Market Making with Decreasing Utility for Information

Miroslav Dudík
Microsoft Research

Rafael Frongillo
Microsoft Research

Jennifer Wortman Vaughan
Microsoft Research

Abstract

We study information elicitation in cost-function-based combinatorial prediction markets when the market maker's utility for information decreases over time. In the *sudden revelation* setting, it is known that some piece of information will be revealed to traders, and the market maker wishes to prevent guaranteed profits for trading on the sure information. In the *gradual decrease* setting, the market maker's utility for (partial) information decreases continuously over time. We design adaptive cost functions for both settings which: (1) preserve the information previously gathered in the market; (2) eliminate (or diminish) rewards to traders for the publicly revealed information; (3) leave the reward structure unaffected for other information; and (4) maintain the market maker's worst-case loss. Our constructions utilize mixed Bregman divergence, which matches our notion of utility for information.

1 INTRODUCTION

Prediction markets have been used to elicit information in a variety of domains, including business [6, 7, 12, 28], politics [4, 29], and entertainment [25]. In a prediction market, traders buy and sell *securities* with values that depend on some unknown future outcome. For example, a market might offer securities worth \$1 if Norway wins a gold medal in Men's Moguls in the 2014 Winter Olympics and \$0 otherwise. Traders are given an incentive to reveal their beliefs about the outcome by buying and selling securities, e.g., if the current price of the above security is \$0.15, traders who believe that the probability of Norway winning is more than 15% are incentivized to buy and those who believe that the probability is less than 15% are incentivized to sell. The equilibrium price reflects the market consensus about the security's expected payout (which here coincides with the probability of Norway winning the medal).

There has recently been a surge of research on the design of prediction markets operated by a centralized authority called a *market maker*, an algorithmic agent that offers to buy or sell securities at some current price that depends on the history of trades in the market. Traders in these markets can express their belief whenever it differs from the current price by either buying or selling, regardless of whether other traders are willing to act as a counterparty, because the market maker always acts as a counterparty, thus "providing the liquidity" and subsidizing the information collection. This is useful in situations when the lack of interested traders would negatively impact the efficiency in a traditional exchange. Of particular interest to us are *combinatorial prediction markets* [8–10, 17–19, 26] which offer securities on various related events such as "Norway wins a total of 4 gold medals in the 2014 Winter Olympics" and "Norway wins a gold medal in Men's Moguls." In combinatorial markets with large, expressive security spaces, such as an Olympics market with securities covering 88 nations participating in 98 events, the lack of an interested counterparty is a major concern. Only a single trader may be interested in trading the security associated with a specific event, but we would still like the market to incorporate this trader's information.

Most market makers considered in the literature are implemented using a pricing function called the *cost function* [11]. While such markets have many favorable properties [1, 2], the current approaches have several drawbacks that limit their applicability in real-world settings. First, existing work implicitly assumes that the outcome is revealed all at once. When concerned about "just-in-time arbitrage," in which traders closer to the information source make last-minute guaranteed profits by trading on the sure information before the market maker can adjust prices, the market maker can prevent such profits by closing the entire market just before the outcome is revealed. This approach is undesirable when partial information about the outcome is revealed over time, as is often the case in practice, including the Olympics market. For instance, we may learn the results of Men's Moguls before Ladies' Figure Skating has taken place. Closing a large combinatorial market whenever a small portion of the outcome is determined seems to be an unreasonably large intervention.

Second, in real markets, the information captured by the market's consensus prices often becomes less useful as the revelation of the outcome approaches. Consider a market over the event "Unemployment in the U.S. falls below 5.8% by the end of 2015." Although there may be a particular moment when the unemployment rate is publicly revealed, this information becomes gradually less useful as that moment approaches; the government may be less able to act on the information as the end of the year draws near. In the Olympics market, the outcome of a particular competition is often more certain as the final announcement approaches, e.g., if one team is far ahead by the half-time of a hockey game, market forecasts become less interesting. Existing market makers fail to take this diminishing utility for information into account, with the strength of the market incentives remaining constant over time.

To address these two shortcomings of existing markets, we consider two settings:

- a *sudden revelation* setting in which it is known that some piece of information (such as the winner of Men's Moguls) will be publicly revealed at a particular time, driving the market maker's utility for this information to zero; crucially, in this setting we assume that the market maker *does not* have direct access to this information at the time it is revealed, which is realistic in the case of the Olympics where a human might not be available to input winners for all 98 events in real time;
- a *gradual decrease* setting in which the market maker has a diminishing utility for a piece of information (such as the unemployment rate for 2015) over time and therefore is increasingly unwilling to pay for this information even while other information remains valuable.

The sudden revelation setting can be viewed as a special case of the gradual decrease setting. In both cases, we model the relevant information as a variable X, representing a partly determined outcome such as the identity of the gold medal winner in a single sports event.

We consider cost-function-based market makers in which the cost function switches one or many times, and aim to design switching strategies such that: (1) information previously gathered in the market is not lost at the time of the switch, (2) a trader who knows the value of X but has no additional information is unable to profit after the switch (for the sudden revelation setting) or is able to profit less and less over time (in the gradual decrease setting), and (3) the market maker maintains the same reward structure for any other information that traders may have. To formalize these objectives, we define the notion of the market maker's utility (Sec. 2) and show how it corresponds to the *mixed Bregman divergence* [13, 15] (Sec. 2.5).

For the sudden revelation setting (Sec. 3), we introduce a generic cost function switching technique which in many cases removes the rewards for "just-in-time arbitragers" who know only the value of X, while allowing traders with other information to profit, satisfying our objectives.

For the gradual decrease setting (Sec. 4), we focus on *linearly constrained market makers* (LCMMs) [13], proposing a time-sensitive market maker that gradually decreases liquidity by employing the cost function of a different LCMM at each point in time, again meeting our objectives.

Others have considered the design of cost-function-based markets with adaptive liquidity [3, 21–24]. That line of research has typically focused on the goal of slowing down price movement as more money enters the market. In contrast, we adjust liquidity to reflect the current market maker's utility which can be viewed as something external to trading in the market. Additionally, we change liquidity only in the "low-utility" parts of the market, whereas previous work considered market-wide liquidity shifts. Brahma et al. [5] designed a Bayesian market maker that adapts to perceived increases in available information. Our market maker does not try to infer high information periods, but assumes that a schedule of public revelations is given a priori. Our market makers have guaranteed bounds on worst-case loss whereas those of Brahma et al. [5] do not.

2 SETTING AND DESIDERATA

We begin by reviewing cost-function-based market making before describing our desiderata. Here and throughout the paper we make use of many standard results from convex analysis, summarized in Appendix A. All of the proofs in this paper are relegated to the appendix. [1]

2.1 COST-FUNCTION-BASED MARKET MAKING

Let Ω denote the *outcome space*, a finite set of mutually exclusive and exhaustive states of the world. We are interested in the design of cost-function-based market makers operating over a set of K *securities* on Ω specified by a *payoff function* $\boldsymbol{\rho} : \Omega \to \mathbb{R}^K$, where $\boldsymbol{\rho}(\omega)$ denotes the vector of security payoffs if the outcome $\omega \in \Omega$ occurs. Traders may purchase *bundles* $\boldsymbol{r} \in \mathbb{R}^K$ of securities from the market maker, with r_i denoting the quantity of security i that the trader would like to purchase; negative values of r_i are permitted and represent short selling. A trader who purchases a bundle $\boldsymbol{r}$ of securities pays a specified cost for this bundle up front and receives a (possibly negative) payoff of $\boldsymbol{\rho}(\omega) \cdot \boldsymbol{r}$ if the outcome $\omega \in \Omega$ occurs.

Following Chen and Pennock [11] and Abernethy et al. [1, 2], we assume that the market maker initially prices securities using a convex potential function $C : \mathbb{R}^K \to \mathbb{R}$,

[1] The full version of this paper on arXiv includes the appendix.

called the *cost function*. The current state of the market is summarized by a vector $\boldsymbol{q} \in \mathbb{R}^K$, where q_i denotes the total number of shares of security i that have been bought or sold so far. If the market state is $\boldsymbol{q}$ and a trader purchases the bundle $\boldsymbol{r}$, he must pay the market maker $C(\boldsymbol{q}+\boldsymbol{r}) - C(\boldsymbol{q})$. The new market state is then $\boldsymbol{q}+\boldsymbol{r}$. The *instantaneous price* of security i is $\partial C(\boldsymbol{q})/\partial q_i$ whenever well-defined; this is the price per share of an infinitesimally small quantity of security i, and is frequently interpreted as the traders' collective belief about the expected payoff of this security. Any expected payoff must lie in the convex hull of the set $\{\boldsymbol{\rho}(\omega)\}_{\omega\in\Omega}$, called *price space*, denoted $\mathcal{M}$.

