm (2010)

As far as we can tell, Othman and Sandholm (2010) are the first to consider the problem of designing scoring rules for decision making. They study a simplified case in which the set of outcomes Ω has only two elements, one with a utility of 1, the other with a utility of 0. Note that the two-outcome-case is special because the mean of a binary random variable fully determines its distribution.

In Section 2.3.2, they give a characterization of differentiable scoring rules with good incentives, which is a special case of our Corollary 1.⁹

6.2 Chen, Kash, et al. (2014)

Chen, Kash, et al. (2014) also characterize scoring rules for decision making (an alternative proof of this characterization from a general result on elicitation is given by Frongillo and Kash, 2014, Section E.1). Their setting is more general than ours in that they allow arbitrary *decision rules*. That is, they allow principals who do not choose the best action, but, for instance, randomize over the best few actions with probabilities depending on the expert’s prediction. Randomizing over *all* actions in particular, even if not uniformly, allows any proper scoring rule for mere prediction to be used to construct strictly proper scoring rules for decision making (Chen, Kash, et al., 2014, Section 4).

They also characterize a much larger class of scoring rules: they merely require that the expert honestly reports the probability distribution that the recommendation gives rise to and allow scoring rules which strictly incentivize misreporting what the best action is. For principals who (like those in our setting) deterministically choose the best action according to the expert’s report, their result is especially easy to derive and understand. To elicit an honest report about the probability distribution resulting from taking the expert’s recommended action, the principal can use any (strictly) proper scoring rule for mere prediction (as defined and characterized by, e.g., Gneiting and Raftery, 2007, Section 2). As an example, consider the quadratic (Brier) scoring rule discussed at the end of Section 2. Interest in this larger class of scoring rules is harder to motivate – usually the principal will primarily want to ensure that the best action is taken, and eliciting honest and accurate predictions about the best action will only be a secondary priority.

Chen, Kash, et al. (2014, Section 5) do also consider the goal of characterizing preferences over lotteries for which making the best recommendation can be incentivized. But they do not give a characterization of (right-action) proper scoring rules for utility-maximizing principals or of what information can be extracted along with the best action.

⁹ We briefly show that the two characterizations are equivalent. They characterize differentiable scoring rules s as ones where A) $s(p, 1) > s(p, 0)$ and B) $\frac{s'(p, 1)}{s'(p, 0)} = \frac{p-1}{p}$ for all p . A is equivalent to being able to write $s(p, 1) = f(p) - g(p)$ and $s(p, 0) = -g(p)$ for some differentiable g and positive, differentiable f . With that we can re-write B:

$$\frac{s'(p, 1)}{s'(p, 0)} = \frac{p-1}{p} \Leftrightarrow \frac{f'(p) - g'(p)}{g'(p)} = \frac{p-1}{p} \Leftrightarrow g'(p) = pf'(p), \quad (49)$$

which is exactly the relationship between f', g' stated in our Corollary 1, except that Othman and Sandholm seem to overlook the necessity of $s'(p, 1) > 0$ (which is equivalent to $f' \geq 0$ in our framework). Note that Othman and Sandholm use different names for scoring rules. In particular, they use “ $f(p)$ ” for $s(p, 1)$ and “ $g(p)$ ” for $s(p, 0)$.

That said, the main result from Chen et al., or Gneiting and Raftery’s characterization (or even the more general characterization of Frongillo and Kash, 2014) can be used to replace part of our proof. Specifically, Corollary 2 is structurally very similar to these results. The only difference is that we impose additional constraints on the constituent functions, for instance, that f only depends on the mean. Our proof works by first showing that only the reported mean matters and then deriving the structure of s . Alternatively, we could have used Gneiting and Raftery’s characterization to first derive the structure without using the results from Section 3 and only then derived and incorporated the additional requirements necessary to ensure that s elicits honest recommendations (and not only honest reporting of the distribution for the recommended action).

