Budget Constraints in Prediction Markets

Abstract

An automated market maker is a natural and common mechanism to subsidize information acquisition, revelation, and aggregation in a prediction market. The sought-after prediction aggregate is the equilibrium price. However, traders with budget constraints are restricted in their ability to impact the market price on their own. We give a detailed characterization of optimal trades in the presence of budget constraints in a prediction market with a cost-function-based automated market maker. As a concrete application of our characterization, we give sufficient conditions for a property we call *budget additivity*: two traders with budgets B and B' and the same beliefs would have a combined impact equal to a single trader with budget $B+B'$. That way, even if a single trader cannot move the market much, a crowd of like-minded traders can have the same desired effect. We show that a generalization of the heavily-used logarithmic market scoring rule is budget additive for affinely independent payoffs, but the quadratic market scoring rule is not. Our results may be used both descriptively, to understand if a particular market maker is affected by budget constraints or not, and prescriptively, as a recipe to construct markets.

1 Introduction

A prediction market is a central clearinghouse for people with differing opinions about the likelihood of an event—say Barack Obama to win the U.S. Presidential election—to trade monetary stakes in the outcome with one another. At equilibrium, the price to buy a contract paying $1 if Obama wins reflects a consensus of sorts on the probability of the event. At that price, and given the wagers already placed, no agent is willing to push the price further up or down. Prediction markets have a good track record of forecast accuracy in many domains [12, 20].

The design of *combinatorial* markets spanning multiple logically-related events raises many interesting questions. What information can be elicited—the full probability distribution, or specific properties of the distribution? What securities can the market allow traders to buy and sell? How can the market support and ensure a variety of trades? For example, in addition to the likelihood of Hillary Clinton winning the 2016 U.S. Presidential election, we may want to elicit information about the distribution of her electoral votes.[1] If we create one security for each possible outcome between 0 and 538, each paying $1 iff Clinton gets exactly that many electoral votes, the market is called *complete* and allows us to elicit any probability distribution. Alternatively, if we create just two securities, one paying out $\$x$ if Clinton wins x electoral votes, and the other paying out $\$x^2$, we cannot elicit all possible distributions, but we can still elicit the mean and variance of the number of electoral votes.

When agents are constrained in how much they can trade only by risk aversion, prediction market prices can be interpreted as a weighted average of traders' beliefs [2, 21], a natural reflection of the "wisdom of the crowd" with a good empirical track record [15] and theoretical support [2]. However, when agents are budget constrained, discontinuities and idiosyncratic results can arise [8, 17] that call into question whether the equilibrium price can be trusted to reflect any kind of useful aggregation.

We consider prediction markets with an automated market maker [1, 5, 14] that maintains standing offers to trade every security at some price. Unlike a peer-to-peer exchange, all transactions route through the market maker. The common market makers have bounded loss and are (myopically) incentive compatible: the best (immediate) strategy is for a trader to move the market price to equal

[1] A U.S. Presidential candidate receives an integer number of electoral votes between 0 and 538. The candidate getting the majority of electoral votes wins the election.

his own belief. The design of such an automated market maker boils down to choosing a convex *cost function* [1]. This amount of design freedom presents an opportunity to seek cost functions that satisfy additional desiderata such as computational tractability [1, 7].

Most of the literature assumes either risk-neutral or risk-averse traders with unbounded budgets. In this paper, we consider how agents with budget constraints trade in such markets, a practical reality in almost all prediction markets denominated in both real and virtual currencies. Our results also help with a systematic study of the market's *liquidity parameter*, or the parameter controlling the sensitivity of price to trading volume. Setting the liquidity is a nearly universal practical concern and, at present, is more (black) art than science. We adopt the notion of the "natural budget constraint" introduced by Fortnow and Sami [9]: the agent is allowed only those trades for which the maximum loss for any possible outcome does not exceed the budget.

The main contribution of this paper is a rich, geometric characterization of the impact of budget constraints. In particular we consider, for a fixed belief, the locus of the resulting price of an optimal trade as a function of the budget. We show that the price moves in the convex hull of the belief and the set of tight outcomes, in a direction that is "perpendicular" to the set of tight outcomes. We also introduce the concept of *budget additivity*: two agents with budgets B and B' and the same beliefs have the same power to move the price as a single agent with the same belief and budget $B + B'$. An absence of budget additivity points to an inefficiency in incorporating information from the traders. We show that budget additivity is a non-trivial property by giving examples of market makers that do not satisfy budget additivity. We give a set of sufficient conditions on the market maker and the set of securities offered which guarantee budget additivity. Further, for two of the most commonly used market scoring rules (quadratic score and log score), we show sufficient conditions on the set of securities that guarantee budget additivity.

Previously, Fortnow and Sami [9] considered a different question: do budget-constrained bidders always move the market price in the direction of their beliefs? They showed that the answer to this is no, that there always exist market prices, beliefs and budgets such that the direction of price movement is not towards the belief. We give a richer characterization of how the market price moves in the presence of budget constraints, by charting the path the price takes with increasing budgets. Fortnow and Sami's [2012] impossibility result can be easily derived from our characterization (see Appendix D).

A designer of a prediction market has a lot of freedom but little guidance, and our results can be used both descriptively and prescriptively. As a descriptive tool, our results enable us to analyze commonly used market makers and understand if budget constraints hamper information aggregation in these markets. As a prescriptive tool, our results can be used to construct markets that are budget additive. In particular, we speculate that budget additivity simplifies the choice of the liquidity parameter in the markets, because it allows considering trader budgets in aggregate.

Proof overview and techniques. Our analysis borrows heavily from techniques in convex analysis and builds on the notion of Bregman divergence. We use the special case of Euclidean distance (corresponding to a quadratic-cost market maker) to form our geometric intuition which we then extend to arbitrary Bregman divergences. For the sake of an example, consider a complete market over a finite set of outcomes where the space of market prices is a simplex, exactly coinciding with the space of probability distributions over outcomes. Every possible outcome imposes a constraint on the set of prices that a trader can move the market to, because the trader is not allowed to exceed the budget if that outcome occurs. The prices satisfying this constraint form a ball with the outcome at its center. The set of feasible prices that can be moved to is therefore the intersection of these balls.

The key structural result we obtain is the chart of the price movement. Suppose that there is an infinite sequence of agents with infinitesimally small budgets all with the same belief, then what is the path along which the price moves? This is determined by the agents' belief and the set of budget constraints that are tight at any point, corresponding to the highest risk outcomes (outcomes with the highest potential loss). We show that the price can always be written as a convex combination of these highest risk outcomes and the agents' belief. Further, the market price moves in a direction that is "perpendicular" to the affine space of these outcomes.

The agents' belief partitions the simplex interior into regions, where each region is the interior of the convex hull of the agent belief and a particular subset of outcomes. For a region that is full-dimensional, every interior point can be uniquely written as a convex combination of the agent belief and all except one outcome. In the anticipation of the further development, we call this outcome "profitable" and others "risky". Motivated by the characterization above, we move perpendicular to the risky outcomes in the direction towards the agents' belief. As a result, we increase the risk of risky outcomes (equally for all outcomes), while getting closer to the one profitable outcome (and hence increasing its profit). The characterization then guarantees that the prices along this path are indeed those chosen by traders at increasing budgets, because the risky outcomes yield tight constraints.

We would like the same to be true for the lower dimensional regions as well; that is, for the set of tight constraints to be exactly the corresponding set of outcomes defining the convex hull. In fact, this is the property that distinguishes whether budget additivity is satisfied. The markets for which the tight constraints are exactly the minimal set of outcomes that define the region the price lies in are budget additive. (We conjecture that the converse holds as well.) The entire path is then as follows: w.l.o.g. you start at a full-dimensional region, move along the perpendicular until you hit the boundary of the region and you are in a lower-dimensional region, move along the perpendicular in this lower-dimensional region, and so on until you reach the belief. The set of tight constraints is monotonically decreasing. We show that such markets are characterized by a certain *acute angles assumption* on the set of possible outcomes. Loosely speaking, this assumption guarantees that outcomes outside the minimal set behave as the profitable outcome in the above example.

Other related work. There is a rich literature on scoring rules and prediction markets. Two of the most studied scoring rules are the quadratic scoring rule [4] and the logarithmic market scoring rule [14]. We consider cost-function-based prediction markets, a fully general class under reasonable assumptions [1, 6]. Hanson [13] and Chen and Pennock [5] also studied cost-function-based market makers. Gneiting and Raftery [11] implicitly noted the equivalence between proper scoring rules and convex-cost prediction markets. Beygelzimer et al. [2] and Frongillo et al. [10] study the relationship between utility functions and price dynamics in prediction markets, drawing a parallel to online learning. (6, draw a different parallel between prediction markets and online learning.) Finally, our analysis touches on the problem of setting the market maker's liquidity parameter [16, 18], which determines how (in)sensitive prices are to trading volume. With budget additivity, the market designer can optimize liquidity according to aggregate budgets, without worrying about how budgets are partitioned among traders.

2 Preliminaries

Securities and payoffs. Consider a probability space with a finite set of outcomes $\Omega \subseteq \mathbb{R}^n$. A *security* is a financial instrument whose payoff depends on the realization of an outcome in Ω. In other words, the payoff of a security is a random variable of the probability space. We consider trading with n securities corresponding to n coordinates of the outcomes $\omega \in \Omega$. A security can be traded before the realization is observed with the intention that the price of a security serves as a prediction for the expected payoff, i.e., the expected value of the corresponding coordinate.

Cost function, prices and utilities. An *automated market maker* always offers to trade securities, for the right price. In fact the price is the current prediction of the market maker for the expectation of ω. A cost function based market maker is based on a differentiable convex *cost function*, $C : \mathbb{R}^n \to \mathbb{R}$. It is a scalar function of an n-dimensional vector $q \in \mathbb{R}^n$ representing the *number of outstanding shares*[2] for our n securities. We also refer to q as the *state* of the market.

The vector of *instantaneous prices* of the securities is simply the gradient of C at q, denoted by $p(q) := \nabla C(q)$. The prices of securities change continuously as the securities are traded, so it is useful to consider the cost of trading a given quantity of securities. The cost of buying $\delta \in \mathbb{R}^n$ units of securities (where a negative value corresponds to selling) is determined by the path integral $\int_\pi p(\bar{q}) \cdot d\bar{q} = C(q+\delta) - C(q)$, where π is any smooth curve from q to $q+\delta$.

When the outcome ω is realized, the vector of δ units of securities pays off an amount of $\delta \cdot \omega$. Thus, the realized utility of a trader whose trade δ moved the market state from q to $q' = q + \delta$ is

$$U(q', \omega; q) := (q' - q) \cdot \omega - C(q') + C(q) \ .$$

We make a standard assumption that the maximum achievable utility, which is also the maximum loss of the market maker, is bounded by a finite constant (in Section 4, we show a standard approach how to easily check this). Let $\mathcal{M}$ be the convex hull of the payoff vectors, $\mathcal{M} := \operatorname{conv}(\Omega)$. It is easy to see that $\mathcal{M}$ contains exactly the vectors $\mu \in \mathbb{R}^n$ which can be realized as expected payoffs $\mathbb{E}[\omega]$ for some probability distribution over Ω. For a trader who believes that $\mathbb{E}[\omega] = \mu$, the expected utility takes form

$$U(q', \mu; q) := \mathbb{E}\left[U(q', \omega; q)\right] = (q'-q)\cdot\mu - C(q') + C(q) \ .$$

We consider throughout, a single *myopic* trader who trades as if he were the last to trade. A key property satisfied by expected utility is *path independence*: for any $q, \bar{q}, q' \in \mathbb{R}^n$, $U(q', \mu; \bar{q}) + U(\bar{q}, \mu; q) = U(q', \mu; q)$, that is, risk-neutral traders have no incentive to split their trades. For a risk-neutral trader, $q' \in \mathbb{R}^n$ is an optimal action if and only if $\mu = \nabla C(q') = p(q')$ (this follows from the first-order optimality conditions). In other words, the trader is incentivized to move the market to the prices corresponding to his belief as long as such prices exist. In general, there may be multiple states yielding the same prices, so the inverse map $p^{-1}(\mu)$ returns a *set*, which can be empty if no state yields the price vector μ.

Commonly-used cost functions include the quadratic cost, logarithmic market-scoring rule (LMSR) and the log-partition function. They are described in detail in Appendix A. The quadratic cost is defined by $C(q) = \frac{1}{2}\|q\|_2^2$ and $p(q) = q$. Log-partition function is defined as $C(q) = \ln(\sum_{\omega\in\Omega} e^{q\cdot\omega})$. It subsumes LMSR as a special case for the *complete market* with the outcomes corresponding to

[2] We allow trading fractions of a security. Negative values correspond to short-selling.

vertices of the simplex. The prices under log-partition cost correspond to the expected value of ω under the distribution $P_q(\omega) = e^{q\cdot\omega - C(q)}$ over Ω, i.e., $p(q) = \mathbb{E}_{P_q}[\omega]$.

Budget constraints. Trading in prediction markets needs an investment of capital. It is possible that an agent loses money on the trade, in particular $U(q', \omega; q)$ could be negative for some ω. One restriction on how an agent trades could be that he is unable to sustain a big loss, due to a budget constraint. We consider the notion of *natural budget constraint* defined by Fortnow and Sami [9] which states that the loss of the agent is at most his budget, for all $\omega \in \Omega$. Given a starting market state q_0 and a budget of $B \geq 0$, a trader with the belief $\mu \in \mathcal{M}$ then solves the problem:

$$\max_{q\in\mathbb{R}^n} U(q, \mu; q_0) \text{ s.t. } U(q, \omega; q_0) \geq -B \quad \forall \omega \in \Omega \ . \tag{2.1}$$

For quadratic costs, each constraint corresponds to a sphere with one of the outcomes at its center, so the feasible region is an intersection of these spheres. We will later see that this generalizes to an intersection of balls w.r.t a Bregman divergence for general costs.

In general, there may be multiple q optimizing this objective. In the following definition we introduce notation for various solution sets we will be analyzing. The belief μ is fixed throughout most of the discussion, so we suppress the dependence on μ.

Definition 2.1 (Solution sets)**.** Let $\hat{Q}(B; q_0)$ denote the set of solutions of Convex Program (2.1) for a fixed initial state and budget. Let $\hat{Q}(q_0) = \bigcup_{B\geq 0} \hat{Q}(B; q_0)$ denote the set of solutions of (2.1) for a fixed initial state across all budgets. Let $\hat{Q}(\nu; q_0) = p^{-1}(\nu) \cap \hat{Q}(q_0)$ denote the set of states q that optimize (2.1) for some budget B and yield the market price vector ν.

The next theorem shows that solutions for a fixed initial state and budget always yield the same price vector. It is proved in Appendix B.

Theorem 2.2. *If $q, q' \in \hat{Q}(B; q_0)$, then $p(q) = p(q')$.*

Geometry of linear spaces. We finish this section by reviewing a few standard geometric definitions we use in next sections. Let $X \subseteq \mathbb{R}^n$. Then $\mathrm{aff}(X)$ denotes the *affine hull* of the set X (i.e., the smallest affine space including X). We write $X^\perp$ to denote the *orthogonal complement* of X: $X^\perp := \{u \in \mathbb{R}^n : \ u \cdot (x' - x) = 0 \text{ for all } x, x' \in X\}$. We use the convention $\emptyset^\perp = \mathbb{R}^n$. A set $\mathcal{K} \in \mathbb{R}^n$ is called a *cone* if it is closed under multiplication by positive scalars. If a cone is convex, it is also closed under addition. Since Ω is finite, the realizable set $\mathcal{M} = \mathrm{conv}(\Omega)$ is a polytope. Its boundary can be decomposed into *faces*. More precisely, $X \subseteq \Omega$, $X \neq \emptyset$, forms a *face* of $\mathcal{M}$ if X is the set of maximizers over Ω of some linear function.[3] We also view $X = \emptyset$ as a face of $\mathcal{M}$. With this definition, for any two faces X, X', also their intersection $X \cap X'$ is a face.

