An Axiomatic Characterization of CFMMs and Equivalence to Prediction Markets

Rafael Frongillo
University of Colorado, Boulder, CO, USA

Maneesha Papireddygari
University of Colorado, Boulder, CO, USA

Bo Waggoner
University of Colorado, Boulder, CO, USA

Abstract

Constant-function market makers (CFMMs), such as Uniswap, are automated exchanges offering trades among a set of assets. We study their technical relationship to another class of automated market makers, cost-function prediction markets. We first introduce axioms for market makers and show that CFMMs with concave potential functions characterize "good" market makers according to these axioms. We then show that every such CFMM on n assets is equivalent to a cost-function prediction market for events with n outcomes. Our construction directly converts a CFMM into a prediction market, and vice versa. Using this equivalence, we give another construction which can produce any 1-homogenous, increasing, and concave CFMM, as are typically used in practice, from a cost function.

Conceptually, our results show that desirable market-making axioms are equivalent to desirable information-elicitation axioms, i.e., markets are good at facilitating trade if and only if they are good at revealing beliefs. For example, we show that every CFMM implicitly defines a *proper scoring rule* for eliciting beliefs; the scoring rule for Uniswap is unusual, but known. From a technical standpoint, our results show how tools for prediction markets and CFMMs can interoperate. We illustrate this interoperability by showing how liquidity strategies from both literatures transfer to the other, yielding new market designs.

2012 ACM Subject Classification Theory of computation → Algorithmic game theory; Mathematics of computing → Information theory

Keywords and phrases Convex analysis, Equivalence result, Axiomatic characterization, Market Makers, Prediction markets, Scoring rules, Cost-functions

Digital Object Identifier 10.4230/LIPIcs.ITCS.2024.51

Related Version *Full Version*: `https://arxiv.org/abs/2302.00196`

Funding *Rafael Frongillo*: National Science Foundation under Grant No. IIS-2045347.

Acknowledgements We thank David Pennock, Daniel Reeves, Anson Kahng, and Manifold Markets for collaboration in working out the cost function corresponding to Uniswap. We also thank Scott Kominers, Ciamac Moallemi, Abe Othman, Tim Roughgarden, and Christoph Schlegel for helpful discussions.

1 Introduction

A prediction market is designed to elicit forecasts, such as for elections or sporting events. Participants trade shares in n securities, each tied to one of the n possible outcomes, paying out \$1 if that outcome occurs; the prices form a probability distribution over the outcomes representing an aggregate belief. To resolve thin market problems, [23] and [11] propose introducing *automated market makers*: mechanisms that offer to buy or sell any number of securities, with prices determined by the past history of transactions. A significant amount

15th Innovations in Theoretical Computer Science Conference (ITCS 2024).
Editor: Venkatesan Guruswami; Article No. 51; pp. 51:1–51:21
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

of research has investigated the design of such markets based on convex *cost functions* C [1, 17]. To acquire a bundle of securities $\mathbf{r} \in \mathbb{R}^n$, the trader pays $C(\mathbf{q}+\mathbf{r}) - C(\mathbf{q})$ cash to the market maker, where $\mathbf{q} \in \mathbb{R}^n$ is the total number of securities sold so far. A prominent example is the log market scoring rule (LMSR) due to [23], given by $C(\mathbf{q}) = b \log \sum_{i=1}^n e^{q_i/b}$ for a liquidity parameter $b > 0$.

More recently, inspired by blockchain applications, there has been significant theoretical and practical interest in the design of decentralized financial markets [10, 5]. These markets allow trade between n assets without a fixed unit of exchange ("cash"). Interestingly, like prediction markets, many decentralized exchanges also employ automated market makers, but for different reasons: automated market makers tend to have lower on-chain implementation costs than traditional order books [24, §4.1]. The dominant paradigm is the *constant-function market maker (CFMM)* based on a concave *potential function* φ. Here, if $\mathbf{q} \in \mathbb{R}^n$ is the amount of market maker holdings or "reserves" in the n assets, then a trade $\mathbf{r} \in \mathbb{R}^n$ is accepted if $\varphi(\mathbf{q}+\mathbf{r}) = \varphi(\mathbf{q})$, i.e., it keeps the potential function constant. One of the most popular CFMMs is the constant-product market maker, as used by Uniswap, where $\varphi(\mathbf{q}) = \sqrt{q_1 q_2}$. While most of our examples are 2-dimensional, our results apply generally to any n assets.

Despite clear similarities between these two settings, prior work has considered prediction markets to be only analogous to CFMMs, but not technically related.[1] Indeed, there are obvious differences between the settings: prediction markets rely on "cash", while CFMMs do not require a special unit of account; prediction markets only offer certain types of outcome-dependent securities (and can manufacture as many of these as they want), while CFMMs deal in arbitrary assets in limited supply; and so on. Furthermore, the *goals* of the two types of markets are very different: prediction markets seek to elicit information, whereas CFMMs seek to facilitate trade.

Nonetheless, we show a tight technical equivalence between prediction markets and CFMMs. We give straightforward reductions to transform a cost function C for a prediction market into a potential function φ for a CFMM, and vice versa. The reductions give a one-to-one correspondence between the trade rules and history in both markets while preserving the relevant market properties. An important implication is that markets designed to elicit information are also good at facilitating trade, and vice versa. We discuss the new insights this equivalence provides in § A.2. We illustrate our results with several examples; among them, the prediction market equivalent to Uniswap, a new homogeneous CFMM from LMSR (Figure 1), and a close connection between Brier score and a hybrid CFMM. We hope our results can inspire new market designs and insights in both literatures.

1.1 Outline of results

The three major contributions of this paper are as follows.

1. **Axiomatic Characterization of CFMMs**: In § 2 we show that an automated market maker satisfies a set of basic trading axioms if and only if it can be represented by a CFMM with a increasing, quasiconcave, and continuous φ. In § 4 we add 1-homogeneity to this list with a stronger construction. These results provide a substantial theoretical

[1] To give two examples, [6, §3.2] argue that the two market makers can diverge, and [8] state the following: "In a separate context, [automated market makers] for prediction markets were first proposed in [23]. Such a structure is fundamentally different from the constant function market makers considered herein insofar as a prediction market includes a terminal time at which bets are realized."

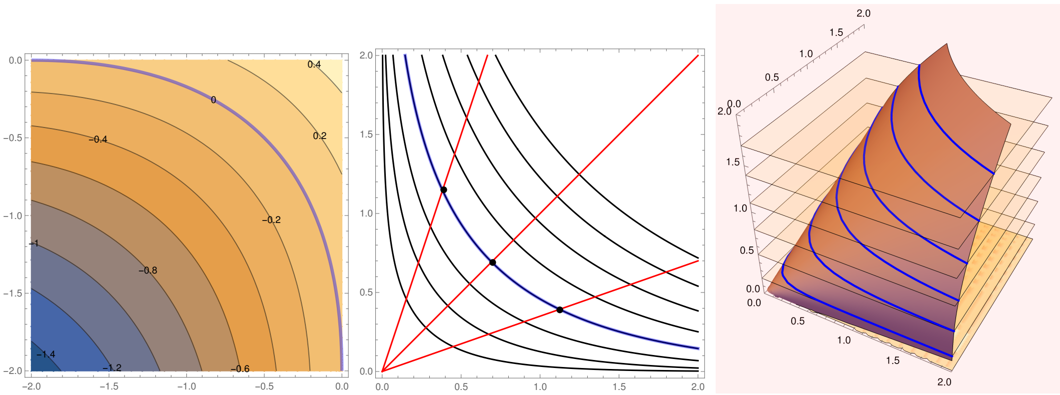

Figure 1 The "reserves-aware" Construction 1 applied to the log market scoring rule (LMSR) of [23]. (L) The level sets of LMSR, with the 0 level set highlighted in blue. (M) Taking the 0 level set of LMSR and reflecting it into the positive orthant, we assign $\varphi = 1$ on this set (again in blue). We then derive the other level sets by shrinking or expanding uniformly towards or away from the origin: the α level set, $\alpha > 0$, is given by scaling the 1 level set by α. As we will see, this value α can be interpreted as the liquidity level of the market. (R) The resulting 1-homogeneous potential φ, which satisfies all of our axioms (§ 4.1).

justification for the use of such CFMMs in decentralized finance. In contast to the existing literature, we do not presuppose CFMM structure, but rather show that this structure follows from the axioms; see § 1.3.

2. **Equivalence between prediction markets and CFMMs**: We prove a strong mechanical equivalence between cost function prediction market and CFMMs (Theorem 7, 8). In particular, the rules are market are the same: both markets have the same set of available trades after a given history of trades. This equivalence holds regardless of the securities or assets that the markets trade, so long as they are equal in number.
3. **Reserve-aware 1-homogeneous construction**: In § 4, we give a construction for CFMMs that satisfy all the axioms in previous sections and also do not run out of reserves. We show that this construction characterizes the set of all 1-homogeneous, concave and increasing CFMMs, which are dominant in practice. The construction takes a cost function whose level sets are parallel (Figure 1(L)) and "stretches" them with curvature adapting to the amount of reserves available (Figure 1(M)).

1.2 Equivalence in a nutshell

Let us illustrate the key intuition for our equivalence. First, suppose we have a cost function C defining a prediction market on n outcomes. We will turn it into a CFMM, i.e. a rule for pricing trades among any set of n assets, without money. The first step is to observe that a unit of cash is equivalent to the "grand bundle", represented as the all-ones vector $\mathbb{1} \in \mathbb{R}^n$, containing one share of each security: exactly one of them will pay off \$1, and the rest \$0, so the grand bundle is worth exactly \$1. We can therefore simulate transactions in a prediction market with only the n securities and no cash: a trade of the bundle $\mathbf{r}' \in \mathbb{R}^n$ in exchange for $c \in \mathbb{R}$ units of cash can be expressed as the combined trade $\mathbf{r} = \mathbf{r}' - c\mathbb{1}$.

