
Vector Optimization with Stochastic Bandit Feedback

Çağın Ararat*

Bilkent University, Ankara, Turkey

Cem Tekin*

Bilkent University, Ankara, Turkey

Abstract

We introduce vector optimization problems with stochastic bandit feedback, in which preferences among designs are encoded by a polyhedral ordering cone C . Our setup generalizes the best arm identification problem to vector-valued rewards by extending the concept of Pareto set beyond multi-objective optimization. We characterize the sample complexity of (ϵ, δ) -PAC Pareto set identification by defining a new cone-dependent notion of complexity, called the *ordering complexity*. In particular, we provide gap-dependent and worst-case lower bounds on the sample complexity and show that, in the worst-case, the sample complexity scales with the square of ordering complexity. Furthermore, we investigate the sample complexity of the naïve elimination algorithm and prove that it nearly matches the worst-case sample complexity. Finally, we run experiments to verify our theoretical results and illustrate how C and sampling budget affect the Pareto set, the returned (ϵ, δ) -PAC Pareto set, and the success of identification.

1 INTRODUCTION

We consider a vector optimization problem with D objectives and K designs, where D and K are positive integers. The mean vector of design $i \in [K]$ is given by $\mu_i \in \mathbb{R}^D$. The mean vectors are partially ordered based on an ordering cone $C \subseteq \mathbb{R}^D$. When these vectors are known, the set of all Pareto optimal designs (Pareto set) can be found by solving the following optimization problem (Jahn, 2011)

$$\text{maximize } \mu_i \text{ with respect to cone } C \text{ over } i \in [K]. \quad (1)$$

*Equal contribution and alphabetical ordering.

In this paper, for the first time in the literature, we investigate vector optimization problems in a pure exploration setting with stochastic bandit feedback, i.e., when μ_i s are unknown and the query of each design i only yields a noisy observation of μ_i . In particular, we seek an answer to the following fundamental question: *What is the sample complexity of Pareto set identification and how does it depend on C ?*

In Section 1.1, by comparing vector optimization with scalar and multi-objective optimization within the context of pure exploration, we argue why one should consider solving problems like (1). In Section 1.2, we detail our main contributions.

1.1 Motivation

Pure exploration problems have attracted significant interest from both machine learning theorists and practitioners (Even-Dar et al., 2006; Bubeck et al., 2009; Audibert et al., 2010; Kalyanakrishnan et al., 2012; Karnin et al., 2013; Kaufmann and Kalyanakrishnan, 2013; Zhou et al., 2014; Laskey et al., 2015). The simplest pure exploration problem takes the form of a K -armed bandit with arm means $\mu_1, \dots, \mu_K \in \mathbb{R}$, in which the learner aims to identify an optimal arm (aka design) by sequential experimentation under noisy feedback. In this setup, a noisy scalar reward is revealed to the learner immediately after the selection of a design. As experimentation consumes resources, the learner seeks to minimize the number of evaluations by adapting its arm selection based on past reward observations. This problem has been formalized in many different ways, with notable examples including (ϵ, δ) -probably approximately correct (PAC) (Even-Dar et al., 2006) (aka Explore-1), fixed confidence and fixed budget (Karnin et al., 2013) best arm identification. Almost all research in this field focuses on devising sample-efficient algorithms that “beat the noise” by adaptive sampling and elimination. Algorithmic contributions include a plethora of techniques such as Median Elimination (Even-Dar et al., 2006), LUCB (Kalyanakrishnan et al., 2012), KL-LUCB (Kaufmann and Kalyanakrishnan, 2013), Sequential Halving (Karnin et al., 2013), Track-and-Stop (Garivier and Kaufmann, 2016), ABA (Hassidim et al., 2020), ϵ -TaS (Garivier and Kaufmann, 2021), several Bayesian sampling

strategies (Qin et al., 2017; Shang et al., 2020; Russo, 2020), and saddle-point methods (Degenne et al., 2019, 2020). The difficulty of the problem depends on the suboptimality gap of each arm i , straightforwardly found as $\max_j \mu_j - \mu_i$. Without noise, everything turns into a simple argmax operation over the design set.

However, not many real-world problems naturally exhibit scalar rewards. Optimization of D -dimensional performance metrics ($D > 1$) is necessary for tasks such as hardware design (Zuluaga et al., 2016), and clinical trials for drug development and dose identification (Lizotte and Laber, 2016). Nevertheless, vector-valued objectives μ_i can be scalarized by weights w that encode the importance of each objective, turning the learning problem into a scalar one. The choice of w is often left to the practitioner, and the right choice might be difficult to come up with.

To tackle this issue, another strand of literature focuses on identifying the set P^* of all Pareto optimal designs, whose mean vectors μ_i are not dominated by others' mean vectors. In particular, when $D = 2$, $\mu_i = (\mu_i^1, \mu_i^2)^\top$ is dominated by $\mu_j = (\mu_j^1, \mu_j^2)^\top$ ($\mu_i \preceq \mu_j$) if and only if (iff) $\forall w \in [0, 1] : w\mu_i^1 + (1-w)\mu_i^2 \leq w\mu_j^1 + (1-w)\mu_j^2$, which is equivalent to the usual componentwise order on \mathbb{R}^2 . This multi-objective optimization problem has been extensively studied in the pure exploration setting with stochastic bandit feedback. Auer et al. (2016) study the sample complexity of Pareto set identification in stochastic K -armed bandit problems. Hernández-Lobato et al. (2016), Shah and Ghahramani (2016), Zuluaga et al. (2016) study Pareto set identification in problems with large design sets. They use Gaussian processes to capture correlations between designs such that a bulk of designs can be explored by a single sample. In another related work, Katz-Samuels and Scott (2018) propose the feasible arm identification problem, in which the goal is to identify arms whose mean rewards belong to a given polyhedron through noisy evaluations.

Apart from pure exploration, multi-objective learning has also been extensively investigated in the bandit regret minimization setting (Drugan and Nowe, 2013; Turgay et al., 2018). This line of research defines gaps that represent the suboptimality of arms, and propose algorithms that minimize the cumulative sum of gaps of the selected arms while ensuring some sort of balance between selections of Pareto optimal arms. Another related line of research is multi-objective reinforcement learning (Van Moffaert and Nowé, 2014; Hayes et al., 2022).

While the multi-objective optimization viewpoint saves the practitioner from fixing a weight vector for the objectives, it ends up returning a large set of designs (see, e.g., $P_{\pi/2}^*$ in Figure 1), which might be even more frustrating. Moreover, the practitioners may want to narrow down the set of alternatives by using their domain knowledge about the relative importance of each objective. For instance, if the




Figure 1: Pareto sets $P_{\pi/4}^*$, $P_{\pi/2}^*$, $P_{3\pi/4}^*$ of SNW dataset used in experiments (Section 7) for the cones $C_{\pi/4}$, $C_{\pi/2}$, $C_{3\pi/4}$ (Example 2.6). By the structure of the cones, $P_{3\pi/4}^* \subseteq P_{\pi/2}^* \subseteq P_{\pi/4}^*$.




Figure 2: Three cones in \mathbb{R}^2 (intersected with balls for illustration purposes). Orange, blue, green cones correspond to $C_{\pi/4}$, $C_{\pi/2}$, $C_{3\pi/4}$ defined in Example 2.6, and used in experiments in Section 7. C' corresponds to a wide cone like the green one, i.e., $C' \supseteq \mathbb{R}_+^2$.

practitioner wants to give at least $100\alpha\%$ relative importance to each objective for $\alpha \in (0, 0.5)$, then this can be achieved by defining a new partial order $\preceq_{C'}$ with $\mu_i \preceq_{C'} \mu_j$ iff $\forall w \in [\alpha, 1 - \alpha] : w\mu_i^1 + (1-w)\mu_i^2 \leq w\mu_j^1 + (1-w)\mu_j^2$. This partial order is equivalent to a partial order induced by the cone¹

$$C' = \{x \in \mathbb{R}^2 \mid \alpha x_1 + (1 - \alpha)x_2 \geq 0, (1 - \alpha)x_1 + \alpha x_2 \geq 0\}. \quad (2)$$

Such a choice usually narrows down the set of all Pareto optimal designs (see, e.g., $P_{3\pi/4}^*$ in Figure 1). In addition, the flexibility of specifying the relative importance of each objective leads to a wide spectrum of Pareto sets that cannot be captured by multi-objective optimization. The aim

¹See Figure 2 for examples of ordering cones in \mathbb{R}^2 , and Definition 2.8 for the precise definition of this order.

of this work is to propose a new framework for pure exploration problems in which the preferences of practitioners are encoded by ordering cones.

One real-world example in which multi-objective optimization does not suffice is intensity-modulated radiation therapy (IMRT) planning (Teichert et al., 2019). Here, the designs represent treatment strategies (e.g., dosage). The objectives are related to the effectiveness of therapy against the tumor and the damage to the surrounding tissue, which are in conflict with each other. Based on the relative importance of each objective encoded by the cone parametrized by $\alpha \in (0, 0.5)$, the practitioner might want to identify a set of Pareto optimal treatments via computer simulations. Note that, in this application, returning the Pareto set of designs according to the componentwise order can be harmful for the patient. For instance, a treatment which achieves a large reduction in tumor size at the cost of extensive damage to the surrounding healthy tissue could be returned if componentwise order were used.

Example 1.1. *We illustrate the limitation of using the componentwise order in IMRT via a simple example. Take $D = 2$, $K = 4$. The first dimension represents the effectiveness against the tumor. Assume that the best value is 1 (e.g., the tumor is completely destroyed) and the worst value is 0 (e.g., the tumor does not shrink at all). The second dimension represents the negative of the damage to the surrounding tissue. Let the best value be 0 (e.g., the surrounding tissue is not damaged at all) and the worst value be -1 (e.g., the surrounding tissue is completely destroyed). Let $\mu_1 = [0.8, -0.3]^\top$, $\mu_2 = [0.9, -0.5]^\top$, $\mu_3 = [0, 0]^\top$, $\mu_4 = [1, -1]^\top$. The Pareto set under the componentwise order (i.e., the classical multi-objective case) includes all four designs for this setting. Design 3 is useless since it has no effect on the tumor. Design 4 is the most effective against the tumor but it causes significant damage to the surrounding tissue; thus, it should not be returned as a treatment option. By returning all four designs, the componentwise order does not provide any useful guidance to the practitioner. On the other hand, the Pareto set with cone C parameterized by relative importance parameter $\alpha \approx 0.29$ (equivalent to cone with $\theta = 3\pi/4$ in Example 2.6 below) includes only designs 1 and 2, both of which provide meaningful compromise between effectiveness and side-effects.*

1.2 Contributions

The subfield of mathematical optimization which generalizes scalar and multi-objective optimization via partial orders induced by cones is called vector optimization (Jahn, 2011; Löhne, 2011). While there exists a plethora of works that focuses on multi-objective optimization through noisy evaluations (Auer et al., 2016; Hernández-Lobato et al., 2016; Zuluaga et al., 2016; Shah and Ghahramani, 2016), our work is the first to consider vector optimization for partial orders in \mathbb{R}^D induced by general polyhedral ordering

cones within the stochastic K -armed bandit framework.

Our first contribution is to generalize the (ϵ, δ) -PAC best arm identification problem in Even-Dar et al. (2006) to vector optimization (Definition 4.6). Our success condition takes into account the directions of improvement specified by C and the uncertainty induced by the norm-subgaussian sampling noise. In line with PAC best arm identification, instead of trying to identify all Pareto optimal designs, we introduce a direction-free, cone-dependent ϵ -covering requirement for the Pareto set. Based on this, we define the (ϵ, δ) -PAC Pareto set identification problem.

Our second contribution is to define the notion of *ordering complexity* β , determined by C (equations (3), (4)). The ordering complexity of C characterizes the sampling budget required for (ϵ, δ) -PAC Pareto set identification. We explain how it can be computed for any polyhedral ordering cone C (Theorem 2.4 in the main text and the theorems in the supplemental document).

Our third contribution is to characterize the learning difficulty of a vector optimization problem by identifying fundamental gaps associated with the geometric properties of the ordering cone C (equations (8), (9)). We show that when mean vectors are known, these gaps can be calculated by solving convex optimization problems involving quadratic and affine functions (Propositions 4.2, 4.3). Motivated by the fact that the sampling noise can shift mean estimates in arbitrary directions, these gaps extend the direction-based gaps proposed for multi-objective problems in Auer et al. (2016) to arbitrary polyhedral ordering cones and arbitrary directions. In particular, we rely on the interpretation of the ordering cone as the set of all directions of improvement, and define two gaps between designs i, j : $m(i, j)$ as the minimum step-length in an arbitrary direction of improvement for i that is sufficient to avoid i being dominated by j , $M(i, j)$ as the minimum step-length in an arbitrary direction of improvement for j that is sufficient to make j dominate i . For the special case of multi-objective problems, while $m(i, j)$ coincides with its analogue in Auer et al. (2016), the analogue of $M(i, j)$ in Auer et al. (2016) gives higher values since the improvement is restricted to the diagonal direction.² Hence, our definition of $M(i, j)$ provides a more flexible gap calculation even in the multi-objective case.

Our fourth contribution is to provide gap-dependent and worst-case lower bounds on the sample complexity. Our lower bounding techniques involve constructing modified problems via perturbations of mean reward vectors along special directions in the objective space related to the

²When $C = \mathbb{R}_+^D$, our $m(i, j)$ is the same as Auer et al. (2016) but our $M(i, j)$ gives $\|(\mu_j - \mu_i)^-\|_2 = \|(\mu_i - \mu_j)^+\|_2$ (by Proposition 4.3(ii)), and it is larger than $\|(\mu_i - \mu_j)^+\|_\infty$, which is the value of $M(i, j)$ in Auer et al. (2016). In the worst-case, their $M(i, j)$ is \sqrt{D} factor more conservative than ours.

geometry of the ordering cone. This makes our lower bound analysis substantially different from the analysis of scalar or multi-objective problems. In particular, we prove that the worst-case sample complexity scales as $\Omega((K\beta_2^2/\epsilon^2)\log(1/\delta))$, where β_2 is one of the two constants whose maximum define the ordering complexity.

As our final contribution, we take a first step towards Pareto set identification in vector optimization (beyond multi-objective optimization) by proposing a naïve elimination algorithm. In particular, we perform sample complexity analysis under the general norm-subgaussian noise assumption. We show that our algorithm successfully identifies an (ϵ, δ) -PAC Pareto set with a sample complexity of $O((K\beta^2/\epsilon^2)\log(1/\delta))$ (Theorem 6.1). In addition, we provide numerical results for parametrized polyhedral cones, and establish a connection between the success rate of Pareto set identification and the angular width of the cone. The details of all proofs are given in the supplemental document.

2 THE ORDERING CONE

In this section, we fix the notation for the rest of the paper, introduce the order structure for vector optimization, and state some important properties of this structure. Let D be a positive integer. We write $[D] := \{1, \dots, D\}$ and denote by \mathbb{R}^D the D -dimensional Euclidean space. The elements of \mathbb{R}^D are denoted by boldface letters. For a vector $\mathbf{v} \in \mathbb{R}^D$, $\|\mathbf{v}\|_2$ represents its ℓ_2 norm. For subsets A, A' of \mathbb{R}^D , $A + A'$ and $A - A'$ denote their Minkowski sum and difference, respectively; $\text{cl}(A)$, $\text{int}(A)$, $\text{bd}(A)$, A^c denote the closure, interior, boundary, and complement of A in \mathbb{R}^D , respectively. For $\mathbf{v} = (v^1, \dots, v^D)^\top \in \mathbb{R}^D$, we write $\mathbf{v}^+ := ((v^1)^+, \dots, (v^D)^+)^\top$ and $\mathbf{v}^- := ((v^1)^-, \dots, (v^D)^-)^\top$, where $r^+ := \max\{0, r\}$ and $r^- := -\min\{0, r\}$ for $r \in \mathbb{R}$. For $\mathbf{v} \in \mathbb{R}^D$ and $r \geq 0$, $B(\mathbf{v}, r)$ represents the ball in \mathbb{R}^D with center \mathbf{v} and radius r . The distance of a vector $\mathbf{v} \in \mathbb{R}^D$ to a set $A \subseteq \mathbb{R}^D$ is defined as $d(\mathbf{v}, A) := \inf_{\mathbf{x} \in A} \|\mathbf{v} - \mathbf{x}\|_2$.

We consider a vector optimization problem with D objectives and a finite set $[K]$ of designs, where D, K are positive integers. The mean vector of design $i \in [K]$ is denoted by $\boldsymbol{\mu}_i = (\mu_i^1, \dots, \mu_i^D)^\top \in \mathbb{R}^D$. To compare mean vectors, we will introduce a partial order on \mathbb{R}^D based on an ordering cone. The latter notion is recalled next.

Definition 2.1. A set $C \subseteq \mathbb{R}^D$ is called a *cone* if $\lambda \mathbf{v} \in C$ for every $\mathbf{v} \in C$ and $\lambda \geq 0$. A cone C is called *pointed* if $C \cap (-C) = \{\mathbf{0}\}$, it is called *solid* if $\text{int}(C) \neq \emptyset$. A closed convex cone that is pointed and solid is called an *ordering cone* (proper cone).

Let $C \subseteq \mathbb{R}^D$ be an ordering cone. As a consequence of the convexity of C , it is immediate that $C + C = C$. In general, an ordering cone can be polyhedral (e.g., the positive

orthant $C = \mathbb{R}_+^D$) or non-polyhedral (e.g., the ice-cream cone $C = \{\mathbf{x} \in \mathbb{R}^D \mid \|(x^1, \dots, x^{D-1})^\top\|_2 \leq x^D\}$ for $D \geq 2$). For the purposes of this paper, we focus on the polyhedral case as detailed in the next definition.

Definition 2.2. A cone C is called *polyhedral* if it can be written as $C = \{\mathbf{x} \in \mathbb{R}^D \mid W\mathbf{x} \geq \mathbf{0}\}$ for some $N \times D$ real matrix W with rows $\mathbf{w}_1^\top, \dots, \mathbf{w}_N^\top$, and positive integer N .

Figure 2 illustrates several examples of polyhedral ordering cones in \mathbb{R}^2 . Throughout the paper, we assume that C is a polyhedral ordering cone as in Definition 2.2.³ Without loss of generality, we assume that this description has no redundancies, that is, W has the minimal number of rows; as well as that $\|\mathbf{w}_n\|_2 = 1$ for each $n \in [N]$.⁴ It follows that the interior of C is given by $\text{int}(C) = \{\mathbf{x} \in \mathbb{R}^D \mid W\mathbf{x} > \mathbf{0}\}$. When all entries of W are nonnegative, it is clear that $C \supseteq \mathbb{R}_+^D$. In financial mathematics, such cones have found applications in multi-asset markets with transaction costs as “solvency cones”, where W is defined in terms of the bid-ask prices of the assets; see Kabanov (1999).

Remark 2.3. As discussed in Section 1.1, practitioners can prefer large cones ($C \supset \mathbb{R}_+^D$) to narrow down the set of Pareto optimal designs. On the other hand, some applications can benefit from using small cones ($C \subset \mathbb{R}_+^D$). One motivating example is the small molecule drug discovery problem (Jayatunga et al., 2022). Consider the optimization of D properties such as solubility, metabolic stability, toxicity. Assume that the final goal is to identify the Pareto optimal molecules in $P_{\pi/2}^*$ from a set of candidates $[K]$ according to cone \mathbb{R}_+^D . The discovery process starts with (relatively cheaper) *in silico* experiments followed by (more expensive) wet lab experiments. Due to distribution mismatch, running *in silico* experiments with $C = \mathbb{R}_+^D$ could result in failing to return an (ϵ, δ) -PAC set for $C = \mathbb{R}_+^D$. By using a smaller cone $C \subset \mathbb{R}_+^D$ for *in silico* experiments, the practitioner can ensure that more designs are passed to the wet lab stage $\hat{P}_C \supseteq \hat{P}_{\pi/2}$ and, since all Pareto optimal designs according to \mathbb{R}_+^D are still in C , one can perform more expensive wet lab experiments on these designs to ensure that an (ϵ, δ) -PAC set is returned. In this case, the knowledge of an upper bound on the distribution mismatch can be used to select C .

In order to characterize the sample complexity of Pareto set identification, we introduce the following constants that are related to the geometry of C :

$$\beta_1 := \sup_{\mathbf{x} \notin C} \frac{d(\mathbf{x}, C \cap (\mathbf{x} + C))}{d(\mathbf{x}, C)}, \quad (3)$$

$$\beta_2 := \sup_{\mathbf{x} \in \text{int}(C)} \frac{d(\mathbf{x}, (\text{int}(C))^c \cap (\mathbf{x} - C))}{d(\mathbf{x}, (\text{int}(C))^c)}. \quad (4)$$

³Cone C represents the preferences of the practitioner. It is treated as an input.

⁴Unless stated otherwise, all of the results in this paper hold when C is a polyhedral ordering cone without redundancy and $\|\mathbf{w}_n\|_2 = 1$ for each $n \in [N]$.

