

General Truthfulness Characterizations Via Convex Analysis

Rafael Frongillo and Ian Kash

Microsoft Research

Abstract. We present a model of truthful elicitation which generalizes and extends mechanisms, scoring rules, and a number of related settings that do not quite qualify as one or the other. Our main result is a characterization theorem, yielding characterizations for all of these settings, including a new characterization of scoring rules for non-convex sets of distributions. We generalize this model to eliciting some property of the agent’s private information, and provide the first general characterization for this setting. We also show how this yields a new proof of a result in mechanism design due to Saks and Yu.

1 Introduction

We examine a general model of information elicitation where a single agent is endowed with some type t that is private information and is asked to reveal it. After doing so, he receives a score $A(t', t)$ that depends on both his report t' and his true type t . We allow A to be quite general, with the main requirement being that $A(t', \cdot)$ is an affine¹ function of the true type t , and seek to understand when it is optimal for the agent to truthfully report his type. Given this truthfulness condition, it is immediately clear why convexity plays a central role—when an agent’s type is t , the score for telling the truth is $A(t, t) = \sup_{t'} A(t', t)$, which is a convex function of t as the pointwise supremum of affine functions.

One special case of our model is mechanism design with a single agent², where the designer wishes to select an outcome based on the agent’s type. In this setting, $A(t', \cdot)$ can be thought of as the allocation and payment given a report of t' , which combine to determine the utility of the agent as a function of his type. In this context, $A(t, t)$ is the consumer surplus function (or indirect utility function), and Myerson’s well-known characterization [50] states that, in single-parameter settings, a mechanism is truthful if and only if the consumer surplus function is convex and its derivative (or subgradient at points of non-differentiability) is the allocation rule. More generally, this remains true in higher

¹ A mapping between two vector spaces is affine if it consists of a linear transformation followed by a translation.

² This is not a real restriction because notions of truthfulness are phrased in terms of holding the behavior of other agents constant. See [4, 23] for additional discussion.

dimensions (see [4]).³

Another special case is a scoring rule, also called a *proper loss* in the machine learning literature, where an agent is asked to predict the distribution of a random variable and given a score based on the observed realization of that variable. In this setting, types are distributions over outcomes, and $A(t', t)$ is the agent's subjective expected score for a report that the distribution is t' when he believes the distribution is t . As an expectation, this score is linear in the agent's type. Gneiting and Raftery [33] unified and generalized existing results in the scoring rules literature by characterizing proper scoring rules in terms of convex functions and their subgradients.

Further, the generality of our model allows it to include settings that do not quite fit into the standard formulations of mechanisms or scoring rules. These include counterfactual scoring rules for decision-making [20, 21, 56], proper losses for machine learning with partial labels [24], mechanism design with partial allocations [17], and responsive lotteries [27].

In many settings, it is difficult, or even impossible, to have agents report an entire type $t \in \mathcal{T}$. For example, when allocating a divisible good (e.g. water), a mechanism that needs to know how much an agent would value each possible allocation requires him to submit an infinite-dimensional type. Even type spaces which are exponential in size, such as those that arise in combinatorial auctions, can be problematic from an algorithmic perspective. Moreover, in many situations, the principal is *uninterested* in all but some small aspect of an agent's private type. For example, the information is often to be used to eventually make a specific decision, and hence only the information directly pertaining to the decision is actually needed—why ask for the agent's entire probability distribution of rainfall tomorrow if a principal wanting to choose between {umbrella, no umbrella} would be content with its expected value, or even just whether she should carry an umbrella or not?

It is therefore natural to consider an indirect elicitation model where agents provide some sort of summary information about their type. Such a model has been studied in the scoring rules literature, where one wishes to elicit some statistic, or *property*, of a distribution, such as the mean or quantile [32, 47, 54, 61]. We follow this line of research, and extend the affine score framework to accept reports from a different (intuitively, much smaller) space than \mathcal{T} .

1.1 Our Contribution

Our main theorem (Section 2) is a general characterization theorem that generalizes and extends known characterization theorems for proper scoring rules (substantially) and truthful mechanisms (slightly, by removing a technical assumption). For scoring rules, this provides the first characterization of proper

³ Note that here the restriction that $A(t', \cdot)$ be affine is without loss of generality, because we view types as functions and function application is a linear operation. (See Section 2.2 for more details.)

scoring rules with non-convex sets of distributions, an idea that has proved useful as a way of separating informed and uninformed experts [9,26], but for which no characterization was known. We also survey applications to related settings and show our theorem can be used to provide characterizations for them as well, including new results about mechanism design with partial allocation and responsive lotteries. Thus, our theorem eliminates the need to independently derive characterizations for such settings.

This unified characterization of mechanisms and scoring rules also clarifies their relationship: both are derived from convex functions in the exact same manner, with mechanisms merely facing additional constraints on the choice of convex function so that it yields a valid allocation rule. This aids in understanding when results or techniques from one setting can be applied in the other. Indeed, the proof of our characterization begins with Gneiting and Raftery’s scoring rule construction [33] and adapts it with a variant of a technique from Archer and Kleinberg [4] for handling mechanisms with non-convex type spaces (see their Theorem 6.1). As an example of the new insights this can provide, results from mechanism design show that a scoring rule is proper if and only if it is locally proper (see Section C.2 and Corollary 3). More broadly, we show how results from mechanism design about implementability and revenue equivalence generalize to our framework.

We then move on to two general results for eliciting a particular property of the agent’s private information. The first is essentially a direct generalization of Theorem 1, which keeps the same general structure but adds the constraint that the convex function must be flat on sets of types which share an optimal report. This is the first general result for arbitrary properties; in addition to serving as our main tool to derive the remainder of our results, this theorem provides several ways to show that a property is not elicitable (by showing that no such convex function can exist). The second result is a transformation of this theorem using *duality*, which shows that there is a strong sense in which properties *are* subgradients of convex functions. We use this result to introduce notions of dual properties and scores, which gives a new construction to convert between scoring rules and randomized mechanisms (see Corollary 9).

We conclude by examining properties that take on a finite number of values, which Lambert and Shoham [48] showed correspond to *power diagrams*. We extend their result to settings where the private information need not be a probability distribution, and give a tight characterization for a particular restricted “simple” case. We also give an explicit construction for generating power diagrams from other measures of distances via a connection to *Bregman Voronoi diagrams* [15]. Finally, we show how these results imply a new proof of an implementability theorem from mechanism design due to Saks and Yu [60].

1.2 Relation to Prior Work

The similarities between mechanisms and scoring rules were noted by (among others) Fiat et al. [29], who gave a construction to convert mechanisms into scoring rules and vice versa, and Feige and Tennenholtz [27], who gave techniques to

convert both to “responsive lotteries.” Further, techniques from convex analysis have a long history in the analysis of both models (see [33,65]). However, we believe that our results use the “right” representation and techniques, which leads to more elegant characterizations and arguments. For example, the construction used by Fiat et al. has the somewhat awkward property that the scoring rule corresponding to a mechanism has one more outcome than the mechanism did, a complication absent from our results. Similarly, the constructions used by Feige and Tennenholtz only handle special cases and they claim “there is no immediate equivalence between lottery rules and scoring rules,” while we can give such an equivalence. So while prior work has understood that there is a connection, the nature of that connection has been far from clear.

A large literature in mechanism design has explored characterizations of when allocation rules can be truthfully implemented; see e.g. [4, 5, 14, 19, 40–42, 49, 51, 60]. Similarly, work on revenue equivalence can be cast in our framework as well [18, 37, 44, 50]. For scoring rules, our work connects to a literature that has used non-convex sets of probability distributions to separate (usefully) informed experts from uninformed experts [9, 26].

Indirect elicitation has a long history in the scoring rules literature, starting with Savage [61]. While the bulk of the literature focuses on specific statistics, such as means and quantiles [32, 33, 35, 54], Osband [55] and Lambert, Pennock, and Shoham [47] first considered the problem of eliciting a more general property Γ . Several authors have made significant contributions toward the general problem for the case where Γ is real-valued [32, 34, 46, 47, 62] and vector-valued [30, 47, 55], but our results are the first for arbitrary multivalued maps. Mechanisms that elicit a ranking over outcomes rather than a utility for each outcome (common in, e.g., matching contexts) are a form of property elicitation, and our results are related to characterizations due to Carroll [19].

Notation. We define $\overline{\mathbb{R}} = \mathbb{R} \cup \{-\infty, \infty\}$ to be the extended real numbers. Given a set of measures M on X , a function $f : X \rightarrow \overline{\mathbb{R}}$ is M -quasi-integrable if $\int_X f(x)d\mu(x) \in \overline{\mathbb{R}}$ for all $\mu \in M$. Let $\Delta(X)$ be the set of all probability measures on X . We denote by $\text{Aff}(X \rightarrow Y)$ and $\text{Lin}(X \rightarrow Y)$ the set of functions from X to Y which are affine and linear, respectively. We write $\text{Conv}(X)$ to denote the convex hull of a set of vectors X , the set of all (finite) convex combinations of elements of X . Some useful facts from convex analysis are collected in §A of the full version [31].

2 Affine Scores

We consider a very general model with an agent who has a given type $t \in \mathcal{T}$ and reports some possibly distinct type $t' \in \mathcal{T}$, at which point the agent is rewarded according to some score $A(t', t)$ which is affine in the true type t . This reward we call an affine score. We wish to characterize all *truthful* affine scores, those which incentivize the agent to report her true type t .

Definition 1. Any function $A : \mathcal{T} \times \mathcal{T} \rightarrow \overline{\mathbb{R}}$, where $\mathcal{T} \subseteq \mathcal{V}$ for some vector space \mathcal{V} over \mathbb{R} and $\mathcal{A} \doteq \{A(t, \cdot) \mid t \in \mathcal{T}\} \subseteq \text{Aff}(\mathcal{T} \rightarrow \overline{\mathbb{R}})$, is a affine score with score set \mathcal{A} . We say A is truthful if for all $t, t' \in \mathcal{T}$,

$$A(t', t) \leq A(t, t). \quad (1)$$

If this inequality is strict for all $t \neq t'$, then A is strictly truthful.

Our characterization uses convex analysis, a central concept of which is the subgradient of a function, which is a generalization of the gradient yielding a linear approximation that is always below the function.

Definition 2. Given some function $G : \mathcal{T} \rightarrow \mathbb{R}$, a function $d \in \text{Lin}(\mathcal{V} \rightarrow \overline{\mathbb{R}})$ is a subgradient to G at t if for all $t' \in \mathcal{T}$,

$$G(t') \geq G(t) + d(t' - t). \quad (2)$$

We denote by $\partial G : \mathcal{T} \rightrightarrows \text{Lin}(\mathcal{V} \rightarrow \overline{\mathbb{R}})$ the multivalued map such that $\partial G(t)$ is the set of subgradients to G at t . We say a parameterized family of linear functions $\{d_t \in \text{Lin}(\mathcal{V} \rightarrow \overline{\mathbb{R}})\}_{t \in \mathcal{T}'}$ for $\mathcal{T}' \subseteq \mathcal{T}$ is a selection of subgradients if $d_t \in \partial G(t)$ for all $t \in \mathcal{T}'$; we denote this succinctly by $\{d_t\}_{t \in \mathcal{T}'} \in \partial G$.

For mechanism design, it is typical to assume that utilities are always real-valued. However, the log scoring rule (one of the most popular scoring rules) has the property that if an agent reports that an event has probability 0, and then that event does occur, the agent receives a score of $-\infty$. Essentially solely to accommodate this, we allow affine scores and subgradients to take on values from the extended reals. In the next paragraph we provide the relevant definitions, but for most purposes it suffices to ignore these and simply assume that all affine scores are real-valued.

It is standard (cf. [33]) to restrict consideration to the “regular” case, where intuitively only things like the log score are permitted to be infinite. In particular, an affine score A is *regular* if $A(t, t) \in \mathbb{R}$ for all $t \in \mathcal{T}$, and $A(t', t) \in \mathbb{R} \cup \{-\infty\}$ for $t' \neq t$. Similarly, a parameterized family of linear functions (e.g. subgradients) $\{d_t \in \text{Lin}(\mathcal{V} \rightarrow \overline{\mathbb{R}})\}_{t \in \mathcal{T}}$ is *\mathcal{T} -regular* if $d_t(t) \in \mathbb{R}$ for all $t \in \mathcal{T}$, and $d_{t'}(t) \in \mathbb{R} \cup \{-\infty\}$ for $t' \neq t$.⁴ Likewise, *\mathcal{T} -regular* affine functions have *\mathcal{T} -regular* linear parts with finite constants (i.e. we exclude the constant functions $\pm\infty$). For the remainder of the paper we assume all affine scores and parameterized families of linear or affine functions are *\mathcal{T} -regular*, where \mathcal{T} will be clear from context.

We now state our characterization theorem. The proof takes Gneiting and Raftery’s [33] proof for the case of scoring rules on convex domains and extends it to the non-convex case using a variant of a technique Archer and Kleinberg [4] introduced for mechanisms with non-convex type spaces. This technique is essentially that used in prior work on extensions of convex functions [57, 66].

⁴ To define linear functions to $\overline{\mathbb{R}}$, we adopt the convention $0 \cdot \infty = 0 \cdot (-\infty) = 0$. Thus, any $l \in \text{Lin}(\mathcal{V} \rightarrow \overline{\mathbb{R}})$ can be written as $l_1 + \infty \cdot l_2$ for some $l_1, l_2 \in \text{Lin}(\mathcal{V} \rightarrow \mathbb{R})$.

Theorem 1. *Let an affine score $A : \mathcal{T} \times \mathcal{T} \rightarrow \overline{\mathbb{R}}$ with score set \mathcal{A} be given. A is truthful if and only if there exists some convex $G : \text{Conv}(\mathcal{T}) \rightarrow \overline{\mathbb{R}}$ with $G(\mathcal{T}) \subseteq \mathbb{R}$, and some selection of subgradients $\{d_t\}_{t \in \mathcal{T}} \in \partial G$, such that*

$$A(t', t) = G(t') + d_{t'}(t - t'). \quad (3)$$

In the remainder of this section, we show how scoring rules, mechanisms, and other related models fit comfortably within our framework.

2.1 Scoring rules for non-convex \mathcal{P}

In this section, we show that the Gneiting and Raftery characterization is a simple special case of Theorem 1, and moreover that we *generalize* their result to the case where the set of distributions \mathcal{P} may be non-convex. We also give a result about local properness derived using tools from mechanism design in the full version [31, § C.2]. To begin, we formally introduce scoring rules and show that they fit into our framework. The goal of a scoring rule is to incentivize an expert who knows a probability distribution to reveal it to a principal who can only observe a single sample from that distribution.

Definition 3. *Given outcome space \mathcal{O} and set of probability measures $\mathcal{P} \subseteq \Delta(\mathcal{O})$, a scoring rule is a function $S : \mathcal{P} \times \mathcal{O} \rightarrow \overline{\mathbb{R}}$. We say S is proper if for all $p, q \in \mathcal{P}$,*

$$\mathbb{E}_{o \sim p}[S(q, o)] \leq \mathbb{E}_{o \sim p}[S(p, o)]. \quad (4)$$

If the inequality in (4) is strict for all $q \neq p$, then S is strictly proper.

To incorporate this into our framework, take the type space $\mathcal{T} = \mathcal{P}$. Thus, we need only construct the correct score set \mathcal{A} of affine functions available to the scoring rule as payoff functions. Intuitively, these are the functions that describe what payment the expert receives given each outcome, but we have a technical requirement that the expert's expected utility be well defined. Thus, following Gneiting and Raftery, we take \mathcal{F} to be the set of \mathcal{P} -quasi-integrable⁵ functions $f : \mathcal{O} \rightarrow \overline{\mathbb{R}}$, and the score set $\mathcal{A} = \{p \mapsto \int_{\mathcal{O}} f(o) dp(o) \mid f \in \mathcal{F}\}$.

We now apply Theorem 1 for our choice of \mathcal{T} and \mathcal{A} , which yields the following generalization of Gneiting and Raftery [33].

Corollary 1. *For an arbitrary set $\mathcal{P} \subseteq \Delta(\mathcal{O})$ of probability measures, a regular⁶ scoring rule $S : \mathcal{P} \times \mathcal{O} \rightarrow \overline{\mathbb{R}}$ is proper if and only if there exists a convex function $G : \text{Conv}(\mathcal{P}) \rightarrow \overline{\mathbb{R}}$ with functions $G_p \in \mathcal{F}$ such that*

$$S(p, o) = G(p) + G_p(o) - \int_{\mathcal{O}} G_p(o) dp(o), \quad (5)$$

where $G_p : q \mapsto \int_{\mathcal{O}} G_p(o) dq(o)$ is a subgradient of G for all $p \in \mathcal{P}$.

⁵ We say that $f : \mathcal{O} \rightarrow \overline{\mathbb{R}}$ is \mathcal{P} -quasi-integrable if $\int_{\mathcal{O}} f(o) dp(o) \in \overline{\mathbb{R}}$ for all $p \in \mathcal{P}$.

⁶ This is the same concept as with affine scores: scores cannot be ∞ and only incorrect reports can yield $-\infty$.

Importantly, Corollary 1 immediately generalizes [33] to the case where \mathcal{P} is not convex, which is new to the scoring rules literature. One direction of this extension is obvious (if S is truthful on the convex hull of a set then it is truthful on that set). The other is not, however, and is an important negative result, ruling out the possibility of new scoring rules arising by restricting the set of distributions (provided the restriction does not alter the convex hull).

In the absence of a characterization, several authors have worked in the non-convex \mathcal{P} case. For example, Babaioff et al. [9] examine when proper scoring rules can have the additional property that uninformed experts do not wish to make a report (have a negative expected utility), while informed experts do wish to make one. They show that this is possible in some settings where the space of reports is not convex. Our characterization shows that, despite not needing to ensure properness on reports outside \mathcal{P} , essentially the only possible scoring rules are still those that are proper on all of $\Delta(\mathcal{O})$. We state the simplest version of such a characterization, for perfectly informed experts, here.

Corollary 2. *Let a non-convex set $\mathcal{P} \subseteq \Delta(\mathcal{O})$ and $\bar{p} \in \Delta(\mathcal{O}) - \mathcal{P}$ be given. A scoring rule S is proper and guarantees that experts with a belief in \mathcal{P} receive a score of at least δ_A while experts with a belief of \bar{p} receive a score of at most δ_R if and only if S is of the form (5) with $G(p) \geq \delta_A \forall p \in \mathcal{P}$ and $G(\bar{p}) \leq \delta_R$.*

With a similar goal to Babaioff et al., Fang et al. [26] find conditions on \mathcal{P} for which every continuous “value function” $G : p \mapsto S(p, p)$ on \mathcal{P} can be attained by some S with the motivation of eliciting the expert’s information when it is known to come from some family of distributions (which in general will not be a convex set). As such, they provide sufficient conditions on particular non-convex sets, as opposed to our result which provides necessary and sufficient conditions for all non-convex sets. Beyond these specific applications, our characterization is useful for answering practical questions about scoring rules. For example, suppose we assume that people have beliefs about probabilities in increments of 0.01. Does that change the set of possible scoring rules? No. What happens if they have finer-grained beliefs but we restrict them to such reports? They will end up picking a “nearby” report (see the discussion of convexity in Section 3.2).