We can now simply set $\varphi(\mathbf{q}) = -C(-\mathbf{q})$ and, perhaps surprisingly, obtain a CFMM where valid trades – that is, trades $\mathbf{r}$ such that $\varphi(\mathbf{q}+\mathbf{r}) = \varphi(\mathbf{q})$ – are equivalent to those trades allowed by the original cost function, after converting cash into the grand bundle. That is, the trades allowed by the CFMM are exactly of the form $\mathbf{r} = \mathbf{r}' - c\mathbb{1}$ for $c = C(\mathbf{q}+\mathbf{r}') - C(\mathbf{q})$, the price charged by the cost function. The negations in the formula $-C(-\mathbf{q})$ resolve a difference in sign conventions (i.e., whether $\mathbf{q}$ is interpreted as a gain or loss) and convert the convex C to a concave φ. We will show that concavity is equivalent to the desirable trading axiom of *demand responsiveness*; and φ also satisfies other nice trading axioms, like *liquidation*: the trader can purchase any bundle of assets for some amount of any given asset. Perhaps surprisingly, these and other market-making axioms come "for free", as long as C is a good prediction-market cost function. One important exception is that ideally, CFMMs do not deplete their reserves. As described above, we will fix this issue with a more sophisticated construction based on the perspective transform that guarantees nonnegative reserves. Figure 1 illustrates that construction applied to LMSR.

In the other direction, one can turn any CFMM into a prediction market with good elicitation properties. Given a CFMM with potential function φ, we show that it has an equivalent cost function defined as follows: Let $C(\mathbf{q}) = c$ for the constant c satisfying $\varphi(c\mathbb{1} - \mathbf{q}) = \varphi(\mathbf{q}_0)$ where $\mathbf{q}_0 \in \mathbb{R}^n$ is the vector of initial reserves and again $\mathbb{1} \in \mathbb{R}^n$ is grand bundle (the all-ones vector). For example, as we recover in § 3.4, the following cost function is equivalent to Uniswap $\varphi(\mathbf{q}) = \sqrt{q_1 q_2}$ when the initial reserves satisfy $\varphi(\mathbf{q}_0) = k$,

$$C(\mathbf{q}) = \frac{1}{2}\left(q_1 + q_2 + \sqrt{4k^2 + (q_1 - q_2)^2}\right) \ . \tag{1}$$

In fact, by standard prediction market facts, this general result implies that every "good" CFMM can be converted into a *proper scoring rule* [20]; we derive the scoring rule for Uniswap in § 3.4.

1.3 Related work

The literature on CFMMs contains several axiomatic results in a spirit related to our work in § 2, e.g. [6, 8, 33]. For example, [8] characterize the set of "ideal" CFMMs from among the class of all CFMMs and [33] axiomatically characterize some sub-classes of CFMMs. A major difference between our axiomatic results in § 2 and the above works is that we do not take the structure of CFMMs as given, for example, the existence of a potential/utility function or its concavity. We derive these attributes directly from the axioms. [22] explore optimal strategies of LP provisioning under certain LP beliefs about future asset prices. Our ratio-of-expectation insight in Appendix A.2 has more of a revealed preference flavor. [27] characterize mechanisms to elicit such price beliefs from traders.

Several authors have noted pieces of the equivalence we present. Most technically related is [11], which proposes the constant-utility market maker for prediction markets. Their market maker accepts trades that, according to some subjective fixed belief (probability distribution) and risk-averse utility function, maintain its expected utility at a constant value. Intriguingly, they observe in their eq. (14) that such a market maker with log utility is equivalent to a cost function of the form 1, which turns out to be the result of our conversion of Uniswap to a prediction market; see § 3.4. We emphasize that [11], unlike the current paper, focused only on prediction markets, i.e. do not show how to use their market maker for more general classes of assets. Within the prediction market context, they also only give equivalences to a certain class of cost functions, those corresponding to weighted pseudo-spherical scoring rules.

Other works have specifically connected market makers in decentralized finance and prediction markets. For example, [32] discusses how to apply the functional format of the LMSR as a potential function in a CFMM. Similarly, [26] and [21] apply Uniswap to the case where assets are contingent securities, obtaining a prediction market. Such works can be described as applying a functional form from one context to a different context, without justifying why they will perform well in the new context. More recent contemporary work [22] explores the direction of converting a cost-function to CFMM. In contrast, this work shows how to convert any cost function to a CFMM and vice versa via a reduction that guarantees to preserve axiomatic properties of the original.

Other technical tools come from the prediction market literature. We rely specifically on the cost-function market formulation [11, 1], axiomatic approaches to prediction markets [1, 3, 17], and the duality between prediction markets and scoring rules [23, 3]. We also make use of constructions from the literature on financial risk measures, specifically obtaining a convex risk measure from a set of "acceptable" positions [18, 19]; the connection between risk measures and prediction markets has been noted several times [29, 12, 4, 16, 2].

1.4 Notation

Vectors are denoted in bold, e.g. $\mathbf{q} \in \mathbb{R}^n$. The all-zeros vector is $\mathbf{0} = (0, \ldots, 0)$ and the all-ones vector is $\mathbb{1} = (1, \ldots, 1)$. The vector $\boldsymbol{\delta}_i = (0, \ldots, 1, \ldots, 0)$, i.e. all-zero except for a one in the i^{th} position. Comparison between vectors is pointwise, e.g. $\mathbf{q} \succ \mathbf{q}'$ if $q_i > q'_i$ for all $i = 1, \ldots, n$, and similarly for $\succeq$. We say $\mathbf{q} \gneq \mathbf{q}'$ when $q_i \geq q'_i$ for all i and $\mathbf{q} \neq \mathbf{q}'$. Define $\mathbb{R}^n_{\geq 0} = \{\mathbf{q} \in \mathbb{R}^n \mid \mathbf{q} \succeq \mathbf{0}\}$, $\mathbb{R}^n_{>0} = \{\mathbf{q} \in \mathbb{R}^n \mid \mathbf{q} \succ \mathbf{0}\}$, et cetera.

Let $f : \mathbb{R}^n \to \mathbb{R}$. We will use the following conditions.

- *increasing*: $f(\mathbf{q}) > f(\mathbf{q}')$ for all $\mathbf{q}, \mathbf{q}' \in \mathbb{R}^n$ with $\mathbf{q} \gneq \mathbf{q}'$.
- *convex*: $\forall x, y \in \mathbb{R}^n, \lambda \in [0, 1]$, $f(\lambda \cdot x + (1 - \lambda)y) \leq \lambda f(x) + (1 - \lambda)f(y)$.
- *concave*: $-f$ is convex.
- *quasiconcave*: $\forall c \in \mathbb{R}$, the set $\{x \mid f(x) \geq c\}$ is convex.
- $\mathbb{1}$*-invariant*: $f(\mathbf{q} + \alpha\mathbb{1}) = f(\mathbf{q}) + \alpha$ for all $\mathbf{q} \in \mathbb{R}^n$, $\alpha \in \mathbb{R}$.
- *1-homogeneous* (on $\mathbb{R}^n_{>0}$): $f(\alpha\mathbf{q}) = \alpha f(\mathbf{q})$ for all $\mathbf{q} \succ \mathbf{0}$ and $\alpha > 0$.

Given a list of vectors, e.g. $h = (\mathbf{r_1}, \ldots, \mathbf{r_t})$, the sum is denoted $\mathsf{sum}(h) = \mathbf{r_1} + \cdots + \mathbf{r_t}$. Concatenation of a new trade $\mathbf{r}$ onto a history h is denoted $h \oplus \mathbf{r}$.

2 A Characterization of Fixed-Liquidity CFMMs

To lay a foundation, we consider market makers that have "fixed liquidity": they do not accept loans nor charge transaction fees. The fixed-liquidity case allows for a simple, universal characterization: if one wants an automated market to satisfy certain natural axioms, it *must* be a CFMM as we show in Theorem 5.

2.1 Automated market makers for general asset markets

We consider automated market makers that process trades sequentially. There are n assets $\{1, \ldots, n\}$. A *trade* or *bundle* $\mathbf{r} \in \mathbb{R}^n$ represents r_i net units of each asset i being given to the market maker by a trader. A positive r_i represents a net transfer from the trader to the market maker, and a negative r_i represents a net transfer from the market maker to the trader. A *history* h is an ordered list of trades. The empty history is denoted $\epsilon = ()$.

A market maker is defined by a function $\mathsf{ValTrades}(h)$ that specifies, for any valid history h, the set of valid trades that it is willing to accept from the next participant. For each arriving participant $t = 1, \ldots$:

- The current market history is h_{t-1}.
- The valid trades are given by $\mathsf{ValTrades}(h_{t-1}) \subseteq \mathbb{R}^n$.
- Participant t selects some trade $\mathbf{r}_t \in \mathsf{ValTrades}(h_{t-1})$.
- The market history is now $h_t = h_{t-1} \oplus \mathbf{r}_t$.

The market maker begins with some *initial reserves* $\mathbf{q}_0$. It will be convenient to include the initial reserves in the history, so in the above we would have $h_0 = (\mathbf{q}_0)$. At any history h, we define the notation

$$\mathbf{q}_h = \mathsf{sum}(h) \ ,$$

meaning the current reserves are equal to the sum of the trades in the history.

Given a market maker defined by $\mathsf{ValTrades}$ and a history h, we let $\mathsf{ValHist}(h)$ denote the *valid histories extending* h, i.e. all valid histories with h as a prefix. Formally, $\mathsf{ValHist}(h)$ is the smallest set containing h such that, for all $h' \in \mathsf{ValHist}(h)$ and all $\mathbf{r} \in \mathsf{ValTrades}(h')$, we have $h' \oplus \mathbf{r} \in \mathsf{ValHist}(h)$. In particular, $\mathsf{ValTrades}(\epsilon)$ may be interpreted as the set of valid initial reserves for this market maker, and $\mathsf{ValHist}(\epsilon)$ the set of valid histories. However, we abuse notation slightly in that $\mathsf{ValHist}(\epsilon)$ is defined not to include ϵ itself, i.e. it only includes nonempty histories.

Examples and CFMMs. The simplest example of an automated market maker is when $\mathsf{ValTrades}(h)$ is a constant set that does not depend on h. Such a market maker would offer a fixed exchange rate between assets. However, we would like the exchange rates to adapt depending on demand. In the context of decentralized finance, a common approach is the following.