3 Characterization of Solution Sets

We start with the optimality (KKT) conditions for the Convex Program (2.1), as characterized by the next lemma. One of the key conditions is that the solution prices must be in the convex hull of the belief μ and all the ω's for which the budget constraints are *tight*. The set of tight constraints is always a face of the polytope $\mathcal{M}$. We allow an empty set as a face, which corresponds to the case when none of the constraints are tight and the solution prices coincide with μ. The proof follows by analyzing KKT conditions (see Appendix C).

Lemma 3.1 (KKT lemma)**.** *Let $q_0 \in \mathbb{R}^n$. Then $q \in \hat{Q}(B; q_0)$ if and only if there exists a face $X \subseteq \Omega$ such that the following conditions hold:*

(a) *$U(q, x; q_0) = U(q, x'; q_0)$, or equivalently $(q - q_0) \cdot (x' - x) = 0$, for all $x, x' \in X$*
(b) *$U(q, \omega; q_0) \geq U(q, x; q_0)$, or equivalently $(q - q_0) \cdot (\omega - x) \geq 0$, for all $x \in X$ and $\omega \in \Omega \backslash X$*
(c) *$p(q) \in \mathrm{conv}(X \cup \{\mu\})$*
(d) *$B = -U(q, x; q_0)$ for some $x \in X$ if $X \neq \emptyset$, or $B \geq \max_{\omega\in\Omega}[-U(q, \omega; q_0)]$ if $X = \emptyset$*

where conditions (a) and (b) hold vacuously for $X = \emptyset$.

The condition (a) requires that $q - q_0$ be orthogonal to the active set X. The set of points satisfying conditions (a) and (c) will be called the Bregman perpendicular and will be defined in the next section. The condition (b) is a statement about acuteness of the angle between $q - q_0$ (the perpendicular) and the outcomes. It will be the basis of our *acute angles* assumption. The condition (d) just states how the budget is related to the active set X.

Witness cones and minimal faces. We now introduce a few notations which allow us to state reinterpretations of the conditions in Lemma 3.1. First of all, given a face X, what are the set of q's that satisfy conditions (a) and (b)? This is captured by what we call the *witness cone*.

Definition 3.2. Given a face $X \subseteq \Omega$, the *witness cone* for X is defined as $\mathcal{K}(X) := \{u \in \mathbb{R}^n : \ u \cdot (\omega - x) \geq 0 \text{ for all } x \in X, \omega \in \Omega\}$ if $X \neq \emptyset$, and $\mathcal{K}(X) := \mathbb{R}^n$ if $X = \emptyset$.

The following two properties of witness cones are immediate from the definition:

- *Anti-monotonicity:* if $X \subseteq X' \subseteq \Omega$, then $\mathcal{K}(X) \supseteq \mathcal{K}(X')$.
- *Orthogonality:* $\mathcal{K}(X) \subseteq X^\perp$.

[3] Strictly speaking, this is the definition of an *exposed face*, but all faces of a polytope are exposed, so the distinction does not matter here. The exposed face is typically defined to be $\mathrm{conv}(X)$, but in the present paper, it is more convenient to work with X directly.

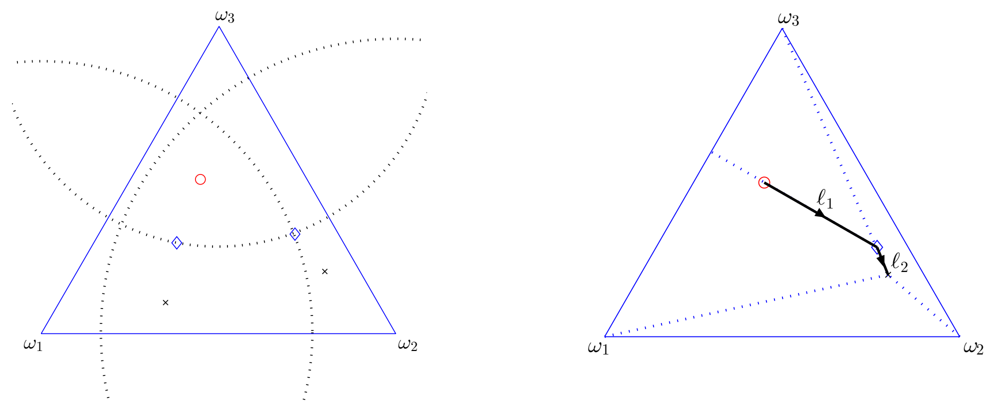

Figure 1: *Left:* $\circ$ —current state, $\times$ —belief, $\diamond$ —optimal action for a given belief and budget. Three circles bound the allowed final states for budget 0.1. We plot optimal actions for two different beliefs. *Right:* A path from the initial state to the belief, consisting of optimal actions for increasing budgets.

A state q satisfies conditions (a) and (b) for a given face X if and only if $q - q_0 \in \mathcal{K}(X)$. Now given a state q, consider the set of faces that could satisfy condition (c). This set has a useful structure, namely that there is a unique minimal face (proved in Appendix C).

Definition 3.3. Given a price vector $\nu \in \mathcal{M}$, the *minimal face* for ν is the minimal face X (under inclusion) s.t. $\nu \in \operatorname{conv}(X \cup \{\mu\})$. The minimal face for ν is denoted as X_ν.

With the existence of a minimal face and the anti-monotonicity of the witness sets, it follows that if q and X satisfy conditions (a), (b) and (c), then so do q and $X_{p(q)}$. Thus we obtain the following version of Lemma 3.1 (proved in Appendix C).

Theorem 3.4 (Characterization of Solution Sets)**.** *$q \in \hat{Q}(q_0)$ if and only if $q \in [q_0 + \mathcal{K}(X_{p(q)})]$.*

Using Theorem 3.4, we immediately obtain a characterization of when a price vector ν could be the price vector of an optimal solution to (2.1).

Corollary 3.5. *$\hat{Q}(\nu; q_0) = p^{-1}(\nu) \cap [q_0 + \mathcal{K}(X_\nu)]$. In particular, ν is the price vector of an optimal solution to (2.1) if and only if $p^{-1}(\nu) \cap [q_0 + \mathcal{K}(X_\nu)] \neq \emptyset$.*

We now study an example using the above characterization. More examples can be found in Appendix E.

Example 3.6 (Quadratic cost on an obtuse triangle; see Example E.2 for details)**.** Consider the following outcome space, belief, and the sequence of market states (depicted in Figure 2):

$$
\begin{aligned}
\omega_1 &= (0.0,\ 0.0) & q_0 &= \nu_0 = \tfrac{11}{14}\omega_2 + \tfrac{3}{14}\omega_3 \\
\omega_2 &= (1.8,\ 0.0) & q_1 &= \nu_1 = \tfrac{1}{3}\omega_2 + \tfrac{2}{3}\mu \\
\omega_3 &= (6.0,\ 4.2) & q_2 &= \nu_2 = \tfrac{1}{9}\omega_1 + \tfrac{8}{9}\mu \\
\mu = q_\mu &= (2.7,\ 1.8) & q_3 &= \nu_3 \approx \tfrac{1}{19}\omega_1 + \tfrac{18}{19}\mu
\end{aligned}
$$

Using the KKT lemma, we can show for $j = 1, 2, 3$, that $q_j = \nu_j$ is an optimal action at $q_{j-1} = \nu_{j-1}$ under belief μ, with the corresponding budgets as:

	ω_1	ω_2	ω_3	
$U(q_1, \cdot; q_0)$	0.45	-0.09	-0.09	$B_{01} = 0.09$
$U(q_2, \cdot; q_1)$	-0.56	-0.56	1.12	$B_{12} = 0.56$
$U(q_3, \cdot; q_2)$	-0.565	$-0.28\ldots$	$0.82\ldots$	$B_{23} = 0.565$
$U(q_\mu, \cdot; q_0)$	-1.215	-1.215	2.565	$B_{0\mu} = 1.215$

The above table also shows that the budget $B_{0\mu} = 1.215$ suffices to move directly from q_0 to q_μ. However, note that the sum $B_{01} + B_{12} + B_{23} = 1.215 = B_{0\mu}$, but $\nu_3 \neq \mu$, i.e., after the sequence of optimal actions with budgets B_{01}, B_{12}, and B_{23}, the market is still not at the belief shared by all agents, even though with the budget $B_{0\mu}$, it would have reached it.

Budget additivity. The above example suggests that multiple traders with the same belief may have less power in moving the market state towards the their belief comparing to a single trader with the same belief and the combined budget. Since prediction markets aim to efficiently aggregate information from agents, it is natural to ask under what conditions multiple traders with the same beliefs do have a combined impact equal to a single trader with the combined budget.

Next, we formally define this property as *budget additivity*. We then define the Euclidean version of the acute angles condition that we show is sufficient for budget additivity.

Definition 3.7 (Budget additivity)**.** We say that a prediction market is budget additive on $\mathcal{M}' \subseteq \mathcal{M}$ if for all beliefs $\mu \in \mathcal{M}'$ and all initial states $q_0 \in p^{-1}(\mathcal{M}')$ the following holds: For any budgets $B, B' \geq 0$ and any sequence of solutions $q \in \hat{Q}(B; q_0)$ and $q' \in \hat{Q}(B'; q)$, we have $p(q), p(q') \in \mathcal{M}'$ and $q' \in \hat{Q}(B + B'; q_0)$.

In other words, the market is budget additive if the sequence of optimal actions of two agents with the same belief and budgets B and B' is also an optimal action of a single agent with the same belief and a larger budget

$B + B'$. Thanks to Theorem 2.2 we also obtain that the market prices following the sequence of optimal actions by the two agents are the same as the market prices after the optimal action by an agent with the combined budget (all with the same beliefs).

We now state the acute angles assumption for the Euclidean case, to give an intuition. Our acute angles assumption (Definition 5.1) is a generalization of this. We later show that the acute angles property is sufficient for budget additivity (Theorem 5.2).

Definition 3.8. We say that the *Euclidean acute angles* hold for a face X of $\mathcal{M}$, if the angle between any point $\bar{\nu} \in \mathcal{M}$, its projection on the affine hull of X and any payoff $\omega \in \Omega$ is non-obtuse (the angle is measured at the projection).

Based on the above example, one may hypothesize that the obtuse angles are to blame for the lack of budget additivity. In the following sections we will show that this is indeed the case, but that the notion of obtuse/acute angles depends on the Bregman divergence. In particular, the above example would have been budget-additive if we used the log-partition cost instead of the quadratic cost.

4 Convex conjugacy, Bregman divergence and perpendiculars

We will see next that the utility function U can be written as the difference of two terms measuring the distance between the belief and the market state before and after the trade. This distance measure is the mixed *Bregman divergence*.[4] To define the Bregman divergence, first let $C^* : \mathbb{R}^n \to (-\infty, \infty]$ be the convex conjugate of C defined as $C^*(\nu) := \sup_{q' \in \mathbb{R}^n} [q' \cdot \nu - C(q')]$. Since C^* is a supremum of linear functions, it is convex lower semi-continuous. Up to a constant, it characterizes the maximum achievable utility on an outcome ω for a fixed initial state q as $\sup_{q' \in \mathbb{R}^n} U(q', \omega; q) = C^*(\omega) + \big[C(q) - q \cdot \omega\big]$. The term in the brackets is always finite, but C^* might be positive infinite. We make a standard assumption that $C^*(\omega) < \infty$ for all $\omega \in \Omega$, i.e., that the maximum achievable utility, which is also the maximum loss of the market maker, is bounded by a finite constant. By convexity, this implies that $C^*(\mu) < \infty$ for all $\mu \in \mathcal{M}$. The *Bregman divergence* derived from C is a function $D : \mathbb{R}^n \times \mathbb{R}^n \to (-\infty, \infty]$ measuring the maximum expected utility under belief μ at a state q

$$D(q, \mu) := C(q) + C^*(\mu) - q \cdot \mu = \sup_{q' \in \mathbb{R}^n} U(q', \mu; q) \ .$$

From the convexity of C and C^* and the definition of C^*, it is clear that: (i) D is convex and lower semi-continuous in each argument separately; (ii) D is non-negative; and (iii) D is zero iff $p(q) = \nabla C(q) = \mu$. By the bounded loss assumption, Bregman divergence is finite on $\mu \in \mathcal{M}$. For $\mu \in \mathcal{M}$, we can write

$$U(q', \mu; q) = D(q, \mu) - D(q', \mu) \ . \tag{4.1}$$

Thus, maximizing the expected utility is the same as minimizing the Bregman divergence between the state q' and the belief μ. From Eq. (4.1) it is also clear that each constraint in (2.1) is equivalent to $D(q, \omega) \le D(q_0, \omega) + B$, and the geometric interpretation is that the agent seeks to find the state closest to his belief, within the intersection of Bregman balls

For the quadratic cost, we have $C^*(\nu) = \frac{1}{2}\|\nu\|^2$ and $D(q, \nu) = \frac{1}{2}\|q - \nu\|^2$, i.e., the Bregman divergence coincides with the Euclidean distance squared. For log-partition cost, we have $C^*(\nu) = \sum_{\omega \in \Omega} P_\nu(\omega) \ln P_\nu(\omega)$ where P_ν is the distribution maximizing entropy among P satisfying $\mathbb{E}_P[\omega] = \nu$. The Bregman divergence is the KL-divergence between P_q and P_ν: $D(q, \nu) = \mathrm{KL}(P_\nu \| P_q)$.

Convex analysis. We overview a few standard definitions and results from convex analysis. For $X \subseteq \mathbb{R}^n$, we write $\operatorname{ri} X$ for the *relative interior* of X (i.e., the interior relative to the affine hull). For a convex function $F : \mathbb{R}^n \to (-\infty, \infty]$, we define its *effective domain* as $\operatorname{dom} F := \{u \in \mathbb{R}^n : F(u) < \infty\}$ (i.e., the set of points where it is finite). The *subdifferential* of F at a point u is the set $\partial F(u) := \{v \in \mathbb{R}^n : F(u') \ge F(u) + (u' - u) \cdot v \text{ for all } u' \in \mathbb{R}^n\}$. We say that F is *subdifferentiable* at u if $\partial F(u) \neq \emptyset$. A standard result of convex analysis states that F is always subdifferentiable on a superset of $\operatorname{ri} \operatorname{dom} F$. If F is not only convex, but also lower semi-continuous, then ∂F and ∂F^* are inverses in the sense that $v \in \partial F(u)$ iff $u \in \partial F^*(v)$. If F is *differentiable* everywhere on $\mathbb{R}^n$, then F^* is strictly convex on $\operatorname{ri} \operatorname{dom} F^*$.

Let $\operatorname{im} p := \{p(q) : q \in \mathbb{R}^n\}$ denote the set of prices that can be expressed by market states. The implications for our setting are that: (i) C^* is subdifferentiable on $\operatorname{im} p$; (ii) $p^{-1}(\nu) = \partial C^*(\nu)$ for all $\nu \in \mathbb{R}^n$; (iii) all beliefs in $\operatorname{ri} \operatorname{dom} C^*$ can be expressed by some state q; (iv) C^* is strictly convex on $\operatorname{ri} \operatorname{dom} C^*$, and similarly $D(q, \nu)$ is strictly convex on $\operatorname{ri} \operatorname{dom} C^*$ as a function of the second argument.

Assumptions on the cost function.

- *Convexity and differentiability on* $\mathbb{R}^n$. C is convex and differentiable on $\mathbb{R}^n$.
- *Finite loss.* $\mathcal{M} \subseteq \operatorname{dom} C^*$, i.e., C^* is finite on $\mathcal{M}$.
- *Inclusion of the relative interior.* $\operatorname{ri} \mathcal{M} \subseteq \operatorname{ri} \operatorname{dom} C^*$.

The first two assumptions are standard. The third assumption is a regularity condition that we require in our results. Here we briefly discuss how it compares with the

[4] Our notion of Bregman divergence is more general than typically assumed in the literature.

