A

Eliciting Predictions for Discrete Decision Making

YILING CHEN, Harvard University
IAN A. KASH, Microsoft Research
MIKE RUBERRY, Harvard University
VICTOR SHNAYDER, Harvard University

We consider a decision maker who can select one of a finite set of possible actions, each of which will result in one of a finite set of possible outcomes. If the decision maker has preferences over these outcomes, it will naturally prefer taking some actions to others, but may be uncertain of the result of each action. In this paper we describe how a decision maker can elicit expert predictions about the outcome of each action, allowing the decision maker to make an informed decision. We show that *strictly proper* decision making, where experts have an incentive to accurately reveal their beliefs about the outcome of each action, allows the decision maker to take a preferred action with probability arbitrarily close—but not equal—to one; with positive probability, the decision maker must take an action at random. Requiring a decision maker to sometimes act randomly is clearly undesirable, so we also describe an alternative where a single expert directly reveals a preferred decision instead of predicting each action's outcome.

Categories and Subject Descriptors: F.0 [**Theory of Computation**]: General; J.4 [**Computer Applications**]: Social and Behavioral Sciences

General Terms: Economics, Theory

Additional Key Words and Phrases: Market design, decision markets, prediction markets, decision making, information elicitation, scoring rules

ACM Reference Format:
Chen, Y., Kash, I., Ruberry, M., and Shnayder, V. 2012. Eliciting Predictions for Discrete Decision Making. ACM V, N, Article A (January YYYY), 23 pages.
DOI = 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

1. INTRODUCTION

When making a decision we are often interested in achieving a specific outcome, but instead of picking outcomes we are limited to picking actions. Worse, we may not know the result of each action: will donating to a politician's campaign improve their odds of winning? Will our company make a profit if we open a new store in Springfield?

To make an informed decision, a decision maker needs to understand the mapping from its actions to the outcomes of interest. Hanson [1999] proposed using experts to predict this relationship. These predictions, which provide a conditional probability distribution over outcomes for each action, can be represented by an *action-outcome* matrix like the one

Parts of this research initially appeared in Chen and Kash [2011] and Chen et al. [2011].
This work is partially supported by NSF grant CCF-0953516. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors alone. The authors are grateful to Craig Boutilier for helpful discussions about this work.
Author's addresses: Y. Chen, M. Ruberry, and V. Shnayder, School of Engineering and Applied Sciences, Harvard University, {yiling, mruberry, shnayder}@seas.harvard.edu; I. A. Kash, Microsoft Research, Cambridge, UK, iankash@microsoft.com.

DOI 10.1145/0000000.0000000 http://doi.acm.org/10.1145/0000000.0000000

in Figure 1. The prediction in Figure 1 contains all the information relevant to making a

		Outcomes $\mathcal{O}$	
		Profit	Loss
Actions $\mathcal{A}$	Springfield	$\frac{2}{3}$	$\frac{1}{3}$
	Greenville	$\frac{2}{5}$	$\frac{3}{5}$

Fig. 1. A prediction of how likely each of two possible actions, building a new store in Springfield or Greenville, would result in a profit or a loss. Entries in the matrix represent the conditional likelihood of their column's outcome given their row's action; the upper-left value of the matrix, for example, indicates a $\frac{2}{3}$ chance the store will make a profit in Springfield. We will assume outcomes are mutually exclusive and exhaustive of all possible futures, so each row is a probability distribution.

decision. And if the decision maker believes the prediction is accurate it will prefer opening a store in Springfield to Greenville.

Hanson [1999] suggested soliciting expert predictions of action-outcome matrices using a *prediction market*-like structure, which he termed a *decision market*. Prediction markets are an effective method of aggregating expert opinion [Wolfers and Zitzewitz 2004; Berg et al. 2001; Chen and Pennock 2010], but are typically used only to predict the likelihood of events independent of the market. A prediction market can, for example, be used to predict how likely a horse is to win a race, or whether it will rain on a particular day[1]. A *strictly proper* prediction market provides an incentive for experts to accurately reveal how likely they think these events are, and (subject to a technical caveat) eventually leads to a consensus prediction that (in theory) reflects the pooled information of every expert [Ostrovsky 2009; Chen et al. 2012].

Unlike a typical prediction market where no decision is made, when we choose a decision based on a market's predictions the realized outcome directly depends on the predictions. This dependency means that the techniques that make prediction markets strictly proper break down for decision markets. Figure 2 is an example of the failure of straightforwardly adapting a strictly proper prediction market to decision making.

Since strict properness is usually considered a critical property of a well-designed prediction market, we describe how to extend it to decision making, providing a characterization of *decision strict properness* analogous to characterizations of strict properness [Gneiting and Raftery 2007]. We show that strictly proper decision making allows the decision maker to take a preferred action with probability arbitrarily close—but not equal—to one; the rest of the time the decision maker must take an action at random. This is an undesirable and likely non-credible restriction in practice, so we also present a new mechanism where a single expert directly reveals a preferred decision instead of describing each action's outcome.

2. BACKGROUND AND RELATED WORK

There has been extensive prior work on scoring rules and prediction markets, and we will, by necessity, only describe a small part. While prediction markets can be operated using continuous double auctions [Forsythe et al. 1992; Berg and Rietz 2003], automated market makers [Othman and Sandholm 2010a; Othman et al. 2010], and other wagering mechanisms [Plott et al. 1997; Pennock 2004; Mangold et al. 2005], we are interested in prediction markets that use scoring rules. It is known that such markets can be equivalently implemented as automated market makers [Chen and Pennock 2007; Chen and Vaughan 2010], but we restrict our discussion to the former for technical tractability.

[1]See intrade.com for active public prediction markets like these.

(a) Most Accurate Prediction

	Profit	Loss
Springfield	$\frac{3}{4}$	$\frac{1}{4}$
Greenville	$\frac{1}{10}$	$\frac{9}{10}$

Expected score $= 1 + \frac{3}{4}\log\frac{3}{4} + \frac{1}{4}\log\frac{1}{4} \approx .44$

(b) Higher Scoring Prediction

	Profit	Loss
Springfield	$\frac{1}{20}$	$\frac{19}{20}$
Greenville	$\frac{1}{10}$	$\frac{9}{10}$

Expected score $= 1 + \frac{1}{10}\log\frac{1}{10} + \frac{9}{10}\log\frac{9}{10} \approx .86$

Fig. 2. A decision market adapted from a strictly proper prediction market is not strictly proper. In this example, the decision maker selects the action more likely to result in a profit and scores the prediction *for the selected action* using a strictly proper logarithmic scoring rule, $1 + \log \mathbf{p}_o$, where $\mathbf{p}_o$ is the prediction for the realized outcome given the selected action (see Section 2 for strictly proper scoring rules). The expert, who is assumed to be the last participant in the market, expects Springfield to result in a profit and Greenville a loss as the action-outcome matrix in (a) indicates but scores higher in expectation by claiming that Springfield is worse than Greenville as shown by the action-outcome matrix in (b). Since the decision maker never builds in Springfield, this deception is never caught.

Scoring rules are fitness functions for predictions. Originally suggested by Brier [1950] to describe the quality of weather forecasts, they are now used to score many kinds of predictions. An in-depth summary of scoring rules appears in [Gneiting and Raftery 2007]. We use the following definition.[2]

Definition 2.1 (*Regular Scoring Rule*). A function

$$s : \Delta(\mathcal{O}) \times \mathcal{O} \to \mathbb{R} \cup \{-\infty\}$$

where $\mathcal{O}$ is a finite set of mutually exclusive and exhaustive outcomes, and $\Delta(\mathcal{O})$ is the set of all probability distributions over outcomes. We often abbreviate $s(\mathbf{p}, o)$ as $s_o(\mathbf{p})$, and require

$$s_o(\mathbf{p}) = -\infty \implies \mathbf{p_o} = 0$$

that is, a (regular) scoring rule can assign a forecast $\mathbf{p}$ a score of negative infinity only if that forecast assigned zero probability to the actual outcome o.

Regularity is a common assumption for scoring rules, and we will restrict our attention to regular scoring rules in the rest of the paper. The interesting scoring rules are (strictly) proper.

Definition 2.2 (*(Strictly) Proper Scoring Rule*). A scoring rule s is proper if and only if

$$\mathbf{q} \in \arg\max_{\mathbf{p} \in \Delta(\mathcal{O})} \sum_o \mathbf{q}_o s_o(\mathbf{p}), \ \forall \mathbf{q} \in \Delta(\mathcal{O})$$

and strictly proper if and only if $\mathbf{q}$ is the unique maximizing argument.

A natural interpretation of a proper scoring rule is that an expert maximizes its expected score for a prediction by predicting what it believes to be true, or, equivalently, predicting as accurately as possible. If the scoring rule is strictly proper then this prediction uniquely maximizes its expected score.

[2]A note on mathemtical notation. For consistency we will use calligraphy ($\mathcal{A}$) for sets, capitals (A) for matrices, bold ($\mathbf{a}$) for vectors and lowercase (a) for functions, individual actions and outcomes, and reals. When we refer to the elements of matrix A or vector $\mathbf{a}$ we will use an index, like $A_{x,y}$ or $\mathbf{a}_y$.

We use in our technical results the following characterizations of (strictly) proper scoring rules, which are given in Gneiting and Raftery [2007] but credited to McCarthy [1956] and Savage [1971].

THEOREM 2.3 (GNEITING AND RAFTERY [2007]). *A scoring rule is (strictly) proper if and only if*

$$s(\mathbf{p}, o) = g(\mathbf{p}) - g^{\star}(\mathbf{p}) \cdot \mathbf{p} + g^{\star}_o(\mathbf{p})$$

where $g : \Delta(\mathcal{O}) \to \mathbb{R}$ *is a (strictly) convex function,* $g^{\star}(\mathbf{p})$ *is a subgradient of* g *at the point* $\mathbf{p}$*, and* $g^{\star}_o(\mathbf{p})$ *is the o-th element of* $g^{\star}(\mathbf{p})$*.*

COROLLARY 2.4 (GNEITING AND RAFTERY [2007]). *Any (strictly) proper scoring rule*

$$s(\mathbf{p}, o) = g(\mathbf{p}) - g^{\star}(\mathbf{p}) \cdot \mathbf{p} + g^{\star}_o(\mathbf{p})$$

satisfies

$$\sum_o \mathbf{p}_o s(\mathbf{p}, o) = g(\mathbf{q}),\ \forall \mathbf{p} \in \Delta(\mathcal{O}).$$

Our characterization of strictly proper decision making in Section 4 and Theorem 2.3 have a similar relationship with convex functions.

