

An Axiomatic Study of Scoring Rule Markets*

Rafael Frongillo¹ and Bo Waggoner²

1 Colorado University, Boulder CO, USA

raf@colorado.edu

2 University of Pennsylvania, Philadelphia PA, USA

bwag@seas.upenn.edu

Abstract

Prediction markets are well-studied in the case where predictions are probabilities or expectations of future random variables. In 2008, Lambert, et al. proposed a generalization, which we call “scoring rule markets” (SRMs), in which traders predict the value of arbitrary statistics of the random variables, provided these statistics can be elicited by a scoring rule. Surprisingly, despite active recent work on prediction markets, there has not yet been any investigation into more general SRMs. To initiate such a study, we ask the following question: in what sense are SRMs “markets”? We classify SRMs according to several axioms that capture potentially desirable qualities of a market, such as the ability to freely exchange goods (contracts) for money. Not all SRMs satisfy our axioms: once a contract is purchased in any market for prediction the median of some variable, there will not necessarily be any way to sell that contract back, even in a very weak sense. Our main result is a characterization showing that slight generalizations of cost-function-based markets are the only markets to satisfy all of our axioms for finite-outcome random variables. Nonetheless, we find that several SRMs satisfy weaker versions of our axioms, including a novel share-based market mechanism for ratios of expected values.

1998 ACM Subject Classification F.2 Algorithmic Game Theory

Keywords and phrases prediction markets, information elicitation, scoring rules

Digital Object Identifier 10.4230/LIPIcs.ITCS.2018.15

1 Introduction

The goal of a prediction market is to collect and aggregate predictions about some future outcome Y taking values in \mathcal{Y} ; common examples arise from sporting, political, meteorological, or financial events. Prediction markets work by offering financial contracts whose payoffs are contingent on the eventually-observed value of Y . The agent’s choices are interpreted, by revealed preference, as predictions about Y , and thus the final state of the market is interpreted as an aggregation of agent beliefs.

Hanson [13] observed that one can design a prediction market using a *proper scoring rule*, which is a contract $S(p', y)$ that scores the accuracy of prediction p' (a probability distribution over outcomes) upon outcome $Y = y$. The relevant guarantee is that the expected score $\mathbb{E}_{Y \sim p} S(p', Y)$ is maximized when the agent reports their true belief $p' = p$. In Hanson’s *scoring rule market (SRM)*,¹ traders arrive sequentially and report their belief p_t , and are eventually paid $S(p_t, y) - S(p_{t-1}, y)$ when $Y = y$ is revealed.

* A full version of the paper is available at <https://arxiv.org/abs/1709.10065>

¹ We will use this term in lieu of the standard *market scoring rule (MSR)*, as the latter could refer to either the scoring rule or the market mechanism.



© Rafael Frongillo and Bo Waggoner;

licensed under Creative Commons License CC-BY

9th Innovations in Theoretical Computer Science Conference (ITCS 2018).

Editor: Anna R. Karlin; Article No. 15; pp. 15:1–15:20

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In many cases, a market designer may only be interested in specific statistics, called *properties*, of trader beliefs, rather than the entire distribution over the outcome Y . Lambert et al. [14] observed that SRMs can be generalized to arbitrary scoring rules $S(r, y)$, wherein traders are paid $S(r_t, y) - S(r_{t-1}, y)$; if the expected value of S is maximized by reporting the value of a particular property, in which case we say S *elicits* the property, then the market should intuitively aggregate trader beliefs about the property in question. For example, one could design a “median market” by leveraging the fact that $S(r, y) = -|r - y|$ elicits the median of the distribution of Y , and we might expect the corresponding SRM to aggregate predictions about the median of Y .

Surprisingly, apart from exploring which properties are elicitable, practically nothing is known about these general SRMs beyond those for expected values [13, 8, 2]. Here SRMs can be rephrased into a dual “cost-function-based” formulation where traders buy and sell *shares* in some securities $\phi_1, \dots, \phi_k \in \mathbb{R}^{\mathcal{Y}}$, and the market prices reveal the trader’s belief about the expected value of $\phi : \mathcal{Y} \rightarrow \mathbb{R}^k$ (§ 2.2, 5). The literature on prediction markets focuses on this setting, in which one can more easily study traditional market quantities like liquidity and depth.

In this paper, we step beyond expected values and study SRMs as a whole. Our first contribution is identifying which questions to ask of general SRMs. In particular, we ask the following.

- (1) In what sense can we think of SRMs as “markets” in the traditional sense?
- (2) Specifically, which SRMs behave like share-based markets, in the sense that traders can always “sell” their shares for some nontrivial price?
- (3) Which SRMs if any maintain the known desirable characteristics of the cost-function-based framework, such as bounded worst-case loss and no arbitrage?

To answer these questions, we introduce axioms, and then study SRMs based on the property they elicit—as we will see, for example, markets eliciting modes will satisfy different axioms than those for medians, regardless of the scoring rules chosen. Our primary focus is the *trade neutralization (TN)* axiom, which captures question (2) above: traders holding a contract from the market maker can “sell” it for a nontrivial price, receiving more from transaction than the worst-case payoff of the contract. (Otherwise, traders would just keep their holdings regardless of their belief.)

Summary and results. After reviewing prediction markets and scoring rules (§ 2), we introduce our axioms (§ 3). We then classify modes, finite properties, medians, quantiles, and expectations, according to what axioms their markets can satisfy (§ 4).

In § 5, we give our main result: *any SRM satisfying trade neutralization must be a certain form of generalized cost-function-based market* (Theorem 16). We note that unlike in prior work on cost functions [2], this result applies to any SRM for any property; so in a sense, it says that only close relatives of expectations can be elicited by any market satisfying trade neutralization.

Our main result raises the question of whether a meaningful relaxation of trade neutralization is satisfied by a broader class of markets. We show that this is indeed the case: markets for ratios of expectations (§ 6) and for expectiles (full version of the paper) both satisfy “weak neutralization” (Axiom 5) but not trade neutralization. While these results suggest that a characterization of weak neutralization is an exciting future direction, they also yield new prediction market mechanisms worthy of attention on their own. For example, as we describe in § 6, our new market framework for ratios of expectations exchanges a security (random variable) not for cash but for units of another security, thereby revealing

trader beliefs about the *ratio* of their expectations.² Finally, we conclude in § 7 with open questions, and a discussion of other properties, elicitation complexity, and alternative market formulations. The Appendix contains all omitted proofs.

Relation to prior work. While we discuss related work in § 2, it is important to distinguish our main result from prior characterizations of cost-function-based markets in the literature. In particular, as mentioned above, cost-function-based markets are known to be equivalent to scoring rules that elicit expected values [13, 8, 2]. In contrast, we ask which among a very broad family of mechanisms satisfy certain basic “market” axioms, and find that such mechanisms must fall into (a minor generalization of) the cost-function-based framework. Indeed, based on the above known equivalence, we further conclude that our axioms imply the elicitation of expected values, though weakening them slightly allows us to elicit ratios of expectations (§ 6) and expectiles (full version) while still retaining some sense of a “market”.

2 Aggregating Information with Prediction Markets

The goal of a prediction market is to crowdsource and aggregate beliefs of participants about some future event. It does so by allowing participants to buy and sell contracts which have different payoffs depending on the outcome of the event, and inferring a prediction from the participants’ choices. In this section, we formally define the class of such markets that we study, *scoring rule markets (SRMs)*, with references to previous work. For other related work, see above and § 7.

2.1 Outcomes and contracts

Let \mathcal{Y} denote the *outcome space* of interest to the market designer and Y a future event or random variable taking values in \mathcal{Y} . The designer will in general allow traders to select *contracts* $d \in \mathbb{R}^{\mathcal{Y}}$ from a list offered by the market, which can be thought of as vectors when \mathcal{Y} is a finite set. The interpretation of a contract d is that, when Y is eventually revealed, the market maker will pay the agent a net payoff $d(Y)$, which may be negative. Its expected payoff under distribution p on Y is written $\mathbb{E}_p d(Y)$, which could be thought of as an inner product between p and d .

Note that under this formulation, any initial payment the agent might make is folded in to d . Specifically, letting $\mathbf{1} \in \mathbb{R}^{\mathcal{Y}}$ denote the “all-ones” contract $\mathbf{1}(y) = 1 \forall y \in \mathcal{Y}$, then a contract d could be written $d = d' - \alpha \mathbf{1}$, interpreted as paying a price of α now for the contract d' , whose payoff will be revealed when Y is observed. We assume agents are indifferent to timing of payments and just wish to maximize total expected payoff. We may occasionally abuse terminology by referring both to contracts d which “fold in” initial payments and to contracts d' having an additional “price” α .

2.2 Automated market makers

When designing prediction markets, rather than a typical continuous double auction (“stock exchange”) mechanism, it is common to employ a centralized *automated market maker*, which

² To see this, note that a trader willing to “buy” one unit of $d \in \mathbb{R}^{\mathcal{Y}}$ in exchange for c units of $b \in \mathbb{R}^{\mathcal{Y}}$ is effectively expressing the belief $\mathbb{E} d > c \mathbb{E} b$, i.e., $\mathbb{E} d / \mathbb{E} b > c$. Similarly, selling at this “price” reveals the belief $\mathbb{E} d / \mathbb{E} b < c$.

offers to buy or sell any available contracts, and through which all trades are executed. See [2] for practical reasons behind this design choice.