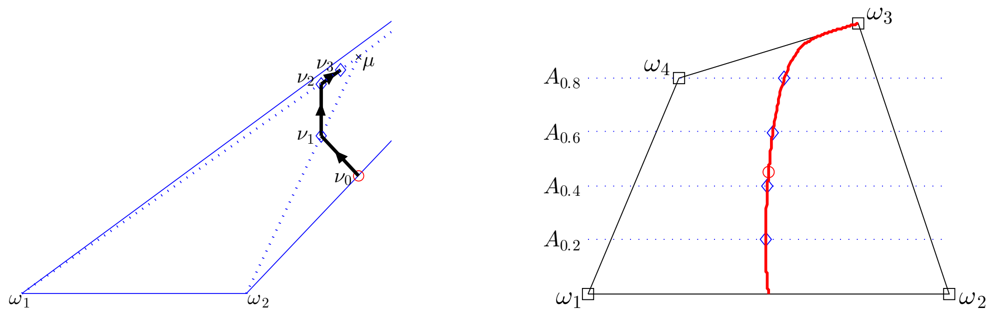

Figure 2: *Left:* An example of non-additive budgets when payoffs form obtuse angles (see Example E.2). *Right:* An examples of a non-linear perpendicular for the log-partition cost.

finite loss assumption. While the two assumptions look similar, neither of them implies the other. For example, if $\operatorname{dom} C^*$ is an n-dimensional simplex and $\mathcal{M}$ is one of its lower dimensional faces, which are lower dimensional simplices, then the finite loss assumption holds, but the inclusion assumption does not. Similarly, for $n = 1$ and $\mathcal{M} = [0, 1]$, the inclusion assumption is satisfied by the conjugate $C^*(\nu) = 1/\nu + 1/(1 - \nu)$ on $\nu \in (0, 1)$ and $C^*(\nu) = \infty$ on $\nu \notin (0, 1)$, but this conjugate does not satisfy the finite loss assumption.

We do not view the inclusion assumption as very restrictive, since it is satisfied by many common cost functions. For instance, it always holds when C is constructed as in [1], because their construction guarantees $\operatorname{dom} C^* = \mathcal{M}$. However, the inclusion assumption might not hold for cost functions that allow arbitrage (e.g., 7).

Our main result relies on strict convexity of C^* on $\operatorname{ri} \operatorname{dom} C^*$, so some of our statements will require that the market prices and beliefs lie in that set. The inclusion assumption above guarantees that at the minimum $\operatorname{ri} \mathcal{M} \subseteq \operatorname{ri} \operatorname{dom} C^*$, but the boundary of $\mathcal{M}$ might not be included. With this motivation in mind, we define the set

$$\tilde{\mathcal{M}} := \begin{cases} \mathcal{M} & \text{if } \mathcal{M} \subseteq \operatorname{ri} \operatorname{dom} C^* \\ \operatorname{ri} \mathcal{M} & \text{otherwise.} \end{cases}$$

In either case we obtain that $\tilde{\mathcal{M}} \subseteq \operatorname{ri} \operatorname{dom} C^* \subseteq \operatorname{im} p$, i.e., beliefs in $\tilde{\mathcal{M}}$ can be expressed by some state q. For the quadratic cost, $\tilde{\mathcal{M}} = \mathcal{M}$. For the log-partition cost, $\tilde{\mathcal{M}} = \operatorname{ri} \mathcal{M}$.

Perpendiculars. We now define the notion of a Bregman perpendicular to an affine space. This is a *constructive* definition, and it plays a central role in the definition of the acute angles assumption, and also in the proof of the main result. The set of optimal price vectors for different budgets will be a sequence of Bregman perpendiculars. Naturally, it is also closely related to the conditions in Lemma 3.1; in particular to the set of q's that satisfy conditions (a) and (c) for a given face X.

For quadratic costs, this notion coincides with the usual Euclidean perpendicular. Consider an affine space and a point not in it. A projection of the point onto the space is the point in the space that is closest in Euclidean distance to the given point. Now consider moving this affine space towards the projected point. The locus of the projection as we move the space is the perpendicular to the space through the given point. We extend this definition to arbitrary Bregman divergences by defining the projection using the corresponding Bregman divergence.

A Bregman perpendicular is determined by three geometric objects within the affine hull $\operatorname{aff}(\operatorname{dom} C^*)$. The first of these is an affine space, say $A_0 \subseteq \operatorname{aff}(\operatorname{dom} C^*)$. The second is a point $a_1 \in \operatorname{aff}(\operatorname{dom} C^*) \backslash A_0$. The affine space $A = \operatorname{aff}(A_0 \cup \{a_1\}) \subseteq \operatorname{aff}(\operatorname{dom} C^*)$ will be the ambient space that will contain the perpendicular. Define parallel spaces to A_0 in A, for an arbitrary point $a_0 \in A_0$, as $A_\lambda := A_0 + \lambda(a_1 - a_0)$ for $\lambda \in \mathbb{R}$. Note that the definition of A_λ is independent of the choice of a_0. The third geometric object is a market state $q \in \mathbb{R}^n$ such that $p(q) \in A$. For technical reasons, we will define a perpendicular at q rather than a more natural notion, which would be at $p(q)$. Our reason for switching into q-space is that inner products, defining optimality of the Bregman projection, are between elements of q-space and ν-space (the two spaces coincide for Euclidean distance). For all $\lambda \in \mathbb{R}$ define a Bregman projection of q onto A_λ as

$$\nu_\lambda := \operatorname*{argmin}_{\nu \in A_\lambda} D(q, \nu) \ .$$

Since $D(q, \nu)$ is bounded from below and lower semicontinuous, the minimum is always attained (but it may be equal to ∞). If it is attained at more than one point, we choose an arbitrary minimizer. Whenever we can choose $\nu_\lambda \in \operatorname{ri} \operatorname{dom} C^*$, this ν_λ must be the unique minimizer by strict convexity of $D(q, \cdot)$ on $\operatorname{ri} \operatorname{dom} C^*$, and the minimum is finite. We use these ν_λ's to define the perpendicular:

Definition 4.1. Given A_0, a_1 and q as above, the a_1-*perpendicular* to A_0 at q is a map $\gamma : \lambda \mapsto \nu_\lambda$ defined over $\lambda \in \Lambda := \{\lambda : \nu_\lambda \in \operatorname{ri} \operatorname{dom} C^*\} \subseteq \mathbb{R}$. We call Λ

the domain of the perpendicular. We define a total order on $\nu_\lambda, \nu_{\lambda'} \in \operatorname{im}\gamma$ as $\nu_\lambda \preceq \nu_{\lambda'}$ iff $\lambda \le \lambda'$.

In Appendix F.2, we show that perpendiculars are continuous maps. The name perpendicular is justified by the following proposition which matches our Euclidean intuition that the perpendiculars can be obtained by intersecting the ambient space A with the affine space which passes through q and is orthogonal to A_0. It also shows that the perpendicular corresponds to the set of prices that satisfy conditions (a) and (c) with the convex hull relaxed to the affine hull (when A_0 is the affine hull of face X, point a_1 coincides with μ and q is the initial state). Recall that for an arbitrary set $X \subseteq \mathbb{R}^n$, its orthogonal complement is defined as $X^\perp := \{u : \ u \cdot (x' - x) = 0 \text{ for all } x, x' \in X\}$.

Proposition 4.2. *Let γ be the a_1-perpendicular to A_0 at q, and let $A = \operatorname{aff}(A_0 \cup \{a_1\})$. The following two statements are equivalent for any $\nu' \in \mathbb{R}^n$:*

(i) $\nu' \in \operatorname{im}\gamma$
(ii) $\nu' \in A \cap (\operatorname{ri}\operatorname{dom} C^*)$, $p^{-1}(\nu') \cap (q + A_0^\perp) \neq \emptyset$

Proposition 4.2 is proved in Appendix F. The perpendiculars have the following closure property which is useful for showing budget additivity (also proved in Appendix F):

Proposition 4.3. *Under the assumptions of Proposition 4.2, γ is also the a_1-perpendicular to A_0 at any $q' \in p^{-1}(\operatorname{im}\gamma) \cap (q + A_0^\perp)$.*

5 Acute angles and budget additivity

We now state the *acute angles* property which links the Bregman perpendicular and Corollary 3.5, and is sufficient for budget additivity.

Definition 5.1. We say that the *acute angles* hold for a face X, if for every μ-perpendicular γ to X at q, such that $\mu \in \tilde{\mathcal{M}}$ and $q \in p^{-1}(\tilde{\mathcal{M}})$, the following holds: If $\nu' \in \operatorname{im}\gamma$ and $\nu' \succeq p(q)$, then $p^{-1}(\nu') \cap [q + \mathcal{K}(X)] \neq \emptyset$.

The motivation for the name "acute angles" comes from the Euclidean distance case, where this assumption is equivalent to Definition 3.8 (see Proposition G.1). The acute angles property is non-trivial and we have seen that without this property, budget additivity need not hold; we conjecture that it is also a necessary condition. After stating the main theorem, we analyze in more detail when the acute angles are satisfied by the quadratic and log-partition costs.

We now state the main result, that the acute angles are sufficient for budget additivity:

Theorem 5.2 (Sufficient conditions for budget additivity)**.** *If acute angles hold for every face $X \subseteq \Omega$, then the prediction market is budget additive on $\tilde{\mathcal{M}}$.*

Sufficient conditions for acute angles. We next give the sufficient conditions when the acute angles hold for the quadratic and log-partition cost functions. We also show that the acute angles hold for all one-dimensional outcome spaces, and that the are preserved by taking *direct sums* of markets. Recall that a set $\mathcal{K} \in \mathbb{R}^n$ is called a cone if it is closed under multiplication by positive scalars. A cone is called *acute*, if $x \cdot y \ge 0$ for all $x, y \in \mathcal{K}$. An *affine cone* with the vertex a_0 is a set $\mathcal{K}'$ of the form $a_0 + \mathcal{K}$ where $\mathcal{K}$ is a cone.

Theorem 5.3 (Sufficient condition for quadratic cost)**.** *Let X be a face and A' be the affine space $a_0 + X^\perp$ for an arbitrary $a_0 \in \operatorname{aff}(X)$. Acute angles hold for the face X and the quadratic cost if and only if the projection of Ω (or, equivalently, $\mathcal{M}$) on A' is contained in an affine acute cone with the vertex a_0.*

Corollary 5.4. *Acute angles hold for the quadratic cost and a hypercube $\Omega = \{0,1\}^n$.*

Corollary 5.5. *Acute angles hold for the quadratic cost and simplex $\Omega = \{e_i : \ i \in [n]\}$ where $[n] = \{1, 2, \ldots, n\}$ and e_i is the i-th vector of the standard basis in $\mathbb{R}^n$.*

Theorem 5.6 (Log-partition over affinely independent outcomes)**.** *If the set Ω is affinely independent then acute angles assumption is satisfied for the log-partition cost.*

Theorem 5.7 (One-dimensional outcome spaces)**.** *Acute angles hold for any cost function if $\mathcal{M}$ is a line segment.*

Let $\Omega_1 \subseteq \mathbb{R}^{n_1}$ and $\Omega_2 \subseteq \mathbb{R}^{n_2}$ be outcome spaces with costs C_1 and C_2. We define the *direct sum* of Ω_1 and Ω_2 to be the outcome space $\Omega = \Omega_1 \times \Omega_2$ with the cost $C : \mathbb{R}^{n_1+n_2} \to \mathbb{R}$ defined as $C(q_1, q_2) = C_1(q_1) + C_2(q_2)$.

Theorem 5.8 (Acute angles for direct sums)**.** *If acute angles hold for Ω_1 with cost C_1, and Ω_2 with cost C_2, then they also hold for their direct sum.*

As a direct consequence of this theorem, we obtain that the log-partition cost function satisfies the acute angles assumption on a hypercube. Also, any direct sum of costs on line segments satisfies the acute angles.

5.1 Proof of the suffucient conditions for budget additivity

In this section we sketch the proof of Thm. 5.2 (for a complete proof see Appendix H). We proceed in several steps. Let $\nu_0 = p(q_0)$. We begin by constructing an oriented curve L joining ν_0 with μ, by sequentially choosing portions of perpendiculars for monotonically decreasing active sets. We then show that budget additivity holds for any solutions with prices in L, and finally show that the curve L is the locus of the optimal prices of solutions $\hat{Q}(q_0)$, as well as optimal prices of solutions $\hat{Q}(q)$ for any $q \in \hat{Q}(q_0)$.

Part 1: Construction of the solution path L. In this part, we construct:

- a sequence of prices $\nu_0, \nu_1, \ldots, \nu_k$ with $\nu_0 = p(q_0)$ and $\nu_k = \mu$

- a sequence of oriented curves $\ell_0, \ldots, \ell_{k-1}$ where each ℓ_i goes from ν_i to ν_{i+1}
- a monotone sequence of sets $\Omega \supseteq X_0 \supset X_1 \supset \cdots \supset X_k = \emptyset$, such that the following *minimality* property holds: X_i is the minimal face for all $\nu \in (\operatorname{im} \ell_i) \backslash \{\nu_{i+1}\}$ for $i \leq k-1$, and X_k is the minimal face for ν_k.
- a sequence of states $q_1, \ldots, q_{k-1}$ such that $q_i \in p^{-1}(\nu_i) \cap [q_{i-1} + \mathcal{K}(X_{i-1})]$

The curves ℓ_i will be referred to as *segments*. The curve obtained by concatenating the segments ℓ_0 through ℓ_{k-1} will be called the *solution path* and denoted L. In the special case that $\nu_0 = \mu$, we have $k = 0$, $X_0 = \emptyset$ and L is a degenerate curve with $\operatorname{im} L = \{\mu\}$.

If $\nu_0 \neq \mu$, we construct the sequence of segments iteratively. Let $X_0 \neq \emptyset$ be the minimal face such that $\nu_0 \in \operatorname{conv}(X_0 \cup \{\mu\})$. By the minimality, $\mu \notin \operatorname{aff}(X_0)$. Let γ be the μ-perpendicular to $\operatorname{aff}(X_0)$ at q_0. The curve γ passes through ν_0 and eventually reaches the boundary of $\operatorname{conv}(X_0 \cup \{\mu\})$ at some ν_1 by continuity of γ. Let segment ℓ_0 be the portion of γ going from ν_0 to ν_1.

This construction gives us the first segment ℓ_0. There are two possibilities:

1. $\nu_1 = \mu$; in this case we are done;
2. ν_1 lies on a lower-dimensional face of $\operatorname{conv}(X_0 \cup \{\mu\})$; in this case, we pick some $q_1 \in p^{-1}(\nu_1) \cap [q_0 + \mathcal{K}(X_0)]$, which can be done by the acute angles assumption, and use the above construction again, starting with q_1, and obtaining a new set $X_1 \subset X_0$ and a new segment ℓ_1; and iterate.

The above process eventually ends, because with each iteration, the size of the active set decreases. This construction yields monotonicity of X_i and the minimality property.

The above construction yields a specific sequence of $q_i \in p^{-1}(\nu_i) \cap [q_{i-1} + \mathcal{K}(X_{i-1})]$. We show in Appendix H that actually $q_i \in p^{-1}(\nu_i) \cap (q_0 + X_{i-1}^{\perp})$ and that the construction of L is independent of the choice of $q_1, q_2, \ldots, q_{k-1}$.

Part 2: Budget additivity for points on L. Let $\nu, \nu' \in \operatorname{im} L$ such that $\nu \preceq \nu'$. Let $q \in \hat{Q}(\nu; q_0)$ and $q' \in \hat{Q}(\nu'; q)$ such that $q \in \hat{Q}(B; q_0)$ and $q' \in \hat{Q}(B'; q)$. In this part we show that $q' \in \hat{Q}(B + B'; q_0)$.

First, consider the case that $\nu' = \mu$. To see that $q' \in \hat{Q}(B + B'; q_0)$, first note that the constraints of Convex Program (2.1) hold, because $U(q', \omega; q_0) = U(q', \omega; q) + U(q, \omega; q_0) \geq -B' - B$ for all ω by path independence of the utility function. As noted in the introduction, in the absence of constraints, the utility $U(\bar{q}, \mu; q_0)$ is maximized at any $\bar{q}$ with $p(\bar{q}) = \mu$. Thus, q' is a global maximizer of the utility and satisfies the constraints, so $q' \in \hat{Q}(B + B'; q_0)$. If $\nu = \mu$, we must also have $\nu' = \mu$ and the statement holds by previous reasoning.