In the full version [31, § C.2], we show how local truthfulness conditions from mechanism design, where one verifies that an affine score is truthful by checking that it is truthful in a small neighborhood around every point, generalize to our framework. In particular, Corollary 11 shows that local properness (i.e. properness for distributions in a neighborhood) is equivalent to global properness for scoring rules on convex \mathcal{P} , an observation that is also new to the scoring rules literature. See [31, § C.2] for the precise meaning of (weak) local properness.

Corollary 3. *For a convex set $\mathcal{P} \subseteq \Delta(\mathcal{O})$ of probability measures, a scoring rule $S : \mathcal{P} \times \mathcal{O} \rightarrow \bar{\mathbb{R}}$ is proper if and only if it is (weakly) locally proper.*

2.2 Mechanism design

We now show how to view a mechanism as an affine score. First, we formally introduce mechanisms in the single agent case (see below for remarks about

multiple agents). Then we show how known characterizations of truthful mechanisms follow easily from our main theorem. This allows us to relax a minor technical assumption from the most general such theorem.

Definition 4. *Given outcome space \mathcal{O} and a type space $\mathcal{T} \subseteq (\mathcal{O} \rightarrow \mathbb{R})$, consisting of functions mapping outcomes to reals, a (direct) mechanism is a pair (f, p) where $f : \mathcal{T} \rightarrow \mathcal{O}$ is an allocation rule and $p : \mathcal{T} \rightarrow \mathbb{R}$ is a payment. The utility of the agent with type t and report t' to the mechanism is $U(t', t) = t(f(t')) - p(t')$; we say the mechanism (f, p) is truthful if $U(t', t) \leq U(t, t)$ for all $t, t' \in \mathcal{T}$.*

Here we suppose that the mechanism can choose an allocation from some set \mathcal{O} of outcomes, and there is a single agent whose type $t \in \mathcal{T}$ is itself the valuation function. That is, the agent's net utility upon allocation o and payment p is $t(o) - p$. Thus, following Archer and Kleinberg [4], we view the type space \mathcal{T} as lying in the vector space $\mathcal{V} = \mathbb{R}^{\mathcal{O}}$. The advantage of this representation is that while agent valuations in mechanism design can generally be complicated functions, viewed this way they are all linear: for any $v_1, v_2 \in \mathcal{V}$, we have $(v_1 + \alpha v_2)(o) = v_1(o) + \alpha v_2(o)$. Thus, we have an affine score $A(t', t) \doteq U(t', t)$, with score set $\mathcal{A} = \{t \mapsto t(o) + c \mid o \in \mathcal{O}, c \in \mathbb{R}\}$, so that every combination of outcome and payment a mechanism could choose is an element of \mathcal{A} .

As an illustration of our theorem, consider the following characterization, due to Myerson [50], for a single parameter setting (i.e. when the agent's type can be described by a single real number). The result states that an allocation rule is implementable, meaning there is some payment rule making it truthful, if and only if it is *monotone* in the agent's type.

Corollary 4 (Myerson [50]). *Let $\mathcal{T} = \mathbb{R}_+$, $\mathcal{O} \subseteq \mathbb{R}$, so that the agent's valuation is $t \cdot o$. Then a mechanism f, p is truthful if and only if: (i) f is monotone non-decreasing in t , and (ii) $p(t) = tf(t) - \int_0^t f(t')dt' + p_0$.*

More generally, applying our theorem gives the following characterization. It is essentially equivalent to that of Archer and Kleinberg [4] (their Theorem 6.1), although our approach allows the relaxation of a technical assumption their version requires when the set of types is non-convex.

Corollary 5. *A mechanism f, p is truthful if and only if there exists a convex function $G : \text{Conv}(\mathcal{T}) \rightarrow \mathbb{R}$ and some selection of subgradients $\{dG_t\}_{t \in \mathcal{T}}$, such that for all $t \in \mathcal{T}$, $f(t) = dG_t$ and $G(t) = t(f(t)) - p(t)$*

Of course, mechanism design asks many questions beyond whether a particular mechanism is truthful, and some of these can be reframed as questions in convex analysis. The study of *implementability* focuses on the question of when there exist payments that make a given allocation rule truthful, whereas *revenue equivalence* asks when all mechanisms with a given allocation rule charge the same prices (up to a constant). By focusing on the subgradient, we recover and extend previous results for both, which we detail in the full version [31, § C, § D].

2.3 Other Applications

There are a number of other application domains that are not quite mechanisms or scoring rules, yet for which our main theorem yields characterization theorems. In the full version [31], we survey four such domains where our theorem could have directly provided the characterization ultimately used rather than requiring effort to conceptualize and prove it. We summarize two of the four below; the others are decision rules [20,21,56] and machine learning with partial labels [24]).

Mechanism design with partial allocation. Cai, Mahdian, Mehta, and Waggoner [17], consider a setting where the mechanism designer wants to elicit two pieces of information: the agent’s (expected) value for an item in an auction and the probability distribution of a random variable conditional on that agent winning, with the goal of understanding how the organizer of a daily deal site can take into account the value that will be created for users (as opposed to just the advertiser) when a particular deal is chosen to be advertised (e.g. the site operator may prefer deals that sell to many users over equally profitable deals that sell only to a few because this keeps users interested for future days). Our approach allows us to provide a characterization (given in [31, §E.3]), of a more general setting where a mechanism designer wishes to elicit two pieces of information, but the second need not be restricted to probability distributions. For example the mechanism designer could have two distinct sets of goods to allocate and want to design a truthful mechanism that is consistent with a partial allocation rule that determines how the primary goods should be allocated given the agent’s preferences over both types of goods.

Responsive lotteries. Feige and Tennenholtz [27] study the problem of how an agent can be incentivized to indirectly reveal his utility function over outcomes by being given a choice of lotteries over those outcomes, an approach with applications to experimental psychology, market research, and multiagent mechanism design. They give a geometric description of how such lotteries can be created with a finite set of outcomes. Our approach allows us to give a complete characterization, which highlights the relationship between natural desiderata and underlying geometric properties of the set of possible lotteries: strict truthfulness and continuity of the lottery rule jointly correspond to strict convexity of the lottery set, and uniqueness of the utility given the optimal lottery corresponds to smoothness of the boundary.

3 Property Elicitation

We wish to generalize the notion of truthful elicitation from eliciting private information from some set \mathcal{T} to accept reports from a space \mathcal{R} which is different from \mathcal{T} . To even discuss truthfulness in this setting, we need a notion of a truthful report r for a given type t . We encapsulate this notion by a general multivalued map which specifies all (and only) the correct values for t .

3.1 Affine Scores for Properties

Definition 5. Let \mathcal{T} be a give type space, where $\mathcal{T} \subseteq \mathcal{V}$ for some vector space \mathcal{V} over \mathbb{R} , and \mathcal{R} be some given report space. A property is a multivalued map $\Gamma : \mathcal{T} \rightrightarrows \mathcal{R}$ which associates a nonempty set of correct report values to each type. We let $\Gamma_r \doteq \{t \in \mathcal{T} \mid r \in \Gamma(t)\}$ denote the set of types t corresponding to report value r .

One can think of Γ_r as the “level set” of Γ corresponding to value r . This concept will be especially useful when we consider finite-valued properties in Section 5. A natural bookkeeping constraint to impose on these level sets is *non-redundancy*, meaning no property value r has a level set entirely contained in another.

We extend the notion of an affine score to this setting, where the report space is \mathcal{R} instead of \mathcal{T} itself. Note that the score set $\mathcal{A} = \{A(r, \cdot) \mid r \in \mathcal{R}\}$ is still a subset of $\text{Aff}(\mathcal{V} \rightarrow \overline{\mathbb{R}})$.

Definition 6. An affine score $A : \mathcal{R} \times \mathcal{T} \rightarrow \overline{\mathbb{R}}$ elicits a property $\Gamma : \mathcal{T} \rightrightarrows \mathcal{R}$ if for all t ,

$$\Gamma(t) = \operatorname{argsup}_{r \in \mathcal{R}} A(r, t). \quad (6)$$

If we merely have $\Gamma(t) \subseteq \operatorname{argsup}_{r \in \mathcal{R}} A(r, t)$, we say A weakly elicits Γ . Property $\Gamma : \mathcal{T} \rightrightarrows \mathcal{R}$ is elicitable if some affine score $A : \mathcal{R} \times \mathcal{T} \rightarrow \overline{\mathbb{R}}$ elicits Γ .

Note that it is certainly possible to write down A such the argsup in (6) is not well defined. This corresponds to some types not having an optimal report, which we view as violating a minimal requirement for a sensible affine score. Thus, in order for A to be an affine score, we require (6) to be well defined for all $t \in \mathcal{T}$.

We now state our property characterization theorem, proved in the full version [31, §F], which in essence says that eliciting a property Γ is equivalent to eliciting subgradients of a convex function G . Intuitively, by truthfulness the linear part of $A(r, \cdot)$ must be a subgradient of G at all $t \in \Gamma_r$. We show that this is equivalent to G being flat along Γ_r , meaning we can calculate G on Γ_r by picking any $t_r \in \Gamma_r$ and following the subgradient. Since all choices of t_r lead to the same value, we could just as easily ask for this subgradient $\varphi(r)$ to be reported directly. As subgradients are functions (in this case from \mathcal{T} to $\overline{\mathbb{R}}$), we use the curried notation $\varphi(r)(t)$ for the application of this function.

Theorem 2. Let non-redundant property $\Gamma : \mathcal{T} \rightrightarrows \mathcal{R}$ and Γ -regular⁷ affine score $A : \mathcal{R} \times \mathcal{T} \rightarrow \overline{\mathbb{R}}$ be given. Then A elicits Γ if and only if there exists some convex $G : \operatorname{Conv}(\mathcal{T}) \rightarrow \overline{\mathbb{R}}$ with $G(\mathcal{T}) \subseteq \mathbb{R}$, some $\mathcal{D} \subseteq \partial G$, and some bijection $\varphi : \mathcal{R} \rightarrow \mathcal{D}$ with $\Gamma(t) = \varphi^{-1}(\mathcal{D} \cap \partial G_t)$, such that for all $r \in \mathcal{R}$ and $t \in \mathcal{T}$,

$$A(r, t) = G(t_r) + \varphi(r)(t - t_r), \quad (7)$$

where $\{t_r\}_{r \in \mathcal{R}} \subseteq \mathcal{T}$ satisfies $r' \in \Gamma(t_{r'})$ for all r' .

⁷ This is defined similarly to regularity; see [31, §F].

3.2 What Properties Are Not Elicitable?

In the remainder of this section, we examine three features that subgradient mappings of convex functions possess and thus that the level sets of elicitable properties must possess.

Convexity. A well-known property of subgradient mappings is that their level sets are convex (for completeness, we provide a proof in the full version [31, § A]). In light of our characterizations, this fact about convex functions immediately applies to elicitable properties:

Corollary 6. *If $\Gamma : \mathcal{T} \rightrightarrows \mathcal{R}$ is elicitable, then Γ_r is convex for all r .*

To see this, just note that $\varphi(r) \in \partial G_t \cap \partial G_{t'}$ implies that $\varphi(r) \in \partial G_{\hat{t}}$ for all $\hat{t} = \alpha t + (1 - \alpha)t'$. Corollary 6 was previously known for special cases [47, 48], where it was used to show variance, skewness, and kurtosis are not elicitable, and was also known in mechanism design (i.e. the set of types for which a given (allocation, payment) pair is optimal is convex).

Cardinality. Combining Theorem 2 with the fact that finite-dimensional convex functions are differentiable almost everywhere (cf. [3, Thm 7.26]) yields the following corollary, which shows that elicitable properties have unique values almost everywhere.

Corollary 7. *Let $\Gamma : \mathcal{T} \rightrightarrows \mathcal{R}$ be an elicitable property with $\mathcal{T} \subseteq \mathcal{V} = \mathbb{R}^n$. If \mathcal{T} is of positive measure in $\text{Conv}(\mathcal{T})$, and Γ is non-redundant, then $|\Gamma(t)| = 1$ almost everywhere.*

Using an appropriate notion of “almost everywhere”, in some cases this holds in infinite-dimensional vector spaces as well (see e.g. [16, p. 195] and [3, p. 274]). One can use this fact to show that $\Gamma(p) = \{(a, b) : \int_a^b p(x)dx = 0.9\}$, the set of 90% confidence intervals for a distribution p , is not an elicitable property. This was previously only known for the case where p has finite support [47].

Topology. Combining Theorem 2 with a closure property of convex functions [59, Thm 24.4] yields the following.

Corollary 8. *Let $\Gamma : \mathcal{T} \rightrightarrows \mathcal{R}$ be an elicitable property with $\mathcal{T} \subseteq \mathcal{V} = \mathbb{R}^n$ convex that can be elicited by a closed, convex G . Then Γ_r is closed for all r .*

Requiring G to be closed is a technical issue regarding the boundary of \mathcal{T} , and is irrelevant for level sets in the relative interior. While [48] showed this for finite report spaces \mathcal{R} , this more general statement shows, for example, that if $\mathcal{T} = \mathbb{R}$ the property $\Gamma(t) = \text{floor}(t) = \max\{z \in \mathbb{Z} \mid z \leq t\}$ is not elicitable. More generally, this often provides a tool for showing that we cannot get around issues of cardinality by finding a tie-breaking rule to make the value unique.

Closure appears is a more delicate property to work with in infinite dimensions, but intuitive violations of it can still be used to show that properties are

not elicitable. As an illustration, we provide a direct proof that a property is not elicitable. We already saw that confidence intervals are not elicitable due to their cardinality, but a natural practical request would be for this “smallest” such interval. Can we elicit this, or even just the length of this interval? The following sketch shows we cannot. For probability distribution F represented by a CDF, let $\Gamma(F) = \inf\{b - a \mid F(b) - F(a) = 1\}$ be the property that is the length of the smallest interval of probability 1⁸. Consider the family of distributions defined by their pdfs as $f_c(x) = 1 - c$ for $0 \leq x \leq 1$ and $f_c(x) = c$ for $1 < x \leq 2$ with corresponding CDFs F_c . Note that for $0 < c < 1$, $\Gamma(F_c) = \{2\}$ but $\Gamma(F_0) = \{1\}$. Suppose we could elicit this with a scoring rule S . Let $X(F) = S(2, F) - S(1, F)$. By elicibility $X(F_c) > 0$ for $0 < c < 1$ but $X(F_0) < 0$, which violates the continuity of X .

4 Duality in property elicitation

In the full version [31, §G], we inspect Theorem 2 further and apply convex duality to reveal two notions of duality between affine scores: *report* duality, which asks the agent to report his desired allocation instead of his type, and *type* duality, which swaps the roles of the type and the allocation. Table 1 gives a particular instantiation of our duality notions, with $\mathcal{T} = \Delta(\mathcal{O})$ and $\mathcal{T}^* = (\mathcal{O} \rightarrow \mathbb{R})$; that is, we construct our affine scores and their duals upon the classic duality between integrable functions and probability measures. Note that G^* is the convex conjugate of G ; see Definition 17.

		Type	
		Primal	Dual
Report	Primal	$A(p', p)$ Scoring rule	$A^*(p', q)$ Menu auction
	Dual	$A(q', p)$ Prediction market	$A^*(q', q)$ Randomized mechanism
		$\sup A(\cdot, p) = G(p)$	$\sup A^*(\cdot, q) = G^*(q)$

Table 1. The duality quadrangle for the duality between distributions and functionals.

As discussed in Section G.2, the columns of Table 1 are well-understood already; the first gives prediction market duality, the well-known fact that market scoring rules are dual to prediction markets, and the second gives the taxation principle, which says that without loss of generality one could think of a mechanism as assigning prices to probability distributions over outcomes o .

The rows of this table, however, are new: in essence, scoring rules are dual mechanisms. In the scoring rule or prediction market setting, an agent has a private distribution (their belief) and the principal gives the agent a utility

⁸ Essentially the same construction can be applied for an $\alpha < 1$ confidence interval.

vector (the score or the bundle of securities), which assigns the agent a real-valued payoff for each possible state of the world. Dually, in a mechanisms, the agent possesses a private type encoding their utility for each state of the world, and the principal assigns a distribution over these states. This observation allows us to give a very simple and natural construction to convert between scoring rules and mechanisms. Unlike previous constructions (e.g., [29]) we do not require any normalization, or even that the set of outcomes be finite.

Corollary 9. *Let $S(p, o) = G(p) + G_p(o) - \int_{\mathcal{O}} G_p(o) dp(o)$ be a proper scoring rule. Then $f(t) = dG^*(t)$ and $p(t) = G^*(t) - t(dG^*(t))$ is a truthful randomized mechanism. Conversely, let (f, p) be a truthful randomized mechanism and $G(t) = t(f(t)) - p(t)$. Then $S(p, o) = G^*(p) + G_p^*(o) - \int_{\mathcal{O}} G_p^*(o) dp(o)$ is a proper scoring rule.*

The connections go much deeper than swapping types, however. To illustrate this with a somewhat whimsical example, suppose a gambler in a casino examines the rules of a dice-based game of chance and forms belief p about the probabilities of possible outcomes, assuming the dice are fair. The gambler then participates in a prediction market A to predict the outcome of the game, and purchases a bundle q . Before the game is played however, the casino informs the gambler that the dice used need not be fair, and offers the gambler the opportunity to select from among different choices of dice using a truthful mechanism where the gambler's private information is q . If the mechanism used is A^* , then the outcome of the mechanism will be using fair dice. The power of duality is that this holds regardless of our choice of A .

5 Finite-valued properties

We now examine the special case where \mathcal{R} is a finite set of reports, using the additional structure to provide stronger characterizations. In the scoring rules literature, Lambert and Shoham [48] view this as eliciting answers to multiple-choice questions. There are also applications to mechanism design, discussed in Section 5.1. Assume throughout that \mathcal{R} is finite and that \mathcal{T} is a convex subset of a vector space \mathcal{V} endowed with an inner product, so that we may write $\langle t, t' \rangle$ and in particular $\|t\|^2 = \langle t, t \rangle$. In this setting, we will use the concept of a power diagram from computational geometry.

