



The price of unfairness in linear bandits with biased feedback

Solenne Gaucher, Alexandra Carpentier, Christophe Giraud

► To cite this version:

Solenne Gaucher, Alexandra Carpentier, Christophe Giraud. The price of unfairness in linear bandits with biased feedback. 2022. hal-03611628v1

HAL Id: hal-03611628

<https://hal.science/hal-03611628v1>

Preprint submitted on 17 Mar 2022 (v1), last revised 29 Nov 2022 (v3)

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

The price of unfairness in linear bandits with biased feedback

Solenne Gaucher ^{*1}, Alexandra Carpentier², and Christophe Giraud¹

¹Laboratoire de Mathématiques d’Orsay, Université Paris-Saclay

²University of Potsdam

March 17, 2022

Abstract

Artificial intelligence is increasingly used in a wide range of decision making scenarios with higher and higher stakes. At the same time, recent work has highlighted that these algorithms can be dangerously biased, and that their results often need to be corrected to avoid leading to unfair decisions. In this paper, we study the problem of sequential decision making with biased linear bandit feedback. At each round, a player selects an action described by a covariate and by a sensitive attribute. She receives a reward corresponding to the covariates of the action that she has chosen, but only observe a biased evaluation of this reward, where the bias depends on the sensitive attribute. To tackle this problem, we design a Fair Phased Elimination algorithm. We establish an upper bound on its worst-case regret, showing that it is smaller than $C\kappa_*^{1/3} \log(T)^{1/3} T^{2/3}$, where C is a numerical constant and κ_* an explicit geometrical constant characterizing the difficulty of bias estimation. The worst case regret is higher than the $dT^{1/2} \log(T)$ regret rate obtained under unbiased feedback. We show that this rate cannot be improved for all instances : we obtain lower bounds on the worst-case regret for some sets of actions showing that this rate is tight up to a sub-logarithmic factor. We also obtain gap-dependent upper bounds on the regret, and establish matching lower bounds for some problem instance. Interestingly, the gap-dependent rates reveal the existence of non-trivial instances where the problem is no more difficult than its unbiased counterpart.

1 Introduction

Artificial intelligence is increasingly used in a wide range of decision making scenario with higher and higher stakes, with application in online advertisement [Perlich et al., 2014], credit [Byanjankar et al., 2015], health care [Fauw et al., 2018], education [Papamitsiou and Economides, 2014] and job interviews [Raghavan et al., 2020], in the hope of improving accuracy and efficiency. However, recent works have shown that the decisions made by algorithms can be dangerously biased against certain categories of people, and have endeavored to mitigate this behavior [Köchling and Wehner, 2020, Fuster et al., 2022, Chawla and Jagadeesan, 2021, Mehrabi et al., 2021]. Studies have shown that the main cause of algorithmic unfairness is the presence of bias in the training set (see [Mehrabi et al., 2021] for a review of the causes of algorithmic bias). This remark led to the development of three types of methods aiming to guarantee the fairness of the algorithms. A first type of methods uses pre-processing of the data to correct for bias, and then trains the usual machine learning algorithms on the data [Oneto and Chiappa, 2020]. A second type of methods relies on implementing fairness constraint in classical learning models. A third type of methods consists in post-processing the output of an unfair algorithm in order to debias it. Our algorithm is somewhere between the first and second class, as it enforces the fairness constraint by estimating during the sequential learning process the bias in the data in order correct it. More precisely, we address the problem of sequential decision making when the feedback on the action taken is provided by an unfair assessment, and does not correspond to the true value of the action.

Statistical fairness has mainly been studied in supervised learning problems, which can be summarized as follows. The statistician has access to a set of n observations $(x_i, z_i, y_i)_{i \leq n}$, where $x_i \in \mathcal{X}$ is a covariate,

^{*}solenne.gaucher@math.u-psud.fr

$y_i \in \mathcal{Y}$ is an outcome that one seeks to predict, and $z_i \in \mathcal{Z}$ is a discrete sensitive attribute corresponding to a group to which x_i belongs. These observations are used to build a prediction function $f : \mathcal{X} \times \mathcal{Z} \rightarrow \mathcal{Y}$ which associates to each point (x, z) a prediction $f(x, z)$ that must be fair towards the different groups. Depending on the objective, different notions of fairness can be used. Two classical notions of group fairness are Demographic Parity and Equalized Odds. Demographic Parity enforces that the predicted outcome $f(x, z)$ is (statistically) independent of the group z . While this notion circumvents biases present in data by ignoring outcomes, it may lead to some issues when correlations exist between the group z and the outcome y . For example, one group may have, on average, more physical strength than the other one, and this skill can be considered as fair when it comes to recruit a piano mover. Such a situation is avoided by Equalized Odds, which ensures that the predicted outcome $f(x, y)$ is independent of the group z conditionally on the outcome y . Yet, Equalized Odds can be very sensitive to biases in the observed outcomes and may reproduce bias from the learning data. Causal fairness provides a notion avoiding these two pitfalls. The influence of sensitive attributes on the rewards and the covariates are captured by structural models, which allows to remove all discriminatory influences in the decisions. The main pitfall with causal fairness is that the notion heavily rely on the causal model, whose accuracy is critical.

In the present paper, we consider the situation where a player must sequentially decide which action to take, and only has access to an unfair assessment that provides a biased evaluation y of the action taken. For example, examiners may be prejudiced against people from a minority group, and give them lower grades; similarly, algorithms trained on biased data may produce unfair assessments of the credit risk of individuals belonging to a minority group. We consider assessments based on a covariate $x \in \mathbb{R}^d$ describing the action and on a sensitive attribute $z_x \in \{-1, +1\}$ indicating the group this action belongs to. The evaluation are systematically biased against a certain group, and we assume that the evaluation y of the action described by (x, z_x) is given by the following biased linear model :

$$\mathbb{E}[y|x, z_x] = x^\top \gamma^* + z_x \omega^*$$

where $\gamma^* \in \mathbb{R}^d$ and $\omega^* \in \mathbb{R}$. We consider a symmetric bias without loss of generality, since we can always add an intercept term in this model.

This biased linear model is a simple example of causal model with linear structural model equations $x = f(z, \xi')$ and $y = x^\top \gamma^* + \omega^* z + \xi$, where ξ and ξ' are noise terms: the covariates x may depend on the sensitive attribute z , and the outcome y of the evaluation mechanism depends on both. Considering the covariate x as a resolving variable [Kilbertus et al., 2017], we see that a fair evaluation of the action is given by $x^\top \gamma^*$, as it does not depend on z conditionally on x . The problem of sequential decision making under biased feedback, called the *biased linear bandit problem*, can be formalized as follows.

Biased linear bandit problem. A player is given a budget $T \in \mathbb{N}$, and is presented with a set of k actions characterized by covariates $x \in \mathcal{X} \subset \mathbb{R}^d$, and by a sensitive attribute $z_x \in \{-1, 1\}$ (we henceforth say that two actions are in the same group if they share the same sensitive attribute). We denote by $\mathcal{A} = \left\{ \begin{pmatrix} x \\ z_x \end{pmatrix} : x \in \mathcal{X} \right\}$ the set of actions and of corresponding sensitive attributes. At time $t \leq T$, the player chooses the action x_t , and receives a reward $x_t^\top \gamma^*$ corresponding to the value of the action that she has chosen. The best action with hindsight knowledge of the parameter γ^* is $x^* \in \operatorname{argmax}_{x \in \mathcal{X}} x^\top \gamma^*$. The player aims at minimizing her regret :

$$R_T = \mathbb{E} \left[\sum_{t \leq T} (x^* - x_t)^\top \gamma^* \right]. \quad (1)$$

By contrast to the classical linear bandit problem, the true reward $x_t^\top \gamma^*$ is never observed. Instead, at each round, the player only observes the outcome y_t of a biased evaluation of the action, given by

$$y_t = x_t^\top \gamma^* + z_{x_t} \omega^* + \xi_t.$$

Here, $\gamma^* \in \mathbb{R}^d$ is the regression parameter specifying the true value of the action, while $\omega^* \in \mathbb{R}$ is a bias parameter, induced by unequal treatment of the groups by the evaluation mechanism, and $\xi_t \stackrel{i.i.d}{\sim} \mathcal{N}(0, 1)$ is a noise term. We denote $\theta^* = \begin{pmatrix} \gamma^* \\ \omega^* \end{pmatrix} \in \mathbb{R}^{d+1}$ the full parameter, and $a_x = \begin{pmatrix} x \\ z_x \end{pmatrix}$ the vector describing an

action and its group. In the rest of the paper, we assume that the reward of all actions are bounded by 1: $\max_{x \in \mathcal{X}} |x^\top \gamma^*| \leq 1$. Moreover, we assume that $|\mathcal{X}| = k$, and that all covariates $x \in \mathcal{X}$ are distinct, which implies that the group z_x is uniquely defined for $x \in \mathcal{X}$. Finally, we consider sets of actions and groups \mathcal{A} such that \mathcal{A} spans \mathbb{R}^{d+1} . Note that this assumption that \mathcal{A} spans \mathbb{R}^{d+1} is enough to guarantee identifiability of the parameters.

Biased linear model. The linear model with systematic bias has been considered in the batch setting in [Chzhen and Schreuder, 2020], where the authors investigate the optimal trade-off between minimax-risk and Demographic Parity. Detection of systematic bias, interpreted as a treatment effect, has been investigated in a batch setting in [Khademi et al., 2019]. In [Barik and Honorio, 2021], the authors consider a similar model, with unobserved sensitive attribute z_i and known bias parameter ω^* . Their goal is to recover both the sensitive attributes z_i and the parameter γ^* , in a high-dimensional setting, under a sparsity assumption on the regression parameter γ^* . This paper also considers a batch setting with a large number of observations at each point, and a large number of points x_i distributed independently of z_i . By contrast, we show that estimation of the bias parameter is one of the main difficulties of the biased bandit problem. Indeed, classical algorithms such as phased elimination algorithms (see, e.g., Algorithm 12 in [Lattimore and Szepesvári, 2020]) or linear Upper Confidence Bound could suffer linear regret in this setting if at some round, the set of potentially optimal actions contains less than d actions, with at least one action in each group. In this case, the true rewards cannot be compared using only these good actions, and sampling sub-optimal actions becomes necessary to gain information on the bias.

Fairness in bandit problems. A recent line of work has focused on the problem of fairness in bandit from the perspective of fair budget allocation between the actions. This problem is motivated by the fact that classical bandit algorithms select all sub-optimal actions but a vanishing fraction of time, which may be undesirable in many situations. To circumvent this problem and guarantee diversity in the actions selected, some papers [Celis et al., 2018] [Patil et al., 2020], [Claure et al., 2020], and [Hadiji et al., 2020] have proposed new algorithms ensuring that the selection frequency of each action belongs to a given set (e.g., the set of vectors whose inputs are positive and bounded away from 0). Alternatively, the authors of [Wang et al., 2021] suggest relaxing this constraint on the selection frequencies, and instead penalize strategies according to their distance to a fair oracle budget allocation. Note that the framework studied in this paper is different: by contrast, we consider here that the mechanism for observing the rewards is unfair, and we aim at correcting it in order to maximize a (fair) true cumulative reward.