In the remainder, we only analyze the case $\nu \preceq \nu' \prec \mu$. This means that $\nu \in (\operatorname{im} \ell_i) \backslash \{\nu_{i+1}\}$ and $\nu' \in (\operatorname{im} \ell_j) \backslash \{\nu_{j+1}\}$ for $i \leq j$. By Theorem 3.4, we therefore must have $q \in [q_0 + \mathcal{K}(X_i)]$ and $q' \in [q + \mathcal{K}(X_j)]$. By anti-monotonicity of witness cones, $\mathcal{K}(X_j) \supseteq \mathcal{K}(X_i)$ and hence, $q' \in [q_0 + \mathcal{K}(X_j)]$, yielding $q' \in \hat{Q}(\nu'; q_0)$.

We now argue that the budgets add up. Let $x \in X_j \subseteq X_i$. By Lemma 3.1, we obtain that $q \in \hat{Q}(B; q_0)$ for $B = -U(q, x; q_0)$, and $q' \in \hat{Q}(B'; q)$ for $B' = -U(q', x; q)$, and finally $q' \in \hat{Q}(\bar{B}; q_0)$ for $\bar{B} = -U(q', x; q_0)$. However, by path independence of the utility function

$$\bar{B} = -U(q', x; q_0) = -U(q', x; q) - U(q, x; q_0) = B' + B \ .$$

Part 3: L as the locus of all solutions. See Appendix H for the proof that

$$\hat{Q}(q_0) = \bigcup_{\nu \in \operatorname{im} L} \hat{Q}(\nu; q_0) \ .$$

Part 3': L as the locus of solutions starting at a midpoint. Let $\nu \in \operatorname{im} L$ and $q \in \hat{Q}(\nu; q_0)$. Since $\hat{Q}(\nu; q_0) \subseteq p^{-1}(\nu) \cap (q_0 + X_\nu^{\perp})$, Part 1' yields that the solution path L' for q coincides with the portion of L starting at ν. Applying the proof of Part 3 to L', we obtain

$$\hat{Q}(q) = \bigcup_{\nu' \in \operatorname{im} L : \nu' \succeq \nu} \hat{Q}(\nu'; q) \ .$$

Part 4: Proof of the theorem. Let $B, B' \geq 0$ and $q \in \hat{Q}(B; q_0)$ and $q' \in \hat{Q}(B'; q)$. From Parts 3 and 3', we know that $q \in \hat{Q}(\nu; q_0)$ and $q' \in \hat{Q}(\nu'; q)$ for some $\nu, \nu' \in \operatorname{im} L$ such that $\nu \preceq \nu'$. By Part 2, we therefore obtain that $q' \in \hat{Q}(B + B'; q_0)$, proving the theorem.

5.2 Necessary conditions for budget additivity

We conjecture that the sufficient conditions are also necessary. We now sketch an approach to prove this and the technical difficulties with it, for the case of the quadratic scoring rule. We can show that in *2 dimensions*, if the acute angles do not hold, then we can essentially replicate Example E.2, which shows lack of budget additivity. We can also show a key technical lemma that any obtuse, pointed cone in n dimensions has a projection onto a 2 dimensional subspace such that the image of the cone is also an obtuse, pointed cone. (See Appendix I for a proof.)

Lemma 5.9. *Let $\mathcal{K}$ be an obtuse, pointed closed convex cone. Then there is a projection of $\mathcal{K}$ onto a two-dimensional subspace that is also obtuse and pointed (and automatically closed convex).*

The idea then is to embed the 2 dimensional example in this subspace and then "lift" it back to the original cone. But this runs into a difficulty, because while the lift maintains conditions (a) and (b) of Lemma 3.1, it is not clear how to simultaneously maintain the condition (c), and some additional structure may be required.

References

[1] Jacob Abernethy, Yiling Chen, and Jennifer Wortman Vaughan. An optimization-based framework for automated market-making. In *ACM Conference on Electronic Commerce*, pages 297–306, 2011.

[2] Alina Beygelzimer, John Langford, and David M. Pennock. Learning performance of prediction markets with Kelly bettors. In *International Conference on Autonomous Agents and Multiagent Systems*, pages 1317–1318, 2012.

[3] Jonathan Borwein and Adrian Lewis. *Convex Analysis and Nonlinear Optimization*. Springer Publishing Company, Incorporated, 2000.

[4] GW Brier. Verication of forecasts expressed in terms of probability. *Monthly Weather Review*, 78:13, 1950.

[5] Yiling Chen and David M. Pennock. A utility framework for bounded-loss market makers. In *UAI*, pages 49–56, 2007.

[6] Yiling Chen and Jennifer Wortman Vaughan. A new understanding of prediction markets via no-regret learning. In *ACM Conference on Electronic Commerce*, pages 189–198, 2010.

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

[8] Edmund Eisenberg and David Gale. Consensus of subjective probabilities: The pari-mutuel method. *Annals of Mathematical Statistics*, 30:165–168, 1959.

[9] Lance Fortnow and Rahul Sami. Multi-outcome and multidimensional market scoring rules. *CoRR*, abs/1202.1712, 2012.

[10] Rafael Frongillo, Nicolas Della Penna, and Mark Reid. Interpreting prediction markets: a stochastic approach. In *Advances in Neural Information Processing Systems 25*, pages 3275–3283. 2012.

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

[12] Sharad Goel, Daniel M. Reeves, Duncan J. Watts, and David M. Pennock. Prediction without markets. In *ACM Conference on Electronic Commerce*, pages 357–366, New York, NY, USA, 2010. ACM.

[13] Robin Hanson. Combinatorial information market design. *Information Systems Frontiers*, 5(1):107–119, 2003.

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

[15] Robert A. Jacobs. Methods for combining experts' probability assessments. *Neural Computation*, 7(5): 867–888, September 1995.

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

[17] Charles F. Manski. Interpreting the predictions of prediction markets. *Economics Letters*, 91(3):425–429, June 2006.

[18] Abraham Othman, Tuomas Sandholm, David M. Pennock, and Daniel M. Reeves. A practical liquidity-sensitive automated market maker. In *ACM Conference on Electronic Commerce*, 2010.

[19] R. Tyrrell Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, New Jersey, 1970.

[20] Justin Wolfers and Eric Zitzewitz. Prediction markets. *Journal of Economic Perspectives*, Winter, 2004.

[21] Justin Wolfers and Eric Zitzewitz. Interpreting prediction market prices as probabilities. IZA Discussion Papers 2092, Institute for the Study of Labor (IZA), April 2006.

A Examples of cost functions

Example A.1 (Quadratic cost)**.** The first example of a cost function, applicable to arbitrary outcome sets Ω, is the quadratic cost function defined by $C(q) = \frac{1}{2}\|q\|^2$. In this case, $p(q) = q$, and $U(q', \mu; q) = \frac{1}{2}\|q-\mu\|^2 - \frac{1}{2}\|q'-\mu\|^2$. It is clear that the expected utility is maximized when $p(q') = q' = \mu$.

Convex conjugate and Bregman divergence: $C^*(\nu) = \frac{1}{2}\|\nu\|^2$ and $D(q, \nu) = \frac{1}{2}\|q - \nu\|^2$, i.e., the Bregman divergence is a monotone transformation of the Euclidean distance.

Example A.2 (LMSR)**.** Our second example is Hanson's logarithmic market-scoring rule (LMSR), which is applied to *complete markets* whose outcomes coincide with basis vectors, i.e., $\Omega = \{e_i : i \in [n]\}$ where e_i denotes the i-th basis vector and $[n]$ denotes the set $\{1, \dots, n\}$. In this case $\mathcal{M}$ is the simplex in $\mathbb{R}^n$ and beliefs μ are in one-to-one correspondence with probability distributions over Ω. The LMSR cost function is

$$C(q) = \ln\left(\sum_{i=1}^n e^{q[i]}\right)$$

where $q[i]$ denotes the i-th coordinate of q. The price vector is

$$p(q)[i] = \frac{\partial C(q)}{\partial q[i]} = \frac{e^{q[i]}}{\sum_{j=1}^n e^{q[j]}} = e^{q[i]-C(q)} \quad .$$

For $\mu \in \mathcal{M}$, the expected utility function takes form

$$\begin{aligned} U(q', \mu; q) &= \sum_{i=1}^n \mu[i]\Big(\ln p(q')[i] - \ln p(q)[i]\Big) \\ &= \mathrm{KL}(\mu\|p(q)) - \mathrm{KL}(\mu\|p(q')) \quad , \end{aligned}$$

where $\mathrm{KL}(\mu\|\nu) = \sum_{i=1}^n \mu[i]\ln(\mu[i]/\nu[i])$ is the KL-divergence. KL-divergence is not symmetric, but it is nonnegative, and zero only if the arguments are equal. Thus, the expected utility is clearly maximized if and only if $\mu = p(q')$.

Convex conjugate and Bregman divergence: $C^*(\nu) = \infty$ if ν is not a probability measure on Ω, and $C^*(\nu) = \sum_{i=1}^n \nu[i]\ln\nu[i]$ otherwise, with the usual convention $0\ln 0 = 0$. The Bregman divergence is $D(q, \nu) = \mathrm{KL}(\nu\|p(q))$.

Example A.3 (Log-partition cost)**.** Next example is the log-partition function, which is applicable to arbitrary outcome sets Ω and which generalizes LMSR:

$$C(q) = \ln\left(\sum_{\omega\in\Omega} e^{q\cdot\omega}\right) \quad .$$

Let P_q be the probability measure over Ω defined by

$$P_q(\omega) = e^{q\cdot\omega - C(q)} \quad .$$

The prices then correspond to expected values of ω under P_q:

$$p(q) = \sum_{\omega\in\Omega} P_q(\omega)\omega \quad .$$

For $\mu \in \mathcal{M}$, let P_μ denote the distribution of maximum entropy among P with $\mathbb{E}_P[\omega] = \mu$ (this distribution is unique and always exists). Note that we are overloading notation on P_q and P_ν and use the "type" of the subscript to indicate which probability distribution we have in mind. The expected utility function can be written as

$$\begin{aligned} U(q', \mu; q) &= (q' - q) \cdot \mathop{\mathbb{E}}_{\omega\sim P_\mu}[\omega] - C(q') + C(q) \\ &= \mathop{\mathbb{E}}_{\omega\sim P_\mu}[\ln P_{q'}(\omega) - \ln P_q(\omega)] \\ &= \mathop{\mathbb{E}}_{\omega\sim P_\mu}\left[\ln\left(\frac{P_\mu(\omega)}{P_q(\omega)}\right) - \ln\left(\frac{P_\mu(\omega)}{P_{q'}(\omega)}\right)\right] \\ &= \mathrm{KL}(P_\mu\|P_q) - \mathrm{KL}(P_\mu\|P_{q'}) \quad . \end{aligned}$$

A standard duality result shows that the infimum of $\mathrm{KL}(P_\mu\|P_{q'})$ over the set $\{P_{q'} : q' \in \mathbb{R}^n\}$ is zero. If there exists q' attaining this minimum, we must have $P_\mu = P_{q'}$ and thus $\mu = p(q')$. We argue that the converse is true as well. Let q', q'' be such that $P_{q'} = P_\mu$ and $p(q') = p(q'') = \mu$. Then by convexity of C, we have $C(q'') - C(q') = (q'' - q') \cdot p(q')$. Therefore,

$$\mathrm{KL}(P_{q'}\|P_{q''}) = (q' - q'') \cdot p(q') - C(q') + C(q'') = 0 \quad ,$$

i.e., $P_{q'} = P_{q''} = P_\mu$. Hence, for any $q \in \mathbb{R}^n$, P_q is exactly the distribution of maximum entropy among those P that satisfy $\mathbb{E}_P[\omega] = p(q)$. In other words, $P_{p(q)} = P_q$.

Convex conjugate and Bregman divergence: $C^*(\nu) = \infty$ if there is no distribution P on Ω such that $\mathbb{E}_P[\omega] = \nu$, and $C^*(\nu) = \sum_{\omega\in\Omega} P_\nu(\omega)\ln P_\nu(\omega)$ otherwise. The Bregman divergence $D(q, \nu) = \mathrm{KL}(P_\nu\|P_q) = \mathrm{KL}(P_\nu\|P_{p(q)})$.

B Proof of Theorem 2.2

Proof of Theorem 2.2. Throughout this proof we use concepts of convex conjugacy and Bregman divergence introduced in Section 4. Let $B \geq 0$ and $\mathcal{B} := \{q : U(q, \omega; q_0) \geq -B \text{ for all } \omega \in \Omega\}$ be the set of states satisfying the constraints of Convex Program (2.1). Using the definition of utility function, we can rewrite Convex Program (2.1) as

$$\underset{q\in\mathbb{R}^n}{\text{Maximize}} \quad [(q - q_0) \cdot \mu - C(q) + C(q_0) - I_{\mathcal{B}}(q)] \tag{B.1}$$

where $I_{\mathcal{B}}(\cdot)$ is the convex indicator function, equal to 0 on the set $\mathcal{B}$ and ∞ outside it. Since the cost function C is convex on $\mathbb{R}^n$, and $\mathcal{B}$ is closed, convex and non-empty, Fenchel's Duality Theorem [19, Theorem 31.1] implies that the supremum of the above objective equals the following minimum

$$\min_{\nu\in\mathbb{R}^n} [C^*(\nu) - q_0 \cdot \mu + C(q_0) + I^*_{\mathcal{B}}(\mu - \nu)] \tag{B.2}$$

and this minimum is attained at some $\hat{\nu} \in \mathbb{R}^n$. Now, let $\hat{q} \in \hat{Q}(B; q_0)$ be a solution of Eq. (B.1). By Fenchel's Duality, the gap between the objectives of Eq. (B.2) and Eq. (B.1) at $\hat{\nu}$ and $\hat{q}$ must be zero:

$$
\begin{aligned}
0 &= C^*(\hat{\nu}) - q_0 \cdot \mu + C(q_0) + I^*_{\mathcal{B}}(\mu - \hat{\nu}) \\
&\quad - (\hat{q} - q_0) \cdot \mu + C(\hat{q}) - C(q_0) + I_{\mathcal{B}}(\hat{q}) \\
&= C^*(\hat{\nu}) - \hat{q} \cdot \hat{\nu} + C(\hat{q}) + I^*_{\mathcal{B}}(\mu - \hat{\nu}) - \hat{q} \cdot (\mu - \hat{\nu}) \\
&\quad + I_{\mathcal{B}}(\hat{q}) \\
&= D(\hat{q}, \hat{\nu}) + \left[I^*_{\mathcal{B}}(\mu - \hat{\nu}) - \left(\hat{q} \cdot (\mu - \hat{\nu}) - I_{\mathcal{B}}(\hat{q}) \right) \right] \quad .
\end{aligned}
$$

The term in the brackets is non-negative from the definition of the convex conjugate. Since $D(\hat{q}, \hat{\nu})$ is also non-negative, we obtain that it must be zero, i.e., $p(\hat{q}) = \hat{\nu}$. Since this reasoning is independent of the choice $\hat{q} \in \hat{Q}(B; q_0)$, the theorem follows. $\square$

C KKT lemma and the minimal face

This section contains proofs of Lemma 3.1 and Theorem 3.4, and shows that minimal faces are well defined.

Proof of Lemma 3.1. We begin by forming a Lagrangian of Convex Program (2.1), with non-negative multipliers $\lambda = (\lambda_\omega)_{\omega \in \Omega}$:

$$
L(q, \lambda) = U(q, \mu; q_0) + \sum_\omega \lambda_\omega \left(U(q, \omega; q_0) + B \right) \quad .
$$

By differentiability and concavity of the objective and constraints, KKT conditions are both necessary and sufficient for optimality. KKT conditions state that q and λ solve the above problem if and only if the following hold:

- *primal feasibility*: $U(q, \omega; q_0) \geq -B$ for all $\omega \in \Omega$;
- *dual feasibility*: $\lambda \geq 0$;
- *first-order optimality*: $\nabla_1 L(q, \lambda) = 0$;
- *complementary slackness*: $\lambda_\omega \left(U(q, \omega; q_0) + B \right) = 0$; for all $\omega \in \Omega$.

