Eliciting information for decision making

CASPAR OESTERHELD AND VINCENT CONITZER, JANUARY 19, 2020

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 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.

We extend our analysis to eliciting recommendations and predictions from multiple experts. With a few modifications, the above characterization for the single-expert case carries over. Among other implications, this characterization implies that no expert should be able to "short-sell" shares in the principal's project and thereby profit if the project goes poorly. Inspired by our results, we finally give a detail-free auction-based mechanism for eliciting information for decision making.

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? What compensation schemes are effective if the firm can consult multiple experts?

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 his[1] 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

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

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), and as such can be viewed as analogous to characterizations of incentive compatible choice functions in mechanism design, including characterizations based on cycle monotonicity or weak monotonicity.

In Section 2, 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 a prediction. 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 2.2). Intuitively, rewarding the expert for getting anything else about the distribution right will make him recommend actions whose outcome is easy to predict as opposed to actions with high expected utility. 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 2.3). 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 2.4.1). 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 utility of the share, thereby revealing the expected utility. Moreover, once the expert has some positive share in the principal's utility, he will be (strictly) incentivized to recommend an optimal action.

Our work on the multi-expert case is also inspired by work on eliciting mere predictions. Much of this work has focused on a particular type of elicitation mechanism known as *prediction markets* [e.g., Hanson, 2003; Pennock and Sami, 2007]. Prediction markets can be constructed using single-expert scoring rules as a building block. For example, in a *market scoring rule*, agents successively update the probability estimate, and an agent that updated the estimate from $\hat{P}_t$ to $\hat{P}_{t+1}$ is eventually rewarded $s(\hat{P}_{t+1}, \omega) - s(\hat{P}_t, \omega)$. But prediction markets can also be implemented to resemble real-world securities markets. In particular, a common approach is to let the experts trade *Arrow-Debreu securities* that each pay out a fixed amount – say, \$1 – if a given event happens, and \$0 otherwise. Then, at any point, the price at which this security trades can be seen as the current market consensus of the probability that the event takes place.[2]

Our goal is, again, to lay the foundations for a theory of eliciting recommendations and predictions from multiple experts that is in some sense analogous to existing work on prediction markets. Indeed, existing work (as reviewed in Section 3.4) has proposed *decision markets* (or conditional prediction markets) as a variation of prediction markets that can advise a decision maker [e.g. Hanson, 1999; Hanson, 2002; Berg and Rietz, 2003; Hanson, 2006; Hanson, 2013, Section IV]. These do not set good incentives in their most commonly described form. Roughly, the problem is (as in the single-expert case) that experts may misreport in order to make the principal take an action that they can predict well (relative to the market) [Othman and Sandholm, 2010, Section 3]. Chen et al. [2014, Section 4] show that choosing suboptimal actions with some small probability can solve the incentive problems while still ensuring that we take the best action most of the time. But in practice, it would be challenging to commit to, based on the outcome of a coin toss, pursuing a course of action that is likely to doom the entire firm. Hence, in this paper, we consider only mechanisms that deterministically (rather than with high probability) take an optimal action.

To find out what properties elicitation mechanisms for decision making with multiple experts must possess, we provide a characterization of such mechanisms (Section 3). We build on the

[2]There is a close correspondence between Arrow-Debreu securities markets and market scoring rules [Hanson, 2003; Hanson, 2007; Pennock and Sami, 2007, Section 4; Chen and Pennock, 2007; Agrawal et al., 2009; Chen and Vaughan, 2010].

single-expert characterization by designing each expert's scoring function on the assumption that everyone else reports honestly (i.e., non-strategically). If all scoring rules are designed in this way, honest reporting is a Nash equilibrium. The resulting characterization is structurally the same as that for the single-expert case, but significantly richer, because the multi-expert setting allows for a number of new possibilities. For example, one may (but is not required to) reward experts for accurately predicting other experts' reports. Moreover, an expert's recommended action and prediction of the utility to be expected from that action are, in general, conditioned on other experts' reports.

Besides the impossibility to elicit anything but the expected utility along with an honest recommendation, the characterization implies an interesting impossibility about elicitation mechanisms for decision making (Section 3.3). In prediction markets, traders who change the market probabilities significantly and toward the truth – and therefore can bee seen as contributing a significant amount of accurate information – reap higher rewards. Our characterization shows that, other things being equal, an expert gets the same reward for honestly reporting a specific mean, regardless of whether the report helps in identifying the optimal action or not.

2 ELICITING FROM A SINGLE EXPERT

2.1 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 $u\colon \Omega \to \mathbb{R}$.

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.[3]

The principal may ask the expert to recommend some action $\hat{a}$ – with the intention of getting the expert to report one that maximizes $\mathbb{E}_P\left[u(O) \mid \hat{a}\right]$ – 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 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 et al., 2014; cf. Zermeño, 2011, 2012]. This can help incentivize the expert to report honestly. For instance, Chen 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 thus learn from the

[3]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 $e \in H$. 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.

expert about the effects of all actions, not just the one she ends up taking.[4] 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 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 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 more effective when they are chosen with low probability as a result of randomization over suboptimal actions, because they will then catch adversaries and competitors by surprise.

In our setting, the principal has no way of verifying the information she receives. We also assume that the expert has no intrinsic interests in the principal's endeavors.[5] To nevertheless incentivize the expert to report his private information honestly, the principal may therefore use a *decision scoring rule* (DSR) $s\colon \Delta(\Omega) \times \Omega \to \mathbb{R}$, 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 constructed as part of the proof of Lemma 2.3. 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 make good recommendations a^* but also to take unobservable actions in the principal's favor whenever he can[6] [cf. the notion of "principal alignment" in prediction markets, as discussed by Shi et al., 2009].

[4]Besides intrinsic curiosity, the principal may also be interested in how large the benefit is of asking an expert – "how better off am I now than I would have been had I taken the action that I thought was best before asking the expert?"

[5]Of course, if the expert wants the principal to succeed, things become easier for the principal. However, the expert's preferences may also conflict with the principal's. Boutilier [2012] discusses this problem in the context of using prediction markets for decision making. There is also a broad literature on various forms of "persuasion games" [e.g. Crawford and Sobel, 1982]. However, this literature generally assumes that the expert is interested *only* in the principal's decision and not additionally motivated by a payment scheme that the principal can select.

[6]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.

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 2.1. 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

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

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

We limit our attention to designing proper DSRs.[7] 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 2.4), no DSR is strictly proper in this sense. We will therefore define partially strict versions of propriety.

Definition 2.2. 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 $\hat{P}_{\hat{a}} \in \Delta(\Omega)$,

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

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

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

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

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

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 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 $\hat{P}_{\hat{a}}$ 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 define no other form of (partially) strict propriety here.

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

[7] Of course, there may not be a problem with using a mechanism that strictly incentivizes some "harmless" form of misrepresentation – in particular misrepresentation that is guaranteed not to affect our decision. But we could modify any such mechanism to obtain a proper one by "doing the misrepresentation for the player". This is essentially a single-player version of the revelation principle.

company, buying said apartment, starting a family, etc.).[8] 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 *ceteris paribus* 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" [Chen et al., 2014, Sections 3-4]. It is informative to work through an example of such a scoring rule for prediction and why it fails to be right-action proper. 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. 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 right-action proper. It generally incentivizes the expert to recommend an action a that is easiest to predict. For instance, suppose that $\Omega = \{\omega_1, ..., \omega_m\}$, that the optimal action a^* leads to the uniform distribution $O_{a^*} = \frac{1}{m} * \omega_1 + ... + \frac{1}{m} * \omega_m$, while a' leads to $O_{a'} = 1 * \omega_1$ deterministically. Then the expert will (assuming $m > 1$) always prefer recommending the suboptimal a', since

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

2.2 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 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 the optimal (recommended) action and the expected utility of that action. That is, proper scoring rules can be strictly proper w.r.t. the recommendation and mean, but nothing else.