While our cost function might not be differentiable at all states $\boldsymbol{q}$, it is always *subdifferentiable* thanks to convexity, i.e., its subdifferential $\partial C(\boldsymbol{q})$ is non-empty for each $\boldsymbol{q}$ and, if it is a singleton, it coincides with the gradient. Let $\boldsymbol{p}(\boldsymbol{q}) \coloneqq \partial C(\boldsymbol{q})$ be called the *price map*. The set $\boldsymbol{p}(\boldsymbol{q})$ is always convex and can be viewed as a multi-dimensional version of the "bid-ask spread". In a state $\boldsymbol{q}$, a trader can make an expected profit if and only if he believes that $\mathbb{E}[\boldsymbol{\rho}(\omega)] \notin \boldsymbol{p}(\boldsymbol{q})$. If C is differentiable at $\boldsymbol{q}$, we slightly abuse notation and also use $\boldsymbol{p}(\boldsymbol{q}) \coloneqq \nabla C(\boldsymbol{q})$.

We assume that the cost function satisfies two standard properties: *no arbitrage* and *bounded loss*. The former means that as long as all outcomes ω are possible, there are no market transactions with a guaranteed profit for a trader. The latter means that the worst-case loss of the market maker is a priori bounded by a constant. Together, they imply that the cost function C can be written in the form $C(\boldsymbol{q}) = \sup_{\boldsymbol{\mu}\in\mathcal{M}}[\boldsymbol{\mu}\cdot\boldsymbol{q} - R(\boldsymbol{\mu})]$, where R is the convex conjugate of C, with $\operatorname{dom} R = \mathcal{M}$. See Abernethy et al. [1, 2] for an analysis of the properties of such markets.

Example 1. *Logarithmic market-scoring rule (LMSR).* The LMSR of Hanson [18, 19] is a cost function for a *complete market* where traders can express any probability distribution over Ω. Here, for any $K \geq 1$, $\Omega = [K] \coloneqq \{1, \ldots, K\}$ and $\rho_i(\omega) = \mathbf{1}[i=\omega]$ where $\mathbf{1}[\cdot]$ is a 0/1 indicator, i.e., the security i pays out \$1 if the outcome i occurs and \$0 otherwise. The price space $\mathcal{M}$ is the simplex of probability distributions in K dimensions. The cost function is $C(\boldsymbol{q}) = \ln\bigl(\sum_{i=1}^K e^{q_i}\bigr)$, which is differentiable and generates prices $p_i(\boldsymbol{q}) = e^{q_i}/\bigl(\sum_{j=1}^K e^{q_j}\bigr)$. Here R is the negative entropy function, $R(\boldsymbol{\mu}) = \sum_{i=1}^K \mu_i \ln \mu_i$.

Example 2. *Square.* The square market consists of two independent securities ($K = 2$) each paying out either \$0 or \$1. This can be encoded as $\Omega = \{0,1\}^2$ with $\rho_i(\boldsymbol{\omega}) = \omega_i$ for $i = 1, 2$. The price space is the unit square $\mathcal{M} = [0,1]^2$. Consider the cost function $C(\boldsymbol{q}) = \ln\bigl(1+e^{q_1}\bigr) + \ln\bigl(1+e^{q_2}\bigr)$, which is differentiable and generates prices $p_i(\boldsymbol{q}) = e^{q_i}/(1+e^{q_i})$ for $i = 1, 2$. Using this cost function is equivalent to running two independent binary markets, each with an LMSR cost function. We have $R(\boldsymbol{\mu}) = \sum_{i=1}^2 \mu_i \ln \mu_i + (1-\mu_i)\ln(1-\mu_i)$.

PROTOCOL 1: Sudden Revelation Market Makers

Input: initial cost function C, initial state $\boldsymbol{s}^{\texttt{ini}}$, switch time t, update functions $\texttt{NewCost}(\boldsymbol{q})$, $\texttt{NewState}(\boldsymbol{q})$

Until time t:
  sell bundles $\boldsymbol{r}^1, \ldots, \boldsymbol{r}^N$ priced using C
    for the total cost $C(\boldsymbol{s}^{\texttt{ini}}+\boldsymbol{r}) - C(\boldsymbol{s}^{\texttt{ini}})$ where $\boldsymbol{r} = \sum_{i=1}^N \boldsymbol{r}^i$
  let $\boldsymbol{s} = \boldsymbol{s}^{\texttt{ini}} + \boldsymbol{r}$
At time t:
  $\tilde{C} \leftarrow \texttt{NewCost}(\boldsymbol{s})$
  $\tilde{\boldsymbol{s}} \leftarrow \texttt{NewState}(\boldsymbol{s})$
After time t:
  sell bundles $\tilde{\boldsymbol{r}}^1, \ldots, \tilde{\boldsymbol{r}}^{\tilde{N}}$ priced using $\tilde{C}$
    for the total cost $\tilde{C}(\tilde{\boldsymbol{s}}+\tilde{\boldsymbol{r}}) - \tilde{C}(\tilde{\boldsymbol{s}})$ where $\tilde{\boldsymbol{r}} = \sum_{i=1}^{\tilde{N}} \tilde{\boldsymbol{r}}^i$
  let $\tilde{\boldsymbol{s}}^{\texttt{fin}} = \tilde{\boldsymbol{s}} + \tilde{\boldsymbol{r}}$
Observe ω
Pay $(\boldsymbol{r}+\tilde{\boldsymbol{r}})\cdot\boldsymbol{\rho}(\omega)$ to traders

PROTOCOL 2: Gradual Decrease Market Makers

Input: time-sensitive cost function $\mathbf{C}(\boldsymbol{q};\, t)$, initial state $\boldsymbol{s}^0$, initial time t^0, update function $\texttt{NewState}(\boldsymbol{q};\, t, t')$

For $i = 1, \ldots, N$ (where N is an unknown number of trades):
  at time $t^i \geq t^{i-1}$: receive a request for a bundle $\boldsymbol{r}^i$
  $\tilde{\boldsymbol{s}}^{i-1} \leftarrow \texttt{NewState}(\boldsymbol{s}^{i-1};\, t^{i-1}, t^i)$
  sell the bundle $\boldsymbol{r}^i$
    for the cost $\mathbf{C}(\tilde{\boldsymbol{s}}^{i-1}+\boldsymbol{r}^i;\, t^i) - \mathbf{C}(\tilde{\boldsymbol{s}}^{i-1};\, t^i)$
  $\boldsymbol{s}^i \leftarrow \tilde{\boldsymbol{s}}^{i-1} + \boldsymbol{r}^i$
Observe ω
Pay $\sum_{i=1}^N \boldsymbol{r}^i\cdot\boldsymbol{\rho}(\omega)$ to traders

Example 3. *Piecewise linear cost.* Here we describe a non-differentiable cost function for a single binary security ($K = 1$). Let $\Omega = \{0,1\}$ and $\rho(\omega) = \omega$, so $\mathcal{M} = [0,1]$. The cost function is $C(q) = \max\{0, q\}$. It gives rise to the price map such that $p(q) = 0$ if $q < 0$, and $p(q) = 1$ if $q > 0$, but at $q = 0$, we have $p(q) = [0,1]$, i.e., because of non-differentiability we have a bid-ask spread at $q = 0$. Here, $R(\mu) = \mathbb{I}\bigl[\mu \in [0,1]\bigr]$ where $\mathbb{I}[\cdot]$ is a $0/\infty$ indicator, equal to 0 if true and ∞ if false. This market is uninteresting on its own, but will be useful to us in Sec. 3.3.

2.2 OBSERVATIONS AND ADAPTIVE COSTS

We study two settings. In the *sudden revelation setting*, it is known to both the market maker and the traders that at a particular point in time (the observation time) some information about the outcome (an observation) will be publicly revealed to the traders, but not to the market maker. More precisely, let any function on Ω be called a *random variable* and its value called the *realization* of this random variable. Given a random variable $X : \Omega \to \mathcal{X}$, we assume that its realization is revealed to the traders at the observation time. For a random variable X and a possible realization x, we define the *conditional outcome space* by $\Omega^x \coloneqq \{\omega \in \Omega : X(\omega) = x\}$. After observing $X = x$ (where, using standard random variable shorthand, we write X for $X(\omega)$), the traders can conclude that

$\omega \in \Omega^x$. Note that the sets $\{\Omega^x\}_{x\in\mathcal{X}}$ form a partition of Ω.

We design *sudden revelation market makers* (Protocol 1) that replace the cost function C with a new cost function $\tilde{C}$, and the current market state $\boldsymbol{s}$ (i.e., the current value of $\boldsymbol{q}$ in the definition above) with a new market state $\tilde{\boldsymbol{s}}$ in order to reflect the decrease in the utility for information about X. Such a switch would typically occur just before the observation time. Note that we allow the new cost function $\tilde{C}$ as well as the new state $\tilde{\boldsymbol{s}}$ to be chosen adaptively according to the last state $\boldsymbol{s}$ of the original cost function C.

In the *gradual decrease* setting, the utility for information about a future observation X is decreasing continuously over time. We use a *gradual decrease market maker* (Protocol 2) with a time-sensitive cost function $\mathbf{C}(\boldsymbol{q};\, t)$ which sells a bundle $\boldsymbol{r}$ for the cost $\mathbf{C}(\boldsymbol{q}+\boldsymbol{r};\, t)-\mathbf{C}(\boldsymbol{q};\, t)$ at time t, when the market is in a state $\boldsymbol{q}$. We place no assumptions on $\mathbf{C}$ other than that for each t, the function $\mathbf{C}(\cdot;\, t)$ should be an arbitrage-free bounded-loss cost function. The market maker may modify the state between the trades.