We first show that KKT conditions imply (a)–(d). Assume that KKT conditions hold. Let X be the set of outcomes with tight constraints, i.e., $X = \{x \in \Omega : U(q, x; q_0) = -B\}$. For this X, the conditions (a) and (b) hold by primal feasibility and our definition of X. Note that we have either $X = \emptyset$ or $X = \operatorname{argmin}_{x \in \Omega} (q - q_0) \cdot x$, i.e., X is a face of $\mathcal{M}$. If $X \neq \emptyset$, then (d) follows from our definition of X. If $X = \emptyset$, then (d) follows by primal feasibility. We prove (c) by analyzing first-order optimality. First note that:

$$
\nabla_1 U(q, \nu; q_0) = \nu - \nabla C(q) = \nu - p(q) \quad .
$$

Thus, first-order optimality is equivalent to

$$
\begin{aligned}
\nabla_1 U(q, \mu; q_0) + \sum_\omega \lambda_\omega \nabla_1 U(q, \omega; q_0) &= 0 \\
\mu - p(q) + \sum_\omega \lambda_\omega (\omega - p(q)) &= 0 \\
p(q) &= \frac{\mu + \sum_\omega \lambda_\omega \omega}{1 + \sum_\omega \lambda_\omega} \quad .
\end{aligned}
$$

By complementary slackness, $\lambda_\omega = 0$ for $\omega \in \Omega \backslash X$, so this shows (c).

Now assume that (a)–(d) hold. Then KKT conditions hold for λ_ω representing $p(q)$ as a convex combination of μ and elements of X. $\square$

Proof of Theorem 3.4. If $q \in [q_0 + \mathcal{K}(X_{p(q)})$ then we have $q \in \hat{Q}(q_0)$ by Lemma 3.1 with $X = X_{p(q)}$. For the converse, assume that $q \in \hat{Q}(q_0)$. Lemma 3.1 then implies that there exists a face X such that $q \in [q_0 + \mathcal{K}(X)]$ and $p(q) \in \operatorname{conv}(X \cup \{\mu\})$. By minimality of $X_{p(q)}$, we must have $X_{p(q)} \subseteq X$. By anti-monotonicity of witness cones, we then have $q \in [q_0 + \mathcal{K}(X_{p(q)})]$, finishing the proof. $\square$

Proposition C.1. *Fix $\mu \in \mathcal{M}$. Then for any $\nu \in \mathcal{M}$, there exists the minimal face X_ν with the following property: for any face X such that $\nu \in \operatorname{conv}(X \cup \{\mu\})$, we must have $X_\nu \subseteq X$.*

Proof. If $\nu = \mu$ then $X_\nu = \emptyset$ and the statement holds. Otherwise, consider the ray ρ from μ towards ν, and let ν' be the last point on the ray that is contained in $\mathcal{M}$. Let X_ν be the unique face such that ν' lies in the relative interior of $\operatorname{conv}(X_\nu)$.[5] We will argue that this face satisfies the condition stated in the proposition. Let X be any face such that $\nu \in \operatorname{conv}(X \cup \{\mu\})$. Then $\nu = \lambda\mu + (1 - \lambda)\nu_X$ for $\nu_X \in \operatorname{conv}(X)$ and $\lambda \in [0, 1)$. Since $\nu_X \in \mathcal{M}$, it must lie on the ray ρ at some point between ν and ν'. We next argue than $\nu' \in \operatorname{conv}(X)$. Suppose not, this means that $\nu' \neq \nu_X$, and ν_X maximizes some linear function, say $u \cdot \bar{\nu}$, over $\bar{\nu} \in \mathcal{M}$, and $u \cdot \nu' < u \cdot \nu_X$, i.e.,

$$
u \cdot (\nu_X - \nu') > 0 \quad .
$$

Since $\nu \neq \mu$ and $\nu' \neq \nu_X$, and the points μ, ν, ν_X, ν' lie on the ray ρ (in that order), there exists $\eta > 0$ such that $\mu - \nu_X = \eta(\nu_X - \nu')$ and thus

$$
u \cdot (\mu - \nu_X) = \eta u \cdot (\nu_X - \nu') > 0
$$

implying that $u \cdot \mu > u \cdot \nu_X$ and contradicting the assumption that ν_X is the maximizer. Thus, $\nu' \in \operatorname{conv}(X)$. By a similar reasoning, we can also show that for any $x \in X_\nu$,

[5] The existence of such a unique face follows by the standard result stating that relative interiors of $\operatorname{conv}(X)$ across non-empty faces X form a disjoint partition of $\mathcal{M}$.

we must have $x \in X$. Again, for the sake of contradiction assume that there is u such that ν' is a maximizer of $u \cdot \bar{\nu}$ over $\bar{\nu} \in \mathcal{M}$, but x is not. Then $x \neq \nu'$, and since $\nu' \in \operatorname{ri}\operatorname{conv}(X_\nu)$, for sufficiently small η, we have $\nu'' := \nu' + \eta(\nu' - x) \in \mathcal{M}$, and $u \cdot \nu'' > u \cdot \nu'$ contradicting the maximizer property of ν'. Thus, $X_\nu \subseteq X$. □

D Impossibility result of Fortnow and Sami

We can use the KKT lemma and the continuity of the perpendiculars (Theorem F.3) to derive the impossibility result of [9]. The result states that in the presence of budget constraints, there is no market scoring rule guaranteeing that the market prices move towards the agent belief along the connecting straight line (unless $\operatorname{aff}(\mathcal{M})$ is a line or a point). Assume the dimension of $\operatorname{aff}(\mathcal{M})$ is $d \geq 2$. Choose $\mu \in \operatorname{ri}\mathcal{M}$ and q_0 such that $p(q_0)$ does not lie in any $\operatorname{aff}(X \cup \{\mu\})$ of dimension less than d. Since q in the KKT lemma must be on a perpendicular, continuity of the perpendiculars implies that $p(q)$ is arbitrarily close to $p(q_0)$ for a small enough budget. Thus, for a small enough budget, X must give the full-dimensional $\operatorname{aff}(X \cup \{\mu\})$, and $p(q) \in \operatorname{ri}\operatorname{conv}(X \cup \{\mu\})$. The latter will remain true for small changes in μ without affecting conditions (a) and (b) of Lemma 3.1. Thus, the direction of movement of market prices is independent of small changes in μ.

E Budget additivity: examples

Using the KKT lemma, we illustrate on examples that budget additivity sometimes holds and sometimes does not. Recall that budget additivity states that if several agents have the same belief and limited budgets, the sequence of their actions is equivalent to the action of a single agent with the same belief and the sum of the budgets. In the first example, we give an illustration of when this property holds. In the second example, we show how this property can be violated, and the single agent with the sum of budgets has more power in the market.

Example E.1 (Quadratic cost on a square)**.** Consider the following outcome space and belief:

$$\begin{aligned} \omega_{00} &= (0,\ 0) \\ \omega_{01} &= (0,\ 1) \\ \omega_{10} &= (1,\ 0) \\ \omega_{11} &= (1,\ 1) \\ \mu &= (0.9,\ 0.3) \end{aligned}$$

Further, consider the following market states:

$$\begin{aligned} q_0 &= \nu_0 = (0.5,\ 0.1) \\ q_1 &= \nu_1 = (0.6,\ 0.2) = \tfrac{1}{3}\omega_{00} + \tfrac{2}{3}\mu \\ q_\mu &= \mu \end{aligned}$$

The divergence of these states (and the belief μ) from individual outcomes is:

$\frac{1}{2}\lVert\cdot - \cdot\rVert^2$	ω_{00}	ω_{01}	ω_{10}	ω_{11}
ν_0	0.13	0.53	0.13	0.53
ν_1	0.2	0.5	0.1	0.4
μ	0.45	0.65	0.05	0.25

With these in hand, we can now use the KKT lemma and show that $q_1 = \nu_1$ is an optimal action at $q_0 = \nu_0$ under belief μ for a specific budget. Since q_1 is a convex combination of ω_{00} and μ, we need to show that the only tight budget constraint is due to ω_{00}. We also calculate budgets required to move from q_0 and q_1 to q_μ:

	ω_{00}	ω_{01}	ω_{10}	ω_{11}	
$U(q_1, \cdot; q_0)$	-0.07	0.03	0.03	0.13	$B_{01} = 0.07$
$U(q_\mu, \cdot; q_0)$	-0.32	-0.12	0.08	0.28	$B_{0\mu} = 0.32$
$U(q_\mu, \cdot; q_1)$	-0.25	-0.15	0.05	0.15	$B_{1\mu} = 0.25$

Hence, a sequence of moves with budgets B_{01} and $B_{1\mu}$ is equivalent to a single move with the budget $B_{0\mu} = B_{01} + B_{1\mu}$. While we have shown this only for a specific sequence of budgets, results of Section 5 show that budget additivity holds for any sequence of budgets and any belief $\mu \in \mathcal{M}$.

Example E.2 (Quadratic cost on an obtuse triangle.)**.** Now, we work out an example where the budget additivity does not hold. Consider the following outcome space and belief:

$$\begin{aligned} \omega_1 &= (0.0,\ 0.0) \\ \omega_2 &= (1.8,\ 0.0) \\ \omega_3 &= (6.0,\ 4.2) \\ \mu &= (2.7,\ 1.8) \end{aligned}$$

Further, consider the following set of market states:

$$\begin{aligned} q_0 &= \nu_0 = (2.7,\ 0.9) \\ q_1 &= \nu_1 = (2.4,\ 1.2) = \tfrac{1}{3}\omega_2 + \tfrac{2}{3}\mu \\ q_2 &= \nu_2 = (2.4,\ 1.6) = \tfrac{1}{9}\omega_1 + \tfrac{8}{9}\mu \\ q_3 &= \nu_3 = \left(0.9\sqrt{\tfrac{105}{13}},\ 0.6\sqrt{\tfrac{105}{13}}\right) \\ &= \left(1 - \tfrac{1}{3}\sqrt{\tfrac{105}{13}}\right)\omega_1 + \left(\tfrac{1}{3}\sqrt{\tfrac{105}{13}}\right)\mu \approx \tfrac{1}{19}\omega_1 + \tfrac{18}{19}\mu \\ q_\mu &= \mu \end{aligned}$$

The divergence of these states (and the belief μ) from individual outcomes is:

$\frac{1}{2}\lVert\cdot - \cdot\rVert^2$	ω_1	ω_2	ω_3
ν_0	4.05	0.81	10.89
ν_1	3.6	0.9	10.98
ν_2	4.16	1.46	9.86
ν_3	4.725	1.74...	9.04...
μ	5.265	2.025	8.325

Again as before, we can use the KKT lemma and show for $j = 1, 2, 3$, that $q_j = \nu_j$ is an optimal action at $q_{j-1} =$

ν_{j-1} under belief μ, with the corresponding budgets as:

	ω_1	ω_2	ω_3	
$U(q_1,\cdot;\,q_0)$	0.45	-0.09	-0.09	$B_{01}=0.09$
$U(q_2,\cdot;\,q_1)$	-0.56	-0.56	1.12	$B_{12}=0.56$
$U(q_3,\cdot;\,q_2)$	-0.565	$-0.28\ldots$	$0.82\ldots$	$B_{23}=0.565$
$U(q_\mu,\cdot;\,q_0)$	-1.215	-1.215	2.565	$B_{0\mu}=1.215$

The above table also shows that the budget $B_{0\mu} = 1.215$ suffices to move directly from q_0 to q_μ. However, note that the sum

$$B_{01} + B_{12} + B_{23} = 1.215 = B_{0\mu} \ ,$$

but $\nu_3 \neq \mu$, i.e., after the sequence of optimal actions with budgets B_{01}, B_{12}, and B_{23}, the market is still not at the belief shared by all agents, even though with the budget $B_{0\mu}$, it would have reached it. Note that it is possible to achieve budget additivity by using log-partition cost instead of quadratic cost (Theorem 5.6).

F Perpendiculars

F.1 Proofs of Propositions 4.2 and 4.3

Proof of Proposition 4.2. We will show that condition (i) is equivalent to condition (ii) by analyzing the first order optimality conditions. Consider the problem

$$\operatorname*{Minimize}_{\nu' \in A_\lambda} \quad D(q,\nu') \tag{F.1}$$

used to define ν_λ. Assume that the minimum is attained at some $\nu' \in \operatorname{ri}\operatorname{dom} C^*$. Thus, $\nu' \in A \cap (\operatorname{ri}\operatorname{dom} C^*)$. Since $D(q,\nu')$ is subdifferentiable at ν', the first order optimality implies that

$$\bigl(\partial_2 D(q,\nu')\bigr) \cap A_\lambda^\perp \neq \emptyset \ . \tag{F.2}$$

Since $\partial_2 D(q,\nu') = \partial C^*(\nu') - q = p^{-1}(\nu') - q$, and and $A_\lambda^\perp = A_0^\perp$, we have

$$p^{-1}(\nu') \cap (q + A_0^\perp) \neq \emptyset \ ,$$

proving that (i)$\Rightarrow$(ii). Conversely, assume that $\nu' \in A \cap (\operatorname{ri}\operatorname{dom} C^*)$ and $p^{-1}(\nu') \cap (q + A_0^\perp) \neq \emptyset$. Then we can pick λ such that $\nu' \in A_\lambda$ and for this λ, we obtain that condition (F.2) holds and hence ν' solves problem (F.1). Since $\nu' \in \operatorname{ri}\operatorname{dom} C^*$, we obtain that $\nu' \in \operatorname{im}\gamma$. □

Proof of Proposition 4.3. Let γ' be the a_1-perpendicular to A_0 at q'. Since the ambient space A for both perpendiculars is the same, by Proposition 4.2(ii), it suffices to show that $q + A_0^\perp = q' + A_0^\perp$. However, this follows by the assumption of the theorem, since $q' - q \in A_0^\perp$. □

F.2 Continuity of perpendiculars

In this section, we prove two important properties of perpendiculars: (a) they are continuous maps; (b) intersections of perpendiculars with compact convex sets correspond to compact sets of market states up to certain "irrelevant displacements". To define these irrelevant displacements, let $\mathcal{L}$ be the linear space parallel to $\operatorname{aff}(\operatorname{dom} C^*)$. Then the displacements of market state within $\mathcal{L}^\perp$ are *irrelevant* in the sense that they have no effect on the Bregman divergence and hence by Eq. (4.1) also no effect on the utility function. Specifically, $D(q+u,\nu) = D(q,\nu)$ for all $u \in \mathcal{L}^\perp$ (see next proposition). For instance, for LMSR over a simplex, the irrelevant displacements are of the form $\lambda\vec{1}$ where $\lambda \in \mathbb{R}$ and $\vec{1}$ is the all-ones vector.

Proposition F.1. *Let $\mathcal{L}$ be the linear space parallel to $\operatorname{aff}(\operatorname{dom} C^*)$. Then for all $q \in \mathbb{R}^n$ and $u \in \mathcal{L}^\perp$*

$$D(q+u,\nu) = D(q,\nu) \ .$$

Proof. If $\nu \notin \operatorname{dom} C^*$ then the statement obviously holds. Pick $\nu \in \operatorname{dom} C^*$, $q \in \mathbb{R}^n$ and $u \in \mathcal{L}^\perp$. By the Mean Value Theorem, we can write

$$C(q+u) - C(q) = u \cdot \nabla C(\bar{q})$$

for some $\bar{q}$. Let $\bar{\nu} := \nabla C(\bar{q}) \in \operatorname{dom} C^*$. Then we can write

$$D(q+u,\nu) - D(q,\nu) = C(q+u) - C(q) - u\cdot\nu = u\cdot(\bar{\nu}-\nu) = 0$$

since $u \perp (\bar{\nu} - \nu)$. □

The following result of Rockafellar [19] will be instrumental in proving continuity properties of the perpendicular. It is paraphrased for our setting. The notation int refers to the topological interior of the set.

Theorem F.2 (Theorem 24.7 of 19)**.** *Let $G : \mathbb{R}^n \to (-\infty,\infty]$ be a lower semi-continuous convex function, and let K be a non-empty, closed and bounded subset of $\operatorname{int}(\operatorname{dom} G)$. Then the set*

$$\partial G(K) = \bigcup_{u \in K} \partial G(u)$$

is non-empty, closed and bounded.