Formally, a sequence of participants (traders) arrive at times $t = 1, \dots, T$ and each selects a contract from a list offered by the market at that time. It will be convenient to let the set of contracts available be indexed at each time by some *report space* $\mathcal{R} \subseteq \mathbb{R}^k$. Following Abernethy et al. [5], we consider a generic market making algorithm, termed a *mechanism*, that specifies the set of contracts available at each time. In general, this may depend on the entire past history of the market, and is represented as a mapping F that, given a report $r \in \mathcal{R}$ of the participant, returns the corresponding contract. More formally, the contract given to a participant who chooses report r_t given the past history of reports r_1, \dots, r_{t-1} is denoted $\vec{F}(r_t|r_1, \dots, r_{t-1}) \in \mathbb{R}^{\mathcal{Y}}$. The net payoff to the trader upon outcome $y \in \mathcal{Y}$ will be denoted $F(r_t, y|r_1, \dots, r_{t-1})$.

A popular instantiation of such a mechanism is the *cost-function-based market maker*, in which $\mathcal{R} = \mathbb{R}^k$ and $F(r_t, y|r_1, \dots, r_{t-1}) = (r_t - r_{t-1}) \cdot \phi(y) - (C(r_t) - C(r_{t-1}))$, where $C : \mathbb{R}^k \rightarrow \mathbb{R}$ is convex and $\phi : \mathcal{Y} \rightarrow \mathbb{R}^k$ [2]. This payoff function corresponds to a trader making a fixed payment $C(r_t) - C(r_{t-1})$ in return for a bundle $r = r_t - r_{t-1} \in \mathbb{R}^k$ of *shares* in the securities $\phi_1, \dots, \phi_k \in \mathbb{R}^{\mathcal{Y}}$, i.e. r_i units of security ϕ_i , for an up-front cost paid to the market maker in terms of C . Among the several nice properties of this market maker, one can see that a trader who believes $\mathbb{E}\phi = x$ (a vector in \mathbb{R}^k , the component-wise expectation) has an incentive to buy or sell securities until $\nabla C(r_t) = x$, thereby revealing their belief.

2.3 Scoring rule markets for properties

The goal of the market is to incentivize a good prediction for some property or statistic of Y , such as the median or mean. Thus, much work considers prediction markets relying on *proper scoring rules*, which are contracts designed to elicit a single agent's belief (i.e. the case $T = 1$ of a market) [7, 12]. While originally designed to elicit an entire distribution over the outcome Y , in many cases, for example when \mathcal{Y} is very large or even infinite, one may be interested in obtaining only summary information about this distribution. It is therefore natural to consider scoring rules which elicit such statistics, or *properties*, of distributions.

Here and throughout the paper, \mathcal{P} is a set of distributions of interest on the domain \mathcal{Y} , for example, the distributions with full support, with finite expectation, or so on.

► **Definition 1.** A *property* is a function $\Gamma : \mathcal{P} \rightrightarrows \mathcal{R}$, which associates a set of correct report values to each distribution. We require Γ to be *non-redundant*, meaning $\Gamma^{-1}(r) \not\subseteq \Gamma^{-1}(r')$ for any reports $r, r' \in \mathcal{R}$ (i.e. we cannot have $r \in \Gamma(p) \implies r' \in \Gamma(p)$ for all p). A property is *single-valued* if each p maps to exactly one report, in which case we write Γ as a function $\Gamma : \mathcal{P} \rightarrow \mathcal{R}$.

A scoring rule $S(r, y)$ simply provides a payoff based on a reported value of the property and the observed outcome y . We say the scoring rule *elicits* the property if the correct report is incentivized.

► **Definition 2.** A scoring rule $S : \mathcal{R} \times \mathcal{Y} \rightarrow \mathbb{R}$ *elicits* a property $\Gamma : \mathcal{P} \rightrightarrows \mathcal{R}$ if for all $p \in \mathcal{P}$, $\Gamma(p) = \operatorname{argmax}_{r \in \mathcal{R}} \mathbb{E}_p S(r, Y)$ where $Y \sim p$. When Γ is single-valued, this condition becomes $\{\Gamma(p)\} = \operatorname{argmax}[\dots]$. A property is *elicitable* if some scoring rule elicits it.

Some properties are not elicitable, meaning there is no way to score reports of their value based on an observed outcome in a manner which incentivizes truthfulness. A classic example is the variance of a distribution, which does not have convex level sets (mixtures of distributions with the same variance have higher variance in general), a necessary condition

for elicibility [14]. (For such non-elicitable statistics, one could still discuss their elicitation *complexity*, the number of reports needed to compute the desired property post-hoc; we discuss this concept in § 7.) However, several well-known statistics are elicitable properties, including expected values, medians, quantiles, expectiles, and ratios of expectations.

Combining the concept of scoring rules for properties with scoring rule markets yields the following natural prediction market mechanism for an arbitrary property $\Gamma : \mathcal{P} \rightrightarrows \mathbb{R}$ elicited by S [13, 14, 4]. Initialize the market state at some $r_0 \in \mathcal{R}$. When trader $t = 1, \dots, T$ arrives, they can choose to update the market state to any $r_t \in \mathcal{R}$, and once the outcome $y \in \mathcal{Y}$ is revealed, the market maker will pay the trader $S(r_t, y) - S(r_{t-1}, y)$. We can again express this mechanism using our F notation above (we will later relate F to Γ).

► **Definition 3.** A *scoring rule market* for scoring rule $S : \mathcal{R} \times \mathcal{Y} \rightarrow \mathbb{R}$ and initial state $r_0 \in \mathcal{R}$ is the mechanism $F(r_t, y | r_1, \dots, r_{t-1}) = S(r_t, y) - S(r_{t-1}, y)$.

For brevity, we will simply write $F(r', y | r) = S(r', y) - S(r, y)$, or using contract notation $\vec{F}(r' | r) = \vec{S}(r') - \vec{S}(r)$, where of course $\vec{S}(r)_y = S(r, y)$.

3 Axioms and Preliminaries

To motivate our choice of SRMs, we consider two preliminary axioms in the full version of the paper that turn out to characterize SRMs; here, we briefly summarize. *Incentive-compatibility (IC)* states that agents maximize expected payment by choosing contracts that reveal their true belief about $\Gamma(p)$, for the property Γ chosen by the market maker. We note that previous work [2, 5] did not explicitly consider IC, but it arose as a consequence of cost-function markets revealing expected values. In this work, however, IC is a primary axiom as we consider markets designed to elicit particular properties, but also a weak one in that essentially every scoring rule elicits some (potentially set-valued) property. *Path-independence (PI)*, an axiom appearing in prior work on prediction markets [2], ensures that a participant cannot gain more by making multiple trades in a row than by simply making the single optimal trade immediately. In the full version, we show that any mechanism which satisfies both PI and IC is an SRM for a property Γ . It is conceptually similar to results for markets eliciting the mean [2, 5], but more general as it holds for any property.

► **Theorem 4.** A mechanism satisfies PI and IC for property Γ if and only if it is a scoring rule based market (SRM) with some scoring rule S that elicits Γ .

Henceforth, unless otherwise noted we assume PI and IC, which is to say we focus on SRMs.

3.1 Arbitrage and Bounded Loss

We now express two time-honored axioms in our notation: no arbitrage, and bounded worst-case loss. Recall that we write $\mathbf{1} \in \mathbb{R}^{\mathcal{Y}}$ to mean the contract paying out $\mathbf{1}(y) = 1$ for each $y \in \mathcal{Y}$. We also will write $\inf d = \inf_{y \in \mathcal{Y}} d(y)$ and $\sup d = \sup_{y \in \mathcal{Y}} d(y)$ to denote the payout bounds of contract $d \in \mathbb{R}^{\mathcal{Y}}$; note that we use infima and suprema as \mathcal{Y} may be infinite, for example when $\mathcal{Y} = \mathbb{R}$.

The following axiom has been a desiderata since the inception of prediction markets: the market maker should not risk losing an unbounded amount of money.

► **Axiom 1 (Worst-Case Loss (WCL)).** An SRM F , initialized at r_0 , satisfies WCL with bound $B \geq 0$ if for all $r \in \mathcal{R}$, $\sup \vec{F}(r | r_0) \leq B$.

As of course $\vec{F}(r|r_0) = \vec{S}(r) - \vec{S}(r_0)$, WCL simply says that the scoring rule S is bounded relative to the initial score $S(r_0, \cdot)$.

Another classic condition for a market mechanism is that a trader should never have an opportunity to make a guaranteed (risk-free) profit. In our contract notation, this means traders should never be able to purchase a contract d for less than its minimum payoff $\inf d$, or equivalently, the net contract provided to traders should not be unconditionally positive.

► **Axiom 2 (No Arbitrage (ARB)).** SRM F satisfies ARB if $\forall r, r' \in \mathcal{R}, \inf \vec{F}(r'|r) \leq 0$.

Equivalently, F satisfies ARB if for every trade, there is an outcome yielding a profit at most 0. The no-arbitrage condition was crucial in deriving cost-function-based markets in terms of convex conjugate duality [2]. Surprisingly, it turns out that any SRM satisfies no-arbitrage automatically if it is incentive-compatible.

► **Proposition 5.** SRM F satisfies ARB if it is IC for some property $\Gamma \rightrightarrows \mathcal{R}$.

Proof. If $0 < \inf \vec{F}(r'|r) = \inf[\vec{S}(r') - \vec{S}(r)]$, then for all $y \in \mathcal{Y}, S(r', y) > S(r, y)$. Thus, for any $p \in \mathcal{P}, \mathbb{E}_p S(r', Y) > \mathbb{E}_p S(r, Y)$, so r cannot be an optimal report for any p , contradicting non-redundancy (as $\Gamma^{-1}(r) = \emptyset \subseteq \Gamma^{-1}(r')$). ◀

3.2 New Axioms

We now identify several new axioms that capture desirable characteristics of prediction markets. The first is motivated by traditional markets, wherein a trader can always “neutralize” or “liquidate” their holdings. For example, if a trader buys a bar of gold, apart from apocalyptic scenarios, she can always sell it at any time for a strictly positive price. As another example, a trader holding an Arrow-Debreu contract, paying \$1 upon event E and \$0 otherwise, should be able to sell it for some nonzero (perhaps very low) price. More generally, if the contract purchased is $d \in \mathbb{R}^{\mathcal{Y}}$ and is non-constant, the trader should receive strictly more than $\inf d$ in cash (in both examples above, $\inf d = 0$). This condition is the *trade neutralization (TN)* axiom, which we now give along with two variants, one stronger and one weaker.