Protocol 2 alternates between trades and cost-function switches akin to those in Protocol 1. In each iteration i, the cost function $\mathbf{C}(\cdot;\, t^{i-1})$ is replaced by the cost function $\mathbf{C}(\cdot;\, t^{i})$ while simultaneously replacing the state $\boldsymbol{s}^{i-1}$ by the state $\tilde{\boldsymbol{s}}^{i-1}$. Crucially, unlike Protocol 1, the cost-function switch here is *state independent*, so any state-dependent adaptation happens through the state update. [2]

At a high level, within each of the protocols, our goal is to design switch strategies that satisfy the following criteria:

- Any information that has already been gathered from traders about the relative likelihood of the outcomes in the conditional outcome spaces is preserved.
- A trader who has information about the observation X but has no additional information about the relative likelihood of outcomes in the conditional outcome spaces is unable to profit from this information (for sudden revelation), or the profits of such a trader are decreasing over time (for gradual decrease).
- The market maker continues to reward traders for new information about the relative likelihood of outcomes in the conditional outcome spaces as it did before, with prices reflecting the market maker's utility for information within these sets of outcomes.

To reason about these goals, it is necessary to define what we mean by the information that has been gathered in the market and the market maker's utility.

2.3 MARKET MAKER'S UTILITY

By choosing a cost function, the market maker creates an incentive structure for the traders. Ideally, this incentive structure should be aligned with the market maker's subjective utility for information. That is, the amount the market maker is willing to pay out to traders should reflect the market maker's utility for the information that the traders have provided. In this section, we study how the traders are rewarded for various kinds of information, and use the magnitude of their profits to define the market maker's implicit "utility for information" formally.

We start by defining the market maker's utility for a belief, where a *belief* $\boldsymbol{\mu} \in \mathcal{M}$ is a vector of expected security payoffs $\mathbb{E}[\boldsymbol{\rho}(\omega)]$ for some distribution over Ω.

Definition 1. *The market maker's* utility for a belief $\boldsymbol{\mu} \in \mathcal{M}$ *relative to the state* $\mathbf{q}$ *is the maximum expected payoff achievable by a trader with belief* $\boldsymbol{\mu}$ *when the current market state is* $\boldsymbol{q}$*:*

$$\mathtt{Util}(\boldsymbol{\mu};\boldsymbol{q}) \coloneqq \sup_{\boldsymbol{r}\in\mathbb{R}^K} \left[\boldsymbol{\mu}\cdot\boldsymbol{r} - C(\boldsymbol{q}+\boldsymbol{r}) + C(\boldsymbol{q})\right] \ .$$

Any subset $\mathcal{E} \subseteq \Omega$ is referred to as an *event*. Observations $X = x$ correspond to events Ω^x. Suppose that a trader has observed an event, i.e., a trader knows that $\omega \in \mathcal{E}$, but is otherwise uninformed. The market maker's utility for that event can then be naturally defined as follows.

Definition 2. *The* utility for a (non-null) event $\mathcal{E} \subseteq \Omega$ *relative to the market state* $\mathbf{q}$ *is the largest guaranteed payoff that a trader who knows* $\omega \in \mathcal{E}$ *(and has only this information) can achieve when the current market state is* $\boldsymbol{q}$*:*

$$\mathtt{Util}(\mathcal{E};\boldsymbol{q}) \coloneqq \sup_{\boldsymbol{r}\in\mathbb{R}^K} \min_{\omega\in\mathcal{E}} \Big[\boldsymbol{\rho}(\omega)\cdot\boldsymbol{r} - C(\boldsymbol{q}+\boldsymbol{r}) + C(\boldsymbol{q})\Big] \ .$$

Finally, consider the setting in which a trader has observed an event $\mathcal{E}$, and also holds a belief $\boldsymbol{\mu}$ consistent with $\mathcal{E}$. Specifically, let $\mathcal{M}(\mathcal{E})$ denote the convex hull of $\{\boldsymbol{\rho}(\omega)\}_{\omega\in\mathcal{E}}$, which is the set of beliefs consistent with the event $\mathcal{E}$, and assume $\boldsymbol{\mu} \in \mathcal{M}(\mathcal{E})$. Then we can define the "excess utility for the belief $\boldsymbol{\mu}$" as the excess utility provided by $\boldsymbol{\mu}$ over just the knowledge of $\mathcal{E}$.

Definition 3. *Given an event* $\mathcal{E}$ *and a belief* $\boldsymbol{\mu} \in \mathcal{M}(\mathcal{E})$*, the* excess utility of $\boldsymbol{\mu}$ over $\mathcal{E}$*, relative to the state* $\boldsymbol{q}$ *is:*

$$\mathtt{Util}(\boldsymbol{\mu} \mid \mathcal{E};\boldsymbol{q}) = \mathtt{Util}(\boldsymbol{\mu};\boldsymbol{q}) - \mathtt{Util}(\mathcal{E};\boldsymbol{q}) \ .$$

Note that in these definitions a trader can always choose not to trade ($\boldsymbol{r} = \mathbf{0}$), so the utility for a belief and an event is non-negative. Also it is not too difficult to see that $\mathtt{Util}(\boldsymbol{\mu};\boldsymbol{q}) \geq \mathtt{Util}(\mathcal{E};\boldsymbol{q})$ for any $\boldsymbol{\mu} \in \mathcal{M}(\mathcal{E})$, so the excess utility for a belief is also non-negative.

In Sec. 2.5, we show that given a state $\boldsymbol{q}$ and a non-null event $\mathcal{E}$, there always exists a (possibly non-unique) belief $\boldsymbol{\mu} \in \mathcal{E}$ such that $\mathtt{Util}(\boldsymbol{\mu} \mid \mathcal{E};\boldsymbol{q}) = 0$. Thus, a trader with such a "worst-case" belief is able to achieve in expectation no reward beyond what any trader that just observed $\mathcal{E}$ would receive. We show that these worst-case beliefs correspond to certain kinds of "projections" of the current price $\boldsymbol{p}(\boldsymbol{q})$ onto $\mathcal{M}(\mathcal{E})$. For LMSR, the projections are with respect to KL divergence and correspond to the usual conditional probability distributions. Moreover, for sufficiently

[2] This simplifying restriction matches our solution concept in Sec. 4, but it could be dropped for greater generality.

Table 1: Information Desiderata

PRICE	*Preserve prices:* $\tilde{\boldsymbol{p}}(\tilde{\boldsymbol{s}}) = \boldsymbol{p}(\boldsymbol{s})$.
CONDPRICE	*Preserve conditional prices:* $\tilde{\boldsymbol{p}}(X{=}x; \tilde{\boldsymbol{s}}) = \boldsymbol{p}(X{=}x; \boldsymbol{s}) \quad \forall x \in \mathcal{X}$.
DECUTIL	*Decrease profits for uninformed traders:* $\tilde{\texttt{U}}\texttt{til}(X{=}x; \tilde{\boldsymbol{s}}) \leq \texttt{Util}(X{=}x; \boldsymbol{s}) \quad \forall x \in \mathcal{X}$, with sharp inequality if $\texttt{Util}(X{=}x; \boldsymbol{s}) > 0$.
ZEROUTIL	*No profits for uninformed traders:* $\tilde{\texttt{U}}\texttt{til}(X{=}x; \tilde{\boldsymbol{s}}) = 0 \quad \forall x \in \mathcal{X}$.
EXUTIL	*Preserve excess utility:* $\tilde{\texttt{U}}\texttt{til}(\boldsymbol{\mu} \vert X{=}x; \tilde{\boldsymbol{s}}) = \texttt{Util}(\boldsymbol{\mu} \vert X{=}x; \boldsymbol{s})$ for all $x \in \mathcal{X}$ and $\boldsymbol{\mu} \in \mathcal{M}(X{=}x)$.

smooth cost functions (including LMSR) they correspond to market prices that result when a trader is optimizing his guaranteed profit from the information $\omega \in \mathcal{E}$ as in Definition 2 (see Appendix E). Because of this motivation, such beliefs are referred to as "conditional price vectors."

Definition 4. *A vector $\boldsymbol{\mu} \in \mathcal{M}(\mathcal{E})$ is called a* conditional price vector*, conditioned on $\mathcal{E}$, relative to the state $\boldsymbol{q}$ if* $\texttt{Util}(\boldsymbol{\mu}; \boldsymbol{q}) = \texttt{Util}(\mathcal{E}; \boldsymbol{q})$. *The set of such conditional price vectors is denoted*

$$\boldsymbol{p}(\mathcal{E}; \boldsymbol{q}) := \{\boldsymbol{\mu} \in \mathcal{M}(\mathcal{E}) : \texttt{Util}(\boldsymbol{\mu}; \boldsymbol{q}) = \texttt{Util}(\mathcal{E}; \boldsymbol{q})\} \ .$$

See Appendix F for additional motivation for our definitions of utility and conditioning. With these notions defined, we can now state our desiderata.