Now we are ready to state and prove the continuity of perpendiculars:

Theorem F.3. *Let γ be the a_1-perpendicular to A_0 at q, and $K \subseteq \operatorname{ri}\operatorname{dom} C^*$ be a closed bounded convex set intersecting $\operatorname{im}\gamma$.*

(a) *The map γ is continuous.*
(b) *The intersection $M := \{(\nu,q) : \nu \in (\operatorname{im}\gamma) \cap K,\ q \in p^{-1}(\nu)\}$ can be written as $M = \mathcal{C} + (0 \oplus \mathcal{L}^\perp)$ where $\mathcal{C}$ is compact and $\oplus$ denotes a direct sum of vector spaces.*

Proof. Throughout the proof, let $F(\nu) := D(q, \nu) = C(q) + C^*(\nu) - q \cdot \nu$. Note that F is strictly convex on $\operatorname{ri}\operatorname{dom} C^*$. We will be also making frequent use of the fact that F is continuous on $\operatorname{ri}\operatorname{dom} C^*$ (because C^* is continuous on $\operatorname{ri}\operatorname{dom} C^*$ by Theorem 10.1 of 19). Let $\|\cdot\|$ denote the usual Euclidean norm. Let $a_0 = \operatorname{argmin}_{a \in A_0} \|a_1 - a\|$, i.e., $(a_1 - a_0) \in A_0^\perp$. Let $A = \operatorname{aff}(A_0 \cup \{a_1\})$ and recall that

$$A_\lambda = A_0 + \lambda(a_1 - a_0)$$

and

$$\nu_\lambda = \operatorname*{argmin}_{A_\lambda} F(\nu) \ .$$

We use the notation $\mathcal{B}(\nu, r; M) := \{\nu' \in M : \|\nu' - \nu\| \leq r\}$ for the Euclidean ball relative to set M, and $\mathcal{S}(\nu, r; M) := \{\nu' \in M : \|\nu' - \nu\| = r\}$ for the Euclidean sphere relative to set M.

Part (a). We need to show that γ is continuous. Let $\lambda \in \Lambda$, i.e., $\nu_\lambda \in \operatorname{ri}\operatorname{dom} C^*$. Choose a sufficiently small $r > 0$ such that the ball $\mathcal{B} := \mathcal{B}(\nu_\lambda, r; A)$ is contained in $\operatorname{ri}\operatorname{dom} C^*$. To show the continuity of γ at λ, it suffices to show that if λ' is close enough to λ then $\nu_{\lambda'} \in \mathcal{B}$.

Let $\varepsilon = r/\sqrt{2}$. Consider the sphere $\mathcal{S}_\lambda := \mathcal{S}(\nu_\lambda, \varepsilon; A_\lambda) \subseteq \operatorname{ri}\operatorname{dom} C^*$. This sphere is a compact set, so F attains the minimum on $\mathcal{S}_\lambda$. By strict convexity of F and the optimality of ν_λ, this minimum must be bounded away from $F(\nu_\lambda)$. Thus, there exists $\delta > 0$ such that

$$F(\nu) \geq F(\nu_\lambda) + \delta \text{ for all } \nu \in \mathcal{S}_\lambda \ . \tag{F.3}$$

Let $\delta' = \delta/3$. Since F is continuous on $\operatorname{ri}\operatorname{dom} C^*$, it is uniformly continuous on $\mathcal{B}$ and thus there exists $\varepsilon' \in (0, \varepsilon]$ such that

$$|F(\nu') - F(\nu)| \leq \delta' \text{ for all } \nu, \nu' \in \mathcal{B} \text{ such that } \|\nu' - \nu\| \leq \varepsilon' \ . \tag{F.4}$$

Let $\mathcal{B}_\lambda := \mathcal{B}(\nu_\lambda, \varepsilon; A_\lambda)$ be the closed ball with $\mathcal{S}_\lambda$ as the border. For any λ', let $\mathcal{S}_{\lambda'} := \mathcal{S}_\lambda + (\lambda' - \lambda)(a_1 - a_0) \subseteq A_{\lambda'}$ and similarly $\mathcal{B}_{\lambda'}$. Let $\tilde{\nu}_{\lambda'} = \nu_\lambda + (\lambda' - \lambda)(a_1 - a_0) \in A_{\lambda'}$.

Note that if $|\lambda' - \lambda| \leq \varepsilon'$, then $\mathcal{B}_{\lambda'} \subseteq \mathcal{B}$, because $\sqrt{(\lambda' - \lambda)^2 + \varepsilon^2} \leq \varepsilon\sqrt{2} = r$. So we can use the above uniform continuity result and write:

- $F(\nu') \geq F(\nu_\lambda) + \delta - \delta' = F(\nu_\lambda) + 2\delta'$ for all $\nu' \in S_{\lambda'}$ by Eqs. (F.3) and (F.4)
- $F(\tilde{\nu}_{\lambda'}) \leq F(\nu_\lambda) + \delta'$ by Eq. (F.4)

By convexity of F, this means that $\nu_{\lambda'} \in \mathcal{B}_{\lambda'}$. This proves that $\nu_{\lambda'} \in \mathcal{B}$ provided that $|\lambda' - \lambda| \leq \varepsilon'$, thus proving the continuity of γ at λ.

Part (b). We first show that the set M is closed and then that it is bounded, except for directions in $0 \oplus \mathcal{L}^\perp$. Since $K \subseteq \operatorname{ri}\operatorname{dom} C^*$, we can use Proposition 4.2 to write the set M as

$$\begin{aligned} M = \{(\nu', q') :\ & \nu' \in \mathbb{R}^n,\ q' \in \partial C^*(\nu')\} \\ & \cap \left(\mathbb{R}^n \times (q + A_0^\perp)\right) \cap (K \times \mathbb{R}^n) \ , \end{aligned} \tag{F.5}$$

where we used the identity $p^{-1}(\nu') = \partial C^*(\nu')$ valid for all ν'. The closedness follows, because the set of pairs $\{(\nu', q') : \nu' \in \mathbb{R}^n, q' \in \partial C^*(\nu')\}$ is closed by Rockafellar [19, Theorem 24.4].

Denote the projections of M on its two components as

$$\begin{aligned} M_1 &:= \{\nu' : (\nu', q') \in M \text{ for some } q'\} \ , \\ M_2 &:= \{q' : (\nu', q') \in M \text{ for some } \nu'\} \ . \end{aligned}$$

To show boundedness, we only need to analyze M_2 since $M_1 \subseteq K$. By Eq. (F.5), it in fact suffices to show that the set $\partial C^*(K) = \bigcup_{\nu \in K} \partial C^*(\nu)$ is bounded except for directions in $\mathcal{L}^\perp$. We would like to appeal to Theorem F.2, but we cannot do it directly, because it is stated for the *interior* rather than the *relative interior*. For $\nu \in \operatorname{ri}\operatorname{dom} C^*$, we have $\partial C^*(\nu) \neq \emptyset$, and using the fact that $C^*(\nu) = \infty$ over $\nu + (\mathcal{L}^\perp \backslash \{0\})$, we obtain that

$$\partial C^*(\nu) = S + \mathcal{L}^\perp$$

for some set $S \subseteq \mathcal{L}$. This set S coincides with subdifferential when $C^*(\nu)$ is only viewed as a function over $\operatorname{aff}(\operatorname{dom} C^*)$. By applying Theorem F.2 to this restriction, we then indeed obtain that

$$\partial C^*(K) = \mathcal{C} + \mathcal{L}^\perp$$

for a non-empty closed and bounded set $\mathcal{C}$. Note that $\mathcal{L}^\perp \subseteq A_0^\perp$, so $\mathcal{L}^\perp$ survives taking the intersection in Eq. (F.5) and hence part (b) of the theorem follows. ☐

G Proofs of sufficient conditions for acute angles

Proposition G.1. *For the quadratic cost, Definition 5.1 is equivalent to Definition 3.8.*

Proof. We first show that Definition 5.1 (general acute angles) implies Definition 3.8 (Euclidean acute angles). Assume that the general acute angles hold for X. Let $\bar{\nu} \in \mathcal{M}$ and ν be its projection on $\operatorname{aff}(X)$. If $\bar{\nu} = \nu$ then angles between $\bar{\nu}$, ν and $\omega \in \Omega$ are non-obtuse in the sense that $(\bar{\nu} - \nu) \cdot (\omega - \nu) \geq 0$. If $\bar{\nu} \neq \nu$, then let γ be the $\bar{\nu}$-perpendicular to X at ν (note that p is the identity map, so ν is both a state and the corresponding price vector). Note that $\bar{\nu} \succeq \nu$ and thus by the general acute angles assumption $p^{-1}(\bar{\nu}) \cap [\nu + \mathcal{K}(X)] \neq \emptyset$. Since $p^{-1}(\bar{\nu}) = \{\bar{\nu}\}$, this is equivalent to

$$(\bar{\nu} - \nu) \cdot (\omega - x) \geq 0 \text{ for all } x \in X, \omega \in \Omega \ .$$

Since $\nu \in \operatorname{aff}(X)$, we obtain that $(\bar{\nu}-\nu)\cdot(\omega-\nu) \geq 0$, i.e., the Euclidean acute angles hold.

Conversely, assume that the Euclidean acute angles hold. Let γ be a μ-perpendicular to a face X for some $\mu \in \mathcal{M}$ and $\nu' \succeq \nu$ be two points in $\operatorname{im}\gamma$ such that $\nu \in \mathcal{M}$. We need to show that $p^{-1}(\nu') \cap [\nu + \mathcal{K}(X)] \neq \emptyset$, which is equivalent to

$$(\nu'-\nu)\cdot(\omega - x) \geq 0 \text{ for all } x \in X, \omega \in \Omega \ . \tag{G.1}$$

If $\nu' = \nu$ then (G.1) holds. Otherwise, we can write $\nu' = \nu + u_X + \lambda(\mu - \hat{\nu})$ for an arbitrary $\hat{\nu} \in \operatorname{aff}(X)$, a suitable $\lambda > 0$ and u_X from the linear space parallel with $\operatorname{aff}(X)$. Pick $\hat{\nu} \in \operatorname{ri}\operatorname{conv}(X)$ (and the corresponding λ and u_X). We claim that there is a small enough $\eta > 0$ such that $\bar{\nu} := \hat{\nu} + \eta(\nu' - \nu) \in \mathcal{M}$. This follows, because from our previous reasoning,

$$\bar{\nu} = \hat{\nu} + \eta u_X + \eta\lambda(\mu - \hat{\nu}),$$

and for sufficiently small $\eta > 0$, we must have $[\hat{\nu} + \eta\lambda(\mu - \hat{\nu})] \in \operatorname{ri}\operatorname{conv}(X \cup \{\mu\})$ and then also for sufficiently small η, $[\hat{\nu}+\eta\lambda(\mu-\hat{\nu})+\eta u_X] \in \operatorname{ri}\operatorname{conv}(X \cup \{\mu\}) \subseteq \mathcal{M}$. Thus, by the Euclidean acute angles,

$$(\bar{\nu}-\hat{\nu})\cdot(\omega-\hat{\nu}) \geq 0 \text{ for all } \omega \in \Omega.$$

Since $\bar{\nu} - \hat{\nu} = \eta(\nu' - \nu)$ and $(\nu'-\nu) \perp (x - \hat{\nu})$ for all $x \in X$, we also obtain

$$\eta(\nu'-\nu)\cdot(\omega - x) \geq 0 \text{ for all } x \in X, \omega \in \Omega$$

proving (G.1) and finishing the proof. $\square$

Proof of Theorem 5.3. Let L be the linear space parallel to $\operatorname{aff}(X)$. First show that the acute angles imply the inclusion of the projection in an acute cone. Note that the inclusion is either true for all $a_0 \in \operatorname{aff}(X)$ or none, so we can without loss of generality choose $a_0 \in \operatorname{conv}(X)$. Let $\omega_1, \omega_2 \in \Omega$ and let ω_1' and ω_2' be their projections to A', thus

$$\omega_1' - \omega_1 \in L \ , \quad \omega_2' - \omega_2 \in L \ .$$

We need to show that

$$(\omega_1' - a_0)\cdot(\omega_2' - a_0) \geq 0 \ .$$

If $\omega_1 \in \operatorname{aff}(X)$ then $\omega_1' = a_0$ and the statement holds. Assume that $\omega_1 \notin \operatorname{aff}(X)$ and let γ be the ω_1-perpendicular to X at a_0. Let $\omega_1'' \in \operatorname{im}\gamma$ be the projection of ω_1 on $\operatorname{im}\gamma$. Thus, we also have

$$\omega_1'' - \omega_1 \in L$$

and also $\omega_1'' \succeq a_0$. Now by the acute angles assumption, $\omega_1'' - a_0 \in \mathcal{K}(X)$, i.e., for any $x \in X$,

$$0 \leq (\omega_1'' - a_0)\cdot(\omega_2 - x) \ .$$

Combining this with the previous identities, we obtain

$$\begin{aligned}0 &\leq (\omega_1'' - a_0)\cdot(\omega_2 - x) = (\omega_1'' - a_0)\cdot(\omega_2' - a_0)\\ &= (\omega_1' - a_0)\cdot(\omega_2' - a_0)\end{aligned}$$

where the first equality follows because $\omega_1'' - a_0 \in L^\perp$ and

$$\omega_2' - \omega_2 \in L \ , \quad a_0 - x \in L \ ,$$

the second equality follows because $\omega_2' - a_0 \in L^\perp$ and

$$\omega_1'' - \omega_1' = [(\omega_1'' - \omega_1) - (\omega_1' - \omega_1)] \in L \ .$$

For the converse, assume that the inclusion of the projection of $\mathcal{M}$ in an affine acute cone holds. Let γ be the μ-perpendicular to X at ν for some $\mu, \nu \in \mathcal{M}$ and let $\nu' \succeq \nu$. We need to show that $\nu' - \nu \in \mathcal{K}(X)$. Note that $0 \in \mathcal{K}(X)$, so we only analyze $\nu' \neq \nu$. Let μ' be the projection of μ on $\operatorname{im}\gamma$ and a_0 be the intersection of $\operatorname{im}\gamma$ with $\operatorname{aff}(X)$. Note that $\nu' - \nu = \eta(\mu' - a_0)$ for a suitable $\eta > 0$, so it suffices to show that $\mu' - a_0 \in \mathcal{K}(X)$. Pick $\omega \in \Omega$ and $x \in X$ and let ω' be the projection of ω into $A' := a_0 + X^\perp$. Since the projection of $\mathcal{M}$ into A' is contained in an affine acute cone with the vertex a_0, we obtain

$$(\mu' - a_0)\cdot(\omega' - a_0) \geq 0$$

Since $\omega' - \omega \in L$ and $x - a_0 \in L$, whereas $\mu' - a_0 \in L^\perp$, we obtain

$$(\mu' - a_0)\cdot(\omega - x) \geq 0$$

showing that the acute angles hold. $\square$

Proof of Corollary 5.4. We will show that the assumption of Theorem 5.3 holds. Since the assumption is invariant under rigid transformations, we can just consider the case $a_0 = 0 \in X$. In this case, the projection of Ω is a lower dimensional hypercube (corresponding to a subset of Ω). Note that Ω lies in the non-negative orthant and the non-negative orthant is an acute cone with the vertex $a_0 = 0$, so the assumption of Theorem 5.3 holds and hence the acute angles hold for the hypercube. $\square$

Proof of Corollary 5.5. Again, by symmetry, it suffices to consider faces of the form $X = \{e_i \,:\, i \in [k]\}$ for $k \in \{1, \ldots, n\}$. Let $a_0 = e_1$. The affine space A' is described by

$$\begin{aligned}A' &= \{a \in \mathbb{R}^n \,:\, (a - e_1)\cdot(e_i - e_1) = 0 \text{ for all } i \in \{2, \ldots, k\}\}\\ &= \{a \in \mathbb{R}^n \,:\, a[i] = a[1] - 1 \text{ for all } i \in \{2, \ldots, k\}\}\end{aligned}$$

where we use notation $a[i]$ to denote the i-th coordinate. The projection of e_j for $j > k$ into A' is of the form

$$e_j' = e_j + \sum_{i=2}^{k} \alpha_j[i]\,(e_i - e_1) \ ,$$

for some $\alpha_j[i] \in \mathbb{R}$, i.e.,

$$e'_j[i] = \begin{cases} -\sum_{i'=2}^{k} \alpha_j[i'] & \text{if } i = 1 \\ \alpha_j[i] & \text{if } i \in \{2, \ldots, k\} \\ 1 & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases}$$