► **Axiom 3.** An SRM satisfies *Trade Neutralization (TN)* if for all trades $r_1 \rightarrow r'_1$, and all reports r_2 , there is a trade $r_2 \rightarrow r'_2$ such that $\vec{F}(r'_1|r_1) + \vec{F}(r'_2|r_2) = c\mathbf{1}$ for some scalar $c > \inf \vec{F}(r'_1|r_1)$.

► **Axiom 4.** An SRM satisfies *Portfolio Neutralization (PN)* if for all sets of trades $r_i \rightarrow r'_i, 1 \leq i \leq m$, and all reports r , there is a trade $r \rightarrow r'$ such that $\vec{F}(r'|r) + \sum_i \vec{F}(r'_i|r_i) = c\mathbf{1}$ for some scalar $c > \inf \left[\sum_{i=1}^m \vec{F}(r'_i|r_i) \right]$.

► **Axiom 5.** An SRM satisfies *Weak Neutralization (WN)* if for all trades $r_1 \rightarrow r'_1$, and all reports r_2 , there is a trade $r_2 \rightarrow r'_2$ such that $\inf [F(r'_1|r_1) + F(r'_2|r_2)] > \inf F(r'_1|r_1)$.

Portfolio neutralization (PN) is stronger than TN: for any *set* of purchased contracts, a trader should be able to sell the entire portfolio for more than its minimum payoff, all in one go, at any time. Surprisingly, we will see that in fact TN is equivalent to PN (Theorem 16).³ Similarly, weak neutralization (WN) is weaker than TN, as it only requires that a trader

³ Note that this equivalence does not follow from path independence (PI), which only refers to *consecutive* sequences of trades. PN is most interesting for nonconsecutive trades, e.g. the same contract $r \rightarrow r'$ purchased multiple times.

be able to improve her worst-case payoff, but not necessarily render her holdings outcome-independent. More precisely, WN states that traders holding a non-constant contract d should be able to purchase a contract d' such that $\inf[d + d'] > \inf d$. Note that WN is not far off from TN in the sense that a WN market maker allowing traders to exchange any contract for the cash value of its worst-case payout, a favorable option for the market maker, would satisfy TN: a trader can purchase d' as above, and then exchange $d + d'$ for $(\inf(d + d'))\mathbf{1}$, thus neutralizing while improving the worst-case payout. In fact, path independence (PI) ensures that the trader can do both in a single trade.

Note that the inequalities above are all strict. If we replaced them by weak inequalities, then WN would be trivially satisfied by essentially every conceivable mechanism: by just keeping your current contract d , you are guaranteed a payoff of at least $\inf d$, by definition. This makes intuitive sense, as unless the contract d is constant, it strictly dominates $(\inf d)\mathbf{1}$, so no rational trader would ever consent to trading d for $(\inf d)\mathbf{1}$. The strict inequality thus captures a reasonable possibility for traders to “sell back” their previously-purchased contracts, in the sense that there is some belief a trader could hold where such a trade would be beneficial. More importantly, the strict inequalities ensure that the market maker does indeed “make a market”: the defining quality of a market maker is a willingness to buy or sell securities at some price, but prices equal to the best- or worst-case payoff of the securities are equivalent to refusing to buy or sell, as no rational trader regardless of belief would accept such a trade.

Our final axiom says that traders with arbitrarily small budgets who disagree with the current prediction can always make some profitable trade. In other words, as the overall market scales larger relative to the maximum allowable loss ϵ of individual traders, they still have incentives to participate (although naturally, the “size” of their updates will be small relative to the entire market). Thus, the following axiom captures scalability of a market both in terms of large numbers of individually small trader budgets and informational updates.


► **Axiom 6.** A market F satisfies *bounded trader budget (BTB)* if, for all market states r and beliefs p with $\Gamma(p) \neq r$, and for all $\epsilon > 0$, there exists a contract $F(r'|r)$ with $\inf F(r'|r) > -\epsilon$ and $\mathbb{E}_p F(r', Y|r) > 0$.

4 Central Examples

With our axioms in hand, we now turn to specific properties: the mode, median, quantiles, and expectations. For each, we wish to understand which axioms the corresponding SRMs can satisfy. As we will see, each of these properties has a unique signature with respect to our axioms.

4.1 Mode and Finite Property Markets

Perhaps the simplest example of any SRM is the canonical mode market for a distribution on k outcomes $\mathcal{Y} = \{1, \dots, k\}$, with report space $\mathcal{R} = \mathcal{Y}$ and $S(r, y) = \mathbf{1}\{y = r\}$ (“\$1 if you guess correctly”). Here we find that even the weakest version of neutralization, WN, is violated: a trade moving r from 1 to 2 yields payoff $\mathbf{1}\{y = 2\} - \mathbf{1}\{y = 1\}$, meaning this contract will lose the owner 1 if $Y = 1$. But if the current market state is, say, 3, then no report the agent can make will avoid this loss of 1 when $Y = 1$. For any report r' , the final payoff is $\mathbf{1}\{y = 2\} - \mathbf{1}\{y = 1\} - \mathbf{1}\{y = 3\} + \mathbf{1}\{y = r'\}$. See Figure 1.



■ **Figure 1 Visualizing trader “position” in a mode market.** Each curve gives the payoff of a given scoring rule or contract as a function of the outcome y . Here $S(r, y) = 1 \iff r = y$. Left: A trader who has reported $r = 1$ stands to gain 1 if $Y = 1$ and gain 0 otherwise. Center: Similarly for $r = 3$. Right: A trader who moves the market from r to r' gets the function $F(r', \cdot|r) = S(r', \cdot) - S(r, \cdot)$. She stands to gain 1 if $Y = r'$, lose 1 if $Y = r$, and gain 0 otherwise.

Intuitively, the lack of WN has negative implications for agents, who must take on risk that cannot be mitigated later. It also violates BTB: The potential agent loss from any trade is a constant 1, so an agent with a budget smaller than 1 will not be able to participate. This causes significant problems in practice for market designers as well, because the only possible solution, scaling down the scoring rule, is also unpleasant: many agents will be able to participate without much risk, so the market prediction will flip from outcome to outcome without necessarily aggregating information.

We emphasize that these characteristics are inherent to the mode, in the sense that they hold for any other scoring rule eliciting it. In fact, the same conclusions hold for markets eliciting any *finite property* $\Gamma : \mathcal{P} \rightarrow \{1, \dots, k\}$, i.e. a property with a finite set of possible values; see [15] for motivation and examples of such properties.

► **Theorem 6.** *Any market for a finite property satisfies WCL, but not TN or BTB.*


4.2 Mean and expectation markets

Consider now the mean of a random variable over the reals, $\mathcal{R} = \mathcal{Y} = \mathbb{R}$. (In this setting, we naturally take \mathcal{P} to be the set of beliefs with well-defined expectations.) As an illustrative example, consider the scoring rule $S(r, y) = 2ry - r^2$, which elicits the mean. The corresponding SRM takes the form $F(r', y|r) = r^2 - (r')^2 + 2y(r' - r)$. Perhaps the first observation one makes is that, because $y \in \mathbb{R}$, any nontrivial trade leaves the trader exposed to unbounded potential loss, as well as unbounded potential gain. This implies that BTB and WCL cannot hold.

What about trade neutralization? Consider a trader holding the contract $\vec{F}(r'_1, y|r_1) = \alpha_1 + \alpha_2 y$ for constants $\alpha_1 = (r_1)^2 - (r'_1)^2$ and $\alpha_2 = 2(r'_1 - r_1)$. She would like to neutralize this position, so she must purchase some other contract of the form $\alpha_3 - \alpha_2 y$, so that her net position will be the constant $\alpha_1 + \alpha_2$. If the current market state is r_2 , our hero can neutralize her previous trade by choosing $r'_2 = (r_1 - r'_1)$, so that $2y(r'_2 - r_2) = 2y(r'_1 - r_1) = 0$. Her worst-case liability decreases from $-\infty$ to a constant, showing that both WN and TN are satisfied. And in fact, even if she holds a set of contracts, her net position is simply the sum and can still be written in the form $\alpha_1 + \alpha_2 y$ and the same argument goes through; this shows that PN is also satisfied.

This example raises the questions of whether other markets eliciting the mean (if any) would have similar properties, and how this might depend on the random variable in question.

Motivated by this example, we now consider a much more general setting. Let $\phi : \mathcal{Y} \rightarrow \mathbb{R}^k$ be a given “random variable” and let $\Gamma(p) = \mathbb{E}_p[\phi(Y)]$ be its expected value. Such a Γ is called a *linear property*. We call SRM S an *expectation market* if S is IC for such a Γ . For



■ **Figure 2 Trader position in a mean market.** $S(r, y) = -(r - y)^2$. Left: Moving the market from $r = -1$ to $r' = 1$ gives a position $F(r', \cdot|r)$ that pays off linearly in y . Center: At r_t , the trader chooses from a set of possible contracts, depending on if she reports 1.5, 2.5, etc. Right: choosing the contract r'_t that neutralizes the previous trade $r \rightarrow r'$.

ease of exposition (i.e. to avoid relative interiors) we will assume that the range $\mathcal{R} = \Gamma(\mathcal{P})$ is full-dimensional in \mathbb{R}^k .