2.4 DESIDERATA

Recall that we aim to design mechanisms which replace a cost function C at a state $\boldsymbol{s}$, with a new cost function $\tilde{C}$ at a state $\tilde{\boldsymbol{s}}$. Let $\texttt{Util}$ denote the utility for information with respect to C and $\tilde{\texttt{U}}\texttt{til}$ with respect to $\tilde{C}$, and let $\boldsymbol{p}$ and $\tilde{\boldsymbol{p}}$ be the respective price maps. In our mechanisms, we attempt to satisfy (a subset of) the conditions on information structures as listed in Table 1.

Conditions PRICE and CONDPRICE capture the requirement to preserve the information gathered in the market. The current price $\boldsymbol{p}(\boldsymbol{q})$ is the ultimate information content of the market at a state $\boldsymbol{q}$ *before* the observation time, but it is not necessarily the right notion of information content *after* the observation time. When we do not know the realization x, we may wish to set up the market so that any trader who has observed $X = x$ and would like to maximize the guaranteed profit would move the market to the same conditional price vector as in the previous market. This is captured by CONDPRICE.

DECUTIL models a scenario in which the utility for information about X decreases over time, and ZEROUTIL represents the extreme case in which utility decreases to zero. These conditions are in friction with EXUTIL, which aims to maintain the utility structure over the conditional outcome spaces. A key challenge is to satisfy EXUTIL and ZEROUTIL (or DECUTIL) simultaneously.

Apart from the information desiderata of Table 1, we would like to maintain an important feature of cost-function-based market makers: their ability to bound the worst-case loss to the market maker. Specifically, we would like to show that there is some *finite* bound (possibly depending on the initial state) such that no matter what trades are executed and which outcome ω occurs, the market maker will lose no more than the amount of the bound. It turns out that the solution concepts introduced in this paper maintain the same loss bound as guaranteed for using just the market's original cost function C, but since the focus of the paper is on the information structures, worst-case loss analysis is relegated to Appendix H.

In Sec. 3, we study in detail the sudden revelation setting with the goal of instantiating Protocol 1 in a way that achieves ZEROUTIL while satisfying CONDPRICE and EXUTIL. Our key result is a characterization and a geometric sufficient condition for when this is possible.

In Sec. 4, we examine instantiations of Protocol 2 for the gradual decrease setting. Our construction focuses on linearly-constrained market makers (LCMM) [13], which naturally decompose into submarkets. We show how to achieve PRICE, CONDPRICE, DECUTIL and EXUTIL in LCMMs. We also show that it is possible to simultaneously decrease the utility for information in each submarket according to its own schedule, while maintaining PRICE.

Before we develop these mechanisms, we introduce the machinery of Bregman divergences, which helps us analyze notions of utility for information.

2.5 BREGMAN DIVERGENCE AND UTILITY

To analyze the market maker's utility for information, we show how it corresponds to a specific notion of distance built into the cost function, the *mixed (or generalized) Bregman divergence* [13, 15]. Let R be the conjugate of C.[3] The mixed Bregman divergence between a belief $\boldsymbol{\mu}$ and a state $\boldsymbol{q}$ is defined as $D(\boldsymbol{\mu} \Vert \boldsymbol{q}) := R(\boldsymbol{\mu}) + C(\boldsymbol{q}) - \boldsymbol{q} \cdot \boldsymbol{\mu}$. The conjugacy of R and C implies that $D(\boldsymbol{\mu} \Vert \boldsymbol{q}) \geq 0$ with equality iff $\boldsymbol{\mu} \in \partial C(\boldsymbol{q}) = \boldsymbol{p}(\boldsymbol{q})$, i.e., if the price vector "matches" the state (see Appendix A). The geometric interpretation of mixed Bregman divergence is as a gap between a tangent and the graph of the function R (see Fig. 1).

To see how the divergence relates to traders' beliefs, consider a trader who believes that $\mathbb{E}[\boldsymbol{\rho}(\omega)] = \boldsymbol{\mu}'$ and moves the market from state $\boldsymbol{q}$ to state $\boldsymbol{q}'$. The expected payoff to this trader is $(\boldsymbol{q}' - \boldsymbol{q}) \cdot \boldsymbol{\mu}' - C(\boldsymbol{q}') + C(\boldsymbol{q}) = D(\boldsymbol{\mu}' \Vert \boldsymbol{q}) - D(\boldsymbol{\mu}' \Vert \boldsymbol{q}')$. This payoff increases as $D(\boldsymbol{\mu}' \Vert \boldsymbol{q}')$

[3] The conjugate is also, less commonly, called the "dual".

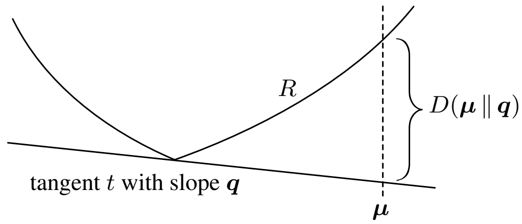

Figure 1: The mixed Bregman divergence $D(\boldsymbol{\mu}\|\boldsymbol{q})$ derived from the conjugate pair C and R measures the distance between the tangent with slope $\boldsymbol{q}$ and the value of R evaluated at $\boldsymbol{\mu}$. By conjugacy, the tangent t is described by $t(\boldsymbol{\mu}) = \boldsymbol{\mu}\cdot\boldsymbol{q} - C(\boldsymbol{q})$. Note that the divergence is well defined even when R is not differentiable, because each slope vector determines a unique tangent.

decreases. Thus, subject to the trader's budget constraints, the trader is incentivized to move to the state $\boldsymbol{q}'$ which is as "close" to his/her belief $\boldsymbol{\mu}'$ as possible in the sense of a smaller value $D(\boldsymbol{\mu}'\|\boldsymbol{q}')$, with the largest expected payoff when $D(\boldsymbol{\mu}'\|\boldsymbol{q}') = 0$. This argument shows that $D(\cdot\|\cdot)$ is an implicit measure of distance used by traders.

The next theorem shows that the Bregman divergence also matches the concepts defined in Sec. 2.3. Specifically, we show that (1) the utility for a belief coincides with the Bregman divergence, (2) the utility for an event $\mathcal{E}$ is the smallest divergence between the current market state and $\mathcal{M}(\mathcal{E})$, and (3) the conditional price vector is the (Bregman) projection of the current market state on $\mathcal{M}(\mathcal{E})$, i.e., it is a belief in $\mathcal{M}(\mathcal{E})$ that is "closest to" the current market state.

Theorem 1. *Let $\boldsymbol{\mu} \in \mathcal{M}$, $\boldsymbol{q} \in \mathbb{R}^K$ and $\emptyset \neq \mathcal{E} \subseteq \Omega$. Then*

$$\texttt{Util}(\boldsymbol{\mu};\boldsymbol{q}) = D(\boldsymbol{\mu}\|\boldsymbol{q}) \ , \tag{1}$$
$$\texttt{Util}(\mathcal{E};\boldsymbol{q}) = \min_{\boldsymbol{\mu}'\in\mathcal{M}(\mathcal{E})} D(\boldsymbol{\mu}'\|\boldsymbol{q}) \ , \tag{2}$$
$$\boldsymbol{p}(\mathcal{E};\boldsymbol{q}) = \operatorname{argmin}_{\boldsymbol{\mu}'\in\mathcal{M}(\mathcal{E})} D(\boldsymbol{\mu}'\|\boldsymbol{q}) \ . \tag{3}$$

We finish this section by characterizing when ExUtil is satisfied and showing that it implies CondPrice. Recall that $\Omega^x = \{\omega : \ X(\omega) = x\}$ and let $\mathcal{M}^x := \mathcal{M}(\Omega^x)$.

Proposition 1. ExUtil *holds if and only if for all $x \in X$, there exists some c^x such that for all $\boldsymbol{\mu} \in \mathcal{M}^x$, $D(\boldsymbol{\mu}\|\boldsymbol{s}) - \tilde{D}(\boldsymbol{\mu}\|\tilde{\boldsymbol{s}}) = c^x$. Moreover,* ExUtil *implies* CondPrice.

3 SUDDEN REVELATION

In this section, we consider the design of sudden revelation market makers (Protocol 1). In this setting, partial information in the form of the realization of X is revealed to market participants (but not to the market maker) at a predetermined time, as might be the case if the medal winners of an Olympic event are announced but no human is available to input this information into the automated market maker on behalf of the market organizer. The random variable X and the observation time are assumed to be known, and the market maker wishes to "close" the submarket with respect to X just before the observation time, without knowing the realization x, while leaving the rest of the market unchanged.

Stated in terms of our formalism, we wish to find functions `NewState` and `NewCost` from Protocol 1 such that the desiderata CondPrice, ExUtil, and ZeroUtil from Table 1 are satisfied. This implies that traders who know only that $X = x$ are not rewarded after the observation time, but traders with new information about the outcome space conditioned on $X = x$ are rewarded exactly as before. As a result, trading immediately resumes in a "conditional market" on $\mathcal{M}(\Omega^x)$ for the correct realization x, without the market maker needing to know x and without any other human intervention. We refer to the goal of simultaneously achieving CondPrice, ExUtil, and ZeroUtil as achieving *implicit submarket closing*.

For convenience, throughout this section we write $\mathcal{M}^x := \mathcal{M}(\Omega^x)$ to denote the conditional price space, and $\mathcal{M}^\star := \bigcup_{x\in\mathcal{X}} \mathcal{M}^x$ to denote prices possible after the observation.