Biased linear bandit as a partial monitoring problem. The problem of linear bandit with biased feedback can be viewed as a stochastic partial monitoring game with linear feedback and linear reward [Kirschner et al., 2020]. In this problem, each action x belongs to a compact set $\mathcal{X} \subset \mathbb{R}^d$, and is described by a linear observation operator $a_x \in \mathbb{R}^{d \times k}$. When the player chooses action x_t , she receives the reward $x_t^\top \theta^*$, and observes a noisy outcome $a_{x_t}^\top \theta^* + \xi_t$, where ξ_t is e.g. assumed to be an i.i.d, centered, sub-gaussian random variable. Depending on the action set and the operator, partial monitoring games can be classified either as globally observable, locally observable, or not observable. A game is said to be globally observable if for all $x, x' \in \mathcal{X}$, $x - x' \in \text{Span}(a_{x''} : x'' \in \mathcal{X})$: in this case, the best action $x^* \in \arg\max_{x \in \mathcal{X}} x^\top \theta^*$ can be identified using the observations $(a_x)_{x \in \mathcal{X}}$. A game which is not globally observable is said to be not observable, and the regret of any algorithm must be linear for some parameter θ^* (see [Kirschner et al., 2020], Theorem 31). Finally, a globally observable game is said to be locally observable if for all convex set $\mathcal{C} \subset \mathbb{R}^d$, $x - x' \in \text{Span}(a_{x''} : x'' \in x^*(\mathcal{C}))$ for all $x, x' \in x^*(\mathcal{C})$, where $x^*(\theta) = \arg\max_{x \in \mathcal{X}} x^\top \theta$, and $x^*(\mathcal{C}) = \cup_{\theta \in \mathcal{C}} x^*(\theta)$. Local observability, verified for example in classical linear bandits, ensures in words that the best action can be distinguished by selecting only good actions: classical algorithms then typically achieve regrets of the order $d\sqrt{T}$, up to logarithmic factors. When a game is globally observable but not locally observable, it has been shown that a worst case regret of order $T^{2/3}$ is unavoidable in some cases (see [Kirschner et al., 2020], Theorem 32).

The biased linear bandit can be formulated as a bandit problem with partial linear monitoring if we extend the actions $x \in \mathbb{R}^d$ to actions in \mathbb{R}^{d+1} by adding a $d+1$ -th coordinate equal to 0, and setting $a_x = \begin{pmatrix} x \\ z_x \end{pmatrix}$. With the terminology of partial monitoring, we note that under the assumption that the

set of actions and labels \mathcal{A} spans \mathbb{R}^{d+1} , the biased problem considered in the present paper is globally observable but not locally observable. Algorithms matching the $\tilde{O}(T^{2/3})$ distribution-independent regret rate have been proposed for the related problem of combinatorial linear bandits with partial monitoring in [Lin et al., 2014, Chaudhuri and Tewari, 2016]. Those papers also establish distribution-dependent regret rates. Closest to our setting is the stochastic linear bandits problem with partial monitoring studied in [Kirschner et al., 2020]. In this paper, the authors propose an Information Directed Sampling algorithm matching both the $\tilde{O}(T^{2/3})$ worst-case regret rate in globally observable games, and the $\tilde{O}(T^{1/2})$ worst-case regret rate in locally observable games. However, it is unclear if the multiplicative terms hidden in the $\tilde{O}(T^{1/2})$ rate are optimal (in terms of dependence on the dimension d , on the alignment constant α , and on the number of actions k). By contrast, in the present paper, we present a new algorithm based on phased elimination with additional exploration rounds to estimate the bias parameter ω^* . We obtain distribution-dependent and distribution-independent upper bounds on this regret. We also show that these bounds are optimal on some instance problems, respectively up to multiplicative constants and sub-logarithmic factors.

1.1 Outline and contribution

In this paper, we introduce a new model of linear bandit with biased feedback. In Section 3, we present a Fair Phased Elimination algorithm based on optimal design. We begin by recalling some known result on the related problems of G-optimal and c-optimal design in Section 2, and by discussing the problem of Δ -optimal design. In Section 3, we present a sketch of the Fair Phased Elimination algorithm, detailed in Appendix 3. Then, in Section 4, we establish an upper bound on its worst-case regret, i.e. on the regret of the algorithm for the worst possible parameter θ^* . In Section 5, we derive a gap-dependent upper bound on the regret of our algorithm, i.e. an upper bound on the regret of the algorithm for all parameters θ^* under assumptions on the difference between the rewards of the best and second best action, and between the best actions of each group. We establish lower bounds on some action sets \mathcal{A} for both the worst-case and the gap-dependent regret, showing that these rates are sharp up respectively up to a sub-logarithmic factor and an absolute multiplicative constant. In order to provide intuition without burdening the reader with details, we sketch the main ideas of our algorithm and results, but defer all details and proofs to the Appendix.

Our upper and lower bounds on the worst-case regret show that it is of order $\kappa_*^{1/3} \log(T)^{1/3} T^{2/3}$ for T large, where κ_* is a constant depending on the geometry of the action set \mathcal{A} , related to c-optimal design, as well as to the problem of optimal estimation against the worst parameter. Interestingly, this large T worst-case regret depends on the geometry of the problem only through κ_* , and not through the dimension d . Note that the regret of order $\tilde{O}(T^{2/3})$ is higher than the classical regret of order $\tilde{O}(dT^{1/2})$ obtained for linear bandits : we show that this increase in the regret corresponds to the price to pay for debiasing the unfair evaluations.

We also characterize the gap-depend regret, showing that it is of order $\left(\frac{d}{\Delta_{\min}} \vee \frac{\kappa(\Delta)}{\Delta_{\neq}^2}\right) \log(T)$, where Δ_{\min} is the minimum gap, Δ_{\neq} is the gap between the best actions of the two groups, and $\kappa(\Delta)$ corresponds to the minimum regret to pay in order to estimate the bias with a variance lower than 1. This bound underlines the relative difficulties of the classical d -dimensional linear bandit problem, and of the bias estimation : when $\frac{d}{\Delta_{\min}} \geq \frac{\kappa(\Delta)}{\Delta_{\neq}^2}$, i.e. when one group dominates the other, and contains all near-optimal actions, the difficulty is dominated by that of the corresponding standard linear bandit problem. On the other hand, when both groups contain near-optimal actions, and $\frac{d}{\Delta_{\min}} \leq \frac{\kappa(\Delta)}{\Delta_{\neq}^2}$, the regret corresponds to the price of debiasing the reward.

1.2 Notations

Recall that the agent plays the linear bandit with biased feedback game for T rounds. For a given set of actions \mathcal{A} and a parameter θ , we denote $R_T^{\mathcal{A}, \theta}$ the regret defined in Equation (1). When the dependence of the regret on \mathcal{A} and θ is obvious from context, we relax notations and simply denote it R_T . For an event \mathcal{F} such that $\mathbb{P}(\mathcal{F}) > 0$, we denote by $\mathbb{E}_{|\mathcal{F}}$ (resp. $\mathbb{P}_{|\mathcal{F}}$) the expectation (resp. the probability) conditionally on \mathcal{F} . For any matrix M , we denote by M^+ a generalized inverse of M , that is a matrix such that $MM^+M = M$. Though theory is blind to the choice of the generalized inverse, for numerical stability it is wise to choose the Moore–Penrose generalized inverse in practice. We denote by $\mathcal{P}^{\mathcal{X}}$ the set of probabilities measures on \mathcal{X} ,

$\mathcal{N}^{\mathcal{X}} = \{\mu : \mathcal{X} \mapsto \mathbb{N}\}$, and $\mathcal{M}^{\mathcal{X}} = \{\mu : \mathcal{X} \mapsto \mathbb{R}_+\}$. For a function $\mu \in \mathcal{P}^{\mathcal{X}}$, $\mu \in \mathcal{N}^{\mathcal{X}}$, or $\mu \in \mathcal{M}^{\mathcal{X}}$, we denote $V(\mu) = \sum_{x \in \mathcal{X}} \mu(x) \begin{pmatrix} x \\ z_x \end{pmatrix} \begin{pmatrix} x \\ z_x \end{pmatrix}^\top$ the covariance matrix corresponding to this allocation. Finally, for $i \leq d+1$, we denote by e_i the i -th vector of the canonical basis of \mathbb{R}^{d+1} .

2 Optimal design for bandits

Before presenting the algorithm introduced to solve the linear bandit problem, we emphasize the following point: classical techniques from linear bandits can be used to bound the prediction error of the biased evaluations $\hat{\theta}^\top \begin{pmatrix} x \\ z_x \end{pmatrix}$ for the actions $x \in \mathcal{X}$. The main difficulty in the biased linear bandit problem is to bound the prediction error of the debiased reward, i.e., to bound the error of $\hat{\theta}^\top \begin{pmatrix} x \\ 0 \end{pmatrix}$ when no observation can be made at the points $\begin{pmatrix} x \\ 0 \end{pmatrix}$. To obtain an unbiased estimate of the rewards, we propose to independently estimate the evaluations $\begin{pmatrix} x \\ z_x \end{pmatrix}^\top \theta^*$, and the bias ω^* . To do this, we rely on results from optimal design.

Optimal design theory addresses the following problem. A scientist must design a set of n experiments $\{x_1, \dots, x_n\} \in \mathcal{X}^n$ so as to estimate a parameter of interest, where each experiment $x \in \mathcal{X}$ corresponds to a point $a_x \in \mathbb{R}^{d+1}$. The aim of the scientist is to choose a design, i.e. a function $\mu \in \mathcal{M}^{\mathcal{X}}$ indicating the budget $\mu(x)$ to be allocated to each experiment $x \in \mathcal{X}$. Each experiment x is then repeated exactly $\mu(x)$ times, and the corresponding observations $y_{x,1}, \dots, y_{x,\mu(x)}$ are collected for each $x \in \mathcal{X}$. The law of the observations corresponding to experiment x at point a_x is given by

$$y_{x,i} = a_x^\top \theta^* + \xi_{x,i},$$

where $\xi_{x,i} \sim \mathcal{N}(0,1)$ are independent noise terms, and $\theta^* \in \mathbb{R}^{d+1}$ is an unknown parameter. The aim of the scientist is to choose the design μ so as to best estimate (some features of) the parameter θ^* , under a constraint on the total number of experiments $\sum_{x \in \mathcal{X}} \mu(x) \leq n$ for some $n \in \mathbb{N}$.

Different criteria can be used to characterize the optimality of a design μ . For example, in the biased linear bandit problem, we need to estimate the evaluation $a_x^\top \theta^*$ of the different actions x , where $a_x = \begin{pmatrix} x \\ z_x \end{pmatrix}$. To obtain estimates with a small error uniformly over the set of actions \mathcal{X} , we rely on a G-optimal design. To debias these evaluations, we can then estimate the bias ω^* by using methods from c -optimal design.

2.1 G-optimal design

G-optimal design aims at minimizing the maximal prediction error over the set of experiments \mathcal{X} . Recall that the ordinary least square estimator is given by

$$\hat{\theta} = V(\mu)^+ \sum_{x \in \mathcal{X}} a_x \left(\sum_{i \leq \mu(x)} y_{x,i} \right) \quad \text{for} \quad V(\mu) = \sum_{x \in \mathcal{X}} \mu(x) a_x a_x^\top \quad (2)$$

where $V(\mu)^+$ is any generalized inverse of the covariance matrix $V(\mu)$. Note that when the image of $V(\mu)$ does not span \mathbb{R}^{d+1} , the least square estimator $\hat{\theta}$ may not be uniquely defined.

When a_x belongs to the image of $V(\mu)$, the best linear unbiased estimator of the outcome $a_x^\top \theta^*$ of experiment x is given by $a_x^\top \hat{\theta}$, where $\hat{\theta}$ is the ordinary least square (OLS) estimator defined in (2). We underline that the estimator $a_x^\top \hat{\theta}$ does not depend on the choice of the generalized inverse $V(\mu)^+$ when a_x belongs to the image of $V(\mu)$ (which holds as soon as $\mu(x) \neq 0$). Then, the variance of the prediction $a_x^\top \hat{\theta}$ is given by $a_x^\top V(\mu)^+ a_x$. Let us denote by $\mathcal{N}_{\mathcal{X}}^{\mathcal{X}}$ the set of functions $\mu \in \mathcal{M}^{\mathcal{X}}$ such that $\{a_x : x \in \mathcal{X}\} \in \text{Range}(V(\mu))$. The exact G-optimal design solves the following problem :

$$\underset{\mu \in \mathcal{N}_{\mathcal{X}}^{\mathcal{X}}}{\text{minimize}} \max_{x \in \mathcal{X}} a_x^\top V(\mu)^+ a_x \quad \text{such that} \quad \sum_{x \in \mathcal{X}} \mu(x) \leq n.$$

To remove the dependence of the objective on the budget n , let us consider the probability measure π given by $\pi(x) = \mu(x)/n$ for all $x \in \mathcal{X}$. For all $x \in \mathcal{X}$, the prediction error at experiment x is now given by $n^{-1} a_x^\top V(\pi)^+ a_x$. Therefore, the best design under the budget constraint minimizes $\max_x a_x^\top V(\pi)^+ a_x$.

