# Electronic Companion: Appendix for Contextual Inverse Optimization: Offline and Online Learning

Omar Besbes[5], Yuri Fonseca[6], and Ilan Lobel[7]

## A Examples of problem classes that fall under our formulation

Next, to illustrate the framework introduced in Section 2, we provide a few prototypical classes of problems that fall under it and relate these more precisely to some existing papers discussed in Section 1.2.

A first special case is one in which the decision-maker faces a sequence of optimization problems of the form $\min_{x\in\mathcal{X}_t} x'c^\star$. This formulation was studied in Bärmann et al. (2017). A natural extension of this problem, also studied in Bärmann et al. (2017), is to consider linear context functions $f_t$. In this case, the sequence of optimization problems becomes solving $\min_{x\in\mathcal{X}_t} x'Z_tc^\star$ for the context function $f_t(x) = x'Z_t$.

As a concrete application of the above, consider a problem where the goal is to learn preferences from observing consumers' behavior. At each time $t$, the consumer faces an arbitrary bundle $S_t$ of $J_t$ products, each with features $z_t^j \in \mathbb{R}^d$ and price $p_t^j$. The consumer chooses products in the assortment in order to maximize his/her utility function subject to a budget constraint of $b_t$. Let $Z_t \in \mathbb{R}^{J_t} \times \mathbb{R}^d$ be the matrix where each row $j$ is given by the vector of features of the product $j$, i.e., $z_t^j$. Let also $p_t = (p_t^1, ..., p_t^{J_t})$ be the vector of prices of products $j = 1, ..., J_t$ available at time $t$. At each time, the consumer solves

$$x_t^\star = \arg\min_{x\in\mathcal{X}_t} f_t(x)'c^\star, \qquad \text{where } \mathcal{X}_t = \{x \in \{0,1\}^{J_t} \ : \ x'p_t \le b_t\},\ f_t(x) = -x'Z_t,$$

which is a sequence of adversarial knapsack problems where the universe of products, the budget constraint and the prices are allowed to change and are arbitrarily selected (by nature). This relates to the formulation explored in the learning from revealed preferences literature.

When there is no budget and $|S_t| = 1$, the problem becomes one of sequential customer buy/no buy decisions. The question then is how to leverage such binary feedback. The consumer is rational and buys the product if and only if the utility of buying it, modeled as $(z_t'\theta)$ satisfies $z_t'\theta \ge p_t$. Here the vector $\theta$ is unknown. The consumer contextual optimization problem is given by $x^\star = \arg\max_{x\in\{0,1\}}(z_t'\theta - p_t)x$. In our notation, this would correspond to context function $f_t(x) = -x(z_t, p_t)$, feasible sets $\mathcal{X} = \{0, 1\}$ and cost vector $c^\star = (\theta, -1)$. Even though the structure of the feedback and the contextual optimization problem faced by the consumer would be the same as in the contextual pricing literature (see, e.g., Cohen et al. (2020)), the problem described above is of different nature since the decision-maker is not allowed to select the price of the product and can't affect the feedback directly. Here, the decision-maker's action

[5]Columbia Business School — ob2015@gsb.columbia.edu.
[6]Columbia Business School — yfonseca23@gsb.columbia.edu.
[7]NYU Stern School of Business — ilobel@stern.nyu.edu.

is a passive one, where instead of setting the price, the decision-maker merely "guesses" if the product would be bought or not by the consumer. This endows the decision-maker with significantly less control on the information collected.

The flexibility of our contextual optimization formulation allows to encompass works in imitation learning problems. In Ward et al. (2019), the authors consider an online problem where the goal is to be able to mimic the optimal scheduling policy of an expert for the following dynamic problem. At each time $t$, the length (states) of $n$ queues $z_t \in \mathbb{R}^n$ is observed and the expert solves $x_t^\star = \arg\max_{x \in \mathcal{X}} x' B z_t$, where $\mathcal{X}$ is the set of admissible schedule configurations that is assumed to be bounded. The matrix $B$ is unknown to the decision-maker. This problem falls under our formulation by defining $f_t$ to be the Kronecker product between $x$ and $z_t$, $f_t(x) = x \otimes z_t = (x_{[1]} z_t, x_{[2]} z_t, \cdots, x_{[n]} z_t)$, and letting $\mathrm{vec}(B)$ be the operator that stacks the columns of $B$. We then have that at each time $t$ the expert solves

$$x_t^\star = \arg\min_{x \in \mathcal{X}} f_t(x)' c^\star, \qquad \text{where } f_t(x) = x \otimes z_t, \text{ and } c^\star = -\mathrm{vec}(B).$$

Next, we illustrate that the framework also captures versions of problems in inverse reinforcement learning. A version of the problem studied in Ratliff et al. (2006) can be described as follows. At each period $t$, the decision-maker faces a new Markov decision process (MDP) and the objective is to match the state-action frequency of the expert. The initial state distribution and the transition probabilities for the period $t$ MDP are known. We let the set of feasible actions for the decision-maker $\mathcal{X}_t$ to be the set of feasible state-action pair frequencies for the MDP. At the end of the period (after a full run of the period $t$ MDP), the decision-maker observes the optimal state-action pair frequency for that period's MDP. At period $t$, let $\mathcal{S}_t$ and $\mathcal{A}_t$ denote the spaces for state and action, respectively. For each state-action pair $(s, a) \in \mathcal{S}_t \times \mathcal{A}_t$, we have an associated $d$-dimensional vector of features $\phi_t(s, a)$. We assume that the cost function $r(s, a)$ associated with taking action $a$ in state $s$ is a linear with respect to the vector of features, i.e., $r(s, a) = \phi_t(s, a)' c^\star$. We also denote $r \in \mathbb{R}^{|\mathcal{S}_t| \times \mathcal{A}_t}$ to be the cost vector for each state-action pair. Let $\Phi_t \in \mathbb{R}^{|\mathcal{S}_t| \times \mathcal{A}_t} \times \mathbb{R}^d$ be the matrix where each row is the feature vector associated with each of the state-action pairs. Using the dual of the LP formulation for the MDP (cf. Sutton and McCallum (2006)), we have that $V_\pi = x_\pi' r = x_\pi' \Phi_t c^\star$, where $x_\pi$ is the vector of state-action pair frequency implied by policy $\pi$. Therefore, at each time $t$ we would like to solve

$$x_t^\star = \arg\min_{x \in \mathcal{X}_t} f_t(x)' c^\star, \qquad \text{where } f_t(x) = x' \Phi_t.$$

Finally, recall that a structured prediction problem is one where we observe some input and we would like to predict a multidimensional output. Let the inputs be denoted by $z \in \mathcal{Z}$, the outputs by $x \in \mathcal{X}(z)$ and the score function by $g(z, x)$. Our problem formulation includes special cases of this class of problems with a regret objective; when it is possible to assume a parametric form for $g(x, z) = -f(x, z)' c^\star$ for a known $f(x,z)$, then the forward problem is given by $x^\star = \arg\min_{x\in\mathcal{X}(z)} f(x,z)'c^\star$ for an unknown $c^\star$. Applications of structured predictions include, e.g., natural language processing and image recognition (Taskar et al., 2005).

The examples above illustrate the generality of the formulation that we study, but also that there is a potential to lift up some existing formulations within a general framework of contextual inverse optimization with regret objective.

# B Appendix: Proofs

***Proof of Lemma 1.*** We first consider the cases where $\theta(c^\star, c^\pi) \geq \pi/2$. In this case, nature can choose $f$ as the identity function and a set $\mathcal{X} = \{x_1, x_2\}$ such that $x_2 - x_1 = c^\star$. Note that $(x_2-x_1)'c^\star = \|c^\star\|^2 = 1 \geq 0$, and therefore $x_2'c^\star \geq x_1'c^\star$, implying that $x_1 \in \psi(c^\star, \mathcal{X}, f)$. We also have that $(x_2-x_1)'c^\pi = c^{\star\prime}c^\pi \leq 0$, where the inequality follows from $\theta(c^\star, c^\pi) \geq \pi/2$. Thus, $x_1'c^\pi \geq x_2'c^\pi$, implying that $x_2 \in \psi(c^\pi, \mathcal{X}, f)$. Therefore, $\mathcal{L}(c^\pi, c^\star) \geq (x_2 - x_1)'c^\star = \|c^\star\|^2 = 1$.

We now turn our attention to the case where $0 \leq \theta(c^\star, c^\pi) < \pi/2$. The proof is organized as follows. We first introduce a relaxation of the maximization problem in the definition of $\mathcal{L}$ (see Eq. (8)). Second, we show that the relaxation leads to a semi-definite programming (SDP) formulation based on the realizability of Gram matrices. Third, we derive an upper bound for the relaxed problem. The final step is to construct an instance that attains this upper bound.

**Step 1.** Recall from Eq. (8) that $\mathcal{L}(c^\pi, c^\star) = \sup_{\mathcal{X}\in\mathcal{B},\, f\in\mathcal{F}, x^\pi\in\psi(c^\pi,\mathcal{X},f)} \Big(f(x^\pi) - f(x^\star)\Big)' c^\star$. For any $\mathcal{X} \in \mathcal{B}$, $f \in \mathcal{F}$, $c^\star, c^\pi \in S^d$, $x^\star \in \psi(c^\star, \mathcal{X}, f)$, and $x^\pi \in \psi(c^\pi, \mathcal{X}, f)$, we have, by the optimality of $x^\star$ and $x^\pi$ (for their respective problems), that $(f(x^\pi) - f(x^\star))'c^\star \geq 0$ and $(f(x^\pi) - f(x^\star))'c^\pi \leq 0$.

In particular, we have the following

$$
\begin{aligned}
\mathcal{L}(c^\pi, c^\star) &= \sup\left\{ (f(x^\pi) - f(x^\star))'c^\star \;:\; \mathcal{X}\in\mathcal{B},\ f\in\mathcal{F}, x^\pi\in\psi(c^\pi,\mathcal{X},f),\ x^\star\in\psi(c^\star,\mathcal{X},f) \right\} \\
&\leq \sup\left\{ \delta' c^\star \;:\; \delta\in\mathbb{R}^d,\ \|\delta\|\leq 1,\ \delta' c^\pi \leq 0 \right\} \\
&= \sup\left\{ \delta' c^\star \;:\; \delta\in\mathbb{R}^d,\ \|\delta\| = 1,\ \delta' c^\pi \leq 0 \right\},
\end{aligned}
$$

where the inequality follows from noting that for any feasible $x^\star$, $x^\pi$, if one sets $\delta = f(x^\star) - f(x^\pi)$, we must have that $\|\delta\| \leq \|x^\star - x^\pi\| \leq 1$, and $\delta'c^\pi \leq 0$. The last equality follows from the fact that an optimal solution will always have a vector $\delta$ with maximal norm. Note that the set of $\delta \in \mathbb{R}^d$ such that $\|\delta\| = 1$ is the set $S^d$. We therefore switch our attention to the optimization problem

$$
\sup\left\{ \delta' c^\star \;:\; \delta\in S^d,\ \delta' c^\pi \leq 0 \right\}, \tag{B-1}
$$

which is an upper bound on the value of $\mathcal{L}(c^\pi, c^\star)$.

**Step 2.** We now analyze problem (B-1). In particular, we show how it can be written as an SDP and solved explicitly. For any $c^\star$, $c^\pi$ and $\delta \in S^d$, let us define the following matrices

$$A = \begin{bmatrix} \delta & c^\star & c^\pi \end{bmatrix}, \quad \text{and} \quad B = A'A.$$

Then, $B$ is equal to

$$B = \begin{bmatrix} 1 & \delta' c^\star & \delta' c^\pi \\ \delta' c^\star & 1 & c' c^\pi \\ \delta' c^\pi & c^{\star\prime} c^\pi & 1 \end{bmatrix}.$$

Note that $B$ is the Gram matrix associated with the matrix $A$. Therefore, it must belong to the set of symmetric positive semi-definite matrices (see Chapter 8.3.1 of Boyd and Vandenberghe (2004)). Moreover, Cauchy-Schwartz implies that all its entries are in $[-1, 1]$. We also have that for any feasible solution of Problem (B-1), the entry $B_{1,3} = \delta' c^\pi \leq 0$ and thus $B_{1,3} \in [-1, 0]$.

Let $\mathcal{S}^3_+$ denote the set of symmetric positive semi-definite matrices in $\mathbb{R}^{3\times 3}$ and let $\mathcal{M}^3$ denote the set of symmetric matrices in $\mathbb{R}^{3\times 3}$. In turn, we have the following relaxation:

$$\begin{aligned}
&\sup\left\{ \delta' c^\star \ : \ \delta \in S^d, \delta' c^\pi \leq 0 \right\} \\
&\quad \leq \ \sup\left\{ B_{1,2} \ : \ B \in \mathcal{S}^3_+, \quad B_{1,2} \in [-1,1], \quad B_{1,3} \in [-1,0], \quad B_{2,3} = c^{\star\prime} c^\pi \right\} \\
&\quad \leq \ \sup\{ B_{1,2} \ : \ B \in \mathcal{M}^3, \quad \det(B) \geq 0, \quad B_{1,2} \in [-1,1], \\
&\qquad\qquad B_{1,3} \in [-1,0], \quad B_{2,3} = c^{\star\prime} c^\pi \} \\
&\quad = \ \sup\{ B_{1,2} \ : \ B \in \mathcal{M}^3, \quad 1 + 2B_{1,2}B_{1,3}B_{2,3} - B_{1,2}^2 - B_{1,3}^2 - B_{2,3}^2 \geq 0, \\
&\qquad\qquad B_{1,2} \in [-1,1], \quad B_{1,3} \in [-1,0], \quad B_{2,3} = c^{\star\prime} c^\pi \} \\
&\quad = \ \sup\{ r \ : \ 1 + 2\rho z r - z^2 - \rho^2 - r^2 \geq 0, \\
&\qquad\qquad -1 \leq \rho \leq 0, \quad -1 \leq r \leq 1, \quad z = c^{\star\prime} c^\pi \},
\end{aligned} \tag{B-2}$$

where the last inequality follows from the relaxation to symmetric matrices with non-negative determinants. To simplify notation, we denoted $r = \delta' c^\star, \rho = \delta' c^\pi$ and $z = c^{\star\prime} c^\pi$.

**Step 3.** In the final step we upper bound the problem in Eq. (B-2). Let $h(r) = 1 + 2\rho z r - z^2 - \rho^2 - r^2$ and note that $h$ is a quadratic function for any $\rho, z$. Also, $h$ is concave and admits two roots

$$\begin{aligned}
r_+ &= z\rho + \sqrt{z^2\rho^2 + (1 - z^2 - \rho^2)}, \\
r_- &= z\rho - \sqrt{z^2\rho^2 + (1 - z^2 - \rho^2)},
\end{aligned}$$

such that $h(r) \geq 0$ if and only if $r \in [r_-, r_+]$. Hence, for any feasible $\rho, z$, the maximal achievable value of $r$ is $r_+$. Solving the problem in Eq. (B-2) reduces to finding the feasible values of $\rho$ and $z$ that maximize

$r_+$. When $\theta(c^\star, c_t^\pi) = 0$, $z = 1$, and $r_+ = 0$, so there is no regret. Now we consider the case where $0 < \theta(c^\star, c_t^\pi) < \pi/2$. Note that $r_+$ is differentiable with respect to $\rho$ and its derivative is given by

$$\frac{\partial r_+}{\partial \rho} = z + 2\rho(z^2 - 1)\frac{1}{2\sqrt{z^2\rho^2 + (1 - z^2 - \rho^2)}}.$$

Since $0 < \theta(c^\star, c^\pi)$ implies $z < 1$, the term inside the square root is always greater than zero and the derivative is always well-defined on the feasible set. Moreover, since $\theta(c^\star, c^\pi) < \pi/2$ implies $z > 0$ and $\rho \le 0$ implies $\rho(z^2 - 1) \ge 0$, we get that the derivative is non-negative on the feasible set. Hence, independently of the value of $z$, the value of $r_+$ is maximized on the feasible set when $\rho$ achieves its maximum value, 0.