These constants, which will appear in the crux of our sample complexity analysis, are novel to our work. We will provide intuitive explanations for these constants after defining fundamental gaps associated with Pareto set identification in Section 4 (see Remark 4.9). In this section, we analyze their technical properties. Note that we have $\beta_1 \geq 1$ and $\beta_2 \geq 1$. Let $\beta := \max\{\beta_1, \beta_2\}$. We call β the *ordering complexity* of C .⁵ The following theorem summarizes its main properties. To that end, for each $n \in [N]$, let us introduce the constant $\alpha_n := \sup_{\mathbf{u} \in B(\mathbf{0}, 1) \cap C} \mathbf{w}_n^\top \mathbf{u}$; note that $\alpha_n \in (0, 1]$ since we assume that $\sup_{\mathbf{u} \in B(\mathbf{0}, 1)} \mathbf{w}_n^\top \mathbf{u} = \|\mathbf{w}_n\|_2 = 1$.

Theorem 2.4. (i) When C is a polyhedral ordering cone without redundancy and $\|\mathbf{w}_n\|_2 = 1$ for each $n \in [N]$, it holds $\beta_1 < +\infty$ and $\beta_2 \leq (\min_{n \in [N]} \alpha_n)^{-1} < +\infty$. (ii) Suppose further that $C \supseteq \mathbb{R}_+^D$. Then, $\beta_1 = \beta_2 = 1$.

Note that Theorem 2.4 establishes the finiteness of the ordering complexity β . In the sample complexity analysis, we will assume that β is known. Hence, we are also interested in the calculation of (an upper bound) β . While Theorem 2.4(i) already gives an easy-to-calculate upper bound on β_2 , the calculation of β_1 is more involved and requires maximizing a Lipschitz function over a compact set. Moreover, in Theorem 2.4(ii), one can relax the condition $C \supseteq \mathbb{R}_+^D$ slightly. These technical points are discussed in detail with proofs and additional results in the supplemental document.

Remark 2.5. The algorithm we propose in Section 6 requires β as input, which can be computed offline before starting evaluations. In addition, we would like to point out that the calculation of β for large values of D is not of primary concern in our setting. In both multi-objective and vector optimization, usually a small number $2 \leq D \leq 5$, and very typically only $D = 2$ or $D = 3$ of conflicting objectives are considered (Kováčová and Rudloff, 2021; Ararat et al., 2022). While it is valuable to provide the decision-maker with a whole frontier of Pareto optimal designs for small number of objectives, the Pareto set becomes intractable for larger number of objectives. In the latter case, it would be more of a burden than flexibility for the decision-maker to choose from the Pareto set of designs.

We calculate β_1, β_2 for some standard ordering cones in \mathbb{R}^2 in the next example.

Example 2.6. Let $D = 2$. Given $\mathbf{x} \in \mathbb{R}^2$, let $\alpha(\mathbf{x}) \in [0, 2\pi)$ denote the angle in the polar coordinates of \mathbf{x} . Let $\theta \in (0, \pi/2)$ and define the ordering cone $C_\theta := \{\mathbf{x} \in \mathbb{R}^2 \mid \alpha(\mathbf{x}) \in [\pi/4 - \theta/2, \pi/4 + \theta/2]\}$. Let $\mathbf{x} \notin C_\theta$. Using elementary planar geometry, it can be checked that $\beta_1 = \beta_2 = \csc(\theta)$. The details of the calculation are given in the supplemental document. If we take $\theta \in (\pi/2, \pi)$,

⁵ β only depends on C . It does not depend on the mean rewards.

then $\beta_1 = \beta_2 = 1$ by Theorem 2.4(ii). Note that, up to a rotation of the angle bisector, this is the general form of an ordering cone for $D = 2$; the current choice of the angle bisector is for convenience.

Remark 2.7. Cone C' in (2) can also be represented in the form given in Example 2.6. For instance, $C_{3\pi/4}$ in Figure 1 is such a cone with $\alpha = \frac{-\sin(\frac{\pi}{4} - \frac{\theta}{2})}{\sqrt{2} \sin(\frac{\theta}{2})} \approx 0.29$, where $\theta = 3\pi/4$. Indeed, given any relative importance level $\alpha \in (0, 0.5)$ specified by the practitioner, one can find θ such that $\frac{-\sin(\frac{\pi}{4} - \frac{\theta}{2})}{\sqrt{2} \sin(\frac{\theta}{2})} = \alpha$.

The polyhedral ordering cone C induces two non-total order relations on \mathbb{R}^D as defined next.

Definition 2.8. For every $\boldsymbol{\mu}, \boldsymbol{\mu}' \in \mathbb{R}^D$, we write $\boldsymbol{\mu} \preceq_C \boldsymbol{\mu}'$ if $\boldsymbol{\mu}' \in \boldsymbol{\mu} + C$, and we write $\boldsymbol{\mu} \prec_C \boldsymbol{\mu}'$ if $\boldsymbol{\mu}' \in \boldsymbol{\mu} + \text{int}(C)$.

It can be checked that both \preceq_C and \prec_C are partial order relations on \mathbb{R}^D . If $C = \mathbb{R}_+^D$, then \preceq_C coincides with the usual componentwise order on \mathbb{R}^D , which is the partial order used in multi-objective optimization. To explain the motivation for using an ordering cone that is different from the positive orthant, let us first observe that

$$\boldsymbol{\mu} \preceq_C \boldsymbol{\mu}' \Leftrightarrow \forall n \in [N]: \mathbf{w}_n^\top \boldsymbol{\mu} \leq \mathbf{w}_n^\top \boldsymbol{\mu}' \quad (5)$$

for every $\boldsymbol{\mu}, \boldsymbol{\mu}' \in \mathbb{R}^D$. We denote by C^+ the convex cone generated by the rows of W , that is,

$$C^+ := \left\{ \sum_{n=1}^N \lambda_n \mathbf{w}_n \mid \lambda_1, \dots, \lambda_N \geq 0 \right\}. \quad (6)$$

It can be shown that C^+ is also a polyhedral ordering cone and it coincides with the so-called *dual cone* of C , that is, $C^+ = \{\mathbf{w} \in \mathbb{R}^D \mid \forall \mathbf{x} \in C: \mathbf{w}^\top \mathbf{x} \geq 0\}$. It is well-known that the positive orthant is self-dual, that is, $(\mathbb{R}_+^D)^+ = \mathbb{R}_+^D$. By (6), we may rewrite (5) as

$$\boldsymbol{\mu} \preceq_C \boldsymbol{\mu}' \Leftrightarrow \forall \mathbf{w} \in C^+: \mathbf{w}^\top \boldsymbol{\mu} \leq \mathbf{w}^\top \boldsymbol{\mu}'. \quad (7)$$

Here, each $\mathbf{w} \in C^+$ can be considered as a “weight vector” since the partial order is determined by comparing the weighted sums of the components of $\boldsymbol{\mu}, \boldsymbol{\mu}'$ for all $\mathbf{w} \in C^+$. Moreover, if C' is another ordering cone, then we have $C \supseteq C'$ iff $C^+ \subseteq (C')^+$. In particular, if W has non-negative entries so that $C \supseteq \mathbb{R}_+^D$, then we have $C^+ \subseteq \mathbb{R}_+^D$. Consequently, \preceq_C is weaker than the componentwise order in this case, thanks to (7). In other words, in vector optimization, the decision-maker may have a more relaxed requirement when comparing two vectors, which is encoded by having a smaller set C^+ of weight vectors.

The above orders on \mathbb{R}^D induce further orders on the design space $[K]$.

Definition 2.9. Let $i, j \in [K]$. Design i is said to be *weakly dominated* by design j ($i \preceq_C j$) if $\boldsymbol{\mu}_i \preceq_C \boldsymbol{\mu}_j$. Design i is said to be *dominated* by design j ($i \preceq_{C \setminus \{0\}} j$) if $\boldsymbol{\mu}_i \preceq_{C \setminus \{0\}} \boldsymbol{\mu}_j$. Design i is *strongly dominated* by design j ($i \prec_C j$) if $\boldsymbol{\mu}_i \prec_C \boldsymbol{\mu}_j$.

3 THE VECTOR OPTIMIZATION PROBLEM

Equipped with the definitions in Section 2, the vector optimization problem can be expressed as

$$\text{maximize } \boldsymbol{\mu}_i \text{ with respect to cone } C \text{ over } i \in [K].$$

The solution of this problem is called the Pareto set, as we define next.

Definition 3.1. A design $i \in [K]$ is called *Pareto optimal* if it is not dominated by any other design with respect to the ordering induced by C . The *Pareto set* $P^* := \{i \in [K] \mid \nexists j \in [K]: i \preceq_{C \setminus \{0\}} j\}$ consists of all Pareto optimal designs.

In vector optimization, a Pareto optimal design $i \in P^*$ is sometimes called an efficient solution or a maximizer of the design space $[K]$, and the corresponding objective vector $\boldsymbol{\mu}_i$ is called a C -maximal element of the objective set $\{\boldsymbol{\mu}_j \mid j \in [K]\}$; see Jahn (2011, Definition 3.1), Heyde and Löhne (2011, Definition 2.1).

Remark 3.2. Let us write $P^*(C) := P^*$ to emphasize the dependence on the ordering cone. Then, switching to a larger ordering cone $C' \supseteq C$ makes the dominance relation weaker. Consequently, the Pareto set becomes smaller under the larger cone: $P^*(C') \subseteq P^*(C)$.

4 PARETO SET IDENTIFICATION

We consider a learning problem where the mean vectors $\boldsymbol{\mu}_i$, $i \in [K]$, are not known beforehand. Our goal is to identify P^* from noisy observations of the objective values of chosen designs in as few evaluations as possible. We assume that the total number of evaluations, $T \geq 1$, is finite. Evaluations are done in a sequential manner with $t \in [T]$ representing the round in which the t th evaluation is made, and we denote by I_t the random variable representing the design evaluated in round t . We consider a noisy feedback setting in which the evaluation in round t yields a random reward vector $\mathbf{X}_t = \boldsymbol{\mu}_{I_t} + \mathbf{Y}_{I_t, t}$, where $\mathbf{Y}_{i, t}$ is the random noise vector associated to this evaluation when the evaluated design is $I_t = i$. We assume that the family $(\mathbf{Y}_{i, t})_{i \in [K], t \in [T]}$ consists of independent and centered random vectors, and it is independent of the family $(I_t)_{t \in [T]}$. We also assume that all members of this family are norm-subgaussian with a common parameter $\sigma \geq 0$ as stated in the next definition.

Definition 4.1. (Jin et al., 2019, Definition 3) A centered random vector \mathbf{Y} is called *norm-subgaussian* with parameter $\sigma \geq 0$ if $\mathbb{P}\{\|\mathbf{Y}\|_2 \geq \epsilon\} \leq 2e^{-\frac{\epsilon^2}{2\sigma^2}}$ for every $\epsilon \geq 0$.

Examples of norm-subgaussian random vectors include bounded random vectors and subgaussian random vectors

(up to a scaling of the parameter); see Jin et al. (2019, Lemma 1).

Due to sampling noise, the sample complexity of identifying P^* depends on the hardness of distinguishing the designs in P^* from the designs that are not in P^* . We quantify the hardness by the following gaps.

Given two designs $i, j \in [K]$, we define $m(i, j)$ as the minimum increment in $\boldsymbol{\mu}_i$ in an arbitrary direction of increase that makes design i not strongly dominated by design j . Formally, we have

$$m(i, j) := \inf\{s \geq 0 \mid \exists \mathbf{u} \in B(\mathbf{0}, 1) \cap C: \boldsymbol{\mu}_i + s\mathbf{u} \notin \boldsymbol{\mu}_j - \text{int}(C)\}. \quad (8)$$

The next proposition lists important properties of this gap.

Proposition 4.2. Let $i, j \in [K]$. (i) It holds $m(i, j) < +\infty$. (ii) It holds $m(i, j) = d(\boldsymbol{\mu}_j - \boldsymbol{\mu}_i, (\text{int}(C))^c \cap (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i - C))$. (iii) We have $m(i, j) > 0$ iff $i \prec_C j$. (iv) It holds $m(i, j) = \min_{n \in [N]} (w_n^\top (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i))^+ / \alpha_n$.

Similarly, given $i, j \in [K]$, we define $M(i, j)$ as the minimum increment in $\boldsymbol{\mu}_j$ in an arbitrary direction of increase that makes design i weakly dominated by design j , that is,

$$M(i, j) := \inf\{s \geq 0 \mid \exists \mathbf{u} \in B(\mathbf{0}, 1) \cap C: \boldsymbol{\mu}_j + s\mathbf{u} \in \boldsymbol{\mu}_i + C\}. \quad (9)$$

We state the fundamental properties of this gap next.

Proposition 4.3. Let $i, j \in [K]$. (i) It holds $M(i, j) < +\infty$. (ii) It holds $M(i, j) = d(\boldsymbol{\mu}_j - \boldsymbol{\mu}_i, C \cap (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i + C))$. (iii) We have $M(i, j) = 0$ iff $i \preceq_C j$.

Remark 4.4. From (9), it is clear that using a larger cone would result in smaller values of $M(i, j)$.

As an immediate consequence of Propositions 4.2(ii, iii) and 4.3(ii, iii), we obtain the following corollary.

Corollary 4.5. Let $i, j \in [K]$. Then, exactly one of the following cases holds.

- (i) We have both $i \prec_C j$ and $i \preceq_C j$ iff $\boldsymbol{\mu}_j - \boldsymbol{\mu}_i \in \text{int}(C)$ iff $m(i, j) > 0$ and $M(i, j) = 0$.
- (ii) We have both $i \not\prec_C j$ and $i \preceq_C j$ iff $\boldsymbol{\mu}_j - \boldsymbol{\mu}_i \in \text{bd}(C)$ iff $m(i, j) = M(i, j) = 0$.
- (iii) We have both $i \not\prec_C j$ and $i \not\preceq_C j$ iff $\boldsymbol{\mu}_j - \boldsymbol{\mu}_i \in C^c$ iff $m(i, j) = 0$ and $M(i, j) > 0$.

In particular, we have $m(i, j) = 0$ or $M(i, j) = 0$.

For each $i \in [K]$, let $\Delta_i^* := \max_{j \in P^*} m(i, j)$. By Proposition 4.3(iii), it is clear that $\Delta_i^* = 0$ if $i \in P^*$. We consider an (ϵ, δ) -PAC (probably approximately correct) Pareto set identification setup under which the estimated Pareto set $P \subseteq [K]$ returned by the learner needs to satisfy the condition in the next definition.

Definition 4.6. (Success condition) Let $\epsilon > 0$, $\delta \in (0, 1)$. A set $P \subseteq [K]$ is called an (ϵ, δ) -PAC Pareto set if the

following properties hold with probability at least $1 - \delta$: (i) $\cup_{i \in P} (\mu_i + (B(\mathbf{0}, \epsilon) \cap C) - C) \supseteq \cup_{i \in P^*} (\mu_i - C)$; (ii) for every $i \in P \setminus P^*$, it holds $\Delta_i^* \leq \epsilon$.

Remark 4.7. Note that property (i) in Definition 4.6 is equivalent to (i') $\cup_{i \in P} (\mu_i + (B(\mathbf{0}, \epsilon) \cap C) - C) \supseteq \{\mu_i \mid i \in P^*\}$. The implication (i) \Rightarrow (i') is obvious. To see (i') \Rightarrow (i), suppose that (i') holds and let $v \in \mu_i - C$ for some $i \in P^*$. Then, by (i'), $\mu_i \in \mu_j + (B(\mathbf{0}, \epsilon) \cap C) - C$ for some $j \in P$. Moreover, $C + C = C$ since C is a convex cone. Hence, $v \in \mu_i - C \subseteq \mu_j + (B(\mathbf{0}, \epsilon) \cap C) - C - C = \mu_j + (B(\mathbf{0}, \epsilon) \cap C) - C$, showing (i). A similar covering property is used in Ararat et al. (2022, Definition 3.5) in the context of deterministic convex vector optimization.

Property (i) in Definition 4.6 is an ϵ -covering requirement for the Pareto set P^* ; in view of Remark 4.7, it is equivalent to the following: for every $i \in P^*$, there exist $j \in P$ and $u \in B(\mathbf{0}, 1) \cap C$ such that $\mu_i \preceq_C \mu_j + \epsilon u$. Roughly speaking, although P might not contain all Pareto optimal designs, it is required to include a close-enough design for each Pareto optimal design. Moreover, although some designs in P might be suboptimal, property (ii) in Definition 4.6 bounds the gaps of such designs; hence, it controls the quality of all returned designs.

Remark 4.8. Ignoring the differences in gap definitions, let us compare our success condition with the two alternative success conditions (SC1 and SC2) in Auer et al. (2016) at a structural level. All three conditions require small suboptimality gap: $\Delta_i^* \leq \epsilon$. While SC1 requires P^* to be a subset of the returned set, SC2 and ours impose the weaker ϵ -covering requirement, which is in line with Even-Dar et al. (2006); Zuluaga et al. (2016). Finally, SC2 has a sparsity requirement for the returned set of designs, which we do not have in Definition 4.6. For instance, when $D = 1$, our success condition is satisfied when all returned designs are ϵ -optimal, which might still violate SC2 depending on the configuration of the reward vectors. However, due to the different definitions of the gaps, our success condition is mathematically incomparable with SC1, SC2.

There is an innate connection between gaps defined in this section and the ordering complexity terms defined in Section 2. The following remark summarizes this relation, which plays an important role in characterizing the sample complexity.

Remark 4.9. (Intuitive explanations of β_1, β_2) Let $\Delta_{ij} = \mu_j - \mu_i$. Consider i, j such that $\Delta_{ij} \notin C$, i.e., $\mu_i \not\preceq_C \mu_j$. By definition, $M(i, j) = d(\Delta_{ij}, C \cap (\Delta_{ij} + C))$ (numerator in (3)) is the minimum change in Δ_{ij} in an arbitrary direction of increase that makes i weakly dominated by j , while $d(\Delta_{ij}, C)$ (denominator in (3)) is the minimum change in Δ_{ij} that makes i weakly dominated by j . In scalar and multi-objective problems, we have $M(i, j) = d(\Delta_{ij}, C)$. However, the ratio $M(i, j)/d(\Delta_{ij}, C)$ can be as large as β_1 in vector optimization. This makes learning fundamentally more difficult in vector optimization when

$\beta_1 > 1$. In the scalar or multi-objective case, problems with larger gaps $M(i, j)$ are easier (success conditions can be met with fewer samples). However, in general vector optimization, larger gaps do not imply easier problems, what is important is the ratio $M(i, j)/d(\Delta_{ij}, C)$. A similar argument can be made for β_2 .

5 LOWER BOUNDS

We provide gap-dependent and worst-case sample complexity bounds for any algorithm that satisfies the success condition in Definition 4.6. For each $i \in P^*$, we define $\Delta_i^+ := \min_{j \in P^* \setminus \{i\}} M(i, j)$ and $\tilde{\Delta}_i^\epsilon := \max\{\Delta_i^+, \epsilon\}$. For each $i \notin P^*$, we have $\Delta_i^* = \max_{j \in P^*} m(i, j)$ and define $\tilde{\Delta}_i^\epsilon := \max\{\Delta_i^*, \epsilon\}$. Finally, let \mathcal{A} be the class of all algorithms that return, for any given $\epsilon > 0$ and $\delta \in (0, 1)$, a set $P \subseteq [K]$ satisfying Definition 4.6 for every given distribution of \mathbf{X}_t with $\mathbf{X}_t \in [0, 1]^D$ almost surely for each $t \in [T]$.

Theorem 5.1. (*Gap-dependent lower bound*) *Suppose that $\mu_i \in [\frac{1}{4}, \frac{3}{4}]^D$ for each $i \in [K]$ and $\epsilon \leq 1/8$. Then, there exist such distributions with the property that every algorithm in \mathcal{A} requires $\Omega(\sum_{i \in [K]} \frac{1}{(\tilde{\Delta}_i^\epsilon)^2} \log(\frac{1}{\delta}))$ samples to work correctly.*

Remark 5.2. (a) Theorem 5.1 is analogous to Theorem 17 in Auer et al. (2016). However, this result works under their SC1, which is structurally much stricter than and mathematically incomparable with our success condition in Definition 4.6; see Remark 4.8. Ours is closer to their SC2, for which there is no lower bound analysis in Auer et al. (2016). Hence, by relaxing the structure of their SC2, we are able to provide lower bounds on sample complexity for a success condition with an ϵ -covering requirement, as opposed to the full coverage requirement in their SC1.