Let $\mathcal{P}_{\mathcal{X}}^{\mathcal{X}}$ be the set of probability measures on \mathcal{X} such that $\{a_x : x \in \mathcal{X}\} \in \text{Range}(V(\pi))$. In the limit $n \rightarrow +\infty$, solving the exact G-optimal design problem is equivalent to finding the probability measure π^* solution of the following approximate G-optimal design problem

$$\min_{\pi \in \mathcal{P}_{\mathcal{X}}^{\mathcal{X}}} \max_{x \in \mathcal{X}} a_x^\top (V(\pi))^+ a_x. \quad (\text{G-optimal design}) \quad (3)$$

It has been shown that this problem is equivalent to the following D-optimal design problem

$$\max_{\mu \in \mathcal{P}_{\mathcal{X}}^{\mathcal{X}}} \det \left(\left(P^\top (V(\pi))^+ P \right)^{-1} \right), \quad (\text{D-optimal design})$$

where $\det(M)$ is the determinant of the matrix M and $P \in \mathbb{P}^{d \times d'}$ is the restriction of the orthonormal projector onto the span of $\{a_x : x \in \mathcal{X}\}$ (whose dimension is denoted d') to the span of $\{a_x : x \in \mathcal{X}\}$. Then, the General Equivalence Theorem of Kiefer [Kiefer, 1974] and Pukelsheim [Pukelsheim, 1980] shows that any design π^* solution to the G-optimal design problem (3) satisfies

$$\max_{x \in \mathcal{X}} a_x^\top (V(\pi^*))^+ a_x \leq d',$$

(see Theorem 2.4.4 in [Sagnol, 2010] for a statement of this result). The G-optimal design π^* can be chosen so that it is supported by at most $(d+1)(d+2)/2$ points (see, e.g., Section 2.3.3 in [Sagnol, 2010]). Polynomial-time algorithms computing the G-optimal design have been devised, using for example interior points method [Vandenberghe et al., 1998], or mixed integer second-order cone programming [Sagnol and Harman, 2013].

A rounding procedure is needed in order to derive an exact design $\mu \in \mathcal{N}^{\mathcal{X}}$ from an approximate G-optimal design $\pi \in \mathcal{P}^{\mathcal{X}}$. A classical trick is to use the design $\mu(x) = \lceil m\pi(x) \rceil$ where m is chosen in order to get the desired precision level.

Lemma 1. *Let π be a solution of the G-optimal design problem (3). Then, for any $m > 0$ and any $x \in \mathcal{X}$, the estimator $a_x^\top \hat{\theta}_\mu$ computed from the design $\mu : x \mapsto \lceil m\pi(x) \rceil$ is an unbiased estimator of the evaluation $a_x^\top \theta$, and it has a variance*

$$a_x^\top V(\mu)^+ a_x \leq m^{-1}(d+1).$$

2.2 c -optimal design.

Optimal design theory can also be used to estimate the bias parameter ω^* with a minimal number of samples. Let c be a vector in \mathbb{R}^{d+1} . Then, c -optimal design aims at finding the best design $\mu \in \mathcal{N}^{\mathcal{X}}$ for estimating the scalar product $c^\top \theta^*$ under a budget constraint $\sum_{x \in \mathcal{X}} \mu(x) \leq n$. Unbiased linear estimation of $c^\top \theta^*$ is possible only when c belongs to the image of $V(\mu)$, and in this case the best linear unbiased estimator of the scalar product $c^\top \theta^*$ is given by $c^\top \hat{\theta}$, where $\hat{\theta}$ is the least-square estimator defined in (2). The variance of the estimator $c^\top \hat{\theta}$ is then equal to $c^\top V(\mu)^+ c$.

Exact c -optimal design aims at choosing the allocation $\mu \in \mathcal{N}^{\mathcal{X}}$ minimizing the variance of $c^\top \hat{\theta}$ for a given budget $\sum_x \mu(x) \leq n$, under the constraint that $c \in \text{Range}(V(\mu))$. Again, let us define the normalized design $\pi : x \in \mathcal{X} \mapsto \mu(x)/n$, and let us underline that π defines a probability on \mathcal{X} . The variance of $c^\top \hat{\theta}$ is then equal to $n^{-1} c^\top V(\pi)^+ c$. In the limit $n \rightarrow +\infty$, the problem is equivalent to the problem of approximate c -optimal design (sometimes simply referred to as *c-optimal design*), that aims at finding a probability measure $\pi \in \mathcal{P}_c^{\mathcal{X}} := \{\pi \in \mathcal{P}^{\mathcal{X}} : c \in \text{Range}(V(\pi))\}$ solution to the following problem

$$\min_{\pi \in \mathcal{P}_c^{\mathcal{X}}} c^\top V(\pi)^+ c. \quad (c\text{-optimal design})$$

Next theorem characterizes solutions to the c -optimal design problem.

Theorem 1 ([Elfving, 1952]). *Let $\mathcal{S} = \text{convex hull}\{+a_x, -a_x : x \in \mathcal{X}\}$ be the Elfving's set of $\{a_x : x \in \mathcal{X}\} \subset \mathbb{R}^{d+1}$, and let $\partial\mathcal{S}$ denote the boundary of \mathcal{S} . A design $\pi \in \mathcal{P}_c^{\mathcal{X}}$ is c -optimal for $c \in \mathbb{R}^{d+1}$ if and only if there exists $\zeta \in \{-1, +1\}^{\mathcal{X}}$ and $t > 0$ such that*

$$tc = \sum_{x \in \mathcal{X}} \pi(x) \zeta_x a_x \in \partial\mathcal{S}.$$

Moreover, $t^{-2} = c^\top (V(\pi))^+ c$ is the minimal variance of the bias estimator.

Note that when $\{a_x : x \in \mathcal{X}\}$ spans \mathbb{R}^{d+1} , for any $c \in \mathbb{R}^{d+1}$, there exists a design π such that $c \in \text{Range}(V(\pi))$, and hence the c -optimal design problem admits a solution. In addition, there exists a c -optimal design supported by at most $d + 1$ points (see, e.g., [Pázman, 1986, Harman and Jurík, 2008] for a proof of this result). Moreover, while finding an exact optimal allocation $\mu \in \mathcal{N}^{\mathcal{X}}$ under the constraint that $\sum_{x \in \mathcal{X}} \mu(x) \leq n$ is NP-complete, finding an approximate optimal design $\pi \in \mathcal{P}_c^{\mathcal{X}}$ can be done in polynomial time [Černý and Hladík, 2012]. Several algorithms, including multiplicative algorithms [Fellman, 1974] and a simplex method of linear programming [Harman and Jurík, 2008], have been proposed to iteratively approximate the optimal design. More recently, [Pronzato and Sagnol, 2021] suggested using screening tests to remove inessential points for the designs to accelerate optimization algorithms.

A rounding procedure is needed in order to derive a design in $\mathcal{N}^{\mathcal{X}}$ from a c -optimal design π . Again, we use the design $\mu : x \mapsto \lceil m\pi(x) \rceil$ where m is chosen in order to get the desired precision level. Let $\mathcal{M}_{e_{d+1}}^{\mathcal{X}} = \{\mu \in \mathcal{M}^{\mathcal{X}} : e_{d+1} \in \text{Range}(V(\mu))\}$ denote the set of positive measures over \mathcal{X} such that e_{d+1} belongs to the image of $V(\mu)$.

Lemma 2. *For any $\pi \in \mathcal{M}_{e_{d+1}}^{\mathcal{X}}$ and any $m > 0$, the estimator $e_{d+1}^\top \hat{\theta}_\mu$ computed from the design $\mu : x \mapsto \lceil m\pi(x) \rceil$ is an unbiased estimator of $e_{d+1}^\top \theta$ and it has a variance at most $m^{-1} e_{d+1}^\top V(\pi)^+ e_{d+1}$.*

Recall that we need to estimate $e_{d+1}^\top \theta^* = \omega^*$ in order to debias the evaluations $a_x^\top \theta^*$. In the following, we denote by v the value of the solution of the corresponding e_{d+1} -optimal design problem, and refer to it as the *minimal variance of the bias estimator*:

$$v = \min_{\pi \in \mathcal{P}_{e_{d+1}}^{\mathcal{X}}} e_{d+1}^\top (V(\pi))^+ e_{d+1},$$

where we recall that $V(\pi) = \sum_{x \in \mathcal{X}} \pi(x) \begin{pmatrix} x \\ z_x \end{pmatrix} \begin{pmatrix} x \\ z_x \end{pmatrix}^\top$ and $\mathcal{P}_{e_{d+1}}^{\mathcal{X}} = \{\pi \in \mathcal{P}^{\mathcal{X}} : e_{d+1} \in \text{Range}(V(\pi))\}$.

The e_{d+1} -optimal design can be equivalently defined as the solution of the problem

$$\text{minimize } \sum_{x \in \mathcal{X}} \mu(x) \quad \text{such that } \mu \in \mathcal{M}_{e_{d+1}}^{\mathcal{X}} \text{ and } e_{d+1}^\top V(\mu)^+ e_{d+1} \leq v. \quad (4)$$

The characterization given in Equation (4) underlines that the e_{d+1} -optimal design provides (up to discretization issues) the minimal number of samples required for estimating ω^* with a variance v , or equivalently with a precision level \sqrt{v} . Let us denote by μ^* the optimal design for estimating ω^* with a precision level 1, defined as

$$\begin{aligned} \mu^* &= \underset{\mu}{\operatorname{argmin}} \sum_{x \in \mathcal{X}} \mu(x) \\ &\quad \text{such that } \mu \in \mathcal{M}_{e_{d+1}}^{\mathcal{X}} \\ &\quad \text{and } e_{d+1}^\top V(\mu)^+ e_{d+1} \leq 1. \end{aligned}$$

By definition of v , we see that $\sum_x \mu^*(x) = v$. A first (naive) approach to obtain an estimate of ω^* with precision level $\epsilon > 0$ would then consist in sampling actions according to $\epsilon^{-2} \mu^*$, rounded according to the procedure defined in Lemma 2. Let us denote by Δ_x the gap $\Delta_x = \max_{x' \in \mathcal{X}} (x' - x)^\top \gamma^*$ between the (non-observed) reward of the best action and the reward of the action x . The regret corresponding to this estimation phase would then be

$$\epsilon^{-2} \sum_{x \in \mathcal{X}} \mu^*(x) \Delta_x,$$

which can be as large as $v \epsilon^{-2} \max_x \Delta_x$. We show below that we can estimate ω^* with the same level of precision at a smaller cost by taking the gaps Δ_x into account while choosing the design.

2.3 Δ -optimal design

To estimate the bias $\omega^* = e_{d+1}^\top \theta^*$ with an error smaller than some $\epsilon > 0$, we can sample the actions according to any design $\mu \in \mathcal{M}_{e_{d+1}}^{\mathcal{X}}$ such that $e_{d+1}^\top V(\mu)^+ e_{d+1} \leq \epsilon^2$. Equivalently, we can choose the actions according

to $\epsilon^{-2}\mu$, where $\mu \in \mathcal{M}_{e_{d+1}}^{\mathcal{X}}$ is such that $e_{d+1}^\top V(\mu)^+ e_{d+1} \leq 1$. For a vector of gaps $\Delta = (\Delta_x)_{x \in \mathcal{X}}$, let us denote by μ^Δ the Δ -optimal design, defined as the solution of the following problem

$$\begin{aligned} \mu^\Delta &= \underset{\mu}{\operatorname{argmin}} \sum_{x \in \mathcal{X}} \mu(x) \Delta_x && (\Delta\text{-optimal design}) \\ &\text{such that } \mu \in \mathcal{M}_{e_{d+1}}^{\mathcal{X}} \\ &\text{and } e_{d+1}^\top V(\mu)^+ e_{d+1} \leq 1. \end{aligned} \tag{5}$$

If we knew the gaps Δ_x , we could sample the actions according to the Δ -optimal design μ^Δ , and pay the regret $\epsilon^{-2}\kappa(\Delta)$ (up to rounding error) for estimating ω^* with an error smaller than ϵ , where

$$\kappa(\Delta) = \sum_{x \in \mathcal{X}} \mu^\Delta(x) \Delta_x.$$

In practice, the Δ -optimal design can be computed using algorithms designed for finding the e_{d+1} -optimal design. Indeed, the next lemma shows that the computation of the Δ -optimal design amounts to computing an e_{d+1} -optimal design for some rescaled features.