Note that for all $\rho \le 0$, we have $r_+ \le \sqrt{1 - z^2}$, with equality when $\rho = 0$. Recall that $z = c^{\star\prime} c^\pi$. Using the identity $\sin^2 x + \cos^2 x = 1$, we obtain $r_+ \le \sqrt{1 - (c^{\star\prime} c^\pi)^2} = \sin\theta(c^\star, c^\pi)$. To summarize, we have established that $\mathcal{L}(c^\pi, c^\star) \le \sin\theta(c^\star, c^\pi)$.

**Step 4.** We now construct an instance to show that, for any $c^\star$ and $c^\pi$ with $\theta(c^\star, c^\pi) < \pi/2$, one may construct an instance $\mathcal{X}$ and $f$ such that the regret is given by $\sin\theta(c^\star, c^\pi)$. Let $\Pi_{c^\pi}(c^\star)$ denote the orthogonal projection of $c^\star$ onto $c^\pi$. Since $\theta(c^\star, c^\pi) < \pi/2$ and both have unity norm, we have that $\Pi_{c^\pi}(c^\star) = \cos\theta(c^\star, c^\pi) \cdot c^\pi$. We now define $r = c^\star - \Pi_{c^\pi}(c^\star)$. Note that $r$ defines the residual of the projection of $c^\star$ onto $c^\pi$. Therefore, it is orthogonal to $c^\pi$ and $\|r\| = \sin\theta(c^\star, c^\pi)$. Then, by summing the angles within the triangle, we have that $\theta(c^\star, c^\pi) + \theta(r, c^\star) + \pi/2 = \pi$, and thus $\theta(r, c^\star) = \pi/2 - \theta(c^\star, c^\pi)$.

Now we set $\delta = r/\|r\|$ and let nature pick $f(x) = x$ and $\mathcal{X} = \{x_1, x_2\}$ such that $x_2 - x_1 = \delta$. By construction, $\delta$ is parallel to $r$ and must be orthogonal to $c^\pi$. Then, $\delta' c^\pi = (x_2 - x_1)' c^\pi = 0$, which implies that $x_1, x_2 \in \psi(c^\pi, \mathcal{X}, f)$. Moreover,

$$\delta' c^\star = \frac{r' c^\star}{\|r\|} = \frac{\|r\|\|c^\star\| \cos\theta(r, c^\star)}{\|r\|} = \cos\theta(r, c^\star) = \cos(\pi/2 - \theta(c^\star, c^\pi)) = \sin\theta(c^\star, c^\pi),$$

completing the proof.

□

***Proof of Lemma 2.*** If $C = \{0\}$, then, for any $\hat{c} \in S^d$, $\sup_{c \in C} \theta(c, \hat{c}) = 0$ and any $\hat{c} \in S^d$ is a minimizer. Next we consider the nontrivial case where $C \setminus \{0\}$ is nonempty.

Define the set $\tilde{C} = \{\tilde{c} \in \mathbb{R}^d \ : \ \tilde{c} = \frac{c}{\|c\|} \text{ for some } c \in C \setminus \{0\}\}$ to be a set of normalized vectors from $C \setminus \{0\}$. Then, the uncertainty angle of $C$ satisfies:

$$\inf_{\hat{c} \in S^d} \sup_{c \in C} \theta(c, \hat{c}) = \inf_{\hat{c} \in S^d} \sup_{c \in C \setminus \{0\}} \theta(c, \hat{c}) = \inf_{\hat{c} \in S^d} \sup_{c \in C \setminus \{0\}} \arccos \frac{c' \hat{c}}{\|c\|} = \inf_{\hat{c} \in S^d} \sup_{c \in \tilde{C}} \arccos c' \hat{c},$$

where the first equality follows from the origin being a suboptimal solution of the maximization $\sup_{c \in C}$ since $\theta(0, \hat{c}) = 0$ and $\theta$ is a non-negative function, the second equality follows from the definition of an angle, and the third equality follows from replacing $C \setminus \{0\}$ with the normalized $\tilde{C}$.

Since arccos is a decreasing continuous function on $[-1, 1]$, we have that the uncertainty angle satisfies:

$$\inf_{\hat{c}\in S^d} \sup_{c\in\tilde{C}} \arccos c'\hat{c} = \inf_{\hat{c}\in S^d} \arccos \inf_{c\in\tilde{C}} c'\hat{c} = \arccos \sup_{\hat{c}\in S^d} \inf_{c\in\tilde{C}} c'\hat{c},$$

and the circumcenter $\hat{c}$ that optimizes the uncertainty angle is the same one that solves the problem:

$$\sup_{\hat{c}\in S^d} \inf_{c\in\tilde{C}} c'\hat{c}. \tag{B-3}$$

Define the function $g(\hat{c}) = \inf_{c\in C} c'\hat{c}$. Then, for any $r > 0$ and $u \in S^d$, we have that:

$$g(\hat{c} + ru) = \inf_{c\in\tilde{C}} c'(\hat{c} + ru) = \inf_{c\in\tilde{C}} \{c'\hat{c} + rc'u\},$$

From the Cauchy-Schwarz inequality, we know that $|c'u| \leq \|c\| \cdot \|u\| = 1$ since both $c$ and $u$ belong to $S^d$. This implies that $-r \leq rc'u \leq r$. Then,

$$g(\hat{c}) - r \leq g(\hat{c} + ru) \leq g(\hat{c}) + r \implies |g(\hat{c} + ru) - g(\hat{c})| \leq r,$$

and we have that $g$ is a continuous function in $\mathbb{R}^d$. Therefore, the problem $\sup_{\hat{c}\in S^d} g(\hat{c})$ from Eq. (B-3) is an optimization problem with a continuous objective function over a compact space. By the Weierstrass theorem, the sup is attained. Since the maximizer of Eq. (B-3) is also the minimizer of the uncertainty angle, the infimum of that problem (the circumcenter) is also attained.

Now we prove uniqueness under the assumption that $\alpha(C) < \pi/2$. Suppose for a moment that the circumcenter is not unique and that $\hat{c}_1$ and $\hat{c}_2$ are two distinct optimal solutions of Eq. (B-3). We will show that we can construct a new solution $\hat{c}_3$ which is strictly better than $\hat{c}_1$ and $\hat{c}_2$, leading to a contradiction.

Let $\hat{c}_3 = (\hat{c}_1 + \hat{c}_2)/2$. First we argue that $\hat{c}_3 \neq 0$. If $\hat{c}_3 = 0$, this would imply that $\hat{c}_1 = -\hat{c}_2$. However, Since $\alpha(C) < \pi/2$, then, it must be the case that $c'\hat{c}_1 > 0$, for all $c \in C \setminus \{0\}$ and $c'\hat{c}_2 > 0$, for all $c \in C \setminus \{0\}$, which implies that $c \in C \setminus \{0\}$ is empty, violating the assumption of the lemma. Thus, $\|\hat{c}_3\| > 0$. We also have that $\|\hat{c}_3\| < 1$ since $\hat{c}_3$ is a convex combination of two distinct vectors on the unit sphere.

Let $z$ be the optimal value of Eq. (B-3), i.e., $z = \inf_{c\in\tilde{C}} c'\hat{c}_1 = \inf_{c\in\tilde{C}} c'\hat{c}_2$. Then,

$$z = \frac{1}{2} \inf_{c\in\tilde{C}} c'\hat{c}_1 + \frac{1}{2} \inf_{c\in\tilde{C}} c'\hat{c}_2 \overset{(a)}{<} \frac{\frac{1}{2}\inf_{c\in\tilde{C}} c'\hat{c}_1 + \frac{1}{2}\inf_{c\in\tilde{C}} c'\hat{c}_2}{\|\hat{c}_3\|} \overset{(b)}{\leq} \frac{\inf_{c\in\tilde{C}} c'\left(\frac{1}{2}\hat{c}_1 + \frac{1}{2}\hat{c}_2\right)}{\|\hat{c}_3\|} = \inf_{c\in\tilde{C}} c' \frac{\hat{c}_3}{\|\hat{c}_3\|},$$

where $(a)$ follows from $\|\hat{c}_3\| < 1$ and $(b)$ follows from combining two infinimums. Since $\hat{c}_3/\|\hat{c}_3\|$ is a feasible solution of Eq. (B-3) and its objective value is strictly above $z$, this would violate the optimality of $\hat{c}_1$ and $\hat{c}_2$. This is a contradiction, and the circumcenter must be unique. □

***Proof of Theorem 1.*** First we show how to construct the upper bound by leveraging the result provided in Lemma 1. Next we how the construct a set $C$ such that any policy incurs a worst-case regret at least as high as the upper bound provided for the circumcenter policy.

*Step 1.* We start by showing that for any policy $c^\pi \in \mathcal{P}'$, we can bound the objective function from Eq. (4) using the worst-case regret loss function $\mathcal{L}$ for specific choices of nature $(c^\star, \mathcal{X}, f)$, defined in Eq. (8). The following inequality holds for any knowledge set $C \subseteq S^d$:

$$\begin{aligned}
\inf_{\pi\in\mathcal{P}} \sup_{\mathcal{X}\in\mathcal{B},\, f\in\mathcal{F},\, c^\star\in C} \Big(f(x^\pi) - f(x^\star)\Big)' c^\star &\overset{(a)}{\leq} \inf_{\pi\in\mathcal{P}'} \sup_{\mathcal{X}\in\mathcal{B},\, f\in\mathcal{F},\, c^\star\in C} \Big(f(x^\pi) - f(x^\star)\Big)' c^\star \\
&\overset{(b)}{=} \inf_{c^\pi\in S^d} \sup_{\mathcal{X}\in\mathcal{B},\, f\in\mathcal{F},\, c^\star\in C} \Big(f(x^\pi) - f(x^\star)\Big)' c^\star \\
&\overset{(c)}{\leq} \inf_{c^\pi\in S^d} \sup_{c^\star\in C} \mathcal{L}(c^\pi, c^\star). \qquad \text{(B-4)}
\end{aligned}$$

where $(a)$ follows from restricting $\pi$ to $\mathcal{P}'$, $(b)$ follows from representing $\pi \in \mathcal{P}'$ in terms of its cost vector $c^\pi$, and $(c)$ follows from the definition of $\mathcal{L}$ (it would be an equality if $x^\pi$ were unique).

Let us define:

$$g(x) \quad = \quad \begin{cases} \sin x & \text{if } 0 \leq x < \pi/2, \\ 1 & \text{if } x \geq \pi/2. \end{cases}$$

Lemma 1 shows that $\mathcal{L}(c^\pi, c^\star) = g(\theta(c^\star, c^\pi))$. Combining with Eq. (B-4), we have:

$$\begin{aligned}
\inf_{\pi\in\mathcal{P}} \sup_{\mathcal{X}\in\mathcal{B},\, f\in\mathcal{F},\, c^\star\in C} \Big(f(x^\pi) - f(x^\star)\Big)' c^\star &\leq \inf_{c^\pi\in S^d} \sup_{c^\star\in C} g(\theta(c^\star, c^\pi)) \\
&= g\left(\inf_{c^\pi\in S^d} \sup_{c^\star\in C} \theta(c^\star, c^\pi)\right) = g(\alpha(C)),
\end{aligned}$$

where the first equality follows from $g(\cdot)$ being nondecreasing and continuous, the second equality follows from the definition of the uncertainty angle and the third equality follows from applying the tighest lower bound by using the circumcenter policy. Therefore,

$$\inf_{\pi\in\mathcal{P}} \sup_{\mathcal{X}\in\mathcal{B},\, f\in\mathcal{F},\, c^\star\in C} \Big(f(x^\pi) - f(x^\star)\Big)' c^\star \leq g(\alpha(C)) = g(\bar{\alpha}).$$

*Step 2.* To show that no policy can be uniformly better than the circumcenter, we construct instances of the problem where any policy incurs at least the regret $g(\bar{\alpha})$. We first consider the case where $\bar{\alpha} \leq \pi/2$. Let $e = \frac{1}{\sqrt{d}}(1, \cdots, 1)'$ and $\tilde{C} = \{c \in S^d : \theta(e, c) \leq \bar{\alpha}\}$, which is the intersection between a sphere and the revolution cone with axis $e$ and aperture angle of $\bar{\alpha}$. Note that if $\bar{\alpha} = \pi/2$ there is no revolution cone, but it suffices to consider the halfspace $\tilde{C} = \{c \in S^d : c'e \geq 0\}$. Moreover, let $f(x) = x$ and $\mathcal{X} = \{x_1, x_2\}$ such that $\delta = x_2 - x_1$ is orthogonal to $e$ and $\|\delta\| = 1$.

Define $c_1^\star = \sin\bar{\alpha} \cdot \delta + \cos\bar{\alpha} \cdot e$. We now argue that $c_1^\star$ belongs to $\tilde{C}$. We first show that $c_1^\star \in S^d$:

$$\|c_1^\star\|^2 = \sin^2\bar{\alpha} \cdot \|\delta\|^2 + \cos^2\bar{\alpha} \cdot \|e\|^2 = \sin^2\bar{\alpha} + \cos^2\bar{\alpha} = 1,$$

where the first equality follows from $\delta$ and $e$ being orthogonal vectors, and the second equality follow from $\delta$ and $e$ having norm 1. We now show that $\theta(e, c_1^\star) = \bar{\alpha}$:

$$\theta(e, c_1^\star) = \arccos\frac{e'c_1^\star}{\|e\|\|c_1^\star\|} = \arccos e'c_1^\star = \arccos\cos\bar{\alpha} = \bar{\alpha},$$

where the second equality follows from $e$ and $c_1^\star$ having norm 1, and the third equality follows from $\delta$ and $e$ being orthogonal.

We now construct a second vector in $\tilde{C}$, $c_2^\star = -\sin\bar{\alpha} \cdot \delta + \cos\bar{\alpha} \cdot e$. By the same argument as above, $c_2^\star$ also belongs to $\tilde{C}$. Note that $x_1$ is the optimal action if $c_1^\star$ is the true cost and $x_2$ is the optimal action if $c_2^\star$ is the true cost since $(x_2 - x_1)'c_1^\star = \sin(\bar{\alpha}) \geq 0$ and $(x_2 - x_1)'c_2^\star = -\sin(\bar{\alpha}) \leq 0$. Since for any policy $\pi \in \mathcal{P}$ the decision-maker must choose either $x_1$ or $x_2$, we have:

$$\begin{aligned}
\inf_{\pi\in\mathcal{P}} \sup_{\mathcal{X}\in\mathcal{B},\, f\in\mathcal{F},\, c^\star\in\tilde{C}} \Big(f(x^\pi) - f(x^\star)\Big)' c^\star &\overset{(a)}{\geq} \inf_{x^\pi\in\{x_1,x_2\}} \sup_{c^\star\in\tilde{C}} \Big(x^\pi - x^\star\Big)' c^\star \\
&\overset{(b)}{\geq} \inf_{x^\pi\in\{x_1,x_2\}} \sup_{c^\star\in\{c_1^\star,c_2^\star\}} \Big(x^\pi - x^\star\Big)' c^\star \\
&\overset{(c)}{=} \min\{(x_2 - x_1)'c_1^\star, (x_1 - x_2)'c_2^\star\}, \\
&\overset{(d)}{=} \min\{\delta' c_1^\star, -\delta' c_2^\star\} \\
&\overset{(e)}{=} \min\{\sin\bar{\alpha}, \sin\bar{\alpha}\} = \sin\bar{\alpha},
\end{aligned}$$

where $(a)$ follows from our choice of instance $(\mathcal{X}, f)$, $(b)$ follows from restricting the choice of $c^\star \in C$ to $\{c_1^\star, c_2^\star\}$, $(c)$ follows from the fact that $x_1$ is optimal for $c_1^\star$ and $x_2$ is optimal for $c_2^\star$, $(d)$ follows from the definition of $\delta$, and $(e)$ follows from the definitions of $c_1^\star$ and $c_2^\star$ and the fact that $\delta$ and $e$ are orthogonal. This completes the argument for the case where $\bar{\alpha} \leq \pi/2$.