(b) Different from their proof in the multi-objective setting, our proof constructs a special direction vector $z^* \in \text{int}(C)$ along which the rewards and their (empirical) means for a fixed design are ordered linearly. Precisely, z^* is chosen in such a way that $z^* \in u + C$ for every $u \in B(\mathbf{0}, 1) \cap C$. Using this vector, we analyze the following four cases: (i) When $i \in P^* \setminus P$, we modify the mean reward of i as $\mu'_i := \mu_i + 2\tilde{\Delta}_i^\epsilon z^*$ and show that i is not covered (in the sense of Definition 4.6) by any $j \in [K]$ that covers i in the original case. Hence, i must be returned by the algorithm in the modified case. (ii) When $i \in P^* \cap P$, we have $\Delta_i^+ = M(i, j)$ for some $j \in P^* \setminus \{i\}$. We take $\mu'_i := \mu_i - 3k^* \tilde{\Delta}_i^\epsilon z^*$ with $k^* := \max_{n \in [N]} \frac{\alpha_n}{w_n^1 z^*}$ and show that $m(i, j) \geq 2\epsilon$ in the modified case. Hence, i cannot be returned by the algorithm. (iii) The case $i \in P \setminus P^*$ is similar to (ii). (iv) When $i \notin P^* \cup P$, we take $\mu'_i := \mu_i + (1 + 2k^*) \tilde{\Delta}_i^\epsilon z^*$ and show that i must be returned. Since we are able to change the returned set in each possibility, the rest of the proof follows as in Auer et al. (2016). The main technical novelties in the proof are

twofold: 1) The direction vector \mathbf{z}^* is constructed by using the structure of the ordering cone. 2) Compared to Auer et al. (2016), case (i) is new to our setting as $P^* \subseteq P$ is required there, which is relaxed by our Definition 4.6.

Theorem 5.3. (Worst-case lower bound) *Suppose that there are $K \geq 4$ designs. Given $\epsilon > 0$, there exist mean reward vectors and norm-subgaussian noise distributions under which any algorithm in \mathcal{A} requires $\Omega((K\beta_2^2/\epsilon^2) \log(1/\delta))$ samples to work correctly.*

The proof of Theorem 5.3 requires utilizing the geometry of the ordering cone to construct a special worst-case problem instance that is hard to distinguish from other similar problem instances. In this special instance, while mean reward vectors of most designs are identical and Pareto optimal, and there is another Pareto optimal design that is in a specially perturbed direction. We show that, when the noise is in a special direction determined by β_2 , a change in its distribution that is proportional to ϵ/β_2 will result in a change in the returned set of designs. This lower bound manifests a distinct feature of vector optimization which is not present in the scalar or multi-objective setting under noisy feedback, i.e., geometry of the ordering cone characterizes the hardness of the Pareto set identification problem. Note that we indeed have $\beta = \beta_2$ for all ordering cones that are at least as large as \mathbb{R}_+^D (Theorem 2.4(ii)) and all ordering cones in the biobjective case (Example 2.6).

6 NAÏVE ELIMINATION AND ITS SAMPLE COMPLEXITY

In this section, we introduce the naïve elimination algorithm that is used for (ϵ, δ) -PAC Pareto set identification. The algorithm operates in the same fashion as the naïve elimination algorithm used for (ϵ, δ) -PAC best arm identification (Even-Dar et al., 2006). For $D = 1$, it is known that naïve elimination is a simple yet provable and practical algorithm that is known to work reasonably well when K is not very large (see the discussion in Hassidim et al. (2020)).⁶

Naïve elimination takes as inputs (ϵ, δ) and the polyhedral ordering cone C defined by matrix W . It evaluates each design L times to form empirical means $\hat{\boldsymbol{\mu}}_i$, $i \in [K]$, where $L \geq 1$ is a positive integer that is used as the exploration parameter and it is set depending on ϵ, δ, K and the ordering complexity β of C .⁷ Hence, $T = LK$ evaluations are

⁶Our aim is to show that even such a simple algorithm can be successful when its sampling budget is tuned according to ordering complexity. We think that many of the successive elimination and adaptive sampling techniques developed for the scalar case can be adapted to devise sample-efficient algorithms for vector optimization. We leave this interesting direction as future work.

⁷The computation of β , which is offline and done once at the beginning, is trivially simple when $C \supseteq \mathbb{R}_+^D$ (Theorem 2.4). For general cones, we show how β can be computed in the supplement

Algorithm 1 Naïve Elimination

Inputs: $\epsilon, \delta, [K], C, \beta, L$
Initialize: $\hat{\boldsymbol{\mu}}_i = N_i = 0, i \in [K]$
for $t = 1, 2, \dots, LK$ **do**
 Evaluate $I_t = (t - 1 \bmod K) + 1$
 Observe reward \mathbf{X}_t
 $N_{I_t} = N_{I_t} + 1; \hat{\boldsymbol{\mu}}_{I_t} = ((N_{I_t} - 1)\hat{\boldsymbol{\mu}}_{I_t} + \mathbf{X}_t)/N_{I_t}$
end for
Return: $P := \{i \in [K] \mid \nexists j \in [K]: i \hat{\succeq}_{C \setminus \{0\}} j\}$

made in total. Then, the algorithm computes and returns a random set P by using Definition 3.1. This is done by checking for each design i whether there exists another design j with $i \hat{\succeq}_{C \setminus \{0\}} j$, where $\hat{\succeq}_{C \setminus \{0\}}$ is the random partial order that is defined by using $\hat{\boldsymbol{\mu}}_i$ in place of $\boldsymbol{\mu}_i$, $i \in [K]$, in the standard expression of $\preceq_{C \setminus \{0\}}$; see Definitions 2.8, 2.9 (the random relations $\hat{\succ}_C$ and $\hat{\succeq}_C$ are defined similarly). Hence, we set $P := \{i \in [K] \mid \nexists j \in [K]: i \hat{\succeq}_{C \setminus \{0\}} j\}$. One possible implementation of naïve elimination is given in Algorithm 1.

For a given (ϵ, δ) , we define the per-design sampling budget as $g(\epsilon, \delta) := \left\lceil \frac{4\beta^2 \epsilon^2 \sigma^2}{\epsilon^2} \log\left(\frac{4D}{\delta}\right) \right\rceil$. The following theorem characterizes the sample complexity of naïve elimination for (ϵ, δ) -PAC Pareto set identification.

Theorem 6.1. *When naïve elimination is run with $L = g(\epsilon, 2\delta/(K(K-1)))$, the returned Pareto set P is an (ϵ, δ) -PAC Pareto set.*

Theorem 6.1 gives an upper bound on the sample complexity that scales as $\tilde{O}(K\beta^2/\epsilon^2)$ which nearly matches the lower bound in Theorem 5.3.⁸ It can also be shown that, when $D = 1$, the sample complexity of naïve elimination matches the one given in Even-Dar et al. (2006).

Let us provide a sketch of the proof of Theorem 6.1. Given $i, j \in [K]$, let us introduce empirical estimates for $m(i, j)$ and $M(i, j)$, defined as $\hat{m}(i, j) := d(\hat{\boldsymbol{\mu}}_j - \hat{\boldsymbol{\mu}}_i, (\text{int}(C))^c \cap (\hat{\boldsymbol{\mu}}_j - \hat{\boldsymbol{\mu}}_i - C))$ and $\hat{M}(i, j) := d(\hat{\boldsymbol{\mu}}_j - \hat{\boldsymbol{\mu}}_i, C \cap (\hat{\boldsymbol{\mu}}_j - \hat{\boldsymbol{\mu}}_i + C))$, respectively. As a first step, we show that $M(i, j) > \epsilon$ implies $\hat{M}(i, j) > 0$ whenever $i \in P^*$, $j \in [K] \setminus \{i\}$, as well as that $m(i, j) > \epsilon$ implies $\hat{m}(i, j) > 0$ whenever $i \notin P^*$, $j \in P^*$. In the second step, we show that these implications are sufficient to conclude that P is an (ϵ, δ) -PAC Pareto set. Then, for $i, j \in [K]$ with $i \neq j$, we define $\Delta_{ij} = \boldsymbol{\mu}_j - \boldsymbol{\mu}_i$, $\hat{\Delta}_{ij} = \hat{\boldsymbol{\mu}}_j - \hat{\boldsymbol{\mu}}_i$, and show that the following conditions in terms of the deviation of $\hat{\Delta}_{ij}$ from Δ_{ij} are sufficient for (ϵ, δ) -PAC Pareto set identification:

Condition a. *For every $i \in P^*$ and $j \in [K] \setminus \{i\}$,*

tal document; however, we do not have a computationally efficient procedure for it as the problem is not convex. Nevertheless, as we mention in Remark 2.5, problems in vector optimization generally have small objective dimension (not to be confused with design space dimension).

⁸Indeed, the sample complexity is $O(K\beta^2 \log(KD/\delta)/\epsilon^2)$.

Table 1: Gap Statistics of Designs $i \in [K] \setminus P_\theta^*$.

	$\Delta_i^* (C_{\pi/4})$	$\Delta_i^* (C_{\pi/2})$	$\Delta_i^* (C_{3\pi/4})$
count	153	180	196
mean	0.666	0.906	1.343
std	0.520	0.711	0.869
min	0.001	0.004	0.018
max	2.576	3.544	4.726

Table 2: Success Rate (%) of Naïve Elimination.

L	ϵ	$C_{\pi/4}$	$C_{\pi/2}$	$C_{3\pi/4}$
10^3	10^{-2}	0	0	29
	10^{-1}	78	99	100
10^4	10^{-2}	22	24	85
	10^{-1}	100	100	100
10^5	10^{-2}	100	99	100
	10^{-1}	100	100	100

$d(\Delta_{ij}, C \cap (\Delta_{ij} + C)) > \epsilon$ implies that $\|\hat{\Delta}_{ij} - \Delta_{ij}\|_2 < d(\Delta_{ij}, C)$.

Condition b. For every $i \notin P^*$ and $j \in P^*$, $d(\Delta_{ij}, (\text{int}(C))^c \cap (\Delta_{ij} - C)) > \epsilon$ implies that $\|\hat{\Delta}_{ij} - \Delta_{ij}\|_2 < d(\Delta_{ij}, (\text{int}(C))^c)$.

Let us also define $\theta_{ij} := \frac{d(\Delta_{ij}, C)}{d(\Delta_{ij}, C \cap (\Delta_{ij} + C))}$ if $\Delta_{ij} \notin C$, $\theta_{ij} := \frac{d(\Delta_{ij}, (\text{int}(C))^c)}{d(\Delta_{ij}, (\text{int}(C))^c \cap (\Delta_{ij} - C))}$ if $\Delta_{ij} \in \text{int}(C)$, and $\theta_{ij} := 1$ if $\Delta_{ij} \in \text{bd}(C)$. In the third step, using properties of norm-subgaussian noise vectors, we show the existence of a constant $c > 0$ (free of all problem parameters) such that, for all $i \neq j$, $\|\hat{\Delta}_{ij} - \Delta_{ij}\|_2 \leq \epsilon \theta_{ij}$ with probability at least $1 - \delta$ provided that naïve elimination is run with $L = g(\epsilon, 2\delta/(K(K-1)))$. Finally, we complete the proof by checking that Conditions a and b hold with probability at least $1 - \delta$.

7 NUMERICAL RESULTS

We use SNW dataset from Zuluaga et al. (2016). It consists of 206 different hardware implementations of a sorting network. The objectives are the area and throughput of the network when synthesized on an FPGA ($D = 2$). Since we consider maximization problems, we use the negative of area as objective value. The mean rewards of designs are taken as the objective values in the dataset. The reward vector of a design is formed by adding independent zero mean Gaussian noise with variance $\sigma^2 = 1$ to the mean value of each objective of the design.

We consider the polyhedral cone C_θ in Example 2.6, which is parametrized by an angle $\theta \in (0, \pi)$. We use $C_{\pi/4}$, $C_{\pi/2}$ and $C_{3\pi/4}$ in our simulations (Figure 2). P_θ^* represents the true Pareto optimal set under the ordering induced by C_θ . The Pareto sets for $C_{\pi/4}$, $C_{\pi/2}$ and $C_{3\pi/4}$ are shown in Figure 1. Some statistics of the gaps Δ_i^* of designs $i \in [K] \setminus P_\theta^*$ are given in Table 1.

The simulation code is available in the supplemental material. As there is no work that considers vector optimization with stochastic bandit feedback, we do not compare with any other method. Instead, we illustrate how the performance of naïve elimination varies as a function of the number of samples from each design and shape of the ordering cone. We set $\delta = 0.01$ in all simulations. The reported results correspond to the average of 100 independent runs.

As typically observed in the best arm identification literature, the theoretical value of L is very large. For instance, for $\theta = \pi/2$, if $\epsilon = 0.1$, then $L \approx 38.8 \times 10^3$; if $\epsilon = 0.01$, then $L \approx 38.8 \times 10^5$. Therefore, instead of using theoretical values, we evaluate the results for different L and ϵ values. The results are provided in Table 2. For fixed L and ϵ , the success rate increases as θ increases. For fixed θ and L , the success rate increases as ϵ increases because both success conditions become easier to satisfy. For a fixed θ and ϵ , the success rate increases as L increases since more samples mean less noise in the estimates. For some (ϵ, θ) pairs, all runs are successful even when L is much smaller than its theoretical value.

8 CONCLUSIONS

We propose vector optimization problems with stochastic bandit feedback. We identify fundamental cone-dependent gaps that characterize the learning difficulty of the Pareto set. We derive sample complexity bounds for (ϵ, δ) -PAC Pareto set identification. Our introduction of noisy bandit feedback to vector optimization brings forth many interesting future research directions. In particular, the design of sample-efficient adaptive algorithms remains as an open problem that we seek to address in the future. For instance, when learning the Pareto set in large design spaces, one can incorporate powerful surrogate models such as Gaussian processes. It should also be possible to extend entropy search (Hernández-Lobato et al., 2016) and hypervolume-based methods (Shah and Ghahramani, 2016) proposed for multi-objective Pareto set identification problems to vector optimization.

Acknowledgements

This work is supported by the Scientific and Technological Research Council of Turkey (TÜBİTAK) under Grant 121E159.

References

- Ç. Ararat, F. Ulus, and M. Umer. A norm minimization-based convex vector optimization algorithm. *Journal of Optimization Theory and Applications*, 192:681–712, 2022.
- J.-Y. Audibert, S. Bubeck, and R. Munos. Best arm identification in multi-armed bandits. In *Proc. Conference on Learning Theory*, pages 41–53, 2010.
- P. Auer, C.-K. Chiang, R. Ortner, and M. Drugan. Pareto front identification from stochastic bandit feedback. In *Proc. 19th International Conference on Artificial Intelligence and Statistics*, pages 939–947, 2016.
- S. Bubeck, R. Munos, and G. Stoltz. Pure exploration in multi-armed bandits problems. In *Proc. International Conference on Algorithmic Learning Theory*, pages 23–37, 2009.
- R. Degenne, W. M. Koolen, and P. Ménard. Non-asymptotic pure exploration by solving games. *Advances in Neural Information Processing Systems*, 32, 2019.
- R. Degenne, P. Ménard, X. Shang, and M. Valko. Gamification of pure exploration for linear bandits. In *Proc. International Conference on Machine Learning*, pages 2432–2442, 2020.
- M. M. Drugan and A. Nowe. Designing multi-objective multi-armed bandits algorithms: A study. In *The 2013 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8, 2013.
- E. Even-Dar, S. Mannor, Y. Mansour, and S. Mahadevan. Action elimination and stopping conditions for the multi-armed bandit and reinforcement learning problems. *Journal of Machine Learning Research*, 7(6), 2006.
- A. Garivier and E. Kaufmann. Optimal best arm identification with fixed confidence. In *Proc. Conference on Learning Theory*, pages 998–1027, 2016.
- A. Garivier and E. Kaufmann. Nonasymptotic sequential tests for overlapping hypotheses applied to near-optimal arm identification in bandit models. *Sequential Analysis*, 40(1):61–96, 2021.
- A. Hassidim, R. Kupfer, and Y. Singer. An optimal elimination algorithm for learning a best arm. *Advances in Neural Information Processing Systems*, 33:10788–10798, 2020.
- C. F. Hayes, R. Rădulescu, E. Bargiacchi, J. Källström, M. Macfarlane, M. Reymond, T. Verstraeten, L. M. Zintgraf, R. Dazeley, F. Heintz, et al. A practical guide to multi-objective reinforcement learning and planning. *Autonomous Agents and Multi-Agent Systems*, 36(1):1–59, 2022.
- D. Hernández-Lobato, J. Hernandez-Lobato, A. Shah, and R. Adams. Predictive entropy search for multi-objective Bayesian optimization. In *Proc. International Conference on Machine Learning*, pages 1492–1501, 2016.
- F. Heyde and A. Löhne. Solution concepts in vector optimization: a fresh look at an old story. *Optimization*, 60: 1421–1440, 2011.
- J. Jahn. *Vector Optimization*. Springer, 2nd edition, 2011.
- M. K. Jayatunga, W. Xie, L. Ruder, U. Schulze, and C. Meier. AI in small-molecule drug discovery: A coming wave. *Nat. Rev. Drug Discov*, 21:175–176, 2022.
- C. Jin, P. Netrapalli, R. Ge, S. M. Kakade, and M. I. Jordan. A short note on concentration inequalities for random vectors with subgaussian norm. *arXiv preprint arXiv:1902.03736*, 2019.
- Y. M. Kabanov. Hedging and liquidation under transaction costs in currency markets. *Finance and Stochastics*, 3: 237–248, 1999.
- S. Kalyanakrishnan, A. Tewari, P. Auer, and P. Stone. PAC subset selection in stochastic multi-armed bandits. In *Proc. International Conference on Machine Learning*, volume 12, pages 655–662, 2012.
- Z. Karnin, T. Koren, and O. Somekh. Almost optimal exploration in multi-armed bandits. In *Proc. International Conference on Machine Learning*, pages 1238–1246, 2013.
- J. Katz-Samuels and C. Scott. Feasible arm identification. In *Proc. International Conference on Machine Learning*, pages 2535–2543, 2018.
- E. Kaufmann and S. Kalyanakrishnan. Information complexity in bandit subset selection. In *Proc. Conference on Learning Theory*, pages 228–251, 2013.
- G. Kováčová and B. Rudloff. Time consistency of the mean-risk problem. *Operations Research*, 69(4):1100–1117, 2021.
- M. Laskey, J. Mahler, Z. McCarthy, F. T. Pokorny, S. Patil, J. Van Den Berg, D. Kragic, P. Abbeel, and K. Goldberg. Multi-armed bandit models for 2d grasp planning with uncertainty. In *Proc. 2015 IEEE International Conference on Automation Science and Engineering (CASE)*, pages 572–579, 2015.
- D. J. Lizotte and E. B. Laber. Multi-objective Markov decision processes for data-driven decision support. *The Journal of Machine Learning Research*, 17(1):7378–7405, 2016.
- A. Löhne. *Vector Optimization with Infimum and Supremum*. Springer, 2011.
- C. Qin, D. Klabjan, and D. Russo. Improving the expected improvement algorithm. *Advances in Neural Information Processing Systems*, 30, 2017.
- D. Russo. Simple Bayesian algorithms for best-arm identification. *Operations Research*, 68(6):1625–1647, 2020.

- A. Shah and Z. Ghahramani. Pareto frontier learning with expensive correlated objectives. In *Proc. International Conference on Machine Learning*, pages 1919–1927, 2016.
- X. Shang, R. Heide, P. Menard, E. Kaufmann, and M. Valko. Fixed-confidence guarantees for Bayesian best-arm identification. In *Proc. International Conference on Artificial Intelligence and Statistics*, pages 1823–1832, 2020.
- K. Teichert, G. Currie, K.-H. Küfer, E. Miguel-Chumacero, P. Süß, M. Walczak, and S. Currie. Targeted multi-criteria optimisation in IMRT planning supplemented by knowledge based model creation. *Operations Research for Health Care*, 23:100185, 2019.
- E. Turgay, D. Oner, and C. Tekin. Multi-objective contextual bandit problem with similarity information. In *Proc. International Conference on Artificial Intelligence and Statistics*, pages 1673–1681, 2018.
- K. Van Moffaert and A. Nowé. Multi-objective reinforcement learning using sets of Pareto dominating policies. *The Journal of Machine Learning Research*, 15(1): 3483–3512, 2014.
- Y. Zhou, X. Chen, and J. Li. Optimal PAC multiple arm identification with applications to crowdsourcing. In *Proc. International Conference on Machine Learning*, pages 217–225, 2014.
- M. Zuluaga, A. Krause, and M. Püschel. ϵ -pal: an active learning approach to the multi-objective optimization problem. *The Journal of Machine Learning Research*, 17(1):3619–3650, 2016.