The only solution of the above form that lies in A' is obtained by setting $\alpha_j[i] = -1/k$, yielding

$$e'_j[i] = \begin{cases} 1 - 1/k & \text{if } i = 1 \\ -1/k & \text{if } i \in \{2, \ldots, k\} \\ 1 & \text{if } i = j \\ 0 & \text{otherwise.} \end{cases}$$

Therefore, for any pair of projections e'_j, $e'_{j'}$ for $j, j' > k$, and $j \neq j'$, we have

$$(e'_j - e_1) \cdot (e'_{j'} - e_1) = 1/k > 0 \ ,$$

so the projection of Ω is in an acute cone, i.e., acute angles hold. $\square$

Proof of Theorem 5.6. We begin by characterizing an a_1-perpendicular to a face $X \neq \Omega$ at q. Let $\nu := p(q) \notin \operatorname{aff}(X)$, so the ambient space of the perpendicular is $\operatorname{aff}(X \cup \{\nu\})$. Thus, for a given X and q, we will have the same $\operatorname{im}\gamma$ and the same order on $\nu \in \operatorname{im}\gamma$, regardless of $a_1 \in \tilde{\mathcal{M}}$ chosen. Recall that P_q is the probability measure over Ω defined by

$$P_q(\omega) = e^{q \cdot \omega - C(q)}$$

and note that $P_q(\omega) > 0$ for all $\omega \in \Omega$. Recall that

$$\nu = \sum_{\omega \in \Omega} P_q(\omega)\omega \ .$$

Let $X^c = \Omega \backslash X$. Separate ν into components corresponding to $x \in X$ and $\omega \in X^c$:

$$\nu_X = \frac{1}{P_q(X)} \sum_{x \in X} P_q(x)x \ , \quad \nu_{X^c} = \frac{1}{P_q(X^c)} \sum_{\omega \in X^c} P_q(\omega)\omega \ ,$$

i.e.,

$$\nu = P_q(X)\nu_X + P_q(X^c)\nu_{X^c} \ .$$

Since $\nu_{X^c} \notin \operatorname{aff}(X)$, we have

$$\operatorname{aff}(X \cup \{\nu\}) = \operatorname{aff}(X \cup \{\nu_{X^c}\}) \ . \tag{G.2}$$

We will show that $\operatorname{im}\gamma$ consists exactly of the points $(1 - \tilde{\alpha})\nu_X + \tilde{\alpha}\nu_{X^c}$ for $\tilde{\alpha} \in (0, 1)$.

Consider $\nu' \in \operatorname{im}\gamma$ and $q' \in p^{-1}(\nu') \cap (q + X^{\perp})$. For any $x', x \in X$, we have

$$(q' - q) \cdot (x' - x) = 0 \ ,$$

so

$$q' \cdot (x' - x) = q \cdot (x' - x) \ ,$$

and hence

$$\frac{P_{q'}(x')}{P_{q'}(x)} = e^{q' \cdot (x' - x)} = e^{q \cdot (x' - x)} = \frac{P_q(x')}{P_q(x)} \ .$$

Since this holds for arbitrary $x, x' \in X$, we obtain

$$\frac{P_{q'}(x)}{P_q(x)} = \frac{P_{q'}(X)}{P_q(X)} \text{ for all } x \in X \ . \tag{G.3}$$

Since ν' is in the ambient space of the perpendicular, which is $\operatorname{aff}(X \cup \{\nu\})$, by Eq. (G.2), we obtain

$$\nu' \in \operatorname{aff}(X \cup \{\nu_{X^c}\}) \ ,$$

so ν' can be written in the form

$$\nu' = \sum_{x \in X} \alpha(x)x + \left(1 - \sum_{x \in X} \alpha(x)\right)\nu_{X^c} \tag{G.4}$$

for some $\alpha(x) \in \mathbb{R}$ for $x \in X$. Also,

$$\nu' = \sum_{x \in X} P_{q'}(x)x + \sum_{\omega \in X^c} P_{q'}(\omega)\omega \ .$$

By the affine independence of Ω, we therefore must have $\alpha(x) = P_{q'}(x)$. Plugging this into Eq. (G.4), we obtain

$$\begin{aligned} \nu' &= \left(\sum_{x \in X} P_{q'}(x)x\right) + P_{q'}(X^c)\nu_{X^c} \\ &= \left(\frac{P_{q'}(X)}{P_q(X)} \sum_{x \in X} P_q(x)x\right) + P_{q'}(X^c)\nu_{X^c} \\ &= P_{q'}(X)\nu_X + P_{q'}(X^c)\nu_{X^c} \end{aligned}$$

where the second equality follows by Eq. (G.3). Thus, indeed $\nu' = (1-\tilde{\alpha})\nu_X + \tilde{\alpha}\nu_{X^c}$ for $\tilde{\alpha} = P_{q'}(X^c)$. Conversely, for any $\tilde{\alpha} \in (0, 1)$, note that

$$\left((1 - \tilde{\alpha})\nu_X + \tilde{\alpha}\nu_{X^c}\right) \in \left(\operatorname{aff}(X) + \tilde{\alpha}(\nu_{X^c} - \nu_X)\right) = A_\lambda$$

for a suitable λ. For this λ, we can take $\nu_\lambda := \operatorname{argmin}_{A_\lambda} D(q, \nu)$. We can write ν_λ as a unique convex combination described by a measure P:

$$\nu_\lambda = \sum_{\omega \in \Omega} P(\omega)\omega \ .$$

To finish the proof we just need to argue that $P(\omega) > 0$ for all ω, which will imply that $\nu_\lambda \in \operatorname{ri}\operatorname{dom} C^*$, i.e., C^* is subdifferentiable at ν_λ, and thus $\lambda \in \Lambda$. However, this follows by noting that $D(q, \nu_\lambda) = \mathrm{KL}(P \| P_q)$ and the latter equals ∞ if there is any point ω such that $P(\omega) = 0$, because $P_q(\omega) > 0$ for all $\omega \in \Omega$.

Now we are ready to prove the theorem. Let $\nu' \in \operatorname{im}\gamma$ such that $\nu' \succeq \nu$, and let $q \in p^{-1}(\nu)$, $q' \in p^{-1}(\nu')$. Use the notation $P := P_q$ and $P' := P_{q'}$, and write

$$
\begin{aligned}
(q'-q)\cdot(\omega-x) &= \ln\left(\frac{e^{q'\cdot\omega}}{e^{q'\cdot x}}\cdot\frac{e^{q\cdot x}}{e^{q\cdot\omega}}\right) \\
&= \ln\left(\frac{P'(\omega)}{P'(x)}\cdot\frac{P(x)}{P(\omega)}\right) \\
&= \ln\left(\frac{P'(\omega)}{P(\omega)}\cdot\frac{P(x)}{P'(x)}\right) \quad .
\end{aligned}
$$

For $\omega \in X^c$, our characterization of the perpendicular implies that $P(x) \geq P'(x)$ and $P(\omega) \leq P(\omega')$ since ν' has a larger (or equal) coefficient $\tilde{\alpha}$ than ν, because it is further (or equally) away from ν_X. Thus, the above expression is non-negative, yielding the acute angles property. □

Proof of Theorem 5.7. Let $\mu \in \tilde{\mathcal{M}}$, $q \in p^{-1}(\tilde{\mathcal{M}})$ and let $\operatorname{im}\gamma$ be the μ-perpendicular to X at q. Note that the perpendicular is well defined only if X is a singleton, say $X = \{x\}$, and $\mu \neq x$. Let $\{x'\}$ be the other singleton face of $\mathcal{M}$. Thus, $\operatorname{im}\gamma = \operatorname{aff}(\{x,x'\}) \cap (\operatorname{ri}\operatorname{dom} C^*)$ with the direction from x towards x'. Note that $\mathcal{K}(X) = \{u : u\cdot(x'-x) \geq 0\}$. Let $\nu = p(q)$ and let $\nu' \succeq \nu$, i.e., $\nu' - \nu = \lambda(x'-x)$ for some $\lambda > 0$. Pick $q' \in p^{-1}(\nu')$, which exists, because $\nu' \in \operatorname{ri}\operatorname{dom} C^*$. By convexity, we have

$$
\begin{aligned}
0 &\leq (q'-q)\cdot(\nabla C(q') - \nabla C(q)) = (q'-q)\cdot(\nu'-\nu) \\
&= \lambda(q'-q)\cdot(x'-x) \quad ,
\end{aligned}
$$

i.e., $q'-q \in \mathcal{K}(X)$. □

Proof of Theorem 5.8. Let $\mathcal{M}_1 = \operatorname{conv}(\Omega_1)$ and $\mathcal{M}_2 = (\operatorname{conv}\Omega_2)$. We first argue that $\mathcal{M} = \operatorname{conv}(\Omega_1\times\Omega_2) = (\operatorname{conv}\Omega_1)\times(\operatorname{conv}\Omega_2) = \mathcal{M}_1\times\mathcal{M}_2$. For $i \in \{1,2\}$, let $\nu_i \in \mathcal{M}_i$, i.e., for some probability measure P_i on Ω_i, we have $\mathbb{E}_{\omega_i\sim P_i}[\omega_i] = \nu_i$. Defining the probability measure P on Ω by $P(\omega_1,\omega_2) = P_1(\omega_1)P_2(\omega_2)$, we obtain $\mathbb{E}_{(\omega_1,\omega_2)\sim P}[(\omega_1,\omega_2)] = (\nu_1,\nu_2)$, i.e., $(\nu_1,\nu_2) \in \mathcal{M}$. Conversely, let $(\nu_1,\nu_2) \in \mathcal{M}$, i.e., for some measure P on Ω, $\mathbb{E}_{(\omega_1,\omega_2)\sim P}[(\omega_1,\omega_2)] = (\nu_1,\nu_2)$. But this also means that $\mathbb{E}_{(\omega_1,\omega_2)\sim P}[\omega_i] = \nu_i$ for $i \in \{1,2\}$, so $\nu_i \in \mathcal{M}_i$ for $i \in \{1,2\}$. Thus, $\mathcal{M} = \mathcal{M}_1\times\mathcal{M}_2$.

We also have $\operatorname{dom} C = (\operatorname{dom} C_1)\times(\operatorname{dom} C_2)$ and $\operatorname{ri}\operatorname{dom} C = (\operatorname{ri}\operatorname{dom} C_1)\times(\operatorname{ri}\operatorname{dom} C_2)$, which implies that $\tilde{\mathcal{M}} \subseteq \tilde{\mathcal{M}}_1\times\tilde{\mathcal{M}}_2$.

We next show that X is a face of $\mathcal{M}$ if and only if $X = X_1\times X_2$ where X_1 is a face of $\mathcal{M}_1$ and X_2 is a face of $\mathcal{M}_2$. A face X of $\mathcal{M}$ is characterized by a vector u and a scalar c such that

$$
\begin{aligned}
u\cdot x &= c \text{ for } x \in X \\
u\cdot\omega &> c \text{ for } \omega \in \Omega\backslash X \quad .
\end{aligned}
$$

If X_1 is a face of $\mathcal{M}_1$ characterized by u_1 and c_1, and X_2 is a face of $\mathcal{M}_2$ characterized by u_2 and c_2, then we immediately obtain that $X_1\times X_2$ is a face of $\mathcal{M}$ characterized by $u = (u_1,u_2)$ and $c = c_1+c_2$. Conversely, assume X is a face of $\mathcal{M}$ characterized by $u = (u_1,u_2)$ and c. We first show that X is a Cartesian product. We proceed by contradiction and assume that $(x_1,x_2) \in X$ and $(x_1',x_2') \in X$, but $(x_1,x_2') \notin X$. By assumption:

$$
\begin{aligned}
u_1\cdot x_1 + u_2\cdot x_2 &= c \\
u_1\cdot x_1' + u_2\cdot x_2' &= c \\
-u_1\cdot x_1 - u_2\cdot x_2' &< -c
\end{aligned}
$$

Summing the above three yields:

$$
u_1\cdot x_1' + u_2\cdot x_2 < c
$$

which is a contradiction with X being a face. By symmetry, we also obtain $(x_1',x_2) \in X$, hence $X = X_1\times X_2$ for some $X_1 \subseteq \Omega_1$ and $X_2 \subseteq \Omega_2$. Let (x_1,x_2) be some element of X. Note that X_1 must be a face of $\mathcal{M}_1$ characterized by u_1 and $c_1 := u_1\cdot x_1$, because otherwise, there would exist $\omega_1 \in \Omega_1\backslash X_1$ such that $u_1\cdot\omega_1 \leq c_1 = u_1\cdot x_1$, i.e.,

$$
u_1\cdot\omega_1 + u_2\cdot x_2 \leq u_1\cdot x_1 + u_2\cdot x_2
$$

which would contradict X being a face, since $(\omega_1,x_2) \notin X$. By symmetry, we also obtain that X_2 is a face of $\mathcal{M}_2$.

Let γ be the a-perpendicular to X at q, where $a = (a_1,a_2)$, $X = X_1\times X_2$ and $q = (q_1,q_2)$. Note that $a \in \tilde{\mathcal{M}}$, $p(q) \in \tilde{\mathcal{M}}$ implies that $a_i \in \tilde{\mathcal{M}}_i$, $p_i(q_i) \in \tilde{\mathcal{M}}_i$ because $\tilde{\mathcal{M}} \subseteq \tilde{\mathcal{M}}_1\times\tilde{\mathcal{M}}_2$. Pick a point $\nu' \in \operatorname{im}\gamma$ such that $\nu' \succeq p(q)$. From the definition of a perpendicular, we have that for $i \in \{1,2\}$ the components ν_i' lie on the a_i-perpendicular to X_i at q_i and $\nu_i' \succeq p_i(q_i)$. Let $x = (x_1,x_2) \in X$ and $\omega = (\omega_1,\omega_2) \in \Omega$. Then by acute angles assumption, we can choose $q_i' \in p_i^{-1}(\nu_i') \cap [q_i + \mathcal{K}(X_i)]$. Let $q' = (q_1',q_2')$. Note that $q' \in p^{-1}(\nu')$. We will argue that also $(q'-q) \in \mathcal{K}(X)$:

$$
(q'-q)\cdot(\omega-x) = \sum_{i\in\{1,2\}} (q_i'-q_i)\cdot(\omega_i-x_i) \geq 0 \quad ,
$$

where the last inequality follows, because $(q_i'-q_i) \in \mathcal{K}(X_i)$. Thus, the acute angles assumption holds for C and Ω. □

H Proof of Theorem 5.2

In this section we give the complete proof of Theorem 5.2. The proof proceeds in several steps. Let $\nu_0 = p(q_0)$. We begin by constructing an oriented curve L joining ν_0 with μ, by sequentially choosing portions of perpendiculars for monotonically decreasing active sets. We then show that budget additivity holds for any solutions with prices in L, and finally show that the curve L is the locus of the optimal prices of solutions $\hat{Q}(q_0)$, as well as optimal prices of solutions $\hat{Q}(q)$ for any $q \in \hat{Q}(q_0)$.