Finally, we now consider the case where $\pi/2 < \bar{\alpha} \leq \pi$. Similarly to the previous case, let $\tilde{C} = \{c \in S^d : \theta(e, c) \leq \bar{\alpha}\}$, however, notice that this set is not a revolution cone anymore. Despite that, it is still well defined and $\alpha(C) = \bar{\alpha}$, $\hat{c}(C) = e$ by construction. Moreover, let $f(x) = x$ and $\mathcal{X} = \{x_1, x_2\}$ such that $\delta = x_2 - x_1$ is orthogonal to $e$ and $\|\delta\| = 1$. But note that in this case, $x_1 \in \tilde{C}$ since $x_1'e = 0$ which implies that $\theta(x_1, e) = \pi/2 < \bar{\alpha}$. By the same argument, $x_2 \in \tilde{C}$. Therefore, no matter the policy $\pi$ used to choose between $x_1$ and $x_2$, nature can always pick a vector in $\tilde{C}$ to cause maximum regret. This completes the proof.

□

***Proof of Theorem 2.*** Let $\pi = \pi_{greedy}$ and let $e_i$ denote the $i$-th vector of the canonical basis in $\mathbb{R}^d$. For $c \in \mathbb{R}^d$, we use $c_{[i]}$ to denote the $i$-th entry of the vector $c$.

The proof strategy is organized as follows. In step 1, we define a useful family of sets that will characterize our sequence of knowledge sets. In step 2, we construct choices of sets $\mathcal{X}_t$ and context functions $f_t$ so that the greedy policy implies that the sequence of knowledge sets $C(\mathcal{I}_t)$ always belongs to the family of sets that we defined in step 1. For this construction, we will focus on a 3-dimensional case, i.e., $d = 3$. In step 3, we show that the regret in every time period must be uniformly bounded away from zero regardless of the time horizon, leading to the linear regret.

*Step 1.* We first define a family of sets that will be central in the construction of instances with linear regret.

Define $h_1 = (2\sin^2\bar{\alpha}, -\sin 2\bar{\alpha}, \sin 2\bar{\alpha})$, $h_2 = (2\sin^2\bar{\alpha}, -\sin 2\bar{\alpha}, -\sin 2\bar{\alpha})$ and $h_3 = (-\varepsilon, 1, 0)$. Consider the following family of sets indexed by $\bar{\alpha}$ and $\varepsilon$ with $0 < \bar{\alpha} < \pi/2$ and $0 \le \varepsilon \le (\tan\bar{\alpha})/2$:

$$C_{\varepsilon,\bar{\alpha}} = \{c \in \mathbb{R}^d : \ h_1' c \ge 0\} \cap \{c \in \mathbb{R}^d : \ h_2' c \ge 0\} \cap \{c \in \mathbb{R}^d : \ h_3' c \ge 0\} \cap S^d.$$

Note that $C_{\varepsilon,\bar{\alpha}}$ is the intersection of a polyhedral cone with the unit sphere. Using the halfpsaces defined by $h_1$, $h_2$ and $h_3$, we can compute the generators of such a polyhedral cone. The generators of the cone are $g_1 = (1, \varepsilon, \varepsilon - \tan\bar{\alpha})$, $g_2 = (1, \varepsilon, \tan\bar{\alpha} - \varepsilon)$ and $g_3 = (\cos\bar{\alpha}, \sin\bar{\alpha}, 0)$. For simplicity, the generators were not normalized. The set $C_{\varepsilon,\bar{\alpha}}$ is always nonempty if $\varepsilon \le \tan\bar{\alpha}$ and $\bar{\alpha} < \pi/2$ and we fix $c^\star = g_3 = (\cos\bar{\alpha}, \sin\bar{\alpha}, 0)$. Using Definition 2, one can also show that $\hat{c}(C_{\varepsilon,\bar{\alpha}}) = (1, \varepsilon, 0)$. In Figure 8a, we depict an example of initial knowledge set $C_0$ for $\varepsilon = 0$ and $\bar{\alpha} = \pi/4$.

*Step 2.* Fix $0 < \bar{\alpha} < \pi/2$ and define a sequence of instances as follows. We let $C_0 = C_{0,\bar{\alpha}}$ (Figure 8a depicts the set $C_0$ for $\bar{\alpha} = \pi/4$). Suppose that $f_t$ is the identity for all $t \ge 1$ and let

$$\varepsilon_1 = \frac{\tan\bar{\alpha}}{2T}, \qquad \varepsilon_t = t\varepsilon_1, \quad t \ge 2,$$
$$\mathcal{X}_t = \{\bar{x}_t, 0\}, \quad t \ge 1,$$

where

$$\bar{x}_t = (e_2 - \varepsilon_t e_1)/\|e_2 - \varepsilon_t e_1\|.$$

In Figure 8b, we provide an illustration for $\bar{\alpha} = \pi/4$ and $T$ taken to be 5, so $\varepsilon_1 = 0.1$.

Next, we establish by induction on $t$ that, under the greedy circumcenter policy, $C(\mathcal{I}_t) = C_{\varepsilon_{t-1},\bar{\alpha}}$ for $t \ge 2$. Note that the result is trivial for $t = 1$ since $C(\mathcal{I}_t) = C_0 = C_{0,\bar{\alpha}}$ by construction. Next we establish

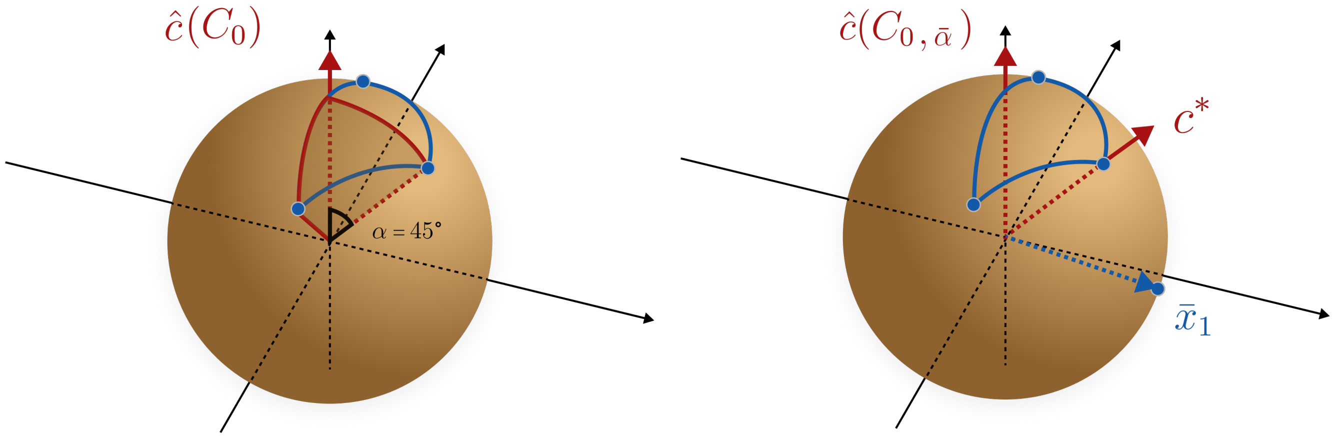


(a) (b)

Figure 8: In $(a)$, we depict the initial knowledge set $C_0$ and its circumcenter. In $(b)$, we depict the first instance of the optimization problem faced by the decision-maker and the true cost vector $c^\star$.

the base case ($t = 2$).

By definition (Eq. (5)) we have that

$$C(\mathcal{I}_2) = C_0 \cap \{c \in \mathbb{R}^d : c' f_1(\bar{x}_1) \leq c' f_1(x),\ \forall\, x \in \mathcal{X}_1\}.$$

Moreover, since $C(\mathcal{I}_1) = C_{0,\bar{\alpha}}$, $f_1$ is the identity and $\mathcal{X}_1 = \{\bar{x}_1, 0\}$, we must have that $\hat{c}(C(\mathcal{I}_1)) = e_1$, $\psi(e_1, f_1, \mathcal{X}_1) = \bar{x}_1$ and $\psi(c^\star, f_1, \mathcal{X}_1) = 0$. Which leads to

$$C(\mathcal{I}_2) = C_0 \cap \{c \in \mathbb{R}^d\ :\ -\varepsilon_1 c_{[1]} + c_{[2]} \geq 0\} = C_{\varepsilon_1,\bar{\alpha}},$$

and the base case is established. Next we show the induction step. Suppose that the result holds for $t$. Then, the circumcenter of $C(\mathcal{I}_t)$ is given by $(1, \epsilon_{t-1}, 0)$. Therefore, $\psi(c^\pi_t, f_t, \mathcal{X}_t) = \bar{x}_t$ and $\psi(c^\star, f, \mathcal{X}_t) = 0$, which leads to the update:

$$\begin{aligned} C(\mathcal{I}_{t+1}) &= C(\mathcal{I}_t) \cap \{c \in \mathbb{R}^d\ :\ -\varepsilon_t c_{[1]} + c_{[2]} \geq 0\} \\ &= C_0 \cap \{c \in \mathbb{R}^d\ :\ -\varepsilon_{t-1} c_{[1]} + c_{[2]} \geq 0\} \cap \{c \in \mathbb{R}^d\ :\ -\varepsilon_t c_{[1]} + c_{[2]} \geq 0\} \\ &= C_0 \cap \{c \in \mathbb{R}^d\ :\ -\varepsilon_t c_{[1]} + c_{[2]} \geq 0\} = C_{\varepsilon_t,\bar{\alpha}}, \end{aligned}$$

which concludes the proof by induction. Having established the above, we now analyze the regret in each period $t$.

*Step 3.* From step 2, we have for every time $t$ that $C(\mathcal{I}_{t+1}) = C_{\varepsilon_t,\bar{\alpha}}$, $\hat{c}(C(\mathcal{I}_{t+1})) = \frac{1}{\sqrt{1+\varepsilon_t^2}}(1, \varepsilon_t, 0)$ and the regret at period $t$ is given by

$$\delta_t^{\pi\prime} c^\star = \frac{(-\epsilon_t e_1 + e_2)'(\cos\bar{\alpha}, \sin\bar{\alpha}, 0)}{\sqrt{1+\varepsilon_t^2}} = \frac{\sin\bar{\alpha} - \varepsilon_t \cos\bar{\alpha}}{\sqrt{1+\varepsilon_t^2}} \geq \frac{T\sin\bar{\alpha} - (t/2T)\sin\bar{\alpha}}{\sqrt{2}} \geq \frac{\sin\bar{\alpha}}{4},$$

where the first inequality follows from the fact that $\varepsilon^2 \leq 1$. Therefore, the cumulative regret must be $\Omega(T)$.

In Figure 9a, we can see the initial knowledge set $C_0$ for $\bar{\alpha} = \pi/4$. In Figure 9b, we have the updated knowledge set $C(\mathcal{I}_2) = C_{\varepsilon,\bar{\alpha}}$ after solving the first optimization instance. In Figure 9c we have the final set after collecting the feedback of the last time period $T = 5$. No matter the horizon $T$, nature can always adjust $\epsilon_1$ as a function of $\bar{\alpha}$ and $T$ in in order to ensure that the updates are not enough to make the circumcenter and the true cost vector sufficiently close to each other.

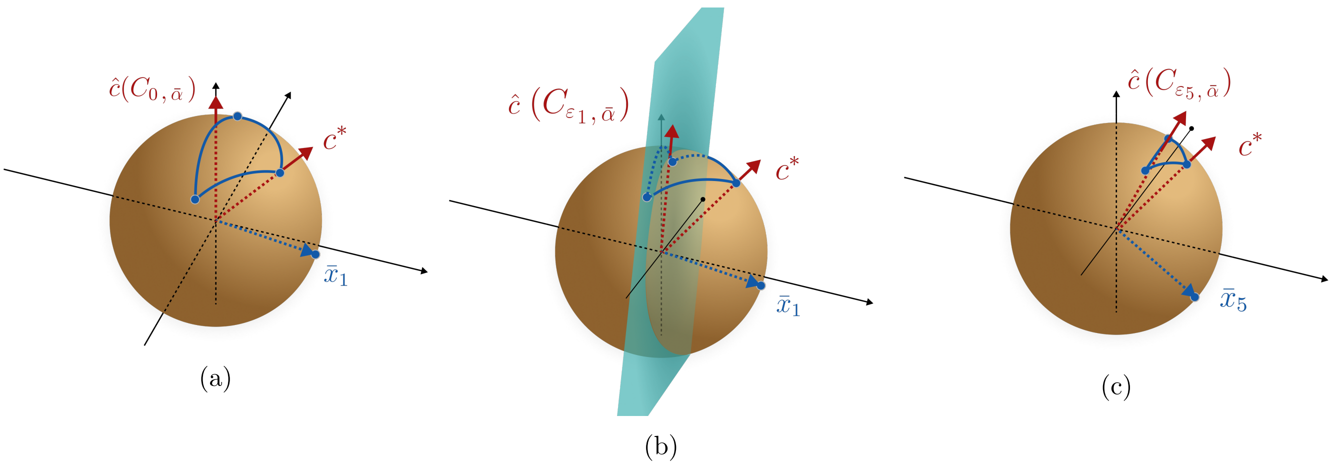


Figure 9: In $(a)$ we have the initial knowledge set and the first instance of the optimization problem. In $(b)$, the updated set $C(\mathcal{I}_2)$, the new circumcenter $\hat{c}(C(\mathcal{I}_2))$, the true cost vector and the previous action $\bar{x}_1$. In $(c)$, the knowledge set after collecting the feedback of period $T$.