Capitalizing on the known characterization of scoring rules for expectations [1, 10] (Theorem 18), we know that any scoring rule S eliciting Γ takes the form,

$$S(r, y) = G(r) + dG_r \cdot (\phi(y) - r) + f(y), \tag{1}$$

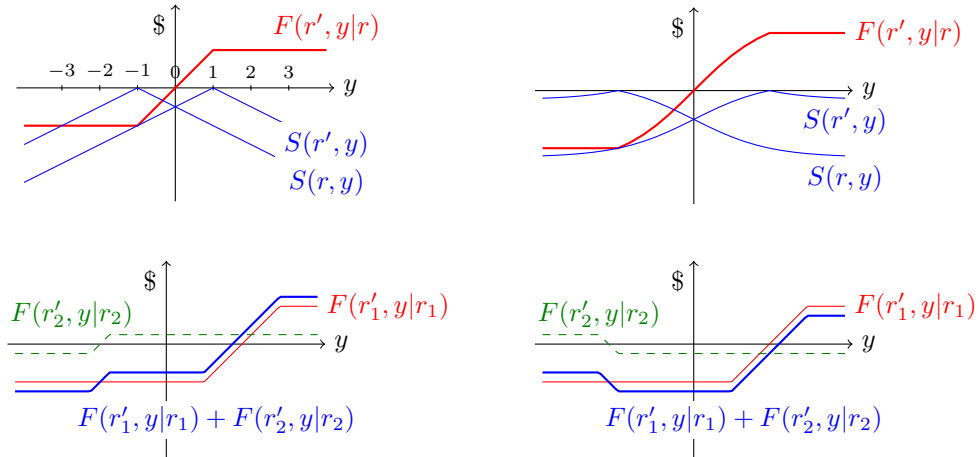
for G strictly convex with subgradients $\{dG_r\}_{r \in \mathcal{R}}$ and f an arbitrary \mathcal{P} -integrable function. Note however that the f term will vanish in the definition of $F(r', y|r)$, so without loss of generality we may take $f(y) = 0$ for all y . We thus refer to expectation markets as being *defined by G* if they satisfy Equation 1 for that G (for any f). Letting $\text{conv}(\phi(\mathcal{Y}))$ be the convex hull of the set of outcomes of the random variable ϕ , we have the following characterization of worst-case loss for expectation markets.

► **Proposition 7.** *The linear property $\Gamma(p) = \mathbb{E}_p \phi(Y)$ has an expectation market satisfying WCL if and only if its domain $\text{conv}(\phi(\mathcal{Y}))$ is bounded, in which case the market defined by any bounded G satisfies WCL.*

► **Proposition 8.** *The linear property $\Gamma(p) = \mathbb{E}_p \phi(Y)$ has an expectation market satisfying BTB if its domain $\text{conv}(\phi(\mathcal{Y}))$ is bounded, in which case the market defined by any differentiable G satisfies BTB.*

To illustrate, let us revisit the example at the beginning of this subsection, i.e. $\phi(y) = y$ and $S(r, x) = 2ry - r^2$. (This corresponds to the convex function $G(r) = r^2$, up to a shift.) We saw that this market satisfies neither BTB nor WCL for $\mathcal{Y} = \mathbb{R}$, the entire real line. However, for $\mathcal{Y} = (0, 1)$, it satisfies both BTB and WCL, as $\text{conv}(\phi(\mathcal{Y})) = (0, 1)$ and G is bounded and differentiable on this domain. Finally, recall that it also (intuitively) satisfied TN, and furthermore, PN. These qualities can be shown to generalize; in particular, prior work has shown that linear properties have nice *cost function* based markets. Such markets take the form $F(r', y | r) = C(r) - C(r') + (r - r') \cdot \phi(y)$, for some convex function C and $\mathcal{R} = \mathbb{R}^k$. It then follows relatively directly that such markets satisfy PN, as any position (set of contracts held by a trader) can be written in the form $d(y) = \alpha_1 + \alpha_2 \cdot \phi(y)$. This can be interpreted as a fixed payment of α and $\alpha_{2,i}$ “shares” in a “security” $\phi(Y)_i$. To neutralize, when the current market state is r' , it turns out to suffice to select the r satisfying $r - r' = -\alpha_2$.

► **Theorem 9.** *On a finite outcome space, i.e. $|\mathcal{Y}| < \infty$, for any linear property there exists an expectation market satisfying PN, WCL, and BTB.*



■ **Figure 3** Trader position in a median market. Here $S(r, y) = -|r - y|$. Top left: The payoffs for $S(r, \cdot)$ and $S(r', \cdot)$, and the contract that moves the market $r \rightarrow r'$, namely $F(r', \cdot|r) = S(r', \cdot) - S(r, \cdot)$. Top right: The same example but with $S(r, y) = -|g(r) - g(y)|$ with g the sigmoid function. Bottom: Two examples where a trader with position “red” $F(r'_1, y|r_1)$ considers a potential contract “green” $F(r'_2, y|r_2)$. If purchased, the net position will be the blue curve. It sometimes falls below the original position, meaning the trader’s worst case has gotten worse.

4.3 Median and Quantile Markets

We previously gave an example of a “mean market” given by $S(r, y) = 2ry - r^2$, the scoring rule analog of squared loss $L(r, y) = (r - y)^2$. Perhaps the most natural statistic to investigate next is the median, elicited by the analog of “absolute loss”, $S(r, y) = -|r - y|$. What we find is surprising: unlike squared loss, the absolute loss market does not satisfy PN, and in fact does not even satisfy WN. We show a general version of this result: no market for the median, or indeed any quantile, can satisfy WN. On the other hand, while the mean market could not satisfy WCL except on bounded domains, there are median and quantile markets that can.

Our setting in this subsection is as follows. Letting $\mathcal{R} = \mathcal{Y} = \mathbb{R}$ and $\alpha \in (0, 1)$, the α -quantile of probability distribution p with continuous CDF is the q_α satisfying $\Pr_p[Y \leq q_\alpha(p)] = \alpha$. (The continuity assumption is dropped in the full version of the paper.) Of course, the median is simply q_α for $\alpha = 1/2$.

We first show that quantile markets do not satisfy WN: there may not be trades which improve the trader liability at all. Figure 3 gives an example with absolute loss.

► **Theorem 10.** *No SRM for any α -quantile satisfies WN.*

Despite this negative result, quantile markets satisfy a surprisingly strong positive property. Recall that the squared loss market with $S(r, y) = 2ry - r^2$, which elicits the mean of Y , could not hope to satisfy bounded worst-case loss if $\mathcal{Y} = \mathbb{R}$. And indeed, the absolute loss market $S(r, y) = -|r - y|$ shares this issue. There is an elegant work-around, however: one can use the sigmoid function $g(r) = e^r / (1 + e^r)$, or another strictly monotone transformation, to map reports continuously into the interval $(0, 1)$. Then $S(r, y) = -|g(r) - g(y)|$ is clearly bounded, but still proper, as strictly monotone functions commute with the median: for any y , all $y' \leq y$ are mapped below y and all $y' \geq y$ are mapped above y , so the quantiles are simply mapped as well.

► **Theorem 11.** *For all $\alpha \in (0, 1)$, there is an SRM for the α -quantile satisfying WCL.*

Finally, quantiles also behave nicely with respect to bounded-budget traders. To see this for $S(r, y) = -|r - y|$, notice that a small trade $r \rightarrow r + \epsilon$ carries a liability of only ϵ , in the case $Y < r$. (See Figure 3.) But any trade smaller than ϵ can carry positive expected payoff for a trader if it moves the market closer to her belief.

► **Theorem 12.** *If distributions in \mathcal{P} do not contain point masses, then any α -quantile market satisfies BTB.*

5 Characterizing Trade Neutralization

In this section, we will characterize mechanisms that satisfy trade neutralization (TN) when the set of outcomes is finite, $|\mathcal{Y}| = n < \infty$. Recall that TN captures the quality of traditional markets, particularly in the sense of a “market maker”: the mechanism is always willing to buy back a contract it has previously sold, for a reasonable price.

TN may seem intuitively easy to satisfy, or at least might not appear to impose much structure on the market maker. In fact, we will see that the opposite is the case. To gain intuition for why TN might impose structure, recall that trades made by a participant must be interpretable as beliefs about the property the market is predicting. This applies to both the purchase of a contract and its sale back to the mechanism. Furthermore, this “canceling” trade must be available regardless of the current state of the market. We will leverage this intuition to show that the mechanisms satisfying TN are quite special and closely related to expectation markets.

In our analysis, it will prove useful to separate out “contracts” $d \in \mathbb{R}^{\mathcal{Y}}$ from “cash” $\mathbf{1} \in \mathbb{R}^{\mathcal{Y}}$, even though technically both reside in the same space. In particular, we would like a canonical way to take a contract d and separate out its “cash” component as $d = d_0 + c\mathbf{1}$, where intuitively $\$c$ is always paid regardless of the outcome, but d_0 depends on outcome. To do this, we simply define d_0 as the projection of d onto the hyperplane normal to $\mathbf{1}$.


We define the range of a scoring rule S by $\mathcal{S} = \{\vec{S}(r) : r \in \mathcal{R}\} \subseteq \mathbb{R}^{\mathcal{Y}}$, and define the corresponding “cashless” contract space by $\mathcal{H} \doteq \mathcal{S}/\mathbf{1} = \{\vec{S}(r) - (\vec{S}(r) \cdot \mathbf{1}/|\mathcal{Y}|)\mathbf{1} : r \in \mathcal{R}\}$. (Recall that one can define $\vec{S}(r) = \vec{F}(r|\emptyset)$. We indicate throughout where we assume IC.) Again, \mathcal{H} is simply the projection of the range of \vec{S} onto the hyperplane normal to $\mathbf{1}$. Similarly, we define $\mathcal{D} \doteq \mathcal{H} - \mathcal{H} = \{h_2 - h_1 : h_1, h_2 \in \mathcal{H}\}$ to be the set of possible score differences modulo $\mathbf{1}$. Equivalently, \mathcal{D} is the projection of the range of \vec{F} onto the same hyperplane.