One well-known strictly proper scoring rule is the logarithmic scoring rule, $s_{\log}(\mathbf{p}, o) = a_o + b \log \mathbf{p}_o$ with $\log 0$ defined as equal to $-\infty$ and $b > 0$ and a_o being parameters. This rule can be constructed via Theorem 2.3 by letting the convex function be the prediction's negative Shannon entropy. This scoring rule already made a brief appearance in Figure 2 with $a_o = b = 1$.

Hanson adapted scoring rules for use in prediction markets [Hanson 2003; Hanson 2007], where experts sequentially improve over previous predictions and are scored for the improvement. Such a mechanism is called a *market scoring rule*. A prediction market opens with an initial prediction $\mathbf{p}^0$ and accepts a series of expert predictions $\mathbf{p}^1, \mathbf{p}^2, \ldots$. Each prediction receives a *net score* equal to the difference of its prediction's score and that of the immediately preceeding prediction. For example, if the outcome o^* occurs the net score for prediction p^t is

$$s(p^t, o^*) - s(p^{t-1}, o^*).$$

Since the market's outcome o^* and previous prediction p^{t-1} are fixed, maximizing the expected net score of a prediction is the same as maximizing its expected prediction score; thus (strictly) proper scoring rules make (strictly) proper prediction markets. Markets cheaply aggregate information from many experts, since each prediction except the last is paid for by the following one.

Hanson [1999] suggested adapting prediction markets to decision making by running a standard market scoring rule prediction market for each action and voiding markets for actions that are not taken. Such markets, as shown in Figure 2, are not strictly proper, however. We create strictly proper decision markets by extending market scoring rules to account for a decision maker's decision policy.

Othman and Sandholm [2010b] first formally developed the idea of using a single expert to predict the results of a finite set of actions and two outcomes of interest, "good" and "bad." Their decision maker wishes to choose the action maximizing the likelihood of the "good" outcome, and they developed *quasi*-strictly proper scoring rules that incentivized the expert to accurately reveal both the action most likely to cause the "good" outcome and how likely that action was to do so. The "quasi-" in quasi-strictly proper is due to the fact that the scoring rule allows the score-maximizing prediction for the non-optimal actions to be arbitrary. In Section 5.2 we generalize this result and characterize all types of preferences that have *right-action* rules, which incentivize an expert to reveal the decision maker's most preferred action. As a corollary, we also characterize the set of preferences that have quasi-strictly proper rules.

Like Othman and Sandholm [2010b], we only consider the decision maker as having preferences over the outcomes. Our experts are indifferent to the action taken and the eventual outcome, and they also have no means of changing what each action does. Other work related to decision making has relaxed these assumptions. Shi et al. [2009] considered a setting where experts can take actions to change the outcome of a prediction market, and defined *principal-aligned* scoring rules that incentivize them to take only "helpful" actions that are more likely to cause outcomes preferred by the decision maker. These rules are extremely similiar to our right-action rules, though we consider a broader class of preferences for the decision maker. Shi et al. [2009] also discuss how these rules might work with multiple experts instead of the single expert we consider in Section 5.2, although this discussion does not include decision markets.

More recently, Boutilier [2012] has described an expert with preferences over the outcomes and a decision maker who, similar to Othman and Sandholm [2010b], always takes the "best" reported action. To incentivize the expert to make an accurate prediction he introduces *compensation functions* that compensate the expert for any loss of utility it may incur when a less desirable outcome is achieved. He also details some realistic complexities of this setting, e.g. the decision maker may not precisely know the expert's utility function.

Adapting these additional considerations to strictly proper decision making is an interesting opportunity for future work.

3. A FRAMEWORK FOR DECISION MAKING

In this section we introduce our formal definition of strict properness for decision making, first describing how predictions are elicited, a decision is made, and experts are scored. This provides a basis for our adaptation of strict properness to decision making. In a standard prediction market—where no decision is made—whether the market is strictly proper or not depends only on its choice of scoring rule. Strictly proper decision making, on the other hand, depends on the relationship between the decision maker's method of making a decision and its method of scoring predictions. Thus, a decision making strategy in our setting is described by a pair (decision method, scoring method). Some such pairs are strictly proper *for a single expert*, others are strictly proper only *for a market*, and some are strictly proper for any number of experts. We will argue that restricting our attention to this last set of pairs is without loss of generality when we consider decision markets, and with only a slight loss of generality when working with a single expert. For completeness, we characterize strict properness *for an expert* separately in Section 5.1 after characterizing strict properness (for any number of experts) in Section 4.

3.1. Predictions and their Elicitation

We begin by describing how the decision maker elicits predictions and uses them to make a decision.

When making a decision we naturally think of a helpful prediction describing what is likely to occur for each choice. Concretely, if we let $\mathcal{A}$ be our finite set of possible actions, and $\mathcal{O}$ a finite set of outcomes, these predictions can be represented by an $\mathcal{A} \times \mathcal{O}$ action-outcomes matrix like the one in Figure 1. Each row is associated with an action and each column an outcome, and a row describes how likely each outcome is to occur if that row's action is taken. An outcome can be any observable property of the future the decision maker is interested in, but in this paper we assume the outcomes are chosen to be mutually exclusive and exhaustive so each row of the matrix is a conditional probability distribution. We let $\mathcal{P}$ be the set of all such matrices.

These predictions are elicited from either a single expert or many experts participating in a decision market. A single expert makes a single prediction P, and experts in a market make a series of predictions $P^1, P^2, \ldots$ after the market opens with an initial prediction P^0. We make no further assumptions about either process except that (1) experts can only observe prior predictions before making their own and (2) both processes produce a final prediction. This latter assumption is trivially true with a single expert. Whether a market produces a consensus prediction with rational participants[3] is beyond the scope of this paper, and we make no additional assumptions about market dynamics.

After the final prediction is produced, the decision maker must choose a single action, and we make two assumptions about this choice: first that the decision maker can draw an action stochastically, second that its method of decision making can be described as a function of the final prediction.

We define a decision rule as the function mapping from (consensus) predictions to distributions the action is drawn from, which we will also refer to as *decision policies.*

Definition 3.1 (*Decision Rule*). A function

$$d : \mathcal{P} \to \Delta(\mathcal{A})$$

mapping predictions in $\mathcal{P}$ to probability distributions over actions in $\Delta(\mathcal{A})$. We describe a decision rule as having full support if $d_a(P) > 0,\ \forall P, a$, and write $\mathbf{d} \in d(\cdot)$ for a decision policy in the image of d.

Decision rules "with full support" always create policies with full support, and in Section 4 we show that a decision rule with full support is necessary and sufficient for a strictly proper decision market. When working with a single expert a decision rule with full support is not required for strict properness, but the decision rules of interest will usually create policies with full support nonetheless; this caveat is discussed further in Section 5.1.

Once the action is picked an outcome is observed. As mentioned above, we assume that no expert can influence what outcome occurs except indirectly by changing the action taken.

3.2. Scoring a Single Expert

After a single expert makes its prediction, the decision maker applies its decision rule and an outcome is observed, the expert is scored. Instead of a scoring rule we use a *decision scoring rule* to assign the prediction a score.

Definition 3.2 (*Regular Decision Scoring Rule*). A function

$$s : \mathcal{A} \times \mathcal{O} \times \Delta(\mathcal{A}) \times \mathcal{P} \to \mathbb{R} \cup \{-\infty\}$$

mapping an action, outcome, decision policy and prediction to the extended reals; as a shorthand we let $s_{a,o}(\mathbf{d}, P) = s(a, o, \mathbf{d}, P)$. We also require

$$s_{a,o}(\mathbf{d}, P) = -\infty \to d_a = 0 \text{ or } P_{a,o} = 0$$

As with scoring rules, we will restrict our attention to regular decision scoring rules for the rest of the paper.

Given a decision rule d and decision scoring rule s, an expert with beliefs Q has an expected score for a prediction P of

$$\sum_{a,o} d_a(P) Q_{a,o} s_{a,o}(d(P), P)$$

[3]See [Ostrovsky 2009; Chen et al. 2012] for a theoretical discussion of information aggregation in prediction markets.

the sum of possible scores $(s_{a,o}(d(P), P))$ weighted by how likely each score is to be realized $(d_a(P)Q_{a,o})$. We now define strict properness[4] *for an expert.*

Definition 3.3 (*Strictly Proper for an Expert*). A pair (d, s) is strictly proper *for an expert* if and only if a single expert always uniquely maximizes its expected score by predicting its beliefs Q

$$\{Q\} = \arg\max_{P \in \mathcal{P}} \sum_{a,o} d_a(P) Q_{a,o} s_{a,o}(d(P), P), \ \forall Q \in \mathcal{P}$$

We will continue to emphasize that this notion of strict properness *applies only to a single expert* for easier reading, and we characterize all pairs that are strictly proper *for an expert* in Section 5.1.