□

***Proof of Lemma 3.*** Lemma 3 is a direct application of a more general result, described in Lemma B-1, stated and proved in Appendix B.1. In particular, if $t$ is a cone update period, we take $p = d$, $\eta = 0$, $\delta = \delta_t^\pi$ and $E(W,U) = E(W_t, U_t)$. We have that $\delta'\hat{c}(E(W_t,U_t)) = \delta_t^{\pi\prime} c_t^\pi \leq 0$ where the inequality holds by definition of the effective difference $\delta_t^\pi$ (Eq. 11) and the circumcenter policy applied to ellipsoidal cones. □

***Proof of Proposition 1.*** Note that the circumcenter of $E(W,U)$ is $Ue_1$. Let $E_H(W) = E(W) \cap H$, where $H = \{c \in \mathbb{R}^d : \ c'e_1 = 1\}$. Next, recall that for any $c \in E(W) \setminus \{0\}$, the first component is always greater than zero. Then,

$$\alpha(E(W,U)) = \sup_{c \in E(W,U)} \theta(c, Ue_1) = \sup_{c \in E(W)} \theta(c, e_1) = \sup_{c \in E(W)\setminus\{0\}} \theta(c, e_1) = \sup_{c \in E_H(W)} \theta(c, e_1),$$

where the second equality follows from the the fact that angles are preserved by orthonormal transformations, the third equality follows from the suboptimality of $\{0\}$ and the last equality follows from the fact that scaling a vector by a positive constant does not affect the angle. Moreover, for every $c \in E_H(W)$, we have that

$$\tan\theta(c, e_1) = \tan\arccos\frac{c'e_1}{\|c\|} = \tan\arccos\frac{1}{\|c\|},$$

where the first equality follows from the definition of an angle and the second one follows from $E_H(W) \subset H$. Hence,

$$\tan\theta(c, e_1) = \sqrt{\|c\|^2 - 1} = \|c_{[2:d]}\|,$$

where we used the trigonometric identity $\tan\arccos x = \frac{\sqrt{1-x^2}}{x}$ for $x \in [0, \pi/2)$. Using the fact that $E_H(W)$ is an ellipsoid one gets that $\|c_{[2:d]}\| \leq \sqrt{\lambda_{\max(M)}}$ and the inequality is tight for some $c \in E_H(W)$. Since $\tan(\cdot)$ is continuous and monotone increasing on $[0, \pi/2)$, we get

$$\alpha(E(W,U)) = \sup_{c \in E_H(W)} \theta(c, e_1) = \sup_{c \in E_H(W)} \arctan\left(\tan\theta(c, e_1)\right) = \arctan\sqrt{\lambda_{\max(M)}}.$$

This completes the proof. □

***Proof of Lemma 4.*** Consider any $c \in E(W_t, U_t) \cap S^d$. Let $\tilde{c} = U_t^{-1}c$ and $\tilde{\delta}_t^\pi = U_t^{-1}\delta_t^\pi$. Recalling that for an orthonormal matrix $U$, $U^{-1} = U'$, the regret in period $t$ if the true underlying cost is $c$ is given by

$$\delta_t^{\pi\prime} c \;=\; \delta_t^{\pi\prime} U_t U_t^{-1} c \;=\; (U_t^{-1}\delta_t^\pi)' U_t^{-1} c \;=\; (\tilde{\delta}_t^\pi)'\tilde{c}.$$

Note that $\tilde{c} \in E(W_t) \cap S^d$ and hence $\tilde{c}_{[1]} \in (0, 1]$. Let $\tilde{\nu} = \tilde{c}_{[2:d]}/\tilde{c}_{[1]}$ and note that $(1, \tilde{\nu}) \in E_H(W)$. In turn, we have

$$\begin{aligned}
\delta_t^{\pi\prime} c &= \tilde{c}_{[1]}\left(\tilde{\delta}_{t,[1]}^\pi + (\tilde{\delta}_{t,[2:d]}^\pi)'\tilde{\nu}\right) \\
&\leq \tilde{c}_{[1]}\left(\tilde{\delta}_{t,[1]}^\pi + \sup_{(1,\nu)\in E_H(W)} (\tilde{\delta}_{[2:d]}^\pi)'\nu\right) \\
&\overset{(a)}{=} \tilde{c}_{[1]}\left(\tilde{\delta}_{t,[1]}^\pi + \sqrt{(\tilde{\delta}_{t,[2:d]}^\pi)' W_t \tilde{\delta}_{t,[2:d]}^\pi}\right) \\
&\overset{(b)}{\leq} \tilde{c}_{[1]}\left(\sqrt{(\tilde{\delta}_{t,[2:d]}^\pi)' W_t \tilde{\delta}_{t,[2:d]}^\pi}\right) \\
&\overset{(c)}{\leq} \epsilon,
\end{aligned}$$

where $(a)$ follows from the fact that $E_H(W)$ is an ellipsoid and the optimization of a linear function over an ellipsoid has an analytical solution as given above (see, for instance, Boyd and Vandenberghe (2004)), $(b)$ holds due to the fact that $\tilde{\delta}^{\pi}_{t,[1]} = (\tilde{\delta}^{\pi})' e_1 = \delta^{\pi}_t{}' U_t e_1 = \delta^{\pi}_t{}' c^{\pi}_t \leq 0$ (see Eq. (11)), and $(c)$ follows from $c^{\star}_{[1]} \leq 1$ and the assumption of the lemma. This completes the proof.

□

***Proof of Lemma 5.*** Let $\lambda_i(W)$ denote the $i$-th eigenvalue of $W$ in nondecreasing order. Suppose first that $\lambda_1(W_1) > \left(\frac{\epsilon}{10(d-1)}\right)^2$. By Lemma 3, we have that that for every update period $t$,

$$\prod_{i=1}^{d-1} \lambda_i(W_{t+1}) \quad \leq \quad e^{-1/(d-1)} \prod_{i=1}^{d-1} \lambda_i(W_t).$$

For any time $t$ after exactly $I^{\pi}_T$ updates took place, we can apply the latter recursively to obtain

$$\prod_{i=1}^{d-1} \lambda_i(W_t) \quad \leq \quad e^{-I^{\pi}_T/(d-1)} \prod_{i=1}^{d-1} \lambda_i(W_1) \ \leq \ e^{-I^{\pi}_T/(d-1)} (\lambda_{max}(W_1))^{d-1}. \tag{B-5}$$

Next we lower bound, for any $t$, the eigenvalues $\lambda_i(W_t)$. We will establish that $\lambda_1(W_t) \geq \left(\frac{\epsilon}{10d}\right)^2$ for all $t \geq 1$.

If $I^{\pi}_T = 0$, then we have that $\lambda_1(W_t) = \lambda_1(W_1) > \left(\frac{\epsilon}{10(d-1)}\right)^2 > \left(\frac{\epsilon}{10d}\right)^2$.

Suppose now that $I^{\pi}_T > 0$. We will show by induction that $\lambda_1(W_t) \geq \left(\frac{\epsilon}{10d}\right)^2$ for all $t \geq 1$.

This clearly true for $t = 1$. Suppose that it is true at time $s$. If there is no update, then this is trivially true for $s + 1$. If there is an update $s$, two cases can happen.

Case 1: $\lambda_1(W_s) > \left(\frac{\epsilon}{10(d-1)}\right)^2$. Note that the largest decrease possible in any eigenvalue after applying the ellipsoid method in $\mathbb{R}^{d-1}$ is $(d-1)^2/d^2$ (this is the decrease that happens when the cut is along that particular eigenvector). As a result, we must have $\lambda_1(W_{s+1}) \geq ((d-1)^2/d^2)\left(\frac{\epsilon}{10(d-1)}\right)^2 = \left(\frac{\epsilon}{10d}\right)^2$.

Case 2: $\lambda_1(W_s) \leq \left(\frac{\epsilon}{10(d-1)}\right)^2$. In this case, given that updates are only performed when $\delta'_{s,[2:d]} W_s \delta_{s,[2:d]} > \epsilon^2$, we can use (Cohen et al., 2020, Lemma 4). Indeed, in an update period, the second to the fourth update equations in algorithm `ConeUpdate` are precisely the update equations for the ellipsoid method for an ellipsoid in $\mathbb{R}^{d-1}$. Cohen et al. (2020) introduce a version of the ellipsoid method where updates are not performed if the length of the ellipsoid along the direction to be cut is smaller than a certain threshold $\epsilon$. The condition in Eq. (12) for an update in our paper is essentially the same (our condition is on the distance from the center to the edge of the ellipsoid, but this is simply half the length).

Recall the notation used in algorithm `ConeUpdate` and note that every time that we update the matrix $W_t$ with the ellipsoid method, we get a matrix $N$. Let $\lambda_1(N)$ denote the smallest eigenvalue of the matrix $N$. By construction, if $t$ is an update period, we have that $\lambda_1(W_{t+1}) = \lambda_1(N)$. Lemma 4 from Cohen et al. (2020) says that if we update our ellipsoidal cone according to algorithm `ConeUpdate`, if $\lambda_1(W_s) \le \left(\frac{\epsilon}{10(d-1)}\right)^2$ and $\delta'_{s,[2:d]} W_s \delta_{s,[2:d]} > \epsilon^2$, then $\lambda_1(W_{s+1}) \ge \lambda_1(W_s)$, i.e., the smallest eigenvalue does not decrease after the update.

In this case, we deduce that $\lambda_1(W_{s+1}) \ge \lambda_1(W_s) \ge \left(\frac{\epsilon}{10d}\right)^2$, where the last inequality follows from the induction hypothesis. This concludes the induction.

Combining this result with (B-5), we obtain

$$\lambda_{max}(W_t) \le \left(\frac{10d}{\epsilon}\right)^{2(d-2)} \left(\lambda_{max}(W_1)\right)^{(d-1)} e^{-I_T^\pi/(d-1)}.$$

Suppose now that $\lambda_1(W_1) \le \left(\frac{\epsilon}{10(d-1)}\right)^2$. By construction, $W_1$ is a revolution cone and $\lambda_i(W_1) = \lambda_1(W_1)$ for $i = 1, ..., d-1$. For any vector $\delta \in \mathbb{R}^{d-1}$ with $\|\delta\| \le 1$, we have $\delta' W_1 \delta \le \|\delta\|^2 \lambda_{max}(W_1) \le \left(\frac{\epsilon}{10(d-1)}\right)^2 \le \epsilon^2$. Hence, in this case, no update takes place, $I_T^\pi = 0$, $W_t = W_1$ for all $t$, and by Lemma 4, the per period regret is bounded above by $\epsilon$. This concludes the proof. $\square$

***Proof of Theorem 3.*** We prove Theorem 3 in four steps. In the first step, we prove that $C_0 \subseteq E(W_1, U_1)$ and, that for every $t \ge 1$, $C(\mathcal{I}_t) \subseteq E(W_{t+1}, U_{t+1})$. In the second step, we decompose the cumulative regret based on periods in which the condition in Eq. (12) is satisfied or violated. In the third step, we upper bound the number of times that Eq. (12) can be violated, to obtain a bound on the cumulative regret. In the fourth step, we prove the running time claim.

**Step 1.** Note that $E(W_1, U_1)$, by construction, is precisely the revolution cone with aperture angle $\alpha(C_0)$ that contains $C_0$. The inclusions $C(\mathcal{I}_t) \subseteq E(W_{t+1}, U_{t+1})$, $1 \le t \le T$ follows from our construction and Lemma 3.

**Step 2.** Recall that $I_T^\pi = \sum_{t=1}^T \mathbf{1}\{\delta'_{t,[2:d]} W_t \delta_{t,[2:d]} > \epsilon^2\}$ is the number of cone-update periods. We have the following regret bound.

$$\begin{aligned}
\mathcal{R}_T^\pi\left(c^\star, \vec{\mathcal{X}}_T, \vec{f}_T\right) &= \sum_{t=1}^T \left(f_t(x_t^\pi) - f_t(x_t^\star)\right)' c^\star \\
&= \sum_{t=1}^T \|f_t(x_t^\pi) - f_t(x_t^\star)\| \delta_t^{\pi\prime} c^\star \\
&\overset{(a)}{\leq} \sum_{t=1}^T \left(\mathbf{1}\{\delta'_{t,[2:d]} W_t \delta_{t,[2:d]} > \epsilon^2\} + \mathbf{1}\{\delta'_{t,[2:d]} W_t \delta_{t,[2:d]} \leq \epsilon^2\}\right) \delta_t^{\pi\prime} c^\star \\
&\overset{(b)}{\leq} \sum_{t=1}^T \mathbf{1}\{\delta'_{t,[2:d]} W_t \delta_{t,[2:d]} > \epsilon^2\} \delta_t^{\pi\prime} c^\star + \sum_{t=1}^T \mathbf{1}\{\delta'_{t,[2:d]} W_t \delta_{t,[2:d]} \leq \epsilon^2\} \epsilon \\
&\overset{(c)}{\leq} \sum_{t=1}^T \mathbf{1}\{\delta'_{t,[2:d]} W_t \delta_{t,[2:d]} > \epsilon^2\} + \sum_{t=1}^T \mathbf{1}\{\delta'_{t,[2:d]} W_t \delta_{t,[2:d]} \leq \epsilon^2\} \epsilon \\
&= I_T^\pi + (T - I_T^\pi)\,\epsilon \\
&\leq I_T^\pi + T\epsilon, \qquad \text{(B-6)}
\end{aligned}$$

where $(a)$ follows from the fact that $\|f_t(x_t^\pi) - f_t(x_t^\star)\| \leq 1$ since $f_t \in \mathcal{F}$, $\mathcal{X}_t \in \mathcal{B}$, for all $t \leq T$, $(b)$ follows from Lemma 4, and $(c)$ follows from Cauchy-Schwarz inequality: $\delta_t^{\pi\prime} c^\star \leq \|\delta_t^\pi\| \|c^\star\| = 1$.

**Step 3.** We now provide an upper bound for $I_T^\pi$. We need to consider two separate cases as a function of $\alpha(C_0)$.

If $\alpha(C_0) \leq \arctan\epsilon$, an application of Theorem 1 shows that the regret of every period is less than $\sin\arctan\epsilon = \epsilon/(1+\epsilon)$ which is enough to ensure a performance of at least $\epsilon$ when using the circumcenter as the proxy cost vector. Moreover, Proposition 1 and our choice of $E(W_1, U_1)$ in Algorithm `EllipsoidalCones` implies that $\lambda_{max}(W_1) \leq \epsilon^2$, and hence, for every possible vector $\delta$, $\delta'_{[2:d]} W_1 \delta_{[2:d]} \leq \epsilon^2$. In this case, the algorithm never has a cone-update period and $I_T^\pi = 0$. Therefore, Eq. (B-6) implies that the regret is bounded by $T\epsilon$, which is equal to $d$ by our choice of $\epsilon$.

When $\alpha(C_0) > \arctan\epsilon$, we may have cone-update periods. We show that after an $\mathcal{O}(\log 1/\epsilon)$ amount of update steps, it must be the case that $\alpha(E(W_t, U_t)) < \arctan\epsilon$. Then, from this time onward, it must be the case that the algorithm never updates again and Lemma 4 implies that the regret is upper bounded by $\epsilon$ for every period from this time onward.

Suppose that we updated the ellipsoidal cone $\tau$ times. Lemma 5 and the revolution cone initialization, $\lambda_1(W_1) = \cdots = \lambda_{d-1}(W_1) = \tan^2\alpha(C_0)$ gives us that

$$\lambda_{max}(W_t) \leq \left(\frac{10d}{\epsilon}\right)^{2(d-2)} \left(\lambda_{max}(W_1)\right)^{d-1} e^{-\frac{\tau}{(d-1)}} = \left(\frac{10d}{\epsilon}\right)^{2(d-2)} \left(\tan\alpha(C_0)\right)^{2(d-1)} e^{-\frac{\tau}{(d-1)}},$$

for every $t$ after $\tau$ updates. Note that if we have $\tau = 2(d-1)^2 \ln\left(\frac{10d \tan\alpha(C_0)}{\epsilon}\right)$, then

$$\begin{aligned}\lambda_{max}(W_t) &\leq \left(\frac{10d}{\epsilon}\right)^{2(d-2)} (\tan\alpha(C_0))^{2(d-1)} e^{-\frac{2(d-1)^2 \ln\left(\frac{10d\tan\alpha(C_0)}{\epsilon}\right)}{(d-1)}} \\ &\leq \left(\frac{10d}{\epsilon}\right)^{2(d-1)-2} (\tan\alpha(C_0))^{2(d-1)} \left(\frac{10d\tan\alpha(C_0)}{\epsilon}\right)^{-2(d-1)} \\ &\leq \left(\frac{10d}{\epsilon}\right)^{-2}.\end{aligned}$$

In this case, starting after $\tau$ updates, we are as in case *ii.*) of Lemma 5; there will be no cone-update periods onward, and the per period regret will bounded by $\epsilon$ in future periods. In turn,

$$I_T^{\pi} \leq 2(d-1)^2 \ln\left(\frac{10d\tan\alpha(C_0)}{\epsilon}\right). \tag{B-7}$$

Combining Eqs. (B-6) and (B-7), and selecting $\epsilon = d/T$ leads to:

$$\mathcal{R}_T^{\pi}\left(c^\star, \vec{\mathcal{X}}_T, \vec{f}_T\right) \leq 2(d-1)^2 \ln\left(\frac{10d\tan\alpha(C_0)}{\epsilon}\right) + T\epsilon \leq 2(d-1)^2 \ln\left(10T\tan\alpha(C_0)\right) + d.$$

**Step 4.** This algorithm runs in polynomial time in $d$ and $T$ since every period's computation is a function only of $d$. Low-regret periods are computationally very cheap, while each cone-update period requires a spectral decomposition, which has a running time upper bound of $\mathcal{O}(d^3)$.