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

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

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

[8] 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) 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.

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, $\mathbb{E}\left[s(Q_p, R_p)\right]$ is non-decreasing in p. We claim that $\mathbb{E}\left[s(Q_p, R_p)\right]$ 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, $\mathbb{E}\left[s(Q_p, R_p)\right] = \max_{p'\in[0,1]} \mathbb{E}\left[s(Q_{p'}, R_p)\right]$. So $\mathbb{E}\left[s(Q_p, R_p)\right]$ 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 B.

We now use continuity to show that $\mathbb{E}_{O\sim P_a}\left[s(P_a, O)\right]$ 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 $\mathbb{E}_{O\sim P_a}\left[s(P_a, O)\right] \geq \mathbb{E}\left[s(Q_p, R_p)\right]$ for all p with $\mu > pH + (1-p)L$ and $\mathbb{E}_{O\sim P_a}\left[s(P_a, O)\right] \leq \mathbb{E}\left[s(Q_p, R_p)\right]$ for all p with $\mu < pH + (1-p)L$. But because $\mathbb{E}\left[s(Q_p, R_p)\right]$ is continuous, this implies $\mathbb{E}_{O\sim P_a}\left[s(P_a, O)\right] = \mathbb{E}\left[s(Q_p, R_p)\right]$ 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 $\mathbb{E}\left[s(Q_p, R_p)\right]$. □

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 reporting his prediction about 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 right-action proper DSR without affecting his expected score.

Lemma 2.4. *Let s be a proper DSR and $P_a, \hat{P}_a \in \Delta(\Omega)$. Then if*

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

and $\operatorname{supp}(P_a) \subseteq \operatorname{supp}(\hat{P}_a)$*, it must be the case that* $\mathbb{E}_{P_a}\left[s(P_a, O)\right] = \mathbb{E}_{P_a}\left[s(\hat{P}_a, O)\right]$.

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

$$\begin{aligned}
\mathbb{E}_{\hat{P}_a}\left[s(\hat{P}_a, O)\right] &= p\mathbb{E}_{P_a}\left[s(\hat{P}_a, O)\right] + (1-p)\mathbb{E}_{P'_a}\left[s(\hat{P}_a, O)\right] && (8)\\
&\underset{s\text{ is proper}}{\leq} p\mathbb{E}_{P_a}\left[s(\hat{P}_a, O)\right] + (1-p)\mathbb{E}_{P'_a}\left[s(P'_a, O)\right] && (9)\\
&\underset{s\text{ is proper}}{\leq} p\mathbb{E}_{P_a}\left[s(P_a, O)\right] + (1-p)\mathbb{E}_{P'_a}\left[s(P'_a, O)\right] && (10)\\
&\underset{\text{Lemma 2.3}}{=} \mathbb{E}_{\hat{P}_a}\left[s(\hat{P}_a, O)\right]. && (11)
\end{aligned}$$

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

Lemma 2.4 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 $\min/\max_{\hat{\omega}\in\Omega} u(\hat{\omega})$ 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.[9] 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 2.5. *Let $s\colon \mathbb{R}\times\Omega \to \mathbb{R}$ 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:

$$ps(\hat{\mu}, \omega_1) + (1 - p)s(\hat{\mu}, \omega_3) = \mathbb{E}[s(\hat{\mu}, Y_1)] \underset{\text{Lemma 2.3}}{=} \mathbb{E}[s(\hat{\mu}, Y_2)] = ps(\hat{\mu}, \omega_2) + (1 - p)s(\hat{\mu}, \omega_3).$$

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

Hence, except for degenerate cases[10] a proper DSR can depend only on the *utility* of the obtained outcome. So from now on, we will without loss of generality consider scoring rules $s\colon \mathbb{R}\times\mathbb{R} \to \mathbb{R}$ that map a reported mean $\hat{\mu}$ and the obtained utility y onto a score $s(\hat{\mu}, y)$.

2.3 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. (Structurally, the characterization resembles a few existing results on proper elicitation. We discuss these in Section 2.5, after giving some alternative forms of the result in Section 2.4.)

Theorem 2.6. *A DSR $s\colon \mathbb{R}\times\mathbb{R} \to \mathbb{R}$ is proper iff*

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

[9]For example, in the characterization given in Theorem 2.6, 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 can change $f(\hat{\mu})$ exactly at the jump points.

[10]Note that if $\hat{\mu} = \min/\max_{\omega\in\Omega} u(\omega)$ and $u(\omega_1) = u(\omega_2) \neq \hat{\mu}$ then the score $s(\hat{\mu}, \cdot)$ *can* be different for ω_1 and ω_2, because in that case both are predicted to occur with probability 0. We ignore this case, e.g., because in this case it is not very useful to treat ω_1 and ω_2 differently – both prove dishonest reporting just the same. We may also not want our scoring rule to make use of information about what $\min/\max_{\hat{\omega}\in\Omega} u(\hat{\omega})$ is. Similarly, we ignore the case in which $\hat{\mu} < \min_{\omega\in\Omega} u(\omega)$ or $\hat{\mu} > \max_{\omega\in\Omega} u(\omega)$, which similarly allows varying scores between outcomes of equal utility.

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.

For completeness, a stand-alone proof of Theorem 2.6 is given in Appendix C. Besides the results of Section 2.2, the proof relies only on elementary techniques in real analysis (such as Riemann integration). It uses no existing results from the mechanism design literature. Equivalent results can also be obtained by combining results from the literature on the shape of elicitation functions – which in turn are usually proven with more advanced techniques from convex analysis – with the results of Section 2.2. For example, every proper DSR is also a proper scoring rule for eliciting predictions. To obtain a characterization of proper DSRs we can therefore use a characterization of proper scoring rules. We still need to implement additional restrictions, of course, such as dependence on nothing but the mean (Lemma 2.4) and f being non-decreasing. This leads to the characterization we give in Corollary 2.8, which is equivalent to the one given here. In Section 2.5, we will give a more detailed overview of related work that can be used to replace parts of the proof in Appendix C, if only for special types of (e.g., differentiable) DSRs.

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 $f = c_1$. Again, this s is not strictly proper w.r.t. the reported mean. But it is easy to see why it is right-action proper, as the expert can influence his expected payoff only by recommending an action that yields the highest utility y in expectation. Imagine the principal manages a company with the utility being the profit generated by the company. Then for, say, $c_1 = 1/5$, scoring according to eq. 4 is like promising the expert 20% of the profits of the company. Of course, the expert will now share the principal's goal of maximizing the company's profit.

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

$$s(\hat{\mu}, y) = \hat{\mu} y - \frac{1}{2}\hat{\mu}^2. \tag{13}$$

Because s is linear 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.

In the next section, we will give some alternative interpretations of Theorem 2.6, one of which (Section 2.4.1) in particular will make it easy to see why the scoring rules of Theorem 2.6 are strictly proper w.r.t. the mean of the optimal action.

2.4 Interpretation and alternative statements

2.4.1 Selling shares at different prices. We can interpret the proper scoring rules of Theorem 2.6 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.[11] By Theorem 1 of Key [1994] (a well-known and intuitive formula for the integral of the inverse), we have

$$f(\hat{\mu})\hat{\mu} - \int_0^{\hat{\mu}} f(x)dx = \int_{f(0)}^{f(\hat{\mu})} f^{-1}(z)dz. \tag{14}$$

Hence

$$s(\hat{\mu}, y) = f(\hat{\mu})y - \int_0^{f(\hat{\mu})} f^{-1}(z)dz + C' \tag{15}$$

[11] As noted in the proof of Theorem 2.6 in Appendix C, if f is not strictly increasing and not continuous, this analysis still works but only with a few cumbersome complications.