A TABLE OF NOTATIONS

Notation	Description
D	Number of objectives
K	Number of designs
$[K] = \{1, \dots, K\}$	Design space
$\boldsymbol{\mu}_i \in \mathbb{R}^D$	Mean reward vector of design $i \in [K]$
$C \subseteq \mathbb{R}^D$	Ordering cone
$\text{int}(C)$	Interior of C
$C^+ = \{\boldsymbol{w} \in \mathbb{R}^D \mid \forall \boldsymbol{x} \in C : \boldsymbol{w}^\top \boldsymbol{x} \geq 0\}$	Dual cone of C
$W = [\boldsymbol{w}_1, \dots, \boldsymbol{w}_N]^\top$	$N \times D$ parameter matrix of the polyhedral ordering cone C
\mathbb{R}_+^D	Positive orthant in \mathbb{R}^D
$B(\boldsymbol{v}, r)$	Closed ball in \mathbb{R}^D with center $\boldsymbol{v} \in \mathbb{R}^D$ and radius $r \geq 0$
\preceq_C, \prec_C	Partial order relations on \mathbb{R}^D and on $[K]$ given in Definitions 2.8 and 2.9
$d(\boldsymbol{v}, r) = \inf_{\boldsymbol{x} \in A} \ \boldsymbol{v} - \boldsymbol{x}\ _2$	Distance of vector $\boldsymbol{v} \in \mathbb{R}^D$ to set $A \subseteq \mathbb{R}^D$
$\beta = \max\{\beta_1, \beta_2\}$	Ordering complexity of cone C (β_1, β_2 are defined in (3), (4), respectively.)
$\alpha_n = \sup_{\boldsymbol{u} \in B(\mathbf{0}, 1) \cap C} \boldsymbol{w}_n^\top \boldsymbol{u}$	
P^*	The (true) Pareto set according to cone C
$m(i, j)$	Minimum increment in $\boldsymbol{\mu}_i$ in an arbitrary direction of increase that makes design $i \in [K]$ not strongly dominated by design $j \in [K]$
$M(i, j)$	Minimum increment in $\boldsymbol{\mu}_j$ in an arbitrary direction of increase that makes design $i \in [K]$ weakly dominated by design $j \in [K]$
$\Delta_i^* = \max_{j \in P^*} m(i, j)$	
σ	Subgaussian parameter of observation noise
I_t	Design evaluated in round $t \in [T]$ in Algorithm 1
$\hat{\boldsymbol{\mu}}_i$	Empirical mean of the random reward of design $i \in [K]$ in Algorithm 1
P	Set of designs returned by Algorithm 1

B SUPPLEMENTAL PROOFS

B.1 Detailed version of Example 2.6

We take $D = 2$. Given $\boldsymbol{x} \in \mathbb{R}^2$, let $\alpha(\boldsymbol{x}) \in [0, 2\pi)$ denote the angle in the polar coordinates of \boldsymbol{x} . Let $\theta \in (0, \pi/2]$ and define the ordering cone $C_\theta := \{\boldsymbol{x} \in \mathbb{R}^2 \mid \alpha(\boldsymbol{x}) \in [\pi/4 - \theta/2, \pi/4 + \theta/2]\}$. Let $\boldsymbol{x} \notin C_\theta$. Using elementary planar geometry, it can be checked that $C_\theta \cap (\boldsymbol{x} + C_\theta) = C_\theta$ so that $d(\boldsymbol{x}, C_\theta) = d(\boldsymbol{x}, C_\theta \cap (\boldsymbol{x} + C_\theta))$ if $\alpha(\boldsymbol{x}) \in [5\pi/4 - \theta/2, 5\pi/4 + \theta/2]$. If $\alpha(\boldsymbol{x}) \in [0, \pi/4 - \theta/2) \cup (\pi/4 + \theta/2, 3\pi/4 + \theta/2] \cup [7\pi/4 - \theta/2, 2\pi)$, then $d(\boldsymbol{x}, C_\theta \cap (\boldsymbol{x} + C_\theta)) = d(\boldsymbol{x}, C_\theta) \csc(\theta)$. If $\alpha(\boldsymbol{x}) \in (3\pi/4 + \theta/2, 5\pi/4 - \theta/2)$, then $d(\boldsymbol{x}, C_\theta \cap (\boldsymbol{x} + C_\theta)) = d(\boldsymbol{x}, C_\theta) \csc(\theta) \sin(5\pi/4 + \theta/2 - \alpha(\boldsymbol{x})) \leq d(\boldsymbol{x}, C_\theta) \csc(\theta)$. Similarly, if $\alpha(\boldsymbol{x}) \in (5\pi/4 + \theta/2, 7\pi/4 - \theta/2)$, then $d(\boldsymbol{x}, C_\theta \cap (\boldsymbol{x} + C_\theta)) = d(\boldsymbol{x}, C_\theta) \csc(\theta) \sin(\alpha(\boldsymbol{x}) - 5\pi/4 + \theta/2) \leq d(\boldsymbol{x}, C_\theta) \csc(\theta)$. Hence, $\beta_1 = \csc(\theta)$. By a similar calculation, it can be checked that $\beta_2 = \csc(\theta)$ as well. If we take $\theta \in (\pi/2, \pi)$, then $\beta_1 = \beta_2 = 1$ by Theorem 2.4(ii).

B.2 Proof of Theorem 2.4 and additional results for Section 2

B.2.1 Computation of β_1 and β_2 when the ordering cone is large

Recall the definitions of β_1, β_2 in (4), which depend on the geometry of the ordering cone C . We start by investigating the special case where the ordering cone is at least as large as the positive orthant (or a rotation of it).

Proposition B.1. *Suppose that $\boldsymbol{w}_n^\top \boldsymbol{w}_k \geq 0$ for each $n, k \in [N]$. Then, $d(\boldsymbol{x}, C \cap (\boldsymbol{x} + C)) = d(\boldsymbol{x}, C)$ for every $\boldsymbol{x} \notin C$. In particular, $\beta_1 = 1$.*

Proof. Let us fix $\boldsymbol{x} \notin C$. Note that $C = \{\boldsymbol{y} \in \mathbb{R}^D \mid W\boldsymbol{y} \geq 0\}$ and $\boldsymbol{x} + C = \{\boldsymbol{y} \in \mathbb{R}^D \mid W\boldsymbol{y} \geq W\boldsymbol{x}\}$. Hence, $C \cap (\boldsymbol{x} + C) = \{\boldsymbol{y} \in \mathbb{R}^D \mid W\boldsymbol{y} \geq (W\boldsymbol{x})^+\}$. Since $\boldsymbol{x} \notin C$, there exists $n \in [N]$ such that $\boldsymbol{w}_n^\top \boldsymbol{x} < 0$. Let $I(\boldsymbol{x}) :=$

$\{n \in [N] \mid \mathbf{w}_n^\top \mathbf{x} \leq 0\} \neq \emptyset$. We have $(\mathbf{w}_n^\top \mathbf{x})^+ = 0$ for each $n \in I(\mathbf{x})$ and $(\mathbf{w}_n^\top \mathbf{x})^+ = \mathbf{w}_n^\top \mathbf{x} > 0$ for each $n \in I(\mathbf{x})^c := [N] \setminus I(\mathbf{x})$. Hence, $d(\mathbf{x}, C \cap (\mathbf{x} + C))^2$ can be written as the optimal value of a quadratic optimization problem as follows:

$$d(\mathbf{x}, C \cap (\mathbf{x} + C))^2 = \inf\{\|\mathbf{y} - \mathbf{x}\|_2^2 \mid \forall n \in I(\mathbf{x}): \mathbf{w}_n^\top \mathbf{y} \geq 0, \forall n \in I(\mathbf{x})^c: \mathbf{w}_n^\top \mathbf{y} \geq \mathbf{w}_n^\top \mathbf{x}\}. \quad (10)$$

As this is a convex optimization problem with affine constraints and finite optimal value, the standard Karush-Kuhn-Tucker conditions are necessary and sufficient for optimality. Hence, a vector $\mathbf{y} \in \mathbb{R}^D$ that satisfies the constraints of (10) is optimal if and only if there exists a Lagrange multiplier vector $\boldsymbol{\lambda} \in \mathbb{R}_+^N$ such that the following conditions are satisfied:

$$2(\mathbf{y} - \mathbf{x}) - \sum_{n \in [N]} \lambda_n \mathbf{w}_n = 0, \quad (11)$$

$$\forall n \in I(\mathbf{x}): \lambda_n \mathbf{w}_n^\top \mathbf{y} = 0, \quad (12)$$

$$\forall n \in I(\mathbf{x})^c: \lambda_n (\mathbf{w}_n^\top \mathbf{y} - \mathbf{w}_n^\top \mathbf{x}) = 0. \quad (13)$$

Here, condition (11) is the first order condition for the Lagrangian of (10) with respect to the primal variable \mathbf{y} ; conditions (12) and (13) are the complementary slackness conditions.

Let $\mathbf{y} \in \mathbb{R}^D$ be an optimal solution of (10) with an associated Lagrange multiplier vector $\boldsymbol{\lambda} \in \mathbb{R}_+^N$. We claim that $\lambda_n = 0$ for every $n \in I(\mathbf{x})^c$. To get a contradiction, suppose that $\lambda_{\bar{n}} > 0$ for some $\bar{n} \in I(\mathbf{x})^c$. By (13), we have $\mathbf{w}_{\bar{n}}^\top (\mathbf{y} - \mathbf{x}) = \mathbf{w}_{\bar{n}}^\top \mathbf{y} - \mathbf{w}_{\bar{n}}^\top \mathbf{x} = 0$. On the other hand, $\mathbf{y} - \mathbf{x} = \frac{1}{2} \sum_{n \in [N]} \lambda_n \mathbf{w}_n$ by (11). Combining these, we get

$$\sum_{n \in [N]} \lambda_n \mathbf{w}_n^\top \mathbf{w}_{\bar{n}} = 0. \quad (14)$$

Note that $\mathbf{w}_n \in C^+$ for each $n \in [N]$. By the assumption on $\mathbf{w}_1, \dots, \mathbf{w}_N$, (14) implies that $\lambda_n \mathbf{w}_n^\top \mathbf{w}_{\bar{n}} = 0$ for each $n \in [N]$. In particular, taking $n = \bar{n}$ gives $\lambda_{\bar{n}} \|\mathbf{w}_{\bar{n}}\|_2^2 = \lambda_{\bar{n}} = 0$, which is a contradiction. Therefore, the claim holds.

Thanks to the above claim, the pair $(\mathbf{y}, \boldsymbol{\lambda})$ satisfies the system

$$2(\mathbf{y} - \mathbf{x}) - \sum_{n \in [N]} \lambda_n \mathbf{w}_n = 0, \quad (15)$$

$$\forall n \in [N]: \lambda_n \mathbf{w}_n^\top \mathbf{y} = 0. \quad (16)$$

Moreover, $\mathbf{w}_n^\top \mathbf{y} \geq 0$ for each $n \in [N]$ by the feasibility of \mathbf{y} for (10). Hence, Karush-Kuhn-Tucker conditions are established for the quadratic optimization problem

$$d(\mathbf{x}, C)^2 = \inf\{\|\mathbf{y} - \mathbf{x}\|_2^2 \mid \forall n \in [N]: \mathbf{w}_n^\top \mathbf{y} \geq 0\} \quad (17)$$

and we conclude that \mathbf{y} is optimal for (17). Therefore, $d(\mathbf{x}, C) = d(\mathbf{x}, C \cap (\mathbf{x} + C)) = \|\mathbf{y} - \mathbf{x}\|_2$. \square

Proposition B.2. *Suppose that $\mathbf{w}_n^\top \mathbf{w}_k \geq 0$ for each $n, k \in [N]$. Then, $d(\mathbf{x}, (\text{int}(C))^c \cap (\mathbf{x} - C)) = d(\mathbf{x}, (\text{int}(C))^c)$ for every $\mathbf{x} \in \text{int}(C)$. In particular, $\beta_2 = 1$.*

Proof. Let us fix $\mathbf{x} \in \text{int}(C)$. We have $\mathbf{x} - C = \{\mathbf{y} \in \mathbb{R}^D \mid \forall n \in [N]: \mathbf{w}_n^\top \mathbf{x} \geq \mathbf{w}_n^\top \mathbf{y}\}$. Moreover, we have $(\text{int}(C))^c = \{\mathbf{y} \in \mathbb{R}^D \mid \exists k \in [N]: \mathbf{w}_k^\top \mathbf{y} \leq 0\}$. Hence,

$$(\text{int}(C))^c \cap (\mathbf{x} - C) = \bigcup_{k \in [N]} \{\mathbf{y} \in \mathbb{R}^D \mid \forall n \in [N]: \mathbf{w}_n^\top \mathbf{x} \geq \mathbf{w}_n^\top \mathbf{y}, \mathbf{w}_k^\top \mathbf{y} \leq 0\},$$

which is a finite union of convex polyhedra. In particular, we may write

$$d(\mathbf{x}, (\text{int}(C))^c \cap (\mathbf{x} - C))^2 = \min_{k \in [N]} f_k(\mathbf{x}), \quad (18)$$

$$\text{where } f_k(\mathbf{x}) := \inf\{\|\mathbf{y} - \mathbf{x}\|_2^2 \mid \forall n \in [N]: \mathbf{w}_n^\top \mathbf{x} \geq \mathbf{w}_n^\top \mathbf{y}, \mathbf{w}_k^\top \mathbf{y} \leq 0\}.$$

Similarly, we have $\text{int}(C) = \bigcup_{k \in [N]} \{\mathbf{y} \in \mathbb{R}^D \mid \mathbf{w}_k^\top \mathbf{y} \leq 0\}$ so that

$$d(\mathbf{x}, (\text{int}(C))^c)^2 = \min_{k \in [N]} g_k(\mathbf{x}), \text{ where } g_k(\mathbf{x}) := \inf\{\|\mathbf{y} - \mathbf{x}\|_2^2 \mid \mathbf{w}_k^\top \mathbf{y} \leq 0\}. \quad (19)$$

Let us fix $k \in [N]$. Note that the optimization problems that define $f_k(\mathbf{x}), g_k(\mathbf{x})$ are quadratic and Karush-Kuhn-Tucker conditions characterize optimality. Let $\mathbf{y} \in \mathbb{R}^D$ be optimal for the calculation of $f_k(\mathbf{x})$. Hence, there exists $\boldsymbol{\lambda} \in \mathbb{R}_+^{N+1}$ such that the following conditions are satisfied:

$$2(\mathbf{y} - \mathbf{x}) + \sum_{n \in [N]} \lambda_n \mathbf{w}_n + \lambda_{N+1} \mathbf{w}_k = 0, \quad (20)$$

$$\forall n \in [N]: \lambda_n \mathbf{w}_n^\top (\mathbf{y} - \mathbf{x}) = 0, \quad (21)$$

$$\lambda_{N+1} \mathbf{w}_k^\top \mathbf{y} = 0. \quad (22)$$

We claim that $\lambda_n = 0$ for each $n \in [N]$. To see this, suppose that $\lambda_{\bar{n}} > 0$ for some $\bar{n} \in [N]$. Then, (21) implies that $\mathbf{w}_{\bar{n}}^\top (\mathbf{y} - \mathbf{x}) = 0$. Combining this with (20), we get

$$\mathbf{w}_{\bar{n}}^\top (\mathbf{x} - \mathbf{y}) = \frac{1}{2} \left(\sum_{n \in [N] \setminus \{k\}} \lambda_n \mathbf{w}_n^\top \mathbf{w}_n + (\lambda_{N+1} + \lambda_k) \mathbf{w}_{\bar{n}}^\top \mathbf{w}_k \right) = 0.$$

By the assumption on $\mathbf{w}_1, \dots, \mathbf{w}_N$, we have $\mathbf{w}_n^\top \mathbf{w}_n \geq 0$ for every $n \in [N]$. Hence, all terms in the above sum are nonnegative. It follows that these terms are indeed zero, that is,

$$\forall n \in [N] \setminus \{k\}: \lambda_n \mathbf{w}_n^\top \mathbf{w}_n = 0, \quad (\lambda_{N+1} + \lambda_k) \mathbf{w}_{\bar{n}}^\top \mathbf{w}_k = 0.$$

If $\bar{n} \neq k$, then we have $\lambda_{\bar{n}} \mathbf{w}_{\bar{n}}^\top \mathbf{w}_n = \lambda_{\bar{n}} \|\mathbf{w}_{\bar{n}}\|_2^2 = \lambda_{\bar{n}} = 0$, which is a contradiction. If $\bar{n} = k$, then we have $(\lambda_{N+1} + \lambda_k) \mathbf{w}_{\bar{n}}^\top \mathbf{w}_k = (\lambda_{N+1} + \lambda_k) \|\mathbf{w}_{\bar{n}}\|_2^2 = \lambda_{N+1} + \lambda_k = 0$ so that $\lambda_{N+1} = \lambda_k = 0$, which is also a contradiction. Hence, the claim holds.

As a consequence of the claim, (20) and (22) yield the equations

$$2(\mathbf{y} - \mathbf{x}) + \lambda_{N+1} \mathbf{w}_k = 0, \quad (23)$$

$$\lambda_{N+1} \mathbf{w}_k^\top \mathbf{y} = 0, \quad (24)$$

which establish Karush-Kuhn-Tucker conditions for the calculation of $g_k(\mathbf{x})$. In particular, $f_k(\mathbf{x}) = g_k(\mathbf{x}) = \|\mathbf{y} - \mathbf{x}\|_2^2$. It follows that $d(\mathbf{x}, (\text{int}(C))^c \cap (\mathbf{x} - C)) = d(\mathbf{x}, (\text{int}(C))^c)$. \square

Corollary B.3. (Theorem 2.4(ii)) Suppose that $C \supseteq \mathbb{R}_+^D$. Then, $\beta_1 = \beta_2 = 1$.

Proof. Since $C \supseteq \mathbb{R}_+^D$, it follows from the definition in (6) that $C^+ \subseteq \mathbb{R}_+^D$. Since $\mathbf{w}_1, \dots, \mathbf{w}_N \in C^+$, we have $\mathbf{w}_n^\top \mathbf{w}_k \geq 0$ for each $n, k \in [N]$. Propositions B.1 and B.2 yield the result. \square

B.2.2 Computation of β_1 for general ordering cones

In the general case, without any additional assumptions on the ordering cone C (besides polyhedrality), we will prove that β_1 can be calculated by solving a global optimization problem over a compact set. We begin with a duality lemma.

Lemma B.4. For each $\mathbf{x} \in \mathbb{R}^D$, we have

$$d(\mathbf{x}, (\mathbf{x} + C) \cap C) = \sup \left\{ \sum_{n \in [N]} \lambda_n (\mathbf{w}_n^\top \mathbf{x})^- \mid \left\| \sum_{n \in [N]} \lambda_n \mathbf{w}_n \right\|_2 \leq 1, \boldsymbol{\lambda} \in \mathbb{R}_+^N \right\}, \quad (25)$$

In particular, $\mathbf{x} \mapsto d(\mathbf{x}, (\mathbf{x} + C) \cap C)$ is a continuous convex function on \mathbb{R}^D .

Proof. Let $\mathbf{x} \in \mathbb{R}^D$. Recall that

$$d(\mathbf{x}, (\mathbf{x} + C) \cap C) = \inf \{ \|\mathbf{y} - \mathbf{x}\|_2 \mid \forall n \in [N]: \mathbf{w}_n^\top \mathbf{y} \geq (\mathbf{w}_n^\top \mathbf{x})^+ \}.$$

By strong duality for convex optimization, we have

$$d(\mathbf{x}, (\mathbf{x} + C) \cap C) = \sup_{\boldsymbol{\lambda} \in \mathbb{R}_+^N} h_1(\boldsymbol{\lambda}),$$

where

$$h_1(\boldsymbol{\lambda}) := \inf_{\mathbf{y} \in \mathbb{R}^D} \left(\|\mathbf{y} - \mathbf{x}\|_2 - \sum_{n \in [N]} \lambda_n \mathbf{w}_n^\top \mathbf{y} + \sum_{n \in [N]} \lambda_n (\mathbf{w}_n^\top \mathbf{x})^+ \right).$$

With a change-of-variables through $\mathbf{u} = \mathbf{y} - \mathbf{x}$, we get

$$\begin{aligned} h_1(\boldsymbol{\lambda}) &= \inf_{\mathbf{u} \in \mathbb{R}^D} \left(\|\mathbf{u}\|_2 - \sum_{n \in [N]} \lambda_n \mathbf{w}_n^\top (\mathbf{u} + \mathbf{x}) + \sum_{n \in [N]} \lambda_n (\mathbf{w}_n^\top \mathbf{x})^+ \right) \\ &= \inf_{\mathbf{u} \in \mathbb{R}^D} \left(\|\mathbf{u}\|_2 - \sum_{n \in [N]} \lambda_n \mathbf{w}_n^\top \mathbf{u} \right) + \sum_{n \in [N]} \lambda_n (\mathbf{w}_n^\top \mathbf{x})^- \\ &= \begin{cases} \sum_{n \in [N]} \lambda_n (\mathbf{w}_n^\top \mathbf{x})^- & \text{if } \left\| \sum_{n \in [N]} \lambda_n \mathbf{w}_n \right\|_2 \leq 1, \\ -\infty & \text{else,} \end{cases} \end{aligned}$$

where the last equality involves the routine calculation of the conjugate of the ℓ_2 -norm. Hence, (25) follows. Clearly, $\mathbf{x} \mapsto \sum_{n \in [N]} \lambda_n (\mathbf{w}_n^\top \mathbf{x})^-$ is a convex function for every $\boldsymbol{\lambda} \in \mathbb{R}_+^N$. Then, by (25), $\mathbf{x} \mapsto d(\mathbf{x}, (\mathbf{x} \cap C) \cap C)$ is convex as a supremum of convex functions. This function is also finite by definition. Hence, by Rockafellar (1970, Corollary 10.1.1), it is continuous as a convex finite function on \mathbb{R}^D . \square

Proposition B.5. *It holds*

$$\beta_1 = \sup\{d(-\mathbf{w}, (-\mathbf{w} + C) \cap C) \mid \|\mathbf{w}\|_2 = 1, \mathbf{w} \in C^+\}. \quad (26)$$

In particular, $\beta_1 < +\infty$.