Lemma 3. *For any vector $\Delta \in (0, +\infty)^{\mathcal{X}}$, let π^Δ be the e_{d+1} -optimal design relative to the set $\mathcal{A}^\Delta = \left\{ \Delta_x^{-1/2} \begin{pmatrix} x \\ z_x \end{pmatrix} : x \in \mathcal{X} \right\}$ and let $\kappa^\Delta = e_{d+1}^\top V(\pi^\Delta)^+ e_{d+1}$ be the e_{d+1} -optimal variance relative to \mathcal{A}^Δ . Then, the Δ -optimal design μ^Δ is given by $\mu^\Delta(x) = \kappa^\Delta \pi^\Delta(x) \Delta_x^{-1}$ for all $x \in \mathcal{X}$. In addition, the support of μ^Δ can be chosen to be of cardinality at most $d+1$.*

Obviously, the gaps $\Delta = (\Delta_x)_{x \in \mathcal{X}}$ are unknown, so we cannot sample according to the Δ -optimal design μ^Δ . Instead, we propose to rely on some upper confidence bounds $\hat{\Delta}_x$ on Δ_x to estimate them. Note that when $\hat{\Delta}_x \geq \Delta_x$ for all $x \in \mathcal{X}$, if we sample actions according to a rounding of $\epsilon^{-2}\mu^{\hat{\Delta}}$, where $\mu^{\hat{\Delta}}$ is the $\hat{\Delta}$ -optimal design, the corresponding regret will be upper bounded by $\epsilon^{-2}\kappa(\hat{\Delta})$. The function κ verifies the following properties.

Lemma 4. *For two vectors of gaps Δ, Δ' , denote by $\Delta \wedge \Delta'$ (respectively $\Delta \vee \Delta'$) the vector of gaps given by $(\Delta \wedge \Delta')_x = \Delta_x \wedge \Delta'_x$ (respectively $(\Delta \vee \Delta')_x = \Delta_x \vee \Delta'_x$) for all $x \in \mathcal{X}$. Moreover, denote $\Delta \leq \Delta'$ if $\Delta_x \leq \Delta'_x$ for all $x \in \mathcal{X}$. Then, the following properties hold :*

- i) for all $c > 0$, $\kappa(c\Delta) = c\kappa(\Delta)$;
- ii) if $\Delta \leq \Delta'$, then $\kappa(\Delta) \leq \kappa(\Delta')$;
- iii) $\kappa(\Delta \vee \Delta') \geq \kappa(\Delta) \vee \kappa(\Delta')$;
- iv) the function $\epsilon \mapsto \kappa(\Delta \vee \epsilon)$ is continuous at 0.

Let us denote by $\mathcal{C}(\mathcal{X}) = \{\gamma \in \mathbb{R}^d : \forall x \in \mathcal{X}, |x^\top \gamma| \leq 1\}$ the set of admissible parameters. Let $\Delta_{\max} = \max_{\gamma \in \mathcal{C}(\mathcal{X})} \max_{x, x' \in \mathcal{X}} (x - x')^\top \gamma$ be an upper bound on all gaps Δ_x , and let $\kappa_* = \Delta_{\max} v$. Then, Lemma 4 implies that $\kappa(\Delta) \leq \kappa_*$ for all Δ such that $\Delta_x \leq \Delta_{\max}$ for all $x \in \mathcal{X}$.

3 The Fair Phased Elimination algorithm

Let us now present the Fair Phased Elimination algorithm, and outline the main ideas behind it. The Fair Phased Elimination algorithm, sketched in Algorithm 3, relies on the following three remarks.

First, note that within a group z , the difference between the true rewards $x^\top \gamma^*$ and $(x')^\top \gamma^*$ of any two actions x and x' is equal to the difference between their biased evaluations $a_x^\top \theta^*$ and $a_{x'}^\top \theta^*$. Hence, within a group, classical algorithms for linear bandit can be used to choose actions : in the following, we will use a variant of the Phased Elimination algorithm of [Lattimore and Szepesvári, 2020], based on G -optimal Exploration and Elimination.

Second, to compare actions belonging to different groups, we need to estimate the bias parameter ω^* . This estimation can be performed at a given precision level with a Δ -optimal sampling. If all the actions of one group appear to be sub-optimal with high confidence, then this group can be dropped, and we can proceed by solely running a classical linear bandit algorithm on the remaining group.

To determine which group contains the best action, bias estimation must be performed up to a precision smaller than the gap between the reward of the best actions of the two groups. When this gap is small, it can be overly costly to run the bias estimation up to the precision level required to identify the best group. To avoid this, we use a stopping criterion, which guarantees that we stop estimating the bias beyond a well-chosen precision level.

Below, we describe the G -Explore-And-Eliminate routine, the Δ -Explore-And-Eliminate routine, and the stopping criteria which are the three key ingredients of the Fair Phased Elimination algorithm (FPE). The FPE algorithm is sketched in Algorithm 3, page 11, and described in full details in Algorithm 4 in Appendix A.

G -Exploration and Elimination. Within a group, we can run a Phased Elimination algorithm [Lattimore and Szepesvári, 2017] in order to sample the actions with a controlled within-group regret. For a sequence $(\epsilon_l)_{l \geq 1}$ of precision levels, typically $\epsilon_l = 2^{1-l} \Delta_{\max}$, the Phased Elimination algorithm proceeds by phases $l = 1, 2, \dots$. At phase l the $O(\epsilon_l)$ -suboptimal actions are removed with high-probability. Let us denote by $\mathcal{Z}_l \subset \{-1, 1\}$ and $\mathcal{X}_l^{(z)}$ the remaining active groups and the remaining actions in group $z \in \mathcal{Z}_l$. At phase l , for each active group $z \in \mathcal{Z}_l$, the G -Explore-And-Eliminate Routine 1 is applied to the set of active actions $\mathcal{X}_l^{(z)}$. This routine first estimates the biased rewards $a_x^\top \theta^*$ at precision level ϵ_l by running a G -optimal sampling of the actions and by computing the Ordinary Least Square (OLS) estimator $\hat{\theta}_l^{(z)}$ based on these samples. Then, the actions with estimated gaps larger than $3\epsilon_l$ are removed from $\mathcal{X}_l^{(z)}$. The remaining actions are gathered in the set $\mathcal{X}_{l+1}^{(z)}$.

Let us briefly explain why running the G -Explore-And-Eliminate routine 1 with $n = 2\epsilon_l^{-2}(d+1) \log(kl(l+1)/\delta)$ and $\epsilon = \epsilon_l$ provides an estimation of the biased rewards $a_x^\top \theta^*$ at precision level ϵ_l for all actions $x \in \mathcal{X}_l^{(z)}$. Let $V_l^{(z)}$ be the covariance matrix of the G -optimal design. The OLS estimator $\hat{\theta}_l^{(z)}$ then verifies

$$\mathbb{P} \left(|a_x^\top (\hat{\theta}_l^{(z)} - \theta^*)| \geq \sqrt{2 \|a_x\|_{(V_l^{(z)})^+}^2 \log(kl(l+1)/\delta)} \right) \leq \frac{\delta}{kl(l+1)}.$$

By Lemma 1, we see that for all $x \in \mathcal{X}_l^{(z)}$,

$$a_x^\top \left(V_l^{(z)} \right)^+ a_x \leq \frac{\epsilon_l^2}{2(d+1) \log(kl(l+1)/\delta)}.$$

This implies that

$$\mathbb{P} \left(|a_x^\top (\hat{\theta}_l^{(z)} - \theta^*)| \geq \epsilon_l \right) \leq \frac{\delta}{kl(l+1)}.$$

An union bound then ensures that $a_x^\top \hat{\theta}_l^{(z)}$ provides an estimation of $a_x^\top \theta^*$ at level ϵ_l for all phases $l \geq 1$ and all actions $x \in \mathcal{X}_l^{(z)}$ with probability $1 - \delta$.

At the final Elimination step, the ϵ_l -accurate estimation of $a_x^\top \theta^*$ ensures that all the actions x with within-group gap larger than ϵ_l are removed. Hence, only actions with gap $O(\epsilon_l)$ remains in $\mathcal{X}_{l+1}^{(z)}$. We observe a slight divergence from the classical Explore-and-Eliminate algorithm. Usually, any action with positive gap is removed. Instead, we require here a ϵ_l -gap for removing an action. The reason for this modification is that we need to keep each action x up to a precision level comparable to Δ_x , in order to ensure a precision level $O(\Delta_x \vee \epsilon_l)$ at phase l for the estimation of the gap Δ_x in the next subroutine.

Δ -Explore-And-Eliminate. The estimators $\hat{\theta}_l^{(z)} = \begin{pmatrix} \hat{\gamma}_l^{(z)} \\ \hat{\omega}_l^{(z)} \end{pmatrix}$ of the G -Explore-And-Eliminate routine provide some estimators $\hat{\omega}_l^{(z)}$ of the bias ω^* . We emphasize that, even if the biased rewards $a_x^\top \theta^*$ are estimated at precision ϵ_l for $x \in \mathcal{X}_l^{(z)}$, we have no control on the precision of the bias estimation $\hat{\omega}_l^{(z)}$. So we cannot rely on these estimators for debiasing the rewards. Instead, we estimate the bias based on a

Routine 1 G-EXPLORE-AND-ELIMINATE (\mathcal{X}, n, ϵ)

- 1: Compute G-optimal design π solution of (3) on \mathcal{X} , with $|\text{supp}(\pi)| \leq (d+1)(d+2)/2$
 - 2: Sample $\lceil n\pi(x) \rceil$ times each action a_x for $x \in \mathcal{X}$ ▷ G-optimal parameter estimation
 - 3: Compute the least square estimator $\hat{\theta}$ defined by (2)
 - 4: $\mathcal{X}' \leftarrow \left\{ x \in \mathcal{X} : \max_{x' \in \mathcal{X}} (x' - x)^\top \hat{\theta} \leq 3\epsilon \right\}$ ▷ Suboptimal actions elimination
 - 5: **return** $\hat{\theta}$ and \mathcal{X}'
-

Routine 2 $\hat{\Delta}$ -EXPLORE-AND-ELIMINATE ($\mathcal{X}, (\mathcal{X}^{(z)}, \hat{\theta}^{(z)})_{z \in \{-1, 1\}}, \hat{\Delta}, n, \epsilon$)

- 1: Compute $\hat{\Delta}$ -optimal design $(\hat{\mu}, \kappa(\hat{\Delta}))$ solution of (5) on \mathcal{X} , with $|\text{supp}(\hat{\mu})| \leq d+1$
 - 2: Sample $\lceil n\hat{\mu}(x) \rceil$ times each action a_x for $x \in \mathcal{X}$ ▷ $\hat{\Delta}$ -optimal bias estimation
 - 3: Compute $\hat{\omega} = e_{d+1}^\top \hat{\theta}^{(0)}$, where $\hat{\theta}^{(0)}$ is the OLS estimator (2)
 - 4: **for** $z \in \{-1, 1\}$ **and** $x \in \mathcal{X}^{(z)}$ **do**
 - 5: $\hat{m}_x \leftarrow a_x^\top \hat{\theta}^{(z)} - z\hat{\omega}$ ▷ Debiased rewards estimation
 - 6: **end for**
 - 7: **if** $\exists z \in \{-1, 1\}$ such that $\max_{x \in \mathcal{X}^{(z)}} \hat{m}_x \geq \max_{x \in \mathcal{X}^{(-z)}} \hat{m}_x + 4\epsilon$ **then**
 - 8: $\mathcal{Z} \leftarrow \{z\}$ ▷ Suboptimal group elimination
 - 9: **else**
 - 10: $\hat{\Delta}_x \leftarrow \Delta_{\max} \wedge (\max_{x' \in \mathcal{X}^{(-1)} \cup \mathcal{X}^{(1)}} \hat{m}_{x'} - \hat{m}_x + 4\epsilon)$ for all $x \in \mathcal{X}^{(-1)} \cup \mathcal{X}^{(1)}$ ▷ Gaps UCB
 - 11: **end if**
 - 12: **return** \mathcal{Z} and $\hat{\Delta}$
-

$\hat{\Delta}$ -optimal sampling ensuring a precision at level ϵ_l , where $\hat{\Delta}$ is an upper-confidence bound for Δ . Actions are sampled according to the design $\hat{\mu}_l$ solution of (5) and the bias estimator $\hat{\omega}_l$ is computed from the sampled actions. Then, reward estimations $a_x^\top \hat{\theta}_l^{(z)}$ for $x \in \mathcal{X}_{l+1}^{(z)}$ are debiased by subtracting $z\hat{\omega}_l$. Similarly as for the G-Explore-And-Eliminate routine, this provides debiased reward estimations $\hat{m}_x = a_x^\top \hat{\theta}_l^{(z)} - z\hat{\omega}_l$ of $x^\top \gamma^*$ at precision $2\epsilon_l$ for all $x \in \mathcal{X}_{l+1}^{(z)}$.