We now show that if an SRM satisfies TN, its corresponding set \mathcal{D} must have considerable structure: it must form an additive subgroup of $\mathbb{R}^{\mathcal{Y}}$.

► **Lemma 13.** *If SRM F satisfies TN, then \mathcal{D} is an additive subgroup of $\mathbb{R}^{\mathcal{Y}}$ in the sense that $d, d' \in \mathcal{D}$ implies $-d, d + d' \in \mathcal{D}$. Moreover, the entire set \mathcal{D} of contracts is available at all times: for any $h \in \mathcal{H}$, we have $\mathcal{D} = \mathcal{H} - \{h\} = \{h' - h \mid h' \in \mathcal{H}\}$.*

Proof. From the definition of TN, we must have the following: $\forall r_1, r_2, r_3 \exists r_4$ with $F(r_2|r_1) + F(r_4|r_3) = c\mathbf{1}$ for some $c \in \mathbb{R}$. In terms of \mathcal{H} , as we have taken everything modulo $\mathbf{1}$, this implies $\forall h_1, h_2, h_3 \in \mathcal{H} \exists h_4 \in \mathcal{H}$ such that $h_2 - h_1 + h_4 - h_3 = 0$. In other words, $h_1, h_2, h_3 \in \mathcal{H} \implies h_3 + h_1 - h_2 \in \mathcal{H}$.

To show that \mathcal{D} is a group under vector addition, we must show closure under addition and the existence of additive inverses. The latter is simple, as given any $d \in \mathcal{D}$ we can write $d = h_2 - h_1$, so of course $-d = h_1 - h_2 \in \mathcal{D}$. We also have $0 = h_1 - h_1 \in \mathcal{D}$. For closure, consider $d, d' \in \mathcal{D}$, and write $d = h_2 - h_1$ and $d' = h'_2 - h'_1$. By the above, $(h_2 + h'_2 - h'_1) \in \mathcal{H}$, so $h_2 - (h_1 + h'_1 - h'_2) = d + d' \in \mathcal{D}$.



■ **Figure 4** The “discretized” LMSR. This market for $\Pr[Y = 1]$ allows only integer trades of a single security, using the LMSR cost function $C(q) = \log(1 + e^q)$ and its conjugate $G(p) = p \log(p) + (1 - p) \log(1 - p)$. Traders may move the market from any “dot” to any other. The possible market states are equally spaced in share space (horizontal axis, left figure), translating to *non*-equally spaced *prices* (horizontal axis, right figure). Theorem 16 implies that all mechanisms satisfying PI, IC, and TN have this general format, being discretizations of cost-function markets.

Finally, we must show that for any $h \in \mathcal{H}$, we have $\mathcal{D} = \mathcal{H} - \{h\}$. Clearly $\mathcal{H} - \{h\} \subseteq \mathcal{D}$; for the converse, consider $d = h_1 - h_2$. By the above, $h + h_1 - h_2 \in \mathcal{H}$, and thus $d = (h + h_1 - h_2) - h \in \mathcal{H} - \{h\}$. \blacktriangleleft

The implication of Lemma 13 is that under TN, we must have a highly structured set \mathcal{D} of possible contracts (modulo cash) available at all times with only the cash price for each contract changing depending on history. For example, a construction allowed by Lemma 13 is $\mathcal{D} = \mathbb{Z}^k$ (the integer lattice of dimension k), and this is indeed possible. Looking ahead, we will see that this would be interpreted as offering to buy or sell an integer number of “shares” in each of k different “securities” at any time. We give a $k = 1$ example in Figure 4.

We will need the following generalization of standard cost-function based markets.

► **Definition 14.** The *generalized cost-function based market* parameterized by a convex “cost function” $C : \mathbb{R}^k \rightarrow \mathbb{R}$, “securities” $\phi : \mathcal{Y} \rightarrow \mathbb{R}^k$, and “share space” $\mathcal{Q} \subseteq \mathbb{R}^k$ is the SRM of the form $F(r', y | r) = C(r) - C(r') + (r' - r) \cdot \phi(y)$ with report space at market state r given by $\mathcal{R} = \{r + q : q \in \mathcal{Q}\}$.

Recall that in standard cost-function markets, $\mathcal{Q} = \mathbb{R}^k$, so that any trade $d = r - r' \in \mathbb{R}^k$ is allowable. Furthermore, here d_i is interpreted as the number of shares purchased of security ϕ_i . (For example, in the $\mathcal{D} = \mathbb{Z}^k$ construction mentioned above, we also have $\mathcal{Q} = \mathbb{Z}^k$, so traders may only purchase integer numbers of shares in this example.)

► **Definition 15.** A cost-function based market defined by C and ϕ is *open* if C is differentiable and⁴ $\{\nabla C(q) : q \in \mathbb{R}^k\} = \text{int}(\text{conv}(\phi(\mathcal{Y})))$. It is *quasi-open* if for every pair of “share vectors” $q, v \in \mathcal{Q}$ and subgradient $x \in \partial C(q)$, we have $x \cdot v < \max_{y \in \mathcal{Y}} v \cdot \phi(y)$.

To see that quasi-openness is a relaxation of openness, we comment that if we required the condition to hold for all $v \in \mathbb{R}^k$, then it would be equivalent to $\cup_{q \in \mathcal{Q}} \partial C(q) \subseteq \text{int}(\text{conv}(\phi(\mathcal{Y})))$, which is a clear relaxation of openness; and quasi-openness requires somewhat less, as it only must hold for all $v \in \mathcal{Q}$.

⁴ We write $\text{int}(A)$ for the interior of the set A and $\text{conv}(A)$ for its convex hull.

► **Theorem 16.** *Let SRM F which is IC for Γ be given. Then F satisfies TN if and only if it is a generalized cost-function-based market for securities $\phi : \mathcal{Y} \rightarrow \mathbb{R}^k$, where:*

- (i) *The share space \mathcal{Q} of possible purchases is an additive subgroup of \mathbb{R}^k , and*
- (ii) *The market is quasi-open.*

Moreover, TN implies PN.

Proof. For the forward direction, that TN implies cost-function-based and PN, we proceed in five parts: (1) construct a basis for \mathcal{D} , which will be the securities ϕ ; (2) rewrite \vec{F} in terms of that basis; (3) show that the map from reports to “shares” is bijective; (4) extract the cost function C from the share representation of F ; (5) show that C satisfies our relaxed version of openness. For the converse, we will appeal to Theorem 22.

1. *Basis of \mathcal{D} .* By Lemma 13, \mathcal{D} is an additive subgroup of $\mathbb{R}^{\mathcal{Y}}$. Let $d^1, \dots, d^k \in \mathcal{D}$ be a basis for the linear span of \mathcal{D} . We can now write our securities in terms of this basis: define $\phi : \mathcal{Y} \rightarrow \mathbb{R}^k$ by $\phi(y)_i = d_y^i$. It will be convenient to work with ϕ in matrix form $\Phi \in \mathbb{R}^{n \times k}$, which naturally we define as $\Phi_{y,i} = \phi(y)_i$. Thus, we now have $\mathcal{D} = \{\Phi \cdot q : q \in \mathcal{Q}\}$ where \mathcal{Q} is an additive subgroup of \mathbb{R}^k .

2. *Rewrite \vec{F} .* By the definition of \mathcal{D} as the range of F modulo $\mathbb{1}$, and the decomposition above, we know that for every $r, r' \in \mathcal{R}$ we can write $\vec{F}(r'|r) = \Phi \cdot v + c\mathbb{1}$ for some $v \in \mathbb{R}^k$ and $c \in \mathbb{R}$. Thus, letting r_0 be the initial state of the given SRM, we have functions $g : \mathcal{R} \rightarrow \mathbb{R}$ and $v : \mathcal{R} \rightarrow \mathbb{R}^k$ such that for all $r \in \mathcal{R}$, $\vec{F}(r|r_0) = \Phi \cdot v(r) + g(r)\mathbb{1}$. By definition $\Phi \cdot v(r) \in \mathcal{D}$ and thus we must have $v(r) \in \mathcal{Q}$, meaning we can write $v : \mathcal{R} \rightarrow \mathcal{Q}$. Note that the expected payoff when $Y \sim p$ can now be written $\mathbb{E}_p F(r, Y|r_0) = p^\top \vec{F}(r|r_0) = p^\top \Phi v(r) + g(r)p^\top \mathbb{1} = \mathbb{E}_p \phi \cdot v(r) + g(r)$. In particular, Γ can only depend on p through $\mathbb{E}_p \phi$, so letting $\mathcal{X} = \text{conv}(\phi(\mathcal{Y}))$ as before, we have some $\psi : \mathcal{X} \rightrightarrows \mathcal{R}$ such that $\Gamma(p) = \psi(\mathbb{E}_p \phi)$.

3. *The map $v : \mathcal{R} \rightarrow \mathcal{Q}$ is a bijection.* Surjectivity of v follows from the fact that transformation by the change of basis matrix Φ is a bijection, as by definition we are simply rewriting elements of \mathcal{D} via elements of \mathcal{Q} . For injectivity, we will use IC. Suppose $v(r) = v(r')$. If $g(r) = g(r')$, Γ is not non-redundant as $\vec{F}(r|r_0) = \vec{F}(r'|r_0)$, so clearly $\Gamma(p) = r \iff \Gamma(p) = r'$. Thus without loss of generality $g(r) > g(r')$, which implies $r' \notin \text{argmax}_{x \in \mathcal{R}} v(x) \cdot \mathbb{E}_{p'} \phi(Y) + g(x) = \text{argmax}_{x \in \mathcal{R}} \mathbb{E}_{p'} F(x, Y|r_0)$ for any $p' \in \mathcal{P}$, so r' cannot be in the range of Γ at all (i.e. the report r dominates r'), a contradiction. Thus, v is injective and hence a bijection.