3.3. Scoring Decision Markets

In a typical prediction setting a strictly proper scoring rule is strictly proper for both a single expert and many experts participating in a market, but a pair (d, s) that is strictly proper *for an expert* is not necessarily strictly proper *for a market* because the market has an additional complication—the decision rule is not applied to an expert's own prediction but the final prediction in the market. Experts in a market are rewarded for improving over the previous prediction made, and to avoid overloading the term "score" we define an expert's expected *net score* as the difference

$$\sum_{a,o} \mathbf{d}_a Q_{a,o} \big(s_{a,o}(\mathbf{d}, P) - s_{a,o}(\mathbf{d}, P')\big)$$

when an expert has beliefs Q, the decision policy is $\mathbf{d}$, the prior prediction is P' and the expert predicts P. We now define strict properness *for a market.*

Definition 3.4 (*Strictly Proper for a Market*). A pair (d, s) is strictly proper *for a market* if and only if an expert in a market always uniquely maximizes its expected net score by predicting its beliefs

$$\sum_{a,o} \mathbf{d}_a Q_{a,o} \big(s_{a,o}(\mathbf{d}, Q) - s_{a,o}(\mathbf{d}, P')\big) \geq \sum_{a,o} \mathbf{d'}_a Q_{a,o} \big(s_{a,o}(\mathbf{d'}, P) - s_{a,o}(\mathbf{d'}), P')\big),$$
$$\forall Q, P, P' \in \mathcal{P}, \mathbf{d}, \mathbf{d'} \in d(\cdot)$$

with the inequality strict if $P \neq Q$. A market using such a pair is called a strictly proper decision market.

Like with the notion of strict properness *for an expert* we will continue to stress the caveat that this notion of strict properness applies only *for a market.*

3.4. Strictly Proper Pairs

While a pair (d, s) may only be strictly proper *for an expert* or only strictly proper *for a market*, we can drop these caveats and simply talk about strictly proper pairs with little loss of generality. We begin by giving a formal definition of a strictly proper pair, then discuss its properties.

[4]We focus on strict properness rather than properness in this paper because properness as a property can be trivially achieved by some uninteresting scoring rules such as a constant scoring rule that assigns a constant value to all predictions.

Definition 3.5 (*Strictly Proper Pair*). A pair (d, s) is strictly proper if and only if a prediction's expected score is independent of the decision policy

$$\sum_{a,o} \mathbf{d}_a Q_{a,o} s_{a,o}(\mathbf{d}, P) = \sum_{a,o} \mathbf{d}'_a Q_{a,o} s_{a,o}(\mathbf{d}', P),\ \forall Q, P \in \mathcal{P}, \mathbf{d}, \mathbf{d}' \in d(\cdot) \tag{1}$$

and uniquely maximized when an expert predicts its beliefs

$$\{Q\} = \arg\max_{P \in \mathcal{P}} \sum_{a,o} \mathbf{d}_a Q_{a,o} s_{a,o}(\mathbf{d}, P),\ \forall Q \in \mathcal{P}, \mathbf{d} \in d(\cdot) \tag{2}$$

The first condition of strict properness, Equation 1, implies the expected scores of prior predictions in a decision market appear constant, and so independently maximizing a prediction's expected score also maximizes its expected net score. As we will discuss at the end of Section 4, the independence between a prediction's expected score and the decision policy also means that experts do not have to know the exact form of the decision rule *ex ante*. The second condition, Equation (2), taken together with the first, implies that an expert maximizes its expected score and expected net score by predicting its beliefs. The following proposition formalizes our argument that strictly proper pairs are strictly proper *for an expert* and *for a market* (the proof appears in the appendix).

PROPOSITION 3.6. *Every strictly proper pair (d, s) is strictly proper for both an expert and a market.*

As mentioned, there are pairs that are strictly proper *for a market* that are not strictly proper as defined here, but these pairs are an uninteresting technical caveat. Every pair that is strictly proper *for a market* can be replaced with a corresponding strictly proper pair without changing the expected net score of any predictions.

PROPOSITION 3.7. *For every pair $(d, \bar{s})$ that is strictly proper* for a market*, there exists a strictly proper pair (d, s) such that every prediction has the same expected net score*

$$\sum_{a,o} \mathbf{d}_a Q_{a,o} \big(\bar{s}_{a,o}(\mathbf{d}, P) - \bar{s}_{a,o}(\mathbf{d}, P')\big) \ = \ \sum_{a,o} \mathbf{d}_a Q_{a,o} \big(s_{a,o}(\mathbf{d}, P) - s_{a,o}(\mathbf{d}, P')\big),$$
$$\forall Q, P, P' \in \mathcal{P}, \mathbf{d} \in d(\cdot)$$

The proof appears in the appendix.

The set difference of strictly proper pairs and those that are strictly proper *for an expert* is not so readily described, and in Section 5.1 we characterize all pairs that are strictly proper *for an expert* separately. In short, while strictly proper pairs always have a decision rule with full support, some pairs that are strictly proper *for an expert* do not. However, pairs that are strictly proper *for an expert* create decision policies with full support *for most predictions*, so this distinction is unlikely to be useful in practice. We focus on strictly proper pairs as a technically incomplete but salient description of strictly proper decision making.

4. STRICTLY PROPER DECISION MAKING

In this section we characterize all strictly proper pairs, describing most of strictly proper decision making. It turns out that having full support is a necessary and sufficient condition for a decision rule to be part of a strictly proper pair. Returning to our example, if our firm runs a strictly proper decision market to decide whether to open a store in Springfield or Greenville, this implies that it can make its preferred decision with probability arbitrarily close—but never equal—to one. With some $\epsilon > 0$ chance it must build in the less preferred location. Even this ϵ risk may be unacceptable in practice, and in Section 5.2 we describe how a preferred decision policy can be directly elicted from a single expert, without learning the complete mapping from actions to outcomes.

4.1. Strictly Proper Pairs have Decision Rules with Full Support

We begin by showing that full support is necessary for the decision rules of strictly proper pairs.

THEOREM 4.1 (FULL SUPPORT IS NECESSARY FOR A STRICTLY PROPER PAIR). *If a pair (d, s) is strictly proper, d has full support.*

PROOF. Assume, for a contradiction, that d is a decision rule without full support and s is a decision scoring rule such that (d, s) is strictly proper. Let P^* be a prediction such that $d_{a'}(P^*) = 0$ for some action a', which must exist by our assumption that d does not have full support, and let Q and Q' be two action-outcome matrices differing only on action a'. Then we have

$$\begin{aligned} & \sum_{a,o} d_a(P^*)Q_{a,o}\big(s_{a,o}(d(P^*), P) - s_{a,o}(d(P^*), \bar{P})\big) \\ = & \sum_{a,o} d_a(P^*)Q'_{a,o}\big(s_{a,o}(d(P^*), P) - s_{a,o}(d(P^*), \bar{P})\big),\ \forall P, \bar{P} \in \mathcal{P} \end{aligned}$$

implying the same prediction maximizes the expected value of an expert who believes Q or Q', and since this prediction cannot be both Q and Q' at the same time, the pair (d, s) violates Equation (2) and so must not be strictly proper, a contradiction as desired. □

Simply put, experts have no incentive to be accurate on actions that are never tested, so a decision rule without full support cannot be strictly proper.

4.2. Constructing Strictly Proper Pairs

As mentioned, a decision rule with full support is a necessary and sufficient condition for it to be part of a strictly proper pair. In fact, given any decision rule d with full support we can construct a strictly proper pair (d, s) using a strictly proper scoring rule, $\bar{s}$, and letting

$$s_{a,o}(\mathbf{d}, P) = \frac{1}{\mathbf{d}_a}\bar{s}_o(P_a). \tag{3}$$

The expected score for a prediction is then

$$\sum_{a,o} \mathbf{d}_a Q_{a,o}\big(\frac{1}{\mathbf{d}_a}\bar{s}_o(P_a)\big) = \sum_{a,o} Q_{a,o}\bar{s}_o(P_a)$$

the same expected score as if an expert were participating in $|\mathcal{A}|$ independent and strictly proper prediction markets, one for each action. The following formal characterization shows that all strictly proper pairs can be described by a similar construction.

Some additional notation is needed before stating the theorem. We use a colon between two matrices to denote their Frobenius inner product, $A : B = \sum_{i,j} A_{ij}B_{ij}$, and let $g^\star(P)$ be a subgradient of the convex function g at P. The subgradient of a real-valued convex function is usually considered a vector, but since P is a matrix we also index the subgradient as a matrix.

THEOREM 4.2 (STRICTLY PROPER PAIR CHARACTERIZATION). *A pair (d, s) is strictly proper if and only if d has full support and there exists a strictly convex function g such that*

$$s_{a,o}(\mathbf{d}, P) = g(P) - g^\star(P) : P + \frac{g^\star_{a,o}(P)}{\mathbf{d}_a} \tag{4}$$

PROOF. We begin by showing that given a decision rule d with full support and a strictly convex g, defining a decision scoring rule s as in Equation 4 makes (d, s) a strictly proper pair.

An expert's expected score for predicting P with beliefs Q and decision policy $\mathbf{d}$ is

$$\begin{aligned}
&\sum_{a,o} \mathbf{d}_a Q_{a,o} s_{a,o}(\mathbf{d}, P) \\
=\ & \sum_{a,o} \mathbf{d}_a Q_{a,o} \big(g(P) - g^\star(P) : P + \frac{g^\star_{a,o}(P)}{\mathbf{d}_a}\big) \\
=\ & \sum_{a,o} \{\mathbf{d}_a Q_{a,o} \big(g(P) - g^\star(P) : P\big)\} + Q : g^\star(P) \\
=\ & g(P) - g^\star(P) : P + Q : g^\star(P) && (\text{since } \textstyle\sum_{a,o} d_a Q_{a,o} = 1) \\
=\ & g(P) + (Q - P) : g^\star(P)
\end{aligned}$$

which is independent of the decision policy, and the expert's expected score for accurately predicting Q is then

$$g(Q) + (Q - Q) : g^\star(Q) = g(Q)$$

and applying the subgradient inequality we have

$$g(Q) > g(P) + (Q - P) : g^\star(P),\ \forall P \neq Q \in \mathcal{P}$$

implying (d, s) is a strictly proper pair.

Now we show that given a strictly proper pair (d, s) it is necessary that d have full support and there exists a strictly convex g such that s is as defined in Equation 4. Since Theorem 4.1 proved the necessity of d having full support, we only need prove the latter condition.