The proof is complete. □

***Proof of Lemma 6.*** The subspace $\Delta_{t_0+1}$ is a one-dimensional object and, by the definition of $\Delta_{t_0+1}$, we have that for every element $c \in \Pi_{\Delta_{t_0+1}}(C(\mathcal{I}_{t_0+1}))$, $c = \gamma\delta_{t_0}^{\pi}$ for some $\gamma \in \mathbb{R}$. Moreover, for every $t \geq t_0$, Eq. (10) implies that we must have for every $c \in C(\mathcal{I}_t)$ that $\delta_{t_0}^{\pi\prime} c \geq 0$. Since $\delta_{t_0}^{\pi} \in \Delta_{t_0+1}$, we have that $\delta_{t_0}^{\pi\prime} c \geq 0 \iff \delta_{t_0}^{\pi\prime}\Pi_{\Delta_{t_0+1}}(c) \geq 0$. Hence,

$$\Pi_{\Delta_{t_0+1}}(C(\mathcal{I}_{t_0+1})) \subseteq \{c \in \mathbb{R}^d : \ c = \gamma\delta_{t_0}^{\pi}, \ \gamma \geq 0\} \implies \alpha(\Pi_{\Delta_{t_0+1}}(C(\mathcal{I}_{t_0+1}))) = 0,$$

and $\Pi_{\Delta_{t_0+1}}(C(\mathcal{I}_{t_0+1}))$ lives in a pointed cone.

The second affirmative follows directly from the fact that whenever $f_t(x_t^{\pi}) = f_t(x_t^\star)$, there is no suboptimality gap and the regret is zero. □

***Proof of Lemma 7.*** For every time $t$, we have that

$$\sup_{x_t^\star \in \psi(c^\star, \mathcal{X}_t, f_t)} (f_t(x_t^{\pi}) - f_t(x_t^\star))' c^\star \overset{(a)}{\leq} \delta_t^{\pi\prime} c^\star = r_t^{\pi\prime} c^\star + \Pi_{\Delta_t}(\delta_t^{\pi})' c^\star \overset{(b)}{\leq} \eta + \Pi_{\Delta_t}(\delta_t^{\pi})' \Pi_{\Delta_t}(c^\star),$$

where $(a)$ follows from the fact that $f_t \in \mathcal{F}$, $\mathcal{X}_t \in \mathcal{B}$ and $(b)$ follows from bounding $r_t^{\pi\prime} c^\star$ with $\eta$ since we assumed $\|r_t^\pi\| \leq \eta$, and replacing $c^\star$ with its projection onto $\Delta_t$.

Next, we show that $\Pi_{\Delta_t}(\delta_t^\pi)'\Pi_{\Delta_t}(c^\star) \leq \epsilon$ when the assumptions of the lemma are satisfied. For that, note that $\Pi_{\Delta_t}(\delta_t^\pi) \in \Delta_t$ and $E(W_t, U_t) \subset \Delta_t$. Moreover, we have by assumption that $\Pi_{\Delta_t}(c^\star) \in \Pi_{\Delta_t}(C(\mathcal{I}_t)) \subseteq E(W_t, U_t)$. Hence, an application of Lemma 4 in the subspace $\Delta_t$ completes the proof.

□

***Proof of Lemma 8***. First we show that we do not exclude any feasible vector when we update the ellipsoidal cone using the projected effective difference. By the definition of the residual, we have that, for any $c \in C(\mathcal{I}_{t+1})$, that $\delta_t^{\pi\prime} c = \Pi_{\Delta_{t+1}}(\delta_t^\pi)'c + r_t^{\pi\prime} c \geq 0$, where the inequality follows from Eq. (10). Therefore,

$$\begin{aligned}
\Pi_{\Delta_{t+1}}\left(C(\mathcal{I}_{t+1})\right) &\subseteq \Pi_{\Delta_{t+1}}\Big(C(\mathcal{I}_t) \cap \{c \in \mathbb{R}^d :\ \delta_t^{\pi\prime} c \geq 0\}\Big) \\
&\subseteq \Pi_{\Delta_{t+1}}\Big(C(\mathcal{I}_t) \cap \{c \in \mathbb{R}^d :\ \Pi_{\Delta_t}(\delta_t^\pi)'c \geq -\eta\}\Big) \\
&= \Pi_{\Delta_t}\Big(C(\mathcal{I}_t) \cap \{c \in \mathbb{R}^d :\ \Pi_{\Delta_t}(\delta_t^\pi)'\Pi_{\Delta_t}(c) \geq -\eta\}\Big) \\
&\subseteq \Pi_{\Delta_t}\left(C(\mathcal{I}_t)\right) \cap \{c \in \Delta_t :\ \Pi_{\Delta_t}(\delta_t^\pi)'c \geq -\eta\} \\
&\subseteq E(W_t, U_t) \cap \{c \in \Delta_t :\ \Pi_{\Delta_t}(\delta_t^\pi)'c \geq -\eta\},
\end{aligned}$$

where the first inclusion follows from the fact that including elements in a set can only include elements in the projection of that set, the second inclusion follows from applying Chauchy-Schwarz for $r_t^{\pi\prime} c$ and using that $\|r_t^\pi\| \leq \eta$, the equality follows the fact that $\Pi_{\Delta_t}(\delta_t^\pi) \in \Delta_t$, which implies that $\Pi_{\Delta_t}(\delta_t^\pi)'c = \Pi_{\Delta_t}(\delta_t^\pi)'\Pi_{\Delta_t}(c)$ and the fact that $\Delta_{t+1} = \Delta_t$, the third inclusion follows from the fact that the orthogonal projection of the intersection of sets must belong to the orthogonal projection of each of the sets and, finally, the last inclusion follows from assumption of the lemma.

The next step is a direct application of Lemma 8. For every cone update period $t$, we focus on the subspace $\Delta_t$ which has dimension $2 \leq p \leq d$. We represent the vectors in $\Delta_t$ under its basis representation by calculating $B_{\Delta_t} c \in \mathbb{R}^p$ for every $c \in \Delta_t^\pi$. Therefore, it suffices to apply Lemma B-1 with $p = p$, $\delta = B_{\Delta_t}\Pi_{\Delta_t}(\delta_t^\pi)$, $\eta = \eta$ and $E(W_t, U_t)$, where the matrices $W_t, U_t$ are already written in the basis representation of $\Delta_t$ by construction.

The final step of this proof is to show that our choice of $\eta$ and $\delta$ satisfy the assumptions of the lemma, meaning that the shallow-cut is sufficiently deep to ensure a volume reduction. Since

$$\Pi_{\Delta_t}(\delta_t)'_{[2:p]} W_t \Pi_{\Delta_t}(\delta_t)_{[2:p]} > \epsilon^2 \implies \frac{1}{\sqrt{\Pi_{\Delta_t}(\delta_t)'_{[2:p]} W_t \Pi_{\Delta_t}(\delta_t)_{[2:p]}}} < \frac{1}{\epsilon},$$

we have that, for $\eta = \epsilon/2d$ and $1 < p \leq d$:

$$\eta \leq \frac{\sqrt{\Pi_{\Delta_t}(\delta_t)'_{[2:p]} W_t \Pi_{\Delta_t}(\delta_t)_{[2:p]}}}{2d} \leq \frac{\sqrt{\Pi_{\Delta_t}(\delta_t)'_{[2:p]} W_t \Pi_{\Delta_t}(\delta_t)_{[2:p]}}}{2(p-1)} = \frac{\sqrt{(B_{\Delta_t}\Pi_{\Delta_t}(\delta_t)_{[2:p]})' W_t B_{\Delta_t}\Pi_{\Delta_t}(\delta_t)_{[2:p]}}}{2(p-1)}$$

The desired inclusion and the reduction in the product of the eigenvalues follows directly from Lemma B-1. □

***Proof of Lemma 9.*** In order to prove the result, we will follow the same strategy as in the proof of Lemma 5. For the boundary cases where $t = t_0^p$ for $2 \leq p \leq d$, the ellipsoidal cone $E(W_t, U_t)$ is constructed by a dimension update period, and we must have that $\lambda_1(W_{t_0^p}) = \lambda_{max}(W_{t_0^p})$.

Suppose $\lambda_1(W_{t_0^p}) > \left(\frac{\epsilon}{10p}\right)^2$. The shallow-cut equivalent of Lemma 4 from Cohen et al. (2020) says that if we update our ellipsoidal cone according to Eq. (13), if $\lambda_1(W_t) \leq \frac{\epsilon^2}{100(p-1)^2}$ and $\delta'_{t,[2:p]} W_t \delta_{t,[2:p]} > \epsilon^2$, then $\lambda_1(W_{t+1}) \geq \lambda_1(W_t)$, i.e., the smallest eigenvalue does not decrease after the update. Since the largest decrease possible in any eigenvalue after applying the ellipsoid method in $\mathbb{R}^{p-1}$ is $(p-1)^2/p^2$ (this is the decrease that happens when the cut is along that particular eigenvector), we have that for all $t_0^p < t < t_0^{p+1}$ (all time periods where the dimensionality of $\Delta_t$ is $p$) that

$$\lambda_1(W_{t+1}) \geq \frac{(p-1)^2}{p^2} \frac{\epsilon^2}{100(p-1)^2} = \left(\frac{\epsilon}{10p}\right)^2.$$

We omit the induction argument here since it mimics the proof of Lemma 5. Hence,

$$\lambda_{p-1}(W_t) \geq \left(\frac{\epsilon}{10p}\right)^2.$$

Moreover, the shrinking factor of Lemma 8 gives us that

$$\prod_{i=1}^{p-1} \lambda_i(W_{t+1}) \leq e^{-1/20(p-1)} \prod_{i=1}^{p-1} \lambda_i(W_t).$$

Using the lower bound for the eigenvalues of $W_t$ gives us that

$$\lambda_{max}(W_t) \leq \left(\frac{10p}{\epsilon}\right)^{2(p-2)} (\lambda_{max}(W_1))^{(p-1)} e^{-\frac{I_T^{\pi,p}}{20(p-1)}},$$

similarly to the proof of step 3 of Theorem 3, we must have that

$$I_T^{\pi,p} \leq 20(p-1)^2 \ln\left(\frac{10p \tan\alpha(E(W_{t_0^p}, U_{t_0^p}))}{\epsilon}\right),$$

where $t_0^p$ is the first time such that the dimension of $\Delta_t = p$ and $E(W_{t_0}, U_{t_0})$ the initial ellipsoidal cone.

If $\lambda_1(W_{t_0^p}) \leq \left(\frac{\epsilon}{10p}\right)^2$, then noting that $\lambda_{max}(W_{t_0^p}) = \lambda_1(W_{t_0^p})$, the same argument from Lemma 5 holds and $I_T^{\pi,p} = 0$, while Lemma 7 ensures a upper bound of $\epsilon$ for each period regret $t$ such that $\dim(\Delta_t) = p$. □

***Proof of Lemma 10.*** The proof is divided in three parts. First we show the inclusion, then we show how to compute the circumcenter, and finally we prove the bound for the uncertainty angle.

**Inclusion.** The definition of the knowledge set implies that for every $c \in C(\mathcal{I}_t)$, $\delta_t^{\pi\prime} c \geq 0$. In particular, this holds for all $i \in \tau(t)$. Moreover, by definition, $\delta_i^{\pi} \in \Delta_t$ for all $i \in \tau(t)$. Hence, for every $c \in C(\mathcal{I}_t)$, we have that $\delta_i^{\pi\prime} c = \delta_i^{\pi\prime} \Pi_{\Delta_t}(c) \geq 0$ for all $i \in \tau(t)$, which implies that $\Pi_{\Delta_t}(C(\mathcal{I}_t)) \subseteq K_t$.

**Computation of circumcenter.** The computation of the circumcenter of $K_t$ is done via algorithm `PolyCenter` (see Algorithm 3). The subspace updating rule implies that $\delta_i^{\pi}$, $i \in \tau(t)$ are linearly independent. Hence, the system of equations described has one and only one solution for each iteration $k$. In addition, we have that $z$ is a vector in the interior of the cone $\{c \in \mathbb{R}^p : \ c = \sum \gamma_i \bar{\delta}_i, \ \gamma_i \geq 0, \ i \in \tau(t)\}$, which is the dual cone of $K_t$ (Boyd and Vandenberghe, 2004). Hence, $c'z > 0$, for all $c \in K_t$ and $K_t \cap \{c \in \mathbb{R}^p : \ c'z \leq 1\}$ is bounded with extreme points given by the rays (not normalized) of $K_t$ and the origin. To see why the quadratic program in `PolyCenter` yields to the solution of the circumcenter, we refer to Seeger and Vidal-Nuñez (2017). The algorithm runs in polynomial time in $p$ since it contains $p$ linear systems with $p$ equations and one quadratic programming formulation.

**Upper bound for the uncertainty angle.**

For any convex, closed and pointed cone $K \subset \mathbb{R}^d$, we define the circumradius of the cone $\mu(K)$ to be equal to sine of its uncertainty angle, i.e., $\mu(K) = \sin \alpha(K)$. In this proof, we will show that

$$\mu(K_{t+1}) \leq \sqrt{1 - \frac{\eta^{2(p-1)}}{p^3}}. \tag{B-8}$$

Once we prove that Eq. (B-8) is true, it follows that

$$\cos \alpha(K_{t+1}) = \sqrt{1 - \mu^2(K_{t+1})} \geq \frac{\eta^{(p-1)}}{p^{3/2}} \geq \frac{\eta^{(d-1)}}{d^{3/2}},$$

where the second inequality follows from the facts that $0 < \eta < 1$ and $p \leq d$, completing the result.

Just as circumradius (and uncertainty angle) are defined by the smallest revolution cone that contains our cone of interest, we also need to define the largest revolution cone that fits within a cone of interest. For any convex, closed and solid cone $K$, we define the inradius of $K$ to be:

$$\rho(K) = \max_{x \in S^d \cap K} \min_{y \in \partial K} \|x - y\|.$$

We denote $x_\rho = \arg\max_{x \in S^d \cap K} \min_{y \in \partial K} \|x - y\|$ as the incenter of $K$, which is analogous to the circumcenter, but referring to the axis of the largest revolution cone inside our cone of interest. The inradius and circumradius of a cone $K$ are dual quantities in the sense that for every closed convex cone, we have that:

$$\mu^2(K) + \rho^2(K^\star) = 1, \tag{B-9}$$

where $K^\star = \{c \in \mathbb{R}^d : \ c'x \geq 0, \ \forall x \in K\}$ denotes the dual cone of $K$ (Henrion and Seeger, 2010, Theorem 1.4). We will develop a lower bound on $\rho(K^\star_{t+1})$ and then obtain through Eq. (B-9) our desired upper bound on $\mu(K_{t+1})$.