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 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 $f^{-1}(q)$ (and therefore $f^{-1}(f(\hat{\mu})) = \hat{\mu}$) 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 15.

As an example, we can rewrite the scoring rule resulting from $f(\hat{\mu}) = \hat{\mu}$ (see eq. 13) 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.

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].

2.4.2 A characterization of differentiable scoring rules. In the proof of Theorem 2.6 in Appendix C, we first showed that $s(\hat{\mu}, y) = f(\hat{\mu})y - g(\hat{\mu})$ and then inferred 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 2.5.1 of our paper), as well as to some results on direct elicitation of properties (see Section 2.5.3 and in particular Footnote 13).

Corollary 2.7. *A differentiable DSR $s \colon \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ 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 $f' > 0$.*

2.4.3 Characterization in terms of convex functions and subgradients. Existing work on elicitation often uses the terminology of convex functions and their subgradients (see Section 2.5.2). Indeed, our result can be put in these terms as well.

Corollary 2.8. *A DSR $s \colon \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ 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 2.6 and the equivalence of convex functions and integrals over subderivatives, see e.g., Theorem 24.2 and Corollary 24.2.1 of Rockafellar [1970]. □

2.5 Related work

2.5.1 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 2.7.[12]

2.5.2 Chen et al. [2014]. Chen 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 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 2.1. 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 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.

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 2.8 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 2.2 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).

2.5.3 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. Abernethy and Frongillo, 2012, Bellini and Bignozzi, 2015, Gneiting, 2011, Lambert et al., 2008]. Of course, in principle, one could elicit entire distributions and would thereby elicit all properties. But eliciting, say, a single-valued

[12]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:

$$\frac{s'(p,1)}{s'(p,0)} = \frac{p-1}{p} \Leftrightarrow \frac{f'(p) - g'(p)}{g'(p)} = \frac{p-1}{p} \Leftrightarrow g'(p) = pf'(p), \tag{16}$$

which is exactly the relationship between f', g' stated in our Corollary 2.7, 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 $s(p,0)$.

point forecast may be required “for reasons of decision making, market mechanisms, reporting requirements, communications, or tradition, among others” [Gneiting, 2011, Section 1]. Lemma 2.4 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 2.3) 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.[13]

3 ELICITING FROM MULTIPLE EXPERTS

3.1 Setting

Again, we consider a principal who selects from a set of actions A. After she has taken an action, an outcome from Ω is obtained. The principal would like to select the action that maximizes the expectation of the value of some utility function $u\colon \Omega \longrightarrow \mathbb{R}$. This time the principal consults n different experts. Again, she asks for information, then takes the best action given the information submitted and finally rewards the experts based on the submitted information and the outcome obtained.

When it comes to the format and reporting of information, however, switching to the multi-expert setting poses a few additional challenges compared to the single-expert setting. The first (and easier to defuse) issue is that for our definition of propriety we need to adapt our model of how the experts select their report. After all, what report maximizes an expert’s expected score depends, in general, on what the other experts report. In this paper, we go with one of the standard answers to this question from the mechanism design literature: Bayes-Nash equilibrium (or Bayes-Nash incentive compatibility). That is, we will call the scoring rule s_i for expert i proper if given that everyone else reports honestly, i also weakly prefers reporting honestly.[14]

The second challenge is trickier and a heavier burden to carry throughout this section: of what types are the beliefs and reports of the experts? We do not want to simply let each expert’s beliefs be some conditional probability distribution $\Delta(\Omega)^A$ again, because it would be unclear how one would aggregate these beliefs. Again, we make use of a standard solution from the economic literature: the common prior model.

We assume that each expert $i = 1, .., n$ has access to a private piece of information from some set H_i. The experts (and principal) share a common prior $Q \in \Delta(H)$ over $H := \bigtimes_i H_i$. For simplicity, we assume that $Q > 0$ so that no two pieces of evidence are ever inconsistent. We imagine also that

[13]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 2.7. Here, we must have $g'(\mu) = \mu f'(\mu)$. Integrating both sides yields $g(\mu) = C + \int_{\mu_0}^{\mu} tf'(t)dt$ for some constant $C \in \mathbb{R}$. Corollary 2.7 thus becomes

$$s(\mu, y) = f(\mu)y - \int_{\mu_0}^{\mu} tf'(t)dt - C = \int_{\mu_0}^{\mu} f'(t)(y - t)dt + C'y - C, \tag{17}$$

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.

[14]Note that the stronger goal of making honest reporting a dominant strategy (i.e., of making our mechanism dominant-strategies incentive-compatible) is unattainable. Roughly, if the other experts misreport, an expert may be able to improve his payoff by misreporting in a way that improves the quality of the recommendation and prediction resulting from the overall report.

the experts (and principal) share a common conditional distribution $P \in \Delta(\Omega)^{A\times H}$, which for any action $a \in A$ and evidence vector $\mathbf{e} \in H$ specifies a probability distribution $P(\cdot \mid a, \mathbf{e})$ over outcomes given that $\mathbf{e}$ is observed and action a is taken. As a report, each expert – after observing $e_i \in H_i$ – submits $\hat{e}_i \in H_i$ and the principal chooses an action that is best given the reported evidence, i.e., an action from $\mathrm{opt}_P(\hat{\mathbf{e}}) := \arg\max_{a\in A} \mathbb{E}_P [u(O) \mid \hat{\mathbf{e}}, a]$.

It would now be natural to consider as scoring rules functions from $H \times \Omega$ to $\mathbb{R}$, i.e., functions that take as input everyone's report and the outcome obtained after taking the experts' collective recommendation. But the set of allowed scoring rules for player i will then in general depend a lot on the structure of H, Q and P. Instead, we will again work toward generic DSRs, i.e., ones that work without assuming anything about the beliefs that the experts could have.[15]

To make our scoring rules generic, we have to return – partially – to scoring probability distributions. In particular, for the purposes of scoring expert i, we imagine that expert i can recommend/believe in

- any probability distribution $\hat{Q}_i \in \Delta(H_{-i})$ over what evidence everyone else submits,
- any policy $\hat{\alpha}\colon H_{-i} \to A$ mapping the others' evidence onto a recommendation, and
- any probability distribution $\hat{P}_{i,\hat{\alpha}} \in \Delta(\Omega)^{H_{-i}}$ that describes what outcomes to expect for any report of the other experts if expert i's recommendation is implemented.

At the same time we imagine that all experts other than i still observe and submit (honestly) from H_{-i}. For propriety we then require that even with this much broader set of possible beliefs and reports, expert i always prefers honest reporting assuming that everyone else submits their H_{-i} honestly.