As a shorthand, we define an expected score function

$$v(\mathbf{d}, Q, P) = \sum_{a,o} \mathbf{d}_a Q_{a,o} s_{a,o}(\mathbf{d}, P)$$

and recall from Definition 3.5 that

$$v(\mathbf{d}, Q, P) = v(\mathbf{d}', Q, P), \forall Q, P \in \mathcal{P}, \mathbf{d}, \mathbf{d}' \in d(\cdot)$$

allowing us to write simply $v(Q, P)$; our strictly convex function g will be $g(P) = v(P, P)$, which is convex (and we will verify is strictly convex shortly), and we'll use

$$g^\star_{a,o}(P) = \bar{\mathbf{d}}_a s_{a,o}(\bar{\mathbf{d}}, P)$$

where $\mathbf{d}$ is any decision policy in $d(\cdot)$, as our subgradient at P. We verify it is a subgradient by checking the subgradient inequality:

$$\begin{aligned}
& g(P) + (Q - P) : g^\star(P) \\
& = v(P, P) + \sum_{a,o} (Q_{a,o} - P_{a,o}) \bar{\mathbf{d}}_a S_{a,o}(\bar{\mathbf{d}}, P) \\
& = v(P, P) + v(Q, P) - v(P, P) \\
& = v(Q, P) \\
& < v(Q, Q)
\end{aligned}$$

for all $P \neq Q \in \mathcal{P}$. The strict inequality following since (d, s) is a strictly proper pair and this strict inequality implies g is strictly convex [Hendrickson and Buehler 1971].

Before concluding, we note that since (d, s) is a strictly proper pair

$$\mathbf{d}_a s_{a,o}(\mathbf{d}, P) = \mathbf{d}'_a s_{a,o}(\mathbf{d}', P),\ \forall \mathbf{d}, \mathbf{d}' \in d(\cdot), P \in \mathcal{P}, a \in \mathcal{A}, o \in \mathcal{O}$$

(otherwise there exist beliefs Q such that $v(\mathbf{d}, Q, P) \neq v(\mathbf{d}', Q, P)$), and we use this fact to verify that g with subgradients as given is, in fact, equal to s

$$
\begin{aligned}
& g(P) - g^\star(P) : P + \frac{g^\star_{a,o}(P)}{\mathbf{d}_a} \\
= \ & v(P,P) - \sum_{a,o}\{P_{a,o}\bar{\mathbf{d}}_a s_{a,o}(\bar{\mathbf{d}}, P)\} + \frac{\bar{\mathbf{d}} s_{a,o}(\bar{\mathbf{d}}, P)}{\mathbf{d}_a} \\
= \ & v(P,P) - v(P,P) + \frac{1}{\mathbf{d}_a}\big(\bar{\mathbf{d}} s_{a,o}(\bar{\mathbf{d}}, P)\big) \\
= \ & \frac{1}{\mathbf{d}_a}\big(\mathbf{d} s_{a,o}(\mathbf{d}, P)\big) \\
= \ & s_{a,o}(\mathbf{d}, P)
\end{aligned}
$$

So from any strictly proper pair we can construct a strictly convex g satisfying Equation 4. □

Theorem 4.2 shows that while a decision maker can take a preferred action with probability arbitrarily close to one, it cannot deterministically take its preferred action and still run a strictly proper decision market. Note, however, that it is sufficient for experts to believe they will be scored in a strictly proper fashion, and the decision maker does not have to *ex ante* design its decision rule. Instead it can simply review the final prediction, construct any decision policy with full support, and then score the experts using an appropriate decision scoring rule to create a strictly proper decision market. This seems a preferable approach in practice to accounting for every contingency in advance.

A good analogy to running a strictly proper decision market is to an overwhelmed teaching assistant grading a midterm. The teaching assistant does not have the time to grade every question, but must pick one from each test. If some questions are more likely to be graded than others then students will spend more time on those and neglect the rest, biasing their scores. Only by (1) possibly grading any question and (2) weighting that question's score by the inverse likelihood its graded will the teaching assistant create an unbiased estimator, where the student's expected grade is the same as if every question were reviewed. This encourages students to pay equal attention to each question and not "game the system."

5. WORKING WITH A SINGLE EXPERT

When there is only one expert making a prediction its prediction is also the last prediction, and this coupling provides some additional structure to exploit. If we are still eliciting the likely outcomes of each action, this distinction allows strictly proper pairs *for an expert* to have a decision policy without full support. But more interestingly, it can allow a decision maker to directly request and then take a preferred action. This latter technique doesn't reveal what *would have happened* if other actions were taken, but allows the decision maker to deterministically take its preferred action instead of risking taking any action.

5.1. Pairs that are Strictly Proper for an Expert

As discussed in Section 3, we call a pair strictly proper if it is strictly proper *for an expert* and strictly proper *for a market*. When describing strictly proper decision markets this distinction is without loss of generality, but the same is not true when working with only a single expert. In this section we characterize strict properness *for an expert*. These pairs, unlike strictly proper pairs, can have decision rules without full support. Such decisions rules will still tend to create *decision policies* with full support, however, suggesting that these pairs are unlikely to be preferred over pairs that are simply strictly proper. Formally, the set of predictions a pair that is strictly proper *for an expert* maps to a decision policy without full support is nowhere dense in the set of all possible predictions. The proofs of this section's theorems appear in the appendix.

THEOREM 5.1. *For any pair (d, s) that is strictly proper* for an expert*, the set of action-outcome matrices that d maps to distributions without full support is nowhere dense in the set of all action-outcome matrices with their natural Euclidean topology.*[5]

Intuitively, this means that for any $d(P)$ without full support, there is a P' arbitrarily close to P, such that $d(P')$ does have full support. Thus, while the set of pairs that are strictly proper *for an expert* is larger the set of pairs that are just strictly proper, this additional latitude is rarely useful.

We conclude with a complete characterization of pairs that are strictly proper *for an expert*. The statement and its proof are similar to those of Theorem 4.2.

THEOREM 5.2 (STRICTLY PROPER FOR AN EXPERT CHARACTERIZATION). *A pair (d, s) is strictly proper* for an expert *if and only if there exist a strictly convex function g and a subgradient $g^\star(P)$ where $g^\star_a(P) = \mathbf{0}$ whenever $d_a(P) = 0$ such that*

$$s_{a,o}(d(P), P) = g(P) - g^\star(P) : P + \frac{g^\star_{a,o}(P)}{d_a(P)},\ \forall d_a(P) > 0. \tag{5}$$

Following this characterization, if the decision maker was averse to actions with a particular distribution over the outcomes, $\mathbf{p} \in \Delta(\mathcal{O})$, it may pick a strictly convex function $g : \mathbb{R}^{|\mathcal{O}|} \to \mathbb{R}$ with a minima at $\mathbf{p}$. It may then construct a decision scoring rule from g using the simple construction (3) from the previous section, and if the expert predicts a conditional distribution of $\mathbf{p}$ for some action (or set of actions) avoid it completely. Note also that Theorem 5.2 only restricts $s_{a,o}(d(P), P)$ when $d_a(P) > 0$. When $d_a(P) = 0$, $s_{a,o}(d(P), P)$ can take any value as long as the decision scoring rule is regular.

This concludes our discussion of strictly proper decision making, where a decision maker solicits the complete mapping from actions to outcomes. In the next section we discuss an alternative where, instead of this mapping, a decision policy is directly solicited. This alternative allows the decision maker to deterministically take a preferred action, instead of doing so with high probability.

5.2. Optimal Decision Making

As shown in Sections 4 and 5.1, strictly proper decision making (generally) requires the decision maker risk taking any action, even if it learns that some actions will result in undesirable outcomes. This is not ideal and in practice may be non-credible. It is hard to imagine a real firm learning that Springfield is a better location for a new store than Greenville, but still building in the latter because of a coin flip.

In this section we describe a new method for making a decision, still derived from scoring rules, that can incentivize a single expert to directly reveal a preferred decision policy, instead of describing the complete mapping from actions to outcomes. This allows the decision maker to take deterministic actions, but not to learn what would have happened if other actions were chosen (e.g. the expert tells our firm to build in Springfield, and not what would happen if it built in Greenville). This approach works for many—but not all—types of decision maker preferences, including the commonly considered preferences of an expected value maximizer who assigns a value to each outcome.

[5]A set is nowhere dense in a topological space if the interior of its closure, with respect to the topological space, is empty.

We assume the decision maker's preferences can be represented by a utility function[6]

$$u : \Delta(\mathcal{O}) \to \mathbb{R},$$

mapping probability distributions over the outcomes, or lotteries, to the reals. Figure 3 shows some example utility functions.

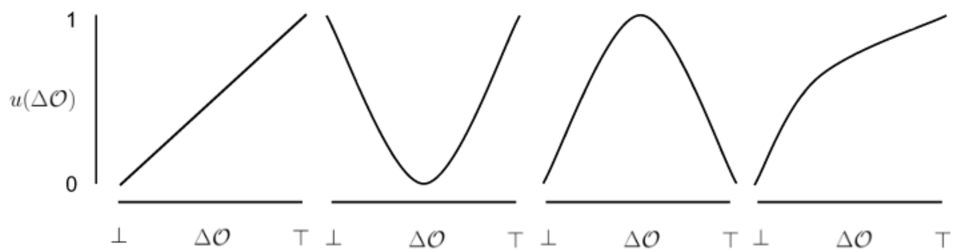

Fig. 3. Four utility functions of lotteries over two outcomes, $\top$ and $\perp$. The x-axis represents the probability of outcome $\top$ and the y-axis is the utility for the lotteries. At left is the utility function of an expected utility maximizer, with the outcome utility of $\perp$ zero and $\top$ one. Middle-left is an extremal utility function that prefers "certainty," middle-right is a utility function that values "uncertainty" or "fairness" and at right is the utility function of a risk averse decision maker.