Proof. Let $\mathbf{p}_C : \mathbb{R}^D \rightarrow C$ be the projection mapping onto C , that is, for each $\mathbf{x} \in \mathbb{R}^D$, $\mathbf{p}_C(\mathbf{x}) \in \mathbb{R}^D$ is the unique point in C for which $d(\mathbf{x}, C) = \|\mathbf{p}_C(\mathbf{x}) - \mathbf{x}\|_2$. Clearly, for $\mathbf{x} \in C^c$, we have $\mathbf{p}_C(\mathbf{x}) \in \text{bd}(C)$.

Let $\mathbf{x} \in C^c$. Then, (15), (16) hold with $\mathbf{y} = \mathbf{p}_C(\mathbf{x})$ for some $\boldsymbol{\lambda} = \boldsymbol{\lambda}^L \in \mathbb{R}_+^N$. Let us define $\bar{\mathbf{x}} := \frac{1}{d(\mathbf{x}, C)} \mathbf{x}$. Note that $\mathbf{w}_n^\top \bar{\mathbf{x}} = \frac{1}{d(\mathbf{x}, C)} \mathbf{w}_n^\top \mathbf{x}$ for every $n \in [N]$. Hence, $I(\mathbf{x}) = I(\bar{\mathbf{x}})$. Let us define $\bar{\mathbf{y}}_L := \frac{1}{d(\mathbf{x}, C)} \mathbf{p}_C(\mathbf{x})$ and $\bar{\boldsymbol{\lambda}}^L := \frac{1}{d(\mathbf{x}, C)} \boldsymbol{\lambda}^L$. By (15) for $(\mathbf{p}_C(\mathbf{x}), \boldsymbol{\lambda}^L)$, we have

$$\bar{\mathbf{y}}_L - \bar{\mathbf{x}} = \frac{1}{d(\mathbf{x}, C)} (\mathbf{p}_C(\mathbf{x}) - \mathbf{x}) = \frac{1}{2} \sum_{n \in [N]} \frac{\lambda_n^L}{d(\mathbf{x}, C)} \mathbf{w}_n = \frac{1}{2} \sum_{n \in [N]} \bar{\lambda}_n^L \mathbf{w}_n. \quad (27)$$

Moreover, for each $n \in [N]$, by (16) for $(\mathbf{p}_C(\mathbf{x}), \boldsymbol{\lambda}^L)$, we have

$$\bar{\lambda}_n^L \mathbf{w}_n^\top \bar{\mathbf{y}} = \frac{\lambda_n^L}{d(\mathbf{x}, C)^2} \mathbf{w}_n^\top \mathbf{p}_C(\mathbf{x}) = 0. \quad (28)$$

Hence, (27) and (28) establish Karush-Kuhn-Tucker conditions for $(\bar{\mathbf{y}}_L, \bar{\boldsymbol{\lambda}}^L)$ so that $\bar{\mathbf{y}}_L = \mathbf{p}_C(\bar{\mathbf{x}})$ and

$$d(\bar{\mathbf{x}}, C) = \|\bar{\mathbf{y}}_L - \bar{\mathbf{x}}\|_2 = \frac{1}{d(\mathbf{x}, C)} \|\mathbf{p}_C(\mathbf{x}) - \mathbf{x}\|_2 = 1.$$

Next, note that (11), (12), (13) hold for some $\mathbf{y} = \mathbf{y}_U \in \mathbb{R}^D$, $\boldsymbol{\lambda} = \boldsymbol{\lambda}^U \in \mathbb{R}_+^N$. Let us define $\bar{\mathbf{y}}_U := \frac{1}{d(\mathbf{x}, C)} \mathbf{y}_U$ and $\bar{\boldsymbol{\lambda}}^U := \frac{1}{d(\mathbf{x}, C)} \boldsymbol{\lambda}^U$. By (11) for $(\mathbf{y}_U, \boldsymbol{\lambda}^U)$, we have

$$\bar{\mathbf{y}}_U - \bar{\mathbf{x}} = \frac{1}{d(\mathbf{x}, C)} (\mathbf{y}_U - \mathbf{x}) = \frac{1}{2} \sum_{n \in [N]} \frac{\lambda_n^U}{d(\mathbf{x}, C)} \mathbf{w}_n = \frac{1}{2} \sum_{n \in [N]} \bar{\lambda}_n^U \mathbf{w}_n. \quad (29)$$

Let $n \in I(\bar{\mathbf{x}}) = I(\mathbf{x})$. Then, by (12) for $(\mathbf{y}_U, \boldsymbol{\lambda}^U)$, we have

$$\bar{\lambda}_n^U \mathbf{w}_n^\top \bar{\mathbf{y}}_U = \frac{\lambda_n^U}{d(\mathbf{x}, C)^2} \mathbf{w}_n^\top \mathbf{y}_U = 0. \quad (30)$$

Let $n \in I(\bar{\mathbf{x}})^c = I(\mathbf{x})^c$. Then, by (13) for $(\mathbf{y}_U, \boldsymbol{\lambda}^U)$, we have

$$\bar{\lambda}_n^U \mathbf{w}_n^\top (\bar{\mathbf{y}}^U - \bar{\mathbf{x}}) = \frac{\lambda_n^U}{d(\mathbf{x}, C)^2} \mathbf{w}_n^\top (\mathbf{y}_U - \mathbf{x}) = 0. \quad (31)$$

Hence, (29), (30), (31) establish Karush-Kuhn-Tucker conditions for $(\bar{\mathbf{y}}_U, \bar{\boldsymbol{\lambda}}^U)$ so that

$$d(\bar{\mathbf{x}}, (\bar{\mathbf{x}} + C) \cap C) = \|\bar{\mathbf{y}}_U - \bar{\mathbf{x}}\|_2 = \frac{\|\mathbf{y}_U - \mathbf{x}\|_2}{d(\mathbf{x}, C)} = \frac{d(\mathbf{x}, (\mathbf{x} + C) \cap C)}{d(\mathbf{x}, C)}.$$

Since we also have $d(\bar{\mathbf{x}}, C) = 1$, we obtain

$$\frac{d(\mathbf{x}, (\mathbf{x} + C) \cap C)}{d(\mathbf{x}, C)} = \frac{d(\bar{\mathbf{x}}, (\bar{\mathbf{x}} + C) \cap C)}{d(\bar{\mathbf{x}}, C)}.$$

Therefore, we may restrict the calculation of β_1 to the set of points that are of unit distance to the ordering cone, i.e.,

$$\beta_1 = \sup\{d(\mathbf{x}, (\mathbf{x} + C) \cap C) \mid \mathbf{x} \in C^c, d(\mathbf{x}, C) = 1\}.$$

Note that a vector $\mathbf{y} \in \mathbb{R}^D$ is in $\text{bd}(C)$ if and only if $\mathbf{y} \in C$ and $\mathbf{w}_n^\top \mathbf{y} > 0$ for at least one $n \in [N]$. In view of this observation, we may write $\text{bd}(C)$ as the union of the faces of C , i.e.,

$$\text{bd}(C) = \bigcup_{I \in 2^{[N]} \setminus \{\emptyset\}} F(I), \text{ where } F(I) := \{\mathbf{y} \in C \mid \forall n \in I: \mathbf{w}_n^\top \mathbf{y} = 0, \forall n \in I^c: \mathbf{w}_n^\top \mathbf{y} > 0\}.$$

For each $I \in 2^{[N]} \setminus \{\emptyset\}$, let

$$\mathcal{X}(I) := \{\mathbf{x} \in C^c \mid \mathbf{p}_C(\mathbf{x}) \in F(I)\}, \quad C^+(I) := \left\{ \sum_{n \in I} \lambda_n \mathbf{w}_n \mid \boldsymbol{\lambda} \in \mathbb{R}_+^N \right\}.$$

We claim that

$$\mathcal{X}(I) = F(I) - C^+(I). \quad (32)$$

To see this, first let us fix $\mathbf{x} \in \mathcal{X}(I)$, i.e., $\mathbf{p}_C(\mathbf{x}) \in F(I)$. Note that $\mathbf{y} = \mathbf{p}_C(\mathbf{x})$ satisfies (15), (16) for some $\boldsymbol{\lambda} \in \mathbb{R}_+^N$. Since $\mathbf{p}_C(\mathbf{x}) \in F(I)$, it follows from (16) that $\lambda_n = 0$ for each $n \in I^c$. Then, $\mathbf{p}_C(\mathbf{x}) - \mathbf{x} = \frac{1}{2} \sum_{n \in I} \lambda_n \mathbf{w}_n \in C^+(I)$. Hence, $\mathbf{x} - \mathbf{p}_C(\mathbf{x}) \in -C^+(I)$. Conversely, let $\mathbf{y} \in F(I)$, $\boldsymbol{\lambda} \in \mathbb{R}_+^N$, and define $\mathbf{x} := \mathbf{y} - \sum_{n \in I} \lambda_n \mathbf{w}_n$. Let us define $\bar{\boldsymbol{\lambda}} \in \mathbb{R}_+^N$ by $\bar{\lambda}_n := \lambda_n$ for each $n \in I$, and by $\bar{\lambda}_n := 0$ for each $n \in I^c$. Then, it is clear that the pair $(\mathbf{y}, \bar{\boldsymbol{\lambda}})$ satisfies (15), (16) for the point \mathbf{x} . Hence, we conclude that $\mathbf{p}_C(\mathbf{x}) = \mathbf{y}$. Therefore, $\mathbf{x} \in \mathcal{X}(I)$, which completes the proof of (32). The last part of the proof also shows that $d(\mathbf{x}, C) = \|\mathbf{w}\|_2$ whenever $\mathbf{x} = \mathbf{y} - \mathbf{w}$ with $\mathbf{y} \in F(I)$ and $\mathbf{w} \in C^+(I)$.

Note that we may write $C^c = \bigcup_{I \in 2^{[N]} \setminus \{\emptyset\}} \mathcal{X}(I)$ as a disjoint finite union. Accordingly, we have

$$\beta_1 = \max_{I \in 2^{[N]} \setminus \{\emptyset\}} \beta_1(I), \text{ where } \beta_1(I) := \sup\{d(\mathbf{x}, (\mathbf{x} + C) \cap C) \mid d(\mathbf{x}, C) = 1, \mathbf{x} \in \mathcal{X}(I)\}.$$

Let us fix $I \in 2^{[N]} \setminus \{\emptyset\}$ and $\mathbf{x} \in \mathcal{X}(I)$ with $d(\mathbf{x}, C) = 1$. We may write $\mathbf{x} = \mathbf{p}_C(\mathbf{x}) - \mathbf{w}$ for some $\mathbf{w} \in C^+(I)$ and $d(\mathbf{x}, C) = \|\mathbf{w}\|_2 = 1$. Let us define $\bar{\mathbf{x}} := -\mathbf{w}$. Let $n \in [N]$. Since $\mathbf{p}_C(\mathbf{x}) \in F(I)$, we have

$$\mathbf{w}_n^\top \bar{\mathbf{x}} = \mathbf{w}_n^\top (\mathbf{x} - \mathbf{p}_C(\mathbf{x})) = \mathbf{w}_n^\top \mathbf{x}$$

if $n \in I$, and

$$\mathbf{w}_n^\top \bar{\mathbf{x}} = \mathbf{w}_n^\top (\mathbf{x} - \mathbf{p}_C(\mathbf{x})) < \mathbf{w}_n^\top \mathbf{x}$$

if $n \in I^c$. Hence, for every $n \in [N]$, we have $\mathbf{w}_n^\top \bar{\mathbf{x}} \leq \mathbf{w}_n^\top \mathbf{x}$ so that $(\mathbf{w}_n^\top \bar{\mathbf{x}})^- \geq (\mathbf{w}_n^\top \mathbf{x})^-$. Hence, by Lemma B.4,

$$\begin{aligned} d(\mathbf{x}, (\mathbf{x} + C) \cap C) &= \sup \left\{ \sum_{n \in [N]} \lambda_n (\mathbf{w}_n^\top \mathbf{x})^- \mid \left\| \sum_{n \in [N]} \lambda_n \mathbf{w}_n \right\|_2 \leq 1, \boldsymbol{\lambda} \in \mathbb{R}_+^N \right\} \\ &\leq \sup \left\{ \sum_{n \in [N]} \lambda_n (\mathbf{w}_n^\top \bar{\mathbf{x}})^- \mid \left\| \sum_{n \in [N]} \lambda_n \mathbf{w}_n \right\|_2 \leq 1, \boldsymbol{\lambda} \in \mathbb{R}_+^N \right\} \end{aligned}$$

$$= d(\bar{\mathbf{x}}, (\bar{\mathbf{x}} + C) \cap C).$$

Moreover, note that $p_C(\bar{\mathbf{x}}) = 0 \in F([N])$, $\bar{\mathbf{x}} \in \mathcal{X}([N])$, and $d(\bar{\mathbf{x}}, C) = \|\mathbf{w}\|_2 = 1$. As $\bar{\mathbf{x}} \in \mathcal{X}([N])$ provides a higher objective value than $\mathbf{x} \in \mathcal{X}([I])$, we can ignore all choices of I but the case $I = [N]$ in the computation of β_1 :

$$\beta_1 = \sup\{d(\mathbf{x}, (\mathbf{x} + C) \cap C) \mid d(\mathbf{x}, C) = 1, \mathbf{x} \in \mathcal{X}([N])\}. \quad (33)$$

Note that $F([N]) = \{0\}$ since C is assumed to be pointed, and $C^+([N]) = C^+$. Hence, $\mathcal{X}([N]) = F([N]) - C^+(I) = -C^+$ by (32). Therefore, (33) becomes

$$\begin{aligned} \beta_1 &= \sup\{d(-\mathbf{w}, (-\mathbf{w} + C) \cap C) \mid d(-\mathbf{w}, C) = 1, \mathbf{w} \in C^+\} \\ &= \sup\{d(-\mathbf{w}, (-\mathbf{w} + C) \cap C) \mid \|\mathbf{w}\|_2 = 1, \mathbf{w} \in C^+\}. \end{aligned}$$

Hence, (26) follows.

Finally, note that $\{\mathbf{w} \in C^+ \mid \|\mathbf{w}\|_2 = 1\}$ is a compact set. Moreover, $\mathbf{w} \mapsto d(-\mathbf{w}, (-\mathbf{w} + C) \cap C)$ is a continuous function by Lemma B.4. Hence, there exists $\mathbf{w}^* \in C^+$ such that $\|\mathbf{w}^*\|_2 = 1$ and $\beta_1 = d(-\mathbf{w}^*, (-\mathbf{w}^* + C) \cap C) < +\infty$. \square

The optimization problem in (26) consists of maximizing a finite convex objective function over a compact set; hence, it is a global (non-convex) optimization problem. Nevertheless, the objective function is indeed Lipschitz over $\{\mathbf{w} \in C^+ \mid \|\mathbf{w}\|_2 = 1\}$ by Rockafellar (1970, Theorem 10.4). Hence, some classical techniques of global optimization (e.g., Goldstein (1977)) can be used to calculate the actual value of β_1 (approximately). To be able to use these methods, one still needs to find a Lipschitz constant for the function $\mathbf{w} \mapsto d(-\mathbf{w}, (-\mathbf{w} + C) \cap C)$ over $\{\mathbf{w} \in C^+ \mid \|\mathbf{w}\|_2 = 1\}$. A more convenient alternative would be to use the same methods of global optimization for the following dual representation of β_1 .

Theorem B.6. *It holds*

$$\begin{aligned} \beta_1 &= \sup_{\substack{\mathbf{w} \in C^+ \\ \|\mathbf{w}\|_2 = 1}} \sup_{\substack{\boldsymbol{\lambda} \in \mathbb{R}_+^N \\ \|\sum_{n \in [N]} \lambda_n \mathbf{w}_n\|_2 \leq 1}} \sum_{n \in [N]} \lambda_n (\mathbf{w}_n^\top \mathbf{w})^+ \\ &= \max_{I \in 2^{[N]}} \sup_{\substack{\mathbf{w} \in C^+ \cap C(I) \\ \|\mathbf{w}\|_2 = 1}} \sup_{\substack{\boldsymbol{\lambda} \in \mathbb{R}_+^N \\ \|\sum_{n \in [N]} \lambda_n \mathbf{w}_n\|_2 \leq 1}} \sum_{n \in I} \lambda_n \mathbf{w}_n^\top \mathbf{w}, \end{aligned}$$

where

$$C(I) := \{\mathbf{y} \in \mathbb{R}^D \mid \forall n \in I: \mathbf{w}_n^\top \mathbf{y} \geq 0, \forall n \in I^c: \mathbf{w}_n^\top \mathbf{y} \leq 0\}.$$

Proof. By Lemma B.4, for each $\mathbf{w} \in C^+$, we have

$$d(-\mathbf{w}, (-\mathbf{w} + C) \cap C) = \sup \left\{ \sum_{n \in [N]} \lambda_n (\mathbf{w}_n^\top \mathbf{w})^+ \mid \left\| \sum_{n \in [N]} \lambda_n \mathbf{w}_n \right\|_2 \leq 1, \boldsymbol{\lambda} \in \mathbb{R}_+^N \right\}.$$

Combining this and Proposition B.5, the first equality in the theorem follows. The second equality follows immediately by decomposing C^+ as $C^+ = \bigcup_{I \in 2^{[N]}} (C^+ \cap C(I))$ and noting that $\sum_{n \in [N]} \lambda_n (\mathbf{w}_n^\top \mathbf{w})^+ = \sum_{n \in I} \lambda_n \mathbf{w}_n^\top \mathbf{w}$ whenever $\mathbf{w} \in C(I)$. \square

We complete this subsection by showing that, for each $I \in 2^{[N]}$, the objective function of the second problem in Theorem B.6 is jointly Lipschitz with a computable Lipschitz constant. This will be achieved by the next proposition. In what follows, $\|\cdot\|_1$ denotes the ℓ_1 -norm.

Proposition B.7. (i) *Let $v^* := \inf\{\|\sum_{n \in [N]} \lambda \mathbf{w}_n\|_2 \mid \|\boldsymbol{\lambda}\|_1 = 1, \boldsymbol{\lambda} \in \mathbb{R}_+^N\}$. Then, $v^* > 0$.*

(ii) *The set $\{\boldsymbol{\lambda} \in \mathbb{R}_+^N \mid \|\sum_{n \in [N]} \lambda_n \mathbf{w}_n\|_2 \leq 1\}$ is a nonempty compact subset of \mathbb{R}^N .*

(iii) *The function $(\mathbf{w}, \boldsymbol{\lambda}) \mapsto \sum_{n \in I} \lambda_n \mathbf{w}_n^\top \mathbf{w}$ is $(\frac{1}{v^*} \vee 1)$ -Lipschitz with respect to the $\ell_2 \times \ell_1$ -norm on $\mathbb{R}^D \times \mathbb{R}^N$.*

Proof. (i) Let us define $\Lambda := \{\boldsymbol{\lambda} \in \mathbb{R}_+^N \mid \|\boldsymbol{\lambda}\|_1 = 1\}$, which is a compact set. The function $\boldsymbol{\lambda} \mapsto \|\sum_{n \in [N]} \lambda_n \mathbf{w}_n\|_2$ is continuous, hence it attains its minimum over Λ at some $\boldsymbol{\lambda}^* \in \Lambda$. We claim that the corresponding minimum value $v^* = \|\sum_{n \in [N]} \lambda_n^* \mathbf{w}_n\|_2$ is nonzero. Indeed, having $v^* = 0$ would imply $\sum_{n \in [N]} \lambda_n^* \mathbf{w}_n = 0$, which is not possible since $\mathbf{w}_1, \dots, \mathbf{w}_N$ are the generating vectors of the dual cone C^+ and C^+ is pointed (as C is solid). Therefore, the claim holds.

(ii) By (i), we have

$$\left\| \sum_{n \in [N]} \lambda_n \mathbf{w}_n \right\|_2 \geq v^* > 0 \quad (34)$$

for every $\boldsymbol{\lambda} \in \Lambda$. Finally, for an arbitrary vector $\boldsymbol{\lambda} \in \mathbb{R}_+^N \setminus \{\mathbf{0}\}$, we have

$$\left\| \sum_{n \in [N]} \lambda_n \mathbf{w}_n \right\|_2 \geq v^* \|\boldsymbol{\lambda}\|_1, \quad (35)$$

which follows by applying (34) to $\boldsymbol{\lambda}/\|\boldsymbol{\lambda}\|_1 \in \Lambda$. Moreover, (35) holds for $\boldsymbol{\lambda} = \mathbf{0}$ trivially. Therefore, (35) holds for every $\boldsymbol{\lambda} \in \mathbb{R}_+^N$. In particular,

$$\sup \left\{ \|\boldsymbol{\lambda}\|_1 \mid \left\| \sum_{n \in [N]} \lambda_n \mathbf{w}_n \right\|_2 \leq 1, \boldsymbol{\lambda} \in \mathbb{R}_+^N \right\} \leq \frac{1}{v^*} < +\infty,$$

which shows the boundedness of $\{\boldsymbol{\lambda} \in \mathbb{R}_+^N \mid \|\sum_{n \in [N]} \lambda_n \mathbf{w}_n\|_2 \leq 1\}$. Trivially, this set is also closed. Hence, compactness follows.