Definition 3.1. A *multi-expert DSR* is a collection $(s_i)_{i=1,\dots,n}$ of functions

$$s_i\colon \Delta(H_{-i}) \times \Delta(\Omega)^{H_{-i}} \times \Delta(H) \times \Delta(\Omega)^{H\times A} \times H_{-i} \times \Omega \to \mathbb{R}, \tag{18}$$

each of which maps

- expert i's reported distribution over H_{-i},
- expert i's reported distribution over Ω under following expert i's conditional recommendation,
- the prior Q over H,
- the posterior $P(\cdot \mid \cdot, \cdot)$ which specifies a distribution over outcomes for each evidence vector and action,
- the submitted evidence of the other experts, and
- the observed outcome

onto a score for expert i. We say that s_i is *proper* if for any priors P, Q and true beliefs $Q_i \in \Delta(H_{-i})$, $P_i \in \Delta(\Omega)^{A\times H_{-i}}$, all possible reports $\hat{Q}_i \in \Delta(H_{-i})$, $\hat{P}_{i,\hat{\alpha}} \in \Delta(\Omega)^{H_{-i}}$ with recommended policy $\hat{\alpha}$

$$\xi_{Q_i,P_i}(\hat{\alpha}, \hat{P}_{i,\hat{\alpha}}, \hat{Q}_i) \leq \max_{\alpha^*\in\mathrm{opt}(P_i)} \xi_{Q_i,P_i}(\alpha^*, P_i(\cdot \mid \alpha^*, \cdot), Q_i), \tag{19}$$

where

$$\xi_{Q_i,P_i}(\hat{\alpha}, \hat{P}_{i,\hat{\alpha}}, \hat{Q}_i) = \mathbb{E}_{\mathbf{E}_{-i}\sim Q_i, O\sim P(\cdot\mid\hat{\alpha}(\mathbf{E}_{-i}),\mathbf{E}_{-i})} \left[s_i(\hat{Q}_i, \hat{P}_{i,\hat{\alpha}}, P, Q, \mathbf{E}_{-i}, O) \right] \tag{20}$$

is expert i's expected score of reporting $\hat{\alpha}, \hat{P}_{i,\hat{\alpha}}, \hat{Q}_i$ under the belief Q_i, P_i; and $\mathrm{opt}(P_i)$ is the set of optimal deterministic policies $A^{H_{-i}}$ given the belief P_i. We say that s_i is *right-action proper* if for all $P, Q, P_i, Q_i, \hat{Q}_i, \hat{P}_{i,\hat{\alpha}}$ as above it is

$$\xi_{Q_i,P_i}(\hat{\alpha}, \hat{P}_{i,\hat{\alpha}}, \hat{Q}_i) = \max_{\alpha^*\in\mathrm{opt}(P_i)} \xi_{Q_i,P_i}(\alpha^*, P_i(\cdot \mid \alpha^*), Q_i) \tag{21}$$

[15] Compare the discussion of genericism in footnote 3. Indeed, the non-generic multi-expert setting reduces cleanly to the non-generic single-expert setting described in that footnote. Since we have no characterization for the non-generic single-expert setting, though, this helps us little here.

only when $\hat{\alpha} \in \text{opt}(P_i)$. We call s_i *strictly proper w.r.t. the means* if eq. 21 implies that for all $\mathbf{e}_{-i} \in H_{-i}$

$$\mathbb{E}_{\hat{P}_i}\left[u(O) \mid \mathbf{e}_{-i}\right] = \mathbb{E}_{P_i}\left[u(O) \mid \mathbf{e}_{-i}, \hat{\alpha}(\mathbf{e}_{-i})\right]. \tag{22}$$

There are a few things to point out. First, notice that the score of i depends only on the distribution given the recommended action. Predictions about what would happen if a suboptimal action were to be taken are not scored. And the score cannot depend on which action is recommended (other than through the outcome that the recommendation gives rise to). All of this can be justified in the same way as we have justified it for the single-expert case.

The above definitions are somewhat schizophrenic. The score of expert i *can* depend on the other experts' reports $\mathbf{e}_{-i} \in H_{-i}$ themselves (not only the probabilistic beliefs that the $\mathbf{e}_{-i}$ give rise to). Moreover, DSRs can depend on the prior probability distributions P and Q. This stands in apparent contradiction to the goal of genericism! After all, it allows the scoring rule to know exactly what probability distributions the expert can submit. However, allowing dependences of this sort is necessary for many things we may like generic decision scoring rules to do in practice. For example, we may like to scale i's reward up or down depending on the value of expert i's information to the principal (perhaps given the other experts' information). We could apply some restrictions on how s_i can depend on the priors P, Q; for instance, we may require that s_i only depends on conditionals like $Q(\cdot \mid e_j)$ and not on the (unconditioned) prior. However, our general characterization below needs no such restrictions.

Rather than the definition of a multi-expert DSR itself, the genericism in Definition 3.1 lies primarily in the definitions of propriety. For propriety, we assume that the scoring rule (or the principal using that scoring rule) does not know (despite the information submitted by the other experts and the priors P, Q) what the possible reports H_i of expert i are or what beliefs these give rise to. We require that s_i sets good incentives even if the expert had beliefs that given $P, Q, \mathbf{e}_{-i}$ he cannot have.

Like single-expert DSRs, multi-expert DSRs are a type of direct-revelation mechanism for information elicitation. In general, there are many other complex games one could set up to get the experts to reveal their information. For instance, much of the existing literature has been concerned with "decision markets", in which experts repeatedly report probabilistic estimates or even trade some kind of securities (see Section 3.4). However, by a revelation principle, the class of direct-revelation mechanisms in some sense characterizes the set of all proper mechanisms for eliciting information for decision making.

3.2 Characterizing right-action proper multi-expert scoring rules for decision making

Definition 3.1 (in particular the underlying solution concept of Bayes-Nash equilibrium) makes propriety a property of single-player games – for judging whether s_i is proper, we need only consider the case in which everyone else submits non-strategically. What's more, this single-player game resembles the elicitation situations studied in Section 2: the expert makes a recommendation which the principal implements; the expert further submits a prediction; finally the expert is scored based on prediction and outcome. Unfortunately, we nevertheless cannot directly apply Theorem 2.6 to obtain a characterization of multi-expert DSRs. This is because the corresponding single-expert elicitation situations have some additional structure that we assumed not to exist in our characterization: we know that part of the "outcome" – namely the evidence submitted by the other experts $\mathbf{E}_{-i}$ – is not affected by what action the principal chooses. Contrary to the general case, incentivizing an expert to reveal honest predictions about $\mathbf{E}_{-i}$ along with good recommendations is easy and can be done using regular scoring rules for prediction.

To apply our existing results, consider for any DSR s_i and any honest distribution Q_i over H_{-i} the function

$$\tilde{s}_i(\hat{P}_{i,\hat{\alpha}}, Q, P, \mathbf{e}_{-i}, \omega) = s_i(\hat{P}_{i,\hat{\alpha}}, Q_i, Q, P, \mathbf{e}_{-i}, \omega). \tag{23}$$

Here, $\hat{P}_{i,\hat{\alpha}}$ is any (reported) probability distribution over Ω. Intuitively, $\tilde{s}_i$ is the scoring function resulting from s_i if Q_i was somehow already common knowledge or if the principal had some way of forcing expert i to report Q_i honestly. It is left to incentivize i to give honest recommendations and predictions conditional on each $\mathbf{e}_{-i} \in H_{-i}$. We can view this as there being a separate decision for each $\mathbf{e}_{-i} \in H_{-i}$. The agent gives a recommendation and a prediction for each of those decisions, but each of them is played only with some probability (which for now we assume to be known). It can be shown (e.g., by transferring the arguments of Section 2.2) that for $\tilde{s}_i$ to be proper, it can (apart from degenerate cases) only depend on the reported means $\mathbb{E}_{O\sim\hat{P}_{i,\hat{\alpha}}}[u(O) \mid \mathbf{e}_{-i}]$ for each $\mathbf{e}_{-i}$ and the obtained utility. Write $\hat{\boldsymbol{\mu}}_i \in \mathbb{R}^{H_{-i}}$ for the collection of such reported means and $\hat{\mu}_{\mathbf{e}_{-i}}$ for the mean for any particular $\mathbf{e}_{-i} \in H_{-i}$. Then we can write $\tilde{s}_i$ as $\tilde{s}_i(\hat{\boldsymbol{\mu}}_i, Q, P, \mathbf{e}_{-i}, y)$.