The decision maker's goal is to elicit a decision policy whose resultant lottery (the distribution over outcomes conditional on that action being taken) maximizes its utility function. Formally, if an expert has beliefs Q the decision maker wants to find a decision policy $\mathbf{d}^*$ solving

$$\max_{\mathbf{d} \in \Delta(\mathcal{A})} u(Q^T \mathbf{d})$$

where $Q^T\mathbf{d}$ is the distribution over outcomes conditional on drawing an action according to $\mathbf{d}$ (we use the transpose of the action-outcome matrix Q since we treat all vectors as column vectors). Also, while a decision maker has preferences over lotteries of outcomes, we will describe a decision policy $\mathbf{d}$ as preferred or utility maximizing, too, with the understanding we are describing its conditional lottery $Q^T\mathbf{d}$.[7]

Instead of predicting an action-outcome matrix we now have the expert report a decision policy and a lottery $(\mathbf{d}, \mathbf{p}) \in \Delta(\mathcal{A}) \times \Delta(\mathcal{O})$, and we let $\mathcal{R}$ be the set of such reports. After the expert makes its report, the decision maker draws an action according to the reported decision policy $\mathbf{d}$, observes an outcome o^* and scores the expert using $s_{o^*}(\mathbf{p})$, where s is a scoring rule. What we will show is that if the decision maker's preferences can be ordered by a convex utility function, this function can be used to create a scoring rule that aligns the expert's incentives with the decision maker's. The expert will want to be scored for conditional distributions more preferred by the decision maker, and so will report an optimal decision policy (although not necessarily the correct conditional distribution).

[6]We adopt the generic textbook definition of a utility function, implicitly assuming the decision maker's preferences are complete and transitive. See Mas-Colell et al. [1995] for details. This definition includes the Von Neumann-Morgenstern utility function for expected value maximizers as a special case but is more general. The utility function of an expected utility maximizer who has value v_o for outcome o is $u(\mathbf{p}) = \sum_{o \in \mathcal{O}} \mathbf{p}_o v_o$ for lottery $\mathbf{p}$.
[7]We consider eliciting a decision policy, and not just a single action, for generality, because some preferred lotteries may require a convex combination of several actions to construct. We will see, however, that we cannot elicit preferred decision policies for preferences like this, and when we can incentivize an expert to reveal a preferred decision policy, a single action will always maximize the decision maker's utility.

5.2.1. Right-action Rules and Making an Optimal Decision. Before proving this result we need some new notation. Given a scoring rule s, an expert with beliefs Q maximizes its expected score by reporting in the set

$$\mathcal{R}^s(Q) = \arg\max_{(\mathbf{d},\mathbf{p})\in\mathcal{R}} \sum_{a,o} \mathbf{d}_a Q_{a,o} s_o(\mathbf{p})$$

and we use $\mathcal{R}^s_{\mathbf{d}}(Q) = \{\mathbf{d} | \exists \mathbf{p}\ (\mathbf{d},\mathbf{p}) \in \mathcal{R}^s(Q)\}$ for the set of decision policies that are part of such pairs. If an expert always maximizes its expected score by revealing a utility-maximizing decision policy, we describe the scoring rule s as a *right-action rule* for the utility function u.

Definition 5.3 (*Right-Action Rule*). A scoring rule s is a right-action rule for a utility function u if

$$u(Q^T\mathbf{d}^*) \geq u(Q^T\mathbf{d}),\ \forall Q \in \mathcal{P}, \mathbf{d} \in \Delta(\mathcal{A}), \mathbf{d}^* \in \mathcal{R}^s_{\mathbf{d}}(Q)$$

This section's primary result is that a right-action rule exists for a utility function u if and only if u is convexible.

Definition 5.4 (*(Strictly) Convexible*). A function f is (strictly) convexible if there exists a (strictly) convex function g such that if $f(x) > f(y)$ then $g(x) > g(y)$.

In other words, a right-action rule for a decision maker's utility function exists if and only if that function's strict preference ordering can be represented by a convex utility function. We stress that this convex utility function may not represent the decision maker's actual preferences because it does not need to preserve indifferences. Intuitively, if we are equally happy with an apple or an orange we have the latitude of "preferring" one to the other without loss of utility, and this flexibility may be useful in constructing a convex ordering of our preferences. This allows, for example, a constant utility function to be strictly convexible with any strictly convex function satisfying the criteria.

Many but not all preferences are convexible. All convex utility functions are trivially convexible, and these include the utility functions of expected value maximizers. Some strictly concave utility functions are convexible as well. Of the four example utility functions in Figure 3, the two at left are convex and thus trivially convexible, but the rightmost and concave utility function is convexible, too, since it's strictly increasing. Only the middle-right function is not convexible.

We can now state our result that a right-action rule exists if and only if the expert's utility function is convexible. The formal statement adds the caveat that the decision maker consider at least two actions, since otherwise it is running a prediction market and any scoring rule is trivially a right-action rule.

THEOREM 5.5 (RIGHT-ACTION RULE CHARACTERIZATION). *If the decision maker is considering at least two actions, a utility function u has a right-action rule if and only if u is convexible.*

PROOF. We first show that any convexible utility function has a right-action rule. Since u is convexible there exists a convex function g such that $u(x) > u(y)$ implies $g(x) > g(y)$. We use this to construct a scoring rule

$$s_o(\mathbf{p}) = g(\mathbf{p}) - g^\star(\mathbf{p}) \cdot \mathbf{p} + g^\star_o(\mathbf{p})$$

and by Theorem 2.3 this scoring rule is proper (for a typical prediction market) with expected score function

$$g(\mathbf{p}) = \sum_o \mathbf{p}_o s_o(\mathbf{p})$$

Since s is proper an expert with beliefs Q maximizes its expected score by solving

$$\max_{\mathbf{d},\mathbf{p}} \sum_o (Q^T \mathbf{d})_o s_o(\mathbf{p})$$

and since s is a proper, Corollary 2.4 and our construction imply

$$\arg\max_{\mathbf{d}} \max_{\mathbf{p}} \sum_o (Q^T \mathbf{d})_o s_o(\mathbf{p}) \in \arg\max_{\mathbf{d}} u(Q^T \mathbf{d}) \tag{6}$$

implying s is a right-action rule for u.

Now assume, for a contradiction, that u is not convexible and has a right-action rule s. The expert's expected maximizing score function over lotteries is

$$v(\mathbf{q}) = \sup_{\mathbf{p}} \sum_o \mathbf{q}_o s_o(\mathbf{p})$$

which is convex, implying there exists $\mathbf{q}$ and $\mathbf{q}'$ such that

$$v(\mathbf{q}) > v(\mathbf{q}')$$
$$u(\mathbf{q}) < u(\mathbf{q}')$$

since u is not convexible. Because v is convex

$$v(\mathbf{q}) \geq v(\alpha\mathbf{q} + (1-\alpha)\mathbf{q}'), \forall \alpha \in (0,1)$$

so an expert prefers to report policies with conditional distribution $\mathbf{q}$, even when policies with conditional distribution $\mathbf{q}'$, which is preferred by the decision maker, or convex combinations of the two are available. This implies that if the expert has beliefs Q with $Q_a = \mathbf{q}$ and $Q_{a'} = \mathbf{q}', \forall a' \neq a$, then it maximizes its expected score by reporting a policy assigning probability one to action a, even though the decision maker prefers policies that assign zero probability to a, contradicting our assumption that s is a right-action rule for u. $\square$

There are two important points to be made about this proof. First, the result is not constructive because it provides no means of finding a convex function that preserves the decision maker's strict preferences, and we leave this problem for future work. Second, the reported lottery $\mathbf{p}$ may not be an accurate prediction of what will occur when an action is drawn according to the report decision policy $\mathbf{d}$. If, for example, the decision maker is an expected value maximizer with a utility for each outcome, then simply paying the expert a portion of the realized utility is a right-action rule. This rule ignores the lottery $\mathbf{p}$ when scoring and so an expert can maximize its score while setting it arbitrarily.

While we have allowed the expert to report an entire decision policy, convexible preferences imply reporting a single action can always maximize a decision maker's utility, too.

PROPOSITION 5.6 (REPORTING ONE ACTION IS ENOUGH). *For any utility function u with a right-action rule s, there exists an action a and decision policy $\mathbf{d} \in \mathcal{R}^s_d(Q)$, $\forall Q \in \mathcal{P}$ where $\mathbf{d}_a = 1$.*

PROOF. An expert with beliefs Q maximizes their expected score by reporting in

$$\mathcal{R}^s(Q) = \arg\max_{(\mathbf{d},\mathbf{p}) \in \mathcal{R}} \sum_{a,o} \mathbf{d}_a Q_{a,o} s_o(\mathbf{p}).$$

Let $(\mathbf{d}^*, \mathbf{p}^*)$ be an element of this set. If $\mathbf{d}^*$ already has an action a where $\mathbf{d}^*_a = 1$ then we're done. Otherwise, for all actions a and a' with $\mathbf{d}^*_a > 0$ and $\mathbf{d}^*_{a'} > 0$, we will show that

$$\sum_o Q_{a,o} s_o(\mathbf{p}) = \sum_o Q_{a',o} s_o(\mathbf{p})$$

which implies that $(\mathbf{d}', \mathbf{p}^*)$ with $\mathbf{d}'_a = 1$ also maximizes the expert's expected score and completes the proof. Suppose the statement is false, then

$$\sum_o Q_{a,o} s_o(\mathbf{p}) > \sum_o Q_{a',o} s_o(\mathbf{p})$$

and reporting $(\mathbf{d}', \mathbf{p}^*)$ has a higher expected value than reporting $(\mathbf{d}', \mathbf{p}^*)$ has a higher expected value than reporting $(\mathbf{d}^*, \mathbf{p}^*)$, contradicting our assumption that the latter pair maximized the expert's expected score. □

In short, although there may be optimal decision policies that assign positive probability to multiple actions, the decision maker must be indifferent between them. E.g. if opening a store in Springfield or Greenville will result in the same profit then any combination of the actions has the same expected value for the decision maker, but the decision maker loses nothing by requiring the expert specify one city or the other.