Definition 7. *Given a set of points $P = \{p_i\}_{i=1}^m \subset \mathcal{V}$, called sites, and weights $w \in \mathbb{R}^m$, a power diagram $D(P, w)$ is a collection of cells $\text{cell}(p_i) \subseteq \mathcal{T}$ defined by*

$$\text{cell}_{P,w}(p_i) = \{t \in \mathcal{T} \mid i \in \text{argmin}_j \{\|p_j - t\|^2 - w_j\}\}. \quad (8)$$

The following result is a straightforward generalization of Theorem 4.1 of Lambert and Shoham [48], and is essentially a restatement of results due to Aurenhammer [6, 8]. See the full version [31, § H] for a discussion about Bregman Voronoi digrams and the role of $\|\cdot\|^2$ in Theorem 3.

Theorem 3. *A property $\Gamma : \mathcal{T} \rightrightarrows \mathcal{R}$ for finite \mathcal{R} is elicitable if and only if the level sets $\{\Gamma_r\}_{r \in \mathcal{R}}$ form a power diagram $D(P, w)$.*

We have now seen what kinds of finite-valued properties are elicitable, but how can we elicit them? More precisely, as the proof above gives sufficient conditions, what are all ways of eliciting a given power-diagram? In general, it is difficult to provide a “closed form” answer, so we restrict to the *simple* case, where essentially the cells of a power diagram are as constrained as possible.

Definition 8 ([7]). *A j -polyhedron is the intersection of dimension j of a finite number of closed halfspaces of \mathcal{V} , where $0 \leq j \leq \dim(\mathcal{V}) < \infty$. A cell complex C in \mathcal{V} is a covering of \mathcal{V} by finitely many j -polyhedra, called j -faces of C , whose (relative) interiors are disjoint and whose non-empty intersections are faces of C . C is called simple if each of its j -faces is in the closure of exactly $(d - j + 1)$ d -faces (cells).*

Theorem 4. *Let finite-valued, elicitable, simple property $\Gamma : \mathcal{T} \rightrightarrows \mathcal{R}$ be given. Then there exist points $\{p_r\}_{r \in \mathcal{R}} \subseteq \mathcal{V}$ such that an affine score $A : \mathcal{R} \times \mathcal{T} \rightarrow \mathbb{R}$ elicits Γ if and only if there exist $\alpha > 0$, and $p_0 \in \mathcal{V}$ such that*

$$A(r, t) = 2 \langle \alpha p_r + p_0, t \rangle - \|\alpha p_r + p_0\|^2 + w_r, \quad (9)$$

where the choice $w \in \mathbb{R}^{\mathcal{R}}$ is determined by α and p_0 .

5.1 Finite Properties in Mechanism Design

Mechanisms with a finite set of allocations are common. Carroll [19] examines them and observes they give rise to polyhedral typespaces. Theorem 3 strengthens this characterization to power diagrams, which rules out polyhedral examples such as the one shown in Figure 2. Suppose we have are in a such a mechanism design setting with a finite set of allocations \mathcal{X} and we have picked an allocation rule a . Under what circumstances is a implementable (i.e. having a payment rule that makes the resulting mechanism truthful)? If the set of types is convex, Saks and Yu [60] showed that the following condition is necessary and sufficient.

Definition 9. *Allocation rule a satisfies weak monotonicity (WMON) if $a(t) \cdot (t' - t) \leq a(t') \cdot (t' - t)$ for all $t, t' \in \mathcal{T}$.*

From Theorem 1, we know that a being implementable means that there exists a G such that a is a selection of its subgradients. But this is equivalent to saying that the property $\Gamma(t) = \mathcal{X} \cap dG_t$ is directly elicitable! Leveraging results from computational geometry, this gives us a new proof of this theorem by showing that WMON characterizes power diagrams.

Theorem 5. *A cell complex C is a power diagram with sites $\{p_1, \dots, p_n\}$ iff for all $t \in Z_i$ and $t' \in Z_j$ we have $p_i \cdot (t' - t) \leq p_j \cdot (t' - t)$ (i.e. C satisfies WMON)*

Corollary 10 (Saks and Yu). *If \mathcal{X} is finite, \mathcal{T} is convex, and a satisfies WMON, then a is implementable.*

References

1. Abernethy, J., Frongillo, R.: A characterization of scoring rules for linear properties. In: Proceedings of the 25th Conference on Learning Theory (2012)
2. Abernethy, J., Chen, Y., Vaughan, J.W.: Efficient market making via convex optimization, and a connection to online learning. *ACM Transactions on Economics and Computation* 1(2), 12 (2013)
3. Aliprantis, C.D., Border, K.C.: *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer (2007)
4. Archer, A., Kleinberg, R.: Truthful germs are contagious: a local to global characterization of truthfulness. In: Proceedings of the 9th ACM Conference on Electronic Commerce. pp. 21–30 (2008)
5. Ashlagi, I., Braverman, M., Hassidim, A., Monderer, D.: Monotonicity and implementability. *Econometrica* 78(5), 1749–1772 (2010)
6. Aurenhammer, F.: Power diagrams: properties, algorithms and applications. *SIAM Journal on Computing* 16(1), 78–96 (1987)
7. Aurenhammer, F.: Recognising polytopical cell complexes and constructing projection polyhedra. *Journal of Symbolic Computation* 3(3), 249–255 (Jun 1987)
8. Aurenhammer, F.: A criterion for the affine equivalence of cell complexes in r d and convex polyhedra in r $d+1$. *Discrete & Computational Geometry* 2(1), 49–64 (Dec 1987)
9. Babaioff, M., Blumrosen, L., Lambert, N.S., Reingold, O.: Only valuable experts can be valued. In: Proceedings of the 12th ACM conference on Electronic commerce. pp. 221–222 (2011)
10. Babaioff, M., Sharma, Y., Slivkins, A.: Characterizing truthful multi-armed bandit mechanisms. In: Proceedings of the 10th ACM conference on Electronic commerce. pp. 79–88 (2009)
11. Baldwin, E., Klemperer, P.: Tropical geometry to analyse demand. Tech. rep., Working paper, Oxford University (2012)
12. Berger, A., Mller, R., Naeemi, S.: Characterizing incentive compatibility for convex valuations. *Algorithmic Game Theory* pp. 24–35 (2009)
13. Berger, A., Mller, R., Hossein, N.S.: Path-monotonicity and incentive compatibility. Tech. rep., Maastricht: METEOR, Maastricht Research School of Economics of Technology and Organization (2010)
14. Bikhchandani, S., Chatterji, S., Lavi, R., Mu'alem, A., Nisan, N., Sen, A.: Weak monotonicity characterizes deterministic dominant-strategy implementation. *Econometrica* 74(4), 1109–1132 (2006)
15. Boissonnat, J.D., Nielsen, F., Nock, R.: Bregman voronoi diagrams: Properties, algorithms and applications. CoRR abs/0709.2196 (2007)
16. Borwein, J.M., Vanderwerff, J.D.: *Convex Functions: Constructions, Characterizations and Counterexamples*. Cambridge University Press (Jan 2010)
17. Cai, Y., Mahdian, M., Mehta, A., Waggoner, B.: Designing markets for daily deals. Preprint (2013)
18. Carbajal, J.C., Ely, J.C.: Mechanism design without revenue equivalence. *Journal of Economic Theory* (2012)
19. Carroll, G.: When are local incentive constraints sufficient? *Econometrica* 80(2), 661–686 (2012)
20. Chen, Y., Kash, I., Ruberry, M., Shnayder, V.: Decision markets with good incentives. *Internet and Network Economics* pp. 72–83 (2011)
21. Chen, Y., Kash, I.A.: Information elicitation for decision making. *AAMAS* (2011)

22. Chen, Y., Ruberry, M., Wortman Vaughan, J.: Cost function market makers for measurable spaces. In: Proceedings of the fourteenth ACM conference on Electronic commerce. pp. 785–802 (2013)
23. Chung, K.S., Ely, J.C.: Ex-post incentive compatible mechanism design. URL <http://www.kellogg.northwestern.edu/research/math/dps/1339.pdf>. Working Paper (2002)
24. Cid-Sueiro, J.: Proper losses for learning from partial labels. In: Advances in Neural Information Processing Systems 25. pp. 1574–1582 (2012)
25. Diakonikolas, I., Papadimitriou, C., Pierrakos, G., Singer, Y.: Efficiency-revenue trade-offs in auctions. In: Automata, Languages, and Programming, pp. 488–499. Springer (2012)
26. Fang, F., Stinchcombe, M.B., Whinston, A.B.: Proper scoring rules with arbitrary value functions. *Journal of Mathematical Economics* 46(6), 1200–1210 (2010)
27. Feige, U., Tennenholtz, M.: Responsive lotteries. *Algorithmic Game Theory* 6386, 150–161 (2010)
28. Feldman, J., Muthukrishnan, S.: Algorithmic methods for sponsored search advertising. In: Liu, Z., Xia, C.H. (eds.) *Performance Modeling and Engineering*, pp. 91–122. Springer US (Jan 2008)
29. Fiat, A., Karlin, A., Koutsoupias, E., Vidali, A.: Approaching utopia: strong truthfulness and externality-resistant mechanisms. In: Proceedings of the 4th conference on Innovations in Theoretical Computer Science. pp. 221–230 (2013)
30. Frongillo, R., Kash, I.: Vector-valued property elicitation. Preprint (2014)
31. Frongillo, R.M., Kash, I.A.: General truthfulness characterizations via convex analysis. arXiv:1211.3043 [cs] (Nov 2012)
32. Gneiting, T.: Making and evaluating point forecasts. *Journal of the American Statistical Association* 106(494), 746–762 (2011)
33. Gneiting, T., Raftery, A.: Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association* 102(477), 359–378 (2007)
34. Gneiting, T., Katzfuss, M.: Probabilistic forecasting. *Annual Review of Statistics and Its Application* 1, 125–151 (2014)
35. Grant, K., Gneiting, T.: Consistent scoring functions for quantiles. In: *From Probability to Statistics and Back: High-Dimensional Models and Processes* A Festschrift in Honor of Jon A. Wellner, pp. 163–173. Institute of Mathematical Statistics (2013)
36. Halpern, J.: Reasoning about uncertainty. MIT Press (2003)
37. Heydenreich, B., Mller, R., Uetz, M., Vohra, R.V.: Characterization of revenue equivalence. *Econometrica* 77(1), 307–316 (2009)
38. Ioffe, A.D., Tikhomirov, V.M.: *Theory of extremal problems*. North-Holland Pub. Co. ; sole distributors for the U.S.A. and Canada, Elsevier North-Holland, Amsterdam; New York; New York (1979)
39. Iyer, R., Bilmes, J.: The lovsz-bregman divergence and connections to rank aggregation, clustering, and web ranking: Extended version. *Uncertainty in Artificial Intelligence* (2013)
40. Jehiel, P., Moldovanu, B.: Efficient design with interdependent valuations. *Econometrica* 69(5), 1237–1259 (2001)
41. Jehiel, P., Moldovanu, B., Stacchetti, E.: How (not) to sell nuclear weapons. *The American Economic Review* 86(4), pp. 814–829 (1996)
42. Jehiel, P., Moldovanu, B., Stacchetti, E.: Multidimensional mechanism design for auctions with externalities. *Journal of Economic Theory* 85(2), 258 – 293 (1999)
43. Kos, N., Messner, M.: Extremal incentive compatible transfers. *Journal of Economic Theory* (2012)

44. Krishna, V., Maenner, E.: Convex potentials with an application to mechanism design. *Econometrica* 69(4), 1113–1119 (Jul 2001)
45. Lai, H.C., Lin, L.J.: The fenchel-moreau theorem for set functions. *Proceedings of the American Mathematical Society* pp. 85–90 (1988)
46. Lambert, N.: Elicitation and evaluation of statistical forecasts. Preprint (2011)
47. Lambert, N., Pennock, D., Shoham, Y.: Eliciting properties of probability distributions. In: *Proceedings of the 9th ACM Conference on Electronic Commerce*. pp. 129–138 (2008)
48. Lambert, N., Shoham, Y.: Eliciting truthful answers to multiple-choice questions. In: *Proceedings of the 10th ACM conference on Electronic commerce*. pp. 109–118 (2009)
49. McAfee, R., McMillan, J.: Multidimensional incentive compatibility and mechanism design. *Journal of Economic Theory* 46(2), 335–354 (Dec 1988)
50. Myerson, R.B.: Optimal auction design. *Mathematics of operations research* pp. 58–73 (1981)
51. Miller, R., Perea, A., Wolf, S.: Weak monotonicity and bayes-nash incentive compatibility. *Games and Economic Behavior* 61(2), 344–358 (2007)
52. Negahban, S.N., Ravikumar, P., Wainwright, M.J., Yu, B.: A unified framework for high-dimensional analysis of m-estimators with decomposable regularizers. *arXiv preprint arXiv:1010.2731* (2010)
53. Ok, E.A.: *Real analysis with economic applications*, vol. 10. Princeton University Press (2007)
54. Osband, K., Reichelstein, S.: Information-eliciting compensation schemes. *Journal of Public Economics* 27(1), 107–115 (Jun 1985)
55. Osband, K.H.: *Providing Incentives for Better Cost Forecasting*. University of California, Berkeley (1985)
56. Othman, A., Sandholm, T.: Decision rules and decision markets. In: *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*. vol. 1, pp. 625–632 (2010)
57. Peters, H.J.M., Wakker, P.P.: Convex functions on non-convex domains. *Economics Letters* 22(2-3), 251–255 (1987)
58. Rochet, J.C.: A necessary and sufficient condition for rationalizability in a quasi-linear context. *Journal of Mathematical Economics* 16(2), 191–200 (1987)
59. Rockafellar, R.: *Convex analysis*, Princeton Mathematics Series, vol. 28. Princeton University Press (1997)
60. Saks, M., Yu, L.: Weak monotonicity suffices for truthfulness on convex domains. In: *Proceedings of the 6th ACM conference on Electronic commerce*. pp. 286–293 (2005)
61. Savage, L.: Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association* pp. 783–801 (1971)
62. Steinwart, I., Pasin, C., Williamson, R., Zhang, S.: Elicitation and identification of properties. *Technical Report Stuttgart University* (2014)
63. Urruty, J.B.H., Lemarchal, C.: *Fundamentals of Convex Analysis*. Springer (2001)
64. Van Manen, M., Siersma, D.: Power diagrams and their applications. *arXiv preprint math/0508037* (2005)
65. Vohra, R.V.: *Mechanism design: a linear programming approach*. Cambridge University Press, Cambridge; New York (2011)
66. Yan, M.: Extension of convex functions. *arXiv preprint arXiv:1207.0944* (2012)

Acknowledgements

We would like to thank Aaron Archer, Gabriel Carroll, Yiling Chen, Philip Dawid, Felix Fischer, Hu Fu, Philippe Jehiel, Nicolas Lambert, Peter Key, David Parkes, Colin Rowat, Greg Stoddard, Matus Telgarsky, and Rakesh Vohra for helpful discussions and anonymous reviewers for helpful feedback.

A Convex Analysis Primer

In this appendix, we review some facts from convex analysis that are used in the paper.

Fact 1 *Let $\{f_t \in \text{Aff}(\mathcal{V} \rightarrow \mathbb{R})\}_{t \in \mathcal{T}}$ be a parameterized family of affine functions. Then $G(t) = \sup_{t' \in \mathcal{T}} f_{t'}(t)$ is convex as the pointwise supremum of convex functions.*

This follows because convex functions are those with convex epigraphs. The epigraph of this supremum is the intersection of the epigraphs of the individual functions, which is a convex set as the intersection of convex sets.

Fact 2 *$d : \mathbb{R} \rightarrow \mathbb{R}$ is a selection of subgradients of a convex function on \mathbb{R} if and only if it is monotone non-decreasing.*

See [59, Theorem 24.3] for a proof of a slightly more general statement.

Fact 3 *For convex G on convex \mathcal{T} , $\{dG_t \in \text{Lin}(\mathcal{V} \rightarrow \mathbb{R})\}_{t \in \mathcal{T}} \in \partial G$ satisfies path independence.*

Since the restriction of G to a line is a one-dimensional convex function, $G(y) - G(x) = \int_{L_{xy}} dG_t(y - x) dt$ [59, Corollary 24.2.1]. Summing along the individual lines in a path from x to y gives that the value of the path integral is $G(y) - G(x)$ regardless of the path chosen.

Fact 4 *For any convex function G , the set $\partial G^{-1}(d) \doteq \{x \in \text{dom}(G) : d \in \partial G_x\}$ is convex.*

Proof. Let $x, x' \in \partial G^{-1}(d)$; then one easily shows (cf. Lemma 2) that $G(x) - G(x') = d(x - x')$. Now let $\hat{x} = \alpha x + (1 - \alpha)x'$; we have,

$$G(\hat{x}) \leq \alpha G(x) + (1 - \alpha)G(x') \tag{10}$$

$$= \alpha(G(x) - G(x')) + G(x')$$

$$= \alpha d(x - x') + G(x')$$

$$= d(\hat{x} - x') + G(x') \tag{11}$$

$$\leq G(\hat{x}), \tag{12}$$

where we applied convexity of G in (10) and the subgradient inequality for d at x' in (12). Hence, by eq. (11) we have shown $G(\hat{x}) - G(x') = d(\hat{x} - x')$, so by Lemma 2, $d \in \partial G_{\hat{x}}$.

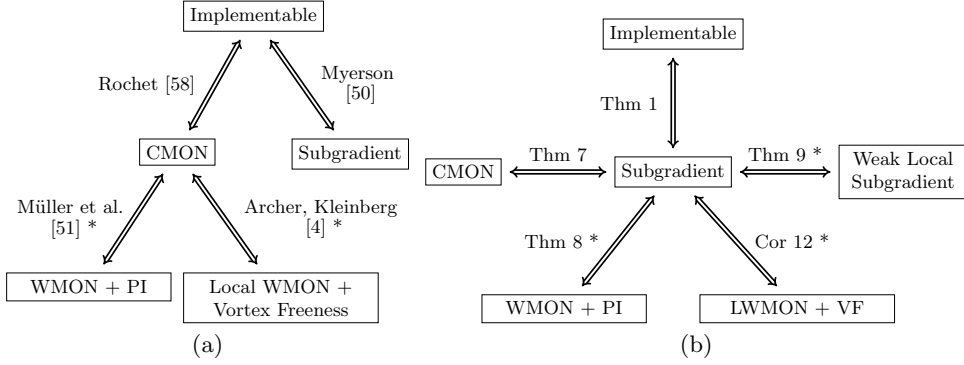


Fig. 1. Proof structure of existing mechanism design literature (a), and the new proof structure presented in this paper (b). Asterisks denote the requirement that \mathcal{T} be convex. We write CMON for cyclic monotonicity, WMON for weak monotonicity, and PI for path independence.




Fig. 2. An allocation rule which cannot be implemented, for any distinct choices of a_i (left), and a rule which could be implemented for appropriate choices of a_i (right).