For simplicity we will henceforth only consider functions $\tilde{s}_i$ where $\tilde{s}_i(\hat{\boldsymbol{\mu}}_i, Q, P, \mathbf{e}_{-i}, y)$ depends only on $\hat{\mu}_{\mathbf{e}_{-i}}$ and none of the other parts of $\hat{\boldsymbol{\mu}}_i$.[16] Theorem 2.6 then implies that for $\tilde{s}_i$ to be proper it has to be

$$\tilde{s}_i(\hat{\boldsymbol{\mu}}_i, Q, P, \mathbf{e}_{-i}, y) = f_{Q,P,\mathbf{e}_{-i}}(\hat{\mu}_{\mathbf{e}_{-i}})(y - \hat{\mu}_{\mathbf{e}_{-i}}) + \int_0^{\hat{\mu}_{\mathbf{e}_{-i}}} f_{Q,P,\mathbf{e}_{-i}}(x)dx + C_{Q,P,\mathbf{e}_{-i}}, \tag{27}$$

where $f_{Q,P,\mathbf{e}_{-i}}$ is non-negative and non-decreasing for every $\mathbf{e}_{-i}, Q, P$. Further, $\tilde{s}_i$ is right-action proper if and only if the $f_{Q,P,\mathbf{e}_{-i}}$ are strictly positive and strictly proper w.r.t. the means if and only if the $f_{Q,P,\mathbf{e}_{-i}}$ are strictly increasing.

For any honestly reported and therefore also any dishonestly reported $\hat{Q}_i$ over H_{-i} we can therefore write s_i as

$$s_i(\hat{\boldsymbol{\mu}}_i, \hat{Q}_i, P, Q, \mathbf{e}_{-i}, y) = f_{\hat{Q}_i,Q,P,\mathbf{e}_{-i}}(\hat{\mu}_{\mathbf{e}_{-i}})(y - \hat{\mu}_{\mathbf{e}_{-i}}) + \int_0^{\hat{\mu}_{\mathbf{e}_{-i}}} f_{\hat{Q}_i,Q,P,\mathbf{e}_{-i}}(x)dx + C_{\hat{Q}_i,Q,P,\mathbf{e}_{-i}}, \tag{28}$$

with the usual constraints on $f_{\hat{Q}_i,Q,P,\mathbf{e}_{-i}}$. It is only left to add to eq. 28 the requirement that honest reporting of $\hat{Q}_i$ is optimal. First notice that regardless of what $\hat{Q}_i$ the expert reports, he is incentivized to report the means $\boldsymbol{\mu}_i$ honestly (potentially up to some interval) s.t. when deciding which $\hat{Q}_i$ to submit, he maximizes

$$\mathbb{E}_{\mathbf{E}_{-i}\sim Q_i}\left[s_i(\boldsymbol{\mu}, \hat{Q}_i, P, Q, \mathbf{e}_{-i}, y)\right] = \sum_{\mathbf{e}_{-i}\in H_{-i}} Q_i(\mathbf{e}_{-i}) \left(\int_0^{\mu_{\mathbf{e}_{-i}}} f_{\hat{Q}_i,Q,P,\mathbf{e}_{-i}}(x)dx + C_{\hat{Q}_i,Q,P,\mathbf{e}_{-i}} \right), \tag{29}$$

[16]More generally, we could have $\tilde{s}_i$ depend on the reported expected utility across multiple $\mathbf{e}_{-i}$. For instance, we could have a rule like

$$\tilde{s}_i(\hat{\boldsymbol{\mu}}_i, y) = f(\mu_i)(y - \mu_i) - g(\mu_i), \tag{24}$$

where

$$\mu_i := \sum_{\mathbf{e}_{-i}\in H_{-i}} Q(\mathbf{e}_{-i} \mid e_i)\hat{\mu}_{\mathbf{e}_{-i}} \tag{25}$$

is the overall mean implied by $\hat{\boldsymbol{\mu}}_i$. Similarly, we could pick a partition of H_{-i} into subsets S across which we use the reported mean within S

$$\mu_S := \frac{1}{Q(\mathbf{E}_{-i} \in S \mid e_i)} \sum_{\mathbf{e}_{-i}\in S} Q(\mathbf{e}_{-i} \mid e_i)\hat{\mu}_{\mathbf{e}_{-i}} \tag{26}$$

for scoring the expert when $\mathbf{e}_{-i} \in S$. This way, we can incentivize $\hat{\mu}_S$ to be the mean for the across $\mathbf{e}_{-i} \in S$ without incentivizing the individual $\hat{\mu}_{\mathbf{e}_{-i}}$ to be accurate. Since the principal knows $\mathbf{e}_{-i}$, he will generally most interested in the mean given specific values of $\mathbf{e}_{-i}$, however.

where Q_i is expert i's honest probability distribution over H_{-i}. Hence, s_i is (strictly) proper w.r.t. Q_i if and only if for any $\boldsymbol{\mu}_i, Q, P$

$$\int_0^{\mu_{\mathbf{e}_{-i}}} f_{\hat{Q}_i, Q, P, \mathbf{e}_{-i}}(x)dx + C_{\hat{Q}_i, Q, P, \mathbf{e}_{-i}} \tag{30}$$

as a function of $\hat{Q}_i, \mathbf{e}_{-i}$ is a (strictly) proper scoring rule for prediction [as defined and characterized by Gneiting and Raftery, 2007].

3.3 A notable impossibility

As noted earlier, we can use a typical revelation principle argument of mechanism design to draw conclusions from these results about what kind of characteristics any generic proper mechanism for eliciting information for decision making from multiple experts must have. For instance, as in the single-expert case, it shows that we cannot incentivize experts to – along with an honest recommendation – reveal anything other than the expected utility of taking the recommended action. Further, no expert may profit from the failure of the principal's project. If we imagine the principal to be a firm maximizing its value, then no expert can be allowed to short-sell shares in the firm. In this section, we consider another, multi-expert-specific, desirable property that as a consequence of our results we cannot obtain.

We might like to reward experts in proportion to how much their report updates the principal's beliefs. This is one of many desirable properties of prediction markets: experts (or traders) are rewarded based on how far they can move the market probabilities toward the truth. For example, an expert who at any point simply agrees with the market probabilities (because he has no relevant private information) can earn no money (in expectation). An expert who updates the market probability for an event from, say, 0.5 to 0.1 receives a high score in expectation (assuming 0.1 represents his true beliefs over the outcome of the random event). Rewarding experts for their impact on the market probability has many advantages. For instance, it sets a natural incentive to acquire relevant information. Therefore, we might want an elicitation mechanism for decision making – perhaps a kind of decision market (see Section 3.4) – that similarly rewards experts for submitting evidence that yields large (justified) changes in the principal's beliefs.

However, from our results it follows immediately that a number of types of changes cannot be rewarded at all. An expert's score cannot depend on how much the trader's report moves the distributions for suboptimal actions. Experts also cannot be rewarded for changing what action is recommended. Generally, if two pieces of information $e_i^1, e_i^2 \in H_i$ have the same implications for the expected utility given the best action, the expert receives the same score in expectation from honestly reporting e_i^1 and honestly reporting e_i^2. This is the case even if e_i^1 affects what the best action is and implies wild changes to the distributions of all actions while e_i^2 does not change the principal's beliefs at all. In particular this implies that a generic strictly proper DSR gives positive expected rewards even to experts i whose private evidence E_i turns out to be of no value to the principal.

How can the principal make sure that despite these impossibilities, experts with more useful information receive higher scores? The only way out, it seems, is to reward experts based on the *ex ante* value of their information. That is, pay expert i (in shares or constant reward) in proportion to how much the principal would be willing to pay to learn E_i. One could also use the willingness to pay given that one already knows or will know $\mathbf{E}_{-i}$. In the extreme case, one could even give a constant score of 0 to experts whose value of information is zero. (The mechanism would then not quite satisfy our generic notion of strict propriety anymore.) This way, obtaining E_i is incentivized to the extent that E_i is useful to the expert.