(iii) For every $\boldsymbol{\lambda}, \bar{\boldsymbol{\lambda}} \in \mathbb{R}_+^N$ such that $\|\sum_{n \in [N]} \lambda_n \mathbf{w}_n\|_2 \leq 1$, $\|\sum_{n \in [N]} \bar{\lambda}_n \mathbf{w}_n\|_2 \leq 1$ and $\mathbf{w}, \bar{\mathbf{w}} \in C^+$ such that $\|\mathbf{w}\|_2 = \|\bar{\mathbf{w}}\|_2 = 1$, we observe that

$$\begin{aligned} \left| \sum_{n \in I} \lambda_n \mathbf{w}_n^\top \mathbf{w} - \sum_{n \in I} \bar{\lambda}_n \mathbf{w}_n^\top \bar{\mathbf{w}} \right| &\leq \left\| \sum_{n \in I} \lambda_n \mathbf{w}_n \right\|_2 \|\mathbf{w} - \bar{\mathbf{w}}\|_2 + \left\| \sum_{n \in I} (\lambda_n - \bar{\lambda}_n) \mathbf{w}_n \right\|_2 \|\bar{\mathbf{w}}\|_2 \\ &\leq \sum_{n \in I} \lambda_n \|\mathbf{w} - \bar{\mathbf{w}}\|_2 + \sum_{n \in I} |\lambda_n - \bar{\lambda}_n| \\ &\leq \frac{1}{v^*} \|\mathbf{w} - \bar{\mathbf{w}}\|_2 + \|\boldsymbol{\lambda} - \bar{\boldsymbol{\lambda}}\|_1 \\ &\leq \left(\frac{1}{v^*} \vee 1 \right) (\|\mathbf{w} - \bar{\mathbf{w}}\|_2 + \|\boldsymbol{\lambda} - \bar{\boldsymbol{\lambda}}\|_1). \end{aligned}$$

Hence, the Lipschitz property follows. \square

The constant

$$v^* = \inf \left\{ \left\| \sum_{n \in [N]} \lambda_n \mathbf{w}_n \right\|_2 \mid \|\boldsymbol{\lambda}\|_1 = 1, \boldsymbol{\lambda} \in \mathbb{R}_+^N \right\} = \inf \left\{ \left\| \sum_{n \in [N]} \lambda_n \mathbf{w}_n \right\|_2 \mid \|\boldsymbol{\lambda}\|_1 \leq 1, \boldsymbol{\lambda} \in \mathbb{R}_+^N \right\},$$

which appears in Proposition B.7, can be calculated by solving a simple minimization problem whose objective function is quadratic and constraints can be linearized. Therefore, in view of Theorem B.6, to evaluate β_1 approximately, one can solve finitely many optimization problems with a compact feasible region and Lipschitz objective function with a known Lipschitz constant.

B.2.3 Computation of β_2 for general ordering cones

We will find an upper bound for β_2 . For this, we first provide a simple formula for the calculation of the distance functions that appear in the definition of β_2 .

Lemma B.8. *Let $\mathbf{x} \in (\text{int}(C))^c$. Then,*

$$d(\mathbf{x}, (\text{int}(C))^c \cap (\mathbf{x} - C)) = \min_{n \in [N]} \frac{\mathbf{w}_n^\top \mathbf{x}}{\alpha_n}, \quad d(\mathbf{x}, (\text{int}(C))^c) = \min_{n \in [N]} \mathbf{w}_n^\top \mathbf{x}.$$

Proof. To see the first identity, note that replacing $\boldsymbol{\mu}_j - \boldsymbol{\mu}_i$ with \mathbf{x} in Proposition 4.2(ii,iv) gives $d(\mathbf{x}, (\text{int}(C))^c \cap (\mathbf{x} - C)) = \min_{n \in [N]} \frac{(\mathbf{w}_n^\top \mathbf{x})^+}{\alpha_n}$. Since $\mathbf{x} \in \text{int}(C)$, we have $\mathbf{w}_n^\top \mathbf{x} > 0$ for each $n \in [N]$. Hence, the first identity follows. Alternatively, one can use (18) and formulate the Lagrange dual problem of $f_n(\mathbf{x})$ for each $n \in [N]$, which yields the same identity.

To prove the second identity, let us use (19). By strong duality for convex optimization, for each $n \in [N]$, we have

$$g_n(\mathbf{x}) = \sup_{\boldsymbol{\lambda} \geq 0} \inf_{\mathbf{y} \in \mathbb{R}^D} (\|\mathbf{y} - \mathbf{x}\|_2 + \boldsymbol{\lambda} \mathbf{w}_n^\top \mathbf{y})$$

$$\begin{aligned}
&= \sup_{\lambda \geq 0} \inf_{\mathbf{u} \in \mathbb{R}^D} (\|\mathbf{u}\|_2 + \lambda \mathbf{w}_n^\top (\mathbf{u} + \mathbf{x})) \\
&= \sup_{\lambda \geq 0} \left(\lambda \mathbf{w}_n^\top \mathbf{x} - \sup_{\mathbf{u} \in \mathbb{R}^D} (-\lambda \mathbf{w}_n^\top \mathbf{u} - \|\mathbf{u}\|_2) \right) \\
&= \sup_{\lambda \geq 0} \begin{cases} \lambda \mathbf{w}_n^\top \mathbf{x} & \text{if } \lambda = \|\lambda \mathbf{w}_n\| \leq 1 \\ -\infty & \text{else} \end{cases} \\
&= \sup_{\lambda \in [0,1]} \lambda \mathbf{w}_n^\top \mathbf{x} = \mathbf{w}_n^\top \mathbf{x},
\end{aligned}$$

where the last equality holds since $\mathbf{w}_n^\top \mathbf{x} > 0$. Hence, the second identity follows. \square

Theorem B.9. *It holds*

$$\beta_2 \leq \frac{1}{\min_{n \in [N]} \alpha_n} < +\infty.$$

Proof. Let $n^* \in [N]$ be such that $\alpha_{n^*} = \min_{n \in [N]} \alpha_n$. By the definition of β_2 and Lemma B.8, we have

$$\begin{aligned}
\beta_2 &= \sup_{\mathbf{x} \in \text{int}(C)} \frac{d(\mathbf{x}, (\text{int}(C))^c \cap (\mathbf{x} - C))}{d(\mathbf{x}, (\text{int}(C))^c)} = \sup_{\mathbf{x} \in \text{int}(C)} \frac{\min_{n \in [N]} \frac{\mathbf{w}_n^\top \mathbf{x}}{\alpha_n}}{\min_{n \in [N]} \mathbf{w}_n^\top \mathbf{x}} \\
&\leq \frac{1}{\alpha_{n^*}} \sup_{\mathbf{x} \in \text{int}(C)} \frac{\min_{n \in [N]} \mathbf{w}_n^\top \mathbf{x}}{\min_{n \in [N]} \mathbf{w}_n^\top \mathbf{x}} = \frac{1}{\alpha_{n^*}}.
\end{aligned}$$

\square

Note that Proposition B.5 and Theorem B.9 yield Theorem 2.4(i).

B.3 Proof of Proposition 4.2

(i) To get a contradiction, suppose that $m(i, j) = +\infty$. Hence, the set whose infimum is calculated in (8) is empty, that is, for every $s \geq 0$ and $\mathbf{u} \in B(\mathbf{0}, 1) \cap C$, we have $\boldsymbol{\mu}_i + s\mathbf{u} \in \boldsymbol{\mu}_j - \text{int}(C)$. Since $\{s\mathbf{u} \mid s \geq 0, \mathbf{u} \in B(\mathbf{0}, 1) \cap C\} = C$, we have $\boldsymbol{\mu}_i + C \subseteq \boldsymbol{\mu}_j - \text{int}(C) \subseteq \boldsymbol{\mu}_j - C$. Hence, $\boldsymbol{\mu}_i - \boldsymbol{\mu}_j + C \subseteq -C$ so that $C \cap (-C) \supseteq C \cap (\boldsymbol{\mu}_i - \boldsymbol{\mu}_j + C)$. The latter intersection consists of all points $\mathbf{x} \in \mathbb{R}^D$ such that $\mathbf{0} \preceq_C \mathbf{x}$ and $\boldsymbol{\mu}_i - \boldsymbol{\mu}_j \preceq_C \mathbf{x}$. Since C is a solid cone, there are infinitely many such points, which contradicts with $C \cap (-C) = \{\mathbf{0}\}$. Hence, $m(i, j) < +\infty$.

(ii) Note that

$$\begin{aligned}
m(i, j) &= \inf \{s \geq 0 \mid \exists \mathbf{u} \in B(\mathbf{0}, 1) \cap C: \boldsymbol{\mu}_i + s\mathbf{u} \in (\boldsymbol{\mu}_j - \text{int}(C))^c\} \\
&= \inf \{s \geq 0 \mid \exists \mathbf{u} \in B(\mathbf{0}, 1) \cap C: \boldsymbol{\mu}_i + s\mathbf{u} \in \boldsymbol{\mu}_j - (\text{int}(C))^c\} \\
&= \inf_{\mathbf{c} \in (\text{int}(C))^c} \inf \{s \geq 0 \mid \exists \mathbf{u} \in B(\mathbf{0}, 1) \cap C: \boldsymbol{\mu}_j - \boldsymbol{\mu}_i - \mathbf{c} = s\mathbf{u}\} \\
&= \inf_{\mathbf{c} \in (\text{int}(C))^c \cap (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i - C)} \inf \{s \geq 0 \mid s = \|\boldsymbol{\mu}_j - \boldsymbol{\mu}_i - \mathbf{c}\|_2\} \\
&= \inf_{\mathbf{c} \in (\text{int}(C))^c \cap (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i - C)} \|\boldsymbol{\mu}_j - \boldsymbol{\mu}_i - \mathbf{c}\|_2 \\
&= d(\boldsymbol{\mu}_j - \boldsymbol{\mu}_i, (\text{int}(C))^c \cap (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i - C)).
\end{aligned}$$

(iii) Since $(\text{int}(C))^c \cap (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i - C)$ is a closed set, by the well-known properties of distance function, we have $\boldsymbol{\mu}_j - \boldsymbol{\mu}_i \in (\text{int}(C))^c \cap (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i - C)$ if and only if $m(i, j) = d(\boldsymbol{\mu}_j - \boldsymbol{\mu}_i, (\text{int}(C))^c \cap (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i - C)) = 0$. Since $\boldsymbol{\mu}_j - \boldsymbol{\mu}_i \in \boldsymbol{\mu}_j - \boldsymbol{\mu}_i - C$ always holds, we have $\boldsymbol{\mu}_j - \boldsymbol{\mu}_i \in (\text{int}(C))^c$ if and only if $m(i, j) = 0$. Since the former condition precisely means that $i \not\prec_C j$, the desired equivalence follows.

(iv) First, suppose that $m(i, j) = 0$, that is, $\boldsymbol{\mu}_j - \boldsymbol{\mu}_i \notin \text{int}(C)$. Since $\text{int}(C) = \{\mathbf{x} \in \mathbb{R}^D \mid \mathbf{W}\mathbf{x} > \mathbf{0}\}$, there exists $\bar{n} \in [N]$ such that $\mathbf{w}_{\bar{n}}^\top (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i) \leq 0$, that is, $(\mathbf{w}_{\bar{n}}^\top (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i))^+ = 0$. Hence, $\min_{n \in [N]} (\mathbf{w}_n^\top (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i))^+ / \alpha_n = 0 = m(i, j)$. Next, suppose that $m(i, j) > 0$, that is, $\boldsymbol{\mu}_j - \boldsymbol{\mu}_i \in \text{int}(C)$. Note that we may write $m(i, j) = \inf R$, where

$$R := \{s \geq 0 \mid \exists \mathbf{u} \in B(\mathbf{0}, 1) \cap C: \boldsymbol{\mu}_i + s\mathbf{u} \notin \boldsymbol{\mu}_j - C\}.$$

We show that R is an interval that is unbounded from above. To that end, let $s \geq 0$ be such that $\boldsymbol{\mu}_i + s\mathbf{u} \notin \boldsymbol{\mu}_j - C$ for some $\mathbf{u} \in B(\mathbf{0}, 1) \cap C$. Let $s' > s$. We claim that $\boldsymbol{\mu}_i + s'\mathbf{u} \notin \boldsymbol{\mu}_j - C$. Suppose otherwise that $\boldsymbol{\mu}_i + s'\mathbf{u} \in \boldsymbol{\mu}_j - C$. Since $\boldsymbol{\mu}_i \in \boldsymbol{\mu}_j - \text{int}(C) \subseteq \boldsymbol{\mu}_j - C$ by supposition, and $\boldsymbol{\mu}_j - C$ is a convex set, we obtain

$$\boldsymbol{\mu}_i + s\mathbf{u} = \left(1 - \frac{s}{s'}\right) \boldsymbol{\mu}_i + \frac{s}{s'}(\boldsymbol{\mu}_i + s'\mathbf{u}) \in \boldsymbol{\mu}_j - C,$$

which is a contradiction to the definition of s . Hence, the claim holds. Since R is an interval in $[0, +\infty)$ that is unbounded from above, we have $\inf R = \sup([0, +\infty) \setminus R)$. Therefore,

$$\begin{aligned} m(i, j) &= \sup\{s \geq 0 \mid \forall \mathbf{u} \in B(\mathbf{0}, 1) \cap C: \boldsymbol{\mu}_i + s\mathbf{u} \in \boldsymbol{\mu}_j - C\} \\ &= \sup\{s \geq 0 \mid \forall \mathbf{u} \in B(\mathbf{0}, 1) \cap C, \forall n \in [N]: \mathbf{w}_n^\top(\boldsymbol{\mu}_j - \boldsymbol{\mu}_i - s\mathbf{u}) \geq 0\} \\ &= \sup\{s \geq 0 \mid \forall \mathbf{u} \in B(\mathbf{0}, 1) \cap C, \forall n \in [N]: \mathbf{w}_n^\top(\boldsymbol{\mu}_j - \boldsymbol{\mu}_i) \geq s\mathbf{w}_n^\top\mathbf{u}\} \\ &= \sup\{s \geq 0 \mid \forall n \in [N]: \mathbf{w}_n^\top(\boldsymbol{\mu}_j - \boldsymbol{\mu}_i) \geq s \sup_{\mathbf{u} \in B(\mathbf{0}, 1) \cap C} \mathbf{w}_n^\top\mathbf{u}\} \\ &= \sup\{s \geq 0 \mid \forall n \in [N]: \mathbf{w}_n^\top(\boldsymbol{\mu}_j - \boldsymbol{\mu}_i) \geq s\alpha_n\} \\ &= \sup\left\{s \geq 0 \mid \min_{n \in [N]} \frac{1}{\alpha_n} \mathbf{w}_n^\top(\boldsymbol{\mu}_j - \boldsymbol{\mu}_i) \geq s\right\} \\ &= \max\left\{0, \min_{n \in [N]} \frac{1}{\alpha_n} \mathbf{w}_n^\top(\boldsymbol{\mu}_j - \boldsymbol{\mu}_i)\right\} \\ &= \min_{n \in [N]} \frac{(\mathbf{w}_n^\top(\boldsymbol{\mu}_j - \boldsymbol{\mu}_i))^+}{\alpha_n}, \end{aligned}$$

which completes the proof.

B.4 Proof of Proposition 4.3

(i) We prove that the set whose infimum is calculated in (9) is nonempty. To get a contradiction, suppose that for every $s \geq 0$ and $\mathbf{u} \in B(\mathbf{0}, 1) \cap C$, we have $\boldsymbol{\mu}_j - \boldsymbol{\mu}_i + s\mathbf{u} \notin C$, that is,

$$\left(\boldsymbol{\mu}_j - \boldsymbol{\mu}_i + \bigcup_{s \geq 0} (B(\mathbf{0}, s) \cap C)\right) \cap C = \emptyset.$$

However, we have $\bigcup_{s \geq 0} (B(\mathbf{0}, s) \cap C) = C$; hence, we get a contradiction to the solidity of C as in the proof of Proposition 4.2(i). It follows that $M(i, j) < +\infty$.

(ii) By elementary calculations, we have

$$\begin{aligned} M(i, j) &= \inf\{s \geq 0 \mid \exists \mathbf{u} \in B(\mathbf{0}, 1) \cap C, \exists \mathbf{c} \in C: \boldsymbol{\mu}_j + s\mathbf{u} = \boldsymbol{\mu}_i + \mathbf{c}\} \\ &= \inf_{\mathbf{c} \in C} \inf\{s \geq 0 \mid \exists \mathbf{u} \in B(\mathbf{0}, 1) \cap C: \boldsymbol{\mu}_j - \boldsymbol{\mu}_i - \mathbf{c} = -s\mathbf{u}\} \\ &= \inf_{\mathbf{c} \in C \cap (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i + C)} \inf\{s \geq 0 \mid s = \|\boldsymbol{\mu}_j - \boldsymbol{\mu}_i - \mathbf{c}\|_2\} \\ &= \inf_{\mathbf{c} \in C \cap (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i + C)} \|\boldsymbol{\mu}_j - \boldsymbol{\mu}_i - \mathbf{c}\|_2 \\ &= d(\boldsymbol{\mu}_j - \boldsymbol{\mu}_i, C \cap (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i + C)). \end{aligned}$$

(iii) Since $C \cap (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i + C)$ is a closed set, we have $M(i, j) = d(\boldsymbol{\mu}_j - \boldsymbol{\mu}_i, C \cap (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i + C)) = 0$ if and only if $\boldsymbol{\mu}_j - \boldsymbol{\mu}_i \in C \cap (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i + C)$. Since $\boldsymbol{\mu}_j - \boldsymbol{\mu}_i \in \boldsymbol{\mu}_j - \boldsymbol{\mu}_i + C$ is always the case, these conditions are also equivalent to $\boldsymbol{\mu}_j - \boldsymbol{\mu}_i \in C$, that is, $i \preceq_C j$.

B.5 Proof of Theorem 5.1

For each $s > 0$, let us define

$$A(s) := \bigcap_{\mathbf{u} \in B(\mathbf{0}, 1) \cap C} (s\mathbf{u} + C), \quad f(s) := d(\mathbf{0}, A(s)) = \inf\{\|\mathbf{z}\| \mid \mathbf{z} \in A(s)\}.$$

Note that $\mathbf{z} \in A(s)$ if and only if $\mathbf{w}_n^\top(\mathbf{z} - s\mathbf{u}) \geq 0$ for every $n \in [N]$ and $\mathbf{u} \in B(\mathbf{0}, 1)$, which is equivalent to $\mathbf{w}_n^\top \mathbf{z} \geq s\alpha_n = s \sup_{\mathbf{u} \in B(\mathbf{0}, 1) \cap C} \mathbf{w}_n^\top \mathbf{u}$ for every $n \in [N]$. It follows that

$$A(s) = \{\mathbf{z} \in \mathbb{R}^D \mid \forall n \in [N]: \mathbf{w}_n^\top \mathbf{z} \geq s\alpha_n\}, \quad A(s) = sA(1), \quad f(s) = sf(1).$$

Since $A(1)$ is a nonempty closed set, it can be checked that there exists $\mathbf{z}^* \in A(1)$ such that $\|\mathbf{z}^*\| = f(1)$, this vector can be calculated by solving a simple quadratic minimization problem with affine constraints. Since $\mathbf{w}_n^\top \mathbf{z}^* \geq \alpha_n > 0$ for each $n \in [N]$, we have $\mathbf{z}^* \in \text{int}(C)$. Let $\mathbf{u}^* := \frac{\mathbf{z}^*}{f(1)} \in B(\mathbf{0}, 1) \cap \text{int}(C)$. Then, for every $s > 0$, we have $sf(1)\mathbf{u}^* \in A(s)$, whose norm yields $f(s)$.

For each $t \in [T]$, we assume that the noise vector \mathbf{Y}_t is either $\frac{1}{4}\mathbf{u}^*$ with probability $p \in (0, 1)$ or it is $-\frac{1}{4}\mathbf{u}^*$ with probability $1 - p$, which implies $\mathbb{E}[\mathbf{Y}_t] = (\frac{p}{2} - \frac{1}{4})\mathbf{u}^*$. Moreover, due to the structure of this distribution, the difference $\hat{\boldsymbol{\mu}}_i - \boldsymbol{\mu}_i$ between the empirical and true means of a design $i \in [K]$ has a discrete distribution over a finite subset of the line segment that connects $-\frac{1}{4}\mathbf{u}^*$ and $\frac{1}{4}\mathbf{u}^*$, hence we may write $\hat{\boldsymbol{\mu}}_i - \boldsymbol{\mu}_i = (\frac{\hat{p}_i}{2} - \frac{1}{4})\mathbf{u}^*$, where \hat{p}_i is the empirical frequency of observing $\mathbf{Y}_t = \frac{1}{4}\mathbf{u}^*$ over all rounds t at which design i is sampled.

As in Auer et al. (2016), we consider the case $p_i = \frac{1}{2}$ for every $i \in [K]$ as the original case, where the noise vectors are centered as in the problem setup. We will introduce modifications on $(p_i)_{i \in [N]}$ in such a way that the algorithm returns an output P' that is different from the original one P . This could be due to a change in the Pareto set or a change in the gap values of the designs.

We consider the following four possibilities for $i \in [K]$.