5.2.2. Quasi-Strict Properness. Taking a preferred action is good, but it may also be important to know the likely result of that action. When the decision maker's utility function is strictly convexible we can construct scoring rules that incentivize revealing not only the preferred policy, but also an accurate conditional lottery. We call these scoring rules *quasi-strictly proper* after Othman and Sandholm [2010b] (see Section 2). In contrast to strictly proper pairs for an expert, which incentivize an expert to accurately reveal the likelihood of the outcomes conditional on each possible action, quasi-strictly proper rules incentivize the accurate revelation of the likelihood of the outcomes conditional on following the preferred decision policy. This is a generalization of Othman and Sandholm [2010b]'s definition, where they consider deterministic decision policies and hence their quasi-strictly proper rules accurately reveal the likelihood of the outcomes conditional on the action that would be taken (but not necessarily for other actions).

Definition 5.7 (*Quasi-Strictly Proper*). A scoring rule s is quasi-strictly proper for a utility function u if it is a right-action rule for u and

$$\mathbf{p} = Q^T \mathbf{d},\ \forall Q \in \mathcal{P}, (\mathbf{d}, \mathbf{p}) \in \mathcal{R}^s(Q)$$

COROLLARY 5.8 (QUASI-STRICTLY PROPER CHARACTERIZATION). *If u is a strictly convexible utility function, there exists a quasi-strictly proper scoring rule for u.*

The proof is immediate from the first part of Theorem 5.5, since a strictly convex function implies Equation 6 is uniquely maximized when the prediction $\mathbf{p}$ is equal to the conditional lottery $Q^T \mathbf{d}$.

In practice quasi-strictly proper scoring rules may be interesting as they suggest the decision maker can "look before they leap" and plan for the effect of their decision.

6. CONCLUSION

We described how a decision maker can create a strictly proper incentive for experts to predict the mapping from its actions to some outcomes of interest. Understanding this mapping allows the decision maker to make an informed decision, although it can only do so with high probability—with some $\epsilon > 0$ chance it must risk taking an action at random. Even this small chance may be too much to accept in practice, so we also propose a method where a single expert directly reveals a preferred decision policy; whether a decision maker can take advantage of this technique depends on the convexibility of its preferences.

While we have discussed our work in the context of decision making, the techniques can be applied to other settings where observations are incomplete or costly. For example, if we would like experts to predict how many barrels of oil are underneath the ocean at five sites, it is costly to drill in each one to verify the predictions (and defeats the purpose of asking for them!). Using a decision market, the market maker can drill in only one location and still expect an accurate prediction from the experts. If this location happens to be the one with lots of oil, all the better; otherwise, the decision maker can simply drill there next, only paying the costs of one additional borehole instead of five.

There are many possible variations for decision making. We restricted our attention to discrete decision making, where a decision maker considers finite sets of actions and outcomes. We not address situations where experts are asked to predict a mapping from an infinite number of actions to an infinite number of outcomes, a situation that might occur if the decision maker is trying to locate a facility, for example. Our model might be extended in other ways, too: as discussed in Section 2, prior work has already considered experts able to take actions affecting the eventual outcome or with their own preferences among the outcomes. Finally, we have left some more immediate questions open, such as describing which types of preferences are convexible, and which are strictly convexible.

7. HISTORY

Portions of this paper originally appeared in [Chen and Kash 2011] and [Chen et al. 2011], which appeared at the Tenth Conference on Autonmous Agents and Mulitagent Systems (AAMAS 2011) and the Seventh Workshop on Internet and Network Economics (WINE 2011), respectively. In particular, Theorem 5.2 appeared in [Chen and Kash 2011] and Theorems 4.1 and 4.2 in [Chen et al. 2011].

This paper's discussion of strict properness *for an expert* and strict properness *for a market*, as well as their formal unification, is new, as are Theorem 5.1 and the entirety of Section 5.2. We have also updated the presentation of prior results for clarity and consistency.

REFERENCES

Berg, J. E., Forsythe, R., Nelson, F. D., and Rietz, T. A. 2001. Results from a dozen years of election futures markets research. In *Handbook of Experimental Economic Results*, C. A. Plott and V. Smith, Eds.

Berg, J. E. and Rietz, T. A. 2003. Prediction markets as decision support systems. *Information Systems Frontier 5*, 79–93.

Boutilier, C. 2012. Eliciting forecasts from self-interested experts: Scoring rules for decision markets. *In Proceedings of the Eleventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS-12).*

Brier, G. W. 1950. Verification of forecasts expressed in terms of probability. *Monthly Weather Review 78,* 1, 1–3.

Chen, Y., Kash, I., Ruberry, M., and Shnayder, V. 2011. Decision markets with good incentives. In *WINE'11: Proceedings of the 7th international conference on Internet and Network Economics*. Springer-Verlag, Berlin, Heidelberg, 72–83.

Chen, Y. and Kash, I. A. 2011. Information elicitation for decision making. In *AAMAS '11: Proceedings of the 10th International Conference on Autonomous Agents and Multiagent Systems.*

Chen, Y. and Pennock, D. M. 2007. A utility framework for bounded-loss market makers. In *UAI '07: Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence*. 49–56.

Chen, Y. and Pennock, D. M. 2010. Designing markets for prediction. *AI Magazine 31,* 4, 42–52.

Chen, Y., Ruberry, M., and Vaughan, J. W. 2012. Designing informative securities. *In Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence (UAI).*

Chen, Y. and Vaughan, J. W. 2010. A new understanding of prediction markets via no-regret learning. In *EC '10: Proceedings of the 11th ACM Conference on Electronic Commerce*. ACM, New York, NY, USA, 189–198.

Forsythe, R., Nelson, F., Neumann, G. R., and Wright, J. 1992. Anatomy of an experimental political stock market. *American Economic Review 82,* 5, 1142–1161.

Gneiting, T. and Raftery, A. E. 2007. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association 102,* 477, 359–378.

Hanson, R. 1999. Decision markets. *IEEE Intelligent Systems 14,* 3, 16–19.

HANSON, R. D. 2003. Combinatorial information market design. *Information Systems Frontiers 5,* 1, 107–119.

HANSON, R. D. 2007. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets 1,* 1, 1–15.

HENDRICKSON, A. D. AND BUEHLER, R. J. 1971. Proper scores for probability forecasters. *Annals of Mathematical Statistics 42,* 6, 1916–1921.

MANGOLD, B., DOOLEY, M., DORNFEST, R., FLAKE, G. W., HOFFMAN, H., KASTURI, T., AND PENNOCK, D. M. 2005. The tech buzz game. *IEEE Computer 38,* 7, 94–97.

MAS-COLELL, A., WHINSTON, M. D., AND GREEN, J. R. 1995. *Microeconomics Theory.* Oxford University Press, New York, NY.

MCCARTHY, J. 1956. Measures of the value of information. *PNAS: Proceedings of the National Academy of Sciences of the United States of America 42,* 9, 654–655.

OSTROVSKY, M. 2009. Information aggregation in dynamic markets with strategic traders. In *EC '09: Proceedings of the tenth ACM conference on Electronic commerce.* ACM, New York, NY, USA, 253.

OTHMAN, A. AND SANDHOLM, T. 2010a. Automated market making in the large: The Gates Hillman prediction market. In *EC '10: Proceedings of the 11th ACM Conference on Electronic Commerce.* ACM, New York, NY, USA, 367–376.

OTHMAN, A. AND SANDHOLM, T. 2010b. Decision rules and decision markets. In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems (AAMAS).* 625–632.

OTHMAN, A., SANDHOLM, T., PENNOCK, D. M., AND REEVES, D. M. 2010. A practical liquidity-sensitive automated market maker. In *EC '10: Proceedings of the 11th ACM Conference on Electronic Commerce.* ACM, New York, NY, USA, 377–386.

PENNOCK, D. M. 2004. A dynamic pari-mutuel market for hedging, wagering, and information aggregation. In *Proceedings of the Fifth ACM Conference on Electronic Commerce (EC'04).*

PLOTT, C. R., WIT, J., AND YANG, W. C. 1997. Parimutuel betting markets as information aggregation devices: Experimental results. Tech. Rep. Social Science Working Paper 986, California Institute of Technology. Apr.

SAVAGE, L. J. 1971. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association 66,* 336, 783–801.

SHI, P., CONITZER, V., AND GUO, M. 2009. Prediction mechanisms that do not incentivize undesirable actions. In *WINE '09: Proceedings of the 5th International Workshop on Internet and Network Economics.* Springer-Verlag, Berlin, Heidelberg, 89–100.

WOLFERS, J. AND ZITZEWITZ, E. 2004. Prediction markets. *Journal of Economic Perspective 18,* 2, 107–126.

Appendix

This appendix contains the proofs of Propositions 3.6 and 3.7 that appeared in Section 3 and the proofs of Theorems 5.1 and 5.2 that appeared in Section 5.1.

7.1. Proof of Proposition 3.6

PROPOSITION 3.6. *Every strictly proper pair (d, s) is strictly proper for an expert and a market.*

PROOF. Let (d, s) be a strictly proper pair. Strict properness *for an expert* requires

$$\arg\max_{P \in \mathcal{P}} \sum_{a,o} d_a(P) Q_{a,o} s_{a,o}(d(P), P) = \{Q\}, \ \forall Q \in \mathcal{P} \tag{7}$$

which is always satisfied since, from the definition of strict properness an expert's true beliefs Q always maximize an expert's expected score with respect to any decision policy $\mathbf{d}$

$$\arg\max_{P \in \mathcal{P}} \sum_{a,o} \mathbf{d}_a Q_{a,o} s_{a,o}(\mathbf{d}, P) = \{Q\}, \ \forall Q \in \mathcal{P}, \mathbf{d} \in d(\cdot)$$

and and the expected score is independent of the decision policy

$$\sum_{a,o} \mathbf{d}_a Q_{a,o} s_{a,o}(\mathbf{d}, P) = \sum_{a,o} \mathbf{d}'_a Q_{a,o} s_{a,o}(\mathbf{d}', P), \ \forall Q, P \in \mathcal{P}, \mathbf{d}, \mathbf{d}' \in d(\cdot)$$

implying Q is also the unique maximizer for Equation 7.