B Omitted Proofs and Figures

Proof (of Theorem 1). It is trivial from the subgradient inequality (2) that the proposed form is in fact truthful, as

$$A(t', t) = G(t') + d_{t'}(t - t') \leq G(t) = G(t) + d_t(t - t) = A(t, t).$$

For the converse, we are given some truthful $A : \mathcal{T} \times \mathcal{T} \rightarrow \overline{\mathbb{R}}$. Note first that for any $\hat{t} \in \text{Conv}(\mathcal{T})$ we may write \hat{t} as a finite convex combination $\hat{t} = \sum_{i=1}^m \alpha_i t_i$ where $t_i \in \mathcal{T}$. Now, as all the elements of the score set \mathcal{A} are affine, we may naturally extend $A(t, \cdot)$ to all of $\text{Conv}(\mathcal{T})$ by defining

$$A(t, \hat{t}) = \sum_{i=1}^m \alpha_i A(t, t_i). \quad (13)$$

One easily checks that this definition coincides with the given A on \mathcal{T} .

Now we let $G(\hat{t}) := \sup_{t \in \mathcal{T}} \mathbf{A}(t, \hat{t})$, which is convex as the pointwise supremum of convex (in our case affine) functions. Since \mathbf{A} is truthful, we in particular have $G(t) = \mathbf{A}(t, t) \in \mathbb{R}$ for all $t \in \mathcal{T}$ by our regularity assumption. Let $\mathbf{A}_\ell(t, \cdot)$ denote the linear part of $\mathbf{A}(t, \cdot)$. Then, also by truthfulness, we have for all $t' \in \mathcal{T}$ and $\hat{t} \in \text{Conv}(\mathcal{T})$,

$$\begin{aligned} G(\hat{t}) &= \sup_{t \in \mathcal{T}} \sum_{i=1}^m \alpha_i \mathbf{A}(t, t_i) \geq \sum_{i=1}^m \alpha_i \mathbf{A}(t', t_i) = \mathbf{A}(t', t') + \sum_{i=1}^m \alpha_i \mathbf{A}_\ell(t', t_i - t') \\ &= G(t') + \mathbf{A}_\ell(t', \hat{t} - t'). \end{aligned}$$

Hence, $\mathbf{A}_\ell(t', \cdot)$ satisfies (2) for G at t' , so \mathbf{A} is of the form (3).

Proof (of Corollary 1). The given form is truthful by the subgradient inequality (2). Let $\mathbf{A} : \mathcal{T} \times \mathcal{T} \rightarrow \mathbb{R}$ be a given truthful affine score. Since $\mathbf{A}(p, \cdot) \in \mathcal{A}$, we have some $f_p \in \mathcal{F}$ generating $\mathbf{A}(p, \cdot)$. We can therefore use $G_p : q \mapsto \int_{\mathcal{O}} f_p(o) dq(o)$ as the subgradients in the proof of Theorem 1, thus giving us the desired form.

Proof (of Corollary 4). By elementary results in convex analysis f is a subgradient of a convex function on \mathbb{R} if and only if it is monotone non-decreasing. By Theorem 1, the mechanism is truthful if and only if f is the subgradient of the particular function $G(t) = U(t, t) = t(f(t)) - p(t)$, which is equivalent to (i) and the condition $G(t) = \int_0^t f(t') dt' + C$.

Proof (of Corollary 9). Apply Theorem 16 with $\Gamma : p \mapsto G_p$ and $\Gamma^* : t \mapsto f(t)$ to get truthfulness of the corresponding affine scores and Corollaries 1 and 5 to get the specified forms.

Proof (of Theorem 3). Let us examine the condition that t is an element of $\text{cell}_{P,w}(p_i)$ for some power diagram $D(P, w)$:

$$\begin{aligned} t \in \text{cell}_{P,w}(p_i) &\iff i \in \underset{j}{\text{argmin}} \{ \|p_j - t\|^2 - w_j \} \\ &\iff i \in \underset{j}{\text{argmin}} \{ \|p_j\|^2 - 2\langle p_j, t \rangle - w_j \}. \end{aligned} \quad (14)$$

Note that eq. (14) is affine in t . Now given some $D = D(P, w)$ with index set \mathcal{R} , we simply let $\mathbf{A}(r, t) = 2\langle p_r, t \rangle + w_r - \|p_r\|^2$. By (14) we immediately have $r \in \text{argsup}_{r'} \mathbf{A}(r', t) \iff t \in \text{cell}_{P,w}(p_r)$, as desired.

Conversely, let an affine score \mathbf{A} eliciting Γ be given. Note that since we are in an inner product space, we may write $\mathbf{A}(r, t) = \langle x_r, t \rangle + c_r$ for $x_r \in \mathcal{V}$ and $c_r \in \mathbb{R}$. Letting $p_r = x_r/2$ and $w_r = \|p_r\|^2 + c_r$, we see by (14) again that $\Gamma_r = \text{cell}(p_r)$ of the diagram $D(\{p_r\}, w)$. Hence, Γ is a power diagram.

Proof (of Theorem 4). A result of Aurenhammer for simple cell complexes, given in Lemma 1 of [6] and the proof of Lemma 4 of [8], states the following: given sites P and P' and weights w , there exist weights w' such that $D(P', w') = D(P, w)$ if and only if P' is a homothet (translated and positively scaled copy) of P . We simply apply this fact to the proof of Theorem 3.

Proof (of Theorem 5). We will make use of the following characterization of power diagrams from the computational geometry literature.

Theorem 6 ([7]). *A cell complex C is a power diagram if and only if there exists a point-set $\{p_1, \dots, p_n\}$ satisfying.*

1. *Orthogonality.* For $Z_i \neq Z_j$, the line L that contains p_i and p_j (and is directed from p_i to p_j) is orthogonal to each face common to Z_i and Z_j .
2. *Orientation:* Any directed line that can be obtained by translating L and that intersects Z_i and Z_j first meets Z_i .

If C is a power diagram, then by definition

$$\begin{aligned} 2p_i \cdot t - w_i &\geq 2p_j \cdot t - w_j \\ 2p_j \cdot t' - w_j &\geq 2p_i \cdot t' - w_i. \end{aligned}$$

Adding these shows C satisfies WMON.

Now suppose C satisfies WMON. We show orthogonality and orientation. For orthogonality, let $t, t' \in Z_i \cap Z_j$. Then $p_i \cdot (t' - t) = p_j \cdot (t' - t)$, or $(p_i - p_j) \cdot (t' - t) = 0$. Thus, the face is orthogonal to L .

For orientation, let $t \in Z_i$ and $t' \in Z_j$ be on such a translated L . That is, we can write $t' = t + c(p_j - p_i)$ for some $c \in \mathbb{R}$. By WMON, $(p_j - p_i) \cdot (t' - t) \geq 0$, or $c(p_j - p_i) \cdot (p_j - p_i) \geq 0$. Thus $c \geq 0$. Therefore such a translated L first meets Z_i .

Proof (of Corollary 10). In order to apply Theorems 3 and 5, it remains to show that an allocation rule a satisfying WMON further implies that it defines a cell complex. This follows by a straightforward geometric argument that has been used in a number of previous proofs (see, e.g., Lemma 4.2 of [4]).

C Characterizing Truthful Mechanisms

While our theorem provides a characterization of truthful mechanisms in terms of convex consumer surplus functions, this is not always the most natural representation for a mechanism. In this section, we examine two other approaches to characterizing truthful mechanisms that have been explored in the literature and show that they have insightful interpretations in convex analysis, which allows us to greatly simplify their proofs. Furthermore, our phrasing of these results is as conditions for a parameterized family of linear functions to be a selection of subgradients of a convex function. We believe this phrasing converts known results in mechanism design into new results in convex analysis. It also shows how any such result in convex analysis would give a characterization of implementable mechanisms. Note that certain results in this section require an assumption that the relevant parameterized families are in fact real-valued, which is natural given our focus on mechanism design.

C.1 Subgradient characterizations

From an algorithmic perspective, it may be more natural to focus on the design of the allocation rule f . There is a large literature that focuses on when there exists a choice of payments p to make f into a truthful mechanism (e.g. [5, 60]). Viewed through our theorem, this becomes a very natural convex analysis question: when is a function f a subgradient of a convex function⁹? Unsurprisingly, the central result in the literature is closely connected to convex analysis.

Definition 10. A family $\{d_t \in \text{Lin}(\mathcal{V} \rightarrow \mathbb{R})\}_{t \in \mathcal{T}}$ satisfies cyclic monotonicity (CMON) if for all finite sets $\{t_0, \dots, t_k\} \subseteq \mathcal{T}$,

$$\sum_{i=0}^k d_{t_i}(t_{i+1} - t_i) \leq 0, \quad (15)$$

where indices are taken modulo $k + 1$. The weaker condition that (15) hold for all pairs $\{t_0, t_1\}$ is known as weak monotonicity (WMON).

A well known characterization from convex analysis is that a function f defined on a convex set is a subgradient of a convex function on that set iff it satisfies CMON [59]. Rochet's [58] proof that such payments exist on a possibly non-convex \mathcal{T} iff f satisfies CMON is effectively a proof of the a generalization of this theorem. Rochet notes that his proof is adapted from the one given in Rockafellar's text [59] of the weaker theorem where \mathcal{T} is restricted to be convex. We adapt Rochet's proof to highlight how its core is a construction of G . As we use this basic construction several times, we first analyze it independently.

Given any family $\{d_t\}_{t \in \mathcal{T}}$ of linear functions in $\text{Lin}(\mathcal{V} \rightarrow \mathbb{R})$, define $P_d : \mathcal{T} \times \mathcal{V} \rightarrow \mathbb{R}$ as follows:¹⁰

$$P_d(t, t') \doteq \sup_{\substack{k \in \mathbb{N}, \{t_1, \dots, t_k\} \subseteq \mathcal{T} \\ t_0 = t, t_{k+1} = t'}} \sum_{i=0}^k d_{t_i}(t_{i+1} - t_i). \quad (16)$$

One way to interpret $P_d(t, t')$ is as the length of the shortest path from t to t' in a graph with edge weights determined by $-d$, and in that form has seen extensive use in mechanism design [65]. We interpret it somewhat differently, as the best lower bound on $G(t') - G(t)$ for an arbitrary convex function G with subgradients d (and infinity if there is no such convex function). In particular, computing the best lower bound at every point yields a convex function.

Lemma 1. Let $\{d_t \in \text{Lin}(\mathcal{V} \rightarrow \mathbb{R})\}_{t \in \mathcal{T}}$ be given. If d satisfies CMON, then for all $t, t' \in \mathcal{T}$ and all $t'' \in \mathcal{V}$, the following hold:

⁹ More precisely, we want for all t the allocation $f(t)$ to be a subgradient at t . Equivalently, we can view f as a parameterized family of functions, which is how we state our results

¹⁰ Note that the second argument of P_d is from \mathcal{V} rather than $T \subset \mathcal{V}$ because we wish to apply this when, e.g., $t \in \text{Conv}(T)$.

1. $P_d(t, t') + P_d(t', t'') \leq P_d(t, t'')$
2. $d_t(t'' - t) \leq P_d(t, t'')$
3. $P_d(t, t) = 0$
4. $P_d(t, t') + P_d(t', t) \leq 0$
5. $P_d(t, \cdot)$ is convex and real-valued on $\text{Conv}(\mathcal{T})$, with $d \in \partial P_d(t, \cdot)$ on \mathcal{T}

Otherwise, $P_d \equiv \infty$ on all inputs.

Proof. If CMON is not satisfied, then there is a cycle $C = t_0, \dots, t_k, t_0$ with positive sum. Then for any t and t' the path $tC^j t'$ that consists of starting at t , going to t_0 , going around the cycle j times, then going to t' has a sum that goes to infinity as j goes to infinity. For the remainder, assume that CMON is satisfied.

1. $P_d(t, t') + P_d(t', t'') \leq P_d(t, t'')$

$$\begin{aligned}
 & P_d(t, t') + P_d(t', t'') \\
 &= \sup_{\substack{k \in \mathbb{N}, \{t_1, \dots, t_k\} \subseteq \mathcal{T} \\ t_0 = t, t_{k+1} = t'}} \sum_{i=0}^k d_{t_i}(t_{i+1} - t_i) + \sup_{\substack{k \in \mathbb{N}, \{t_1, \dots, t_k\} \subseteq \mathcal{T} \\ t_0 = t', t_{k+1} = t''}} \sum_{i=0}^k d_{t_i}(t_{i+1} - t_i) \\
 &= \sup_{\substack{j, k \in \mathbb{N}, \{t_1, \dots, t_k\} \subseteq \mathcal{T} \\ t_0 = t, t_j = t', t_{k+1} = t''}} \sum_{i=0}^k d_{t_i}(t_{i+1} - t_i) \\
 &\leq \sup_{\substack{k \in \mathbb{N}, \{t_1, \dots, t_k\} \subseteq \mathcal{T} \\ t_0 = t, t_{k+1} = t''}} \sum_{i=0}^k d_{t_i}(t_{i+1} - t_i) \\
 &= P_d(t, t'')
 \end{aligned}$$

2. $d_t(t'' - t) \leq P_d(t, t'')$

Taking $k = 0$ shows that $d_t(t'' - t)$ is an element of set over which the supremum is taken.

3. $P_d(t, t) = 0$

By CMON, $P_d(t, t) \leq 0$. By claim (2), $d_t(t - t) = 0 \leq P_d(t, t'')$.

4. $P_d(t, t') + P_d(t', t) \leq 0$

By claims (1) and (3), $P_d(t, t') + P_d(t', t) \leq P(t, t) = 0$.

5. $P_d(t, \cdot)$ is convex and real-valued on $\text{Conv}(\mathcal{T})$, with $d \in \partial P_d(t, \cdot)$ on \mathcal{T}

By CMON, for $t' \in \mathcal{T}$ $P_d(t, t') \leq -d_t(t_0 - t')$. Thus, $P_d(t, t')$ is finite on \mathcal{T} . $P_d(t, \cdot)$ is a pointwise supremum of convex functions, so is convex. By

convexity, it is also finite on $\text{Conv}(\mathcal{T})$. For any $t' \in \mathcal{T}$ and $t'' \in \text{Conv}(\mathcal{T})$,

$$\begin{aligned}
P_d(t, t') + d_{t'}(t'' - t') &= d_{t'}(t'' - t') + \sup_{\substack{k \in \mathbb{N}, \{t_1, \dots, t_k\} \subseteq \mathcal{T} \\ t_0 = t, t_{k+1} = t'}} \sum_{i=0}^k d_{t_i}(t_{i+1} - t_i) \\
&= \sup_{\substack{k \in \mathbb{N}, \{t_1, \dots, t_k\} \subseteq \mathcal{T} \\ t_0 = t, t_k = t', t_{k+1} = t''}} \sum_{i=0}^k d_{t_i}(t_{i+1} - t_i) \\
&\leq \sup_{\substack{k \in \mathbb{N}, \{t_1, \dots, t_k\} \subseteq \mathcal{T} \\ t_0 = t, t_{k+1} = t''}} \sum_{i=0}^k d_{t_i}(t_{i+1} - t_i) \\
&= P(t, t''),
\end{aligned}$$

so d_t satisfies (2).

Having extracted the construction at the core of Rochet's proof, the rephrasing of his result as a statement about convex functions now follows easily.

Theorem 7 (Adapted from Rochet [58]). *A family $\{d_t \in \text{Lin}(\mathcal{V} \rightarrow \mathbb{R})\}_{t \in \mathcal{T}}$ satisfies CMON if and only if there exists a convex $G : \text{Conv}(\mathcal{T}) \rightarrow \mathbb{R}$ such that $\{d_t\}_{t \in \mathcal{T}} \in \partial G$.*

Proof. Given such a G , by (2) we have $d_{t_i}(t_{i+1} - t_i) \leq G(t_{i+1}) - G(t_i)$. Summing gives (15). Given such a family $\{d_t\}_{t \in \mathcal{T}}$, fix some $t_0 \in \mathcal{T}$ and let $t_0 \in \mathcal{T}$ and set $G : t \mapsto P_d(t_0, t)$. The result follows from Lemma 1(5).

A number of papers have sought simpler and more natural conditions than CMON that are necessary and sufficient in special cases, e.g. [4, 5, 60]. These results are typically proven by showing they are equivalent to CMON. However, it is much more natural to directly construct the relevant G . As an example, we show one such result has a simple proof using our framework. This particular proof also has the advantage of providing a characterization of the payments that is more intuitive than the supremum in Rochet's construction.

As in Myerson's [50] construction for the single-parameter case, we construct a G by integrating over d_t . In particular, for any two types x and y our construction makes use of the line integral

$$\int_{L_{xy}} d_t(y - x) dt = \int_0^1 d_{(1-t)x + ty}(y - x) dt.$$

As Berger et al. [12] and Ashlagi et al. [5] observed, if $\{d_t\}_{t \in \mathcal{T}}$ satisfies WMON and \mathcal{T} is convex, this (Riemann) integral is well defined because it is the integral of a monotone function. If these line integrals vanish around all triangles (equivalently $\int_{L_{xy}} d_t(y - x) dt + \int_{L_{yz}} d_t(z - y) dt = \int_{L_{xz}} d_t(z - x) dt$) we say $\{d_t\}$ satisfies *path independence*.

Theorem 8 (adapted from [51]). *For convex \mathcal{T} , a family $\{d_t \in \text{Lin}(\mathcal{V} \rightarrow \mathbb{R})\}_{t \in \mathcal{T}}$ is a selection of subgradients of a convex function if and only if $\{d_t\}_{t \in \mathcal{T}}$ satisfies WMON and path independence.*

Proof. Given a convex function G and selection of subgradients $\{d_t\}$, $\{d_t\}$ satisfies CMON and thus WMON. Path independence also follows from convexity (Rockafellar [59] p. 232). Now given a $\{d_t\}$ that satisfies WMON and path independence, fix a type $t_0 \in \mathcal{T}$ and define $G(t') = \int_{L_{t_0 t'}} d_t(t' - t_0) dt$ (well defined by WMON as the integral of a monotone function). Given $x, y, z \in \mathcal{T}$ such that $z = \lambda x + (1 - \lambda)y$, by path independence and the linearity of d_z we have

$$\begin{aligned} & \lambda G(x) + (1 - \lambda)G(y) \\ &= G(z) + \lambda \int_{L_{zx}} d_t(x - z) dt + (1 - \lambda) \int_{L_{zy}} d_t(y - z) dt \\ &\geq G(z) + \lambda d_z(x - z) + (1 - \lambda)d_z(y - z) = G(z), \end{aligned}$$

so G is convex. Similarly, for $x, y \in \mathcal{T}$, d_t satisfies (2) because

$$d_x(y - x) \leq \int_{L_{xy}} d_t(y - x) dt = G(y) - G(x). \square$$

C.2 Local Characterizations

In many settings, it is easier to reason about the behavior of a mechanism given small changes to its input rather than arbitrary changes, so several authors have sought to characterize truthful mechanisms using local conditions [4, 12, 19]. We show in this section how many of these results are in essence a consequence of a more fundamental statement, that convexity is an inherently local property. For example, in the twice differentiable case it can be verified by determining whether the Hessian is positive semidefinite at each point. We start with a local convexity result, and use it to show that an affine score is truthful if and only if it satisfies a very weak local truthfulness property introduced by Carroll [19]. Afterwards we turn to a characterization by Archer and Kleinberg [4] that proved a similar theorem for a different notion of local truthfulness. Our results (specifically Theorem 9) show that these two notions of local truthfulness are equivalent because Archer and Kleinberg's definition corresponds to the property of being a local subgradient, while Carroll's corresponds to the property of being a weak, local subgradient, which we now define.