6.3 Direct elicitation of properties

Typically, when designing scoring rules for prediction (without the recommendation component) the goal is to elicit entire probability distributions over outcomes. But a recent line of work has explored the direct elicitation of particular properties of the distribution *without* eliciting the entire distribution (e.g. Lambert et al., 2008; Gneiting, 2011; Abernethy and Frongillo, 2012; Bellini and Bignozzi, 2015). Of course, in principle, one could elicit entire distributions and would thereby elicit all properties. But eliciting, say, a single-valued point forecast may be required “for reasons of decision making, market mechanisms, reporting requirements, communications, or tradition, among others” (Gneiting, 2011, Section 1). Lemma 2 gives another reason to study scoring rules for eliciting just the expected utility, albeit with the additional requirements that the expected score under honest reporting must be the same for two variables with equal mean (Lemma 1) and that the expected score under honest reporting must be increasing in the true mean of the random variable. Results from the literature on property elicitation can also be used to replace parts of the proof of our main result.¹⁰

¹⁰ For instance, consider Theorem 3 of Lambert et al. (2008), which characterizes continuously differentiable scoring rules for any elicitable property. Their characterization relies on what they call the “signature” of a property. Up to normalization, we can use as a signature of the expected utility, $v(t, \omega) = u(\omega) - t$. So the class of scoring rules for the mean according to Theorem 3 of Lambert et al. (2008) is the class of functions s that can be written as $s(\mu, y) = s_0(y) + \int_{\mu_0}^{\mu} \lambda(t)(y - t)dt$, where μ_0 is the lowest reportable mean and λ is some non-negative, continuous weight function. In their setting, s_0 can be an arbitrary function. But for s to be a proper *decision* scoring rule, s_0 has to be positive affine in y .

Now take our characterization of proper differentiable scoring rules for decision making in Corollary 1. Here, we must have $g'(\mu) = \mu f'(\mu)$. Integrating both sides yields $g(\mu) = C + \int_{\mu_0}^{\mu} t f'(t) dt$ for some constant $C \in \mathbb{R}$. Corollary 1 thus becomes

$$s(\mu, y) = f(\mu)y - \int_{\mu_0}^{\mu} t f'(t) dt - C = \int_{\mu_0}^{\mu} f'(t)(y - t) dt + C'y - C, \quad (50)$$

Acknowledgements

We thank Ian Kash for comments on this paper. We also thank Carl Shulman and participants of the DRIV(E) research group meeting, the Economics Theory Lunch (especially Atila Abdulkadiroglu, Attila Ambrus and David McAdams) and the CS-Econ seminar (especially Brandon Fain, Kamesh Munagala and Deb-malya Panigrahi) at Duke for helpful discussions. We are thankful for support from the NSF under award IIS-1814056.

References

- Abernethy, Jacob D. and Rafael M. Frongillo (June 2012). “A Characterization of Scoring Rules for Linear Properties”. In: *Proceedings of the 25th Annual Conference on Learning Theory*. Ed. by Shie Mannor, Nathan Srebro, and Robert C. Williamson. Vol. 23. Proceedings of Machine Learning Research. Edinburgh, Scotland: PMLR, pp. 27.1–27.13. URL: <http://proceedings.mlr.press/v23/abernethy12.html>.
- Ashlagi, Itai, Shahar Dobzinski, and Ron Lavi (May 2012). “Optimal Lower Bounds for Anonymous Scheduling Mechanisms”. In: *Mathematics of Operations Research* 37.2, pp. 244–258. DOI: 10.1287/moor.1110.0534.
- Bellini, Fabio and Valeria Bignozzi (2015). “On elicitable risk measures”. In: *Quantitative Finance* 15.5, pp. 725–733. DOI: 10.1080/14697688.2014.946955.
- Bikhchandani, Sushil, Shurojit Chatterji, Ron Lavi, Ahuva Mu’alem, Noam Nisan, and Arunava Sen (July 2006). “Weak Monotonicity Characterizes Deterministic Dominant-Strategy Implementation”. In: *Econometrica* 74.4, pp. 1109–1132. DOI: 10.1111/j.1468-0262.2006.00695.x.
- Boutilier, Craig (2012). “Eliciting Forecasts from Self-interested Experts: Scoring Rules for Decision Makers”. In: *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems*. Ed. by Conitzer, Winikoff, Padgham, and van der Hoek.
- Brier, Glenn W. (Jan. 1950). “Verification of Forecasts Expressed in Terms of Probability”. In: *Monthly Weather Review* 78.1.
- Carroll, Gabriel (May 2019). “Robust incentives for information acquisition”. In: 181, pp. 382–420. DOI: 10.1016/j.jet.2019.03.001.
- Chen, Yiling and Ian A. Kash (2011). “Information elicitation for decision making”. In: *Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2011), Taipei, 2-6 May 2011, edited by Tumer, Yolum, Sonenberg, and Stone*. Springer. URL: https://dash.harvard.edu/bitstream/handle/1/5128480/chen_aamas11_0.pdf?sequence=1.