▶ **Definition 1** (CFMM). *A* constant-function market maker (CFMM) *is a market such that*

$$\mathsf{ValTrades}(h) = \{\mathbf{r} \in \mathbb{R}^n \mid \varphi(\mathbf{r} + \mathbf{q}_h) = \varphi(\mathbf{q}_h)\} \ ,$$

for some function $\varphi : \mathbb{R}^n \to \mathbb{R}$*, called the potential function.*

In a CFMM, the market maker accepts any trade keeping the potential of its reserves $\varphi(\mathbf{q}_h)$ constant. Some common examples are Uniswap or the constant product market where $\varphi(\mathbf{r}) = (r_1 \cdot r_2 \cdots r_n)^{1/n}$, and its generalization Balancer or constant geometric mean market which has $\varphi(\mathbf{r}) = r_1^{\alpha_1} \cdot r_2^{\alpha_2} \dots r_n^{\alpha_n}$ with $\sum_{i=1}^n \alpha_i = 1, \alpha_i \geq 0 \forall i$.

Non-examples. There are two natural classes of AMMs that are not (or are not quite) CFMMs. First, a market could charge a transaction fee. Depending on the implementation, the transaction fee will strictly grow the reserves over time, whereas a CFMM would always allow returning to the original reserves. Second, a market could post fixed prices for the assets, offering to sell up to e.g. one unit of each asset at a given price (in terms of the other assets). More generally, we will see that any market whose prices are history-dependent, in addition to depending on the current reserves, is not a CFMM. However, our results below will reveal a sense in which any "natural" AMM can likely be cast as a variant of a CFMM, as otherwise it fails to satisfy important axioms.

There are also natural markets that are not AMMs at all. An example are order books, where traders submit buy or sell requests over time and the marketplace periodically matches them to each other. In our AMM setting, all participants interact sequentially with the market and complete their transaction before the next participant arrives.

2.2 Axioms

An axiom is a potentially desirable property for a market maker as defined by $\mathsf{ValTrades}$. We begin with a few intuitive axioms.

▶ **Axiom 1** (NoDominatedTrades). *For all* $h \in \mathsf{ValHist}(\epsilon)$*,* $\mathsf{ValTrades}(h)$ *does not contain any* dominated trade*, i.e. an* $\mathbf{r}$ *where, for some* $\mathbf{r}' \in \mathsf{ValTrades}(h)$*, we have* $\mathbf{r}' \gneq \mathbf{r}$*.*

Assuming all assets have nonnegative value, a rational trader would never select a dominated trade. Conversely, as typically $\mathbf{0} \in \mathsf{ValTrades}(h)$, meaning a trader can chose not trade altogether, then under NoDominatedTrades no trader can take assets from the market maker without giving something in return. This latter property is the typical condition of *no-arbitrage* (within the system) in the finance literature.

▶ **Axiom 2** (PathIndependence)**.** *For all $h \in \mathsf{ValHist}(\epsilon)$, for all $\mathbf{r} \in \mathsf{ValTrades}(h)$ and $\mathbf{r}' \in \mathsf{ValTrades}(h \oplus \mathbf{r})$, we have that $\mathbf{r} + \mathbf{r}' \in \mathsf{ValTrades}(h)$.*

PathIndependence ensures that a trader cannot profit by splitting a single large trade into two smaller sequential trades instead. By backward induction, it implies that any *sequence* of trades can just as well be made in a single bundle: if $h' = h \oplus r_1 \oplus \cdots \oplus r_j$ is a valid history, we must have $\mathbf{r}_1 + \cdots + \mathbf{r}_j \in \mathsf{ValTrades}(h)$.

▶ **Observation 2.** *PathIndependence implies the following inductive version: for all $h \in \mathsf{ValHist}(\epsilon)$ and all $h' \in \mathsf{ValHist}(h)$, we have $\mathbf{q}_{h'} - \mathbf{q}_h \in \mathsf{ValTrades}(h)$.*

▶ **Axiom 3** (Liquidation)**.** *For all $h \in \mathsf{ValHist}(\epsilon)$ and all $\mathbf{r}, \mathbf{r}' \gneq \mathbf{0}$, there exists $\beta \geq 0$ such that $\mathbf{r} - \beta \cdot \mathbf{r}' \in \mathsf{ValTrades}(h)$.*

In other words, a participant can supply any nonnegative bundle $\mathbf{r}$ and specify a nonnegative demand bundle $\mathbf{r}'$, and receive some multiple of $\mathbf{r}'$ in return for $\mathbf{r}$. The Liquidation axiom captures the point of a "market maker", i.e., to offer to trade at some exhange rate between any pair of assets (or more generally, bundles).

Any market satisfying Liquidation and NoDominatedTrades must in particular offer the option to trade nothing; to see this set $\mathbf{r} = \mathbf{r}'$ for any nontrivial bundle $\mathbf{r}$.

▶ **Observation 3.** *Liquidation and NoDominatedTrades imply $\mathbf{0} \in \mathsf{ValTrades}(h)$ for all $h \in \mathsf{ValHist}(\epsilon)$. Furthermore, any $\mathbf{r} \in \mathsf{ValTrades}(h) \setminus \{\mathbf{0}\}$ can be written $\mathbf{r} = \mathbf{r}^+ - \mathbf{r}^-$ where $\mathbf{r}^+ = \max(\mathbf{r}, \mathbf{0})$ and $\mathbf{r}^- = \min(\mathbf{r}, \mathbf{0})$ satisfy $\mathbf{r}^+, \mathbf{r}^- \gneq \mathbf{0}$.*

▶ **Axiom 4** (DemandResponsiveness)**.** *For all $h \in \mathsf{ValHist}(\epsilon)$, if $\mathbf{r} - \mathbf{r}' \in \mathsf{ValTrades}(h)$ for some $\mathbf{r}, \mathbf{r}' \succeq \mathbf{0}$, and if there exist $\alpha, \beta > 0$ such that $\alpha\mathbf{r} - \beta\mathbf{r}' \in \mathsf{ValTrades}(h \oplus (\mathbf{r} - \mathbf{r}'))$, then $\beta \leq \alpha$.*

In other words, if a participant supplies $\mathbf{r}$ in return for $\mathbf{r}'$ once, then the "exchange rate" should increase the second time in a row: more of $\mathbf{r}$ is needed to buy a corresponding amount of $\mathbf{r}'$.

2.3 Characterization of fixed-liquidity CFMMs

In this subsection we characterize automated market makers that satisfy the axioms above: CFMMs with increasing, concave potential functions. In addition to setting the stage for the equivalence with prediction markets, this result places limits on the design space of automated market makers, or at least shows what axioms must be relinquished to move beyond the concave-CFMM framework. We summarize these results here, leaving complete details to Appendix A of the full version.

▶ **Proposition 4.** *Fix the initial reserves $\mathbf{q}_0$. A market maker* ValTrades *with $\mathsf{ValTrades}(\epsilon) = \{\mathbf{q}_0\}$ satisfies Liquidation, NoDominatedTrades, PathIndependence, and DemandResponsiveness if and only if it can be implemented as a CFMM with an increasing, concave potential function φ.*

If the market maker has multiple initial feasible reserves, then the market initiated from each one can be implemented as a CFMM with a concave potential. But it is not clear that there is a single concave potential function that implements all of these markets simultaneously. However, there at least exists a single *quasiconcave* potential function that implements the entire market as a CFMM simultaneously. This statement holds as long as we assume the slightly stronger axiom of StrongPathIndependence, which enforces a relationship between the market when starting at different initial reserves.

▶ **Axiom 5** (StrongPathIndependence)**.** *For all* $h, h' \in \mathsf{ValHist}(\epsilon)$ *with* $\mathsf{sum}(h) = \mathsf{sum}(h')$*, we have* $\mathsf{ValTrades}(h) = \mathsf{ValTrades}(h')$*. Furthermore, we have* $\mathbf{q}_h \in \mathsf{ValTrades}(\epsilon)$ *for all* $h \in \mathsf{ValHist}(\epsilon)$*.*

▶ **Theorem 5.** *A market maker, with* $\mathsf{ValTrades}(\epsilon) = \mathbb{R}^n$*, satisfies Liquidation, NoDominatedTrades, StrongPathIndependence, and DemandResponsiveness if and only if it is a CFMM with an increasing, continuous, quasiconcave potential function* φ*.*

As the potential functions in Proposition 4 are concave rather than continuous and quasiconcave, a weaker condition, one might ask whether this weaker condition is necessary in Theorem 5. We conjecture that the answer is yes, in the sense that there are CFMMs with continuous, quasiconcave potential functions which cannot be "concavified" by any monotone transformation; see e.g. [15, 13].

3 Equivalence of CFMMs and Prediction Markets

The goal of CFMMs is to facilitate trade. In contrast, a prediction market is designed to elicit and aggregate information about future events. We will now see that prediction markets (satisfying good elicitation axioms) are in a very strong sense equivalent to CFMMs (satisfying good trade axioms). We give direct reductions to convert each into the other. As a consequence, much existing research on design and properties of prediction markets can now transfer to CFMMs.

We will briefly recall the basics of cost-function prediction markets and summarize our results. In § 3.1, we define cost function prediction markets and recall that they characterize automated market makers satisfying elicitation axioms. Here we take the key step of converting them into "cashless prediction markets". In § 3.2, we give the reductions and equivalence between prediction markets and CFMMs. § 3.3 discusses information-elicitation consequences, and § 3.4 gives examples. Full details appear in Appendix A.

3.1 Prediction markets

An Arrow-Debreu (AD) cost-function prediction market for an event on n outcomes (e.g. to predict the winner of an election with n candidates) is a special case of an automated market maker as defined in Section 2. The $n + 1$ assets in the market consist of "cash", the numeraire that is used to buy and sell, along with n *securities*. Each security is an asset, and the owner of c units of security $i \in \{1, \ldots, n\}$ will receive c units of cash if outcome i occurs, and will receive 0 if i does not occur. Here c can be any positive or negative real number; holding a negative quantity of security i corresponds to having sold that number of units to the market maker.

The market is defined by a *cost function* $C : \mathbb{R}^n \to \mathbb{R}$, which is required to be convex, increasing, and $\mathbb{1}$-invariant. The initial amount of securities sold is denoted $\mathbf{q}_0 = \mathbf{0} \in \mathbb{R}^n$. At each time t, trader t arrives and requests any bundle of securities $\mathbf{r}_t \in \mathbb{R}^n$. The market updates to $\mathbf{q}_t = \mathbf{q}_{t-1} + \mathbf{r}_t$ and gives the trader $\mathbf{r}_t$. The trader pays the following amount of cash to the market:

$$\text{payment} = C(\mathbf{q}_{t-1} + \mathbf{r}_t) - C(\mathbf{q}_{t-1}) = C(\mathbf{q}_t) - C(\mathbf{q}_{t-1}).$$

Thus, the market maintains the invariant that its total amount of cash equals $C(\mathbf{q}_t) - C(\mathbf{0})$. We note that **by convention, $\mathbf{q}_t$ denotes the amount of assets sold by the market maker, rather than its reserves.** So a negative sign is needed to convert $\mathbf{q}_t$ in the prediction market setting to its meaning in the CFMM setting.