These debiased reward estimators \hat{m}_x are then used for comparing the groups and updating the gaps estimators $\hat{\Delta}_x$. First, the Routine 2 checks whether the maximum debiased rewards \hat{m}_x in one group $\mathcal{X}_{l+1}^{(z)}$ exceeds the maximum debiased rewards in the other group $\mathcal{X}_{l+1}^{(-z)}$ by more than $4\epsilon_l$. If so, the algorithm only keep the group $\mathcal{X}_{l+1}^{(z)}$ (suboptimal group elimination), and no more bias estimation is performed. The algorithm then reduces to a simple Phased Elimination algorithm [Lattimore and Szepesvári, 2020] on the remaining group. If not so, for all $x \in \mathcal{X}_{l+1}^{(-1)} \cup \mathcal{X}_{l+1}^{(1)}$, the upper confidence bound $\hat{\Delta}_x$ on the gap is updated using \hat{m}_x , in order to get a $O(\epsilon_l)$ -precision level.

Breaking criteria. We observe that the $\hat{\Delta}$ -optimal Exploration and Elimination samples actions in the whole set of actions \mathcal{X} , not only in the remaining near-optimal action sets $\mathcal{X}_{l+1}^{(1)} \cup \mathcal{X}_{l+1}^{(-1)}$. As a consequence, when the gap between the best two actions of each group is small, finding the best group with $\hat{\Delta}$ -optimal bias estimation can be overly costly in terms of regret. To avoid excessive regret, if the best group has not been found at stage l fulfilling $\epsilon_l \leq \left(\kappa(\hat{\Delta}^l) \log(T)/T \right)^{1/3}$, the bias estimation is stopped and the empirical best action in $\mathcal{X}_{l+1}^{(1)} \cup \mathcal{X}_{l+1}^{(-1)}$ is sampled for the remaining time.

Putting together the three main ingredients described above, we obtained the Fair Phased Elimination algorithm sketched in Algorithm 3, and fully described in Algorithm 4, page 24.

Algorithm 3 FAIR PHASED ELIMINATION (sketched)

```
1: input:  $\delta, T, \mathcal{X}, k = |\mathcal{X}|, \Delta_{\max} \leftarrow \max_{\gamma \in \mathcal{C}(\mathcal{X}), (x, x') \in \mathcal{X}^2} (x - x')^\top \gamma, \epsilon_l = 2^{1-l} \Delta_{\max}$  for  $l \geq 1$ 
2: initialize:  $\mathcal{X}_1^{(+1)} \leftarrow \{x : z_x = 1\}, \mathcal{X}_1^{(-1)} \leftarrow \{x : z_x = -1\}, \mathcal{Z}_1 \leftarrow \{-1, +1\}$ 
3:  $\Delta_{\max} \leftarrow \max_{\gamma \in \mathcal{C}(\mathcal{X}), (x, x') \in \mathcal{X}^2} (x - x')^\top \gamma, \hat{\Delta}^1 \leftarrow (\Delta_{\max}, \dots, \Delta_{\max})$ 
4: Iterate until all the budget is spent:
5: for  $l \geq 1$  do
6:   for  $z \in \mathcal{Z}_l$  do
7:      $(\hat{\theta}^{(z)}, \mathcal{X}_{l+1}^{(z)}) \leftarrow \text{G-EXPLORE-AND-ELIMINATE} \left( \mathcal{X}_l^{(z)}, \frac{2(d+1)}{\epsilon_l^2} \log \left( \frac{kl(l+1)}{\delta} \right), \epsilon_l \right)$ 
8:   end for
9:   if  $\mathcal{Z}_l = \{-1, +1\}$  then
10:    if  $\epsilon_l \leq \left( \kappa(\hat{\Delta}^l) \log(T)/T \right)^{1/3}$  then ▷ Break bias estimation
11:      Sample best action in  $\mathcal{X}_{l+1}^{(-1)} \cup \mathcal{X}_{l+1}^{(+1)}$  for the remaining time
12:    else
13:       $(\mathcal{Z}_{l+1}, \hat{\Delta}^{l+1}) \leftarrow \Delta\text{-EXPLORE-AND-ELIMINATE} \left( \mathcal{X}, \left( \mathcal{X}_{l+1}^{(z)}, \hat{\theta}_l^{(z)} \right)_{z \in \{-1, 1\}}, \hat{\Delta}^l, \frac{2}{\epsilon_l^2} \log \left( \frac{l(l+1)}{\delta} \right), \epsilon_l \right)$ 
14:    end if
15:  end if
16: end for
```

4 Upper bound on the worst-case regret of Fair Phased Elimination

In this section, we provide an upper bound on the worst-case regret of Algorithm 4, showing that it is asymptotically $O\left(\kappa_*^{1/3} T^{2/3} \log(T)^{1/3}\right)$ as T goes to infinity.

Theorem 2. *For the choice $\delta = T^{-1}$, there exists two numerical constants $C, C' > 0$ such that the following bound on the regret of the FAIR PHASED ELIMINATION algorithm 4 holds*

$$\begin{aligned} R_T &\leq C \left(\kappa_*^{1/3} T^{2/3} \log(T)^{1/3} + (d \vee \kappa_*) \log(T) + d^2 + d \kappa_*^{-1/3} T^{1/3} \log(kT) \log(T)^{-1/3} \right) \\ &\leq C' \kappa_*^{1/3} T^{2/3} \log(T)^{1/3}, \quad \text{for } T \geq \frac{((d \vee \kappa_*)^{3/2} \log(T)) \vee d^3}{\sqrt{\kappa_*}} \vee \frac{(d \log(kT))^3}{(\kappa_* \log(T))^2}. \end{aligned}$$

Before discussing the bounds established in Theorem 2, we briefly present a sketch of this result. Detailed proofs are postponed to Appendix C.1.

Sketch of Proof. We begin this proof by introducing some notations.

Notations We denote by L_T the largest integer l such that $\epsilon_l \geq \kappa_*^{1/3} T^{-1/3} \log(T)^{1/3}$. We denote by $L^{(0)}$ the last phase where $\hat{\Delta}^l$ -optimal Exploration and Elimination happens. We denote by $\text{Exp}_l^{(z)}$ the time indices where G-exploration is performed on $\mathcal{X}_l^{(z)}$ and by $\text{Exp}_l^{(0)}$ the time indices where Δ -exploration is performed at phase l . We also denote by Recovery the time indices subsequent to the stopping criteria, this set being empty when the stopping criteria is not activated.

We define a "good" event $\bar{\mathcal{F}}$ such that the errors $\left| a_x^\top (\theta^* - \hat{\theta}_l) \right|$ and $|\omega^* - \hat{\omega}_l^{(0)}|$ are smaller than ϵ_l for all l such that these quantities are defined, and all $x \in \mathcal{X}_l^{(-1)}$ and $\mathcal{X}_l^{(+1)}$. In the following, we use c, c' to denote positive absolute constants, which may vary from line to line. With these notations, we decompose the regret as follows :

$$\begin{aligned}
R_T \leq & \Delta_{\max} T \mathbb{P}(\mathcal{F}) + \mathbb{E}_{|\overline{\mathcal{F}}}\left[\underbrace{\sum_{l \leq L_T} \sum_{z \in \{-1, +1\}} \sum_{t \in \text{Exp}_l^{(z)}} (x^* - x_t)^\top \gamma^*}_{R_T^G}\right] + \mathbb{E}_{|\overline{\mathcal{F}}}\left[\underbrace{\sum_{l \leq L^{(0)}} \sum_{t \in \text{Exp}_l^{(0)}} (x^* - x_t)^\top \gamma^*}_{R_T^\Delta}\right] \\
& + \mathbb{E}_{|\overline{\mathcal{F}}}\left[\underbrace{\sum_{l \geq L_T+1} \sum_{z \in \{-1, +1\}} \sum_{t \in \text{Exp}_l^{(z)}} (x^* - x_t)^\top \gamma^* + \sum_{t \in \text{Recovery}} (x^* - x_t)^\top \gamma^*}_{R_T^{\text{Rec}}}\right].
\end{aligned}$$

Bound on $\Delta_{\max} T \mathbb{P}(\mathcal{F})$. Using arguments based on concentration of Gaussian variables, we show that $\mathbb{P}(\mathcal{F}) \leq 2T^{-1}$. Since $\Delta_{\max} \leq 2$, we find that $\Delta_{\max} T \mathbb{P}(\mathcal{F}) \leq 4$.

Bound on R_T^G . We show that on $\overline{\mathcal{F}}$, only actions with gaps smaller than $c\epsilon_l$ remain in the sets $\mathcal{X}_l^{(-1)}$ and $\mathcal{X}_l^{(+1)}$. The length of each G-optimal Exploration and Elimination phase is of the order $d \log(klT)/\epsilon_l^2$, so the regret of each phase is of the order $d \log(klT)/\epsilon_l$. Summing over the different phases, we find that

$$R_T^G \leq cd \log(kL_T T)/\epsilon_{L_T}. \quad (6)$$

Using the definition of L_T , we find that $R_T^G \leq cd \log(kL_T T) \kappa_*^{-1/3} \log(T)^{-1/3} T^{1/3}$.

Bound on R_T^Δ . We show that on \mathcal{F} , $\hat{\Delta}^l \geq \Delta$ for all $l \geq 1$. Then, our choice of design $\mu_l^{(0)}$ ensures that for $l \leq L^{(0)}$, on $\overline{\mathcal{F}}$,

$$\sum_{t \in \text{Exp}_l^{(0)}} (x^* - x_t)^\top \gamma^* \leq c \left(\frac{\log(l(l+1)T)}{\epsilon_l^2} \kappa(\hat{\Delta}^l) + d + 1 \right)$$

for some constant $c > 0$. Summing over the different phases, we find that

$$R_T^\Delta \leq c\kappa(\hat{\Delta}^{L^{(0)}}) \log(L^{(0)}T)/\epsilon_{L^{(0)}}^2. \quad (7)$$

Now, the algorithm does not enter the Recovery phase before phase $L^{(0)} + 1$, so we must have $\epsilon_{L^{(0)}}^{-2} \leq T^{2/3} \log(T)^{-2/3} \kappa(\hat{\Delta}^{L^{(0)}})^{-2/3}$. This implies that $R_T^\Delta \leq c\kappa(\hat{\Delta}^{L^{(0)}})^{1/3} \log(T)^{1/3} T^{2/3}$. Since $\kappa(\hat{\Delta}^l) \leq \kappa_*$, we find that $R_T^\Delta \leq c' \kappa_*^{1/3} \log(T)^{1/3} T^{2/3}$.