4. *Extracting the cost function.* Now that we have a bijection from reports to “shares”, intuitively, we should just be able to take $C(q) = -g(v^{-1}(q))$. More care is needed, however, as C must be convex with the correct subgradients. Let us define a convex function $G : \mathcal{X} \rightarrow \mathbb{R}$ by $G(x) = \sup_r v(r) \cdot x + g(r)$, which is the optimal expected score when $\mathbb{E}_p \phi = x$. By IC and the definition of ψ above, we have $G(x) = v(r) \cdot x + g(r)$ if and only if $r \in \psi(x)$. Now we can define $C : \mathbb{R}^k \rightarrow \mathbb{R}$ as the conjugate of G , namely $C(q) = \sup_{x \in \mathcal{X}} q \cdot x - G(x)$. Thus, we need only show that $C(v(r)) = -g(r)$, as alluded to above, as this would imply $F(r, y|r_0) = v(r) \cdot \phi(y) - C(v(r))$ as desired.

To make headway, we will appeal to convex analysis. First, we will establish that $G(x) = v(r) \cdot x + g(r) \iff v(r) \in \partial G(x)$, meaning if r is the optimal report for expected value x , then $v(r)$ is a subgradient of G . This follows directly from the subgradient inequality; if r is optimal for x , then for all x' , $G(x') = \sup_{r'} v(r') \cdot x' + g(r') \geq v(r) \cdot x' + g(r) = G(x) + v(r) \cdot (x' - x)$, so $v(r) \in \partial G(x)$. Conversely, suppose $v(r) \in \partial G(x)$ but r is optimal for some x' , so $G(x') = v(r) \cdot x' + g(r)$. The subgradient inequality gives $v(r) \cdot x' + g(r) = G(x') \geq G(x) + v(r) \cdot (x' - x)$, which implies $v(r) \cdot x + g(r) \geq G(x)$, so r must be optimal for x as well: $G(x) = v(r) \cdot x + g(r)$.

In summary, we have $r \in \psi(x) \iff G(x) = v(r) \cdot x + g(r) \iff v(r) \in \partial G(x)$. By [18, Thm E.1.4.1], we immediately have $C(q) = q \cdot x - G(x) \iff q \in \partial G(x)$. Putting these

together, we have $r \in \psi(x) \iff v(r) \in \partial G(x) \iff C(v(r)) = v(r) \cdot x - G(x) = -g(r)$. Thus, as every $r \in \mathcal{R}$ is in the image of ψ , we must have $C(v(r)) = -g(r)$ for all r .

5. *C is (almost) open.* We have established that F can be written as a cost-function-based market with securities ϕ and a potentially nondifferentiable C . To conclude this direction of the proof, suppose that for some $x \in \partial C(q)$ we have $x \cdot v \geq \max_{y \in \mathcal{Y}} v \cdot \phi(y)$. Following the proof of Theorem 22, we start the trader at q selling v (as $(-v) \in \mathcal{Q}$) and place the state back at q . To neutralize, the trader must now purchase $v \in \mathcal{Q}$, but by the same argument as before (using the fact that $\partial C(\mathbb{R}^k) \subseteq \mathcal{X}$ by construction of C), $C(q + v) - C(q) = x \cdot v \geq \max_{y \in \mathcal{Y}} v \cdot \phi(y)$, which contradicts TN.

Converse. Finally, for the reverse direction, we can simply repeat the proof of Theorem 22: if the trades v_i are elements of the group \mathcal{Q} for all v_i , then so is $v = \sum_{i=1}^w v_i$. It only remains to be shown that the price is less than the worst-case payoff of v , which follows from integrating the guarantee in (ii) as in Lemma 21. ◀

Theorem 16 shows that, despite the apparent freedom allowed by the TN axiom, the only SRMs satisfying TN are cost-function-based, or variations thereof (see Figure 4). By way of contrast, note that apart from other axiomatic characterizations with axioms such as “information incorporation”, previous work has only established that cost-function-based market makers elicit expectations and satisfy TN [2], whereas here we show that in fact they are the *only* markets satisfying TN.

Given that a market satisfying TN must elicit expectations (or “discretized” expectations), the following natural question arises: does WN allow for a broader family of markets, for properties beyond expectations? If false, this would mean that TN and WN were equivalent in a formal sense. We now give an example which answers this question in the affirmative: a market for the ratio of expectations which (by virtue of not eliciting a discretized expected value) does not satisfy TN, but does satisfy WN. In the full version we give another such example: expectiles. Together, these examples show that WN is a potentially useful condition for the design of non-expectation prediction markets, and warrants further exploration.

6 Share-Like Market for Ratios of Expectations

From Theorem 16, we know that any SRM satisfying the TN axiom will effectively be a cost-function-based market dealing in shares of some set of securities, perhaps restricting trades to some discretization of the share space. As a corollary, all such TN markets must elicit an expected value either directly or indirectly, or perhaps a discretization thereof. Theorem 16 leaves open the possibility, however, of a “share-like” market for a property other than an expected value which satisfies WN but not TN. We now give one such example: a ratio of expectations.

Before diving into the formalism, let us expand on the intuition given in § 1. In a cost-function-based market, traders purchase bundles $r \in \mathbb{R}^k$ of securities $\phi : \mathcal{Y} \rightarrow \mathbb{R}^k$, in exchange for paying some up-front cost $C(r_t + r) - C(r_t)$ to the market maker in the form of cash. In other words, traders are exchanging the bundle r of securities ϕ for an amount $C(r_t + r) - C(r_t)$ of the security $\mathbf{1}$. The ratio of expectations market can be thought of similarly, but now traders exchange the bundle r of securities ϕ for an amount $C(r_t + r) - C(r_t)$ of some new security b , which by convention will be nonnegative (and bounded away from 0 for infinite \mathcal{Y}). In other words, the new market is exactly the same as the old, but the cost function is in units of the security b rather than cash ($\mathbf{1}$). Immediately we can glean that, just as in a cost-function-based market traders with belief p had an incentive to trade until $\nabla C(r_t) = \mathbb{E}_p \phi$, i.e. the change in cost is equal to the expected security payoff, now traders

will trade until $\nabla C(r_t)\mathbb{E}_p b = \mathbb{E}_p \phi$. Thus, the prices of the new market will be the market's consensus about the ratio of expectations of ϕ and b , $\nabla C(r_t) = \mathbb{E}_p \phi / \mathbb{E}_p b$.

What axioms does this new market satisfy? It would seem that its share-like nature could somehow circumvent Theorem 16 and satisfy TN, but of course this is not the case. As we will show, however, the ratio of expectations market does satisfy WN. First, we must formally define our new market.

Let \mathcal{Y} be a finite set. For $\phi : \mathcal{Y} \rightarrow \mathbb{R}^k$ and $b : \mathcal{Y} \rightarrow \mathbb{R}$ with $\inf b > 0$, we define the ratio of expectations $\Gamma(p) = \mathbb{E}_p \phi / \mathbb{E}_p b$. As usual we assume that ϕ is affinely independent, so $\{\mathbb{E}_p \phi : p \in \mathcal{P}\}$ is full-dimensional, and that $\inf b > 0$, which since \mathcal{Y} is finite is equivalent to the payoffs $b(y)$ being positive for all $y \in \mathcal{Y}$. A characterization of scoring rules eliciting Γ was shown by Frongillo and Kash [10], which extended the real-valued case given by Gneiting [11]: a scoring rule S elicits Γ if and only if,

$$S(r, y) = b(y)G(r) + dG_r \cdot (\phi(y) - rb(y)) + f(y) , \quad (2)$$

where as in Theorem 18, $G : \mathcal{R} \rightarrow \mathbb{R}$ is strictly convex with selection of subgradients dG , and f is arbitrary (and, as usual, irrelevant for SRMs).

Our main result for this section is that essentially all “open” markets for Γ satisfy WN. The proof first applies convex conjugate duality to arrive at the cost-function-like market described above, and then applies facts about the usual cost-function-based framework (e.g. Lemma 21) to show that the prices for trading in a bundle must be strictly less than its worst-case payoff.

► **Theorem 17.** *A ratio-of-expectations market for differentiable G and $\nabla G(\mathcal{R}) = \mathbb{R}^k$ satisfies WN.*

Proof. We construct a “cost function” $C(q) = \sup_{r \in \mathcal{R}} r \cdot q - G(r)$ as the convex conjugate of G . The corresponding scoring rule is $S_C(q, y) = \phi(y) \cdot q - C(q)b(y)$. By assumption, $\nabla G : \mathcal{R} \rightarrow \mathbb{R}^k$ is a bijection, and thus markets F_G with G and F_C with C are equivalent. It therefore suffices to show that F_C satisfies WN.

Consider any trade $q_1 \rightarrow q'_1$ and let $v = q'_1 - q_1$. For any q_2 , we will show that the trade $q'_2 = q_2 - v$ satisfies WN. The difference in net payoff between the first and second trade is $(C(q_2) - C(q_2 - v))b(y) - v \cdot \phi(y)$, so it suffices to show that this is always positive.

By the proof of Lemma 21, we then have for any $q, w \in \mathbb{R}^k$

$$C(q + w) - C(q) < \sup_{x \in \mathcal{R}} x \cdot w = \max_{y \in \mathcal{Y}} (\phi(y)/b(y)) \cdot w . \quad (3)$$

Taking $q = q_2$ and $w = -v$, we have $C(q_2) - C(q_2 - v) < \phi(y) \cdot (-v)/b(y)$ for all y , which is equivalent to $(C(q_2 - v) - C(q_2))b(y) > \phi(y) \cdot v$ for all y , thus establishing WN. ◀

By establishing WN for SRMs eliciting a ratio of expectations, we are essentially saying that these markets are “reasonable” and could plausibly aggregate trader beliefs effectively. It would be interesting to see how such markets perform in practice.