Part 1: Construction of the solution path L In this part, we construct:

- a sequence of prices $\nu_0, \nu_1, \ldots, \nu_k$ with $\nu_0 = p(q_0)$ and $\nu_k = \mu$
- a sequence of oriented curves $\ell_0, \ldots, \ell_{k-1}$ where each ℓ_i goes from ν_i to ν_{i+1}
- a monotone sequence of sets $\Omega \supseteq X_0 \supset X_1 \supset \cdots \supset X_k = \emptyset$, such that the following *minimality* property holds: X_i is the minimal face for all $\nu \in (\operatorname{im} \ell_i) \backslash \{\nu_{i+1}\}$ for $i \leq k-1$, and X_k is the minimal face for ν_k.
- a sequence of states $q_1, \ldots, q_{k-1}$ such that $q_i \in p^{-1}(\nu_i) \cap [q_{i-1} + \mathcal{K}(X_{i-1})]$

The curves ℓ_i will be referred to as *segments*. The curve obtained by concatenating the segments ℓ_0 through ℓ_{k-1} will be called the *solution path* and denoted L. In the special case that $\nu_0 = \mu$, we have $k = 0$, $X_0 = \emptyset$ and L is a degenerate curve with $\operatorname{im} L = \{\mu\}$.

If $\nu_0 \neq \mu$, we construct the sequence of segments iteratively. Let $X_0 \neq \emptyset$ be the minimal face such that $\nu_0 \in \operatorname{conv}(X_0 \cup \{\mu\})$. By the minimality, $\mu \notin \operatorname{aff}(X_0)$. Let γ be the μ-perpendicular to $\operatorname{aff}(X_0)$ at q_0. The curve γ passes through ν_0 and eventually reaches the boundary of $\operatorname{conv}(X_0 \cup \{\mu\})$ at some ν_1 by continuity of γ. Let segment ℓ_0 be the portion of γ going from ν_0 to ν_1.

This construction gives us the first segment ℓ_0. There are two possibilities:

1. $\nu_1 = \mu$; in this case we are done;
2. ν_1 lies on a lower-dimensional face of $\operatorname{conv}(X_0 \cup \{\mu\})$; in this case, we pick some $q_1 \in p^{-1}(\nu_1) \cap [q_0 + \mathcal{K}(X_0)]$, which can be done by the acute angles assumption, and use the above construction again, starting with q_1, and obtaining a new set $X_1 \subset X_0$ and a new segment ℓ_1; and iterate.

The above process eventually ends, because with each iteration, the size of the active set decreases. This construction yields monotonicity of X_i as well as the minimality property.

The above construction yields a specific sequence of $q_i \in p^{-1}(\nu_i) \cap [q_{i-1} + \mathcal{K}(X_{i-1})]$. We will now show that actually $q_i \in p^{-1}(\nu_i) \cap (q_0 + X_{i-1}^{\perp})$ and that the construction of L is independent of the specific $q_1, q_2, \ldots, q_{k-1}$ chosen. To begin, note that from our construction, we can write $q_i = q_0 + u_0 + u_1 + \ldots + u_{i-1}$ for some $u_j \in \mathcal{K}(X_j) \subseteq X_j^{\perp}$. Since $X_i \subseteq X_j$ for $j = 1, \ldots, i-1$, we actually have $u_j \in X_{i-1}^{\perp}$, so $q_i \in (q_0 + X_{i-1}^{\perp})$. Note that $X_{i-1}^{\perp} \subseteq X_i^{\perp}$, and according to Proposition 4.3, any $q_i \in p^{-1}(\nu_i) \cap (q_0 + X_i^{\perp})$ yields the same μ-perpendicular to $\operatorname{aff}(X_i)$ and hence the same segment ℓ_i. By induction it therefore follows that the segments $\ell_0, \ldots, \ell_{k-1}$ are uniquely determined by our construction regardless of the specific $q_1, \ldots, q_{k-1}$.

Part 1': The solution path starting at a midpoint Let $\nu \in \operatorname{im} L$, and $q \in p^{-1}(\nu) \cap (q_0 + X_{\nu}^{\perp})$, and let L' be the solution path if the initial state were q rather than q_0. By a similar reasoning as in the previous paragraph, we see that L' is a restriction of L starting with ν.

Part 2: Budget additivity for points on L Let $\nu, \nu' \in \operatorname{im} L$ such that $\nu \preceq \nu'$. Let $q \in \hat{Q}(\nu; q_0)$ and $q' \in \hat{Q}(\nu'; q)$ such that $q \in \hat{Q}(B; q_0)$ and $q' \in \hat{Q}(B'; q)$. In this part we show that $q' \in \hat{Q}(B + B'; q_0)$.

First, consider the case that $\nu' = \mu$. To see that $q' \in \hat{Q}(B + B'; q_0)$, first note that the constraints of Convex Program (2.1) hold, because

$$U(q', \omega; q_0) = U(q', \omega; q) + U(q, \omega; q_0) \geq -B' - B \quad \forall \omega \in \Omega$$

by path independence of the utility function. As noted in the introduction, in the absence of constraints, the utility $U(\bar{q}, \mu; q_0)$ is maximized at any $\bar{q}$ with $p(\bar{q}) = \mu$. Thus, q' is a global maximizer of the utility and satisfies the constraints, so $q' \in \hat{Q}(B + B'; q_0)$. If $\nu = \mu$, we must also have $\nu' = \mu$ and the statement holds by previous reasoning.

In the remainder, we only analyze the case $\nu \preceq \nu' \prec \mu$. This means that $\nu \in (\operatorname{im} \ell_i) \backslash \{\nu_{i+1}\}$ and $\nu' \in (\operatorname{im} \ell_j) \backslash \{\nu_{j+1}\}$ for $i \leq j$. By Theorem 3.4, we therefore must have $q \in [q_0 + \mathcal{K}(X_i)]$ and $q' \in [q + \mathcal{K}(X_j)]$. By anti-monotonicity of witness cones, $\mathcal{K}(X_j) \supseteq \mathcal{K}(X_i)$ and hence, $q' \in [q_0 + \mathcal{K}(X_j)]$, yielding

$$q' \in \hat{Q}(\nu'; q_0) \ .$$

Let $x \in X_j \subseteq X_i$. Further by Lemma 3.1, we obtain that

$$\begin{aligned} &q \in \hat{Q}(B; q_0) \text{ for } B = -U(q, x; q_0) \\ &q' \in \hat{Q}(B'; q) \text{ for } B' = -U(q', x; q) \\ &q' \in \hat{Q}(\bar{B}; q_0) \text{ for } \bar{B} = -U(q', x; q_0) \end{aligned}$$

However, by path independence of the utility function

$$\bar{B} = -U(q', x; q_0) = -U(q', x; q) - U(q, x; q_0) = B' + B \ .$$

Part 3: L as the locus of all solutions In this part we show that

$$\hat{Q}(q_0) = \bigcup_{\nu \in \operatorname{im} L} \hat{Q}(\nu; q_0) \ .$$

We will begin by defining sets of budgets for which the optimal price is ν and show that their union across all $\nu \in \operatorname{im} L$ is a closed interval. Since both $\nu_0, \mu \in \operatorname{im} L$, this will mean that we have included prices across all possible budgets. The statement of Part 3 will then follow by Theorem 2.2.

Let $x \in X_{k-1} = \bigcap_{i=0}^{k-1} X_i$ and let $B(q) \coloneqq -U(q, x; q_0)$. Further, for $\nu \in \operatorname{im} \ell_i$, let

$$B_i(\nu) \coloneqq \{B(q) : \, q \in p^{-1}(\nu) \cap [q_0 + \mathcal{K}(X_i)]\} \ .$$

From Corollary 3.5, we know that for $\nu \in (\operatorname{im} \ell_i) \backslash \{\nu_{i+1}\}$, $B_i(\nu)$ is exactly the set of budgets for which ν is the optimal price vector. The set $B_i(\nu_{i+1})$ is potentially only a subset of such budgets (corresponding to X_i being the tight set, rather than the actual minimal set X_{i+1}). First we show that $B_i(\nu)$ is non-empty for $\nu \in \operatorname{im} \ell_i$. Let $\nu \in \operatorname{im} \ell_i$. By acute angles assumption, there exists $q \in p^{-1}(\nu) \cap [q_i + \mathcal{K}(X_i)]$. Furthermore, $q_j \in [q_{j-1} + \mathcal{K}(X_{j-1})]$ for $j = 1, \ldots, i$, so we can write $q = q_0 + u_0 + \cdots + u_i$ where $u_j \in \mathcal{K}(X_j)$. By anti-monotonicity of witness cones, $\mathcal{K}(X_j) \subseteq \mathcal{K}(X_i)$ for $j = 1, \ldots, i$, so we actually have $u_j \in \mathcal{K}(X_i)$ and thus $q \in [q_0 + \mathcal{K}(X_i)]$, proving that the set $B_i(\nu)$ is non-empty.

We will next show that

$$B_i(\ell_i) \coloneqq \bigcup_{\nu \in \operatorname{im} \ell_i} B_i(\nu)$$

is an interval.

Consider a fixed $\nu \in \operatorname{im} \ell_i$. For $q \in p^{-1}(\nu)$, we have $C(q) = q \cdot \nu - C^*(\nu)$, i.e., $B(q)$ is linear in q over $q \in p^{-1}(\nu)$. Since the set $p^{-1}(\nu)$ is closed and convex, so is $p^{-1}(\nu) \cap [q_0 + \mathcal{K}(X_i)]$. The latter set is also non-empty, hence the set $B_i(\nu)$ must be a non-empty closed interval. Let $B_i^{\min}(\nu)$ and $B_i^{\max}(\nu)$ be the lower and upper endpoints of $B_i(\nu)$. Since the budget additivity holds along L (by Part 2), we must have that $B_i^{\max}$ is non-decreasing on ℓ_i. Next note that for $\nu \neq \nu'$ the sets $B_i(\nu)$ and $B_i(\nu')$ must be disjoint. This implies that $B_i^{\max}$ is actually increasing and so is $B_i^{\min}$.

We next show that $B_i^{\max}$ is right-continuous on ℓ_i. Let $M \coloneqq \{(\nu, q) : \, \nu \in \operatorname{im} \ell_i, \, q \in p^{-1}(\nu)\}$. By Theorem F.3, the set M can be written $\mathcal{C} + (0 \oplus \mathcal{L}^\perp)$ where $\mathcal{C}$ is compact. Let $M' \coloneqq \{(\nu, q) : \nu \in \operatorname{im} \ell_i, \, q \in p^{-1}(\nu) \cap [q_0 + \mathcal{K}(X_i)]\}$. Since the set $[q_0 + \mathcal{K}(X_i)]$ is closed, the set M' can be written as $\mathcal{C}' + (0 \oplus \mathcal{L}^\perp)$ where $\mathcal{C}'$ is compact. To show that $B_i^{\max}$ is right-continuous, pick $\nu \in \operatorname{im} \ell_i$ and let $\{\nu'_t\}_{t=1}^\infty$ be a sequence of $\nu'_t \in \ell_i$, $\nu'_t \succeq \nu$ such that $\lim_{t \to \infty} \nu'_t = \nu$. Pick q'_t such that $(\nu'_t, q'_t) \in \mathcal{C}'$ and $B(q'_t) = B_i^{\max}(\nu'_t)$. By compactness, the sequence $\{(\nu'_t, q'_t)\}_{t=1}^\infty$ must have a cluster point $(\nu, q) \in \mathcal{C}'$, and by continuity of B, we must have $\lim_{t \to \infty} B_i^{\max}(\nu'_t) = \lim_{t \to \infty} B(q'_t) = B(q) \leq B_i^{\max}(\nu)$. The right continuity of $B_i^{\max}$ now follows by monotonicity. By symmetric reasoning, $B_i^{\min}$ must be left-continuous.

Now for the sake of contradiction, assume that $B_i(\ell_i)$ is not an interval, i.e., assume that there is a value $B^* \notin B_i(\ell_i)$ such that some higher and lower values are in $B_i(\ell_i)$. By monotonicity, there must exist ν^* such that $B_i^{\max}(\nu) < B^*$ for $\nu \prec \nu^*$ and $B_i^{\min}(\nu) > B^*$ for $\nu \succ \nu^*$. However, this means that

$$B_i^{\min}(\nu^*) = \lim_{\nu \uparrow \nu^*} B_i^{\min}(\nu^*) \leq \lim_{\nu \uparrow \nu^*} B_i^{\max}(\nu^*) \leq B^*$$

and

$$B_i^{\max}(\nu^*) = \lim_{\nu \downarrow \nu^*} B_i^{\max}(\nu^*) \geq \lim_{\nu \downarrow \nu^*} B_i^{\min}(\nu^*) \geq B^*$$

which means that $B^* \in B_i(\nu^*)$ yielding a contradiction.

Finally, note that $B_i(\nu_{i+1}) \subseteq B_{i+1}(\nu_{i+1})$ for $i \leq k-1$, hence $\bigcup_{i=1}^{k-1} B_i(\ell_i)$ is an interval as well.

Part 3': L as the locus of solutions starting at a mid-point Let $\nu \in \operatorname{im} L$ and $q \in \hat{Q}(\nu; q_0)$. Since $\hat{Q}(\nu; q_0) \subseteq p^{-1}(\nu) \cap (q_0 + X_\nu^\perp)$, Part 1' yields that the solution path L' for q coincides with the portion of L starting at ν. The reasoning of the previous part applied to L' then yields the following statement:

$$\hat{Q}(q) = \bigcup_{\substack{\nu' \in \operatorname{im} L: \\ \nu' \succeq \nu}} \hat{Q}(\nu'; q) \ .$$

Part 4: Proof of the theorem Let $B, B' \geq 0$ and $q \in \hat{Q}(B; q_0)$ and $q' \in \hat{Q}(B'; q)$. From Parts 3 and 3', we know that $q \in \hat{Q}(\nu; q_0)$ and $q' \in \hat{Q}(\nu'; q)$ for some $\nu, \nu' \in \operatorname{im} L$ such that $\nu \preceq \nu'$. By Part 2, we therefore obtain that $q' \in \hat{Q}(B + B'; q_0)$, proving the theorem.

I Cone Projection Lemma

Proof of Lemma 5.9. We give a constructive proof, by giving explicitly the required projection. We use the result of Borwein and Lewis [3, Theorem 3.3.15] that shows that a cone is pointed if and only if there is a conic section that is bounded. That is, $\mathcal{K}$ is pointed if and only if there is some $u \in \mathcal{K}$ with $\|u\| = 1$ such that the set

$$S \coloneqq \{x \in \mathcal{K} : \, x \cdot u = 1\}$$

is bounded.

Since the given cone $\mathcal{K}$ is pointed, let u be one such vector in $\mathcal{K}$. Since $\mathcal{K}$ is obtuse, we know that there are some two vectors in $\mathcal{K}$ whose dot product is negative; let these be $w, w' \in \mathcal{K}$. Without loss of generality, we may assume that both w, w' lie in the conic section S. In other words, we may assume that $w = u + v$ and $w' = u + v'$ where $u \cdot v = u \cdot v' = 0$. (Also w.l.o.g. $v \neq 0$.) We consider the projection of $\mathcal{K}$ onto $\operatorname{span}\{u, v\}$ and call it $\mathcal{K}'$. We show that $\mathcal{K}'$ is pointed and obtuse.

We use the following easy observation in the rest of the proof.

Observation: Let x' be the projection of $x \in \mathbb{R}^n$ onto some subspace Y. Then for all $y \in Y$, $x \cdot y = x' \cdot y$. Note that x can be written as $x' + x^\perp$ where $x^\perp$ is orthogonal to Y. Since $y \in Y$, it follows that $x^\perp \cdot y = 0$. Therefore $x' \cdot y = x \cdot y$.

$\mathcal{K}'$ **is pointed.** We obtain this by considering the conic section defined by u in $\mathcal{K}'$ and showing that this is bounded. That is, the set $S' = \{x \in K' : x \cdot u = 1\}$ is bounded. This follows from the fact that S' is the projection of S and since S is bounded, so is S'. To see this, take any $x' \in S'$ and let $x \in \mathcal{K}$ be any preimage of x' under the projection. Then from the observation above, $x \cdot u = x' \cdot u = 1$ therefore $x \in S$. Vice versa, it is easy to see that for any $x \in S$, its projection x' is in S'. Thus we have that $\mathcal{K}'$ is pointed.

$\mathcal{K}'$ **is obtuse.** Let $w'' \in \mathcal{K}'$ be the projection of $w' \in \mathcal{K}$. Note that by definition w is already in $\mathcal{K}'$. Again using the observation above, $w'' \cdot w = w' \cdot w < 0$ and $\mathcal{K}'$ is obtuse. $\square$