for some constant $C' \geq 0$. So in the case of differentiable DSRs, Theorem 3 of Lambert et al. can be used to replace part of our proof.

- Chen, Yiling, Ian A. Kash, Mike Ruberry, and Victor Shnayder (June 2014). “Eliciting Predictions and Recommendations for Decision Making”. In: *ACM Transactions on Economics and Computation*. Vol. 2. 2. Chap. 6. URL: <http://yiling.seas.harvard.edu/wp-content/uploads/a6-chen.pdf>.
- D. Hartline, Jason, Yingkai Li, Liren Shan, and Yifan Wu (July 2020). *Optimization of Scoring Rules*. URL: <https://arxiv.org/pdf/2007.02905.pdf>.
- Frongillo, Rafael M. and Ian A. Kash (2014). *General Truthfulness Characterizations Via Convex Analysis*. Short version published in WINE 2014. URL: <https://www.cs.colorado.edu/~raf/media/papers/wine-affine.pdf>.
- Gneiting, Tilmann (2011). “Making and Evaluating Point Forecasts”. In: *Journal of the American Statistical Association* 106.494, pp. 746–762. DOI: 10.1198/jasa.2011.r10138.
- Gneiting, Tilmann and Adrian E. Raftery (Mar. 2007). “Strictly Proper Scoring Rules, Prediction, and Estimation”. In: *Journal of the American Statistical Association* 102.477. DOI: 10.1198/016214506000001437.
- Good, I. J. (1952). “Rational Decisions”. In: *Journal of the Royal Statistical Society. Series B (Methodological)* 14 (1), pp. 107–114. URL: <https://www.jstor.org/stable/2984087>.
- Johnstone, David J., Victor Richmond R. Jose, and Robert L. Winkler (2011). “Tailored Scoring Rules for Probabilities”. In: *Decision Analysis* 8.4, pp. 256–268. DOI: 10.1287/deca.1110.0216.
- Key, Eric (Mar. 1994). “Disks, Shells, and Integrals of Inverse Functions”. In: *The College Mathematics Journal* 25.2, pp. 136–138.
- Lambert, Nicolas, David M. Pennock, and Yoav Shoham (2008). “Eliciting Properties of Probability Distributions”. In: *ACM Conference on Electronic Commerce*. Extended Abstract.
- Lavi, Ron and Chaitanya Swamy (June 2007). “Truthful Mechanism Design for Multi-Dimensional Scheduling via Cycle Monotonicity”. In: *Economics and Computation 2007*. San Diego, California, USA: ACM.
- McCarthy, John (Sept. 1956). “Measures of the Value of Information”. In: *Proceedings of the National Academy of Sciences of the United States of America* 42 (9), pp. 654–655. DOI: 10.1073/pnas.42.9.654. URL: <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC534271/>.
- Neyman, Eric, Georgy Noarov, and S. Matthew Weinberg (Feb. 2020). *Binary Scoring Rules that Incentivize Precision*. URL: <https://arxiv.org/pdf/2002.10669.pdf>.
- Osband, Kent (Oct. 1989). “Optimal Forecasting Incentives”. In: *Journal of Political Economy* 97.5, pp. 1091–1112. URL: <https://www.jstor.org/stable/1831887>.
- Othman, Abraham and Tuomas Sandholm (2010). “Decision Rules and Decision Markets”. In: *Proc. of 9th Int. Conf. on Autonomous Agents and Multiagent Systems (AAMAS 2010)*, van der Hoek, Kaminka, Lespérance, Luck and Sen (eds.), May, 10–14, 2010, Toronto, Canada, pp. 625–632.