Definition 11. *Let \mathcal{T} be convex. A family $\{d_t \in \text{Lin}(\mathcal{V} \rightarrow \overline{\mathbb{R}})\}_{t \in \mathcal{T}}$ is a weak local subgradient (WLSG) of a convex function $G : \mathcal{T} \rightarrow \overline{\mathbb{R}}$ if for all $t \in \mathcal{T}$ there exists an open neighborhood U_t of t such that for all $t' \in U_t$,*

$$G(t) \geq G(t') + d_{t'}(t - t') \quad \text{and} \quad G(t') \geq G(t) + d_t(t' - t). \quad (17)$$

Furthermore, if for every $s \in \mathcal{T}$, eq. (17) holds for all $t, t' \in U_s$, we say $\{d_t\}_{t \in \mathcal{T}}$ is a local subgradient (LSG) of G .

We now show that being a WLSG is a necessary and sufficient condition for a family of functions to be a selection of subgradients. The proof is heavily inspired by Carroll [19].

Theorem 9. *Let \mathcal{T} be convex. A family $\{d_t \in \text{Lin}(\mathcal{V} \rightarrow \overline{\mathbb{R}})\}_{t \in \mathcal{T}}$ is a selection of subgradients of a convex function $G : \mathcal{T} \rightarrow \overline{\mathbb{R}}$ if and only if it is a WLSG of G .*

Proof ((adapted from [19])). As usual, the forward direction is trivial. For the other, let $t, t' \in \mathcal{T}$ be given; we show that the subgradient inequality for $d_{t'}$ holds at t . By compactness of $\text{Conv}(\{t, t'\})$, we have a finite set $t_i = \alpha_i t' + (1 - \alpha_i)t$, where $0 = \alpha_0 \leq \dots \leq \alpha_{k+1} = 1$, such that WLSG holds between each t_i and t_{i+1} . (The cover $\{U_s \mid s \in \text{Conv}(\{t, t'\})\}$ has a finite subcover. Take t_{2i} from the subcover and $t_{2i+1} \in U_{t_{2i}} \cap U_{t_{2i+2}}$.) By the WLSG condition (17), we have for each i ,

$$0 \geq G(t_{i+1}) - G(t_i) + d_{t_{i+1}}(t_i - t_{i+1}) \quad (18)$$

$$0 \geq G(t_i) - G(t_{i+1}) + d_{t_i}(t_{i+1} - t_i). \quad (19)$$

Now using the identity $t_{i+1} - t_i = (\alpha_{i+1} - \alpha_i)(t' - t)$ and adding $\alpha_i/(\alpha_{i+1} - \alpha_i)$ times (18) to $\alpha_{i+1}/(\alpha_{i+1} - \alpha_i)$ times (19), we have

$$0 \geq G(t_i) - G(t_{i+1}) + \alpha_i d_{t_i}(t' - t) - \alpha_{i+1} d_{t_{i+1}}(t' - t). \quad (20)$$

Summing (20) over $0 \leq i \leq k$ gives

$$0 \geq G(t_0) - G(t_{k+1}) + \alpha_0 d_{t_0}(t' - t) - \alpha_{k+1} d_{t_{k+1}}(t' - t),$$

which when recalling our definitions for α_i and t_i yields the result.

The WLSG condition translates to an analogous notion in terms of truthfulness, *weak local truthfulness*.

Definition 12. *An affine score is weakly locally truthful if for all $t \in \mathcal{T}$ there exists some open neighborhood U_t of t , such that truthfulness holds between t and every $t' \in U_t$, and vice versa. That is,*

$$\forall t \in \mathcal{T}, \forall t' \in U_t, \quad A(t', t) \leq A(t, t) \quad \text{and} \quad A(t, t') \leq A(t', t'). \quad (21)$$

Corollary 11 (Generalization of Carroll [19]). *An affine score $A : \mathcal{T} \times \mathcal{T} \rightarrow \overline{\mathbb{R}}$ for convex \mathcal{T} is truthful if and only if it is weakly locally truthful.*

Proof. Defining $G(t) := A(t, t)$, by weak local truthfulness we may write

$$\begin{aligned} G(t) = A(t, t) &\geq A(t', t) = G(t') + A_\ell(t', t - t') \\ G(t') = A(t', t') &\geq A(t, t') = G(t) + A_\ell(t, t' - t), \end{aligned}$$

where t' is local to t and $A_\ell(t, \cdot)$ is the linear part of $A(t, \cdot)$. This says that $d_t = A_\ell(t, \cdot)$ satisfies WLSG for convex function G ; the rest follows from Theorem 9 and Theorem 1.

Finally, in the spirit of Section C.1, Archer and Kleinberg [4] characterized local conditions under which an allocation rule can be made truthful. A key condition from their paper is *vortex-freeness*, which is a condition they show to be equivalent to local path independence (analogous to our terminology of weak local subgradients it can be thought of as weak local path independence). The other condition, local WMON, means that WMON holds in some neighborhood around each type. Their result then follows from the observation that local WMON and local path independence imply local subgradient. While this particular proof is not significantly simpler than the original, we believe it is somewhat more natural and clarifies the connection between the underlying reasons a notion of local truthfulness suffices both here and in Carroll’s setting.

Corollary 12. *Let \mathcal{T} be convex. A family $\{d_t \in \text{Lin}(\mathcal{V} \rightarrow \mathbb{R})\}_{t \in \mathcal{T}}$ is a selection of subgradients of a convex function if and only if it satisfies local WMON and is vortex-free.*

Proof. We prove the reverse direction; suppose $\{d_t\}_{t \in \mathcal{T}}$ satisfies local WMON and is vortex-free. From Lemma 3.5 of [4] we have that vortex-freeness is equivalent to path independence, so by Theorem 8 for all t there exists some open U_t such that $\{d_{t'}\}_{t' \in U_t}$ is the subgradient of some convex function $G^{(t)} : U_t \rightarrow \mathbb{R}$. We need only show the existence of some G such that $\{d_t\}_{t \in \mathcal{T}}$ is the subgradient of G on each U_t ; the rest follows from Theorem 9.

Fix some $t_0 \in \mathcal{T}$ and define $G(t) = \int_{L_{t_0 t}} d_{t'} dt'$, which is well defined by compactness of $\text{Conv}(\{t_0, t\})$ and the fact that a locally increasing real-valued function is increasing. But for each t' and $t \in U_{t'}$ we can also write $G^{(t')}(t) = \int_{L_{t' t}} d_{t''} dt''$ by [59, p. 232], and now by path independence we see that G and $G^{(t')}$ differ by a constant. Hence $\{d_t\}_{t \in \mathcal{T}}$ must be a subgradient of G on $U_{t'}$ as well, for all $t' \in \mathcal{T}$.

D Revenue Equivalence

Perhaps the most celebrated result in auction theory is the revenue equivalence theorem, which states that, in a single item auction, the revenue from an agent (equivalently that agent’s consumer surplus) is determined up to a constant by the equilibrium probability that each possible type of that agent will receive the item [50]. A large body of work has looked for more general conditions under which this property holds (see, e.g., [44]) or what can be said when it does not [18]. One general approach is due to Heydenreich et al. [37], who use a graphical representation related to CMON. Given our main theorem, this is unsurprising. In convex analysis terms, asking whether an implementable allocation rule satisfies revenue equivalence is asking whether all convex functions that have a selection of their subgradients that corresponds to that allocation rule are the same up to a constant. As we saw in the proof of Lemma 1, CMON permits the natural construction of a convex function from its subgradient via (15). Intuitively, if we know the payments we want for some subset of types, we

can check if those are consistent with a desired payment for some other type by checking whether this construction still works, both in terms of the constraints of the existing types on the new one and the new one on the existing ones. The following theorem applies this insight to get a result that is stronger than revenue equivalence as it characterizes the possible payments for *every* mechanism.

Theorem 10. *Let G be a convex function on $\text{Conv}(\mathcal{T})$, $d = \{d_t\}_{t \in \mathcal{T}}$ a selection of its subgradients on \mathcal{T} , $S \subseteq \mathcal{T}$ non-empty, $t^* \in \mathcal{T} \setminus S$, and c be given. Then there exists a convex G' on $\text{Conv}(\mathcal{T})$ agreeing with G on S , with $\{d_t\}_{t \in \mathcal{T}} \in \partial G'$ and $G'(t^*) = c$, if and only if*

$$\sup_{t_0 \in S} G(t_0) + P_d(t_0, t^*) \leq c \leq \inf_{t_0 \in S} G(t_0) - P_d(t^*, t_0) \quad (22)$$

Proof. Given such a G' , the LHS of (22) becomes $\sup_{t_0 \in S} G'(t_0) + P_d(t_0, t^*) \leq G'(t^*)$. Applying the definition of P_d (16) and then repeatedly applying the subgradient inequality (2) yields the desired inequality. Similarly, the RHS of (22) can be rewritten as $G'(t^*) + P_d(t^*, t_0) \leq G'(t_0)$ for all $t_0 \in S$, and the definition and subgradient inequality applied.

Now suppose (22) holds. Let $G'(t) \doteq \max\{c + P_d(t^*, t), \sup_{t_0 \in S} G(t_0) + P_d(t_0, t)\}$. By Theorem 7, d satisfies CMON, so by Lemma 1 G' is convex, finite-valued on $\text{Conv}(\mathcal{T})$, and has $\{d_t\} \in \partial G'$. Hence, we need only show that G' agrees with G on S and has $G'(t^*) = c$.

First, fixing any $t \in S$, we will establish the following:

$$G(t) = \sup_{t_0 \in S} G(t_0) + P_d(t_0, t). \quad (23)$$

As $P_d(t, t) = 0$ from Lemma 1(3), we have $G(t) = G(t) + P_d(t, t) \leq \sup_{t_0 \in S} G(t_0) + P_d(t_0, t)$. Furthermore, $G(t_0) + P_d(t_0, t) \leq G(t)$ for all $t_0 \in \mathcal{T}$ by repeated application of the subgradient inequality (2). Hence, we have $\sup_{t_0 \in S} G(t_0) + P_d(t_0, t) \leq G(t)$ as well.

By eq. (23), we can write $G'(t) = \max\{c + P_d(t^*, t), G(t)\}$ when $t \in S$. But by the RHS of eq. (22), we see $c + P_d(t^*, t) \leq G(t)$, so $G'(t) = G(t)$. Similarly, applying the LHS of eq. (22) and $P_d(t^*, t^*) = 0$ to the definition of $G'(t^*)$, we have $G'(t^*) = c$.

Viewed through Theorem 10, revenue equivalence holds when the upper and lower bounds from (22) match after the value of G is fixed as a single point. This allows us to derive a necessary and sufficient condition for revenue equivalence that is equivalent to that given by Heydenreich et al. [37] and actually applies to all affine scores. For example, this gives a revenue equivalence theorem for mechanisms with partial allocation.

Corollary 13 (Revenue Equivalence). *Let a truthful affine score $A : \mathcal{T} \times \mathcal{T} \rightarrow \mathbb{R}$ be given, and $d = \{d_t\}_{t \in \mathcal{T}}$ be the corresponding selection of subgradients from (3). Then every truthful affine score $A' : \mathcal{T} \times \mathcal{T} \rightarrow \mathbb{R}$ with the same corresponding selection of subgradients differs from A by a constant (i.e. $A(t', t) = A'(t', t) + c$) if and only if $P_d(t', t) + P_d(t, t') = 0$ for all $t, t' \in \mathcal{T}$.*

Proof. We will show that the convex function G from eq. (3) is unique up to a constant if and only if $P_d(t', t) + P_d(t, t') = 0$ for all $t, t' \in \mathcal{T}$.

For the forward direction, let $t_0 \in \mathcal{T}$ be arbitrary. Then for all $t \in \mathcal{T}$, taking (22) with $S = \{t_0\}$ and $G(t) \doteq c + P_d(t_0, t)$ gives the condition $G(t_0) + P_d(t_0, t) \leq G'(t) \leq G(t_0) - P_d(t, t_0)$ for the value of $G'(t)$. But as $P_d(t, t_0) = -P_d(t_0, t)$ we have $G'(t) = P_d(t_0, t) + G(t_0) = G(t)$ for all t .

For the reverse direction, assume $P_d(t^1, t^2) \neq -P_d(t^2, t^1)$ for some $t^1, t^2 \in \mathcal{T}$, and let $G^1(t) \doteq P_d(t^1, t)$ and $G^2(t) \doteq P_d(t^1, t^2) + P_d(t^2, t)$. We easily check from Lemma 1(3) that $G^1(t^2) = G^2(t^2) = P_d(t^1, t^2)$, but we have $G^1(t^1) = 0$ while $G^2(t^1) = P_d(t^1, t^2) + P_d(t^2, t^1) \neq 0$.

We note that these two results are similar to results of Kos and Messner [43]. The main novelties in our version are showing that every value in the interval yields a convex function (as opposed to merely the extremal ones), the ability to characterize possible values after the values at multiple points are fixed (as opposed to a single point), and the framing in terms of convex analysis.

The conditions given by Theorem 10 and Corollary 13, while general, are not particularly intuitive. However, there are a number of special cases where they do have natural interpretations for mechanism design. The first is when the set of types is finite. In this setting (explored in an auction theory context in, e.g., [25]) it is well known that revenue equivalence does not hold. The finite set of constraints (22) can be used in general as a linear program to, e.g., maximize revenue (see Section 6.5.2 of [65] for an example). In particular cases, they may become simple enough to have a nice characterization. For example, in the single-parameter setting only a linear number of paths need be considered. This setting is illustrated in Figure 3.




Fig. 3. Consider a one-dimensional setting with type space $\mathcal{T} = \{0, 1, 2\}$ and $d_0 = 0, d_1 = 1, d_2 = 2$. In (a), we fix $G(0) = 0$, yielding a range of possible values dictated by the subgradients: $0 \leq G(1) \leq 1$ and $1 \leq G(2) \leq 3$. We can pick any point in the resulting set and fix G there. However, we cannot pick any increasing function: in (b), we fix $G(1) = 0.5$, restricting $G(2)$ to the interval $[1.5, 2.5]$.

More broadly, as we saw in the proof of Theorem 8, the (supremum over the) sum can often be interpreted as an integral. In particular, the fact that G is convex guarantees that (under mild conditions) integrals of a selection of its subgradient are path independent and the integral from t to t' gives $G(t') - G(t)$. If \mathcal{T} is connected by smooth paths (e.g. if it is convex), this means that \mathcal{T} satisfies revenue equivalence for all implementable mechanisms (previously shown under a somewhat different notion of the set of types [37]). As it is particularly simple to prove, we state the version for convex \mathcal{T} .

Corollary 14. *Let \mathcal{T} be convex, a truthful affine score $A : \mathcal{T} \times \mathcal{T} \rightarrow \overline{\mathbb{R}}$ be given, and $\{dG_t\}_{t \in \mathcal{T}}$ be the corresponding selection of subgradients from (3). Then any truthful affine score $A' : \mathcal{T} \times \mathcal{T} \rightarrow \mathbb{R}$ with the same corresponding selection of subgradients differs from A by a constant (i.e. $A(t', t) = A'(t', t) + c$).*

Proof. By Theorem 1, we know that A and A' only differ only in their choice of convex function G . However, each choice has the same selection of subgradients, and two convex functions with the same selection of subgradients differ by a constant [59]. For intuition, see the construction of G by integrating its subgradients in the proof of Theorem 8.

E Additional Applications

In this section, we demonstrate the power of our characterization theorem with several additional applications.

E.1 Decision Rules

Theorem 1 also generalizes Gneiting and Raftery's [33] characterization to settings beyond eliciting a single distribution. For example, a line of work has considered a setting where a decision maker needs to select from a finite set \mathcal{D} of decisions and so desires to elicit the distribution over outcomes conditional on selecting each alternative [20, 21, 56]. Since only one decision will be made and so only one conditional distribution can be sampled, simply applying a standard proper scoring rule generally does not result in truthful behavior. Applying Theorem 1 to this setting characterizes what expected scores must be, from which many of the results in these papers follow.

As an illustration, consider the model of proper scoring rules for decision rules [21]. There is a decision maker who will select an action from a set $\mathcal{A} = \{1, \dots, n\}$. Once an action is selected, some outcome from the set $\mathcal{O} = \{o_1, \dots, o_m\}$ will be realized, where the probability of each outcome depends on the action chosen. The decision maker seeks to elicit from an expert the probability $P_{i,o}$ of outcome o occurring given that action i is chosen. The decision maker uses a fixed decision rule $D : \mathcal{P} \rightarrow \Delta(\mathcal{A})$, where $D_i(P)$ is the probability of choosing action i given the expert reported the matrix P . The decision maker rewards the expert using a (regular) scoring rule that depends on the action chosen

$S : \mathcal{A} \times \mathcal{O} \times \mathcal{P} \rightarrow \mathbb{R} \cup \{-\infty\}$. For brevity, we write $S_{i,o}(P)$. Given a belief P and report Q , we can write the expert's expected score as

$$V(Q, P) = \sum_{i,o} D_i(Q) P_{i,o} S_{i,o}(Q)$$

The definition of (strict) properness for a particular decision rule then follows naturally.

Definition 13. A regular scoring rule S is proper for a decision rule D if

$$V(P, P) \geq V(Q, P)$$

for all P and all $Q \neq P$. It is strictly proper for the decision rule if the inequality is strict.

As this V is an affine score, we can immediately apply our theorem to derive Chen and Kash's [21] characterization and even extend it to the case where the set of probability matrices is not convex. In the theorem statement we make use of the Frobenius inner product $P:Q \doteq \sum_{i,o} P_{i,o} Q_{i,o}$.