The lower bound for $\rho(K^\star_{t+1})$ is obtained by an application of three lemmas that are interesting by its own that we prove in the appendix. The first one, Lemma B-2, shows that the inradius of $K^\star_{t+1}$ can be lower bounded by the ratio of the largest and smallest eigenvalue of the gram-matrix (the square matrix) constructed with it's generators. Suppose $\dim(\Delta_{t+1}) = p \leq d$. We denote $g_i = B_{\Delta_{t+1}} \delta^\pi_{\tau(t+1)}(i) \in \mathbb{R}^p$, for $i = 1, ..., p$, where we used the notation $\delta^\pi_{\tau(t+1)}(i)$ to denote the $i - th$ effective difference that belongs to $\tau(t+1)$. Let $G$ be the matrix such that its columns are given by $g_i$, $i = 1, ..., p$. By construction, the columns of $g_i$ are linearly independent and we have $p$ vectors generating $K^\star_{t+1}$ that lives in a subspace with dimension $p$ and the assumptions of Lemma B-2 are satisfied. The lemma gives us that

$$\rho(K^\star_{t+1}) \geq \frac{1}{\sqrt{p}} \sqrt{\frac{\lambda_{max}(G'G)}{\lambda_{min}(G'G)}}.$$

The second lemma, Lemma B-3, allows us to provide an upper-bound for $\frac{\lambda_{max}(G'G)}{\lambda_{min}(G'G)}$. Note that $G$ is full column rank with unit norm vectors by construction (the change of basis through $B_{\Delta_{t+1}}$ do not affect the norm of the effective differences vectors), and the assumption of the lemma is satisfied. Let $\Pi_{g_{-i}}(.)$ denote the projection operator on the subspace generated by $\{g_1, \ldots, g_p\} \setminus \{g_i\}$. Define

$$\varphi = \min_{i \leq p} \|g_i - \Pi_{g_{-i}}(g_i)\|,$$

and note that $\phi$ is the minimum norm obtained by regressing the column $g_i$ on every other columns. Lemma B-3 implies that

$$\frac{\lambda_{max}(G'G)}{\lambda_{min}(G'G)} \leq \left(\frac{p}{\varphi}\right)^2.$$

The final step in our proof is to show that $\varphi$ cannot be arbitrarily small. In Lemma B-4, we show that if the sequence of the $g_i$'s satisfy $\|g_i - \Pi_{g_1,\ldots,g_{i-1}}(g_i)\|^2 \geq \eta^2$, (which is true due to the subspace updating rule) then it must be the case that $\|g_i - \Pi_{g_{-i}}(g_i)\| \geq \eta^{2(p-i)}$ for $i = 1$ or greater than $\eta^{2(p-i)}$ if $i > 1$. Taking $i = 1$ (or $i = 2$) ensures that $\|g_i - \Pi_{g_{-i}}(g_i)\| \geq \eta^{2(p-1)}$ for every $i$. Then, the combination of the three lemmas establishes that the inradius of $K^\star_{t+1}$ is lower bounded by $\frac{\eta^{p-1}}{p^{3/2}} \geq \frac{\eta^{d-1}}{d^{3/2}}$, concluding the proof of Lemma 10.

□

***Proof of Theorem 4.*** We prove the result in four steps. First, we establish that $\Pi_{\Delta_t}(C(\mathcal{I}_t)) \subseteq E(W_t, U_t)$ for every $t$, and thus we never lose track of the true cost. In a second step, we decompose the cumulative regret as a function of the different kinds of periods in algorithm `ProjectedCones` and provide an upper bound for the one-period regret under each case. In the third step, we upper bound the number of periods that we use cone-updates to obtain our regret bound. In the fourth step, we prove the polynomial runtime.

**Step 1.** We show that, for each time $t$, the set $E(W_t, U_t)$ contains $\Pi_{\Delta_t}(C(\mathcal{I}_t))$. We establish the result by induction on $t$. $E(W_2, U_2)$ trivially contains $\Pi_{\Delta_2}(C(\mathcal{I}_2))$, so the base case is satisfied. We next consider $t \geq 2$. Suppose that $E(W_t, U_t)$ is such that $\Pi_{\Delta_t}(C(\mathcal{I}_t)) \subseteq E(W_t, U_t)$. We then analyze $E(W_{t+1}, U_{t+1})$ as a function of the three situations that can happen at time $t$, no update (low-regret period), cone update, or subspace update.

*No update.* This case is trivial since $C(\mathcal{I}_{t+1}) \subseteq C(\mathcal{I}_t)$, $E(W_{t+1}, U_{t+1}) = E(W_t, U_t)$, and $\Delta_{t+1} = \Delta_t$, thus $\Pi_{\Delta_{t+1}}(C(\mathcal{I}_{t+1})) \subseteq E(W_{t+1}, U_{t+1})$.

*Cone update.* This inclusion follows from Lemma 8.

*Subspace update.* Lemma 10 shows that $\Pi_{\Delta_{t+1}}(C(\mathcal{I}_{t+1}))$ is contained in $K_{t+1}$. Our choices of $W_{t+1}$ and $U_{t+1}$ ensure that $K_{t+1}$ is included in $E(W_{t+1}, U_{t+1})$. The same lemma also shows that the constructed ellipsoidal cone is large enough to contain $K_{t+1}$.

**Step 2.** Lemma 7 shows that we incur at most regret $\epsilon + \eta$ in low-regret periods. Therefore, our total regret from low-regret periods is bounded by $T(\epsilon + \eta)$. For all other periods, we use the trivial regret upper bound of 1. There are at most $d$ subspace-update periods, so the total regret from subspace-update periods is bounded by $d$. Let $I_T^{\pi,p}$ be the number of periods where the subspace has dimension $p$ and we use a cone-update. Bounding the regret of these cone-update periods by 1 as well, we have that for any $c^\star \in S^d$, $f_t \in \mathcal{F}$, and $\mathcal{X}_t \in \mathcal{B}$, the total regret is bounded by:

$$\mathcal{R}_T^\pi\left(c^\star, \vec{\mathcal{X}}_T, \vec{f}_T\right) \leq T(\epsilon + \eta) + d + \sum_{p=2}^{d} I_T^{\pi,p}, \tag{B-10}$$

where the last sum starts from $p = 2$ because there are never cone-updates when $p = 1$.

**Step 3.** Lemma 9 shows that for any $p = 2, ..., d$, we have

$$I_T^{\pi,p} \leq 20(p-1)^2 \ln\left(\frac{10p \tan\alpha(E(W_{t_0}, U_{t_0}))}{\epsilon}\right),$$

where $t_0$ refers to the period where the subspace was increased to $p$. For simplicity, we replace $p-1$ and $p$ with the larger value $d$: $I_T^{\pi,p} \leq 20d^2 \ln\left(\frac{10d \tan\alpha(E(W_{t_0}, U_{t_0}))}{\epsilon}\right)$. At period $t_0$, the subspace update constructs a revolution cone such that $\alpha(E(W_{t_0}, U_{t_0})) = \arccos \eta^{d-1}/d^{3/2}$ (see Lemma 10). Since $\tan(x) \leq 1/\cos(x)$, we have that $\tan\alpha(E(W_{t_0},U_{t_0})) \le d^{3/2}/\eta^{d-1}$. Summing over all $p$:

$$\sum_{p=2}^{d} I_T^{\pi,p} \le 20d^3 \ln\left(\frac{10d^{5/2}}{\epsilon\eta^{d-1}}\right).$$

Plugging the bound above into Eq. (B-10) and selecting $\epsilon = d/T$, $\eta = \epsilon/2d$ leads to:

$$\begin{aligned}
\mathcal{R}_T^{\pi}\left(c^\star,\vec{\mathcal{X}}_T,\vec{f}_T\right) &\le T(\epsilon+\eta) + d + 20d^3\ln\left(\frac{10d^{5/2}}{\epsilon\eta^{d-1}}\right)\\
&= d + 2 + d + 20d^3\ln(5d^{3/2}T^d2^d)\\
&= 2 + 2d + 20d^3\ln 5 + 30d^3\ln(d) + 20d^4\ln(2T) = \mathcal{O}(d^4\ln T).
\end{aligned}$$

**Step 4.** This algorithm runs in polynomial time in $d$ and $T$ because every period's computation is a function only of $d$. Low-regret periods are computationally very cheap. Each cone-update period requires a spectral decomposition which has a running time upper bound of $\mathcal{O}(p^3)$. Each subspace-update period requires the computation of a new circumcenter via algorithm `PolyCenter`, which runs in polynomial time in $d$ (cf. Lemma 10). This completes the proof. □

## B.1 Auxiliary Results

**Lemma B-1** (Robust ellipsoidal cone updates)**.** *Consider a diagonal and positive-definite matrix* $W \in \mathbb{D}_{++}^{p-1}$ *and an orthonormal matrix* $U \in \mathbb{R}^p \times \mathbb{R}^p$ *and define the ellipsoidal cone* $E(W,U) \subset \mathbb{R}^p$*. Fix* $\eta \ge 0$ *and a vector* $\delta \in \mathbb{R}^p$ *such that* $\eta \le \sqrt{\delta_{[2:p]}' W \delta_{[2:p]}}(2(p-1))^{-1}$ *and* $\delta'\hat{c}(W,U) \le 0$*. Define*

$$\bar{\delta} = U^{-1}\delta_{[2:p]}/\|U^{-1}\delta_{[2:p]}\|,\ \beta = -\frac{\eta}{\sqrt{\bar{\delta}'W\bar{\delta}}},\ b = \frac{W\bar{\delta}}{\sqrt{\bar{\delta}'W\bar{\delta}}},\ a = \frac{1+(p-1)\beta}{p}b,$$

*and*

$$N = \frac{(p-1)^2}{(p-1)^2-1}(1-\beta^2)\left(W - \frac{2(1+(p-1)\beta)}{p(1+\beta)}bb'\right),\ M = \begin{pmatrix}1 & a'\\ a & aa'-N\end{pmatrix}.$$

*Let* $V\Lambda V'$ *denote the spectral decomposition of the matrix* $M$*. Then*

$$E(W,U)\cap\{c\in\mathbb{R}^d\ :\ \delta'c \ge -\eta\} \subseteq E(\widetilde{W},\widetilde{U})$$

*for* $\widetilde{U} = UV$*, and* $\widetilde{W}$ *is a diagonal matrix such that* $\widetilde{W}_{i,i} = \lambda_i(N)$*,* $i = 1,\cdots,d-1.$*, where the eigenvalues are in a nonincreasing order. Moreover, if* $\eta = 0$*,* $\prod_{i=1}^{d-1}\lambda_i(\widetilde{W}) \le e^{-1/2(d-1)}\prod_{i=1}^{d-1}\lambda_i(W)$*. Otherwise* $\prod_{i=1}^{d-1}\lambda_i(\widetilde{W}) \le e^{-1/20(d-1)}\prod_{i=1}^{d-1}\lambda_i(W)$*.*

***Proof of Lemma B-1.*** We first consider the case where $U$ is the identity matrix, so that $E(W,U) = E(W)$ and $\bar{\delta} = \delta$. The proof strategy is as follows. In the first step, we show that $E(W) \cap \{c \in \mathbb{R}^d : \ \delta' c \geq 0\}$ is contained in a half-ellipsoidal cone. Second, we show that we can use a variation of the ellipsoid method to update this half-ellipsoidal cone and construct an appropriate updated ellipsoidal cone. In a third step, we apply a theorem from Seeger and Vidal-Nuñez (2017) in order to characterize the obtained ellipsoidal cone in terms of a standard-position cone and its orthonormal rotation.

**Step 1.** Let us consider the intersection of $E(W)$ and a hyperplane characterized by its circumcenter $\hat{c}(E(W)) = e_1$. We denote such a hyperplane as $H = \{c \in \mathbb{R}^d : \ e_1' c = 1\}$. Since $c_{[1]} > 0$ and $E(W)$ is a cone, we can scale any element $c \in E(W)$ such that the scaled vector lies in the intersection. Moreover, the intersection will be an ellipsoid given by

$$E_H(W) = \{c \in \mathbb{R}_+ \times \mathbb{R}^{d-1} : \ c'_{[2:d]} W^{-1} c_{[2:d]} \leq 1, \ c_{[1]} = 1\},$$

which is precisely the equation of an ellipsoid in $\mathbb{R}^{d-1}$ that lives in a $(d-1)$- dimensional subspace of $\mathbb{R}^d$. The ellipsoid $E_H(W)$ defined above was obtained by a specific type of projection known in the literature as the perspective projection of $E(W)$ onto the hyperplane $H$.

Let us define also $A = E(W) \cap \{c \in \mathbb{R}^d : \ \delta' c \geq -\eta\}$. Next, consider some $c \neq 0 \in A$. Since $A \subseteq E(W)$, we must have $c_{[1]} > 0$. Moreover, $A$ is a cone since it is the intersection of two cones. Then, $c/c_{[1]} \in A$ and $c/c_{[1]} \in H$. Hence,

$$\begin{aligned}
c/c_{[1]} &\in A \cap H \\
&= E(W) \cap H \cap \{c \in \mathbb{R}^p : \ \delta' c \geq -\eta\} \\
&= E_H(W) \cap \{c \in \mathbb{R}^d : \ \delta' c \geq -\eta\} \\
&\overset{(a)}{=} E_H(W) \cap \{c \in \mathbb{R}^d : \ \delta_{[1]} + {\delta_{[2:d]}}' c_{[2:d]} \geq -\eta\} \\
&\overset{(b)}{\subseteq} E_H(W) \cap \{c \in \mathbb{R}^d : \ \delta_{[1]} + {\delta_{[2:d]}}' c_{[2:d]} \geq -\eta + \delta' \hat{c}(E(W))\} \\
&\overset{(c)}{\subseteq} E_H(W) \cap \{c \in \mathbb{R}^d : \ {\delta_{[2:d]}}' c_{[2:d]} \geq -\eta\},
\end{aligned}$$

where $(a)$ follows from the fact that $E_H(W) \subset H$ so $c_{[1]} = 1$. $(b)$ follows from the fact that $\hat{c}(E(W))'\delta \leq 0$ by assumption and $(c)$ follows from the fact that $= \hat{c}(E(W))'\delta = e_1'\delta = \delta_{[1]} \leq 0$. The set of the last equation is precisely the ellipsoid $E_H(W)$ with a shallow-cut, and hence is a description of a half-ellipsoid.

**Step 2.** We now use the ellipsoid method update to replace this half-ellipsoid with its own Löwner-John ellipsoid. The definitions of $a$, $N$ and $\beta$ are precisely the ones for an ellipsoid update with shallow-cut (See Eq. (3.1.16) and Eq. (3.1.17) of Grötschel et al. (1993)). Therefore,

$$E_H(W) \cap \{c \in \mathbb{R}^d : \ \bar{\delta}'_{[2:d]} c_{[2:d]} \geq -\eta\} \subseteq \{c \in \mathbb{R}^d : \ (c_{[2:d]} - a)' N^{-1} (c_{[2:d]} - a) \leq 1, \ c_{[1]} = 1\}.$$

Furthermore, since $c_{[1]} > 0$, we have that

$$c/c_{[1]} \in \{c \in \mathbb{R}^d : \ (c_{[2:d]} - a)' N^{-1} (c_{[2:d]} - a) \leq 1\}$$

if and only if

$$c \in \{c \in \mathbb{R}^d : \ (c_{[2:d]} - c_{[1]} a)' N^{-1} (c_{[2:d]} - c_{[1]} a) \leq c_{[1]}^2\}.$$

Hence, if $c \in A$, then we must have that $c \in \{c \in \mathbb{R}^d : \ (c_{[2:d]} - c_{[1]} a)' N^{-1} (c_{[2:d]} - c_{[1]} a) \leq c_{[1]}^2\}$.

We finish the second step by showing the contraction in the product of the eigenvalues of the matrix $N$ (or equivalently, its volume).