Bound on R_T^{Rec} . On the one hand, the actions selected during the Phases $\text{Exp}_l^{(-1)}$ and $\text{Exp}_l^{(+1)}$ for $l \geq L_T + 1$ are sub-optimal by a gap at most $c\epsilon_{L_T}$ on the event $\overline{\mathcal{F}}$. On the other hand, if the algorithm enters the Recovery phase at a phase l , then

$$\epsilon_l \leq \kappa(\hat{\Delta}^{L^{(0)}})^{1/3} T^{-1/3} \log(T)^{1/3} \leq \kappa_*^{1/3} T^{-1/3} \log(T)^{1/3},$$

so we must have $l = L^{(0)} + 1 \geq L_T + 1$. Therefore, all actions selected during the Recovery phase are sub-optimal by a gap at most $c\epsilon_{L_T}$. Then, R_T^{Rec} can be bounded as $R_T^{\text{Rec}} \leq c\epsilon_{L_T} T$. This implies in particular that $R_T^{\text{Rec}} \leq c' \kappa_*^{1/3} \log(T)^{1/3} T^{2/3}$.

When $T \geq T_{\kappa_*, d, k}$ for some $T_{\kappa_*, d, k}$ large enough, we find that $\mathbb{R}_T \leq c' \kappa_*^{1/3} \log(T)^{1/3} T^{2/3}$. □

Theorem 2 shows that the worst-case regret of the Fair Phased Elimination algorithm asymptotically grows as $C\kappa_*^{1/3} T^{2/3} \log(T)^{1/3}$. In Section 6.2, we show that this bound is sharp in some settings up to the sub-logarithmic factor $\log(T)^{1/3}$.

The worst-case regret rate of biased linear bandits is higher than the typical rate $Cd \log(T) T^{1/2}$ obtained under unbiased feedback on the rewards (see, e.g., [Abbasi-yadkori et al., 2011]). This increase in the regret

corresponds to the cost of learning from unfair evaluations. It is due to the fact that the algorithm may need to sample actions that are very sub-optimal in order to estimate the bias parameter, and distinguish between two actions of different groups with similar rewards.

Previous bounds on the worst-case regret, established by the authors of [Kirschner et al., 2020] for linear bandit with partial monitoring, rely on a different measure of the complexity of estimating the unbiased reward (called the worst-case alignment constant) :

$$\alpha = \max_{\gamma \in \mathbb{R}^d} \frac{\max_{x, x' \in \mathcal{X}} \left((x - x')^\top \gamma \right)^2}{\max_{\tilde{x} \in \mathcal{X}} \left((z_{\tilde{x}} \tilde{x})^\top \gamma + 1 \right)^2}.$$

The authors of [Kirschner et al., 2020] show that the worst-case regret of their Information Directed Sampling algorithm is bounded by $C\alpha^{1/3}d^{1/2}T^{2/3}\log(kT)^{1/2}$, where C is a numerical constant. We underline that this bound depends on the dimension d , the geometrical constant α , and on the number of actions k , while the large T bound of Theorem 2 only depends on the geometry of the action set through κ_* . It is unclear whether the dependence on $\alpha^{1/3}$, $d^{1/2}$, or even $\log(kT)^{1/2}$ is sharp in some setting, since the corresponding lower bound established in [Kirschner et al., 2020] is $cT^{2/3}$, where the constant c depends on the geometry of $\left\{ \begin{pmatrix} x \\ z_x \end{pmatrix} : x \in \mathcal{X} \right\}$ and is not explicit. Note that our upper bounds grows as $\log(T)^{1/3}$ instead of $\log(T)^{1/2}$.

By contrast, the present paper addresses the related problem of biased linear bandits. Using an algorithm specifically design for this problem, we establish regret bounds that are sharp in some settings in terms of dependence on κ_* . This shows that the constant $\kappa_*^{1/3}$ is a relevant measure of the difficulty of the linear bandit problem for a given set of actions \mathcal{X} .

We underline that the regret rates only depends on the geometry of the problem through the constant κ_* . This constant is defined as $\kappa_* = \Delta_{\max} v$, where v is the variance corresponding to the e_{d+1} -optimal design. While optimal design has been studied extensively, and the connection between G-optimal design and linear bandit algorithm already exploited, it is to the best of our knowledge the first time that c -optimal design is applied to this problem. Interestingly, the constant κ_* also appears naturally when considering the related problem of optimal bias estimation against the worst parameter.

Recall that if we use a design μ to sample actions during the Δ -optimal Exploration and Exploitation Phase, the regret of is given by $\sum_x \mu(x) \Delta_x$, where $\Delta_x = \max_{x' \in \mathcal{X}} (x' - x)^\top \gamma^*$. If we don't know the gaps Δ_x , or equivalently if we don't know the parameter γ^* , we can try to obtain a low regret against the worst parameter γ . Note that, for a given design μ , this worst-case regret is given by

$$\max_{x' \in \mathcal{X}, \gamma \in \mathcal{C}(\mathcal{X})} \sum_x \mu(x) (x' - x)^\top \gamma,$$

where we recall that $\mathcal{C}(\mathcal{X})$ is the set of admissible parameters. To achieve the lowest regret against the worst parameter, we must use the design $\tilde{\mu}$ solution to the problem

$$\begin{aligned} \tilde{\mu} &= \underset{\mu}{\operatorname{argmin}} \max_{x' \in \mathcal{X}, \gamma \in \mathcal{C}(\mathcal{X})} \sum_{x \in \mathcal{X}} \mu(x) (x' - x)^\top \gamma \\ &\text{such that } \mu \in \mathcal{M}_{e_{d+1}}^{\mathcal{X}} \\ &\text{and } e_{d+1}^\top V(\mu)^+ e_{d+1} \leq 1. \end{aligned} \tag{8}$$

The following lemma shows that e_{d+1} -optimal design is related to the problem of minimizing the regret for estimating the bias against the worst parameter.

Lemma 5. *Let $\tilde{\mu}$ be defined as in Equation (8), and let us define the worst case regret $\tilde{\kappa}$ of the design $\tilde{\mu}$ as $\tilde{\kappa} = \max_{\gamma \in \mathcal{C}(\mathcal{X}), x' \in \mathcal{X}} \sum_{x \in \mathcal{X}} \tilde{\mu}(x) (x' - x)^\top \gamma$. Then, $\tilde{\kappa} \leq \kappa_* \leq 2\tilde{\kappa}$.*

Lemma 5 underlines that the regret corresponding to the e_{d+1} -optimal design is no larger than twice the regret of the best design against the worst parameter. Thus, the two problems are essentially equivalent.

5 Gap-dependent regret bounds

In Section 4, we established an upper bound on the regret for a given budget T , for the worst possible parameter θ^* . However, for a given parameter θ^* , if the best action x^* is unique, the Fair Phased Elimination algorithm becomes able to identify it as the budget T increases, subsequently having 0 regret for the remaining rounds. Note that the time necessary to identify the best group intuitively depends on the gap between the best actions of the two groups $\Delta_{\neq} = \min_{x \in \mathcal{X}, z_x = -z_{x^*}} (x^* - x)^\top \theta^*$, and on the minimal regret for estimating the bias $\kappa(\Delta)$. Similarly, the time necessary to identify the best action depends on the gap between the best action and the second best action $\Delta_{\min} = \min_{x \in \mathcal{X} \setminus \{x^*\}} (x^* - x)^\top \theta^*$, as well as on the dimension d . In this section, we establish a gap-dependent upper bound on the regret of Algorithm 4, i.e. a bound depending on Δ_{\neq} , $\kappa(\Delta)$, Δ_{\min} and d .

5.1 Upper bound on the gap-depend regret of Fair Phased Elimination

Theorem 3 provides a gap-dependent bound for the regret of the FAIR PHASED ELIMINATION algorithm 4.

Theorem 3. *Assume that $x^* \in \operatorname{argmax}_{x \in \mathcal{X}} x^\top \gamma^*$ is unique. Then, there exists two numerical constants $C, C' > 0$ such that, for the choice $\delta = T^{-1}$, the following bound on the regret of the FAIR PHASED ELIMINATION algorithm 4 holds*

$$\begin{aligned} R_T &\leq C \left(\left(\frac{d}{\Delta_{\min}} \vee \frac{\kappa(\Delta \vee \Delta_{\neq} \vee \varepsilon_T)}{\Delta_{\neq}^2} \right) \log(T) + d^2 + \frac{d}{\Delta_{\min}} \log(k) \right) \\ &\leq C' \left(\frac{d}{\Delta_{\min}} \vee \frac{\kappa(\Delta \vee \Delta_{\neq} \vee \varepsilon_T)}{\Delta_{\neq}^2} \right) \log(T) \quad \text{for } T \geq k \vee e^{d\Delta_{\min}} \end{aligned}$$

where $\Delta_{\min} = \min_{x \in \mathcal{X} \setminus x^*} \Delta_x$, $\Delta_{\neq} = \min_{x \in \mathcal{X}: z_x = -z_{x^*}} \Delta_x$, and $\varepsilon_T = \left(\frac{\kappa_* \log(T)}{T} \right)^{1/3}$.

Sketch of Proof. The proof of Theorem 3 is close to that of Theorem 2, and we adopt the same notations as in the proof sketch page 11.

Notations We denote by $L^{(0)}$ the last phase where $\widehat{\Delta}^l$ -optimal Exploration and Elimination happens. We denote $\overline{\mathcal{F}}$ some "good" event such that the errors $|a_x^\top(\theta^* - \widehat{\theta}_l^{(z_x)})|$ and $|\omega^* - \widehat{\omega}_l^{(0)}|$ are smaller than ϵ_l for all l such that these quantities are defined, and all $x \in \mathcal{X}_l^{(-1)} \cup \mathcal{X}_l^{(+1)}$. We denote by $\operatorname{Exp}_l^{(z)}$ the time indices where G-exploration is performed on $\mathcal{X}_l^{(z)}$ and by $\operatorname{Exp}_l^{(0)}$ the time indices where Δ -exploration is performed at phase l . We also denote by Recovery the time indices subsequent to the stopping criteria, this set being empty when the stopping criteria is not activated. In the following, we use c, c' to denote positive absolute constants, which may vary from line to line.

Fact 1 Let $l_{\Delta_{\min}}$ be the largest integer such that $\epsilon_{l_{\Delta_{\min}}} \geq C\Delta_{\min}$ for some well-chosen absolute constant $C > 0$. We show that on the good event $\overline{\mathcal{F}}$, no more than $l_{\Delta_{\min}}$ G-optimal Exploration and Elimination phases are needed to find the best action. For all phases $l \geq l_{\Delta_{\min}}$, the algorithm always chooses x^* , and suffers no regret.

Fact 2 We show that on the good event $\overline{\mathcal{F}}$, for each phase l , $\widehat{\Delta}^l \leq c(\Delta \vee \epsilon_l)$ for some constant c . Lemma 4 then implies that for all $l \leq L^{(0)}$ and all $\tau > 0$, $\kappa(\widehat{\Delta}^l) \leq c\kappa(\Delta \vee \epsilon_l) \leq c(1 + \epsilon_l\tau^{-1})\kappa(\Delta \vee \tau)$.

Fact 3 Let $l_{\Delta_{\neq}}$ be the largest integer such that $\epsilon_{l_{\Delta_{\neq}}} \geq C\Delta_{\neq}$ for some well-chosen absolute constant $C > 0$. On the good event $\overline{\mathcal{F}}$, if the algorithm enters the $\widehat{\Delta}^l$ -optimal Exploration and Elimination phase at round $l \geq l_{\Delta_{\neq}}$, we show that the algorithm finds the best group at this phase. This implies that $L^{(0)} \leq l_{\Delta_{\neq}}$.

Fact 4 We denote by L_T the largest integer l such that $\epsilon_l \geq (\kappa_* \log(T)/T)^{1/3}$. Since $\kappa_* \geq \kappa(\widehat{\Delta}^l)$ for all $l \geq 1$, we see that if the algorithm enters the Recovery phase, we must have $L_T \leq L^{(0)}$, and $\epsilon_{L_T} \leq \epsilon_{L^{(0)}} \approx \varepsilon_T$.