It is worth noting that we can get around most of the impossibility results for eliciting decision-relevant information by removing some of our underlying assumptions or weakening our goals. For instance, they disappear if we allow randomization over which action the principal takes [Chen et al., 2011; Chen et al., 2014, Section 4] or if we look for non-generic mechanisms for specific decision problems in which the principal has a detailed model of experts' information structure.

3.4 Related work

As far as we are aware, most work on eliciting decision-relevant information from multiple agents has focused on designing prediction-market-like mechanisms or "decision markets" (as opposed to considering the class of direct-revelation mechanisms discussed in this paper) [e.g. Berg and Rietz, 2003, Hanson, 1999, 2002, 2006, 2013, Section IV]. Othman and Sandholm [2010, Section 3] are the first to point out incentive problems with this model. Our impossibility results can be seen as an extension of their result (though we have limited attention to mechanisms which guarantee full information aggregation, which may not a be a primary goal for the design of decision markets). Inspired by Othman and Sandholm's proof that decision markets sometimes set poor incentives, Teschner et al. [2017] conduct an empirical study in which human subjects took the roles of the experts (or "traders") to show that strategic reporting may not be a problem in decision markets in practice. Chen et al. [2011; 2014] show that by randomizing over all actions (potentially with a strong bias toward the optimal action) decision markets can, in some sense, be made to be analogous to prediction markets.

4 CONCLUSION

We have characterized mechanisms for eliciting decision-relevant information from both individual and multiple experts, thus – we hope – laying the foundations for theories analogous to those of proper scoring rules and prediction markets. We have shown that to strictly incentivize optimal recommendations the principal can only set very limited incentives on the experts' predictions: only the report on the expected utility of taking the recommended action can be strictly incentivized. Moreover, the right-action proper scoring rules allow for a simple, intuitive interpretation. In the single-expert case, the principal offers the expert at different prices shares in the principal's project. From the price above which the expert rejects the offers, it can be inferred what the expert believes the expected value of a share and therefore the expected utility of the principal's project to be. Further, since the expert owns a share in the project, he wants to give the best-possible recommendation. A similar result holds for a setting with multiple experts. By the usual revelation principle, our characterization tells us how any right-action proper mechanism for eliciting decision-relevant information must look like. Many desirable properties cannot be obtained. For instance, short-selling of the principal's project – as is possible in prediction markets – cannot be allowed in a right-action proper mechanisms. The most natural objective of further research therefore is the search for a natural, practical equivalent of a (Arrow-Debreu securities or market scoring rule) prediction market in the decision making case.

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 Debmalya Panigrahi) at Duke for helpful discussions.

REFERENCES

Jacob D. Abernethy and Rafael M. Frongillo. 2012. A Characterization of Scoring Rules for Linear Properties. In *Proceedings of the 25th Annual Conference on Learning Theory (Proceedings of Machine Learning Research)*, Shie Mannor, Nathan Srebro, and Robert C. Williamson (Eds.), Vol. 23. PMLR, Edinburgh, Scotland, 27.1–27.13. http://proceedings.mlr.press/v23/abernethy12.html

Shipra Agrawal, Erick Delage, Mark Peters, Zizhuo Wang, and Yinyu Ye. 2009. A unified framework for dynamic pari-mutuel information market design. In *EC '09 Proceedings of the 10th ACM conference on Electronic commerce*. 255–264.

Fabio Bellini and Valeria Bignozzi. 2015. On elicitable risk measures. *Quantitative Finance* 15, 5 (2015), 725–733. https://doi.org/10.1080/14697688.2014.946955

Joyce E. Berg and Thomas A. Rietz. 2003. Prediction Markets as Decision Support Systems. *Information Systems Frontiers* 5, 1 (2003), 79–93. https://doi.org/10.1023/A:1022002107255

Craig Boutilier. 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*, Conitzer, Winikoff, Padgham, and van der Hoek (Eds.).

Glenn W. Brier. 1950. Verification of Forecasts Expressed in Terms of Probability. *Monthly Weather Review* 78, 1 (1 1950).

Gabriel Carroll. 2019. Robust incentives for information acquisition. 181 (5 2019), 382–420. https://doi.org/10.1016/j.jet.2019.03.001

Yiling Chen 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. https://dash.harvard.edu/bitstream/handle/1/5128480/chen_aamas11_0.pdf?sequence=1

Yiling Chen, Ian A. Kash, Mike Ruberry, and Victor Shnayder. 2011. Decision markets with Good Incentives. In *Lecture Notes in Computer Science 7090*. 72–83. https://dash.harvard.edu/bitstream/handle/1/11693950/Chen_DecisionMarkets.pdf?sequence=1

Yiling Chen, Ian A. Kash, Mike Ruberry, and Victor Shnayder. 2014. Eliciting Predictions and Recommendations for Decision Making. In *ACM Transactions on Economics and Computation*. Vol. 2. Chapter 6. http://yiling.seas.harvard.edu/wp-content/uploads/a6-chen.pdf

Yiling Chen and David M. Pennock. 2007. A utility framework for bounded-loss market makers. In *UAI'07 Proceedings of the Twenty-Third Conference on Uncertainty in Artificial Intelligence*. 49–56.

Yiling Chen and Jennifer Wortman Vaughan. 2010. A new understanding of prediction markets via no-regret learning. In *EC' 10 Proceedings of the 11th ACM conference on Electronic commerce*. 189–198.

Vincent P. Crawford and Joel Sobel. 1982. Strategic Information Transmission. *Econometrica* 50, 6 (11 1982).

Rafael M. Frongillo and Ian A. Kash. 2014. General Truthfulness Characterizations Via Convex Analysis. https://www.cs.colorado.edu/~raf/media/papers/wine-affine.pdf Short version published in WINE 2014.

Tilmann Gneiting. 2011. Making and Evaluating Point Forecasts. *J. Amer. Statist. Assoc.* 106, 494 (2011), 746–762. https://doi.org/10.1198/jasa.2011.r10138

Tilmann Gneiting and Adrian E. Raftery. 2007. Strictly Proper Scoring Rules, Prediction, and Estimation. *J. Amer. Statist. Assoc.* 102, 477 (3 2007). https://doi.org/10.1198/016214506000001437

I. J. Good. 1952. Rational Decisions. *Journal of the Royal Statistical Society. Series B (Methodological)* 14 (1952), 107–114. Issue 1. https://www.jstor.org/stable/2984087

Robin Hanson. 1999. Decision Markets. *IEEE Intelligent Systems* 14, 3 (1999), 16–19. http://mason.gmu.edu/~rhanson/decisionmarkets.pdf

Robin Hanson. 2002. Decision Markets. In *Entrepreneurial Economics – Bright Ideas from the Dismal Science*, Alexander Tabarrok (Ed.). Oxford University Press, Chapter 5, 79–85.

Robin Hanson. 2003. Combinatorial Information Market Design. *Information Systems Frontiers* 5, 1 (2003), 107–119.

Robin Hanson. 2006. Decision Markets for Policy Advice. In *Promoting the General Welfare: New Perspectives on Government Performance*, Alan S. Gerber and Eric M. Patashnik (Eds.). Brookings Institution Press, 151–173.

Robin Hanson. 2007. Logaritmic Market Scoring Rules for Modular Combinatorial Information Aggregation. *The Journal of Prediction Markets* 1 (2007), 3–15.

Robin Hanson. 2013. Shall We Vote on Values, But Bet on Beliefs? *Journal of Political Philosophy* 21, 2 (2013). https://doi.org/10.1111/jopp.12008

David J. Johnstone, Victor Richmond R. Jose, and Robert L. Winkler. 2011. Tailored Scoring Rules for Probabilities. *Decision Analysis* 8, 4 (2011), 256–268. https://doi.org/10.1287/deca.1110.0216

Eric Key. 1994. Disks, Shells, and Integrals of Inverse Functions. *The College Mathematics Journal* 25, 2 (3 1994), 136–138.