Note that if $\eta = 0$, the update equations reduces to the standard update equations of the ellipsoid method, then the standard volume reduction of the holds and we get $\prod_{i=1}^{p-1} \lambda_i(N) \leq \prod_{i=1}^{p-1} \lambda_i(W) e^{1/2(p-1)}$.

If $\eta > 0$, we need to argue that the shallow cuts are still sufficiently deep to induce a reduction in the product of the eigenvalues first need to ensure that $\beta \geq 1/(p-1)$ in order to have $\prod_{i=1}^{p-1} \lambda_i(N) \leq \prod_{i=1}^{p-1} \lambda_i(W) e^{(1+\beta(p-1))/5(p-1)}$ (See Eq. (3.3.21) of Grötschel et al. (1993)). Since in this case

$$0 < \eta \leq \frac{\sqrt{\delta'_{[2:p]} W \delta_{[2:p]}}}{2(p-1)}$$

we have that,

$$\beta = -\frac{\eta}{\sqrt{\delta'_{[2:p]} W \delta_{[2:p]}}} \geq -\frac{1}{2(p-1)},$$

substituting the lower bound of $\beta$ above leads to $\prod_{i=1}^{p-1} \lambda_i(N) \leq \prod_{i=1}^{p-1} \lambda_i(W) e^{1/20(p-1)}$.

**Step 3.** We have now constructed an ellipsoidal cone $\{c \in \mathbb{R}^d : \ (c_{[2:d]} - c_{[1]} a)' N^{-1} (c_{[2:d]} - c_{[1]} a) \leq c_{[1]}^2\}$ that contains the half ellipsoidal cone of interest, which is a cone not in standard position. However, instead of having the set described by a rotation orthonormal basis, we have it described via a translation of the center of the ellipsoid at $H$. In the remainder of this proof, we show how to construct a representation of this ellipsoidal cone that is consistent with our Definition 4. That is, we need to find a mapping from the parameters $N$ and $a$ to the matrices $W$ and $U$.

To find this mapping, we apply a theorem from (Seeger and Vidal-Nuñez, 2017, Theorem 4.4 page 296). The theorem states that for

$$M = \begin{bmatrix} 1 & a' \\ a & aa' - N \end{bmatrix},$$

the matrix $M$ is invertible and the spectral decomposition $M = V \Lambda V'$ allows us to write the ellipsoidal cone $E(\hat{W}, \hat{U})$ that is identical to $\{c \in \mathbb{R}^d : \ (c_{[2:d]} - c_{[1]} a)' N^{-1} (c_{[2:d]} - c_{[1]} a) \leq c_{[1]}^2\}$, but in standard representation. This ellipsoidal cone is given by $\hat{U} = V$ and $\hat{W}$ equal to the diagonal matrix with diagonal entries $\hat{W}_{ii}$ given by $-\lambda_{i+1}(M)/\lambda_1(M)$, for $i = 1, \cdots, d-1$, where the eigenvalues of $M$ are in nonincreasing order.

For our ellipsoidal update, we will use the same orthonormal rotation as the one produced by Seeger and Vidal-Nuñez's theorem: $\widetilde{U} = \hat{U}$. At this point, we could declare the proof done if we had also defined in the algorithm $\widetilde{W}$ to be equal to $\hat{W}$. However, in order to facilitate our analysis, we use a different choice of $\widetilde{W}$ by setting $\widetilde{W}_{ii} = \lambda_{p-i}(N)$ for all $i = 1, ..., p-1$. Hence, in order to conclude the proof, we still need to show that our choice of $\widetilde{W}$ has eigenvalues that are at least as large as $\hat{W}$. To be specific, we need to prove that for all $i = 1, ..., p-1$, $\lambda_i(\hat{W}) = -\lambda_{i+1}(M)/\lambda_1(M) \leq \lambda_i(N) = \lambda_i(\widetilde{W})$. Our definitions of $M$, $N$ and $a$ satisfy:

$$M = \begin{bmatrix} 0 & 0 \\ 0 & -N \end{bmatrix} + \begin{bmatrix} 1 \\ a \end{bmatrix} \begin{bmatrix} 1 \\ a \end{bmatrix}'.$$

Since the matrix $N$ is obtained by one iteration of the ellipsoid method, we must have that $N$ has strictly positive eigenvalues. Hence, the interlacing theorem (Hwang, 2004) implies that $M$ has one positive eigenvalue and $d-1$ negative eigenvalues.

Furthermore, we have that $\lambda_{max}(M) = \lambda_1(M) = \sup_{c \in S^d} c'Mc$. Since $c = e_1$ satisfies $c'Mc = 1$, we must have that $\lambda_1(M) \geq 1$. Finally, another application of the interlacing theorem for $M$ gives us that for $i = 1, \cdots, d-1$, satisfies $\lambda_{i+1}(M) \geq \lambda_i(-N)$, and we get that

$$-\lambda_{i+1}(M)/\lambda_1(M) \leq -\lambda_{i+1}(M) \leq -\lambda_i(N) = \lambda_{p-i}(N) = \lambda_i(\widetilde{W}), \ i = 1, \cdots, p-1.$$

Hence, the choice of $\widetilde{W}$ in algorithm `ConeUpdate` is at least as large as necessary, which concludes the proof for the case where $U$ is the identity matrix.

For the general case (when $U$ is not the identity matrix), it suffices to rotate $\delta$ by $U^{-1}$, which is equivalent to analyze the problem under the basis representation given by the rows of $U$. After that, we return to the canonical basis by rotating $V$ by $U$, which leads to $\widetilde{U} = UV$. □

**Lemma B-2** (Inradius of simplicial cones)**.** *Fix $p \geq 2$ and let $K \subset \mathbb{R}^p$ denote a simplicial cone, i.e., it can be written as $K = \{c \in \mathbb{R}^p : \ c = \sum_{i=1}^p \alpha_i g_i, \ \alpha_i \geq 0\}$, where the $g_i$'s are linearly independent unit norm vectors. Define $G \in \mathbb{R}^p \times \mathbb{R}^p$ to be the matrix where the columns are the generators $g_i$, $i = 1, \cdots, p$. We have that the inradius $\rho(K) := \max_{x \in S^d \cap K} \min_{y \in \partial K} \|x - y\|$ is lower bounded as follows*

$$\rho(K) \geq \frac{1}{\sqrt{p}} \sqrt{\frac{\lambda_{min}(G'G)}{\lambda_{max}(G'G)}}.$$

***Proof of Lemma B-2.*** We denote $x_\rho = \arg\max_{x \in S^d \cap K} \min_{y \in \partial K} \|x - y\|$ as the incenter of $K$, which is analogous to the circumcenter, but referring to the axis of the largest revolution cone inside our cone of interest. Let $P$ denote the nonnegative orthant in $\mathbb{R}^p$. One can check that $\rho(P) = \frac{1}{\sqrt{p}}$ and the incenter of $P$ is given by $x_\rho = \mathbf{1}\frac{1}{\sqrt{p}}$ (see, for instance, Henrion and Seeger (2010)). We have

$$
\begin{aligned}
\rho(K) &\overset{(a)}{=} \max_{\pi\in S^d\cap K}\min_{\delta\in\partial(K)} \|\pi-\delta\| \\
&\overset{(b)}{=} \max_{\|Gx\|=1, x\in P}\min_{Gy\in\partial(K)} \|Gx-Gy\| \\
&\overset{(c)}{=} \max_{\|Gx\|=1, x\in P}\min_{y\in\partial(P)} \|Gx-Gy\| \\
&= \max_{\|Gx\|=1, x\in P}\min_{y\in\partial(P)} \|G(x-y)\| \\
&= \max_{\|Gx\|=1, x\in P}\min_{y\in\partial(P)} \sqrt{(x-y)'G'G(x-y)} \\
&\geq \max_{\|Gx\|=1, x\in P}\min_{y\in\partial(P)} \|x-y\|\sqrt{\lambda_{min}(G'G)}
\end{aligned}
$$

$(a)$ follows from the definition of inradius, $(b)$ follows from the fact that $\pi\in K \iff \pi = Gx$ for some $x\in P$ since $K$ is defined by the nonnegative linear combination of the columns of $G$, and by the fact that $G$ is full rank by assumption, so $y = G^{-1}\delta$ is well defined. $(c)$ holds since $G$ is linear and establishes a bijection between the generators of $P$ and $K$. The inequality follows from the fact that $G'G$ is positive definite (which follows from $G$ being full column-rank).

Moreover, since $G$ has full column-rank, we have that $Gx = 0$ if and only if $x = 0$, implying that $x = \frac{1}{\|G\mathbf{1}\|}\mathbf{1}$ has only positive components and is well defined since $\|G\mathbf{1}\| \neq 0$. Thus, $x = \frac{1}{\|G\mathbf{1}\|}\mathbf{1}$ is feasible for the maximization problem presented above and we have

$$
\begin{aligned}
\rho(K) &\overset{(a)}{\geq} \sqrt{\lambda_{min}(G'G)}\min_{y\in\partial(P)}\left\|\frac{1}{\|G\mathbf{1}\|}\mathbf{1}-y\right\| \\
&= \frac{\sqrt{p}}{\|G\mathbf{1}\|}\sqrt{\lambda_{min}(G'G)}\min_{y\in\partial(P)}\left\|\frac{1}{\sqrt{p}}\mathbf{1}-y\frac{\|G\mathbf{1}\|}{\sqrt{p}}\right\| \\
&\overset{(b)}{=} \frac{\sqrt{p}}{\|G\mathbf{1}\|}\sqrt{\lambda_{min}(G'G)}\min_{y\in\partial(P)}\left\|\frac{1}{\sqrt{p}}\mathbf{1}-y\right\| \\
&\overset{(c)}{\geq} \frac{\sqrt{p}}{\|G\mathbf{1}\|}\sqrt{\lambda_{min}(G'G)}\frac{1}{\sqrt{p}} \\
&= \frac{1}{\|G\mathbf{1}\|}\sqrt{\lambda_{min}(G'G)},
\end{aligned}
$$

where $(a)$ follows from the fact that $\frac{1}{\|G\mathbf{1}\|}$ is feasible for the maximization problem, $(b)$ follows from the fact that we can always scale by a positive constant our choice for $y$ in the minimization problem, and $(c)$ follows from the fact that the minimization problem over $y$ is lower bounded by the definition of the inradius of the nonnegative orthant.

In addition, we have $\|G\mathbf{1}\|^2 = \mathbf{1}'G'G\mathbf{1} \leq \|\mathbf{1}\|^2\lambda_{max}(G'G) = p\lambda_{max}(G'G)$. In turn, we have

$$\rho(K) \quad \geq \quad \frac{1}{\sqrt{p}}\sqrt{\frac{\lambda_{min}(G'G)}{\lambda_{max}(G'G)}}.$$

□

**Lemma B-3** (Condition number of Gram-Matrices)**.** *Let $p \geq 2$ and let $G$ be a full-column rank matrix with unit norm columns denoted by $g_i$, $i = 1, \cdots, p$. Let $\Pi_{g_{-i}}(\cdot)$ denote the projection operator on the subspace generated by $\{g_1, \ldots, g_p\} \setminus \{g_i\}$. Let*

$$\varphi = \min\{\|g_i - \Pi_{g_{-i}}(g_i)\| : i = 2, ..., p\}. \tag{B-11}$$

*Then*

$$\frac{\lambda_{max}(G'G)}{\lambda_{min}(G'G)} \leq \left(\frac{p}{\varphi}\right)^2.$$

***Proof of Lemma B-3.*** We start from an upper bound on the ratio of largest to smallest eigenvalues of $G'G$ in terms of its trace and the trace of its inverse. Recall that $G$ is full column-rank, which implies that $G'G$ is positive definite and we have that

$$\frac{\lambda_{max}(G'G)}{\lambda_{min}(G'G)} = \lambda_{max}(G'G)\lambda_{max}((G'G)^{-1}) \leq \mathrm{tr}(G'G)\mathrm{tr}((G'G)^{-1}), \tag{B-12}$$

where the last inequality follows from the fact that $G'G$ and $(G'G)^{-1}$ are positive definite and all eigenvalues are nonnegative.

Since the diagonal elements of $G'G$ are equal to one, $\mathrm{tr}(G'G) = p$. Next, we establish an upper bound on $\mathrm{tr}((G'G)^{-1})$. In order to do that, we use a fact about the inverse of correlation matrices, commonly denoted as precision matrices. Note that $G'G$ is a correlation matrix because all the columns of $G$ have norm one and we can consider that each $g_i$ is an "observation" of some data in $\mathbb{R}^p$. For the inverse of correlation matrices, we have that the diagonal elements $(a_{ii})$ of $(G'G)^{-1}$ satisfy

$$R_i^2 = 1 - \frac{1}{a_{ii}}, \tag{B-13}$$

where $R_i^2$ is the coefficient of determination of the linear regression problem of the column $g_i$ onto the other columns $g_j$, $j \neq i$ (Raveh (1985)). Note that since the $g_i$'s have unit norm, $R_i^2 = \|\Pi_{g_{-i}}(g_i)\|^2 \leq 1$ and it follows from Eq. (B-13) that

$$a_{ii} = \frac{1}{1 - \|\Pi_{g_{-i}}(g_i)\|^2} \leq \frac{1}{1 - \max_{i=1,\ldots,p}\|\Pi_{g_{-i}}(g_i)\|^2}.$$

Note that $\|g_i\|^2 = \|\Pi_{g_{-i}}(g_i)\|^2 + \|g_i - \Pi_{g_{-i}}(g_i)\|^2$, and using the fact that $\|g_i\| = 1$ and (B-11), we have

$$\max_{i=1,\ldots,p} \|\Pi_{g_{-i}}(g_i)\|^2 \leq 1 - \varphi^2.$$

In turn, we deduce that for $i = 1, ...p$

$$a_{ii} \leq \frac{1}{\varphi^2},$$

which implies that $\mathrm{tr}((G'G)^{-1}) = \sum_i a_{ii} \leq p/\varphi^2$. Hence, returning to Eq. (B-12), we have

$$\frac{\lambda_{max}(G'G)}{\lambda_{min}(G'G)} \leq \left(\frac{p}{\varphi}\right)^2,$$

which concludes the proof. □

**Lemma B-4** (Residuals)**.** *Let $g_i \in \mathbb{R}^p$, be a sequence of $p$ unit norm vectors such that*

$$\|g_i - \Pi_{g_1,\cdots,g_{i-1}}(g_i)\|^2 \geq \eta^2, \quad i = 2, \cdots, p, \tag{B-14}$$

*where $\Pi_{g_1,\cdots,g_{i-1}}(\cdot)$ denote the projection operator on the subspace generated by $g_1, \cdots, g_{i-1}$. Then the residuals of the projection of an arbitrary vector on the subspace generated by all other vectors can be lower bounded as follows*

$$\|g_i - \Pi_{g_{-i}}(g_i)\|^2 \geq \eta^{2(p-i)}(\eta^2)^{\mathbf{1}\{i>1\}}, \quad i = 1, ..., p, \tag{B-15}$$

*where $\Pi_{g_{-i}}(\cdot)$ denotes the projection operator on the subspace generated by $\{g_1, \ldots, g_p\} \setminus \{g_i\}$.*

***Proof of Lemma B-4.*** In order to understand how close $\|g_i - \Pi_{g_{-i}}(g_i)\|^2$ can be to zero, we need to understand how much of the residual $\|g_i - \Pi_{g_1,\ldots,g_{i-1}}(g_i)\|^2$ can be reduced when adding the new vectors $g_j$'s, for $j = i+1, .., p$. The more the vectors $g_j$'s can explain, the closest $\|g_i - \Pi_{g_{-i}}(g_i)\|^2$ gets to zero. We will leverage that the vectors $g_j$'s satisfy Eq. (B-14) to establish that there is a limit for how small $\|g_i - \Pi_{g_{-i}}(g_i)\|^2$ can be.