The instantaneous prices of the n securities are given by $\nabla C(\mathbf{q}_t)$ and correspond to the market prediction; e.g. the price for a unit of security i can be interpreted as the predicted probability of outcome i. More details are given in Appendix A.

All truthful automated markets are cost-function prediction markets. Prior work has shown that automated market makers, if they satisfy basic axioms of path independence and truthfully eliciting predictions, *must* be cost-function prediction markets. A key step is a classic equivalence of cost functions to *proper scoring rules*; details appear in Appendix A. This fact lets us draw a conceptual chain: truthful $\iff$ prediction market $\leftrightarrow$ CFMM $\iff$ facilitating trade.

▶ **Fact 1.** *Let M be any automated market maker offering outcome-contingent securities and cash. If M satisfies (1) Path Independence and (2) Incentive Compatibility, then M is a cost-function AD prediction market.*

Cashless prediction markets. To translate to CFMMs, a key observation is that it is possible to run a cost-function prediction market without any cash transactions until the very final payouts. The point is to consider the "grand bundle" of securities, $\mathbb{1} = (1, \ldots, 1) \in \mathbb{R}^n$. This bundle of assets is equivalent to one unit of cash for any agent, because exactly one outcome will occur, exactly one of the securities will pay out exactly one unit of cash, and the others will become worthless.

Cast as an automated market maker framework, define for any cost function C a "cashless" prediction market on n assets that charges multiples of $\mathbb{1}$ instead of cash:

$$\mathsf{ValTrades}'_C(h) = \{\mathbf{r} + \alpha\mathbb{1} \mid \mathbf{r} \in \mathbb{R}^n\} \text{ where } \alpha = C(-\mathbf{q}_h - \mathbf{r}) - C(-\mathbf{q}_h). \quad (2)$$

A key fact is the following lemma, which states that a $\mathbb{1}$-invariant C maintains a *constant* function value for all legal cashless trades.

▶ **Lemma 6.** $\mathsf{ValTrades}'_C(h) = \{\mathbf{r} \mid C(-\mathbf{q}_h - \mathbf{r}) = C(-\mathbf{q}_h)\}$.

3.2 Equivalence to CFMMs

Using the cashless formulation of a prediction market, we can "convert" it to a CFMM and vice versa. What we mean by convert is that a cost function C, defining a prediction market on any event with n outcomes, is converted to a potential function φ, defining a CFMM on any set of n assets. The "rules of trade" imposed by C are converted to "rules of trade" imposed by φ. Even though the setting may be very different – e.g., instead of eliciting predictions about an election, we are facilitating trade among an unrelated set of tokens – the resulting φ will satisfy good market-making axioms. The proofs rely on convex analysis tools developed in the context of *financial risk measures* and appear in Appendix B.3 of full version.

▶ **Theorem 7.** *Let C define an AD cost-function prediction market. Then the function $\varphi : \mathbb{R}^n \to \mathbb{R}$ defined by*

$$\varphi(\mathbf{q}) = -C(-\mathbf{q})$$

is concave and increasing, i.e. defines a CFMM. Furthermore, $\mathsf{ValTrades}_\varphi = \mathsf{ValTrades}'_C$.

▶ **Theorem 8.** *Let $\varphi : \mathbb{R}^n \to \mathbb{R}$ be concave and increasing, defining a CFMM. For any $\mathbf{q}_0 \in \mathbb{R}^n$, the function $C : \mathbb{R}^n \to \mathbb{R}$ defined by*

$$C(\mathbf{q}) = \inf \{c \in \mathbb{R} \mid \varphi(c\mathbb{1} - \mathbf{q}) \geq \varphi(\mathbf{q}_0)\} \tag{3}$$

is convex, increasing, and ones-invariant, i.e., defines an AD cost-function prediction market. Furthermore, for all $h \in \mathsf{ValHist}_\varphi((\mathbf{q}_0))$, we have $\mathsf{ValTrades}'_C(h) = \mathsf{ValTrades}_\varphi(h)$.

We note that the correspondence between potential functions φ and cost functions C is not one-to-one. For example, $\varphi(\mathbf{q})$ and $\varphi(\mathbf{q}) + 1$ both map to the same cost function C. This is related to the fact that a given market (i.e. a given ValTrades) that can be represented as a CFMM does not have a *unique* representation as a CFMM (and similarly for cost functions). For example, $\alpha\varphi$ and φ define the same CFMM (i.e. same ValTrades function) for any nonzero α.

3.3 CFMMs and scoring rules

Recall that a scoring rule is a function $S(\mathbf{p}, i)$ that evaluates the prediction $\mathbf{p} \in \Delta_\mathcal{Y}$ given an observed outcome i. A scoring rule is *proper* if for all probability distributions $\mathbf{p}$, the expected score when $i \sim \mathbf{p}$ is maximized by report $\mathbf{p}$. It is *strictly* proper if $\mathbf{p}$ is the unique maximizer. It is well-known (see [20]) that a scoring rule is proper if and only if it can be written as an affine approximation to a particular convex *generating function* $G(\mathbf{p})$.

It is well-known that cost function market makers are equivalent to scoring rule market makers; see § A.1. In light of our results, then, it follows that CFMMs are are closely linked to proper scoring rules. This connection emphasizes the surprising nature of the equivalence between cost functions and CFMMs, namely that CFMMs, while designed for a very different context than probabilistic forecasting, *characterize* the properties needed for eliciting forecasts.

▶ **Corollary 9.** *Every CFMM for a concave, increasing φ and initial reserves $\mathbf{q}_0$ is associated with a proper scoring rule.*

To briefly sketch this connection, every cost function C corresponds to the scoring rule S_G generated by $G = C^*$, the convex conjugate of C, and vice versa. In fact, we can say more: given $\mathbf{q}_0$ and a strictly concave, increasing φ, if C is the cost function from Theorem 8 and $G = C^*$, then for all $h \in \mathsf{ValHist}_\varphi((\mathbf{q}_0))$ we may write

$$\mathsf{ValTrades}_\varphi(h) = \{S_G(\mathbf{p}, \cdot) - S_G(\mathbf{p}', \cdot) \in \mathbb{R}^n \mid \mathbf{p}' \in \Delta_\mathcal{Y}\} \ , \tag{4}$$

where $\mathbf{p} \in \partial C(\mathbf{q}_h)$ is a current price. In other words, the set of valid trades is exactly the set of scoring rule difference vectors. Thus, trades implicitly choose a new price vector $\mathbf{p}'$, and in fact we could change the interface to the market entirely so that traders submit a distribution $\mathbf{p}_t$ and the corresponding trade is the (negative) market scoring rule $S_G(\mathbf{p}_{t-1}, \cdot) - S_G(\mathbf{p}_t, \cdot) \in \mathbb{R}^n$. See § 5 for the potential benefits of this price-oriented interface.

As alluded to above, for strict properness, one needs additional conditions on C, namely that (1) C is differentiable, so that there is a unique market price at any given state, and (2) the closure of the set of gradients is equal to the probability simplex, so that any given market price/prediction is achievable by some set of trades. These carry over into corresponding conditions on φ, as we explore in Appendix A.2.

3.4 Examples

We now present two examples to illustrate our equivalence thus far. We discuss these examples again in § 4.1 along with others.

LMSR → CFMM. As a warm-up, let us see how the most popular cost function, the LMSR $C(\mathbf{q}) = b \log \sum_{i=1}^n e^{q_i/b}$, can be interpreted as a CFMM. Letting $\varphi(\mathbf{q}) = C(-\mathbf{q})$, we have $\mathbf{r} \in \mathsf{ValTrades}(h) \iff \varphi(\mathbf{q}+\mathbf{r}) = \varphi(\mathbf{q}_0)$, where $\mathbf{q}_0$ is the initial reserves. If $k = \varphi(\mathbf{q}_0)$, then $\mathbf{r}$ is a valid trade if and only if $b \log \sum_{i=1}^n e^{-(q_i+r_i)/b} = k \iff \sum_{i=1}^n e^{-(q_i+r_i)/b} = e^{k/b}$. For two assets, this equation reduces to $e^{-(q_1+r_1)/b} + e^{-(q_1+r_1)/b} = e^{k/b}$ (cf. [32]). See § 4.1 for a version which scales the liquidity b depending on $\mathbf{q}_0$.

Uniswap → cost function / scoring rule. One of the most iconic CFMMs is Uniswap, with the potential function $\varphi_U : \mathbb{R}^2 \to \mathbb{R}$ given by $\varphi_U(\mathbf{q}) = \sqrt{q_1 q_2}$. Interestingly, φ_U is not defined on all of $\mathbb{R}^n$, and is only increasing on $\mathbb{R}^n_{>0}$; as such, typically one restricts to the latter space. As we will see in § 4, this restriction to $\mathbb{R}^n_{>0}$ is typical for CFMMs used in practice, and does not pose a barrier for our equivalence. Briefly, as long as we require $\mathbf{q}_0 \succ \mathbf{0}$, the market reserves will stay within $\mathbb{R}^n_{>0}$, and the characterization from Theorem 5 and equivalence from Theorem 8 will still apply.

Let $\mathbf{q}_0 \succ \mathbf{0}$ and set $k = \varphi(\mathbf{q}_0) > 0$. Applying Theorem 8, we obtain the cost function from eq. (1),

$$C_k(\mathbf{q}) = \frac{1}{2}\left(q_1 + q_2 + \sqrt{4k^2 + (q_1 - q_2)^2}\right) .$$

Let us verify the construction. One can easily check that C_k is $\mathbb{1}$-invariant and increasing. Consequently, it suffices to verify a single level set of C_k; we will show $\{\mathbf{q} \mid C(-\mathbf{q}) = 0\} = \{\mathbf{q} \mid \mathbf{q} \succ \mathbf{0}, \varphi_U(\mathbf{q}) = k\}$.

$$\begin{aligned} C_k(-\mathbf{q}) = 0 &\iff \sqrt{4k^2 + (q_1 - q_2)^2} = (q_1 + q_2) \\ &\iff 4k^2 + (q_1 - q_2)^2 = (q_1 + q_2)^2 \text{ and } q_1, q_2 > 0 \\ &\iff 4k^2 = 4q_1 q_2 \text{ and } q_1, q_2 > 0 \\ &\iff k = \sqrt{q_1 q_2} \text{ and } q_1, q_2 > 0 . \end{aligned}$$

We conclude that, indeed, $\mathsf{ValTrades}'_{C_k} = \mathsf{ValTrades}_{\varphi_U}$, as promised by Theorem 8. In other words, a prediction market run using C_k would behave exactly the same as Uniswap.