3.1 SIMPLIFYING THE OBJECTIVE

We first show that achieving implicit submarket closing can be reduced to finding a function $\tilde{R}$ satisfying a simple set of constraints, and defining `NewCost` to return the conjugate $\tilde{C}$ of $\tilde{R}$. As a first step, we observe that it is without loss of generality to let `NewState` be an identity map, i.e., to assume that $\tilde{\boldsymbol{s}} = \boldsymbol{s}$; when this is not the case, we can obtain an equivalent market by setting $\tilde{\boldsymbol{s}} = \boldsymbol{s}$ and shifting $\tilde{C}$ so that the Bregman divergence is unchanged.

Lemma 1. *Any desideratum of Table 1 holds for $\tilde{C}$ and $\tilde{\boldsymbol{s}}$ if and only if it holds for $\tilde{C}'(\boldsymbol{q}) = \tilde{C}(\boldsymbol{q} + \tilde{\boldsymbol{s}} - \boldsymbol{s})$ and $\tilde{\boldsymbol{s}}' = \boldsymbol{s}$.*

To simplify exposition, we assume that $\tilde{\boldsymbol{s}} = \boldsymbol{s}$ throughout the rest of the section as we search for conditions on `NewCost` that achieve implicit submarket closing. Under this assumption, Proposition 1 can be used to characterize our goal in terms of $\tilde{R}$. Specifically, we show that ExUtil and CondPrice hold if $\tilde{R}$ differs from R by a (possibly different) constant on each conditional price space $\mathcal{M}^x$.

Lemma 2. *When $\tilde{\boldsymbol{s}} = \boldsymbol{s}$,* ExUtil *and* CondPrice *hold together if and only if there exist constants b^x for $x \in \mathcal{X}$ such that $\tilde{R}(\boldsymbol{\mu}) = R(\boldsymbol{\mu}) - b^x$ for all $x \in \mathcal{X}$ and $\boldsymbol{\mu} \in \mathcal{M}^x$.*

This suggests parameterizing our search for $\tilde{R}$ by vectors $\boldsymbol{b} = \{b^x\}_{x\in\mathcal{X}}$. For $\boldsymbol{b} \in \mathbb{R}^{\mathcal{X}}$, define a function

$$R^{\boldsymbol{b}}(\boldsymbol{\mu}) = \begin{cases} R(\boldsymbol{\mu}) - b^x & \text{if } \boldsymbol{\mu} \in \mathcal{M}^x, x \in \mathcal{X}, \\ \infty & \text{otherwise.} \end{cases}$$

If the sets $\mathcal{M}^x$ overlap, $R^{\boldsymbol{b}}$ is not well defined for all $\boldsymbol{b}$. Whenever we write $R^{\boldsymbol{b}}$, we assume that $\boldsymbol{b}$ is such that $R^{\boldsymbol{b}}$ is well defined. To satisfy Lemma 2 with a specific $\boldsymbol{b}$, it suffices to find a convex function $\tilde{R}$ "consistent with" $R^{\boldsymbol{b}}$ in the following sense.

Definition 5. *We say that a function $\tilde{R}$ is* consistent *with $R^{\boldsymbol{b}}$ if $\tilde{R}(\boldsymbol{\mu}) = R^{\boldsymbol{b}}(\boldsymbol{\mu})$ for all $\boldsymbol{\mu} \in \mathcal{M}^\star$.*

We next simplify our objective further by proving that whenever implicit submarket closing is achievable, it suffices to consider functions `NewCost` that set $\tilde{C}$ to be the conjugate of the largest convex function consistent with $R^{\boldsymbol{b}}$ for some $\boldsymbol{b} \in \mathbb{R}^{\mathcal{X}}$. To establish this, we examine properties of the *convex roof* of $R^{\boldsymbol{b}}$, the largest convex function that lower-bounds (but is not necessarily consistent with) $R^{\boldsymbol{b}}$.

Definition 6. *Given a function $f : \mathbb{R}^K \to (-\infty, \infty]$, the* convex roof *of f, denoted* $(\operatorname{conv} f)$*, is the largest convex function lower-bounding f, defined by*

$$(\operatorname{conv} f)(\boldsymbol{x}) \coloneqq \sup\{g(\boldsymbol{x}) :\ g \in \mathcal{G},\ g \le f\}$$

where $\mathcal{G}$ is the set of convex functions $g : \mathbb{R}^K \to (-\infty, \infty]$, and the condition $g \le f$ holds pointwise.

The convex roof is analogous to a convex hull, and the epigraph of $(\operatorname{conv} f)$ is the convex hull of the epigraph of f. See Urruty and Lemarchal [30, §B.2.5] for details.

Example 4. Recall the square market of Example 2. Let $X(\boldsymbol{\omega}) = \omega_1$, so traders observe the payoff of the first security at observation time. Then $\mathcal{M}^x = \{x\} \times [0, 1]$ for $x \in \{0, 1\}$. For simplicity, let $\boldsymbol{b} = \boldsymbol{0}$. We have $R^{\boldsymbol{b}}(\boldsymbol{\mu}) = \mu_2 \ln \mu_2 + (1-\mu_2)\ln(1-\mu_2)$ for $\boldsymbol{\mu} \in \mathcal{M}^1 \cup \mathcal{M}^2$ and $R^{\boldsymbol{b}}(\boldsymbol{\mu}) = \infty$ for all other $\boldsymbol{\mu}$. Examining the convex hull of the epigraph of $R^{\boldsymbol{b}}$ gives us that for all $\boldsymbol{\mu} \in [0, 1]^2$, we have $(\operatorname{conv} R^{\boldsymbol{b}})(\boldsymbol{\mu}) = \mu_2 \ln \mu_2 + (1-\mu_2)\ln(1-\mu_2)$.

As this example illustrates, the roof of $R^{\boldsymbol{b}}$ is the "flattest" convex function lower-bounding $R^{\boldsymbol{b}}$. Given the geometric interpretation of Bregman divergence (Fig. 1), a "flatter" $\tilde{R}$ yields a smaller utility for information. This flatness plays a key role in achieving ZEROUTIL. Assume that $\tilde{R}$ is consistent with $R^{\boldsymbol{b}}$, so CONDPRICE and EXUTIL hold by Lemma 2. Following the intuition in Fig. 1, to achieve ZEROUTIL, i.e., $\tilde{D}(\hat{\boldsymbol{\mu}}^x \| \boldsymbol{s}) = 0$ across all $x \in \mathcal{X}$ and $\hat{\boldsymbol{\mu}}^x \in \boldsymbol{p}(\Omega^x; \boldsymbol{s})$, it must be the case that for all x and $\hat{\boldsymbol{\mu}}^x$, the function values $\tilde{R}(\hat{\boldsymbol{\mu}}^x)$ lie on the tangent of $\tilde{R}$ with slope $\boldsymbol{s}$. That is, the graph of $\tilde{R}$ needs to be *flat* across the points $\hat{\boldsymbol{\mu}}^x$. This suggests that the roof might be a good candidate for $\tilde{R}$. This intuition is formalized in the following lemma, which states that instead of considering arbitrary convex $\tilde{R}$ consistent with $R^{\boldsymbol{b}}$, we can consider $\tilde{R}$ which take the form of a convex roof.

Lemma 3. *If any convex function $\tilde{R}$ is consistent with $R^{\boldsymbol{b}}$ then so is the convex roof $\tilde{R}' = (\operatorname{conv} R^{\boldsymbol{b}})$. Furthermore, if $\tilde{R}$ satisfies* ZEROUTIL *or* DECUTIL *then so does $\tilde{R}'$.*

3.2 IMPLICIT SUBMARKET CLOSING

We now have the tools to answer the central question of this section: When can we achieve implicit submarket closing? Lemma 1 implies that we can assume that `NewState` is the identity function, and Lemmas 2 and 3 imply that it suffices to consider functions `NewCost` that set $\tilde{C}$ to the conjugate of $\tilde{R} = (\operatorname{conv} R^{\boldsymbol{b}})$ for some $\boldsymbol{b} \in \mathbb{R}^{\mathcal{X}}$. What remains is to find the vector $\boldsymbol{b}$ that guarantees ZEROUTIL. As mentioned above, ZEROUTIL is satisfied if and only if $\big(\hat{\boldsymbol{\mu}}^x, \tilde{R}(\hat{\boldsymbol{\mu}}^x)\big)$ lies on the tangent of $\tilde{R}$ with the slope $\boldsymbol{s}$ for all $x \in \mathcal{X}$ and $\hat{\boldsymbol{\mu}}^x \in \boldsymbol{p}(\Omega^x; \boldsymbol{s})$. This implies that $\tilde{R}(\hat{\boldsymbol{\mu}}^x) = \hat{\boldsymbol{\mu}}^x \cdot \boldsymbol{s} - c$ for all x and $\hat{\boldsymbol{\mu}}^x$ and some constant c. The specific choice of c does not matter since $\tilde{D}$ is unchanged by vertical shifts of the graph of $\tilde{R}$. For convenience, we set $c = C(\boldsymbol{s})$, which makes the tangents of R and $\tilde{R}$ with the slope $\boldsymbol{s}$ coincide. This and Lemma 2 then yield the choice of $\boldsymbol{b} = \hat{\boldsymbol{b}}$, with

$$\hat{b}^x \coloneqq R(\hat{\boldsymbol{\mu}}^x) + C(\boldsymbol{s}) - \hat{\boldsymbol{\mu}}^x \cdot \boldsymbol{s} = D(\hat{\boldsymbol{\mu}}^x \| \boldsymbol{s}) \qquad (4)$$

for all x and any choice of $\hat{\boldsymbol{\mu}}^x \in \boldsymbol{p}(\Omega^x; \boldsymbol{s})$. The resulting construction of $\tilde{R} = (\operatorname{conv} R^{\hat{\boldsymbol{b}}})$ can be described using geometric intuition. First, consider the tangent of R with slope equal to the current market state $\boldsymbol{s}$. For each $x \in \mathcal{X}$, take the subgraph of R over the set $\mathcal{M}^x$ and let it "fall" vertically until it touches this tangent at the point $\hat{\boldsymbol{\mu}}^x$. The set of fallen graphs for all x together describes $R^{\hat{\boldsymbol{b}}}$ and the convex hull of the fallen epigraphs yields $\tilde{R} = (\operatorname{conv} R^{\hat{\boldsymbol{b}}})$.