Using **Fact 1**, we find that the regret can be written as

$$\begin{aligned}
R_T \leq & \Delta_{\max} T \mathbb{P}(\mathcal{F}) + \mathbb{E}_{|\overline{\mathcal{F}}}\left[\underbrace{\sum_{l \leq l_{\Delta_{\min}}} \sum_{z \in \{-1, +1\}} \sum_{t \in \text{Exp}_l^{(z)}} (x^* - x_t)^\top \gamma^*}_{R_T^G} \right] \\
& + \mathbb{E}_{|\overline{\mathcal{F}}}\left[\underbrace{\sum_{l \leq L^{(0)}} \sum_{t \in \text{Exp}_l^{(0)}} (x^* - x_t)^\top \gamma^*}_{R_T^\Delta} \right] + \mathbb{E}_{|\overline{\mathcal{F}}}\left[\underbrace{\sum_{t \in \text{Recovery}} (x^* - x_t)^\top \gamma^*}_{R_T^{\text{Rec}}} \right].
\end{aligned}$$

Bound on R_T^G . We rely on arguments similar to those used in Equation (6) to show that $R_T^G \leq c(d+1) \log(kl_{\Delta_{\min}} T) \epsilon_{l_{\Delta_{\min}}}^{-1}$. Since $\epsilon_{l_{\Delta_{\min}}} \geq C \Delta_{\min}$, this implies that

$$R_T^G \leq \frac{c(d+1) \log(kl_{\Delta_{\min}} T)}{\Delta_{\min}} \leq \frac{c'd \log(T)}{\Delta_{\min}}$$

if $T \geq k$.

Bound on $R_T^\Delta + R_T^{\text{Rec}}$. We begin by bounding R_T^Δ . Recall that Equation (7) states that $R_T^\Delta \leq c\kappa(\widehat{\Delta}^{L^{(0)}}) \log(l_{L^{(0)}} T) \epsilon_{L^{(0)}}^{-2}$. Using **Fact 2**, we find that for any $\tau > 0$,

$$R_T^\Delta \leq c\kappa(\Delta \vee \tau) \log(l_{L^{(0)}} T) (\epsilon_{L^{(0)}}^{-2} + \epsilon_{L^{(0)}}^{-1} \tau^{-1}). \quad (9)$$

Let us now consider two cases, corresponding to $\text{Recovery} = \emptyset$ and $\text{Recovery} \neq \emptyset$.

Case 1: $\text{Recovery} = \emptyset$. On the one hand, our case assumption implies that

$$R_T^{\text{Rec}} = 0.$$

On the other hand, by **Fact 3**, we know that on $\overline{\mathcal{F}}$, $L^{(0)} \leq l_{\Delta_{\neq}}$. Then, using the definition of $l_{\Delta_{\neq}}$ and Equation (9) with $\tau = \Delta_{\neq}$, we find that

$$R_T^\Delta \leq c\kappa(\Delta \vee \Delta_{\neq}) \log(L^{(0)} T) \Delta_{\neq}^{-2}.$$

Case 2: $\text{Recovery} \neq \emptyset$. All actions selected during the Recovery phase belong to $\mathcal{X}_{L^{(0)}+1}^{(-1)} \cup \mathcal{X}_{L^{(0)}+1}^{(+1)}$, so on $\overline{\mathcal{F}}$ these actions are sub-optimal by a gap at most $c\epsilon_{L^{(0)}+1}$, so $R_T^{\text{Rec}} \leq cT\epsilon_{L^{(0)}+1}$. Now, since the algorithm enters the Recovery phase, we must have $\epsilon_{L^{(0)}+1} \leq (\kappa(\Delta^{L^{(0)}+1}) \log(T)/T)^{1/3}$, which implies that

$$R_T^{\text{Rec}} \leq \frac{c\kappa(\widehat{\Delta}^{L^{(0)}+1}) \log(T)}{\epsilon_{L^{(0)}+1}^2}.$$

Using **Fact 2** with $\tau = \epsilon_{L^{(0)}}$ together with Equation (9), we find that

$$R_T^\Delta + R_T^{\text{Rec}} \leq \frac{c\kappa(\Delta \vee \epsilon_{L^{(0)}}) \log(T)}{\epsilon_{L^{(0)}}^2}.$$

On the one hand, **Fact 3** guarantees that, since we entered the Recovery phase before finding the best group, we must have $\epsilon_{L^{(0)}} \leq \epsilon_{l_{\Delta_{\neq}}}$. On the other hand, **Fact 4** ensures that $\epsilon_{L^{(0)}} \geq \varepsilon_T$. Thus,

$$R_T^{\text{Rec}} \leq \frac{c\kappa(\Delta \vee \varepsilon_T) \log(T)}{\Delta_{\neq}^2}.$$

Conclusion Combining these results, we find that

$$R_T \leq c \left(\frac{d}{\Delta_{\min}} \vee \frac{\kappa(\Delta \vee \Delta_{\neq})}{\Delta_{\neq}^2} \vee \frac{\kappa(\Delta \vee \varepsilon_T)}{\Delta_{\neq}^2} \right) \log(T)$$

when $T \geq k$. Using Lemma 4, we get that $\kappa(\Delta \vee \Delta_{\neq}) \vee \kappa(\Delta \vee \varepsilon_T) \leq \kappa(\Delta \vee \Delta_{\neq} \vee \varepsilon_T)$, which concludes the proof of the results. \square

The term $\frac{d}{\Delta_{\min}} \vee \frac{\kappa(\Delta \vee \Delta_{\neq} \vee \varepsilon_T)}{\Delta_{\neq}^2}$ highlights the two sources of difficulty of the problem. On the one hand, the term $\frac{d}{\Delta_{\min}}$ is unavoidable in d -dimensional linear bandits [Abbasi-yadkori et al., 2011]. On the other hand, the term $\frac{\kappa(\Delta \vee \Delta_{\neq} \vee \varepsilon_T)}{\Delta_{\neq}^2}$ is characteristic of the linear bandit problem with systematic bias. Again, this term only depends on the dimension through the geometrical constant $\kappa(\Delta \vee \Delta_{\neq} \vee \varepsilon_T)$. Interestingly, when $\frac{d}{\Delta_{\min}} > \frac{\kappa(\Delta \vee \Delta_{\neq} \vee \varepsilon_T)}{\Delta_{\neq}^2}$, the difficulty of the linear bandit problem with systematic bias is dominated by that of the classical d -linear bandit problem. In this case, the algorithm is able to find the group containing the best action, and the problem reduces to a linear bandit problem in d dimension. Thus, the linear bandit with systematic bias is an example of a globally observable game that can be locally observable around the best action (thus answering a question by the authors of [Kirschner et al., 2020] on the existence of such games).

When $\frac{d}{\Delta_{\min}} \leq \frac{\kappa(\Delta \vee \Delta_{\neq} \vee \varepsilon_T)}{\Delta_{\neq}^2}$, the regret corresponds to the regret of the bias estimation phase. In other words, when both groups contain near-optimal actions, the difficulty of the problem is dominated by the price to pay for debiasing the unfair evaluations.

Next, we show that the gap-dependent bound on the regret provided in Theorem 2 and 3 are tight in some settings, respectively up to sub-logarithmic and constant multiplicative factors. In the following, we show that there exists v -correlated action sets such that any algorithm must suffer a regret of order $d/\Delta_{\min} + \kappa(\Delta)/\Delta_{\neq}^2$ for some parameter θ , where Δ_{\min} is the gap between the reward of the best and second best action, and Δ_{\neq} is the gap between the rewards of the best actions of the two groups.

6 Lower bounds on the regret

In this section, we derive lower bounds on the worst-case regret and the gap-dependent regret that respectively match the upper bounds established in Theorem 2 and 3 up to sub-logarithmic factors or numerical constants. We begin by presenting the gap-dependent lower bound, and derive the worst-case bound as a corollary. Before stating our result, we underline that the difficulty of linear bandits is highly dependent on the geometry of the action set. For this reason, bounds established in the literature adopt either of two complementary approaches. The first approach, adopted for example by the authors of [Kirschner et al., 2020], aims at establishing results that give for any set of actions \mathcal{X} a lower bound on regret of the form $c_{\mathcal{X}} f(T)$, where $c_{\mathcal{X}}$ is a non-explicit constant depending on the set of actions \mathcal{X} , and $f(T)$ is a function depending only on T . The second approach, to which this work belongs, chooses to establish lower bounds for some given set of actions \mathcal{X} that are explicit in terms of their dependence on \mathcal{X} and T . These results can then be used to show that certain upper bounds cannot be improved uniformly on all problems, and that certain algorithms are optimal for given sets of actions.

The upper bounds established previously underline the dependence of the regret on the geometry of the action set. Before stating our result, we therefore begin by introducing the notion of v -correlated action set.

Definition 1 (v -correlated action set). *For $v \geq 1$, a set of actions \mathcal{A} is v -correlated if $\mathcal{A} \in \mathbf{A}_{v,d}$, where*

$$\mathbf{A}_{v,d} = \left\{ \mathcal{A} = \left\{ \begin{pmatrix} x_1 \\ z_{x_1} \end{pmatrix}, \dots, \begin{pmatrix} x_k \\ z_{x_k} \end{pmatrix} \right\} \subset (\mathbb{R}^{(d+1)} \times \{-1, +1\})^k : \right. \\ \left. k \in \mathbb{N}^*, \min_{\pi \in \mathcal{P}_{e_{d+1}}^{\mathcal{A}}} \left\{ e_{d+1}^{\top} \left(\sum_{\begin{pmatrix} x \\ z \end{pmatrix} \in \mathcal{A}} \pi(x) \begin{pmatrix} x \\ z_x \end{pmatrix} \begin{pmatrix} x \\ z_x \end{pmatrix}^{\top} \right) e_{d+1} \right\} \geq v \right\}$$

is the set of actions sets such that the minimal variance of the bias estimator is larger than v .

In the following, we establish a lower bound on the regret valid for all $v \geq 1$. More precisely, for all $v \geq 1$, and all $d \geq 4$, we design a v -correlated set of actions $\mathcal{A} \in \mathbf{A}_{v,d}$, and we obtain lower bounds on the regret of any algorithm on this set of actions.

6.1 Lower bound on the gap-dependent regret

In this section, we present a lower bound on the gap-dependent regret. More precisely, for given values of Δ_{\min} and Δ_{\neq} , we establish a lower bound on the worst case regret among parameters θ verifying $\Delta_{\min} \leq \min_{x \in \mathcal{X} \setminus x^*} \Delta_x$, and $\Delta_{\neq} \leq \min_{x \in \mathcal{X} : z_x = -z_{x^*}} \Delta_x$. Before stating formally the result, let us define the corresponding parameter set. For an action set $\mathcal{A} \in \mathbf{A}_{v,d}$, and for $(\Delta_{\min}, \Delta_{\neq}) \in (0, 1)^2$ such that $\Delta_{\min} \leq \Delta_{\neq}$, let

$$\Theta_{\Delta_{\min}, \Delta_{\neq}}^{\mathcal{A}} = \left\{ \theta = \begin{pmatrix} \gamma \\ \omega \end{pmatrix} : \gamma \in \mathcal{C}(\mathcal{X}), \exists ! \begin{pmatrix} x^* \\ z_{x^*} \end{pmatrix} \in \operatorname{argmax}_{\begin{pmatrix} x \\ z_x \end{pmatrix} \in \mathcal{A}} \{x^\top \gamma\}, \right. \\ \left. \begin{aligned} &\forall \begin{pmatrix} x' \\ z_{x'} \end{pmatrix} \in \mathcal{A} \text{ such that } x' \neq x^*, (x^* - x')^\top \gamma \geq \Delta_{\min}, \\ &\forall \begin{pmatrix} x' \\ z_{x'} \end{pmatrix} \in \mathcal{A} \text{ such that } z_{x'} \neq z_{x^*}, (x^* - x')^\top \gamma \geq \Delta_{\neq} \end{aligned} \right\}$$

be the set of parameters with minimum gap Δ_{\min} , and with gap Δ_{\neq} between the best rewards of the two groups.