As discussed in § 1, the expression for C_k appears in [11, eq. (14)]. There they show that a market maker keeping constant log utility assuming the binary outcome would be drawn from the uniform distribution $\pi = (1/2, 1/2)$. Moreover, as observed by [8] and [30], constant expected log utility under the uniform distribution gives rise to Uniswap: $(1/2)\log(q_1) + (1/2)\log(q_1) = c \iff \sqrt{q_1 q_2} = e^c$. (Changing this distribution π gives rise to more general forms of Uniswap, such as Balancer, which takes the form $\varphi_\pi(\mathbf{q}) = q_1^{\pi_1} q_2^{\pi_2}$.) Amazingly, it appears that no one has yet made these two connections together, that Uniswap is equivalent to a cost-function prediction market.

In light of Corollary 9, each level set of Uniswap must therefore be equivalent to a scoring rule market for some choice of proper scoring rule. Let us now derive this family of scoring rules. Taking the convex conjugate of C_k, we obtain the convex generating function

$$G_k(\mathbf{p}) = -2k\sqrt{p_1 p_2}\ .$$

The corresponding scoring rule is

$$S_k(\mathbf{p}, i) = G(\mathbf{p}) + dG_{\mathbf{p}}(\boldsymbol{\delta}_i - \mathbf{p}) = -k\sqrt{p_i/p_j}\ ,$$

where $\{i, j\} = \{1, 2\}$. This scoring rule is exactly the boosting loss of [9], negated and scaled by k. (This scoring rule also appears in a shifted form as "hs" in [7].) One can check that indeed $\mathsf{ValTrades}_{\varphi_U}(h) = \{S_k(\mathbf{p}, \cdot) - S_k(\mathbf{p}', \cdot) \mid \mathbf{p}' \in \Delta_{\{1,2\}}\}$ for $\mathbf{p} = \nabla C_k(\mathbf{q}_h)$. Interestingly, while in general each level set of φ_U could have corresponded to an entirely different scoring rule, instead each level set merely scales the same scoring rule. As we will see in § A.3, this phenomenon is shared by all 1-homogeneous CFMMs.

4 Bounded reserves and liquidity

CFMMs are defined in terms of the current reserves $\mathbf{q} \in \mathbb{R}^n$. Naturally, one might like to ensure $\mathbf{q} \succeq \mathbf{0}$ at every state of the market, i.e., that the market maker cannot go short on any asset.

▶ **Axiom 6** (BoundedReserves). *$\mathbb{R}^n_{>0} \subseteq \mathsf{ValTrades}(\epsilon)$ and for all $h \in \mathsf{ValHist}(\epsilon)$, $\mathbf{q}_h \succeq \mathbf{0}$.*

In turns out that CFMMs constructed from cost functions in our previous equivalence, Theorem 7, can essentially never satisfy BoundedReserves (Appendix C.2 in full version). Using the perspective transform from convex analysis, however, which is well-known in the prediction market literature to capture liquidity levels, we provide a new construction for this purpose. The magic is in "wrapping up" the level sets of $C_\alpha(\mathbf{q}) := \alpha C(\mathbf{q}/\alpha)$ for different α into a single function φ, so that starting at $\mathbf{q}_0$ implies a liquidity of $\alpha = \varphi(\mathbf{q}_0)$. Furthermore, we ensure that $\varphi(\mathbf{q}) \to 0$ as $\mathbf{q} \to \mathbf{0}$, so that a market starting with very low initial reserves quickly moves the price so as not to run out.

▶ **Construction 1.** *Let $C : \mathbb{R}^n \to \mathbb{R}$ be a convex, increasing, $\mathbb{1}$-invariant cost function satisfying $C(\mathbf{0}) > 0$.[2] Then define $\varphi : \mathbb{R}^n_{>0} \to \mathbb{R}$ by $\varphi(\mathbf{q}) = \alpha$ where $C(-\mathbf{q}/\alpha) = 0$.*

We next show that this construction produces market makers satisfying our five major axioms. Let us see why it is behaving as we would hope. By definition, the α-level set of φ matches the 0-level set of C_α, i.e., $\{\mathbf{q} \mid \varphi(\mathbf{q}) = \alpha\} = \{\mathbf{q} \mid C_\alpha(-\mathbf{q}) = 0\}$; see Figure 2(L). If $\varphi(\mathbf{q}_0) = \alpha$ for $\alpha > 0$, then we have $\mathsf{ValTrades}_\varphi(h) = \mathsf{ValTrades}_{C_\alpha}(h)$. In other words, φ indeed has liquidity level α when $\varphi(\mathbf{q}_0) = \alpha$. Moreover, if $\mathbf{q}_0 = \beta\mathbf{q}$, then $\varphi(\mathbf{q}_0) = \beta\varphi(\mathbf{q})$, so indeed the liquidity level drops to 0 as $\mathbf{q}_0 \to \mathbf{0}$, thereby protecting the reserves.

▶ **Proposition 10.** *Let C satisfy the conditions of Construction 1 as well as $\{\mathbf{q} \mid C(\mathbf{q}) = 0\} \subseteq \mathbb{R}^n_{<0}$. Then the resulting $\varphi : \mathbb{R}^n_{>0} \to \mathbb{R}$ defines a CFMM satisfying Liquidation, NoDominatedTrades, PathIndependence, DemandResponsiveness, and BoundedReserves.*

The scaling property noted above, known as *1-homogeneity*, is crucial to our analysis (§ A.3). As essentially all CFMMs used in practice are 1-homogeneous [5], it is natural to ask whether Construction 1 could have produced them. We now show that the answer is yes.

[2] The assumption $C(\mathbf{0}) > 0$ is without loss of generality; otherwise, add a sufficiently large constant to C.

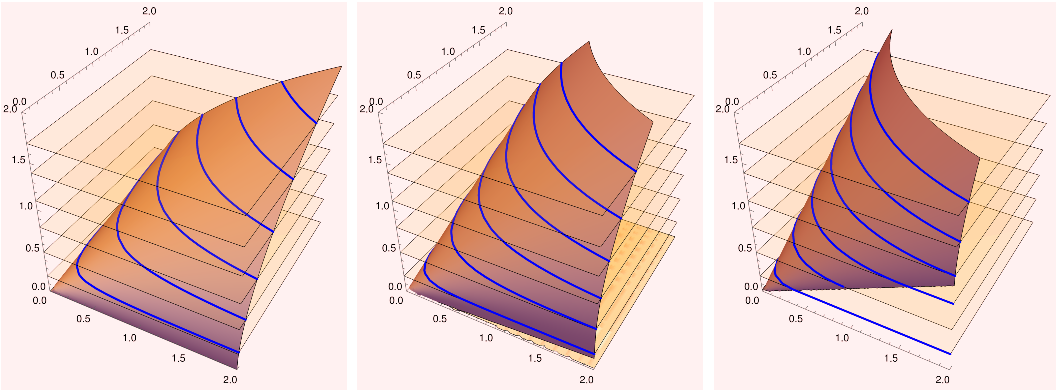

Figure 2 Plots of Construction 1 for Uniswap (L), LMSR (M), and Brier score (R); see § 4.1 for how the construction was applied in each setting. In each plot, the orange surface plots the 0-level set of the perspective transform $T : (q_1, q_2, \alpha) \mapsto \alpha C(-(q_1, q_2)/\alpha)$, i.e., the set of triples (q_1, q_2, α) such that $\alpha C(-(q_1, q_2)/\alpha) = 0$. By design, this set also happens to be the graph of φ resulting from Construction 1, i.e., the set of triples (q_1, q_2, α) such that $\varphi((q_1, q_2)) = \alpha$. To emphasize this connection, the α-level sets φ for $\alpha = 0.2, 0.6, 1.0, 1.4, 1.8$ are shown in blue. (Technically, these level sets should be projected down to the plane on the bottom of the figure.) As T is convex (a fact of perspective transforms), the sublevel set $T(q_1, q_2, \alpha) \leq 0$, the region below the surface, is a convex set. We can see therefore that φ is always a concave function.

▶ **Proposition 11.** *A potential $\varphi : \mathbb{R}^n_{>0}$ is the result of applying Construction 1 if and only if φ is 1-homogeneous, increasing, and concave.*

As a result of Proposition 11, through Corollary 9, every 1-homogeneous CFMM can be associated with a unique proper scoring rule; the various level sets of the potential merely correspond to scaling up the underlying scoring rule. As we will see in § 4.1, this fact can be used to construct useful CFMMs from the vast literature on proper scoring rules. All proofs can be found in Appendix C in full version, as well as several additional results; chief among them is a full characterization of market makers satisfying BoundedReserves in addition to our other four axioms (Theorem 25 in full version).

4.1 Examples

Uniswap. Let us perform Construction 1 on the cost function C_k we derived from Uniswap, $\varphi_U(\mathbf{q}) = \sqrt{q_1 q_2}$, given some level set $k > 0$:

$$C_k(\mathbf{q}) = \frac{1}{2}\left(q_1 + q_2 + \sqrt{4k^2 + (q_1 - q_2)^2}\right) .$$

First, let us check that C_k satisfies the condition of Construction 1: indeed, $C_k(\mathbf{0}) = k > 0$. From § 3.4, we have $C_k(-\mathbf{q}) = 0 \iff \varphi_U(\mathbf{q}) = \sqrt{q_1 q_2} = k$. Thus $C(-\mathbf{q}/\alpha) = 0 \iff \sqrt{q_1 q_2} = \alpha k$. Thus, the construction yields $\varphi_k(\mathbf{q}) = \sqrt{q_1 q_2}/k$, which is exactly φ_U when $k = 1$, and a scaled version for other values.