Case 1: Suppose that $i \in P^* \setminus P$. Then, Remark 3.7 implies that there exist $j \in P$ and $\mathbf{u} \in B(\mathbf{0}, 1) \cap C$ such that $\boldsymbol{\mu}_j + \epsilon\mathbf{u} \in \boldsymbol{\mu}_i + C$. Let $J(i) \subseteq P$ be the set of all such j . Then, by the definition of $M(i, j)$, we have $0 < M(i, j) \leq \epsilon$ for every $j \in J(i)$. In particular, $M(i, j) \leq \epsilon$ for every $j \in J(i)$ and $\tilde{\Delta}_i^\epsilon = \epsilon$. Consider a modification of p_i such that the corresponding mean vector becomes $\boldsymbol{\mu}'_i := \boldsymbol{\mu}_i + 2\tilde{\Delta}_i^\epsilon f(1)\mathbf{u}^* = \boldsymbol{\mu}_i + 2\epsilon f(1)\mathbf{u}^*$ while all the other mean vectors remain the same. Note that $\boldsymbol{\mu}'_i - \boldsymbol{\mu}_i = 2\epsilon f(1)\mathbf{u}^* \in A(2\epsilon)$, i.e., $\boldsymbol{\mu}'_i \in \boldsymbol{\mu}_i + 2\epsilon\mathbf{u} + C$ for every $\mathbf{u} \in B(\mathbf{0}, 1) \cap C$. We claim that i is not covered by any $j \in J(i)$ in the modified case. Indeed, if we had $j \in J(i)$ and $\mathbf{u} \in B(\mathbf{0}, 1) \cap C$ such that $\boldsymbol{\mu}_j + \epsilon\mathbf{u} \in \boldsymbol{\mu}'_i + C$, then the properties of $\boldsymbol{\mu}'_i$ would imply that $\boldsymbol{\mu}_j + \epsilon\mathbf{u} \in \boldsymbol{\mu}_i + 2\epsilon\mathbf{u} + C$, i.e., $\boldsymbol{\mu}_j \in \boldsymbol{\mu}_i + \epsilon\mathbf{u} + C \subseteq \boldsymbol{\mu}_i + C$ so that $\boldsymbol{\mu}_j$ dominates $\boldsymbol{\mu}_i$ and $i \notin P^*$, a contradiction. Hence, the claim holds. On the other hand, we have $i \in (P^*)'$, where $(P^*)'$ denotes the Pareto set in the modified case. Hence, due to Definition 3.6, we must have $i \in P'$, where P' is the set returned by the algorithm in the modified case.

Case 2: Suppose that $i \in P^* \cap P$. Then, $\Delta_i^+ = M(i, j)$ for some $j \in P^* \setminus \{i\}$. Let us modify $\boldsymbol{\mu}_i$ as $\boldsymbol{\mu}'_i := \boldsymbol{\mu}_i - k\tilde{\Delta}_i^\epsilon \mathbf{u}^*$, where $k > 0$ is to be determined. We want to choose k in such a way that $m'(i, j) \geq 2\epsilon$ for the gap $m'(i, j)$ in the modified case so that $i \notin P'$ due to Definition 3.6. Note that $m'(i, j) \geq 2\epsilon$ if and only if $\boldsymbol{\mu}'_i + 2\epsilon\mathbf{u} \in \boldsymbol{\mu}_j - C$ for every $\mathbf{u} \in B(\mathbf{0}, 1) \cap C$. Hence, the minimum value of k is given by

$$\begin{aligned} k_i &:= \inf\{k \geq 0 \mid \forall \mathbf{u} \in B(\mathbf{0}, 1) \cap C: \boldsymbol{\mu}_i - k\tilde{\Delta}_i^\epsilon \mathbf{u}^* + 2\epsilon\mathbf{u} \in \boldsymbol{\mu}_j - C\} \\ &= \inf\{k \geq 0 \mid \forall \mathbf{u} \in B(\mathbf{0}, 1) \cap C: \boldsymbol{\mu}_j - \boldsymbol{\mu}_i + k\tilde{\Delta}_i^\epsilon \mathbf{u}^* \in 2\epsilon\mathbf{u} + C\} \\ &= \inf\{k \geq 0 \mid \boldsymbol{\mu}_j - \boldsymbol{\mu}_i + k\tilde{\Delta}_i^\epsilon \mathbf{u}^* \in A(2\epsilon)\} \\ &= \inf\{k \geq 0 \mid \forall n \in [N]: \mathbf{w}_n^\top(\boldsymbol{\mu}_j - \boldsymbol{\mu}_i + k\tilde{\Delta}_i^\epsilon \mathbf{u}^*) \geq 2\epsilon\alpha_n\} \\ &= \inf\left\{k \geq 0 \mid \forall n \in [N]: k \geq \frac{2\epsilon\alpha_n - \mathbf{w}_n^\top(\boldsymbol{\mu}_j - \boldsymbol{\mu}_i)}{\tilde{\Delta}_i^\epsilon \mathbf{w}_n^\top \mathbf{u}^*}\right\} \\ &= \max_{n \in [N]} \frac{(2\epsilon\alpha_n - \mathbf{w}_n^\top(\boldsymbol{\mu}_j - \boldsymbol{\mu}_i))^+}{\tilde{\Delta}_i^\epsilon \mathbf{w}_n^\top \mathbf{u}^*} \leq \max_{n \in [N]} \frac{2\epsilon\alpha_n + (\mathbf{w}_n^\top(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j))^+}{\tilde{\Delta}_i^\epsilon \mathbf{w}_n^\top \mathbf{u}^*}. \end{aligned}$$

We claim that $\max_{n \in [N]} \frac{(\mathbf{w}_n^\top(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j))^+}{\alpha_n} \leq M(i, j)$. To see this, we observe that, given a number $s \geq 0$, a sufficient condition for $s \leq M(i, j)$ to hold is given by the following:

$$\begin{aligned} &\forall \mathbf{u} \in B(\mathbf{0}, 1) \cap C: \boldsymbol{\mu}_j + s\mathbf{u} \notin \boldsymbol{\mu}_i + \text{int}(C) \\ &\Leftrightarrow \forall \mathbf{u} \in B(\mathbf{0}, 1) \cap C \exists n \in [N]: \mathbf{w}_n^\top(\boldsymbol{\mu}_j + s\mathbf{u}) \leq \mathbf{w}_n^\top \boldsymbol{\mu}_i \\ &\Leftrightarrow \forall \mathbf{u} \in B(\mathbf{0}, 1) \cap C \exists n \in [N]: s \leq \frac{\mathbf{w}_n^\top(\boldsymbol{\mu}_i - \boldsymbol{\mu}_j)}{\mathbf{w}_n^\top \mathbf{u}} \end{aligned}$$

$$\Leftrightarrow \forall \mathbf{u} \in B(\mathbf{0}, 1) \cap C: s \leq \max_{n \in [N]} \frac{(\mathbf{w}_n^\top (\boldsymbol{\mu}_i - \boldsymbol{\mu}_j))^+}{\mathbf{w}_n^\top \mathbf{u}}.$$

Now, taking $s = \max_{n \in [N]} \frac{(\mathbf{w}_n^\top (\boldsymbol{\mu}_i - \boldsymbol{\mu}_j))^+}{\alpha_n}$, the last condition above is verified immediately thanks to the definition of α_n . Hence, the claim holds. Then, we may write

$$\begin{aligned} k_i &\leq \max_{n \in [N]} \frac{2\epsilon\alpha_n + (\mathbf{w}_n^\top (\boldsymbol{\mu}_i - \boldsymbol{\mu}_j))^+}{\tilde{\Delta}_i^\epsilon \mathbf{w}_n^\top \mathbf{u}^*} \\ &\leq \max_{n \in [N]} \frac{2\epsilon\alpha_n + M(i, j)\alpha_n}{\tilde{\Delta}_i^\epsilon \mathbf{w}_n^\top \mathbf{u}^*} \leq \max_{n \in [N]} \frac{2\tilde{\Delta}_i^\epsilon \alpha_n + \tilde{\Delta}_i^\epsilon \alpha_n}{\tilde{\Delta}_i^\epsilon \mathbf{w}_n^\top \mathbf{u}^*} = \max_{n \in [N]} \frac{3\alpha_n}{\mathbf{w}_n^\top \mathbf{u}^*}. \end{aligned}$$

Therefore, choosing $k := \max_{n \in [N]} \frac{3\alpha_n}{\mathbf{w}_n^\top \mathbf{u}^*}$ ensures that $m'(i, j) \geq 2\epsilon$.

Case 3: Suppose that $i \in P \setminus P^*$. Then, there exists $j \in P^*$ such that $\Delta_i^* = m(i, j)$ and we have $m(i, j) \leq \epsilon$ by Definition 3.6. In particular, $\tilde{\Delta}_i^\epsilon = \epsilon$. Let us modify $\boldsymbol{\mu}'$ as $\boldsymbol{\mu}'_i := \boldsymbol{\mu}_i - k\tilde{\Delta}_i^\epsilon \mathbf{u}^*$, where $k > 0$ is to be chosen such that $m'(i, j) \geq 2\epsilon$. Once this is achieved, we have $i \notin P'$ according to Definition 3.6. By the proof of Proposition 3.2(iv), we have $m'(i, j) \geq 2\epsilon$ if and only if $\boldsymbol{\mu}'_i + 2\epsilon \mathbf{u} \in \boldsymbol{\mu}_j - C$ for every $\mathbf{u} \in B(\mathbf{0}, 1) \cap C$. Hence, the minimum value of k is given by

$$\begin{aligned} k_i &:= \inf\{k \geq 0 \mid \forall \mathbf{u} \in B(\mathbf{0}, 1) \cap C: \boldsymbol{\mu}_i - k\tilde{\Delta}_i^\epsilon \mathbf{u}^* + 2\epsilon \mathbf{u} \in \boldsymbol{\mu}_j - C\} \\ &= \inf\{k \geq 0 \mid \forall \mathbf{u} \in B(\mathbf{0}, 1) \cap C: \boldsymbol{\mu}_j - \boldsymbol{\mu}_i + k\tilde{\Delta}_i^\epsilon \mathbf{u}^* \in 2\epsilon \mathbf{u} + C\} \\ &= \inf\{k \geq 0 \mid \boldsymbol{\mu}_j - \boldsymbol{\mu}_i + k\tilde{\Delta}_i^\epsilon \mathbf{u}^* \in A(2\epsilon)\} \\ &= \inf\{k \geq 0 \mid \forall n \in [N]: \mathbf{w}_n^\top (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i + k\tilde{\Delta}_i^\epsilon \mathbf{u}^*) \geq 2\epsilon\alpha_n\} \\ &= \inf\left\{k \geq 0 \mid \forall n \in [N]: k \geq \frac{2\epsilon\alpha_n - \mathbf{w}_n^\top (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i)}{\tilde{\Delta}_i^\epsilon \mathbf{w}_n^\top \mathbf{u}^*}\right\} \\ &= \max_{n \in [N]} \frac{(2\epsilon\alpha_n - \mathbf{w}_n^\top (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i))^+}{\tilde{\Delta}_i^\epsilon \mathbf{w}_n^\top \mathbf{u}^*} \\ &\leq \max_{n \in [N]} \frac{(2\epsilon\alpha_n - m(i, j)\alpha_n)^+}{\tilde{\Delta}_i^\epsilon \mathbf{w}_n^\top \mathbf{u}^*} = \max_{n \in [N]} \frac{2\epsilon\alpha_n - m(i, j)\alpha_n}{\tilde{\Delta}_i^\epsilon \mathbf{w}_n^\top \mathbf{u}^*} \leq \max_{n \in [N]} \frac{2\alpha_n}{\mathbf{w}_n^\top \mathbf{u}^*}. \end{aligned}$$

Hence, choosing $k := \max_{n \in [N]} \frac{2\alpha_n}{\mathbf{w}_n^\top \mathbf{u}^*}$ ensures that $m'(i, j) \geq 2\epsilon$.

Case 4: Suppose that $i \notin P^* \cup P$. Let us modify $\boldsymbol{\mu}'$ as $\boldsymbol{\mu}'_i := \boldsymbol{\mu}_i + k\tilde{\Delta}_i^\epsilon \mathbf{u}^*$, where $k > 0$ is to be chosen such that i is not covered by any $j \in P^*$ in the modified case. For such k , i will automatically be in the Pareto set of the modified case, i.e., $i \in (P^*)'$, hence we must have $i \in P'$ due to Definition 3.6. To determine the desired value of k , note that the largest value of k for which i is covered by $j \in P^*$ in the modified case is given by

$$\begin{aligned} k_{ij} &:= \sup\{k \geq 0 \mid \exists \mathbf{u} \in B(\mathbf{0}, 1) \cap C: \boldsymbol{\mu}'_i \in \boldsymbol{\mu}_j + \epsilon \mathbf{u} - C\} \\ &= \sup_{\mathbf{u} \in B(\mathbf{0}, 1) \cap C} \sup\{k \geq 0 \mid \forall n \in [N]: \mathbf{w}_n^\top \boldsymbol{\mu}_i + k\tilde{\Delta}_i^\epsilon \mathbf{w}_n^\top \mathbf{u}^* \leq \mathbf{w}_n^\top \boldsymbol{\mu}_j + \epsilon \mathbf{w}_n^\top \mathbf{u}\} \\ &= \sup_{\mathbf{u} \in B(\mathbf{0}, 1) \cap C} \sup\left\{k \geq 0 \mid \forall n \in [N]: k \leq \frac{\mathbf{w}_n^\top (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i + \epsilon \mathbf{u})}{\tilde{\Delta}_i^\epsilon \mathbf{w}_n^\top \mathbf{u}^*}\right\} \\ &= \sup_{\mathbf{u} \in B(\mathbf{0}, 1) \cap C} \min_{n \in [N]} \frac{(\mathbf{w}_n^\top (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i + \epsilon \mathbf{u}))^+}{\tilde{\Delta}_i^\epsilon \mathbf{w}_n^\top \mathbf{u}^*} \leq \sup_{\mathbf{u} \in B(\mathbf{0}, 1) \cap C} \min_{n \in [N]} \frac{(\mathbf{w}_n^\top (\boldsymbol{\mu}_j - \boldsymbol{\mu}_i))^+ + \epsilon \mathbf{w}_n^\top \mathbf{u}}{\tilde{\Delta}_i^\epsilon \mathbf{w}_n^\top \mathbf{u}^*} \\ &\leq \sup_{\mathbf{u} \in B(\mathbf{0}, 1) \cap C} \min_{n \in [N]} \frac{m(i, j)\alpha_n + \epsilon \mathbf{w}_n^\top \mathbf{u}}{\tilde{\Delta}_i^\epsilon \mathbf{w}_n^\top \mathbf{u}^*} \leq \sup_{\mathbf{u} \in B(\mathbf{0}, 1) \cap C} \min_{n \in [N]} \frac{\alpha_n + \mathbf{w}_n^\top \mathbf{u}}{\mathbf{w}_n^\top \mathbf{u}^*} \\ &\leq \sup_{\mathbf{u} \in B(\mathbf{0}, 1) \cap C} \max_{n \in [N]} \frac{\alpha_n + \mathbf{w}_n^\top \mathbf{u}}{\mathbf{w}_n^\top \mathbf{u}^*} \leq \max_{n \in [N]} \frac{2\alpha_n}{\mathbf{w}_n^\top \mathbf{u}^*}, \end{aligned}$$

where the passage to the penultimate line follows from Proposition 3.2(iv). Hence, choosing $k := 1 + \max_{n \in [N]} \frac{2\alpha_n}{\mathbf{w}_n^\top \mathbf{u}^*}$ ensures that i is not covered by any $j \in P^*$.

Since we are able to change the returned output in each possibility, as argued in Auer et al. (2016), the algorithm has to sample design i at least $\Omega(\frac{1}{(\tilde{\Delta}_i^\epsilon)^2} \log(\frac{1}{\delta}))$ times in order to detect these changes, which completes the proof.

B.6 Proof of Theorem 5.3

Let $\mathbf{v} \in \mathbb{R}^D$ be such that $\mathbf{v} \notin C \cup (-C)$. Note that $d(\mathbf{v}, (\mathbf{v} + C) \cap C) > 0$. The arguments in the proof of Proposition B.7 show that the function $\mathbf{x} \mapsto d(\mathbf{x}, (\mathbf{x} + C) \cap C)$ is positively homogeneous. Hence, by scaling, we may assume without loss of generality that

$$d(\mathbf{v}, (\mathbf{v} + C) \cap C) = \epsilon.$$

Let $\boldsymbol{\mu} \in C \setminus \{\mathbf{0}\}$. We set the mean reward vectors as

$$\boldsymbol{\mu}_1 = \mathbf{0}, \quad \boldsymbol{\mu}_2 = \boldsymbol{\mu}, \quad \boldsymbol{\mu}_i = \boldsymbol{\mu} - \mathbf{v} \quad \forall i \in [K] \setminus [2].$$

Since $\mathbf{v} \notin C \cup (-C)$, it follows that $P^* = [K] \setminus [1]$.

Let $\mathbf{u}^* \in \text{int}(C)$ be such that

$$\beta_2 = \frac{d(\mathbf{u}^*, (\text{int}(C))^c \cap (\mathbf{u}^* - C))}{d(\mathbf{u}^*, (\text{int}(C))^c)}.$$

It can be checked that, for every $a > 0$, replacing \mathbf{u}^* with $a\mathbf{u}^*$ does not change the value of the above ratio; hence, we assume that $d(\mathbf{u}^*, (\text{int}(C))^c) = 1$ without loss of generality (otherwise, we may take $a = 1/d(\mathbf{u}^*, (\text{int}(C))^c)$). By Lemma B.8, we have

$$d(\mathbf{u}^*, (\text{int}(C))^c \cap (\mathbf{u}^* - C)) = \min_{n \in [N]} \frac{\mathbf{w}_n^\top \mathbf{u}^*}{\alpha_n}. \quad (36)$$

For each $t \in [T]$, we assume that the noise vector $\mathbf{Y}_{i,t}$ is either $\frac{1}{4}\mathbf{u}^*$ with probability $p_i = 1/2$ or it is $-\frac{1}{4}\mathbf{u}^*$ with probability $1 - p_i$, which implies that $\mathbb{E}[\mathbf{Y}_{i,t}] = \mathbf{0}$. Moreover, due to the structure of this distribution, the difference $\hat{\boldsymbol{\mu}}_i - \boldsymbol{\mu}_i$ between the empirical and true means of a design i has a discrete distribution over a finite subset of the line segment that connects $-\frac{1}{4}\mathbf{u}^*$ and $\frac{1}{4}\mathbf{u}^*$, hence we may write $\hat{\boldsymbol{\mu}}_i - \boldsymbol{\mu}_i = (\frac{\hat{p}_i}{2} - \frac{1}{4})\mathbf{u}^*$, where \hat{p}_i is the empirical frequency of observing $\mathbf{Y}_{i,t} = \frac{1}{4}\mathbf{u}^*$ over all rounds t at which design i is sampled.

We proceed by choosing an arbitrary design $i \in [K] \setminus [2]$. Under the original problem $(\boldsymbol{\mu}_j)_{j \in [K]}$, we have $i \in P^*$. Then, we define two modified problems $(\boldsymbol{\mu}'_j)_{j \in [K]}$ and $(\boldsymbol{\mu}''_j)_{j \in [K]}$, where $\boldsymbol{\mu}_i$ is replaced by $\boldsymbol{\mu}'_i$ and $\boldsymbol{\mu}''_i$ respectively, while other designs' mean reward vectors remain unchanged, i.e., $\boldsymbol{\mu}_j = \boldsymbol{\mu}'_j = \boldsymbol{\mu}''_j$ for each $j \neq i$. The new mean reward vectors will be formed by only changing the parameter of noise from p_i to p'_i and p''_i whose values will be specified later. We will carefully set the noise direction to make the Pareto set identification problem difficult. Then, we will argue that any $i \notin P$ should be in P' , the returned set of the modified problem $(\boldsymbol{\mu}'_j)_{j \in [K]}$. We will also argue that any $i \in P$ should not be in P'' , the returned set of the modified problem $(\boldsymbol{\mu}''_j)_{j \in [K]}$. Making an analogy with the coin bias problem as in Auer et al. (2016), we will conclude that design i needs to be sampled $\Omega(\frac{\beta_2^2}{\epsilon^2} \log \frac{1}{\delta})$ times in order to distinguish these cases with at least $1 - \delta$ probability. Since the number of such designs is $K - 2$, this yields $\Omega(\frac{K\beta_2^2}{\epsilon^2} \log \frac{1}{\delta})$ lower bound on the regret.

Case 1: Fix some $i > 2$ and suppose that $i \notin P$. Then, Remark 4.7 implies that there exist $j \in P$ and $\mathbf{u} \in B(\mathbf{0}, 1) \cap C$ such that $\boldsymbol{\mu}_j + \epsilon\mathbf{u} \in \boldsymbol{\mu}_i + C$. Let $J(i) \subseteq P$ be the set of all such j . By the definition of $M(i, j)$, we have $0 < M(i, j) \leq \epsilon$ for every $j \in J(i)$. In particular, $\tilde{\Delta}_i^\epsilon = \epsilon$.