Two key properties that allow us to establish the result can be derived from the Frisch-Waugh-Lovell Theorem (Lovell, 2008). Intuitively, the theorem states that when we add a new regressor to a linear regression model, it suffices to consider only the component of the new regressor that is orthogonal to the linear subspace of the regressors already present. We will use the following properties implied by the theorem:

- Projection decomposition: for $i = 1, ..., p$ and $k = 0, ..., p - i$,

$$\Pi_{g_1,\ldots,g_{i+k}}(g_{i+k+1}) = \Pi_{g_1,\ldots,g_{i-1},g_{i+1},\ldots,g_{i+k}}(g_{i+k+1}) + \Pi_{(g_i - \Pi_{g_1,\ldots,g_{i-1},g_{i+1},\ldots,g_{i+k}}(g_i))}(g_{i+k+1}), \tag{B-16}$$

- $R^2$ decomposition: for $i = 1, ..., p$ and $k = 0, ..., p-i$,

$$\|\Pi_{g_1,...,g_{i+k}}(g_{i+k+1})\|^2 = \|\Pi_{g_1,...,g_{i-1},g_{i+1},...,g_{i+k}}(g_{i+k+1})\|^2 + \|\Pi_{(g_i - \Pi_{g_1,...,g_{i-1},g_{i+1},...,g_{i+k}}(g_i))}(g_{i+k+1})\|^2. \quad \text{(B-17)}$$

The first property states that we can decompose the projection of $g_{i+1}$ in two parts. First, we consider the projection of $g_{i+1}$ onto the subspace without $g_i$. Second, we consider the projection of $g_{i+1}$ on the component of $g_i$ that is orthogonal to the subspace considered in the previous step, i.e., the projection of $g_{i+1}$ onto the vector $g_i - \Pi_{g_1,...,g_{i-1}}(g_i)$. For the second property, since $g_i - \Pi_{g_1,...,g_{i-1}}(g_{i+1})$ is orthogonal to the vectors $g_1, ..., g_{i-1}$, we have a situation in which the triangle equality holds and we can decompose the coefficient of determination $R^2$ in the contribution of each orthogonal set of regressors.

Note that when $i = p$, the result follows from the assumption (cf. Eq.(B-14)). Next, fix $i$ in $\{1, ..., p-1\}$ and define the following sequence of residuals:

$$r_i^k = g_i - \Pi_{g_1,...,g_{i-1},g_{i+1},...,g_{i+k}}(g_i), \quad k = 0, ..., p-i. \quad \text{(B-18)}$$

When $i = 1$ and $k = 0$ we will define by convention $r_1^0 = g_1$.

We will show that the sequence of residuals satisfies $\|r_i^{k+1}\|^2 \geq \|r_i^k\|^2\eta^2$. This recursive relationship allows us to write

$$\|r_i^{k+1}\|^2 \geq \|r_i^0\|^2\eta^{2(k+1)},$$

Moreover, $\|r_i^0\|^2 \geq \eta^2$ for $i > 1$ (cf. Eq. (B-14)), and $\|r_1^0\|^2 = \|g_1\|^2 = 1$, which implies that for $k = 0, ..., p-i$,

$$\|r_i^k\| \geq \begin{cases} \eta^{2k}, & i = 1 \\ \eta^{2(k+1)}, & 2 \leq i \leq p-1, \end{cases}$$

Taking $k = p-i$ leads to (B-15). Next, we establish that $\|r_i^{k+1}\|^2 \geq \|r_i^0\|^2\eta^{2(k+1)}$ holds.

We define the following vectors

$$\nu_i^k = g_{i+k} - \Pi_{g_1,...,g_{i-1},g_{i+1},...,g_{i+k-1}}(g_{i+k}), \quad k = 0, ..., p-i.$$

Note that Eq. (B-17) implies that $\nu_i^k$ is the vector that effectively contributes to reduce the residual $r_i^{k-1}$ when $g_{i+k}$ is used to explain $g_i$ and $\{g_1, ..., g_{i+k-1}\} \setminus \{g_i\}$ were already considered. Eq. (B-17) gives us that

$$\|\Pi_{g_1,...,g_{i-1},g_{i+1},...,g_{i+k+1}}(g_i)\|^2 = \|\Pi_{g_1,...,g_{i-1},g_{i+1},...,g_{i+k}}(g_i)\|^2 + \|\Pi_{\nu_i^{k+1}}(g_i)\|^2,$$

and since $g_i$ has unit norm, by construction, we have

$$\begin{aligned}\|\Pi_{g_1,\ldots,g_{i-1},g_{i+1},\ldots,g_{i+k+1}}(g_i)\|^2 &= 1-\|r_i^{k+1}\|^2\\ \|\Pi_{g_1,\ldots,g_{i-1},g_{i+1},\ldots,g_{i+k}}(g_i)\|^2 &= 1-\|r_i^{k}\|^2,\end{aligned}$$

and hence the following recursive relationship

$$\|r_i^{k+1}\|^2 = \|r_i^k\|^2 - \|\Pi_{\nu_i^{k+1}}(g_i)\|^2. \tag{B-19}$$

Next, we upper bound $\|\Pi_{\nu_i^{k+1}}(g_i)\|^2$. We have that

$$\|\Pi_{\nu_i^{k+1}}(g_i)\|^2 = \frac{(g_i'\nu_i^{k+1})^2}{\|\nu_i^{k+1}\|^2} \overset{(a)}{=} \frac{\left((\Pi_{g_1,\ldots,g_{i-1},g_{i+1},\ldots,g_{i+k}}(g_i)+r_i^k)'\nu_i^{k+1}\right)^2}{\|\nu_i^{k+1}\|^2} \overset{(b)}{=} \frac{\left(r_i^{k\prime}\nu_i^{k+1}\right)^2}{\|\nu_i^{k+1}\|^2}, \tag{B-20}$$

where $(a)$ follows from the definition of $r_i^k$ (Eq. (B-18)), and $(b)$ follows from the fact that $\nu_i^{k+1}$ and $\Pi_{g_1,\ldots,g_{i-1},g_{i+1},\ldots,g_{i+k}}(g_i)$ are orthogonal since we removed the projection of $g_{i+1}$ onto the subspace generated by $\{g_1,\ldots,g_{i+k}\}\setminus\{g_i\}$ when constructing $\nu_i^{k+1}$. Next we develop an alternative representation for $\nu_i^{k+1}$. We have

$$\begin{aligned}\nu_i^{k+1} &= g_{i+k+1}-\Pi_{g_1,\ldots,g_{i-1},g_{i+1},\ldots,g_{i+k}}(g_{i+k+1})\\ &\overset{(a)}{=} \Pi_{g_1,\ldots,g_{i+k}}(g_{i+k+1})+r_{i+k+1}^0-\Pi_{g_1,\ldots,g_{i-1},g_{i+1},\ldots,g_{i+k}}(g_{i+k+1})\\ &\overset{(b)}{=} \Pi_{g_1,\ldots,g_{i-1},g_{i+1},\ldots,g_{i+k}}(g_{i+k+1})+\Pi_{\beta_i^k}(g_{i+k+1})+r_{i+k+1}^0-\Pi_{g_1,\ldots,g_{i-1},g_{i+1},\ldots,g_{i+k}}(g_{i+k+1})\\ &= \Pi_{\beta_i^k}(g_{i+k+1})+r_{i+k+1}^0,\end{aligned}$$

where (a) follows from the definition of $r_{i+k+1}^0$ and in (b), we have used Eq. (B-16) and defined $\beta_i^k = g_i-\Pi_{g_1,\ldots,g_{i-1},g_{i+1},\ldots,g_{i+k}}(g_i)$. Using this new representation for $\nu_i^{k+1}$ in Eq. (B-20) yields

$$\begin{aligned}\left(r_i^{k\prime}\nu_i^{k+1}\right)^2 &= \left(r_i^{k\prime}(\Pi_{\beta_i^k}(g_{i+k+1})+r_{i+k+1}^0)\right)^2\\ &\overset{(a)}{=} \left(r_i^{k\prime}\Pi_{\beta_i^k}(g_{i+k+1})\right)^2\\ &\overset{(b)}{\le} \|r_i^k\|^2\|\Pi_{\beta_i^k}(g_{i+k+1})\Big\|^2,\end{aligned} \tag{B-21}$$

where $(a)$ follows from the fact that $r_{i+k+1}^0$ is orthogonal to the subspace generated by $g_1,\ldots,g_{i+k}$, to which $r_i^k$ belongs and (b) follows from the Cauchy-Schwarz inequality. In addition, we have that $\Pi_{\nu_i^k}(g_{i+k+1})$ and $r_{i+k+1}^0$ are orthogonal since $r_{i+k+1}^0$ is by definition orthogonal to the subspace generated by $g_1,\ldots,g_{i+k}$, which $\Pi_{\beta_i^k}(g_{i+k+1})$ belongs to. Therefore, we have that

$$\|\nu_i^{k+1}\|^2 = \|\Pi_{\beta_i^k}(g_{i+k+1}) + r_{i+k+1}^0\|^2 = \|\Pi_{\beta_i^k}(g_{i+k+1})\|^2 + \|r_{i+k+1}^0\|^2,$$

which implies that

$$\frac{\|\Pi_{\beta_i^k}(g_{i+k+1})\|^2}{\|\nu_i^{k+1}\|^2} = 1 - \frac{\|r_{i+k+1}^0\|^2}{\|\nu_i^{k+1}\|^2} \leq 1 - \|r_{i+k+1}^0\|^2,$$

where the inequality follows from the fact that $\nu_i^{k+1}$ is a projection of a unit vector,so $\|\nu_i^{k+1}\|^2 \leq 1$.

Returning to Eq. (B-19) and Eq. (B-20), we have established that

$$\|r_i^{k+1}\|^2 \geq \|r_i^k\|^2 - \|r_i^k\|^2(1 - \|r_{i+k+1}^0\|^2) = \|r_i^k\|^2\|r_{i+k+1}^0\|^2 \geq \|r_i^k\|^2\eta^2,$$

where the last inequality follows from assumption (cf. Eq. (B-14)). This completes the proof.

□

# C Additional Figures

We provide here additional analysis for the simulation results. For the cases depicted in Figures 6a, 6b, 7a and 7b, we depict the results in log-log scale in Figures 10a, 10b, 11a and 11b, respectively. A dependence of the form $T^\alpha$ would lead to a linear relationship on this graph. Each graph depicts for reference the curve $T^{1/2}$. We can see that our algorithms appear to achieve logarithmic regret when the time horizon is sufficiently long. It is worth noticing that the transition regime (before achieving the logarithmic regret rate) can be affected by the dimension of the problem, which is also observed for the competing methods. Interestingly, for both OGD and EWU, we observe over the numerics that these seem to achieve regret rates better than $O(\sqrt{T})$ for the stochastic case. Whereas the theoretical guarantees for these methods in the adversarial case are currently of order $O(\sqrt{T})$, the numerics highlight an interesting theme for future research: delineating whether or not OGD and EWU can achieve better regret bounds in non-adversarial environments.

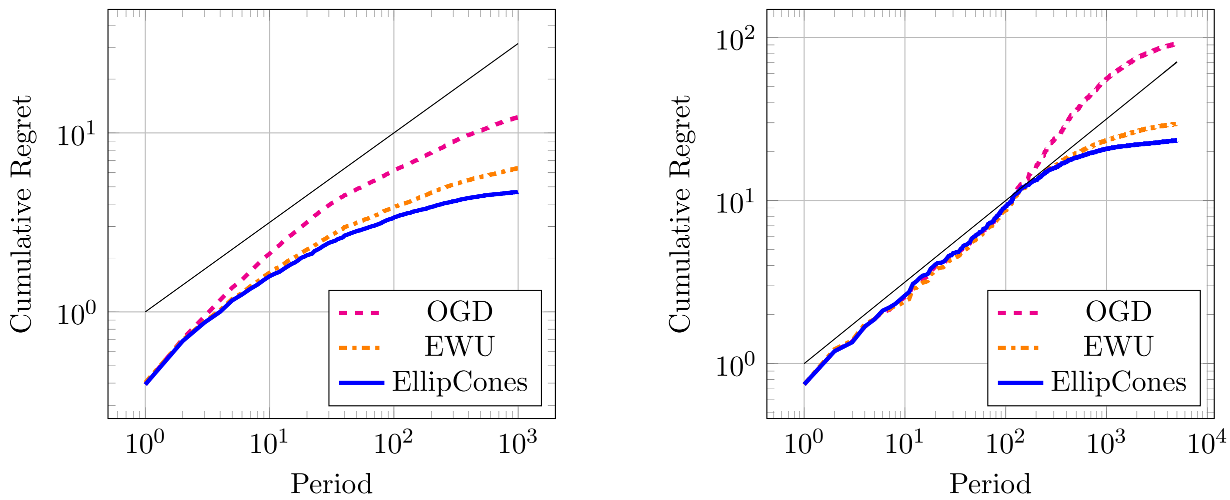


(a) $d = 10$ and $T = 1000$.

(b) $d = 25$ and $T = 5000$

Figure 10: Average cumulative regret over 50 simulations for EWU, OGD and the `EllipsoidalCones` algorithms for the pointed case in loglog scale.

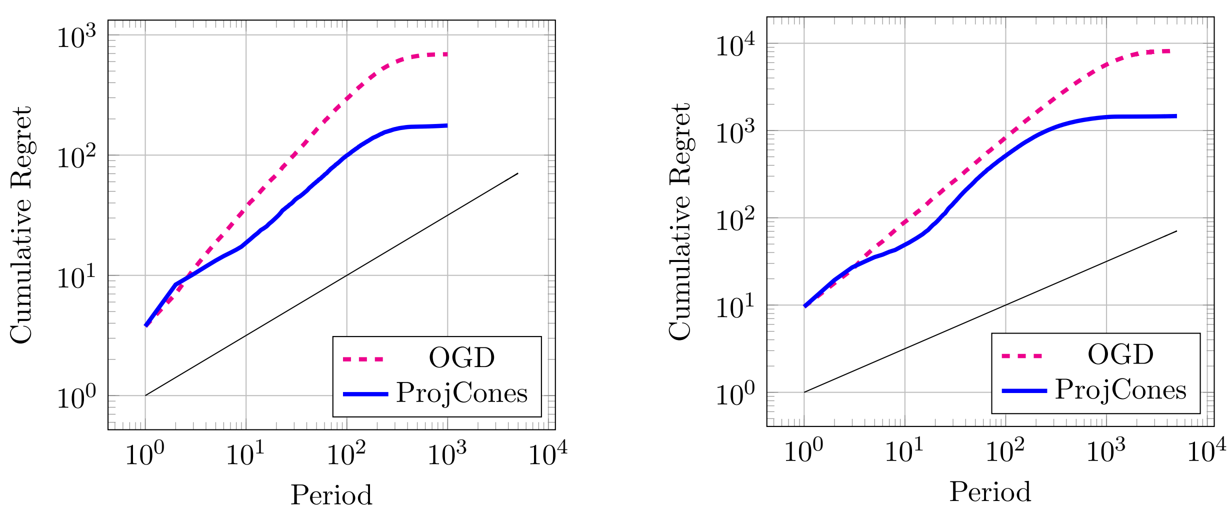


(a) $d = 10$ and $T = 1000$.

(b) $d = 25$ and $T = 5000$.

Figure 11: Average cumulative regret over 50 simulations for OGD and the `ProjectedCones` algorithms for the general case in loglog scale.