LMSR. Let $C(\mathbf{q}) = b \log(\sum_{i=1}^n e^{q_i/b})$ be LMSR. Unlike our choice $\varphi(\mathbf{q}) = -C(-\mathbf{q})$ from § 3.4, let us now consider a "reserves-aware" version, which will automatically scale liquidity with $\mathbf{q}_0$. As $C(\mathbf{0}) = b \log n > 0$, Construction 1 applies, giving a potential φ. See Figure 2(M) for a visualization. Moreover, the condition of Proposition 10 is also satisfied, as $C(\mathbf{q}) = 0 \implies \sum_{i=1}^n e^{q_i/b} = 1 \implies q_i < 0$ for all i. So the CFMM with potential φ satisfies all five of our axioms. Unfortunately, unlike Uniswap, this φ does not permit a closed-form expression. Nonetheless, as we discuss in § 5, we believe it could be practical to deploy in real markets.

CFMM from Brier score. Let us see how to take one of the most popular scoring rules, the Brier (or quadratic) score,

$$S(\mathbf{p}, i) = 2\mathbf{p} \cdot \boldsymbol{\delta}_i - \|\mathbf{p}\|_2^2 \ , \tag{5}$$

and convert it to a CFMM using our chain of equivalences. The corresponding convex generating function is $G(\mathbf{p}) = \|\mathbf{p}\|^2$. The associated cost function C, the convex conjugate of G, does not have a convenient closed form in general; for $n = 2$ assets it can be written

$$C(\mathbf{q}) = \begin{cases} q_1 & q_1 - q_2 \geq 2 \\ q_2 & q_1 - q_2 \leq -2 \\ \frac{1}{8}\left((q_1 - q_2)^2 + 4(1 + q_1 + q_2)\right) & \text{otherwise} \end{cases} . \tag{6}$$

From this form, one can see that C is not always increasing, as e.g. the price for asset 2 becomes 0 when $q_1 - q_2 \geq 2$. The resulting market therefore does not satisfy NoDominatedTrades, as selling asset 2 for a price of 0 is dominated by doing nothing.

It turns out, however, that we may still apply Construction 1 to obtain a sensible CFMM; in this case, for $n = 2$, we obtain

$$\varphi(\mathbf{q}) = \sqrt{q_1 q_2} + \frac{q_1 + q_2}{2} \ , \tag{7}$$

the sum of the geometric and arithmetic means; see Figure 2(R). In fact, this potential function is a hybrid CFMM appearing already in the literature, e.g. [5, § 2.4].

Curve. Finally, let us consider Curve [14], given by $\varphi(\mathbf{q}) = \sum_i q_i + \sum_i (1/q_i)$. As noted by [5], this potential is not 1-homogeneous, and thus could not be produced by Construction 1. As a result, its level sets are not merely scaled copies of each other, but change as the value of φ grows. (For instance, its 0-level set is the same as Uniswap with $k = 1$, but no other level set is in common with Uniswap.) Nonetheless, we may still apply Theorem 25(see full version) to see that φ satisfies all five of our axioms; to check condition (c), note that $\overline{\varphi} = -\infty$ on $\partial\mathbb{R}^n_{\geq 0}$.

5 Adaptive Liquidity and Other Future Directions

The literature on automated market makers for decentralized exchanges has largely proceeded somewhat removed from the expansive literature on prediction markets. We believe the results presented here will allow for these independent lines of research to merge and inform each other. To conclude, we briefly discuss several avenues for future work, with a focus on liquidity adaptation.

Transaction fees. In practice, CFMMs often allow liquidity to change over time, in two ways: (1) they may charge a “transaction fee”, wherein traders must contribute directly to the reserves in addition to their trade, and (2) they may allow liquidity providers to contribute to the reserves in exchange for a dividend. Let us first discuss the transaction fee.

Here the market designer chooses a parameter $\gamma \in (0,1]$, where lower γ corresponds to a higher fee. The market maker then accepts any trade keeping the potential function φ constant after discounting the bundle given to the market maker by γ. Formally, we can write

$$\mathsf{ValTrades}_{\varphi,\gamma}(h) = \{\mathbf{r} \in \mathbb{R}^n \mid \varphi(\mathsf{sum}(h) + \gamma\mathbf{r}_+ - \mathbf{r}_-) = \varphi(\mathsf{sum}(h))\} \ , \tag{8}$$

where $\mathbf{r}_+ = \max(\mathbf{r}, \mathbf{0})$ and $\mathbf{r}_- = \min(\mathbf{r}, \mathbf{0})$. To state this set of valid trades more naturally in terms of transaction fees, relative to the no-fee ValTrades for the vanilla CFMM, we may write

$$\mathsf{ValTrades}_{\varphi,\gamma}(h) = \{\mathbf{r} + \mathsf{fee}(\mathbf{r}) \mid \mathbf{r} \in \mathsf{ValTrades}_{\varphi}(h)\} \ , \tag{9}$$

where $\mathsf{fee}(\mathbf{r}) = \beta\mathbf{r}_+$ for $\beta = (1-\gamma)/\gamma > 0$. That is, a trader may choose any trade $\mathbf{r}$ keeping φ constant, but then must also add $\beta\mathbf{r}_+$ to the reserves.

Now suppose φ is increasing, concave, and 1-homogeneous, as is commonly the case, and as guaranteed by Construction 1. Then as $\varphi(\mathbf{q}_h) = \varphi(\mathbf{q}_h + \mathbf{r})$, and φ is increasing, we will have $\varphi(\mathbf{q}_h + \mathbf{r} + \mathsf{fee}(\mathbf{r})) > \varphi(\mathbf{q}_h + \mathbf{r}) = \varphi(\mathbf{q}_h)$. As we saw in § 4, $\varphi(\mathbf{q})$ represents the liquidity level when the reserves are $\mathbf{q}$, so we conclude that the liquidity increases after each trade. Moreover, as the fee-less CFMM satisfies BoundedReserves, and φ is increasing, it is easy to verify that BoundedReserves will still be satisfied with the transaction fee. Thus, the transaction fee successfully subsidizes the liquidity increase of the market without risking depleting the reserves.

Implicitly defined potential functions. Recall that the φ from Construction 1 is only implicitly defined, as the solution $\varphi(\alpha)$ to $C(-\mathbf{q}/\alpha) = 0$ for the given cost function C. Thus, even if C is given explicitly, φ may not have a closed form. While sometimes one can solve for φ explicitly, as we saw in § 4.1, one cannot for the 1-homogeneous CFMM potential φ we derived from LMSR. The lack of a closed form poses a challenge, as the transaction fee causes φ to change, and thus the value of $\varphi(\mathbf{q}_h)$ would need to be recalculated rather than being fixed ahead of time.

We now propose a straightforward workaround using the fact that φ is still implicitly defined by a known cost function C. Suppose the current value $\alpha = \varphi(\mathbf{q}_h)$ is publicly known, and a trader wishes to purchase $\mathbf{r} \in \mathsf{ValTrades}_{\varphi}(\mathbf{q}_h)$, i.e., such that $\varphi(\mathbf{q}_h + \mathbf{r}) = \alpha$. We will ask the trader to announce $\mathbf{r}$, as well as the value $\alpha' = \varphi(\mathbf{q}_h + \mathbf{r} + \mathsf{fee}(\mathbf{r}))$, up to some error tolerance. As $C(-\mathbf{q}/\alpha)$ is monotone in α, the trader can easily compute α' within the desired accuracy given an expression for C. Moreover, the two relevant conditions can be checked on-chain: $\varphi(\mathbf{q}_h + \mathbf{r}) = \alpha$ by $C(-(\mathbf{q}_h + \mathbf{r})/\alpha) = 0$, and $\varphi(\mathbf{q}_h + \mathbf{r} + \mathsf{fee}(\mathbf{r})) = \alpha'$ by $C(-(\mathbf{q}_h + \mathbf{r} + \mathsf{fee}(\mathbf{r}))/\alpha') = 0$.

To illustrate, let C be LMSR, and φ the result of Construction 1. From the condition $C(-\mathbf{q}/\alpha) = 0$, the level set $\varphi(\mathbf{q}) = \alpha$ is $\{\mathbf{q} \in \mathbb{R}^n_{>0} \mid \sum_{i=1}^n e^{-q_i/\alpha} = 1\}$. The validity of $\mathbf{r}$ and α' can be checked, via $\sum_{i=1}^n \exp(-(q_i + r_i)/\alpha) \approx 1$ and $\sum_{i=1}^n \exp(-(q_i + r_i + \mathsf{fee}(\mathbf{r})_i)/\alpha') \approx 1$. Thus, by asking a trader to compute a valid trade, as well as approximating the next value of φ, the market can proceed even without a closed form for φ.

Alternate interface as a scoring rule market. The typical interface of an automated market maker in decentralized finance has traders submit a desired bundle $\mathbf{r}$, and the check $\mathbf{r} \in \mathsf{ValTrades}(h)$ is done on-chain. To keep gas fees low, the market maker is designed to make this check straightforward, as in the constant-product check $(q_1 + r_1)(q_2 + r_2) = k$ for Uniswap. As we saw in the previous point, however, sometimes changing the interface slightly can ease the computation needed on-chain.

The connection to scoring rule markets suggests an entirely different interface which could similarly ease computation costs, while opening up the door to many more possible invariant curves. The idea is simple: choose a proper scoring rule S, track the current normalized prices $\mathbf{p} \in \Delta_\mathcal{Y}$, and ask traders to report the desired next price vector $\mathbf{p}' \in \Delta_\mathcal{Y}$. The resulting implied trade $S(\mathbf{p}, \cdot) - S(\mathbf{p}', \cdot) \in \mathbb{R}^n$ is automatically an element of $\mathsf{ValTrades}(h)$ by our previous observations; see eq. (4). Yet now there is no check needed on-chain.

To illustrate the power of this approach, consider the 1-homogenous version of LMSR above. It is well-known that LMSR with liquidity parameter $b > 0$ is equivalent to the log scoring rule, $S(\mathbf{p}, i) = b \log p_i$. Thus, any valid trade $\mathbf{r}$ must take the form $r_i = b \log p_i - b \log p'_i = b \log p_i/p'_i$. We can check that indeed, trades of this form satisfy the desired invariant, as we will always have $\mathbf{q} = -S(\mathbf{p}, \cdot)$ for some $\mathbf{p} \in \Delta_\mathcal{Y}$, and thus $C(-\mathbf{q}) = b \log \sum_i e^{-q_i/b} = b \log \sum_i p_i = 0$.