Strict properness *for a market* requires

$$\begin{aligned}&\sum_{a,o} \mathbf{d}_a Q_{a,o}\big(s_{a,o}(\mathbf{d}, Q) - s_{a,o}(\mathbf{d}, P')\big)\\ \geq\ &\sum_{a,o} \mathbf{d'}_a Q_{a,o}\big(s_{a,o}(\mathbf{d'}, P) - s_{a,o}(\mathbf{d'}, P')\big),\ \forall Q, P, P' \in \mathcal{P}, \mathbf{d}, \mathbf{d'} \in d(\cdot)\end{aligned}$$

with the inequality strict if $P \neq Q$. From the definition of strictly proper pair we have

$$\sum_{a,o} \mathbf{d}_a Q_{a,o} s_{a,o}(\mathbf{d}, P) = \sum_{a,o} \mathbf{d}'_a Q_{a,o} s_{a,o}(\mathbf{d'}, P),\ \forall Q, P \in \mathcal{P}, \mathbf{d}, \mathbf{d'} \in d(\cdot)$$

implying the difference

$$\begin{aligned}&\sum_{a,o} \mathbf{d}_a Q_{a,o}\big(s_{a,o}(\mathbf{d}, Q) - s_{a,o}(\mathbf{d}, P')\big) - \mathbf{d'}_a Q_{a,o}\big(s_{a,o}(\mathbf{d'}, P) - s_{a,o}(\mathbf{d'}, P')\big)\\ =\ &\sum_{a,o} \mathbf{d}_a Q_{a,o} s_{a,o}(\mathbf{d}, Q) - \mathbf{d}_a Q_{a,o} s_{a,o}(\mathbf{d}, P)\end{aligned}$$

and

$$\sum_{a,o} \mathbf{d}_a Q_{a,o} s_{a,o}(\mathbf{d}, Q) \geq \sum_{a,o} \mathbf{d}_a Q_{a,o} s_{a,o}(\mathbf{d}, P),\ \forall Q, P \in \mathcal{P}, \mathbf{d} \in d(\cdot)$$

unless $P = Q$ since Q is the unique maximizing argument to the expression. So the strictly proper *for a market* inequality always holds. □

7.2. Proof of Proposition 3.7

PROPOSITION 3.7. *For every pair $(d, \bar{s})$ that is strictly proper* for a market*, there exists a strictly proper pair (d, s) such that every prediction has the same expected net score*

$$\begin{aligned}&\sum_{a,o} \mathbf{d}_a Q_{a,o}\big(\bar{s}_{a,o}(\mathbf{d}, P) - \bar{s}_{a,o}(\mathbf{d}, P')\big)\\ =\ &\sum_{a,o} \mathbf{d}_a Q_{a,o}\big(s_{a,o}(\mathbf{d}, P) - s_{a,o}(\mathbf{d}, P')\big),\ \forall Q, P, P' \in \mathcal{P}, \mathbf{d} \in d(\cdot)\end{aligned}$$

PROOF. We will prove there exists a decision policy $\bar{\mathbf{d}} \in d(\cdot)$ and function $c_{\bar{\mathbf{d}}}$ such that

$$\mathbf{d}_a \bar{s}_{a,o}(\mathbf{d}, P) - c_{\bar{\mathbf{d}}}(a, o, \mathbf{d}) = \bar{\mathbf{d}}_a \bar{s}_{a,o}(\bar{\mathbf{d}}, P),\ \forall P \in \mathcal{P}. \tag{8}$$

In other words, there exists a function, independent of the prediction P, that translates from the expected score under one decision policy to another (in this case, $\bar{d}$). From the definition of strict properness *for a market* we have

$$\begin{aligned}&\sum_{a,o} \mathbf{d}_a Q_{a,o}\big(\bar{s}_{a,o}(\mathbf{d}, Q) - \bar{s}_{a,o}(\mathbf{d}, P')\big)\\ \geq\ &\sum_{a,o} \mathbf{d'}_a Q_{a,o}\big(\bar{s}_{a,o}(\mathbf{d'}, P) - \bar{s}_{a,o}(\mathbf{d'}, P')\big),\ \forall Q, P, P' \in \mathcal{P}, \mathbf{d}, \mathbf{d'} \in d(\cdot)\\ \implies &\sum_{a,o} \mathbf{d}_a Q_{a,o}\big(\bar{s}_{a,o}(\mathbf{d}, Q) - \bar{s}_{a,o}(\mathbf{d}, P)\big) = \sum_{a,o} \mathbf{d}'_a Q_{a,o}\big(\bar{s}_{a,o}(\mathbf{d'}, Q) - \bar{s}_{a,o}(\mathbf{d'}, P)\big),\ \forall \mathbf{d}, \mathbf{d'} \in d(\cdot), Q, P \in \mathcal{P}\\ &\sum_{a,o} Q_{a,o}\big(\mathbf{d}_a \bar{s}_{a,o}(\mathbf{d}, Q) - \mathbf{d}'_a \bar{s}_{a,o}(\mathbf{d'}, Q)\big) = \sum_{a,o} Q_{a,o}\big(\mathbf{d}_a \bar{s}_{a,o}(\mathbf{d}, P) - \mathbf{d}'_a \bar{s}_{a,o}(\mathbf{d'}, P)\big).\end{aligned}$$

Assuming, for a contradiction, that Equation 8 is false and there exists no such function $c_{\bar{\mathbf{d}}}$. Then there exists P and P', decision policies $\mathbf{d}$ and $\mathbf{d}'$ and an action a and outcome o such that

$$\mathbf{d}_a \bar{s}_{a,o}(\mathbf{d}, P) - \mathbf{d}'_a \bar{s}_{a,o}(\mathbf{d}', P) \neq \mathbf{d}_a \bar{s}_{a,o}(\mathbf{d}, P') - \mathbf{d}'_a \bar{s}_{a,o}(\mathbf{d}', P')$$

implying there exists a linear combination (belief) such that

$$\sum_{a,o} Q_{a,o}\big(\mathbf{d}_a \bar{s}_{a,o}(\mathbf{d}, P) - \mathbf{d}'_a \bar{s}_{a,o}(\mathbf{d}', P)\big) \neq \sum_{a,o} Q_{a,o}\big(\mathbf{d}_a \bar{s}_{a,o}(\mathbf{d}, P') - \mathbf{d}'_a \bar{s}_{a,o}(\mathbf{d}', P')\big)$$

but from strict properness *for a market* it must be that

$$\sum_{a,o} Q_{a,o}\big(\mathbf{d}_a \bar{s}_{a,o}(\mathbf{d}, Q) - \mathbf{d}'_a \bar{s}_{a,o}(\mathbf{d}', Q)\big) = \sum_{a,o} Q_{a,o}\big(\mathbf{d}_a \bar{s}_{a,o}(\mathbf{d}, P) - \mathbf{d}'_a \bar{s}_{a,o}(\mathbf{d}', P)\big),\ \forall P$$

which cannot be simultaneously true for both P and P', a contradiction as desired, so there exists a function satisfying Equation 8.

Equation 8 lets us define our new scoring rules as

$$s_{a,o}(\mathbf{P}, P) = \begin{cases} \bar{s}_{a,o}(\mathbf{d}, P) - c_{\bar{\mathbf{d}}}(a, o, \mathbf{d})/\mathbf{d}_a, \ \mathbf{d}_a > 0 \\ 0, \ \text{o.w.} \end{cases}$$

we need to prove (1) that using this scoring rule gives each expert the same expected net score, and (2) that it is strictly proper. Substituting our new rule for the old in the expected net score equation gives

$$\begin{aligned} & \sum_{a,o} \mathbf{d}_a Q_{a,o}\big(s_{a,o}(\mathbf{d}, P) - s_{a,o}(\mathbf{d}, P')\big) \\ = & \sum_{a,o} \mathbf{d}_a Q_{a,o}\big(\bar{s}_{a,o}(\mathbf{d}, P) - \bar{s}_{a,o}(\mathbf{d}, P')\big) \end{aligned}$$

so the two scoring rules have the same expected net score, as desired. We conclude by proving that (d, s) is a strictly proper pair.

First, the expected score is independent of the decision policy, since substitution gives

$$\sum_{a,o} \mathbf{d}_a Q_{a,o} s_{a,o}(\mathbf{d}, P) = \sum_{a,o} \bar{\mathbf{d}}_a Q_{a,o} \bar{s}_{a,o}(\bar{\mathbf{d}}, P),\ \forall \mathbf{d} \in d(\cdot), Q, P \in \mathcal{P}$$

so it's as if the decision policy is fixed. Second, the expected score is uniquely maximized when an expert predicts its beliefs since (again, from the definition of strict properness *for a market*)

$$\arg\max_{P \in \mathcal{P}} \sum_{a,o} \mathbf{d}_a Q_{a,o}\big(\bar{s}_{a,o}(\mathbf{d}, P) - \bar{s}_{a,o}(\mathbf{d}, P')\big) = \{Q\},\ Q, P' \in \mathcal{P}, \mathbf{d} \in d(\cdot)$$

and the expected net score is the same using the new scoring rule s, so

$$\arg\max_{P \in \mathcal{P}} \sum_{a,o} \mathbf{d}_a Q_{a,o}\big(s_{a,o}(\mathbf{d}, P) - s_{a,o}(\mathbf{d}, P')\big) = \{Q\},\ Q, P' \in \mathcal{P}, \mathbf{d} \in d(\cdot)$$

and the score of the previous prediction $Q_{a,o}\mathbf{d}_a s_{a,o}(\mathbf{d}, P')$ is independent of P, so equivalently the maximizing argument of the above expression is

$$\arg\max_{P \in \mathcal{P}} \sum_{a,o} \mathbf{d}_a Q_{a,o} s_{a,o}(\mathbf{d}, P) = \{Q\},\ Q \in \mathcal{P}, \mathbf{d} \in d(\cdot)$$

and the pair (d, s) is strictly proper. $\square$

7.3. Proof of Theorem 5.1

THEOREM 5.1. *For any pair (d, s) that is strictly proper* for an expert*, the set of action-outcome matrices that d maps to distributions without full support is nowhere dense in the set of all action-outcome matrices with their natural Eucdliean topology.*

PROOF. Let (d, s) be any pair that is strictly proper *for an expert* and assume, for a contradiction, that the set of matrices that d maps to distributions assigning zero probability to some action a', $\mathcal{P}^{a'} = \{P \in \mathcal{P} | d_{a'}(P) = 0\}$, is not nowhere dense in $\mathcal{P}$. The set of all matrices d maps to distributions without full support is a finite union of such sets (since we assumed in Section 3 the set of actions was finite), so proving $P^{a'}$ is nowhere dense in $\mathcal{P}$ for arbitrary a' proves the theorem.