Nicolas Lambert, David M. Pennock, and Yoav Shoham. 2008. Eliciting Properties of Probability Distributions. In *ACM Conference on Electronic Commerce*. Extended Abstract.

John McCarthy. 1956. Measures of the Value of Information. *Proceedings of the National Academy of Sciences of the United States of America* 42 (9 1956), 654–655. Issue 9. https://doi.org/10.1073/pnas.42.9.654

Abraham Othman 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.* 625–632.

David M. Pennock and Rahul Sami. 2007. Computational Aspects of Prediction Markets. In *Algorithmic Game Theory*, Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani (Eds.). Cambridge University Press, Chapter 26, 651–675.

R. Tyrrell Rockafellar. 1970. *Convex Analysis.* Princeton University Press, Princeton, NJ.

Walter Rudin. 1976. *Principles of Mathematical Analysis* (3 ed.). McGraw-Hill.

Leonard J. Savage. 1971. Elicitation of Personal Probabilities and Expectations. *J. Amer. Statist. Assoc.* 66 (12 1971), 783–801. Issue 336.

Mark J. Schervish. 1989. A General Method for Comparing Probability Assessors. *The Annals of Statistics* 17, 4 (1989), 1856–1879. https://projecteuclid.org/euclid.aos/1176347398

Peng Shi, Vincent Conitzer, and Mingyu Guo. 2009. Prediction Mechanisms That Do Not Incentivize Undesirable Actions. In *Workshop on Internet and Network Economics (WINE) 2009*, S. Leonardi (Ed.). Lecture Notes in Computer Science, Vol. 5929. Springer.

Florian Teschner, David Rothschild, and Henner Gimpel. 2017. Manipulation in Conditional Decision Markets. *Group Decision and Negotiation* 26, 5 (9 2017), 953–971. https://doi.org/10.1007/s10726-017-9531-0

Luis Zermeño. 2011. A Principal-Expert Model and the Value of Menus. (11 2011). http://economics.mit.edu/files/7299

Luis Zermeño. 2012. The Role of Authority in a General Principal-Expert Model. (3 2012). http://economics.mit.edu/files/7627

A OVERVIEW OF NOTATION

We here introduce some notation that is used throughout this paper. Since much of this is self-explanatory or sufficiently widely used, the reader may skip this section and refer to it later.

For any set Ω, define

$$\Delta(\Omega) := \{P \colon \Omega \to [0,1] \mid \sum_{\omega\in\Omega} P(\omega) = 1\} \tag{31}$$

to be the set of probability distributions over Ω. For probabilities $p_1, ..., p_n \in [0,1]$ with $\sum_{i=1}^n p_i = 1$, we write $p_1 * \omega_1 + ... + p_n * \omega_n$ for the random variable that assumes value ω_i with probability p_i for $i = 1, ..., n$.

For any probability distribution $P \in \Delta(\Omega)$, define the support of P

$$\text{supp}(P) := \{\omega \in \Omega \mid P(\omega) > 0\} \tag{32}$$

to be the set of outcomes to which P assigns positive probability. For any real-valued function $u \colon \Omega \to \mathbb{R}$, we use

$$\mathbb{E}_{O\sim P}\left[u(O)\right] := \sum_{\omega\in\Omega} P(\omega)u(\omega) \tag{33}$$

to denote the expected value of $u(O)$ where O is the random variable distributed according to P. We generally use roman capitals to denote random variables. When it is clear that O is distributed according to P, we may drop the subscript $O \sim P$ for brevity. When it is clear that O is the random variable that the expectation is over, but there is ambiguity as to how O is distributed, we give only P in the subscript.

Since we are dealing with decision problems, we are less interested in the case where the expert's belief P is single probability distribution. Rather, we often use *conditional probability distributions* $P \in \Delta(\Omega)^A$, where A is some set with generally at least two elements and for each $a \in A$, the entry of P corresponding to a is a probability distribution over Ω. We variably denote these conditional probability distributions as $P(\cdot \mid a)$ or $P_a(\cdot)$. Note that in this notation, (unless $|A| = 1$ perhaps) P itself (before conditioning on any a) is not a probability distribution. In particular, we do not require there to be any underlying joint probability distribution over A and Ω; an element of A is chosen, not randomly drawn. We will also use $\mathbb{E}_{O\sim P}\left[u(O) \mid a\right] := \mathbb{E}_{O\sim P(\cdot\mid a)}\left[u(O)\right]$ for conditional expectations.

In the second part of this paper (Section 3), we furthermore use the following standard game-theoretic notation. If $X = \bigtimes_{i=1}^n X_i$ (for sets $X_1, ..., X_n$), then $X_{-j} := \bigtimes_{i\in\{1,...,j-1,j+1,...,n\}} X_i$. We use bold font to denote vectors $\mathbf{x} \in X$, and extend the notation to $\mathbf{x}_{-j} := (x_1, ..., x_{j-1}, x_{j+1}, ..., x_n)$. We will also use $\mathbf{x}_{-j}$ to denote elements of X_{-j} directly.

B AN ELEMENTARY PROOF OF THE CONTINUITY OF PROPER DSRS

We here give an elementary proof of the continuity claim used in the proof of Lemma 2.3.

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 $\mathbb{E}\left[s(Q_p, R_p)\right]$ 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 $\epsilon > 0$

$$\mathbb{E}\left[s(Q_{\tilde{p}}, R_{\tilde{p}})\right] > \mathbb{E}\left[s(Q_{\tilde{p}-\epsilon}, R_{\tilde{p}-\epsilon})\right] + \delta. \tag{34}$$

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

$$\mathbb{E}\left[s(Q_{\tilde{p}}, R_{\tilde{p}})\right] < \mathbb{E}\left[s(Q_{\tilde{p}+\epsilon}, R_{\tilde{p}+\epsilon})\right] - \delta. \tag{35}$$

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.

$$\mathbb{E}\left[s(Q_{\tilde{p}}, R_{\tilde{p}-\epsilon})\right] = (\tilde{p}-\epsilon)s(Q_{\tilde{p}}, \omega_H) + (1-(\tilde{p}-\epsilon))s(Q_{\tilde{p}}, \omega_L) \quad (36)$$

$$\overset{\epsilon\to 0}{\to} \tilde{p}s(Q_{\tilde{p}}, \omega_H) + (1-\tilde{p})s(Q_{\tilde{p}}, \omega_L) \quad (37)$$

$$= \mathbb{E}\left[s(Q_{\tilde{p}}, R_{\tilde{p}})\right]. \quad (38)$$

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

$$\mathbb{E}\left[s(Q_{\tilde{p}}, R_{\tilde{p}-\epsilon})\right] > \mathbb{E}\left[s(Q_{\tilde{p}}, R_{\tilde{p}})\right] - \delta/2 > \mathbb{E}\left[s(Q_{\tilde{p}-\epsilon}, R_{\tilde{p}-\epsilon})\right] + \delta/2. \quad (39)$$

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 i 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 $R_{\tilde{p}}$. Then for small enough ϵ, the expert prefers submitting $Q_{\tilde{p}+\epsilon}$ over submitting $Q_{\tilde{p}}$, because of the following.

$$\mathbb{E}\left[s(Q_{\tilde{p}}, R_{\tilde{p}})\right] = \tilde{p}s(Q_{\tilde{p}+\epsilon}, \omega_H) + (1-\tilde{p})s(Q_{\tilde{p}+\epsilon}, \omega_L) \quad (40)$$

$$\overset{\epsilon\to 0}{\leftarrow} \tilde{p}s(Q_{\tilde{p}+\epsilon}, \omega_H) + (1-\tilde{p})s(Q_{\tilde{p}+\epsilon}, \omega_L) \quad (41)$$

$$+\epsilon(s(Q_{\tilde{p}+\epsilon}, \omega_H) - s(Q_{\tilde{p}+\epsilon}, \omega_L)) \quad (42)$$

$$= \mathbb{E}\left[s(Q_{\tilde{p}+\epsilon}, R_{\tilde{p}+\epsilon})\right]. \quad (43)$$

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

$$\mathbb{E}\left[s(Q_{\tilde{p}+\epsilon}, R_{\tilde{p}})\right] > \mathbb{E}\left[s(Q_{\tilde{p}+\epsilon}, R_{\tilde{p}+\epsilon})\right] - \delta/2 > \mathbb{E}\left[s(Q_{\tilde{p}}, R_{\tilde{p}})\right] + \delta/2. \quad (44)$$

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

C PROOF OF THEOREM 2.6

We here prove our main theorem. For convenience, we first restate the theorem here.