Other forms of liquidity adaptation. The prediction market literature offers many other forms of transaction fees and schemes for liquidity adaptation; see e.g. [31, 28, 25]. Using the equivalence developed in this work, these market makers can be readily transformed into those for decentralized exchanges. It would be particularly interesting to study the CFMM transaction fee within the volume-parameterized market (VPM) framework of [3]; we conjecture that it fails to satisfy their *shrinking spread* axiom (for any choice of volume function) but satisfies their others. Conversely, the perspective market from that same paper may be of interest to the decentralized finance community.

Other directions. As mentioned above, many CFMMs allow liquidity providers to directly contribute to the reserves. In light of our results, it would be interesting and potentially impactful to study the elicitation implications of these protocols from the perspective of liquidity providers. Finally, while the above focuses on potential contributions from the prediction market literature to decentralized finance, contributions in this direction are likely to be fruitful as well.

References

1 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. URL: http://dl.acm.org/citation.cfm?id=2465777.

2 Jacob D. Abernethy, Rafael M. Frongillo, and Sindhu Kutty. On Risk Measures, Market Making, and Exponential Families. *SIGecom Exchanges*, 13(2):21–25, January 2015. doi:10.1145/2728732.2728734.

3 Jacob D. Abernethy, Rafael M. Frongillo, Xiaolong Li, and Jennifer Wortman Vaughan. A General Volume-parameterized Market Making Framework. In *Proceedings of the Fifteenth ACM Conference on Economics and Computation*, EC '14, pages 413–430, New York, NY, USA, 2014. ACM. doi:10.1145/2600057.2602900.

4 Shipra Agrawal, Erick Delage, Mark Peters, Zizhuo Wang, and Yinyu Ye. A unified framework for dynamic prediction market design. *Operations research*, 59(3):550–568, 2011. URL: http://pubsonline.informs.org/doi/abs/10.1287/opre.1110.0922.

5 Guillermo Angeris, Akshay Agrawal, Alex Evans, Tarun Chitra, and Stephen Boyd. *Constant Function Market Makers: Multi-asset Trades via Convex Optimization*, pages 415–444. Springer International Publishing, Cham, 2022. doi:10.1007/978-3-031-07535-3_13.

6 Guillermo Angeris and Tarun Chitra. Improved price oracles: Constant function market makers. In *Proceedings of the 2nd ACM Conference on Advances in Financial Technologies*, pages 80–91, 2020.

7 Shalev Ben-David and Eric Blais. A new minimax theorem for randomized algorithms. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 403–411. IEEE, 2020.

8 Maxim Bichuch and Zachary Feinstein. Axioms for automated market makers: A mathematical framework in fintech and decentralized finance, 2022. doi:10.48550/ARXIV.2210.01227.

9 Andreas Buja, Werner Stuetzle, and Yi Shen. Loss functions for binary class probability estimation and classification: Structure and applications. *Working draft, November*, 3:13, 2005. URL: https://pdfs.semanticscholar.org/d670/6b6e626c15680688b0774419662f2341caee.pdf.

10 Agostino Capponi and Ruizhe Jia. The adoption of blockchain-based decentralized exchanges. *arXiv preprint arXiv:2103.08842*, 2021.

11 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.

12 Y. Chen and J.W. Vaughan. A new understanding of prediction markets via no-regret learning. In *Proceedings of the 11th ACM conference on Electronic commerce*, pages 189–198, 2010.

13 Chris Connell and Eric B Rasmusen. Concavifying the quasiconcave. *Journal of Convex Analysis*, 24(4):1239–1262, 2017.

14 Michael Egorov. Stableswap-efficient mechanism for stablecoin liquidity. *Retrieved Feb*, 24:2021, 2019.

15 W. Fenchel. *Convex Cones, Sets, and Functions*. Princeton University, Department of Mathematics, Logistics Research Project, 1953. URL: https://books.google.com/books?id=qOQ-AAAAIAAJ.

16 Rafael Frongillo and Mark D. Reid. Convergence Analysis of Prediction Markets via Randomized Subspace Descent. In *Advances in Neural Information Processing Systems*, pages 3016–3024, 2015. URL: http://papers.nips.cc/paper/5727-convergence-analysis-of-prediction-markets-via-randomized-subspace-descent.

17 Rafael Frongillo and Bo Waggoner. An Axiomatic Study of Scoring Rule Markets. In Anna R. Karlin, editor, *9th Innovations in Theoretical Computer Science Conference (ITCS 2018)*, volume 94 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 15:1–15:20, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. doi:10.4230/LIPIcs.ITCS.2018.15.

18 Hans Föllmer and Alexander Schied. Convex and coherent risk measures. *October*, 8:2008, 2008. URL: http://alexschied.de/Encyclopedia6.pdf.

19 Hans Föllmer and Stefan Weber. The Axiomatic Approach to Risk Measures for Capital Determination. *Annual Review of Financial Economics*, 7(1), 2015.

20 Tilmann Gneiting and Adrian E. Raftery. Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378, 2007.

21 Gnosis. Automated market makers for prediction markets, 2020. URL: https://docs.gnosis.io/conditionaltokens/docs/introduction3/.

22 Mohak Goyal, Geoffrey Ramseyer, Ashish Goel, and David Mazieres. Finding the right curve: Optimal design of constant function market makers. In *Proceedings of the 24th ACM Conference on Economics and Computation*, EC '23, pages 783–812, New York, NY, USA, 2023. Association for Computing Machinery. doi:10.1145/3580507.3597688.

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

24 Johannes Rude Jensen, Victor von Wachter, and Omri Ross. An introduction to decentralized finance (defi). *Complex Systems Informatics and Modeling Quarterly*, 26:46–54, 2021. doi:10.7250/CSIMQ.2021-26.03.

25 Xiaolong Li and Jennifer Wortman Vaughan. An axiomatic characterization of adaptive-liquidity market makers. In *ACM EC*, 2013.

26 Manifold. Maniswap. https://www.notion.so/Maniswap-ce406e1e897d417cbd491071ea8a0c39, 2022.

27 Jason Milionis, Ciamac C. Moallemi, and Tim Roughgarden. A myersonian framework for optimal liquidity provision in automated market makers, 2023. arXiv:2303.00208.

28 A. Othman and T. Sandholm. Profit-charging market makers with bounded loss, vanishing bid/ask spreads, and unlimited market depth. In *ACM EC*, 2012.
29 A. Othman, T. Sandholm, D. M Pennock, and D. M Reeves. A practical liquidity-sensitive automated market maker. In *Proceedings of the 11th ACM conference on Electronic commerce*, pages 377–386, 2010.
30 Abraham Othman. New invariants for automated market making, 2021. URL: `https://github.com/shipyard-software/market-making-whitepaper/blob/main/paper.pdf`.
31 Abraham Othman, David M Pennock, Daniel M Reeves, and Tuomas Sandholm. A practical liquidity-sensitive automated market maker. *ACM Transactions on Economics and Computation (TEAC)*, 1(3):1–25, 2013.
32 Paradigm. Uniswap v3: The universal amm. `https://www.paradigm.xyz/2021/06/uniswap-v3-the-universal-amm`, 2021.
33 Jan Christoph Schlegel, Mateusz Kwaśnicki, and Akaki Mamageishvili. Axioms for constant function market makers, 2022. `doi:10.48550/ARXIV.2210.00048`.

A Prediction Market and CFMM Equivalence in more detail

A.1 Cost function prediction markets

To begin, let us briefly review prediction markets for Arrow–Debreu (AD) securities, which are designed to elicit a full probability distribution over a future event. Formally, let $\mathcal{Y} = \{1, \dots, n\}$ be a set of outcomes, and let Y be the future event, a random variable taking values in $\mathcal{Y}$. For example, in a tournament with n teams, Y can be the winner of the tournament, and $Y = i$ is the outcome where team i wins. The probability simplex on $\mathcal{Y}$ is denoted $\Delta_{\mathcal{Y}}$. The AD prediction market is a market with n assets $A_1, \dots, A_n$. When the outcome of Y is observed, each unit of A_i pays off to the owner one unit of cash (some fixed currency, such as US dollars) if $Y = i$, and pays off zero otherwise.[3] We allow traders to both buy and sell assets, including holding negative amounts, representing a short position.

By design, a trader's fair price for A_i is therefore the probability $\Pr[Y = i]$ according to their belief. The current market prices for $A_1, \dots, A_n$ form a "consensus" prediction in the form of a probability distribution over $\mathcal{Y}$.

A priori, as with general asset markets, it may seem that the valid trades and associated prices at each moment can be set in essentially any manner whatsoever. It may therefore be surprising that, in order to successfully elicit and aggregate information while maintaining path independence, a market for contingent securities must take the form of a *cost-function market maker*, as we define it next. We then sketch the proof of this well-known characterization in Fact 2.

The classic cost-function market maker for Arrow–Debreu securities, slightly generalized for our setting, operates as follows.

▶ **Definition 12.** *Let $C : \mathbb{R}^n \to \mathbb{R}$ be convex, increasing, and $\mathbb{1}$-invariant. Let $\mathbf{q}_0 \in \mathbb{R}^n$ be an initial state. The* AD market maker *defined by C, with initial state $\mathbf{q}_0$, operates as follows. At round $t = 1, 2, \dots$*

- *A trader can request any bundle of securities $\mathbf{r}_t \in \mathbb{R}^n$.*
- *The market state updates to $\mathbf{q}_t = \mathbf{q}_{t-1} + \mathbf{r}_t$.*
- *The trader pays the market maker $C(\mathbf{q}_t) - C(\mathbf{q}_{t-1})$ in cash.*

After an outcome of the form $Y = i$ occurs, for each round t, the trader responsible for the trade $\mathbf{r}_t$ is paid $(\mathbf{r}_t)_i$ in cash, i.e. the number of shares purchased in outcome i.

[3] In more general prediction markets than the Arrow–Debreu case described here, there can be any number of assets that each pay off according to an arbitrary function of Y.

When C is differentiable, the instantaneous prices of the securities are given by $\mathbf{p}_t = \nabla C(\mathbf{q}_t)$. That is, the approximate cost of any bundle $x \in \mathbb{R}^{\mathcal{Y}}$ is given by $\mathbf{p}_t \cdot x$. (The exact cost is given by integrating this trade from zero to x, resulting in $C(\mathbf{q}_t + x) - C(\mathbf{q}_t)$.) We can therefore regard $\mathbf{p}_t = \nabla C(\mathbf{q}_t)$ as the market prediction. Because of $\mathbb{1}$-invariance, $\mathbf{p}_t$ is always a probability distribution.