We begin by finding two matrices in the interior of the closure of $\mathcal{P}^{a'}$ that differ only on a' and have different expected scores. The existence of points differing only on a' in the interior of the closure of $\mathcal{P}^{a'}$ is immediate from our assumption that the closure has nonempty interior, while the expected score of a prediction given (d, s) can be written as a function

$$v(Q, P) = \sum_{a,o} d_a(P) Q_{a,o} s_{a,o}(d_a(P), P)$$

Let $\hat{P}$ and $\bar{P}$ be matrices differing only on a' in the interior of the closure of $\mathcal{P}^{a'}$, without loss of generality either $v(\hat{P}, \hat{P}) > v(\bar{P}, \bar{P})$ or $v(\hat{P}, \hat{P}) = v(\bar{P}, \bar{P})$. In the former case we're done so we assume, for a contradiction, the latter is true. Now consider the difference

$$\begin{aligned} & v(\hat{P}, \bar{P}) - v(\bar{P}, \bar{P}) \\ = & \sum_{a,o} d_a(\bar{P}) \hat{P}_{a,o} s_{a,o}(d_a(\bar{P}), \bar{P}) - \sum_{a,o} d_a(\bar{P}), \bar{P}_{a,o} s_{a,o}(d_a(\bar{P}), \bar{P}) \\ = & \sum_{o} \big(\hat{P}_{a',o} - \bar{P}_{a',o}\big) d_{a'}(\bar{P}) s_{a',o}(d_a(\bar{P}), \bar{P}) \end{aligned}$$

This is a difference of two convex combinations of the terms $d_{a'}(\bar{P}) s_{a',o}(d_a(\bar{P}), \bar{P})$. If these terms are not all equal then, because of our assumption that $\bar{P}$ is in the interior of the closure of $\mathcal{P}^{a'}$ there exists a P^* also differing from $\bar{P}$ only on action a' with $v(P^*, \bar{P}) > v(\bar{P}, \bar{P})$, implying $v(P^*, P^*) > v(\bar{P}, \bar{P})$ too because we assumed (d, s) was strictly proper *for an expert*. Alternatively all these terms are equal, $d_{a'}(\bar{P}) s_{a',o'}(d_a(\bar{P}), \bar{P}) = d_{a'}(\bar{P}) s_{a',\bar{o}}(d_a(\bar{P}), \bar{P})$, $\forall o', \bar{o} \in \mathcal{O}$, and the argument applies symmetrically to $\hat{P}$ so we assume its corresponding terms are equal, and immediately from the difference and the relationship $v(\hat{P}, \hat{P}) = v(\bar{P}, \bar{P})$ the terms for both predictions must all equal each other, too. This is, however, a contradiction, since it implies $v(\hat{P}, \bar{P}) = v(\bar{P}, \bar{P}) = v(\hat{P}, \hat{P})$ which contradicts our assumption that (d, s) is strictly proper *for an expert*. Thus we conclude there exist points $\hat{P}$ and $\bar{P}$ in the interior of the closure of $\mathcal{P}^{a'}$ differing only on action a' with different expected scores, $v(\hat{P}, \hat{P}) > v(\bar{P}, \bar{P})$.

Next we show the function $v(P) = v(P, P)$ is continuous. Assume, for a contradiction and without loss of generality

$$\lim_{Q \to P} v(Q) < v(P)$$

But the difference

$$\begin{aligned}
& \lim_{Q \to P} v(Q,P) - v(P,P) \\
= & \sum_{a,o} \big(Q_{a,o} - P_{a,o}\big) d_a(P) s_{a,o}(d_a(P), P) \\
= & \ 0
\end{aligned}$$

implying (since we assumed (d, s) was strictly proper *for an expert*) that $\lim_{Q \to P} v(Q) \geq v(P)$, a contradiction. Alternatively

$$\lim_{Q \to P} v(Q) > v(P)$$

and we have a symmetric argument, so the function $v(P) = v(P, P)$ is continuous.

Now let $\{\hat{Q}^k\}$ and $\{\bar{Q}^k\}$ be sequences in $\mathcal{P}^{a'}$ converging to $\hat{P}$ and $\bar{P}$ respectively, and consider the difference of the limits

$$\begin{aligned}
& \lim_{k \to \infty} v(\hat{P}, \hat{P}) - v(\bar{P}, \hat{Q}^k) \\
= & \lim_{k \to \infty} v(\hat{Q}^k, \hat{Q}^k) - v(\bar{P}, \hat{Q}^k) && \text{(by continuity)} \\
= & \lim_{k \to \infty} \sum_{a,o} d_a(\hat{Q}^k)\big(\hat{Q}^k_{a,o} - \bar{P}^k_{a,o}\big) s_{a,o}(\hat{Q}^k) \\
= & \lim_{k \to \infty} d_{a'}(\hat{Q}^k)\big(\hat{P}_{a',o} - \bar{P}_{a',o}\big) s_{a',o}(\hat{Q}^k) \\
= & \ 0 && (\text{since } d_{a'}(\hat{Q}^k) = 0, \ \forall k \text{ because } \hat{Q}^k \text{ is always in } \mathcal{P}^{a'})
\end{aligned}$$

Recalling that we assumed (without loss of generality) that $v(\hat{P}, \hat{P}) > v(\bar{P}, \bar{P})$, this implies there exists a prediction $P^* \in \{\hat{Q}^k\}$ such that $v(\bar{P}, P^*) > v(\bar{P}, \bar{P})$ which, again, contradicts our assumption that (d, s) is strictly proper *for an expert*. We conclude $\mathcal{P}^{a'}$ must be nowhere dense in $\mathcal{P}$, and since a' was chosen arbitrarily the set of all matrices d maps to distributions with full support is also nowhere dense in $\mathcal{P}$. $\square$

7.4. Proof of Theorem 5.2

THEOREM 5.2. *A pair (d, s) is strictly proper* for an expert *if and only if there exist a strictly convex function g and a subgradient $g^\star(P)$ where $g^\star_a(P) = \mathbf{0}$ whenever $d_a(P) = 0$ such that*

$$s_{a,o}(d(P), P) = g(P) - g^\star(P) : P + \frac{g^\star_{a,o}(P)}{d_a(P)}, \ \forall d_a(P) > 0. \tag{5}$$

PROOF. Given a decision rule d, strictly convex function g, and subgradients $g^\star$ of g, with $g^\star_a(P) = \vec{0}$ whenever $d_a(P) = 0$, we prove that the pair (d, s) with s written as in

Equation (5) whenever $d_a(P) > 0$ is strictly proper *for an expert*.

$$\begin{aligned}
&\sum_{a,o} d_a(P) Q_{a,o} s_{a,o}(P) \\
=& \sum_{a \in \{a|d_a(P)>0\},o} d_a(P) Q_{a,o}\big(g(P) - g^\star(P) : P + \frac{g^\star_{a,o}(P)}{d_a(P)}\big) \\
=& \; g(P) - g^\star(P) : P + \sum_{a \in \{a|d_a>0\},o} g^\star_{a,o}(P) Q_{a,o} \\
=& \; g(P) - P : g^\star(P) + Q : g^\star(P) \qquad\qquad (\text{recalling } d_a(P) = 0 \implies g^\star_a(P) = \mathbf{0}) \\
=& \; g(P) + (Q - P) : g^\star(P)
\end{aligned}$$

Exactly as in Theorem 4.2 and the same argument concludes this direction of the proof.

Given a (d, s) pair that is strictly proper *for an expert*, we can construct a strictly convex function g and subgradients of g satisfying the theorem's criteria. As in Theorem 4.2 we consider the expected score function

$$v(Q, P) = \sum_{a,o} d_a(P) Q_{a,o} s_{a,o}(d(P), P),$$

and let $g(P) = v(P, P)$. Define the subgradients of $g(P)$ as

$$g^*_{a,o}(P) = d_a(P) s_{a,o}(d(P), P).$$

We can verify that they satisfy the subgradient inequality:

$$\begin{aligned}
& g(P) + (Q - P) : g^\star(P) \\
&= \sum_{a,o} (P_{a,o} + Q_{a,o} - P_{a,o}) d_a(P) s_{a,o}(d(P), P) \\
&= \sum_{a,o} Q_{a,o} d_a(P) s_{a,o}(d(P), P) \\
&= V(Q, P) < V(Q, Q) = g(Q)
\end{aligned}$$

for any $Q \neq P$, where the inequality holds because (d, s) is strictly proper *for an expert*. This also implies g is strictly convex. Note also that $d_a(P) = 0 \implies g^\star_a(P) = \vec{0}$ for all P, so it only remains to show that s can be written as in Equation (5) using g for any $a \in \{a|d_a(P) > 0\}$:

$$\begin{aligned}
& g(P) - g^\star(P) : P + \frac{g^\star_{a,o}(P)}{d_a(P)} \\
=& \; \frac{g^\star_{a,o}(P)}{d_a(P)} \\
=& \; \frac{d_a(P) s_{a,o}(d(P), P)}{d_a(P)} = s_{a,o}(d(P), P).
\end{aligned}$$

□