- Rochet, Jean (1987). “A necessary and sufficient condition for rationalizability in a quasi-linear context”. In: *Journal of Mathematical Economics* 16.2, pp. 191–200.
- Rockafellar, R. Tyrrell (1970). *Convex Analysis*. Princeton, NJ: Princeton University Press.
- Rudin, Walter (1976). *Principles of Mathematical Analysis*. 3rd ed. McGraw-Hill.
- Savage, Leonard J. (Dec. 1971). “Elicitation of Personal Probabilities and Expectations”. In: *Journal of the American Statistical Association* 66 (336), pp. 783–801.
- Schervish, Mark J. (1989). “A General Method for Comparing Probability Assessors”. In: *The Annals of Statistics* 17.4, pp. 1856–1879. URL: <https://projecteuclid.org/euclid.aos/1176347398>.
- Shi, Peng, Vincent Conitzer, and Mingyu Guo (2009). “Prediction Mechanisms That Do Not Incentivize Undesirable Actions”. In: *Workshop on Internet and Network Economics (WINE) 2009*. Ed. by S. Leonardi. Vol. 5929. Lecture Notes in Computer Science. Springer.
- Zermeño, Luis (Nov. 2011). “A Principal-Expert Model and the Value of Menus”. URL: <http://economics.mit.edu/files/7299>.
- (Mar. 2012). “The Role of Authority in a General Principal-Expert Model”. URL: <http://economics.mit.edu/files/7627>.

A An elementary proof of the continuity of proper DSRs

We here give a proof of the continuity claim used in the proof of Lemma 1. Our proof uses only basic real calculus and no convex analysis.

Claim: Let s be a proper DSR and $\omega_L, \omega_H \in \Omega$ be two outcomes. For $p \in (0, 1)$, let $R_p = p * \omega_H + (1 - p) * \omega_L$ and Q_p be the distribution of that random variable. Then $\mathbb{E}[s(Q_p, R_p)]$ is continuous in p .

Proof. We conduct a proof by contradiction. Assume there is a discontinuity at $p = \tilde{p}$. Then one of two things is true:

1. There is a $\delta > 0$ such that for all $\epsilon > 0$

$$\mathbb{E}[s(Q_{\tilde{p}}, R_{\tilde{p}})] > \mathbb{E}[s(Q_{\tilde{p}-\epsilon}, R_{\tilde{p}-\epsilon})] + \delta. \quad (51)$$

2. There is a $\delta > 0$ such that for all $\epsilon > 0$

$$\mathbb{E}[s(Q_{\tilde{p}}, R_{\tilde{p}})] < \mathbb{E}[s(Q_{\tilde{p}+\epsilon}, R_{\tilde{p}+\epsilon})] - \delta. \quad (52)$$

We derive contradictions from these two cases separately.

1. Imagine that the expert believes that under the optimal action, O is distributed according to $R_{\tilde{p}-\epsilon}$. Then for small enough ϵ , the expert prefers submitting $Q_{\tilde{p}}$ over submitting $Q_{\tilde{p}-\epsilon}$, because of the following.