Theorem 11. Given a set of probability matrices $\mathcal{P} \subseteq \Delta(\mathcal{O})^n$ A regular scoring rule is (strictly) proper for a decision rule D if and only if

$$S_{i,o}(Q) = \begin{cases} G(Q) - G_Q:Q + \frac{G_{Q,i,o}}{D_i(Q)} & D_i(Q) > 0 \\ \Pi_{i,o}(Q) & D_i(Q) = 0 \end{cases}$$

where $G : \text{Conv}(\mathcal{P}) \rightarrow \mathbb{R} \cup \{-\infty\}$ is a (strictly) convex function, G_Q is a subgradient of G at the point Q with $G_{Q,i,o} = 0$ when $D_i(Q) = 0$, and $\Pi_{i,o} : \mathcal{P} \rightarrow \mathbb{R} \cup \{-\infty\}$ is an arbitrary function that can take a value of $-\infty$ only when $Q_{i,o} = 0$.

Proof. By Theorem 1, S is (strictly) proper for D if and only if there exists a (strictly) convex G such that $V(Q, P) = G(Q) + dG_Q(P - Q)$. That is,

$$\sum_{i,o} D_i(Q) P_{i,o} S_{i,o}(Q) = G(Q) - G_Q:Q + \sum_{i,o} G_{Q,i,o} P_{i,o},$$

or for all i such that $D_i(Q) \neq 0$,

$$S_{i,o}(Q) = G(Q) - G_Q:Q + \frac{G_{Q,i,o}}{D_i(Q)}$$

When $D_i(Q) = 0$, S is unconstrained (other than the minimal requirements regarding $-\infty$ for regularity). However, note that our affine score is restricted in that, because $D_i(Q)$ is fixed, some choices in \mathcal{A} are not possible to select as subgradients. In particular, it must be that $G_{Q,i,o} = 0$ when $D_i(Q) = 0$

E.2 Proper losses for partial labels

Several variants of proper losses have appeared in the machine learning literature, one of which is the problem of estimating the probability distribution of labels for an item when the training data may contain several noisy labels, possibly not even including the correct label. (This is frequently the case, for example, when using crowdsourced labels for items.) More formally, one wishes to estimate $p \in \Delta_n$ where the true label $y \in [n]$ is drawn from p . However, instead of observing a sample $y \sim p$ and designing a proper loss $\ell(\hat{p}, y)$, one instead only observes some noisy set of labels $S \subseteq [n]$. Hence, the task is to design a loss $\ell(\hat{p}, S)$ which when minimized over one's data yields accurate estimates of the true p .

Recently this problem was studied in [24] under the assumption that $S \sim q$ where $q = Mp$ for some known $M \in \mathbb{R}^{2^n \times n}$, meaning if the observed label is drawn from p , the noisy set of labels is drawn from Mp (using some indexing of the sets, say lexicographical). Cid-Sueiro (in his Theorem 4.3) provides a characterization of all proper losses for an even more general version of this setting where M is known only to be a member of a class rather than exactly and we want the loss to be proper regardless of which member it is. Note that the (negative) payoff $\mathbb{E}_{S \sim Mp}[\ell(\hat{p}, S)] = \ell(\hat{p}, \cdot)^\top Mp$ is linear in the underlying distribution p , so our Theorem 1 applies and allows us to recover his characterization result. We refrain from introducing the model necessary to explicitly state this result as it would require an excessive number of definitions. Note that this is essentially a latent observation setting, and the fact that what we observe is a set of labels is in no way necessary — any observed outcome whose distribution has a linear (or affine) relationship with the latent outcome would suffice to apply our theorem.

E.3 Mechanism design with partial allocation

Several mechanism design settings considered in the literature have some form of *exogenous* randomization, in that “Nature” chooses some outcome ω according to some (often unknown) distribution, which may in turn depend on the allocation chosen by the mechanism. Examples include sponsored search auctions [28], multi-armed bandit mechanisms [10], and recent work on daily deals [17]. The work of Cai et al. [17] introduces a very general model for such settings, which we now describe.

Let \mathcal{O} be a set of allocations, and for each allocation o and each agent i , let $\Omega^{i,o}$ be some set of outcomes (e.g. which agent wins an auction for the opportunity to advertise a special offer from its business). Agents each have a valuation function $v^i : \mathcal{O} \rightarrow \mathbb{R}$ and a set of beliefs $p^{i,o} \in \Delta(\Omega^{i,o})$ for each allocation $o \in \mathcal{O}$ (an expected value for getting to advertise and a probability distribution over the number of customers who accept the deal). The mechanism aggregates all of this information into a single allocation o , and additionally chooses some payoff function $s^i : \Omega^{i,o} \rightarrow \mathbb{R}$, so that the final utility of agent i is $v^i(o) + \mathbb{E}_{p^{i,o}}[s^i]$ (the winning agent both gets to advertise and accepts a scoring rule contract regarding its prediction of the number of customers). A mechanism

is truthful if for all values of v and p for the other agents, agent i maximizes her total utility by reporting v^i and $p^i \doteq (p^{i,o})_{o \in \mathcal{O}}$ truthfully. For additional examples, the standard sponsored search setting has $\Omega^{i,o} = \{\text{click}, \text{no click}\}$ for o such that i is allocated a slot, and the probabilities $p^{i,o}$ are assumed to be public knowledge. Moreover, the decision rules framework discussed above is a single-agent special case with $v \equiv 0$ and $\Omega^o = \Omega^{o'} = \Omega$ for all $o \in \mathcal{O}$ (of course, unlike the notation above, $o \in \mathcal{O}$ is the allocation/decision, and Ω is the set of outcomes).

We first observe that this model can easily be cast as an affine score, as follows. For simplicity, we fix some agent i and focus on the single-agent case; as discussed several times above, this is essentially without loss of generality. The type space is simply the combined private information of the agent,

$$\mathcal{T} = \left\{ (v, p) : v \in \mathcal{O} \rightarrow \mathbb{R}, p \in \prod_{o \in \mathcal{O}} \Delta(\Omega^{i,o}) \right\}. \quad (24)$$

The utility of the agent upon allocation and payoff o, s is simply $v(o) + \mathbb{E}_{p^o}[s] = \text{Eval}_o[v] + s \mathbb{1}_o^\top p$, which is linear in the type $t = (v, p)$ and therefore affine. (Here we represent p as a matrix in $\mathbb{R}^{\mathcal{O} \times \Omega^{i,o}}$ and $s \in \mathbb{R}^{\Omega^{i,o}}$, and define $\mathbb{1}_o$ to be the standard vector with 1 at entry o and 0 elsewhere.) Thus, letting $t = (v, p)$, we can represent this as an affine score:

$$\mathbf{A}(t', t) = v(o(t')) + \mathbb{E}_{p^{o(t')}}[s(t')]. \quad (25)$$

Motivated by incorporating the utilities of the end consumers in a daily deal setting, Cai et al. [17] ask when one can implement an allocation rule of the form $f(v, p) = \text{argmax}_{o \in \mathcal{O}} v(o) + g^o(p^o)$; in other words, when does there exist some choice of score $s(v, p) \in \mathbb{R}^{\Omega^{i,f(v,p)}}$ making f truthful. They conclude that this can be done if and only if g^o is convex for each $o \in \mathcal{O}$. It is interesting, and perhaps illuminating, to view this question in terms of our affine score framework.

Stepping back for a moment, consider a type space $\mathcal{T} \subseteq \mathcal{V} = \mathcal{V}^X \times \mathcal{V}^Y$ which partitions into two (subsets of) subspaces. (Note that \mathcal{V}^Y no longer need be restricted to probability distributions.) We wish to know when a function $f : \mathcal{T} \rightarrow \text{Lin}(\mathcal{V}^X \rightarrow \mathbb{R})$ is implementable, in the sense that there exists some truthful affine score $\mathbf{A} : \mathcal{T} \times \mathcal{T} \rightarrow \overline{\mathbb{R}}$ with score set $\mathcal{A} \subseteq \text{Aff}(\mathcal{V} \rightarrow \mathbb{R})$, and some $h : \mathcal{T} \rightarrow \text{Aff}(\mathcal{V}^Y \rightarrow \mathbb{R})$ such that $\mathbf{A}(t', t) = f(t')(t^X) + h(t')(t^Y)$, where of course $t = (t^X, t^Y)$. That is, when can we complete the partial ‘‘allocation’’ f into a truthful affine score?

For convenience, for each $a \in \mathcal{A}$ we write $X(a) \in \text{Lin}(\mathcal{V}^X \rightarrow \mathbb{R})$ to be the linear part of a on \mathcal{V}^X , and $Y(a)$ to be the *affine* part of a on \mathcal{V}^Y . Then we have that f is implementable if and only if

$$f(t) \in \text{argsup}_{x \in X(\mathcal{A})} \left\{ x(t^X) + \sup_{\substack{a \in \mathcal{A}(\mathcal{T}) \\ X(a)=x}} \{Y(a)(t^Y)\} \right\} \quad (26)$$

To see this, one direction follows from the fact that an affine score is truthful if and only if

$$\mathbf{A}(t) \in \operatorname{argsup} \{a(t) : a \in \mathbf{A}(\mathcal{T})\}, \quad (27)$$

by taking the supremum first over $X(\mathcal{A})$ and then over the rest. For the other direction, note that taking $\mathbf{A}(t', t) = f(t')(t^X) + y(t')(t^Y)$ where y is in the argsup of the supremum of eq. (26) gives a truthful affine score.

Returning to the special case of daily deals, let us denote by $a_{o,s} \in \mathcal{A}$ the function $(v, p) \mapsto v(o') + \mathbb{E}_{p^{o'}}[s]$. We now see that $f(v, p)$ is implementable if and only if it satisfies

$$f(v, p) \in \operatorname{argsup}_{o \in \mathcal{O}} \left\{ v(o) + \sup_{s: a_{o,s} \in \mathcal{A}(\mathcal{T})} \{\mathbb{E}_{p^o}[s]\} \right\}. \quad (28)$$

Thus, letting $g^o(p^o) = \sup \{\mathbb{E}_{p^o}[s] : a_{o,s} \in \mathcal{A}\}$, we see that g^o is convex as the supremum of affine functions. Moreover, given any collection of convex functions $\{g^o\}_{o \in \mathcal{O}}$, where $g^o : \Delta(\Omega^{i,o}) \rightarrow \mathbb{R}$, we can define $S^o \doteq \{\omega \mapsto g(p) + dg(\mathbb{1}_\omega - p) : p \in \operatorname{dom}(g)\}$ and $\mathcal{A} \doteq \{a_{o,s} : o \in \mathcal{O}, s \in S^o\}$, thus recovering each g^o in the above expression. It then only remains to show that no other nonconvex function can serve in the argsup; for this one may appeal to the argument of Cai et al. [17] which observes that the indifference points between different allocations is fixed, thus determining the function in the argsup up to a constant.

E.4 Responsive Lotteries

Utility functions consistent with particular preferences are only unique up to an affine transformation. (Since there are no payments, multiplying the value of each outcome by a constant or adding a constant to the value for each outcome has no effect on the optimal lottery for an agent). Therefore, we state our theorem for utilities that have been projected onto the unit sphere.

Theorem 12. *Let $\mathcal{T} = \{t \in \mathbb{R}^n : \|t\|_2 = 1\}$ be the unit sphere in \mathbb{R}^n , and let a truthful affine score $\mathbf{A} : \mathcal{T} \times \mathcal{T} \rightarrow \mathbb{R}$ with score set $\mathcal{A} \subseteq \operatorname{Lin}(\mathbb{R}^n \rightarrow \mathbb{R}) \cong \mathbb{R}^n$ be given. Then $S(t) \doteq \mathbf{A}(t, \cdot)$ is surjective and continuous (as a function to \mathbb{R}^n) and \mathbf{A} is strictly truthful if and only if \mathcal{A} is the boundary of a compact and strictly convex set $K \subset \mathbb{R}^n$.¹¹ S is additionally injective if and only if K is additionally smooth.*

Before proving Theorem 12, we state the relevant corollary and then provide useful definitions.

¹¹ We define a convex set C to be *strictly convex* if no point x on the boundary of C can be expressed as a convex combination of other points in C (i.e. x is extreme). C is *smooth* if each point on the boundary of C has a unique unit normal vector. See Appendix E.4 for formal definitions.

Corollary 15. *A lottery rule f satisfies incentive compatibility and rational uniqueness if and only if $f(x) = \operatorname{argmax}_{p \in K} \langle x, p \rangle$ for $K \subset \Delta_n$ compact and strictly convex relative to Δ_n . Moreover, f additionally satisfies rational invertibility (and thus is truthful dominant) if and only if K is additionally smooth.*

Proof. Project the utilities and probability simplex onto the set $V = \{x \in \mathbb{R}^n : \sum_i x_i = 0\}$, which only changes the expected utilities by a constant. Then express these vectors in a basis for $V \cong \mathbb{R}^{n-1}$, and normalize the utilities (only scaling them) to get the unit sphere in V , and apply Theorem 12.

We denote by ∂K the boundary of the set $K \subseteq \mathbb{R}^n$.

Definition 14. *Given a compact convex set $K \subset \mathbb{R}^n$, we define the exposed face $F_K(t)$ in direction $t \neq 0$ and the normal cone $N_K(k)$ at point $k \in \partial K$ by*

$$F_K(t) = \operatorname{argmax}_{k \in K} \langle t, k \rangle, \quad N_K(k) = \{t \in \mathbb{R}^n : k \in F_K(t)\}. \quad (29)$$

Definition 15. *We say K is strictly convex if $F_K(t)$ is a singleton for all $t \neq 0$. Dually, we say K is smooth if $N_K(k)$ is a ray (i.e. $\{\alpha t : \alpha \geq 0\}$ for some $t \neq 0$) for all $k \in \partial K$.*

Proof (of Theorem 12.). We begin with the first part of the theorem. Let K be compact and strictly convex, and $\mathcal{A} = \partial K$. Then as A is truthful, we must have $S(t) \in \operatorname{argsup}_{a \in \mathcal{A}} \langle t, a \rangle$. As $\mathcal{A} = \partial K$, we may also write $S(t) \in \operatorname{argmax}_{k \in K} \langle t, k \rangle$. Now by strict convexity of K , we have for every $a \in \mathcal{A} = \partial K$, there exists some $t \in \mathcal{T}$ such that $\{a\} = \operatorname{argmax}_{k \in K} \langle t, k \rangle$, giving us both surjectivity and strict truthfulness (as $S(t) = a$). Continuity follows immediately from Berge's Maximum Theorem [53].

For the converse, let S be strictly truthful, surjective, and continuous. By standard arguments, since \mathcal{T} is a compact subset of \mathbb{R}^n , we have $\mathcal{A} = S(\mathcal{T})$ is compact as a continuous image of a compact set. Thus, $K \doteq \operatorname{Conv}(\mathcal{A})$ is a compact convex set. Letting $F_K(t) \doteq \operatorname{argmax}_{k \in K} \langle t, k \rangle$ be the exposed face of K in direction t , we will now show $F_K(t) = \{S(t)\}$. First, observe that the extreme points of K , $\operatorname{ext}(K)$, are a subset of \mathcal{A} (otherwise we have $k \in \operatorname{ext}(K) \setminus \mathcal{A}$, so $K \setminus \{k\}$ is a convex set containing \mathcal{A} , contradicting the definition of $K = \operatorname{Conv}(\mathcal{A})$). Now we may apply [63, Proposition A.2.4.6] to express the argmax in terms of the extreme points of K , giving us

$$F_K(t) \doteq \operatorname{argmax}_{k \in K} \langle t, k \rangle = \operatorname{Conv} \left(\operatorname{argmax}_{k \in \operatorname{ext}(K)} \langle t, k \rangle \right) \subseteq \operatorname{Conv} \left(\operatorname{argmax}_{a \in \mathcal{A}} \langle t, a \rangle \right) = \{S(t)\}.$$

As K is compact, $F_K(t)$ is nonempty, and thus $F_K(t) = \{S(t)\}$, and additionally we conclude $S(t) \in \operatorname{ext}(K)$. Hence $\mathcal{A} = S(\mathcal{T}) \subseteq \operatorname{ext}(K)$, and as we concluded the reverse conclusion above, we have $\mathcal{A} = \operatorname{ext}(K)$. We now apply [63, Proposition C.3.1.5] to obtain $\partial K = \bigcup_{t \in \mathcal{T}} F_K(t)$, which in turn gives $\partial K = \mathcal{A}$ by surjectivity. Finally, as $\operatorname{ext}(K) = \mathcal{A} = \partial K$, we have strict convexity of K .

For the final statement of the theorem, we note that by [63, Proposition C.3.1.4], we have $k \in F_K(t) \iff t \in N_K(k)$. By the above, we already have $F_K(t) = \{S(t)\}$ for all $t \in \mathcal{T}$, which implies $N_K(k) \cap \mathcal{T} = \{t : S(t) = k\}$. Hence, $N_K(a)$ is a ray for all $a \in \mathcal{A}$ if and only if S is injective.

F Property Characterization

Here we go through the proof of Theorem 2 (restated as Theorem 13) in detail. To begin, note that the simplest way to come up with an elicitable property is to induce one from an affine score. For any $A : \mathcal{R} \times \mathcal{T} \rightarrow \overline{\mathbb{R}}$ with score set $\mathcal{A} \subset \text{Aff}(\mathcal{V} \rightarrow \overline{\mathbb{R}})$, the property

$$\Gamma^A : t \rightarrow \operatorname{argsup}_{r \in \mathcal{R}} A(r, t) \quad (30)$$

is trivially elicited by A if this argsup is well defined.

Observe also that any affine score A eliciting Γ gives rise to a truthful affine score in the original sense — in fact, this is a version of the *revelation principle* from mechanism design. For each t let $r_t \in \Gamma(t)$ be a report choice for t ; then the affine score $A^\mathcal{T}(t', t) \doteq A(r_{t'}, t)$ is truthful. Moreover, by our choices of $\{r_t\}$, we have

$$G(t) \doteq \sup_{t' \in \mathcal{T}} A^\mathcal{T}(t', t) = \sup_{r \in \mathcal{R}} A(r, t). \quad (31)$$

Of course, in general, $A^\mathcal{T}$ will not be strictly truthful, since by definition, any reports t', t'' with $r_{t'} = r_{t''}$ will have $A^\mathcal{T}(t', \cdot) \equiv A^\mathcal{T}(t'', \cdot)$. Thus we may think of a property as *refining* the notion of strictness for a truthful affine score. The connection we draw in Theorem 13 is that, in light of (31), a property Γ therefore specifies the portions of the domain of \mathcal{T} where G must be “flat”. To get at the connection between properties and “flatness”, we start with a technical lemma which shows that having the same subgradient at two different points is equivalent to G being flat in between.

Lemma 2. *Let $G : \operatorname{Conv}(\mathcal{T}) \rightarrow \overline{\mathbb{R}}$ be convex with $G(\mathcal{T}) \subseteq \mathbb{R}$, and let $d \in \partial G_t$ for some $t \in \mathcal{T}$. Then for all $t' \in \mathcal{T}$,*

$$d \in \partial G_{t'} \iff G(t) - G(t') = d(t - t').$$

Proof. First, the forward direction. Applying the subgradient inequality (2) at t' for $dG_t = d$ and at t for $dG_{t'} = d$, we have

$$\begin{aligned} G(t') &\geq G(t) + d(t' - t) \\ G(t) &\geq G(t') + d(t - t'), \end{aligned}$$

from which the result follows (as $G(t)$ and $G(t')$ are finite).