7 Discussion and Directions

We have presented market axioms for scoring rule markets (SRMs) and studied a handful of well-known statistics/properties to see which of these axioms their corresponding markets can satisfy. Interestingly, we have seen wide variation among the satisfied axioms, most dramatic of which are the neutralization axioms: TN is satisfied essentially exclusively by cost-function-based markets, whereas median/quantile and mode markets do not even satisfy

its weakest version WN. This raised the question of whether any market satisfied WN but not TN, i.e., whether any non-cost-function-based markets satisfied WN. We saw that indeed markets for ratios of expectations are one such example, and in the full version we show that markets for expectiles are another. As discussed in § 3, WN is nearly equivalent to TN, and thus these markets could be expected to perform similarly to cost-function-based markets.

We conclude with future directions and discussion.

Open questions. Our work leaves several questions open. There are of course many other properties to explore and categorize with respect to our axioms, and perhaps the most compelling question along these lines would be a characterization of SRMs satisfying WN, or of properties elicited by WN SRMs. We would also like to show that the share space \mathcal{Q} in Theorem 16 is in fact isomorphic to $\mathbb{R}^{k_1} \oplus \mathbb{Z}^{k_2}$, meaning that shares could be purchased with arbitrary precision in some directions but in units of some smallest purchase in others; this isomorphism would also mean that IC, TN, and BTB would imply cost-function-based markets without qualification. Finally, a construction of pseudo-barriers for arbitrary convex sets would allow Proposition 8 to be an “if and only if”, and would be interesting in their own right.

Other market forms. [19] study prediction markets from a more empirical perspective, and suggest a number of possible market formulations. One which is not covered here is to offer contracts d^r such that $\mathbb{E}_p d^r = 0 \iff \Gamma(p) = r$. (This $V(r, y) = d^r(y)$ is called an *identification function*.) They illustrate this idea with a market eliciting the median of Y , where d^r pays \$1 if $Y > r$ and $-\$1$ otherwise. While the properties of such a contract space would be interesting to study in a continuous double-auction, one may ask how to translate it to an automated market maker. Here the market maker would maintain a centralized median r_t and traders could either buy or sell d^{r_t} , moving $r_{t+1} \leftarrow r_t \pm \epsilon$. Unfortunately, for any $\epsilon > 0$, such a market has unbounded worst-case loss even on a bounded domain (traders buy and sell between $r = 0$ and $r = \epsilon$ and $y = \epsilon/2$). For infinitesimal ϵ however, the market becomes the absolute-loss SRM with $S(r, y) = -|r - y|$, because $d^r = \frac{d}{dr} S(r, y)$.

Complexity. Finally, we remark that the recent concept of *elicitation complexity* brings interesting implications for our study. Here one asks, given a property Γ of interest, how many dimensions k does one need for there to exist an elicitable $\Gamma' : \mathcal{P} \rightarrow \mathbb{R}^k$ from some “nice” class of properties, where one can compute Γ from Γ' via a link function f , i.e. $\Gamma = f \circ \Gamma'$. In our context, one may choose “nice” to mean a property having an SRM satisfying any one of the axioms we discuss. The most natural may be TN, in which case one is essentially asking Γ' to be an expected value, a case studied by [6]. It would be interesting to characterize properties having markets satisfying WN, and identify properties with low complexity with respect to these WN properties.

Acknowledgements. Thanks to Yiling Chen for helpful comments and pointers, and several referees for constructive feedback. This material is based upon work supported by the National Science Foundation under Grant No. 1657598.

References

- 1 J. Abernethy and R. Frongillo. A characterization of scoring rules for linear properties. In *Proceedings of the 25th Conference on Learning Theory*, pages 1–27, 2012. URL: <http://jmlr.csail.mit.edu/proceedings/papers/v23/abernethy12/abernethy12.pdf>.

- 2 Jacob Abernethy, Yiling Chen, and Jennifer Wortman Vaughan. Efficient market making via convex optimization, and a connection to online learning. *ACM Transactions on Economics and Computation*, 1(2):12, 2013.
- 3 Jacob Abernethy, Sindhu Kutty, Sébastien Lahaie, and Rahul Sami. Information aggregation in exponential family markets. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pages 395–412. ACM, 2014.
- 4 Jacob D. Abernethy and Rafael M. Frongillo. A collaborative mechanism for crowdsourcing prediction problems. In *Advances in Neural Information Processing Systems 24*, pages 2600–2608, 2011.
- 5 Jacob D. Abernethy, Rafael M. Frongillo, Xiaolong Li, and Jennifer Wortman Vaughan. A general volume-parameterized market making framework. In Moshe Babaioff, Vincent Conitzer, and David Easley, editors, *ACM Conference on Economics and Computation, EC '14, Stanford, CA, USA, June 8-12, 2014*, pages 413–430. ACM, 2014. doi:10.1145/2600057.2602900.
- 6 Arpit Agarwal and Shivani Agarwal. On consistent surrogate risk minimization and property elicitation. In *JMLR Workshop and Conference Proceedings*, volume 40, pages 1–19, 2015. URL: <http://www.jmlr.org/proceedings/papers/v40/Agarwal15.pdf>.
- 7 G.W. Brier. Verification of forecasts expressed in terms of probability. *Monthly weather review*, 78(1):1–3, 1950.
- 8 Y. Chen and D.M. Pennock. A utility framework for bounded-loss market makers. In *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence*, pages 49–56, 2007.
- 9 Rafael Frongillo and Ian Kash. General truthfulness characterizations via convex analysis. In *Web and Internet Economics*, pages 354–370. Springer, 2014.
- 10 Rafael Frongillo and Ian Kash. Vector-Valued Property Elicitation. In *Proceedings of the 28th Conference on Learning Theory*, pages 1–18, 2015.
- 11 T. Gneiting. Making and Evaluating Point Forecasts. *Journal of the American Statistical Association*, 106(494):746–762, 2011.
- 12 Tilmann Gneiting and Adrian E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.
- 13 R. Hanson. Combinatorial Information Market Design. *Information Systems Frontiers*, 5(1):107–119, 2003.
- 14 Nicolas S. Lambert, David M. Pennock, and Yoav Shoham. Eliciting properties of probability distributions. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, pages 129–138, 2008.
- 15 Nicolas S. Lambert and Yoav Shoham. Eliciting truthful answers to multiple-choice questions. In *Proceedings of the 10th ACM conference on Electronic commerce*, pages 109–118, 2009.
- 16 R.T. Rockafellar. *Convex analysis*, volume 28 of *Princeton Mathematics Series*. Princeton University Press, 1997.
- 17 L.J. Savage. Elicitation of personal probabilities and expectations. *Journal of the American Statistical Association*, pages 783–801, 1971.
- 18 Jean-Baptiste Hiriart Urruty and Claude Lemaréchal. *Fundamentals of Convex Analysis*. Springer, 2001.
- 19 J. Wolfers and E. Zitzewitz. Prediction Markets. *Journal of Economic Perspective*, 18(2):107–126, 2004.

A Proofs for Mean and Expectation Markets

We note the characterization of scoring rules for the expectation, which has been established with varying degrees of generality [17, 1, 10]:

► **Theorem 18.** *A scoring rule S elicits $\Gamma : p \mapsto \mathbb{E}_p[\phi(Y)]$ if and only if $S(r, y) = G(r) + dG_r \cdot (\phi(y) - r) + f(y)$ for some G strictly convex with subgradients $\{dG_r\}_{r \in \mathcal{R}}$, and f an arbitrary \mathcal{P} -measurable function.⁵*

► **Proposition (7).** *The linear property $\Gamma(p) = \mathbb{E}_p \phi(Y)$ has an expectation market satisfying WCL if and only if its domain $\text{conv}(\phi(\mathcal{Y}))$ is bounded, in which case the market defined by any bounded G satisfies WCL.*

Proof. Lemma 19 proves that on a bounded domain, WCL is equivalent to boundedness of G . (Some bounded convex G always exists for a bounded domain.) Lemma 20 proves that WCL cannot be satisfied on an unbounded domain. ◀

► **Lemma 19.** *On a bounded domain $\text{conv}(\phi(\mathcal{Y}))$, an expectation market defined by G satisfies WCL if and only if G is bounded.*

Proof. Let the initial market state be any r where $d_G(r)$ is finite, such as any r in the interior of the domain.⁶ We show that for this fixed r , worst-case loss is bounded by a constant if and only if G is bounded.

Recall from the characterization that $S(r, y) = G(r) + d_G(r) \cdot \phi(y)$. By bounded domain, $\|\phi(y)\|$ is bounded, so there exists a constant B with $-B \leq S(r, y) \leq B$ for all y .⁷

The loss of the market maker when the final state is r' is $S(r', y) - S(r, y)$, and the worst-case loss is $WCL \doteq \sup_{r', y} [S(r', y) - S(r, y)]$. We have $WCL \in (\sup_{r', y} S(r', y)) \pm B$, hence it is bounded if and only if the first term is. For each y , the supremum over r is achieved at $r = \phi(y)$ by properness of the scoring rule (r is the optimal report for the distribution δ_y), so the first term is $\sup_y S(\phi(y), y) = \sup_y G(\phi(y))$. So worst-case loss is bounded if and only if G is. ◀

► **Lemma 20.** *On an unbounded domain, no expectation market satisfies WCL, assuming the initial market prediction lies in the interior of \mathcal{R} .*

Proof. Let r be the market's initial starting point. We assume that r is in the interior of the convex hull of $\{\phi(y) : y \in \mathcal{Y}\}$, and in particular lies in an ϵ -ball contained in the interior. The idea is to pick a “direction” and consider a sequence of outcomes that are farther and farther away. If the final market state is some distance in that direction, then loss will be unbounded.