$$\mathbb{E}[s(Q_{\tilde{p}}, R_{\tilde{p}-\epsilon})] = (\tilde{p} - \epsilon)s(Q_{\tilde{p}}, \omega_H) + (1 - (\tilde{p} - \epsilon))s(Q_{\tilde{p}}, \omega_L) \quad (53)$$

$$\stackrel{\epsilon \rightarrow 0}{\rightarrow} \tilde{p}s(Q_{\tilde{p}}, \omega_H) + (1 - \tilde{p})s(Q_{\tilde{p}}, \omega_L) \quad (54)$$

$$= \mathbb{E}[s(Q_{\tilde{p}}, R_{\tilde{p}})]. \quad (55)$$

This means that there exists an $\epsilon > 0$ such that

$$\mathbb{E}[s(Q_{\bar{p}}, R_{\bar{p}-\epsilon})] > \mathbb{E}[s(Q_{\bar{p}}, R_{\bar{p}})] - \delta/2 > \mathbb{E}[s(Q_{\bar{p}-\epsilon}, R_{\bar{p}-\epsilon})] + \delta/2. \quad (56)$$

This contradicts propriety.

2. This case is actually a little harder. We need the fact that $s(Q_p, \omega_H)$ is monotonically increasing in p and $s(Q_p, \omega_L)$ is monotonically decreasing in p , i.e. that for all $p_2 > p_1$ it is $s(Q_{p_2}, \omega_H) \geq s(Q_{p_1}, \omega_H)$ and $s(Q_{p_2}, \omega_L) \leq s(Q_{p_1}, \omega_L)$. This in turn can be shown by contradiction with different cases. For instance, imagine there were some $p_2 > p_1$ s.t. $s(Q_{p_2}, \omega_H) < s(Q_{p_1}, \omega_H)$ and $s(Q_{p_2}, \omega_L) < s(Q_{p_1}, \omega_L)$. Then the expert always prefers submitting Q_{p_1} over submitting Q_{p_2} , even when the true distribution is Q_{p_2} . Because $s(Q_p, \omega_H)$ and $s(Q_p, \omega_L)$ are monotone in $p \in (0, 1)$, they are bounded on every $[a, b]$ with $0 < a \leq b < 1$.

With this, we can make a similar argument as above. Imagine that the expert believes that under the optimal action, O is distributed according to $R_{\bar{p}}$. Then for small enough ϵ , the expert prefers submitting $Q_{\bar{p}+\epsilon}$ over submitting $Q_{\bar{p}}$, because of the following.

$$\mathbb{E}[s(Q_{\bar{p}}, R_{\bar{p}})] = \tilde{p}s(Q_{\bar{p}+\epsilon}, \omega_H) + (1 - \tilde{p})s(Q_{\bar{p}+\epsilon}, \omega_L) \quad (57)$$

$$\stackrel{\epsilon \rightarrow 0}{\leftarrow} \tilde{p}s(Q_{\bar{p}+\epsilon}, \omega_H) + (1 - \tilde{p})s(Q_{\bar{p}+\epsilon}, \omega_L) \quad (58)$$

$$+ \epsilon(s(Q_{\bar{p}+\epsilon}, \omega_H) - s(Q_{\bar{p}+\epsilon}, \omega_L)) \quad (59)$$

$$= \mathbb{E}[s(Q_{\bar{p}+\epsilon}, R_{\bar{p}+\epsilon})]. \quad (60)$$

The line in the middle is due to the boundedness of $s(Q_{\bar{p}+\epsilon}, \omega_H)$ and $s(Q_{\bar{p}+\epsilon}, \omega_L)$, which implies that $\epsilon(s(Q_{\bar{p}+\epsilon}, \omega_H) - s(Q_{\bar{p}+\epsilon}, \omega_L)) \xrightarrow{\epsilon \rightarrow 0} 0$. This means that there exists an $\epsilon > 0$ such that

$$\mathbb{E}[s(Q_{\bar{p}+\epsilon}, R_{\bar{p}})] > \mathbb{E}[s(Q_{\bar{p}+\epsilon}, R_{\bar{p}+\epsilon})] - \delta/2 > \mathbb{E}[s(Q_{\bar{p}}, R_{\bar{p}})] + \delta/2. \quad (61)$$

This contradicts propriety again. We conclude that for any proper DSR s , the term $\mathbb{E}[s(Q_p, O_p)]$ must be continuous in p . \square