Theorem 2.6. *A DSR $s\colon \mathbb{R}\times\mathbb{R}\to\mathbb{R}$ is proper iff*

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

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.

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 $\mu + d$:

$$\begin{aligned}
\mathbb{E}[s(\mu+d,U)] &= s(\mu+d,\mu) && (45)\\
&= -df(\mu+d) + \int_0^{\mu+d} f(x)dx + C && (46)\\
&= s(\mu,\mu) - df(\mu+d) + \int_\mu^{\mu+d} f(x)dx && (47)\\
&\underset{f\text{ non-decr.}}{\leq} s(\mu,\mu) - df(\mu+d) + df(\mu+d) && (48)\\
&= s(\mu,\mu) && (49)\\
&= \mathbb{E}[s(\mu,U)] && (50)
\end{aligned}$$

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 $\mu \coloneqq \mathbb{E}[U_a]$ 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:

$$\begin{aligned}
\mathbb{E}[s(\mu+d,U_{a^*})] &= s(\mu+d,\mu+d) && (51)\\
&= \int_0^{\mu+d} f(x)dx + C && (52)\\
&= s(\mu,\mu) + \int_\mu^{\mu+d} f(x)dx && (53)\\
&\underset{f\text{ is non-negative}}{\geq} s(\mu,\mu) && (54)\\
&= \mathbb{E}[s(\mu,U_a)]\,. && (55)
\end{aligned}$$

If 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

$$s(\hat{\mu}, y) = f(\hat{\mu})y - g(\hat{\mu}) \tag{56}$$

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:

$$\begin{aligned}
p\mathbb{E}[s(\hat{\mu},X)] + (1-p)\mathbb{E}[s(\hat{\mu},x')] &= \mathbb{E}[s(\hat{\mu},Y)] && (57)\\
&\underset{\text{Lemma 2.3}}{=} \mathbb{E}[s(\hat{\mu},Y')] && (58)\\
&= ps(\hat{\mu},x) + (1-p)\mathbb{E}[s(\hat{\mu},x')]\,. && (59)
\end{aligned}$$

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. 56. 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

$$g(\mu) = f(\mu)\mu - \int_0^\mu f(x)dx - C \tag{60}$$

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 $d > 0$

$$s(\mu + d, \mu) \leq s(\mu, \mu), \tag{61}$$

which – using eq. 56 – we can rewrite as

$$g(\mu + d) - g(\mu) \geq \mu \cdot (f(\mu + d) - f(\mu)). \tag{62}$$

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

$$s(\mu, \mu + d) \leq s(\mu + d, \mu + d), \tag{63}$$

which we can rewrite as

$$g(\mu + d) - g(\mu) \leq (\mu + d) \cdot (f(\mu + d) - f(\mu)). \tag{64}$$

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 62 and 64, it follows that for all positive d

$$\mu \cdot (f(\mu + d) - f(\mu)) \leq (\mu + d) \cdot (f(\mu + d) - f(\mu)), \tag{65}$$

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:

$$g(\hat{\mu}) = g(0) + \sum_{i=1}^{n} g\left(\frac{i\hat{\mu}}{n}\right) - g\left(\frac{(i-1)\hat{\mu}}{n}\right). \tag{66}$$

Since relative to any f, g can only be unique up to a constant, we will write C instead of $g(0)$. From equations 62 and 64, it follows that

$$\begin{aligned} & \sum_{i=1}^{n} \frac{(i-1)\hat{\mu}}{n}\left(f\left(\frac{i\hat{\mu}}{n}\right) - f\left(\frac{(i-1)\hat{\mu}}{n}\right)\right) && (67)\\ \leq \quad & g(\hat{\mu}) - C && (68)\\ \leq \quad & \sum_{i=1}^{n} \frac{i\hat{\mu}}{n}\left(f\left(\frac{i\hat{\mu}}{n}\right) - f\left(\frac{(i-1)\hat{\mu}}{n}\right)\right) && (69) \end{aligned}$$

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.[17] It

[17] In fact, we could immediately interpret them as Riemann sums of the function f^{-1} for the partition $(f\left(\frac{i\hat{\mu}}{n}\right))_{i=1,\ldots,n}$. This works out but leads to a number of technical issues that are cumbersome to deal with: if f is discontinuous, then $(f\left(\frac{i\hat{\mu}}{n}\right))_{i=1,\ldots,n}$ 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].

is

$$\begin{aligned}
& \sum_{i=1}^{n} \frac{i\hat{\mu}}{n} \left(f\left(\frac{i\hat{\mu}}{n}\right) - f\left(\frac{(i-1)\hat{\mu}}{n}\right) \right) && (70)\\
= \; & \sum_{i=1}^{n} \frac{i\hat{\mu}}{n} f\left(\frac{i\hat{\mu}}{n}\right) - \frac{(i-1)\hat{\mu}}{n} f\left(\frac{(i-1)\hat{\mu}}{n}\right) - \sum_{i=1}^{n} \frac{\hat{\mu}}{n} f\left(\frac{(i-1)\hat{\mu}}{n}\right) && (71)\\
= \; & \hat{\mu} f(\hat{\mu}) - \sum_{i=1}^{n} \frac{\hat{\mu}}{n} f\left(\frac{(i-1)\hat{\mu}}{n}\right). && (72)
\end{aligned}$$

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

$$\begin{aligned}
& \sum_{i=1}^{n} \frac{(i-1)\hat{\mu}}{n} \left(f\left(\frac{i\hat{\mu}}{n}\right) - f\left(\frac{(i-1)\hat{\mu}}{n}\right) \right) && (73)\\
= \; & \sum_{i=1}^{n} \frac{i\hat{\mu}}{n} f\left(\frac{i\hat{\mu}}{n}\right) - \frac{(i-1)\hat{\mu}}{n} f\left(\frac{(i-1)\hat{\mu}}{n}\right) - \sum_{i=1}^{n} \frac{\hat{\mu}}{n} f\left(\frac{i\hat{\mu}}{n}\right) && (74)\\
= \; & \hat{\mu} f(\hat{\mu}) - \sum_{i=1}^{n} \frac{\hat{\mu}}{n} f\left(\frac{i\hat{\mu}}{n}\right). && (75)
\end{aligned}$$

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 integral:

$$\sum_{i=1}^{n} \frac{\hat{\mu}}{n} f\left(\frac{i\hat{\mu}}{n}\right) \underset{n\to\infty}{\rightarrow} \int_0^{\hat{\mu}} f(x)dx \underset{n\to\infty}{\leftarrow} \sum_{i=1}^{n} \frac{\hat{\mu}}{n} f\left(\frac{(i-1)\hat{\mu}}{n}\right). \tag{76}$$

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

$$g(\hat{\mu}) = C + \hat{\mu} f(\hat{\mu}) - \int_0^{\hat{\mu}} f(x)dx. \tag{77}$$

From this, eq. 12 follows as claimed. □