Defining `NewCost` using this construction guarantees ZEROUTIL, but CONDPRICE and EXUTIL are achieved only when $\tilde{R}$ is consistent with $R^{\hat{\boldsymbol{b}}}$. Conversely, whenever the three properties are achievable, this construction produces a function $\tilde{R}$ consistent with $R^{\hat{\boldsymbol{b}}}$. This yields a full characterization of when implicit submarket closing is achievable.

Theorem 2. *Let $\hat{\boldsymbol{b}}$ be defined as in Eq.* (4)*.* CONDPRICE*,* EXUTIL*, and* ZEROUTIL *can be satisfied using Protocol 1 if and only if* $(\operatorname{conv} R^{\hat{\boldsymbol{b}}})$ *is consistent with $R^{\hat{\boldsymbol{b}}}$. In this case, they can be achieved with* `NewState` *as the identity and* `NewCost` *outputting the conjugate of $\tilde{R} = (\operatorname{conv} R^{\hat{\boldsymbol{b}}})$.*

3.3 CONSTRUCTING THE COST FUNCTION

Theorem 2 describes how to achieve implicit submarket closing by defining the cost function $\tilde{C}$ output by `NewCost` implicitly via its conjugate $\tilde{R}$. In this section, we provide an explicit construction of the resulting cost function, and illustrate the construction through examples.

Fixing R, for each $x \in \mathcal{X}$ define a function $C^x(\boldsymbol{q}) \coloneqq \sup_{\boldsymbol{\mu} \in \mathcal{M}^x}\big[\boldsymbol{q} \cdot \boldsymbol{\mu} - R(\boldsymbol{\mu})\big]$. Each function C^x can be viewed as a bounded-loss and arbitrage-free cost function for outcomes in Ω^x. The conjugate of each C^x coincides with R on $\mathcal{M}^x$ (and is infinite outside $\mathcal{M}^x$). The explicit expression for $\tilde{C}$ is described in the following proposition.

Proposition 2. *For a given C with conjugate R, define $\hat{\boldsymbol{b}}$ as in Eq.* (4) *and let $\tilde{R} = (\operatorname{conv} R^{\hat{\boldsymbol{b}}})$. The conjugate $\tilde{C}$ of $\tilde{R}$ can be written $\tilde{C}(\boldsymbol{q}) = \max_{x \in \mathcal{X}}\big[\hat{b}^x + C^x(\boldsymbol{q})\big]$. Furthermore, for each $x \in \mathcal{X}$, $\hat{b}^x = C(\boldsymbol{s}) - C^x(\boldsymbol{s})$.*

At any market state $\boldsymbol{q}$ with a unique $\hat{x} \coloneqq \operatorname{argmax}_{x \in \mathcal{X}}\big[\hat{b}^x + C^x(\boldsymbol{q})\big]$, the price according to $\tilde{C}$ lies in the set $\mathcal{M}^{\hat{x}}$. When $\hat{x}$ is not unique, the market has a bid-ask spread. The addition of $\hat{b}^x$ ensures that the bid-ask spread at the market state $\boldsymbol{s}$ contains conditional prices $\hat{\boldsymbol{\mu}}^x$ across all x. To illustrate this construction, we return to the example of a square.

Example 5. Consider again the square market from Examples 2 and 4 with $X(\boldsymbol{\omega}) = \omega_1$. One can verify that $C^x(\boldsymbol{q}) = xq_1 + \ln\bigl(1 + e^{q_2}\bigr)$ for $x \in \{0, 1\}$. Prop. 2 gives

$$\tilde{C}(\boldsymbol{q}) = \max_{x\in\{0,1\}}\Bigl[x(q_1-s_1)+\ln(1+e^{q_2})+\ln(1+e^{s_1})\Bigr]$$
$$= \max\{0, q_1 - s_1\} + \ln(1 + e^{s_1}) + \ln(1 + e^{q_2}).$$

In switching from C to $\tilde{C}$ we have effectively changed the first term of our cost from a basic LMSR cost for a single binary security to the piecewise linear cost of Example 3, introducing a bid-ask spread for security 1 when $q_1 = s_1$; states $\boldsymbol{q} = (s_1, q_2)$ have $\tilde{\boldsymbol{p}}(\boldsymbol{q}) = [0, 1] \times \{e^{q_2}/(1 + e^{q_2})\}$. The market for security 1 has thus implicitly closed; as the new market begins with $\boldsymbol{q} = \boldsymbol{s}$, any trader can switch the price of security 1 to 0 or 1 by simply purchasing an infinitesimal quantity of security 1 in the appropriate direction, at essentially no cost and with no ability to profit.

The example above illustrates our cost function construction, but does not show that $\tilde{R}$ is consistent with $R^{\hat{\boldsymbol{b}}}$ as required by Theorem 2. In fact, it is consistent. This follows from the sufficient condition proved in Appendix G.2. Briefly, the condition is that $\mathcal{M}^\star$ does not contain any price vectors $\boldsymbol{\mu}$ that can be expressed as nontrivial convex combinations of vectors from multiple $\mathcal{M}^x$.

In Appendix G.3, we show that this sufficient condition applies to many settings of interest such as arbitrary partitions of simplex and submarket observations in binary-payoff LCMMs (defined in Sec. 4), which were used to run a combinatorial market for the 2012 U.S. Elections [14].

A case in which the sufficient condition is violated is the square market with $X(\boldsymbol{\omega}) = \omega_1 + \omega_2 \in \{0, 1, 2\}$, where $\mathcal{M}^0 = (0, 0)$ and $\mathcal{M}^2 = (1, 1)$ but $(\frac{1}{2}, \frac{1}{2}) = \frac{1}{2}(0, 0) + \frac{1}{2}(1, 1) \in \mathcal{M}^1$. This particular example also fails to satisfy Theorem 2 (see Appendix G.1), but in general the sufficient condition is not necessary (see Appendix G.4).

4 GRADUAL DECREASE

We now consider gradual decrease market makers (Protocol 2) for the gradual decrease setting in which the utility of information about a future observation X is decreasing continuously over time. We focus on *linearly constrained market makers* (LCMMs) [13], which naturally decompose into submarkets. Our proposed gradual decrease market maker employs a different LCMM at each time step, and satisfies various desiderata of Sec. 2.4 between steps.

As a warm-up for the concepts introduced in this section, we show how the "liquidity parameter" can be used to implement a decreasing utility for information.

Example 6. *Homogeneous decrease in utility for information.* We begin with a differentiable cost function C in a state $\boldsymbol{s}$. Let $\alpha \in (0, 1)$, and define $\tilde{C}(\boldsymbol{q}) = \alpha C(\boldsymbol{q}/\alpha)$, and $\tilde{\boldsymbol{s}} = \alpha\boldsymbol{s}$. $\tilde{C}$ is parameterized by the "liquidity parameter" α. The transformation $\tilde{\boldsymbol{s}}$ guarantees the preservation of prices, i.e., $\tilde{\boldsymbol{p}}(\tilde{\boldsymbol{s}}) = \nabla\tilde{C}(\tilde{\boldsymbol{s}}) = \alpha\nabla C(\tilde{\boldsymbol{s}}/\alpha)/\alpha = \nabla C(\boldsymbol{s}) = \boldsymbol{p}(\boldsymbol{s})$. We can derive that $\tilde{R}(\boldsymbol{\mu}) = \alpha R(\boldsymbol{\mu})$, and $\tilde{D}(\boldsymbol{\mu}\|\boldsymbol{q}) = \alpha D(\boldsymbol{\mu}\|\boldsymbol{q}/\alpha)$, so, for all $\boldsymbol{\mu}$, $\tilde{D}(\boldsymbol{\mu}\|\tilde{\boldsymbol{s}}) = \alpha D(\boldsymbol{\mu}\|\boldsymbol{s})$. In words, the utility for all beliefs $\boldsymbol{\mu}$ with respect to the current state is decreased according to the multiplier α.

This idea will be the basis of our construction. We next define the components of our setup and prove the desiderata.