For the converse, assume $G(t) = G(t') + d(t - t')$ and let $t'' \in \mathcal{T}$ be arbitrary. Note that $d(t) \in \mathbb{R}$ as $d \in \partial G_t$, so $d(t') \in \mathbb{R}$ as well. Then using the subgradient inequality (2),

$$G(t') + d(t'' - t') = G(t') + d(t'' - t) + d(t - t') = G(t) + d(t'' - t) \leq G(t''). \square$$

We are now ready to state our first characterization, which in essence says that eliciting a property Γ is equivalent to eliciting subgradients of a convex function G . Intuitively, by truthfulness the linear part of $A(r, \cdot)$ must be a subgradient of G at all $t \in \Gamma_r$. The lemma shows that this equivalent to flatness, which means we can calculate G on Γ_r set by picking any type $t_r \in \Gamma_r$ and following the subgradient. Since all choices of t_r lead to the same value, we could just as easily ask for this subgradient $\varphi(r)$ to be reported directly. As subgradients are functions (in this case from \mathcal{T} to $\overline{\mathbb{R}}$), we use the curried notation $\varphi(r)(t)$ for the application of this function.

Note that as we allow $A(r, t)$ to take on values in the extended reals to capture scoring rules such as the log score, we again need a notion of regularity — an affine score A is Γ -regular if $A(r, t) < \infty$ always and $A(r, t) \in \mathbb{R}$ whenever $r \in \Gamma(t)$. We define Γ -regular linear and affine families similarly.¹²

Theorem 13. *Let non-redundant property $\Gamma : \mathcal{T} \rightrightarrows \mathcal{R}$ and Γ -regular affine score $A : \mathcal{R} \times \mathcal{T} \rightarrow \overline{\mathbb{R}}$ be given. Then A elicits Γ if and only if there exists some convex $G : \text{Conv}(\mathcal{T}) \rightarrow \overline{\mathbb{R}}$ with $G(\mathcal{T}) \subseteq \mathbb{R}$, some $\mathcal{D} \subseteq \partial G$, and some bijection $\varphi : \mathcal{R} \rightarrow \mathcal{D}$ with $\Gamma(t) = \varphi^{-1}(\mathcal{D} \cap \partial G_t)$, such that for all $r \in \mathcal{R}$ and $t \in \mathcal{T}$,*

$$A(r, t) = G(t_r) + \varphi(r)(t - t_r), \quad (32)$$

where $\{t_r\}_{r \in \mathcal{R}} \subseteq \mathcal{T}$ satisfies $r' \in \Gamma(t_{r'})$ for all r' .

Proof. For the converse, let A be given of the form (32). We show that it elicits Γ , i.e. $\Gamma(t) = \text{argsup}_{r \in \mathcal{R}} A(r, t)$. The third line of the derivation applies Lemma 2.

$$\begin{aligned} r \in \Gamma(t) &\iff r \in \varphi^{-1}(\mathcal{D} \cap \partial G_t) \\ &\iff \varphi(r) \in \mathcal{D} \cap \partial G_t \\ &\iff A(r, t) = G(t) \\ &\iff r \in \text{argsup}_{r' \in \mathcal{R}} A(r', t) \end{aligned}$$

For the forward direction, assume that affine score A elicits Γ . For each r , we may extend $A(r, \cdot)$ to all $\hat{t} \in \text{Conv}(\mathcal{T})$ by linearity as in the proof of Theorem 1, whence we may define $G(\hat{t}) \doteq \sup_{r \in \mathcal{R}} A(r, \hat{t})$, which is finite for $\hat{t} \in \mathcal{T}$ as A is Γ -regular. We wish to show that the choice $\varphi : r \mapsto A_\ell(r, \cdot)$ suffices, where A_ℓ denotes the linear part of A , with \mathcal{D} the range of φ and $\{t_r\}$ arbitrary satisfying the theorem. Given this construction, we need to check each of the following.

1. G is convex with subgradients $\varphi(\Gamma(t)) \subseteq \partial G_t$ Let t and $r \in \Gamma(t)$ be given. We show that $\varphi(r)$ satisfies the property of a subgradient at t , and thus G is convex

¹² The family $\{\ell_r \in \text{Lin}(\mathcal{V} \rightarrow \overline{\mathbb{R}})\}_{r \in \mathcal{R}}$ is Γ -regular if $\ell_r(t) \in \mathbb{R}$ for all $t \in \Gamma_r$, and $\ell_r(t') \in \mathbb{R} \cup \{-\infty\}$ for $t' \notin \Gamma_r$. Likewise for Γ -regular affine functions.

with appropriate subgradients.

$$\begin{aligned}
G(t) + \varphi(r)(t' - t) &= \sup_{r' \in \mathcal{R}} A(r', t) + A_\ell(r, t' - t) \\
&= A(r, t) + A_\ell(r, t' - t) = A(r, t') \\
&\leq \sup_{r' \in \mathcal{R}} A(r', t') = G(t')
\end{aligned} \tag{33}$$

2. *A satisfies eq. (32)* This follows from (33) with $t = t_r$, as $r \in \Gamma(t_r)$.

3. *φ is a bijection* By definition, \mathcal{D} is the range of φ , so we only need to check that it is injective. Suppose for contradiction that $\varphi(r) = \varphi(r')$. Then, by definition, $A_\ell(r, \cdot) = A_\ell(r', \cdot)$. Since A elicits Γ , we have $A(r, \cdot) = A(r', \cdot)$. But then $r \in \Gamma(t) \iff r' \in \Gamma(t)$, contradicting Γ being non-redundant.

4. *$\Gamma(t) = \varphi^{-1}(\mathcal{D} \cap \partial G_t)$* We already know that $\varphi(\Gamma(t)) \subseteq \partial G_t$, so since \mathcal{D} is the range of φ we have $\varphi(\Gamma(t)) \subseteq \mathcal{D} \cap \partial G_t$. For the other direction, $d \in \mathcal{D} \cap \partial G_t$ is $\varphi(r)$ for some r . Then by Lemma 2, $A(r, t) = G(t_r) + \varphi(r)(t - t_r) = G(t)$, so $r \in \Gamma(t)$.

As a corollary, we also obtain a better understanding of weak elicitation, which we will need in the following sections.

Corollary 16. *Let non-redundant property $\Gamma : \mathcal{T} \rightrightarrows \mathcal{R}$ and Γ -regular affine score $A : \mathcal{R} \times \mathcal{T} \rightarrow \overline{\mathbb{R}}$ be given. Then A weakly elicits Γ if and only if A satisfies (32) with the weaker condition that $\Gamma(t) \subseteq \varphi^{-1}(\mathcal{D} \cap \partial G_t)$.*

Proof. Given any affine score A , and defining Γ^A as in (30), we see that A weakly elicits Γ if and only if $\Gamma(t) \subseteq \Gamma^A(t)$ for all t . Now let A weakly elicit Γ . As A trivially elicits Γ^A , we apply Theorem 13 and now have in particular $r \in \Gamma(t) \implies r \in \Gamma^A(t) \implies \varphi(r) \in \partial G_t$. For the converse, simply define $\Gamma^A(t) = \{r \in \mathcal{R} \mid \varphi(r) \in \partial G_t\}$. By Theorem 13, A elicits Γ^A , and by assumption we have $\Gamma(t) \subseteq \Gamma^A(t)$ for all t .

Using Corollary 16, we see that an affine score A is truthful if and only if it weakly elicits $\Gamma : t \mapsto \{t\}$. Hence, Theorem 13 and Corollary 16 are actually generalizations of Theorem 1. Of course, we also obtain the following corollary characterizing non-redundant properties.

Corollary 17. *Non-redundant $\Gamma : \mathcal{T} \rightrightarrows \mathcal{R}$ is elicitable if and only if exists there some convex $G : \text{Conv}(\mathcal{T}) \rightarrow \overline{\mathbb{R}}$ with $G(\mathcal{T}) \subseteq \mathbb{R}$, some $\mathcal{D} \subseteq \partial G$, and some invertible $\varphi : \mathcal{R} \rightarrow \mathcal{D}$ such that $\Gamma(t) = \varphi^{-1}(\mathcal{D} \cap \partial G_t)$.*

An important question which would give stronger characterizations is the following:

Question 1. Given non-redundant elicitable Γ , what are all pairs G, \mathcal{D} such that there exists some bijection φ satisfying Theorem 13? Equivalently (up to redundancy), given a convex function G with subgradient level sets $LS_G(d) = \{t : d \in \partial G_t\}$, what are all the convex functions G' with $LS_{G'} \equiv LS_G$?

In Section 5 we will see that the answer to this question has a lot of structure in the case where \mathcal{R} is finite. In the general case, certainly performing a homothet of the subgradients of G (i.e. scaling G and adding a linear term), will preserve the elicitation structure. However, surely more can be done—the property

$$\Gamma(t) = \begin{cases} \{t - 1\} & \text{if } t < 0 \\ \{0\} & \text{if } t = 0 \\ \{t + 1\} & \text{if } t > 0 \end{cases} \quad (34)$$

can be elicited with both $G(t) = |t| + t^2/2$ and $G(t) = t^2/2$, which is not a homothet transformation.

While we do not have a complete answer to Question 1 our characterization sheds new light on the structure of elicitable properties in two directions. First, in the scoring rules literature, it is common to assume strong conditions on Γ and \mathcal{R} , such as Γ being a function rather than a multivalued map, and Γ being linear [1] or real-valued [47] to achieve characterizations. In contrast, Theorem 13 allows for an extremely general Γ and \mathcal{R} and shows us how to construct affine scores for such properties. Second, we can identify features that all elicitable properties share, which provides a means to prove that specific properties are not elicitable.

G Duality in elicitation

We saw from Theorem 13 that in a strong sense an elicitable property Γ is like a subgradient mapping of a convex function. We now turn to removing the word “like” from the sentence above — we look at properties which *are* subgradient mappings. This exploration has two main benefits. First, it gives us a concrete tool to reason about properties, by working directly with a convex function rather than through some map φ . Second, it gives a new framework to discuss duality in elicitation, as has been observed between scoring rules and prediction markets [1, 22].

G.1 Direct elicitation

Now that we have formalized the relationship between the report space and subgradients of convex functions, we can see what the “canonical” properties look like: those which are (subsets of) subgradient mappings of a convex function. For these properties, we can talk about *direct elicitation*, which roughly speaking means removing the intermediary φ between \mathcal{R} and ∂G . In fact, for such “canonical” properties, we can even talk about a convex function *itself* eliciting Γ .

Definition 16. *A property $\Gamma : \mathcal{T} \rightrightarrows \mathcal{D}$, where $\mathcal{T} \subseteq \mathcal{V}$ and $\mathcal{D} \subseteq \mathcal{V}^* \doteq \text{Lin}(\mathcal{V} \rightarrow \overline{\mathbb{R}})$, is directly elicitable if there exists $G : \text{Conv}(\mathcal{T}) \rightarrow \overline{\mathbb{R}}$ convex with $G(\mathcal{T}) \subseteq \mathbb{R}$ such that $\Gamma(t) \subseteq \partial G_t$. In this case we say G directly elicits, or just elicits, Γ .*

In other words, G elicits $\Gamma : \mathcal{T} \rightrightarrows \mathcal{D}$ if the φ in Theorem 13 and Corollary 17 is the identity. Of course, it remains to be shown that there exists an affine score eliciting such a property, but the proof is trivial.

Proposition 1. *Directly elicitable properties are elicitable.*

Proof. Let $\Gamma : \mathcal{T} \rightrightarrows \mathcal{R}$ and $G : \text{Conv}(\mathcal{T}) \rightarrow \overline{\mathbb{R}}$ convex with $G(\mathcal{T}) \subseteq \mathbb{R}$ be given such that $\Gamma(t) \subseteq \partial G_t$. Then taking $\mathcal{D} = \mathcal{R}$ and $\varphi = \text{id}_{\mathcal{D}}$, we have by Theorem 17 that Γ is elicitable.

Note that this direct elicibility in no way necessary for elicibility, since the report space is not required to have any intrinsic meaning. For example, one can take $\Gamma(t) \doteq -\partial G_t$ for some G , which in general will not be *directly* elicitable, but still elicitable with $\varphi(r) = -r$ and G .

The notion of direct elicitation is often useful for generating intuitive examples, since the report space itself has meaning. In fact, given any convex function G , the property $\Gamma(t) = \partial G_t$ is directly elicitable by G . This is in fact how equation (34) was generated, though at $t = 0$ we selected $\{0\}$ instead of the full subgradient set $\partial G_0 = [-1, 1]$ to make Γ non-redundant.

We can also clarify what we mean when we say direct elicitation is canonical: every elicitable property gives rise to a directly elicitable property.

Proposition 2. *Let Γ be an elicitable property, elicited by $A(r, t) = G(t_r) + \varphi(r)(t - t_r)$. Then $\Gamma^\varphi(t) = \varphi(\Gamma(t))$ is directly elicitable.*

Proof. Simply keep G and take $\text{id}_{\mathcal{D}}$ as the new φ .

In other words, properties are literally just subsets of subderivative mappings, up to some bijection (or *link function*) taking them to some other report space \mathcal{R}' (see the discussion following Theorem 14).

As a final remark, we note a few observations about direct elicitation. One first notices that the G eliciting some Γ is not unique, as $G' \doteq G + c$ will also elicit Γ for any constant c . But these are the *only* convex functions directly eliciting Γ . Moreover, recovering such a G from Γ is easy: simply integrate (a selection of) Γ to obtain G . Testing whether Γ is directly elicitable is less straight-forward, but there are a variety of monotonicity conditions addressing this issue as well (cf Appendix C).

G.2 Report duality

We are now ready to hold up a mirror to properties and their scores, by introducing notions of duality. As we will see, there are actually *two* mirrors, yielding four combinations of dualities (see Table 1). In this subsection we will explore the first, flipping the report from the type to the dual type. For now, we will take our dual vector space to be all linear functions from \mathcal{V} to \mathbb{R} (*not* $\overline{\mathbb{R}}$ as above).¹³ We begin with the fundamental object of convex duality, the convex conjugate.

¹³ When the dual space can take on infinite values, the conjugate is not always well-defined, as values of the form $\infty - \infty$ are encountered.

Definition 17. Let $\mathcal{V}^* \doteq \text{Lin}(\mathcal{V} \rightarrow \mathbb{R})$. The convex conjugate of $G : \mathcal{V} \rightarrow \overline{\mathbb{R}}$, denoted $G^* : \mathcal{V}^* \rightarrow \overline{\mathbb{R}}$, is given by

$$G^*(d) = \sup_{v \in \mathcal{V}} d(v) - G(v). \quad (35)$$

The power of the conjugate is apparent after the following lemma, which says roughly that the convex conjugate “encodes” the subgradients of G . This is a classic result in convex analysis (cf. [63, Thm E.1.4.1]) which we prove for completeness.

Lemma 3. Let $G : \mathcal{V} \rightarrow \overline{\mathbb{R}}$ be convex. Then for all $v \in \mathcal{V}, d \in \mathcal{V}^*$,

$$G^*(d) = d(v) - G(v) \iff d \in \partial G_v.$$

Proof. We can simply break down the conditions step by step:

$$\begin{aligned} G^*(d) = d(v) - G(v) &\iff v \in \text{argsup}_{v' \in \mathcal{V}} d(v') - G(v') \\ &\iff \forall v' \in \mathcal{V}, d(v) - G(v) \geq d(v') - G(v') \\ &\iff \forall v' \in \mathcal{V}, G(v') \geq G(v) + d(v' - v), \end{aligned}$$

where in the last step we merely negated and added $d(v') \in \mathbb{R}$ to both sides.

Lemma 3 lets us further simplify Theorem 13, as follows. Note however that we are making an additional assumption, that $G > -\infty$.

Theorem 14. Let non-redundant property $\Gamma : \mathcal{T} \rightrightarrows \mathcal{R}$ and Γ -regular affine score $A : \mathcal{R} \times \mathcal{T} \rightarrow \overline{\mathbb{R}}$ be given with score set $\mathcal{A} \subseteq \text{Aff}(\mathcal{T} \rightarrow \mathbb{R})$. Then A elicits Γ if and only if there exists some convex $G : \text{Conv}(\mathcal{T}) \rightarrow \mathbb{R}$, and bijective $\varphi : \mathcal{R} \rightarrow \mathcal{D}$ with $\mathcal{D} \subseteq \partial G$ satisfying $\varphi(\Gamma(t)) \subseteq \partial G_t$, such that for all $r \in \mathcal{R}$ and $t \in \mathcal{T}$,

$$A(r, t) = \varphi(r)(t) - G^*(\varphi(r)). \quad (36)$$

Theorem 14 has natural interpretations for both mechanisms and scoring rules. For mechanisms, it captures a version of the *taxation principle*, that a mechanism can be viewed as a menu of possible allocations and payment associated with each allocation. For scoring rules, it captures the relationship between a scoring rule and a prediction market. We discuss these ideas briefly following Table 1, and in more detail in Appendix G.

An immediate consequence of Theorem 14 pertains to optimization qualities of loss functions in machine learning. Given a loss function $L(r, \omega)$, a crucial subroutine in many machine learning applications involves minimizing $L(r, p) \doteq \mathbb{E}_{\omega \sim p} L(r, \omega)$ with respect to r , where p is either known or approximated from data (a practice known as *empirical risk minimization*). To enable the swift computation of this minimum, a typical practice is to choose a *link function* $\psi : \mathcal{R} \rightarrow \mathcal{R}'$ such that the optimization $\min_{r'} L(\psi^{-1}(r'), p)$ is convex. But for which properties Γ and which loss functions L can we choose ψ such that $L(\psi^{-1}(r'), p)$ is convex in r' ? Provided we are in a paired space (see Definition 18 below),

Theorem 14 reveals that the answer is *always*: take $\psi = \varphi$ in Theorem 14. Then $L(r', p) \doteq -A(\varphi^{-1}(r'), p) = G^*(r') - \langle r', p \rangle$, which is convex in r' . There are two important caveats of this result, however: (1) this is a *constrained* optimization, as the minimization is actually over $r' \in \varphi(\mathcal{R})$, and (2) we have mapped a potentially low-dimensional representation r to a necessarily high-dimensional object $\varphi(r)$. As it turns out, (1) is alleviated by noticing that we may equally optimize over $r' \in \partial G$, which is a convex set. Addressing (2), however, requires finding a lower-dimensional representation of \mathcal{R}' to be practical.