Let us consider the modification $\boldsymbol{\mu}'_i = \boldsymbol{\mu}_i + k\epsilon\mathbf{u}^*$, where $k > 0$ is to be chosen such that i is not covered by any $j \in P^*$ in the modified case. For such k , i will automatically be in the Pareto set of the modified case, i.e., $i \in (P^*)'$, hence we must have $i \in P'$ due to Definition 4.6. To determine the desired value of k , note that the largest value of k for which i is covered by $j \in P^*$ in the modified case is given by

$$k_{ij} := \sup\{k \geq 0 \mid \exists \mathbf{u} \in B(\mathbf{0}, 1) \cap C: \boldsymbol{\mu}'_i \in \boldsymbol{\mu}_j + \epsilon\mathbf{u} - C\} \leq \max_{n \in [N]} \frac{2\alpha_n}{\mathbf{w}_n^\top \mathbf{u}^*}.$$

as computed in Case 4 of the proof of Theorem 5.1. Moreover, by (36) and the definition of \mathbf{u}^* , we have

$$\max_{n \in [N]} \frac{2\alpha_n}{\mathbf{w}_n^\top \mathbf{u}^*} = \frac{2}{d(\mathbf{u}^*, (\text{int}(C))^c \cap (\mathbf{u}^* - C))} = \frac{2}{\beta_2} \frac{1}{d(\mathbf{u}^*, (\text{int}(C))^c)} = \frac{2}{\beta_2}. \quad (37)$$

Hence, choosing any k that satisfies

$$k > \frac{2}{\beta_2}$$

does the job. For such k , we automatically have $i \in (P^*)'$, which implies that $i \in P'$ due to Definition 4.6.

Case 2: Fix $i > 2$ and suppose that $i \in P$. Let us consider the modification $\mu_i'' = \mu_i - k\epsilon \mathbf{u}^*$, where $k > 0$ is to be chosen such that $m''(i, 2) \geq 2\epsilon$. Hence, for such k , we have $i \notin P''$ in the modified case thanks to Definition 4.6.

Following a similar derivation as in Case 2 of the proof of Theorem 5.1, the minimum k for which $m''(i, 2) \geq 2\epsilon$ is given by

$$k_i := \inf\{k \geq 0 \mid \forall \mathbf{u} \in B(\mathbf{0}, 1) \cap C: \mu_i - k\epsilon \mathbf{u}^* + 2\epsilon \mathbf{u} \in \mu_2 - C\} \leq \max_{n \in [N]} \frac{2\epsilon \alpha_n + M(i, 2)\alpha_n}{\epsilon \mathbf{w}_n^\top \mathbf{u}^*}.$$

Note that $M(i, 2) = d(\mu_2 - \mu_i, (\mu_2 - \mu_i + C) \cap C) = d(\mathbf{v}, (\mathbf{v} + C) \cap C) = \epsilon$ by construction. Hence,

$$k_i \leq \max_{n \in [N]} \frac{2\epsilon \alpha_n + \epsilon \alpha_n}{\epsilon \mathbf{w}_n^\top \mathbf{u}^*} = \max_{n \in [N]} \frac{3\alpha_n}{\mathbf{w}_n^\top \mathbf{u}^*} = \frac{3}{\beta_2},$$

where the last equality is by (37). Hence, we may simply choose $k = \frac{3}{\beta_2}$ and ensure that $i \notin P''$.

B.7 Proof of Theorem 6.1

Given $i, j \in [K]$, let $\hat{m}(i, j) := d(\hat{\mu}_j - \hat{\mu}_i, (\text{int}(C))^c \cap (\hat{\mu}_j - \hat{\mu}_i - C))$ and $\hat{M}(i, j) := d(\hat{\mu}_j - \hat{\mu}_i, C \cap (\hat{\mu}_j - \hat{\mu}_i + C))$ represent the empirical estimates of $m(i, j)$ and $M(i, j)$, respectively. In the next lemma, we prove that the following conditions in terms of the gaps are sufficient for (ϵ, δ) -PAC Pareto set identification.

Condition 1. For every $i \in P^*$ and $j \in [K] \setminus \{i\}$, $M(i, j) > \epsilon$ implies that $\hat{M}(i, j) > 0$.

Condition 2. For every $i \notin P^*$ and $j \in P^*$, $m(i, j) > \epsilon$ implies that $\hat{m}(i, j) > 0$.

Lemma B.10. If Conditions 1 and 2 hold, then P is an (ϵ, δ) -PAC Pareto set.

Proof. Assume that Condition 1 holds. To get a contradiction, suppose that property (i) in Definition 4.6 does not hold. By Remark 4.7, this implies that there exists $i \in P^*$ such that $\mu_i \notin \mu_j + (B(\mathbf{0}, \epsilon) \cap C) - C$ for each $j \in P$. Hence, $\mu_j - \mu_i \notin (B(\mathbf{0}, \epsilon) \cap (-C)) + C$ for each $j \in P$. Next, we will show that this implies $M(i, j) = d(\mu_j - \mu_i, C \cap (\mu_j - \mu_i + C)) > \epsilon$ for each $j \in P$.

The result follows from a more general argument. Consider $\mathbf{x} \in \mathbb{R}^D$. Let $\mathbf{x}^* \in C \cap (\mathbf{x} + C)$ be such that $d(\mathbf{x}, C \cap (\mathbf{x} + C)) = d(\mathbf{x}, \mathbf{x}^*)$. Note that $\mathbf{x}^* \in C \cap (\mathbf{x} + C)$ implies the existence of unit vectors $\mathbf{u}_1, \mathbf{u}_2 \in C$ and scalars $c_1, c_2 \in \mathbb{R}_+$ such that $\mathbf{x}^* = \mathbf{x} + c_1 \mathbf{u}_1 = c_2 \mathbf{u}_2$. Hence, we have $\mathbf{x} = c_2 \mathbf{u}_2 - c_1 \mathbf{u}_1$. We claim that $d(\mathbf{x}, \mathbf{x}^*) \leq \epsilon$ implies that $\mathbf{x} \in (B(\mathbf{0}, \epsilon) \cap (-C)) + C$. Now assume that $d(\mathbf{x}, \mathbf{x}^*) \leq \epsilon$. Note that $d(\mathbf{x}, \mathbf{x}^*) = c_1 \leq \epsilon$. It immediately follows that $\mathbf{x} = c_2 \mathbf{u}_2 - c_1 \mathbf{u}_1 \in (B(\mathbf{0}, \epsilon) \cap (-C)) + C$, since $-c_1 \mathbf{u}_1 \in (B(\mathbf{0}, \epsilon) \cap (-C))$ and $c_2 \mathbf{u}_2 \in C$. This proves our claim. Therefore, $\mathbf{x} \notin (B(\mathbf{0}, \epsilon) \cap (-C)) + C$ implies that $d(\mathbf{x}, C \cap (\mathbf{x} + C)) = d(\mathbf{x}, \mathbf{x}^*) > \epsilon$. Plugging $\mathbf{x} = \mu_j - \mu_i$ to the argument above gives the desired result.

Then, Condition 1 implies $\hat{M}(i, j) > 0$ for each $j \in P$. By Corollary 4.5(iii) applied to the random partial orders $\hat{\succ}_C, \hat{\succeq}_C$, we get $i \not\hat{\succeq}_C j$ for each $j \in P$. Hence, we must have $i \in P$. In particular, taking $j = i$ gives $\mu_i \notin \mu_i + (B(\mathbf{0}, \epsilon) \cap C) - C$, a contradiction. Hence, property (i) holds.

Assume that Condition 2 holds. Let $i \in P \setminus P^*$. Then, there is no $j \in [K] \setminus \{i\}$ such that $i \hat{\succeq}_C j$. By Corollary 4.5(iii) applied to the random orders $\hat{\succ}_C, \hat{\succeq}_C$, we have $\hat{m}(i, j) = 0$ for every $j \in [K] \setminus \{i\}$, hence, for every $j \in P^*$. Then, Condition 2 implies that $m(i, j) \leq \epsilon$ for every $j \in P^*$. Hence, $\Delta_i^* \leq \epsilon$, i.e., property (ii) in Definition 4.6 holds. \square

For each $i, j \in [K]$ with $i \neq j$, let us define $\Delta_{ij} = \mu_j - \mu_i$, $\hat{\Delta}_{ij} = \hat{\mu}_j - \hat{\mu}_i$. The next lemma shows that the following conditions in terms of the deviation of $\hat{\Delta}_{ij}$ from Δ_{ij} are sufficient for (ϵ, δ) -PAC Pareto set identification.

Condition a. For every $i \in P^*$ and $j \in [K] \setminus \{i\}$, $d(\Delta_{ij}, C \cap (\Delta_{ij} + C)) > \epsilon$ implies that $\|\hat{\Delta}_{ij} - \Delta_{ij}\|_2 < d(\Delta_{ij}, C)$.

Condition b. For every $i \notin P^*$ and $j \in P^*$, $d(\Delta_{ij}, (\text{int}(C))^c \cap (\Delta_{ij} - C)) > \epsilon$ implies that $\|\hat{\Delta}_{ij} - \Delta_{ij}\|_2 < d(\Delta_{ij}, (\text{int}(C))^c)$.

Lemma B.11. *If Conditions a and b hold, then P satisfies the success condition in Definition 4.6.*

Proof. Assume that Condition a holds. We verify Condition 1. Let $i \in P^*$ and $j \in [K] \setminus \{i\}$ such that $M(i, j) > \epsilon$, that is, $d(\Delta_{ij}, C \cap (\Delta_{ij} + C)) > \epsilon$ by Proposition 4.3(ii). Then, Condition a implies that $\|\hat{\Delta}_{ij} - \Delta_{ij}\|_2 < d(\Delta_{ij}, C)$. Since $d(\Delta_{ij}, C \cap (\Delta_{ij} + C)) > 0$ and $C \cap (\Delta_{ij} + C)$ is a closed set, we have $\Delta_{ij} \notin C \cap (\Delta_{ij} + C)$. Hence, $\Delta_{ij} \notin C$. We claim that $\hat{\Delta}_{ij} \notin C$. To get a contradiction, suppose that $\hat{\Delta}_{ij} \in C$. Then, we have $\|\hat{\Delta}_{ij} - \Delta_{ij}\|_2 \geq \inf_{c \in C} \|\Delta_{ij} - c\|_2 = d(\Delta_{ij}, C)$, which is a contradiction. Hence, the claim holds and we have $\hat{\Delta}_{ij} \notin C \cap (\hat{\Delta}_{ij} + C)$. By Proposition 4.3(ii), we obtain $\hat{M}(i, j) = d(\hat{\Delta}_{ij}, C \cap (\hat{\Delta}_{ij} + C)) > 0$. Hence, Condition 1 holds.

Assume that Condition b holds. We verify Condition 2. Let $i \notin P^*$ and $j \in P^*$ such that $m(i, j) > \epsilon$, that is, $d(\Delta_{ij}, (\text{int}(C))^c \cap (\Delta_{ij} - C)) > \epsilon$ by Proposition 4.2(ii). Then, Condition b implies that $\|\hat{\Delta}_{ij} - \Delta_{ij}\|_2 < d(\Delta_{ij}, (\text{int}(C))^c)$. Moreover, since $(\text{int}(C))^c \cap (\Delta_{ij} - C)$ is a closed set, we have $\Delta_{ij} \notin (\text{int}(C))^c \cap (\Delta_{ij} - C)$ and hence $\Delta_{ij} \in \text{int}(C)$. We claim that $\hat{\Delta}_{ij} \in \text{int}(C)$. Supposing otherwise that $\hat{\Delta}_{ij} \in (\text{int}(C))^c$, we get $\|\hat{\Delta}_{ij} - \Delta_{ij}\|_2 \geq \inf_{c \in (\text{int}(C))^c} \|\Delta_{ij} - c\|_2 = d(\Delta_{ij}, (\text{int}(C))^c)$, which is a contradiction. Hence, the claim holds and we have $\hat{\Delta}_{ij} \notin (\text{int}(C))^c \cap (\hat{\Delta}_{ij} - C)$. By Proposition 4.2(ii), we get $\hat{m}(i, j) = d(\hat{\Delta}_{ij}, C \cap (\hat{\Delta}_{ij} + C)) > 0$. Hence, Condition 2 holds. \square

Let $i, j \in [K]$ with $i \neq j$. The next lemma explains how $\hat{\Delta}_{ij}$ concentrates around Δ_{ij} in ℓ_2 norm as a function of the exploration parameter L . Let us introduce the constant

$$\theta_{ij} := \begin{cases} \frac{d(\Delta_{ij}, C)}{d(\Delta_{ij}, C \cap (\Delta_{ij} + C))} & \text{if } \Delta_{ij} \notin C, \\ \frac{d(\Delta_{ij}, (\text{int}(C))^c)}{d(\Delta_{ij}, (\text{int}(C))^c \cap (\Delta_{ij} - C))} & \text{if } \Delta_{ij} \in \text{int}(C), \\ 1 & \text{if } \Delta_{ij} \in \text{bd}(C). \end{cases}$$

Moreover, for a given (ϵ, δ) , we define the per-design sampling budget as

$$g(\epsilon, \delta) := \left\lceil \frac{4\beta^2 c^2 \sigma^2}{\epsilon^2} \log\left(\frac{4D}{\delta}\right) \right\rceil.$$

Lemma B.12. *Suppose that $L = g(\epsilon, \delta)$ noisy observations are evaluated for each design, where $c > 0$ is a constant. Then, there exists a choice of $c > 0$ (free of all problem parameters) such that, for each $i, j \in [K]$ with $i \neq j$, we have $\|\hat{\Delta}_{ij} - \Delta_{ij}\|_2 \leq \epsilon \theta_{ij}$ with probability at least $1 - \delta$.*

Proof. By Jin et al. 2019, Corollary 7, there exists an absolute constant $c > 0$ such that, with probability at least $1 - \delta/2$, we have

$$\begin{aligned} L \|\hat{\mu}_i - \mu_i\|_2 &= L \left\| \frac{\sum_{t=1}^{LK} \mathbf{X}_t \mathbb{I}(I_t = i)}{L} - \mu_i \right\|_2 \\ &= \left\| \sum_{t=1}^{KL} (\mathbf{X}_t - \mu_i) \mathbb{I}(I_t = i) \right\|_2 \leq c \sqrt{L \sigma^2 \log\left(\frac{4D}{\delta}\right)}, \end{aligned}$$

that is, $\|\hat{\mu}_i - \mu_i\|_2 \leq c \sqrt{\frac{\sigma^2}{L} \log\left(\frac{4D}{\delta}\right)}$. Setting $L = \lceil (4\beta^2 c^2 \sigma^2 / \epsilon^2) \log(4D/\delta) \rceil$ ensures that $\|\hat{\mu}_i - \mu_i\|_2 \leq \epsilon / (2\beta)$ with probability at least $1 - \delta/2$. Applying a union bound, we obtain $\mathbb{P}\{\|\hat{\mu}_i - \mu_i\|_2 \leq \frac{\epsilon}{2\beta}, \|\hat{\mu}_j - \mu_j\|_2 \leq \frac{\epsilon}{2\beta}\} \geq 1 - \delta$. Noting that $\|\hat{\Delta}_{ij} - \Delta_{ij}\|_2 = \|\hat{\mu}_j - \mu_j + \mu_i - \hat{\mu}_i\|_2 \leq \|\hat{\mu}_i - \mu_i\|_2 + \|\hat{\mu}_j - \mu_j\|_2$, we have $\mathbb{P}\{\|\hat{\Delta}_{ij} - \Delta_{ij}\|_2 \leq \epsilon/\beta\} \geq 1 - \delta$. If $\Delta_{ij} \notin C$, then

$$\frac{\epsilon}{\beta} \leq \frac{\epsilon}{\beta_1} \leq \frac{\epsilon d(\Delta_{ij}, C)}{d(\Delta_{ij}, C \cap (\Delta_{ij} + C))} = \epsilon \theta_{ij};$$

if $\Delta_{ij} \in \text{int}(C)$, then

$$\frac{\epsilon}{\beta} \leq \frac{\epsilon}{\beta_2} \leq \frac{\epsilon d(\Delta_{ij}, (\text{int}(C))^c)}{d(\Delta_{ij}, (\text{int}(C))^c \cap (\Delta_{ij} - C))};$$

and if $\Delta_{ij} \in \text{bd}(C)$, then $\epsilon/\beta \leq \epsilon = \epsilon \theta_{ij}$; see (4) for the definitions of β_1, β_2 . Hence, the result follows. \square

After defining the lemmas that will be used in the proof, we complete the proof of Theorem 6.1 by the argument below.

With the given choice of L , by Lemma B.12, for every $i, j \in [K]$ with $i \neq j$, we have $\|\hat{\Delta}_{ij} - \Delta_{ij}\|_2 \leq \epsilon\theta_{ij}$ with probability at least $1 - 2\delta/(K(K-1))$. The application of union bound shows that with probability at least $1 - \delta$, $\|\hat{\Delta}_{ij} - \Delta_{ij}\|_2 \leq \epsilon\theta_{ij}$ simultaneously for all $i, j \in [K]$ such that $i \neq j$. Under this event, we verify Conditions a and b. Let $i \in P^*$ and $j \in [K] \setminus \{i\}$ such that $d(\Delta_{ij}, C \cap (\Delta_{ij} + C)) > \epsilon$. In particular, $\Delta_{ij} \notin C$. Hence, we have

$$\|\hat{\Delta}_{ij} - \Delta_{ij}\|_2 \leq \epsilon\theta_{ij} = \frac{\epsilon d(\Delta_{ij}, C)}{d(\Delta_{ij}, C \cap (\Delta_{ij} + C))} < d(\Delta_{ij}, C).$$

This shows that Condition a holds. Let $i \notin P^*$ and $j \in P^*$ such that $d(\Delta_{ij}, (\text{int}(C))^c \cap (\Delta_{ij} - C)) > \epsilon$. In particular, $\Delta_{ij} \in \text{int}(C)$. Hence, we have

$$\|\hat{\Delta}_{ij} - \Delta_{ij}\|_2 \leq \epsilon\theta_{ij} = \frac{\epsilon d(\Delta_{ij}, (\text{int}(C))^c)}{d(\Delta_{ij}, (\text{int}(C))^c \cap (\Delta_{ij} - C))} < d(\Delta_{ij}, (\text{int}(C))^c).$$

This shows that Condition b holds. Therefore, by Lemma B.11, P is an (ϵ, δ) -PAC Pareto set.

When $D = 1$, the sample complexity of naïve elimination will match the one given in Even-Dar et al. (2006). In order to match, the analysis in Lemma B.10 and Theorem 6.1 needs to be updated. In particular, we no longer need concentration of $\hat{\Delta}_{ij}$ s. Showing concentration of $\hat{\mu}_j$ s will suffice.

C SUPPLEMENTAL NUMERICAL RESULTS

The following table contains additional numerical results about the experiments performed in Section 7. Code for the paper is included as non-textual supplementary material (ZIP file).

Table 3: Additional results for experiments conducted in Section 7. θ : cone angle in degrees. SR_θ : success rate (in %) for C_θ . NF1_θ : average number of designs in P_θ^* that fail success condition (i) in Definition 4.6. NF2_θ : average number of designs in $P \setminus P_\theta^*$ that fail success condition (ii) in Definition 4.6. PM_θ : $(|P_\theta^* \setminus P|/|P_\theta^*|) \times 100$.

L	ϵ	SR_{45}	SR_{90}	SR_{135}	NF1_{45}	NF1_{90}	NF1_{135}	NF2_{45}	NF2_{90}	NF2_{135}	PM_{45}	PM_{90}	PM_{135}
10^2	10^{-3}	0	0	0	12.44	7.61	1.83	9.26	7.07	2.52	24.8	29.3	18.3
	10^{-2}	0	0	0	11.52	6.87	1.83	7.44	6.08	2.52	24.8	29.3	18.3
	10^{-1}	1	7	27	4.22	1.21	0.63	1.27	0.98	0.83	24.8	29.3	18.3
10^3	10^{-3}	0	0	29	4.45	3.33	0.32	5.78	3.72	0.89	9.5	12.8	3.2
	10^{-2}	0	0	29	3.68	2.73	0.32	3.57	2.40	0.89	9.5	12.8	3.2
	10^{-1}	78	99	100	0.18	0.00	0.00	0.06	0.01	0.00	9.5	12.8	3.2
10^4	10^{-3}	1	3	85	1.33	1.06	0.00	2.54	1.88	0.15	3.4	4.1	0
	10^{-2}	22	24	85	0.63	0.77	0.00	0.87	0.57	0.15	3.4	4.1	0
	10^{-1}	100	100	100	0.00	0.00	0.00	0.00	0.00	0.00	3.4	4.1	0
10^5	10^{-3}	17	55	100	0.24	0.06	0.00	1.10	0.44	0.00	1.2	0.2	0
	10^{-2}	100	99	100	0.00	0.00	0.00	0.00	0.01	0.00	1.2	0.2	0
	10^{-1}	100	100	100	0.00	0.00	0.00	0.00	0.00	0.00	1.2	0.2	0

References

- A. A. Goldstein. Optimization of Lipschitz continuous functions. *Mathematical Programming*, 13: 14–22, 1977.
- R. T. Rockafellar. *Convex Analysis*. Princeton University Press, 1970.