4.1 LINEARLY CONSTRAINED MARKETS

Recall that $\boldsymbol{\rho} : \Omega \to \mathbb{R}^K$ is the payoff function. Let $\mathcal{G}$ be a system of non-empty disjoint subsets $g \subseteq [K]$ forming a partition of coordinates of $\boldsymbol{\rho}$, so $[K] = \bigcup_{g\in\mathcal{G}} g$. We use the notation $\boldsymbol{\rho}_g(\omega) := \bigl(\rho_i(\omega)\bigr)_{i\in g}$ for the block of coordinates in g, and similarly $\boldsymbol{\mu}_g$ and $\boldsymbol{q}_g$. Blocks g describe groups of securities that are treated as separate "submarkets," but there can be logical dependencies among them.

Example 7. *Medal counts.* Consider a prediction market for the Olympics. Assume that Norway takes part in n Olympic events. In each, Norway can win a gold medal or not. Encode this outcome space as $\Omega = \{0, 1\}^n$. Define random variables $X_i(\boldsymbol{\omega}) = \omega_i$ equal to 1 iff Norway wins gold in the ith Olympic event. Also define a random variable $Y = \sum_{i=1}^n X_i$ representing the number of gold medals that Norway wins in total. We create $K = 2n + 1$ securities, corresponding to 0/1 indicators of the form $\mathbf{1}[X_i = 1]$ for $i \in [n]$ and $\mathbf{1}[Y = y]$ for $y \in \{0, \ldots, n\}$. That is, $\rho_i = X_i$ for $i \in [n]$ and $\rho_{n+1+y} = \mathbf{1}[Y = y]$ for $y \in \{0, \ldots, n\}$. A natural block structure in this market is $\mathcal{G} = \bigl\{\{1\}, \{2\}, \ldots, \{n\}, \{n + 1, \ldots, 2n + 1\}\bigr\}$ with submarkets corresponding to the X_i and Y.

Given the block structure $\mathcal{G}$, the construction of a linearly constrained market begins with bounded-loss and arbitrage-free convex cost functions $C_g : \mathbb{R}^g \to \mathbb{R}$ with conjugates R_g and divergences D_g for each $g \in \mathcal{G}$. These cost functions are assumed to be easy to compute and give rise to a "direct-sum" cost $C_\oplus(\boldsymbol{q}) = \sum_{g\in\mathcal{G}} C_g(\boldsymbol{q}_g)$ with the conjugate $R_\oplus(\boldsymbol{\mu}) = \sum_{g\in\mathcal{G}} R_g(\boldsymbol{\mu}_g)$ and divergence $D_\oplus(\boldsymbol{\mu}\|\boldsymbol{q}) = \sum_{g\in\mathcal{G}} D_g(\boldsymbol{\mu}_g\|\boldsymbol{q}_g)$.

Since $C_\oplus$ decomposes, it can be calculated quickly. However, the market maker $C_\oplus$ might allow arbitrage due to the lack of consistency among submarkets since arbitrage opportunities arise when prices fall outside $\mathcal{M}$ [1]. $\mathcal{M}$ is always polyhedral, so it can be described as $\mathcal{M} = \bigl\{\boldsymbol{\mu} \in \mathbb{R}^K : \mathbf{A}^\top\boldsymbol{\mu} \geq \boldsymbol{b}\bigr\}$ for some matrix $\mathbf{A} \in \mathbb{R}^{K\times M}$ and vector $\boldsymbol{b} \in \mathbb{R}^M$. Letting $\boldsymbol{a}_m$ denote the mth column of $\mathbf{A}$, arbitrage opportunities open up if the price of the bundle $\boldsymbol{a}_m$ falls below b_m. For any $\boldsymbol{\eta} \in \mathbb{R}^M_+$, the bundle $\mathbf{A}\boldsymbol{\eta}$ presents an arbitrage opportunity if priced below $\boldsymbol{b}\cdot\boldsymbol{\eta}$.

A *linearly constrained market maker* (LCMM) is described by the cost function $C(\boldsymbol{q}) = \inf_{\boldsymbol{\eta}\in\mathbb{R}^M_+}\bigl[C_\oplus(\boldsymbol{q}+\mathbf{A}\boldsymbol{\eta})-\boldsymbol{b}\cdot\boldsymbol{\eta}\bigr]$. While the definition of C is slightly involved, the conjugate R has a natural meaning as a restriction of the direct-sum market to the price space $\mathcal{M}$, i.e., $R(\boldsymbol{\mu}) = R_{\oplus}(\boldsymbol{\mu}) + \mathbb{I}[\boldsymbol{\mu} \in \mathcal{M}]$. Furthermore, the infimum in the definition of C is always attained (see Appendix D.1). Fixing $\boldsymbol{q}$ and letting $\boldsymbol{\eta}^\star$ be a minimizer in the definition, we can think of the market maker as automatically charging traders for the bundle $\mathbf{A}\boldsymbol{\eta}^\star$, which would present an arbitrage opportunity, and returning to them the guaranteed payout $\boldsymbol{b} \cdot \boldsymbol{\eta}$. This benefits traders while maintaining the same worst-case loss guarantee for the market maker as $C_{\oplus}$ [13].

Example 8. *LCMM for medal counts.* Continuing the previous example, for submarkets X_i, we can define LMSR costs $C_i(q_i) = \ln(1 + \exp(q_i))$. For the submarket for Y, let $g = \{n+1, \ldots, 2n+1\}$ and use the LMSR cost $C_g(\boldsymbol{q}_g) = \ln\bigl(\sum_{y=0}^{n} \exp(q_{n+1+y})\bigr)$. The submarkets for X_i and Y are linked. One example of a linear constraint is based on the linearity of expectations: for any distribution, we must have $\mathbb{E}[Y] = \sum_{i=1}^{n} \mathbb{E}[X_i]$. This places an equality constraint $\sum_{y=0}^{n} y \cdot \mu_{n+1+y} = \sum_{i=1}^{n} \mu_i$ on the vector $\boldsymbol{\mu}$, which can be expressed as two inequality constraints (see Dudík et al. [13, 14] for more on constraint generation).

4.2 DECREASING LIQUIDITY

We now study the gradual decrease scenario in which the utility for information in each submarket g decreases over time. In the Olympics example, the market maker may want to continuously decrease the rewards for information about a particular event as the event takes place.

We generalize the strategy from Example 6 to LCMMs and extend them to time-sensitive cost functions by introducing the "information-utility schedule" in the form of a differentiable non-increasing function $\beta_g : \mathbb{R} \to (0, 1]$ with $\beta_g(t^0) = 1$. The speed of decrease of β_g controls the speed of decrease of the utility for information in each submarket. (We make this statement more precise in Theorem 3.)

We first define a gradual decrease direct-sum cost function $\mathbf{C}_{\oplus}(\boldsymbol{q}; t) = \sum_{g \in \mathcal{G}} \beta_g(t) C_g\bigl(\boldsymbol{q}_g / \beta_g(t)\bigr)$ which is used to define a gradual decrease LCMM, and a matching `NewState` as follows:

$$\mathbf{C}(\boldsymbol{q}; t) = \inf_{\boldsymbol{\eta} \in \mathbb{R}_+^M} \bigl[\mathbf{C}_{\oplus}(\boldsymbol{q} + \mathbf{A}\boldsymbol{\eta}; t) - \boldsymbol{b} \cdot \boldsymbol{\eta}\bigr]$$

$$\texttt{NewState}(\boldsymbol{q};\, t, \tilde{t}) = \tilde{\boldsymbol{q}}$$

$$\text{such that } \tilde{\boldsymbol{q}}_g = \tfrac{\beta_g(\tilde{t})}{\beta_g(t)}(\boldsymbol{q}_g + \boldsymbol{\delta}_g^\star) - \boldsymbol{\delta}_g^\star$$

$$\text{where } \boldsymbol{\eta}^\star \text{ is a minimizer in } \mathbf{C}(\boldsymbol{q}; t) \text{ and } \boldsymbol{\delta}^\star = \mathbf{A}\boldsymbol{\eta}^\star\ .$$

When considering the state update from time t to time $\tilde{t}$, the ratio $\beta_g(\tilde{t})/\beta_g(t)$ has the role of the liquidity parameter α in Example 6. The motivation behind the definition of `NewState` is to guarantee that $\tilde{\boldsymbol{q}}_g + \boldsymbol{\delta}_g^\star = [\beta_g(\tilde{t})/\beta_g(t)](\boldsymbol{q}_g + \boldsymbol{\delta}_g^\star)$, which turns out to ensure that $\boldsymbol{\eta}^\star$ remains the minimizer and the prices are unchanged. The preservation of prices (PRICE) is achieved by a scaling similar to Example 6, albeit applied to the market state in the direct-sum market underlying the LCMM.

This intuition is formalized in the next theorem, which shows that the above construction preserves prices and decreases the utility for information, as captured by the mixed Bregman divergence, according to the schedules β_g. We use the notation $C^t(\boldsymbol{q}) := \mathbf{C}(\boldsymbol{q}; t)$ and write D_g^t for the divergence derived from $C_g^t(\boldsymbol{q}_g) := \beta_g(t) C_g(\boldsymbol{q}_g / \beta_g(t))$.