We can cast cost-function market makers as a special case of the automated market maker framework in § 2. Specifically, represent cash by an asset $A_\$$, and consider the $n+1$ assets $\mathcal{A} = \{A_1, \ldots, A_n, A_\$\}$. Then we can restate Definition 12 as the following asset market $\mathsf{ValTrades}_C$ over $\mathcal{A}$. An important convention is that **prediction markets interpret states q and trades r as net transfers to the trader**, whereas our automated market maker follows the CFMM convention to interpret $\mathbf{q}$ and $\mathbf{r}$ as net amounts for the market maker. Therefore, a negative sign is always required to translate between the settings.

- The initial reserves are $(-\mathbf{q}_0, 0) \in \mathbb{R}^{n+1}$.
- Define $\mathsf{ValTrades}_C(h) = \{(\mathbf{r}, \alpha) \mid \mathbf{r} \in \mathbb{R}^n, \alpha = C(-\mathbf{q}_h - \mathbf{r}) - C(-\mathbf{q}_h)\}$.

In other words, $\mathsf{ValTrades}_C(h)$ consists of any bundle of securities $\mathbf{r} \in \mathbb{R}^n$ to be sold, along with a payment in cash of $C(-\mathbf{q}_h - \mathbf{r}) - C(-\mathbf{q}_h)$. As mentioned, the input to C is negated because C expects as an argument the net bundle the market has *sold* rather than purchased.

A priori, the design space for prediction markets could include any automated market maker over any set of contingent securities. From this perspective, it is perhaps surprising that, to achieve "good" elicitation of predictions according to a standard set of axioms, one *must* use a cost-function market maker. That is, the following is known:

▶ **Fact 2** ([17, 1]). *Let M be any automated market maker offering contingent securities for cash and satisfying these axioms: (1) Path Independence; (2) Incentive Compatibility for predicting Y, in the sense that: (2a) Any history h of the market defines a probability distribution over Y (the "market prediction"), and (2b) All possible predictions are achievable by some history, and (2c) A trader maximizes expected net payment by moving the market prediction match their belief $p \in \Delta_Y$. Then M is a cost-function AD prediction market.*

We recall that a *scoring rule* is a function $S : \Delta_{\mathcal{Y}} \times \mathcal{Y} \to \mathbb{R} \cup \{-\infty\}$, assigning a score $S(\mathbf{p}, i)$ to each prediction $\mathbf{p}$ when the obeserved outcome is i. The scoring rule is *proper* if the expected score, with respect to any given belief $\mathbf{p}^*$, is maximized by reporting $\mathbf{p} = \mathbf{p}^*$. It is *strictly proper* if $\mathbf{p} = \mathbf{p}^*$ uniquely maximizes the expected score. Examples of strictly proper scoring rules are the log score $S(\mathbf{p}, i) = \log \mathbf{p}_i$ and the Brier or quadratic score $S(\mathbf{p}, i) = 2\mathbf{p}_i - \sum_j \mathbf{p}_j^2$.

We now sketch the proof of Fact 2 for completeness.

Proof sketch for Fact 2. The first key step is Theorem 3.1 of [17], which states that any market maker satisfying Path Independence and Incentive Compatibility must in fact be representable as a "scoring rule market". A scoring rule market, in our terminology, is an automated market maker including an asset of cash where $\mathsf{ValTrades}(h)$ must map one-to-one to predictions $\mathbf{p}_t$, and the net payoff for moving the market prediction from $\mathbf{p}_{t-1}$ to $\mathbf{p}_t$ must be given by the formula of [23]:

$$\text{net payoff} = S(\mathbf{p}_t, i) - S(\mathbf{p}_{t-1}, i),$$

where S is a strictly proper scoring rule. The intuition is that Incentive Compatibility is characterized by use of a strictly proper scoring rule in each round, while Path Independence imposes the requirement of "telescoping sums" for sequences of predictions, giving the above formula.

The second key step is given by classic results of [11, 1] stating that scoring-rule markets are equivalent to cost-function markets. In particular, a strictly proper scoring rule S corresponds via convex duality to a convex cost function C such that

$$S(\mathbf{p}_t, i) - S(\mathbf{p}_{t-1}, i) = (\mathbf{q}_t)_i - (\mathbf{q}_{t-1})_i - [C(\mathbf{q}_t) - C(\mathbf{q}_{t-1})] .$$

(Also, C is differentiable and has a certain set of gradients, namely the probability simplex.) Here convex duality gives a one-to-one correspondence between market states $\mathbf{q}_t$ and predictions $\mathbf{p}_t$, such that the trade $\mathbf{r}$ with $\mathbf{q}_t = \mathbf{q}_{t-1} + \mathbf{r}$ is equivalent to the prediction update $\mathbf{p}_{t-1} \to \mathbf{p}_t$. In other words, the cost-function interface to the market is equivalent to the scoring-rule interface.[4]

To complete the sketch, we observe the properties of C. Convexity follows from the convex duality of scoring rules discussed above. Increasing and ones-invariance follow from Incentive Compatibility, as follows. Because each trade should correspond one-to-one to a feasible prediction, Incentive Compatibility implies there must be no available trades that make a guaranteed profit regardless of the outcome (what the prediction market literature calls an arbitrage opportunity). Otherwise, traders with many different beliefs would all choose such a trade. In particular, nonnegative trade bundles must cost a positive amount of money, which gives that C is increasing. And α units of the grand bundle $(1, \ldots, 1)$ must cost exactly α, since otherwise buying or short-selling it would guarantee profit; this yields ones-invariance. ◀

Fact 2 is essentially a characterization: any cost-function AD prediction market satisfies the above axioms, if C is also differentiable with a certain set of gradients [1]. For completeness, Appendix A includes a sketch of a proof. The following are the key ideas derived from [17, 1]: Incentive Compatibility requires the market to use *proper scoring rules*, functions that induce truthful forecasts from individual experts (see Section 3.3). Path Independence then implies that the market must use "chained" scoring rules as proposed by [23]. Finally, it is known that Hanson's scoring-rule markets are equivalent to cost function based prediction markets via convex duality [1].

A.2 CFMMs elicit ratios of valuations

A consequence of the equivalence is that CFMMs can be viewed as prediction markets. But what exactly do they "predict"? The answer is: **CFMMs elicit ratios of valuations.** Intuitively, a CFMM can reveal that the market currently values one asset at, e.g., 5 times another asset, but not the absolute value of each asset relative to some outside numeraire. In particular, this allows the market to come to a consensus about e.g. this 5-to-1 ratio, even among traders who disagree about the absolute value of the assets.

Our point here is that this intuition is strongly true in a formal sense. To explain, let us return to the "cashless prediction market", in which there are n assets. Instead of payments in cash, each asset had a "price" in units of the grand bundle $(1, \ldots, 1)$. Now generalize: consider a market with n arbitrary assets $A_1, \ldots, A_n$ with some nonnegative value. One can similarly use a cost function C to define a cashless automated market maker with all prices given in units of the grand bundle $(1, \ldots, 1)$.

[4] There is a difference in timing, because as usually implemented, a cost function collects the cash payment at the time of trade while the securities pay off at the end; the scoring rule assesses all payments at the end. For Fact 2, we must either wait to collect the payment until the end, or assume that one unit of cash has constant utility over time (the usual assumption).

When $A_1, \ldots, A_n$ are contingent securities (e.g. A_i pays off equal to the number of points scored by team i during the season), this is known as a *ratio-of-expectations* market [17]. The market prices reflect the *ratio* between the value of A_1 and of the grand bundle $A_1 + \cdots + A_n$, because traders pay in units of the latter to obtain units of the former. However, the market valuation of the grand bundle itself is never revealed. So the prices reveal the relative values of each security, but not their absolute values.

In exactly the same way, the market prices in a CFMM reflect *ratios* of valuations of each asset to the others. In particular, at any history h, a CFMM defines via $\mathsf{ValTrades}_\varphi(h)$ a truthful mechanism to induce a single agent to reveal, not their valuation in cash of the assets, but their ratios of the value of each asset to the value of the grand bundle. (This is an extension of the fact that CFMMs define proper scoring rules.).

It is not surprising that a CFMM's prices reflect ratios of valuations of the assets. What may be surprising is that, *if one designed a market maker with the intention of eliciting ratios of valuations*, it would result in a CFMM. For example, a primary purpose of CFMMs is to "provide liquidity", i.e. offer prices to exchange any asset A_i for another A_j. Paying in units of the grand bundle clearly defeats this purpose, which seems to make a ratio-of-valuations market a poor choice. However, Incentive Compatibility implies that a ratio-of-valuations market should be implemented by a cost-function, which is equivalent to a CFMM that satisfies Liquidation. So good market-making axioms fall out "for free" from elicitation ones.

A.3 Homogeneous CFMMs and unique scoring rules

As essentially all CFMMs used in practice are 1-homogeneous [5], it is natural to ask whether Construction 1 could have produced them. We now show the answer is yes: every 1-homogeneous (increasing, concave) potential function is the result of Construction 1 for some choice of cost function.

▶ **Proposition 11.** *A potential $\varphi : \mathbb{R}^n_{>0}$ is the result of applying Construction 1 if and only if φ is 1-homogeneous, increasing, and concave.*

As a result of Proposition 11, through Corollary 9, every 1-homogeneous CFMM can be associated with a unique proper scoring rule; the various level sets of the potential merely correspond to scaling up the underlying scoring rule. As we will see in § 4.1, this fact can be used to construct useful CFMMs from the vast literature on proper scoring rules.

▶ **Corollary 13.** *Every CFMM for a 1-homogeneous, concave, increasing φ is associated with a unique proper scoring rule.*

Proof of Proposition 11 and Corollary 13 can be found in the full version.

Construction 1 can be regarded geometrically as a gauge function (also known as a Minkowski functional) on a star set, which is known to be a characterization of 1-homogeneous functions. For any set $K \subseteq \mathbb{R}^n$, the gauge function (Minkowski functional) of K is defined as $g_K(\mathbf{q}) = \inf\{c > 0 \mid \mathbf{q} \in cK\}$. Letting $K = \{\mathbf{q} \in \mathbb{R}^n_{>0} \mid C(-\mathbf{q}) \leq 0\}$, we have $\varphi(\mathbf{q}) = \inf\{c > 0 \mid \mathbf{q}/c \in K\} = g_K(\mathbf{q})$.