G.3 Type duality and the duality quadrangle

Beyond dual report spaces, we now define dual *properties* and their scores, where we swap the roles of types and reports. This is the second “mirror,” and with both in hand now we have a full four combinations of dual report and type, which we call the duality quadrangle; see Table 1.

To start, we need a dual vector space with more structure than simply $\text{Lin}(\mathcal{V} \rightarrow \mathbb{R})$. For this we use the notion of a *dual pair*, which is a standard setting for convex analysis in infinite-dimensional spaces.

Definition 18 ([3, §5.14]). *A pair of topological vector spaces $(\mathcal{V}, \mathcal{V}^*)$ is a dual pair if it is equipped with a bilinear form $\langle \cdot, \cdot \rangle : \mathcal{V} \times \mathcal{V}^* \rightarrow \mathbb{R}$ which separates points, in the sense that $\forall v^* \langle v, \cdot \rangle \equiv 0$ implies $v = 0$ and $\forall v \langle \cdot, v^* \rangle \equiv 0$ implies $v^* = 0$.*

Note that the above assumption that $(\mathcal{V}, \mathcal{V}^*)$ is a dual pair implies in particular that for all $v^* \in \mathcal{V}^*$, the map $v \mapsto \langle v, v^* \rangle$ is linear. This isn’t crucial when interpreting $\mathcal{R} \subseteq \mathcal{V}^*$ as the type space, since affine scores must be affine in the type. For the remainder of this section we assume that we have a dual pair $(\mathcal{V}, \mathcal{V}^*)$.

A natural question is to determine the conditions under which we have $G^{**} \doteq (G^*)^* = G$. That is, when is the conjugacy operator an involution? This has been thoroughly studied in convex analysis. We state the classic theorem due to Fenchel and Moreau [38, 45].

Definition 19. *A function $f : X \rightarrow \overline{\mathbb{R}}$ is lower semi-continuous (l.s.c.) if for every x_0 in $\text{dom}(f)$ it holds that $\liminf_{x \rightarrow x_0} f(x) \geq f(x_0)$.*

Theorem 15 (Fenchel–Moreau). *Let X be a Hausdorff locally convex space. For any function $G : X \rightarrow \overline{\mathbb{R}}$, it follows that $G = G^{**}$ if and only if one of the following is true: (1) G is a proper, l.s.c., and convex function, (2) $G \equiv +\infty$, or (3) $G \equiv -\infty$.*

The following corollary will prove very helpful in our discussion of type duality below. The proof follows from applying Theorem 15 (note that as \mathbb{R} is Hausdorff, \mathcal{V} together with the product topology inherited from the dual pair is also Hausdorff and locally convex; see [3, §7] for details), and then Lemma 3 twice, once for G and once for G^* .

Corollary 18. *If G is convex, proper, and l.s.c., then $v^* \in \partial G_v \iff v \in \partial G_{v^*}^*$.*

We now introduce the concept of a *dual property* Γ^* , which essentially swaps the type and the report. That is, an agent has a “true report” r and $\Gamma^*(r)$ encodes all the “correct types” t . We then go on to show the relationship between the direct elicibility of dual properties.

Definition 20. *Let $\Gamma : \mathcal{T} \rightrightarrows \mathcal{R}$ where $\mathcal{R} \subseteq \mathcal{V}^*$. Then the dual of Γ , written $\Gamma^* : \mathcal{R} \rightrightarrows \mathcal{T}$, is defined by $\Gamma^* \doteq \Gamma^{-1}$. In other words, Γ^* satisfies $r \in \Gamma(t) \iff t \in \Gamma^*(r)$.*

Theorem 16. *For dual pair $(\mathcal{V}, \mathcal{V}^*)$, let $\Gamma : \mathcal{T} \rightrightarrows \mathcal{D}$ be given with $\mathcal{T} \subseteq \mathcal{V}$ and $\mathcal{D} \subseteq \mathcal{V}^*$. Let convex proper and l.s.c. G be given. Then G elicits Γ if and only if G^* elicits Γ^* .*

Proof. We apply Corollary 18 to obtain $d \in \partial G_t \iff t \in \partial G_d^*$. If G directly elicits Γ , then we have

$$t \in \Gamma^*(d) \iff d \in \Gamma(t) \iff d \in \partial G_t \iff t \in \partial G_d^*,$$

so G^* directly elicits Γ^* . Clearly the above may be applied in the reverse direction as well, yielding the result.

Note that when G and G^* elicit Γ and Γ^* , respectively, we have by the above discussion that $A(d, t) = \langle t, d \rangle - G^*(d)$ elicits Γ and $A^*(t, d) = \langle t, d \rangle - G(t)$ elicits Γ^* . Moreover, the “consumer surplus” functions of A and A^* are G and G^* , respectively. This curious relationship, combined with the notion of report duality, can be visualized as shown in Table 2. Note that traveling around the table does not necessarily mean arriving at the same choice of G , nor does it imply that $G^{**} = G$. However, when $G^{**} = G$ does hold, the diagram “commutes” in a certain sense.

As an example, in Table 1 we used the dual pair given by distributions and functionals to explore duality between scoring rules, prediction markets, and mechanisms. The prediction market model we need is a condensed version of the standard cost-function framework [2]. Briefly, a centralized market maker chooses a convex cost function C , and traders who wish to buy a bundle of securities $q \in \mathbb{R}^{\mathcal{O}}$ (where the trader will receive $\$q_o$ upon outcome o) pays $C(q^0 + q) - C(q^0)$, where q^0 is the vector of total purchases made so far in the market. Abstracting away q^0 , we set $G^*(q) = C(q^0 + q) - C(q^0)$, yielding affine score $A(q', p) = \langle p, q' \rangle - G^*(q')$.

To illustrate the power of this duality, observe that for any p, q , a prediction market A with cost function G^* and menu auction A^* with price function G satisfy $A(q, p) - A^*(p, q) = G(p) - G^*(q)$. This means that difference between the expected payoff under p for purchasing q from the prediction market, and the the expected utility according to q for selecting menu item p , is equal to the difference between the corresponding consumer surpluses.

		Type	
		Primal	Dual
Report	Primal	$A(t', t)$ $=$ $G(t') + \langle t - t', dG_{t'} \rangle$	$A^*(t', d)$ $=$ $\langle t', d \rangle - G(t')$
	Dual	$A(d', t)$ $=$ $\langle t, d' \rangle - G^*(d')$	$A^*(d', d)$ $=$ $G^*(d) + \langle dG_{d'}^*, d - d' \rangle$

$\sup A(\cdot, t) = G(t)$

$\sup A^*(\cdot, d) = G^*(d)$

Table 2. The duality quadrangle.

G.4 Further identities and remarks

Dual-report mechanisms and the taxation principle. The notion of a dual-report mechanism is already well-known as a consequence of the *taxation principle* — instead of asking the agent for her type, one could simply ask the agent directly for the desired allocation, posting a menu prices (or “taxes”) for each. This is without loss of generality because a mechanism’s prices cannot depend on the agent’s type except through the chosen allocation. In our notation, each allocation d is listed with its price $G^*(d)$. It is worth noting however that this is not always identical to the original mechanism. Specifically, while the equilibrium payoffs for the posted-price mechanism $A(d, t)$ are the same as those of the direct revelation mechanism $A(t', t)$, the *off-equilibrium* payoffs need not be equivalent, as the posted-price mechanism may allow reports $d \in \partial G_t$ which are not $dG_{t'}$ for any t' . In other words, because the primal-report (i.e., direct) mechanism must choose a single subgradient dG_t for every point, if $\{dG_t\}_{\mathcal{T}} \subsetneq \partial G = \mathcal{D}$, the dual-report mechanism can be strictly more expressive.

Dual-report scoring rules and prediction markets. The notion of report duality exactly captures the relationship between scoring rules and prediction markets. Here the scoring rules have the primal report space, and prediction markets the dual, where the optimal share bundle is essentially a subgradient of the scoring rule at the trader’s belief. There we will further discuss conditions for which the duality can be run in reverse without loss of generality, but as mentioned above about mechanisms, in general the “menu” format (dual report) of an affine score can be strictly more expressive than the type format (primal report).

Identities Table 2 shows that the theory of elicitation inherits a lot of structure from convex duality. Ignoring boundary and regularity concerns for the moment, we obtain some nice identities:

$$A(d, t) + A^*(t, d) \geq \langle t, d \rangle \quad (37)$$

$$A(d, t) - A^*(t, d) = G(t) - G^*(d). \quad (38)$$

The first follows from the classic Fenchel-Young inequality [59], the proof of which for G proper follows directly from the definition of the conjugate (Definition 17).

Lemma 4 (Fenchel-Young inequality). $\forall v \in \mathcal{V}, v^* \in \mathcal{V}^*, G(v) + G^*(v^*) \geq \langle v, v^* \rangle$.

The elicitation game Define a two-player game $M(d, t)$, with row strategies $d \in \mathcal{D}$ and column strategies $t \in \mathcal{T}$, as

$$M(d, t) = \left(A(d, t), A^*(t, d) \right) = \left(\langle t, d \rangle - G^*(d), \langle t, d \rangle - G(t) \right). \quad (39)$$

One could think of the column player as choosing the agent’s type, and the row player as choosing the principal’s “allocation.” Interestingly, this interpretation implies that the row is the agent and the column is the principal (they each choose each other’s “type”). Immediately one realizes that the Nash equilibria of this elicitation game M are exactly the set of dual-optimal points (d, t) such that $d \in \partial G_t$ and $t \in \partial G_d^*$. Moreover, the equilibrium payoffs for the Nash (d, t) are $(G(t), G^*(d))$.

It is interesting to note the mixed strategies of this game: if $d \sim P_{\mathcal{D}}$ and $t \sim P_{\mathcal{T}}$ independently, the payoffs are

$$M(P_{\mathcal{D}}, P_{\mathcal{T}}) = \left(\langle \bar{t}, \bar{d} \rangle - \mathbb{E}_{P_{\mathcal{D}}}[G^*(d)], \langle \bar{t}, \bar{d} \rangle - \mathbb{E}_{P_{\mathcal{T}}}[G(t)] \right), \quad (40)$$

and if $(d, t) \sim P$ is supported only on dual points,

$$\mathbb{E}_P[M(d, t)] = \left(\mathbb{E}_{P|_{\mathcal{T}}}[G(t)], \mathbb{E}_{P|_{\mathcal{D}}}[G^*(d)] \right), \quad (41)$$

both of which bear resemblance to quantities in Bayesian or randomized mechanism settings.

Score divergences The score divergence $A(t, t) - A(t', t)$ is a natural notion of “regret” which arises frequently in the scoring rules literature (cf. [33]). Our score divergence, as we define below, is reminiscent of a Bregman divergence.

$$D_{G, dG}(t, t') \doteq A(t, t) - A(t', t) = G(t) - G(t') - \langle t - t', dG_{t'} \rangle. \quad (42)$$

Note that the first argument to D is the true type, as opposed to our A notation. Also note the subscripts to D , which specify both the convex function G and a

selection of subgradients. A Bregman divergence requires G to be continuously differentiable, but our definition (42) is a natural extension, and has been studied before (cf. [39]).

Score divergences have many nice properties, like convexity in the first argument, and (directional) differentiability at $t' = t$. Score divergences also enable reasoning about the magnitude of off-equilibrium payoffs, which can be important in practice, when externalities are often present. For example, Fiat et al. [29] introduce the notion of “strong truthfulness”, where the payoff decays as $\|t - t'\|^2$, to design mechanisms that are robust even when agents care about the utility of other agents.

Turning to our various notions of duality, the following are four divergences corresponding to the duality quadrangle, starting in the (primal, primal) setting and moving counter-clockwise.

$$D_{G, dG}(t, t') = G(t) - G(t') - \langle t - t', dG_{t'} \rangle \quad (43)$$

$$D_G(t, d') = G(t) + G^*(d') - \langle t, d' \rangle \quad (44)$$

$$D_{G^*, dG^*}(d, d') = G^*(d) - G^*(d') - \langle dG_{d'}^*, d - d' \rangle \quad (45)$$

$$D_{G^*}(d, t') = G^*(d) + G(t') - \langle t', d \rangle. \quad (46)$$

Amazingly, we see that $D_G(t, d) = D_{G^*}(d, t)$ for all t, d (not just dual points). In other words, the loss of reporting d in the primal but having type t is the same as reporting t in the dual but having “type” d . In the context of the elicitation game above, this means that at any pure strategy pair, both players have the same regret, so they both stand to gain the same amount in a best response (though a simultaneous best response will *not* lead to an equilibrium point in general).

H Bregman Voronoi digrams and the role of $\|\cdot\|^2$

The squared norm seems fundamental to our derivation; let us dig further to see if this is indeed the case. Observe that the form (9) is simply

$$A(r, t) = 2 \langle t_r, t \rangle - \|t_r\|^2 + w_r,$$

where $t_r = \alpha p_r + p_0$. Consider the case where $w_r = 0$ for all r , which corresponds to Γ being a *Voronoi diagram*. In this case, could think of A as being a special case of the “Brier score” $A^B(t', t) = 2 \langle t', t \rangle - \|t'\|^2$, so that $A(r, t) = A^B(t_r, t)$. In other words, we can think of our finite-report case as just restricting the allowed reports in a general direct-revelation affine score. Note that the score divergence for A^B is just $D_G(t', t) = \|t' - t\|^2$, where $G(t) = \|t\|^2$ is just the square norm. This raises the following interesting question: what do we get when we replace $G = \|\cdot\|^2$ with another convex function on \mathcal{T} , and restrict the reports from \mathcal{T} to just a few points $\{t_r\}_{\mathcal{R}}$? That is, take $A^G(t', t) = G(t') - dG_{t'}(t - t')$ and set $A(r, t) = A^G(t_r, t)$. Surely, for any such G , whatever Γ is elicited by such a modified affine score would have to be a diagram by Theorem 3. But then why does the squared norm seem so fundamental?

As it happens, we are touching on precisely the notion of a *Bregman Voronoi diagram*, introduced by Boissonnat et al. [15, §4]. There, instead of defining $\text{cell}_i = \{t : i \in \operatorname{argmin}_j \|t_j - t\|\}$, the squared norm is replaced by any Bregman divergence D_G , so that $\text{cell}_i = \{t : i \in \operatorname{argmin}_j D_G(t, t_j)\}$.¹⁴ Our conclusion that such diagrams coincide with power diagrams corresponds to their Theorem 8.

Framed in terms of our report duality from §G.2, we can see this yet another way. We can rewrite the Bregman Voronoi cell as

$$\text{cell}_i = \left\{ t : i \in \operatorname{argmax}_j G(t_j) - dG_{t_j}(t - t_j) \right\}. \quad (47)$$

By Lemma 3, this can in turn be written

$$\text{cell}_i = \left\{ t : i \in \operatorname{argmax}_j \langle \tilde{t}_j, t \rangle - G^*(\tilde{t}_j) \right\}, \quad (48)$$

where $\tilde{t}_j = dG_{t_j}$. Hence, for any convex function G , the sites $\{p_j\}$ and weights w of a power diagram corresponding to the D_G Bregman Voronoi diagram with sites $\{t_j\}$ are given by $p_j = \frac{1}{2}dG_{t_j}$ and $w_j = \frac{1}{4}\|dG_{t_j}\|^2 - G^*(dG_{t_j})$.

I Discussion

We have presented a model of truthful elicitation which generalizes and extends both mechanisms and scoring rules. On the mechanism design side, we have seen how our framework provides simpler, more general, or more constructive proofs of a number of known results about implementability and revenue equivalence, some of which lead to new results about scoring rules. On the scoring rules side, we have provided the first characterization for scoring rules for non-convex sets of probability distributions. We have also extended our model to eliciting a property of the agent's private information. This has been studied for specific cases in the scoring rules literature, but we have provided the first general characterization. We also show how results about power diagrams in the scoring rules literature lead to a new proof of the Saks-Yu result in mechanism design.

Our analysis makes use of the fact that $A(t', t)$ is affine in t to ensure that $G(t) = \sup_{t'} A(t', t)$ is a convex function. However, this property continues to hold if $A(t', t)$ is instead a convex function of t . Thus, a natural direction for future work is to investigate characterizations of convex scores. While mechanisms can always be represented as affine functions by taking the types to be functions from allocations to \mathbb{R} , it may be more natural to treat the type as a parameter of a (convex) utility function. While many such utility functions are affine (e.g. dot-product valuations), others such as Cobb-Douglas functions are not. Berger, Müller, and Naeemi [12, 13] have investigated such functions and given characterizations that suggest a more general result is possible. Another potential application is scoring rules for alternate representations of uncertainty, several of which result in a decision maker optimizing a convex function [36].

¹⁴ In [15], three types of diagrams are introduced; here we refer to the first type.

In one sense getting such a characterization is straightforward. In the affine case we want $A(t', t)$ to be an affine function such that $A(t', t) \leq G(t)$ and $A(t', t') = G(t')$. Since we have fixed its value at a point, the only freedom we have is in the linear part of the function, and being such a linear function is exactly the definition of a subgradient. So while our characterization of affine scores is in some sense vacuous, it is also powerful in that it allows us to make use of the tools of convex analysis. A similarly vacuous characterization is possible for the convex case: $A(t', t)$ is a convex function such that $A(t', t) \leq G(t)$ and $A(t', t') = G(t')$. The challenge is to find a way to state it that is useful and naturally handles constraints such as those imposed by the form of a utility function.

Many questions in the literature on properties remain open. Most notable is the characterization of elicitable nonlinear and multidimensional properties — the single dimensional case is covered in [47] and the linear vector-valued case in [1]. We hope that the results and intuition from Section 3 will yield a useful characterization in this case. Another interesting direction is for non-functional properties: aside from the finite R case, all work in the literature to our knowledge assumes that Γ is a function (having a single correct report for each type). The generality of Theorem 13 may prove useful in exploring non-functional settings as well. A result requiring few regularity conditions on Γ would be useful in domains such as statistics where natural properties like the median cannot in general be expressed as functions.

Theorem 3 shows that scoring rules for finite properties are essentially equivalent to the weights and points that induce a power diagram. As power diagrams are known to be connected to the spines of amoebas in algebraic geometry, aspects of toric geometry used by string theorists, and tropical hypersurfaces in tropical geometry [64], there may be useful characterization results in those fields as well. The last is particularly suggestive given the recent use of tropical geometry techniques in mechanism design [11].

While our examples have focused on mechanism design and scoring rules, another interesting direction to pursue is other settings where our results may be applicable. One natural domain is the literature on M-estimators in machine learning, statistics and economics. Essentially, this literature takes a loss function (i.e. a scoring rule) and asks what it elicits. For example, the mean is an M-estimator induced by the squared error loss function. Some work in this literature (e.g. [52]) requires that the loss function satisfy certain conditions, and our results may be useful in characterizing and supplying such loss functions.