Theorem 3. *Let $\mathbf{C}$ be a gradual decrease LCMM, let $t, \tilde{t} \in \mathbb{R}$ and $\boldsymbol{s} \in \mathbb{R}^K$. The replacement of C^t by $\tilde{C} := C^{\tilde{t}}$ and $\boldsymbol{s}$ by $\tilde{\boldsymbol{s}} :=$ `NewState`$(\boldsymbol{s};\, t, \tilde{t})$ satisfies* PRICE. *Also,*

$$\tilde{D}(\boldsymbol{\mu} \| \tilde{\boldsymbol{s}}) = \sum_{g \in \mathcal{G}} \alpha_g D_g^t(\boldsymbol{\mu}_g \| \boldsymbol{s}_g + \boldsymbol{\delta}_g^\star) + (\mathbf{A}^\top \boldsymbol{\mu} - \boldsymbol{b}) \cdot \boldsymbol{\eta}^\star \quad (5)$$

for all $\boldsymbol{\mu} \in \mathcal{M}$, where $\boldsymbol{\eta}^\star$ and $\boldsymbol{\delta}^\star$ are defined by `NewState`$(\boldsymbol{s};\, t, \tilde{t})$*, and $\alpha_g = \beta_g(\tilde{t})/\beta_g(t) > 0$.*

The first term on the right-hand side of Eq. (5) is the sum of divergences in submarkets g, each weighted by a coefficient α_g which is equal to one at $\tilde{t} = t$ and weakly decreases as $\tilde{t}$ grows. The divergences are between $\boldsymbol{\mu}_g$ and the state resulting from the arbitrager action in the direct-sum market. The second term is non-negative, since $\boldsymbol{\mu} \in \mathcal{M}$, and represents expected arbitrager gains beyond the guaranteed profit from the arbitrage in the direct-sum market. The only terms that depend on time $\tilde{t}$ are the multipliers α_g. Since they are decreasing over time, we immediately obtain that the utility for information, `Util`$(\boldsymbol{\mu}; \tilde{\boldsymbol{s}}) = \tilde{D}(\boldsymbol{\mu} \| \tilde{\boldsymbol{s}})$, is also decreasing, with the contributions from individual submarkets decreasing according to their schedules β_g.

When only one of the schedules β_g is decreasing and the other schedules stay constant, we can show that the excess utility and conditional prices are preserved (conditioned on $\boldsymbol{\rho}_g$), and under certain conditions also DECUTIL holds.

For a submarket g, let $\mathcal{X}_g := \{\boldsymbol{\rho}_g(\omega) : \omega \in \Omega\}$ be the set of realizations of $\boldsymbol{\rho}_g$. Recall that $\mathcal{M}(\mathcal{E})$ is the convex hull of $\{\boldsymbol{\rho}(\omega)\}_{\omega \in \mathcal{E}}$. We show that DECUTIL holds if C_g is differentiable and the submarket g is "tight" as follows.

Definition 7. *We say that a submarket g is* tight *if for all $\boldsymbol{x} \in \mathcal{X}_g$ the set $\{\boldsymbol{\mu} \in \mathcal{M} : \boldsymbol{\mu}_g = \boldsymbol{x}\}$ coincides with $\mathcal{M}(\boldsymbol{\rho}_g = \boldsymbol{x})$, i.e., if all the beliefs $\boldsymbol{\mu}$ with $\boldsymbol{\mu}_g = \boldsymbol{x}$ can be realized by probability distributions over states ω with $\boldsymbol{\rho}_g(\omega) = \boldsymbol{x}$. (In general, the former is always a superset of the latter, hence the name "tight" when the equality holds.)*

While this condition is somewhat restrictive, it is easy to see that all submarkets with binary securities, i.e., with $\boldsymbol{\rho}_g(\omega) \in \{0, 1\}^g$, are tight (see Appendix D.4).

Theorem 4. *Assume the setup of Theorem 3. Let $g \in \mathcal{G}$ and assume that $\beta_g(\tilde{t}) < \beta_g(t)$ whereas $\beta_{g'}(\tilde{t}) = \beta_{g'}(t)$ for $g' \neq g$. Then the replacement of C^t by $\tilde{C}$ and $\boldsymbol{s}$ by $\tilde{\boldsymbol{s}}$ satisfies* CONDPRICE *and* EXUTIL *for the random variable $\boldsymbol{\rho}_g$. Furthermore, if C_g is differentiable and the submarket g is tight, we also obtain* DECUTIL.

References

[1] Jacob Abernethy, Yiling Chen, and Jennifer Wortman Vaughan. An optimization-based framework for automated market-making. In *Proceedings of the 12th ACM Conference on Electronic Commerce*, 2011.

[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):12:1–12:39, 2013.

[3] Jacob Abernethy, Rafael Frongillo, Xiaolong Li, and Jennifer Wortman Vaughan. A general volume-parameterized market making framework. In *Proceedings of the 15th ACM Conference on Economics and Computation*, 2014.

[4] Joyce Berg, Robert Forsythe, Forrest Nelson, and Thomas Rietz. Results from a dozen years of election futures markets research. In Charles R. Plott and Vernon L. Smith, editors, *Handbook of Experimental Economics Results*, volume 1, pages 742–751. Elsevier, 2008.

[5] Aseem Brahma, Mithun Chakraborty, Sanmay Das, Allen Lavoie, and Malik Magdon-Ismail. A Bayesian market maker. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, 2012.

[6] Philip Delves Broughton. Prediction markets: Value among the crowd. *Financial Times*, April 2013.

[7] Robert Charette. An internal futures market. *Information Management*, March 2007.

[8] Y. Chen, L. Fortnow, E. Nikolova, and D.M. Pennock. Betting on permutations. In *Proceedings of the 8th ACM Conference on Electronic Commerce*, 2007.

[9] Y. Chen, L. Fortnow, N. Lambert, D. M. Pennock, and J. Wortman. Complexity of combinatorial market makers. In *Proceedings of the 9th ACM Conference on Electronic Commerce*, 2008.

[10] Y. Chen, S. Goel, and D.M. Pennock. Pricing combinatorial markets for tournaments. In *ACM Symposium on Theory of Computing*, 2008.

[11] Yiling Chen and David M. Pennock. A utility framework for bounded-loss market makers. In *Proceedings of the 23rd Conference on Uncertainty in Artificial Intelligence*, 2007.

[12] Bo Cowgill, Justin Wolfers, and Eric Zitzewitz. Using prediction markets to track information flows: Evidence from Google. Working paper, 2008.

[13] Miroslav Dudík, Sébastien Lahaie, and David M. Pennock. A tractable combinatorial market maker using constraint generation. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, 2012.

[14] Miroslav Dudík, Sébastien Lahaie, David M. Pennock, and David Rothschild. A combinatorial prediction market for the U.S. elections. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, 2013.

[15] Geoffrey J. Gordon. Regret bounds for prediction problems. In *Proceedings of the 12th Annual Conference on Computational Learning Theory*, pages 29–40, 1999.

[16] P.D. Grünwald and A.P. Dawid. Game theory, maximum entropy, minimum discrepancy and robust Bayesian decision theory. *The Annals of Statistics*, 32(4):1367–1433, 2004.

[17] M. Guo and D.M. Pennock. Combinatorial prediction markets for event hierarchies. In *International Conference on Autonomous Agents and Multiagent Systems*, 2009.

[18] R. Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):105–119, 2003.

[19] R. Hanson. Logarithmic market scoring rules for modular combinatorial information aggregation. *Journal of Prediction Markets*, 1(1):3–15, 2007.

[20] Jean-Baptiste Hiriart-Urruty and Claude Lemaréchal. *Convex Analysis and Minimization Algorithms*, volume 1. Springer, 1996.

[21] Xiaolong Li and Jennifer Wortman Vaughan. An axiomatic characterization of adaptive-liquidity market makers. In *Proceedings of the 14th ACM Conference on Electronic Commerce*, 2013.

[22] A. Othman and T. Sandholm. Liquidity-sensitive automated market makers via homogeneous risk measures. In *Proceedings of the 7th Workshop on Internet and Network Economics*, 2011.

[23] A. Othman and T. Sandholm. Profit-charging market makers with bounded loss, vanishing bid/ask spreads, and unlimited market depth. In *Proceedings of the 13th ACM Conference on Electronic Commerce*, 2012.

[24] 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*, 2010.

[25] David M. Pennock, Steve Lawrence, C. Lee Giles, and Finn A. Nielsen. The real power of artificial markets. *Science*, 291:987–988, 2002.

[26] D.M. Pennock and L. Xia. Price updating in combinatorial prediction markets with Bayesian networks. In *Proceedings of the 27th Conference on Uncertainty in Artificial Intelligence*, 2011.

[27] R. Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, 1970.

[28] Martin Spann and Bernd Skiera. Internet-based virtual stock markets for business forecasting. *Management Science*, 49 (10):1310–1326, 2003.

[29] Lyle Ungar, Barb Mellors, Ville Satopää, Jon Baron, Phil Tetlock, Jaime Ramos, and Sam Swift. The good judgment project: A large scale test of different methods of combining expert predictions. AAAI Technical Report FS-12-06, 2012.

[30] Jean-Baptiste Hiriart Urruty and Claude Lemarchal. *Fundamentals of Convex Analysis*. Springer, 2001.