To formalize this, let y_1, y_2, \dots be a sequence with $\|\phi(y_i) - r\|$ positive and increasing without bound. Such a sequence must exist by unboundedness of the domain. Consider the associated sequence of size- ϵ vectors $\{v_i = \epsilon \frac{\phi(y_i) - r}{\|\phi(y_i) - r\|} : i = 1, 2, \dots\}$. As a sequence in a compact set (the ϵ -sphere), it has a convergent subsequence with a limit v . This is the

⁵ Typically scores are allowed to take on values in $\mathbb{R} \cup \{\infty\}$, essentially to accommodate the log scoring rule, but we will typically restrict to the relative interior of the domain anyhow, thus avoiding this issue.

⁶ This technical condition only rules out boundary such as the log scoring rule with initial prediction $p = 0$.

⁷ If $f(y) \neq 0$, then we can say $S(r, y) + f(y) \in [-B, B] + f(y)$, and we would see $f(y)$ cancel out later in the proof.

“direction”. By monotonicity of strictly convex functions, we have $\gamma \doteq v \cdot (dG_{r+v} - dG_r) > 0$. Let $K \subseteq \mathbb{N}$ be the indices of this subsequence.

Now consider the following sequence of markets closing states and outcomes, indexed by $i \in K$. Each market i has initial state r and final state $r + v$. The outcome is $Y = y_i$. The initial state is valid because r was fixed at the beginning of the proof, and $r + v$ is a valid final closing state because it lies on the ϵ -ball around r , assumed to lie in the report space because r was in the interior. The worst-case loss is at least the following quantity, which is then rearranged:

$$\begin{aligned} \text{Loss} &= S(r + v, \phi(y_i)) - S(r, \phi(y_i)) \\ &= G(r + v) - G(r) + dG_{r+v} \cdot (\phi(y_i) - r - v) - dG_r \cdot (\phi(y_i) - r) \\ &= [G(r + v) - G(r) - dG_r \cdot (r + v - r)] + (dG_{r+v} - dG_r) \cdot (\phi(y_i) - r - v) \\ &\geq (dG_{r+v} - dG_r) \cdot (\phi(y_i) - r - v) \end{aligned}$$

using the definition of subgradient to conclude that the bracketed term is at least zero. Now we divide both sides by $\|\phi(y_i) - r\|$ and recall our definition of v_i :

$$\frac{\text{Loss}}{\|\phi(y_i) - r\|} \geq \epsilon [(dG_{r+v} - dG_r) \cdot v_i] - \left[(dG_{r+v} - dG_r) \cdot \frac{v}{\|\phi(y_i) - r\|} \right].$$

The first term on the right side converges to $\epsilon\gamma$ because $v_i \rightarrow v$; the second term is equal to $\gamma/\|\phi(y_i) - r\|$ which converges to zero. The ratio on the left therefore either diverges or converges to some constant larger than $\epsilon\gamma$, and in either case, worst-case loss is unbounded. ◀

► **Proposition (8).** *The linear property $\Gamma(p) = \mathbb{E}_p\phi(Y)$ has an expectation market satisfying BTB if its domain $\text{conv}(\phi(\mathcal{Y}))$ is bounded, in which case the market defined by any differentiable G satisfies BTB.*

Proof. Suppose G is differentiable and its domain is bounded, e.g. $\|x\| \leq B$ for all x . Let x^0 be the market state and consider belief μ . By monotonicity of strictly convex functions, any trade $x' = \alpha\mu + (1 - \alpha)x^0$ has strictly positive expected score. Meanwhile, the worst-case score for any trade from x^0 to x' is

$$\begin{aligned} &\sup_x G(x^0) - G(x') + dG_{x'} \cdot (x - x') - dG_{x^0} \cdot (x - x^0) \\ &= \sup_x G(x^0) - G(x') + x \cdot (dG_{x'} - dG_{x^0}) - x' \cdot dG_{x'} - x^0 \cdot dG_{x^0} \\ &\leq B\|dG_{x'} - dG_{x^0}\| + O(\|x^0 - x'\|) \end{aligned}$$

where both terms can be made arbitrarily small with $\alpha \rightarrow 0$, $x' \rightarrow x^0$: G is continuous, and it is a convex differentiable function so dG is as well. ◀

► **Theorem (9).** *On a finite outcome space, i.e. $|\mathcal{Y}| < \infty$, for any linear property there exists an expectation market satisfying PN, WCL, and BTB.*

Proof. We utilize Theorem 22, which says that if a cost-function based market is open (Definition 15), then it satisfies PN.

The primary examples are *exponential-family* markets [3], where

$$C(q) = \log \sum_{y \in \mathcal{Y}} \exp(q \cdot \phi(y)),$$

the “log-partition” function. We recall the key facts behind the construction and refer the reader to [3]. One interprets $q \cdot \phi(y)$ as a “weight” on y and define the (exponential-family) distribution $p \in \text{int}(\Delta_{\mathcal{Y}})$ with $p_y = \frac{\exp(q \cdot \phi(y))}{\sum_{y \in \mathcal{Y}} \exp(q \cdot \phi(y))}$. One obtains

$$\nabla C(q) = \sum_y p_y \phi(y)$$

so the set of gradients is exactly the interior of $\text{conv}(\phi(\mathcal{Y}))$, i.e. the market is open.

By Theorem 22, this implies TN. The convex conjugate $G(\mu)$ is bounded (equaling the negative entropy of p); this implies WCL by Lemma 19. Finally, it and C are both strictly convex and differentiable. This implies BTB by Proposition 8. ◀

► **Lemma 21.** *Given $\phi : \mathcal{Y} \rightarrow \mathbb{R}^k$ such that $\mathcal{X} := \text{conv}(\phi(\mathcal{Y}))$ is full-dimensional in \mathbb{R}^k , let $C : \mathbb{R}^k \rightarrow \mathbb{R}$ be convex with subgradients $\partial C(\mathbb{R}^k) \subseteq \text{int}(\mathcal{X})$. Then For all $q, v \in \mathbb{R}^k$, $\max_{y \in \mathcal{Y}} v \cdot \phi(y) > C(q + v) - C(q)$.*

Proof. As $\partial C(q') \subseteq \text{int}(\mathcal{X})$ for all $q' \in \mathbb{R}^k$, in particular $dC(q') \cdot v < \max_{x \in \mathcal{X}} x \cdot v = \max_{y \in \mathcal{Y}} v \cdot \phi(y)$ [18, Prop A.2.4.6], where dC is a selection of subgradients of C . By [9, Thm B.4], the function $t \mapsto dC(q + tv) \cdot v$ is monotone and therefore integrable, and thus $C(q + v) - C(q) = \int_{t=0}^1 dC(q + tv) \cdot v dt < \max_{y \in \mathcal{Y}} v \cdot \phi(y)$. ◀

► **Theorem 22.** *Cost-function-based markets satisfy TN if and only if they are open. Moreover, if they satisfy TN, they satisfy PN.*

Proof. Suppose a cost-function-based market is open. We will show that it satisfies not only TN but PN. Consider a non-empty set of bundles purchased $v_i = q'_i - q_i$ for $1 \leq i \leq m$, for a total cost of $c = \sum_i C(q'_i) - C(q_i)$, and let $v = \sum_{i=1}^m v_i$ be their sum. Clearly, to neutralize this position from the current market state q , the trader must sell v , for a cost of $C(q - v) - C(q)$. After this trade, the trader’s total contract is $c\mathbf{1} - (C(q - v) - C(q))\mathbf{1}$, so to establish PN, we need only show $c - C(q - v) + C(q) > \inf_{y \in \mathcal{Y}} [c + (-v) \cdot \phi(y)]$. Subtracting c from both sides and then negating, the rest follows from observing $\sup = \max$ as \mathcal{Y} is finite, and applying Lemma 21.

Now suppose a cost-function-based market satisfies TN; we will show it is open. By assumption, such a market has a differentiable C with $\text{cl}(\{\nabla C(q) : q \in \mathbb{R}^k\}) = \text{conv}(\phi(\mathcal{Y})) =: \mathcal{X}$. Suppose the market is not open; this implies that $\nabla C(q)$ lies on the boundary of \mathcal{X} for some q . As \mathcal{X} is a convex polytope, $\nabla C(q)$ must lie on an exposed face of \mathcal{X} , so let $v \in \mathbb{R}^k$ be a direction exposing that face, meaning $\max_{x \in \mathcal{X}} x \cdot v = \nabla C(q) \cdot v$ [18, Sec A.2.4].

Now suppose the trader buys the bundle $(-v)$ at state q , and the market state has returned to q , which corresponds to $q_1 = q$, $q'_1 = q - v$, $q_2 = q$. To satisfy TN, the trader must now purchase v , but we must additionally have $\inf[(C(q) - C(q - v))\mathbf{1} + (C(q) - C(q + v))\mathbf{1}] > \inf[(C(q) - C(q - v))\mathbf{1} + (-v) \cdot \phi]$, which is equivalent to $C(q + v) - C(q) < \max_{y \in \mathcal{Y}} v \cdot \phi(y)$. By weak monotonicity [16, Thm 24.9], the map $t \mapsto \nabla C(q + tv) \cdot v$ is monotone increasing in t , and thus as $\nabla C(q + tv) \in \mathcal{X}$ and $\nabla C(q) \cdot v = \max_{x \in \mathcal{X}} x \cdot v$, we have $\nabla C(q + tv) \cdot v = \nabla C(q) \cdot v$ for all $t \geq 0$. But now we have $C(q + v) - C(q) = \int_{t=0}^1 \nabla C(q + tv) \cdot v dt = \nabla C(q) \cdot v = \max_{y \in \mathcal{Y}} v \cdot \phi(y)$, so the trader’s minimum payoff has not increased, violating TN. ◀