Theorem 4. *For all $v \geq 1$ and all $d \geq 4$, there exists an action set $\mathcal{A} \in \mathbf{A}_{v,d}$, such that for all bandit algorithms, for all $(\Delta_{\min}, \Delta_{\neq}) \in (0, 1/8)^2$ with $\Delta_{\min} \leq \Delta_{\neq}$, and for all budget $T \geq 2$, there exists a problem characterized by $\theta \in \Theta_{\Delta_{\min}, \Delta_{\neq}}^{\mathcal{A}}$ such that the regret of the algorithm on the problem satisfies*

$$R_T^{\mathcal{A}, \theta} \geq \left[\frac{d}{10\Delta_{\min}} \log(T) \left[1 - \frac{\log\left(\frac{8d \log(T)}{\Delta_{\min}^2}\right)}{\log(T)} \right] \right] \vee \left[\frac{v+1}{4\Delta_{\neq}^2} \log(T) \left[1 - \frac{\log\left(\frac{8v \log(T)}{\Delta_{\neq}^3}\right)}{\log(T)} \right] \right] \\ \vee \left[\frac{v}{4\Delta_{\neq}^2} \left[1 \wedge \log\left(\frac{T\Delta_{\neq}^3}{8v}\right) \right] \right]. \quad (10)$$

Moreover, on this problem, $\kappa(\Delta) \in [\frac{v}{8}, 2v]$.

Remark 1. *The assumption $d \geq 4$ is necessary in our proof to design an action set \mathcal{A} such that Equation (10) holds for all $\Delta_{\min}, \Delta_{\neq} \in (0, 1/8)$. On the other hand, for $d \in \{2, 3\}$, for all $\Delta_{\min}, \Delta_{\neq} \in (0, 1/8)$, we can show that there exists action sets \mathcal{A} and $\theta \in \Theta_{\Delta_{\min}, \Delta_{\neq}}^{\mathcal{A}}$ such that the lower bound in Equation (10) still holds, by considering separately the cases $\frac{d}{\Delta_{\min}} > \frac{v}{\Delta_{\neq}^2}$ and $\frac{d}{\Delta_{\min}} \leq \frac{v}{\Delta_{\neq}^2}$.*

Let us discuss the implications of Theorem 4, before sketching the proof of this result. Theorem 4 provides a non-asymptotic lower bound on the regret on some action set \mathcal{A} for all $T \geq 1$ and all $\Delta_{\min}, \Delta_{\neq} \in (0, 1/8)^2$ that depends on the e_{d+1} -optimal variance v . To compare it with the upper bound provided by Theorem 3, we consider asymptotic regimes where $T \rightarrow \infty$, and the gaps Δ_{\min} and Δ_{\neq} are fixed. Moreover, we note that the lower bound of Theorem 4 corresponds to problems such that v and $\kappa(\Delta)$ and of the same order. Thus, the lower bound in Theorem 4 can be rephrased as a function of the Δ -optimal regret $\kappa(\Delta)$. To state this result, we define a class of problems $\Theta_{\Delta_{\min}, \Delta_{\neq}, \kappa}^{\mathcal{A}}$ such that $\kappa(\Delta) \leq \kappa$.

Let us abuse notations, and denote for a parameter $\gamma \in \mathcal{C}(\mathcal{X})$, $\Delta(\gamma)_x = \max_{x' \in \mathcal{X}} (x' - x)^\top \gamma$, and $\Delta(\gamma) = (\Delta(\gamma)_x)_{x \in \mathcal{X}}$. Moreover, for a given set \mathcal{A} , let us denote

$$\Theta_{\Delta_{\min}, \Delta_{\neq}, \kappa}^{\mathcal{A}} = \Theta_{\Delta_{\min}, \Delta_{\neq}}^{\mathcal{A}} \cap \left\{ \theta = \begin{pmatrix} \gamma \\ \omega \end{pmatrix} : \gamma \in \mathcal{C}(\mathcal{X}), \kappa(\Delta(\gamma)) \leq \kappa \right\}.$$

Corollary 1. *For all $\kappa \geq 2$ and all $d \geq 4$, there exists a set of actions $\mathcal{A} \in \mathbb{R}^{d+1}$ such that for all $(\Delta_{\min}, \Delta_{\neq}) \in (0, 1/8)^2$ with $\Delta_{\min} \leq \Delta_{\neq}$,*

$$\liminf_{T \rightarrow \infty} \sup_{\theta \in \Theta_{\Delta_{\min}, \Delta_{\neq}, \kappa}^{\mathcal{A}}} \frac{R_T^{\mathcal{A}, \theta}}{\log(T)} \geq \left[\frac{d}{10\Delta_{\min}} \right] \vee \left[\frac{\kappa + 2}{8\Delta_{\neq}^2} \right].$$

Corollary 1 shows that for some action sets \mathcal{A} , the gap-depend the regret of the Fair Phased Elimination algorithm is asymptotically optimal up to a numerical constant. Again, the assumption $d \geq 4$ can be relaxed if we consider different action sets for the cases $\frac{d}{\Delta_{\min}} \geq \frac{\kappa}{\Delta_{\neq}^2}$ and $\frac{d}{\Delta_{\min}} < \frac{\kappa}{\Delta_{\neq}^2}$.

Sketch of Proof. To prove Theorem 4, we need to construct two close problem instances with optimal actions belonging to different groups - to obtain the part of the lower bound involving Δ_{\neq} - and in addition we must also create confusing instances with different optimal actions belonging to a same group - to obtain the part of the lower bound involving Δ_{\min} . This is done by considering the following set of actions and of problems.

Set $\mathcal{A} = \left\{ \begin{pmatrix} x_1 \\ z_{x_1} \end{pmatrix}, \dots, \begin{pmatrix} x_{d+1} \\ z_{x_{d+1}} \end{pmatrix} \right\}$, where $\begin{pmatrix} x_i \\ z_{x_i} \end{pmatrix} = e_i + e_{d+1}$, for $i \in \{2, \dots, \lfloor d/2 \rfloor\}$, $\begin{pmatrix} x_i \\ z_{x_i} \end{pmatrix} = e_i - e_{d+1}$ for $i \in \{\lfloor d/2 \rfloor + 1, \dots, d\}$, and $\begin{pmatrix} x_{d+1} \\ z_{x_{d+1}} \end{pmatrix} = -\left(1 - \frac{2}{\sqrt{v}+1}\right) e_1 - e_{d+1}$. Then, Lemma 6 shows that $\mathcal{A} \in \mathbf{A}_{v,d}$.

Lemma 6. *It holds that*

$$\min_{\pi \in \mathcal{P}_{e_{d+1}}^{\mathcal{A}}} \left\{ e_{d+1}^\top \left(\sum_{\begin{pmatrix} x \\ z \end{pmatrix} \in \mathcal{A}} \pi(x) \begin{pmatrix} x \\ z \end{pmatrix} \begin{pmatrix} x \\ z \end{pmatrix}^\top \right)^+ e_{d+1} \right\} = v.$$

We also define the following parameters:

$$\begin{aligned} \gamma^{(1)} &= \frac{1 + \Delta_{\neq} - \Delta_{\min}}{2} \left(\sum_{1 \leq j \leq \lfloor d/2 \rfloor} e_j \right) + \frac{1 - \Delta_{\neq} - \Delta_{\min}}{2} \left(\sum_{\lfloor d/2 \rfloor + 1 \leq j \leq d} e_j \right) + \Delta_{\min} e_1 + \Delta_{\min} e_{\lfloor d/2 \rfloor + 1} \\ \gamma^{(i)} &= \gamma^{(1)} + 2\Delta_{\min} e_i + 2\Delta_{\min} e_{\lfloor d/2 \rfloor + i} \quad \forall i \in \{2, \dots, \lfloor d/2 \rfloor\} \\ \gamma^{(\lfloor d/2 \rfloor + 1)} &= \frac{1 - \Delta_{\neq} - \Delta_{\min}}{2} \left(\sum_{1 \leq j \leq \lfloor d/2 \rfloor} e_j \right) + \frac{1 + \Delta_{\neq} - \Delta_{\min}}{2} \left(\sum_{\lfloor d/2 \rfloor + 1 \leq j \leq d} e_j \right) + \Delta_{\min} e_1 + \Delta_{\min} e_{\lfloor d/2 \rfloor + 1}. \end{aligned}$$

The bias parameters are given by $\omega^{(i)} = -\frac{\Delta_{\neq}}{2} \quad \forall i \in \{1, \dots, \lfloor d/2 \rfloor\}$, and $\omega^{(\lfloor d/2 \rfloor + 1)} = \frac{\Delta_{\neq}}{2}$. The parameters $\theta^{(i)} = \begin{pmatrix} \gamma^{(i)} \\ \omega^{(i)} \end{pmatrix}$ characterize $\lfloor d/2 \rfloor + 1$ problems, with noise distribution i.i.d. $\mathcal{N}(0, 1)$. We write **Problem i** for the problem characterized by $\theta^{(i)}$. Note that by construction and for any $i \in \{1, \dots, \lfloor d/2 \rfloor + 1\}$, we have that $\theta^{(i)} \in \Theta_{\Delta_{\min}, \Delta_{\neq}}^{\mathcal{A}}$.

The following facts hold:

- For any $i \in \{1, \dots, \lfloor d/2 \rfloor + 1\}$, action x_i is the unique optimal action in **Problem i**. Since $1/2 \geq \Delta_{\neq} \geq \Delta_{\min}$, sampling any other (sub-optimal) action leads to an instantaneous regret of at least Δ_{\min} . Moreover, choosing an action in the group $-z_i$ leads to an instantaneous regret of at least Δ_{\neq} .
- In **Problem i** for any $i \in \{1, \dots, \lfloor d/2 \rfloor + 1\}$, action $d+1$ is very sub-optimal and sampling it leads to an instantaneous regret higher than $(1 - 2/(\sqrt{v}+1))(1 - \Delta_{\neq} + \Delta_{\min}) + (1 + \Delta_{\neq} + \Delta_{\min})/2 \geq 1/2$, since $v \geq 1$ and $1/2 \geq \Delta_{\neq} \geq \Delta_{\min}$. This action is the worst action in all problems.
- Many actions are such that their distributions are the same across problems. More specifically:
 - For any $i \in \{2, \dots, \lfloor d/2 \rfloor\}$, between **Problem 1** and **Problem i**, the only actions that provide different evaluations when sampled are action i and action $\lfloor d/2 \rfloor + i$, and the mean difference between the evaluations in both cases is $2\Delta_{\min}$.
 - Between **Problem 1** and **Problem** $\lfloor d/2 \rfloor + 1$, the only actions that provide different evaluations when sampled is action $d+1$, and the mean gap in this case is $\frac{2}{\sqrt{v}+1} \Delta_{\neq} := \alpha \Delta_{\neq}$.

The proof is then divided in two parts, one part for proving the part of the bound depending on Δ_{\min} and one part for proving the part of the bound depending on Δ_{\neq} .

Part of the bound depending on Δ_{\min} . This part of the proof is obtained using classical arguments for K -armed bandit problems. For $i \in \{2, \dots, \lfloor d/2 \rfloor\}$, all actions but x_i and $x_{\lfloor d/2 \rfloor + i}$ have the same feedback under **Problem 1** and **Problem i**. On the other hand, the average feedback for actions x_i and $x_{\lfloor d/2 \rfloor + i}$ differs by $2\Delta_{\min}$, so either action needs to be selected approximately $\frac{\log(T)}{\Delta_{\min}^2}$ times in order to identify the problem at hand with high enough probability. In **Problem 1**, the simple regret for choosing x_i or $x_{\lfloor d/2 \rfloor + i}$ is larger than Δ_{\min} , so the total regret obtained when doing this is at least of the order $\frac{\log(T)}{\Delta_{\min}}$. Summing over the different actions i leads to a lower bound of the order $\frac{d \log(T)}{\Delta_{\min}}$.

Part of the bound depending on Δ_{\neq} . To obtain the second part of the lower bound, we note that all actions but x_{d+1} have the same feedback under **Problem 1** and **Problem $\lfloor d/2 \rfloor + 1$** . The average feedback for actions x_{d+1} differs by $\alpha\Delta_{\neq}$ under these parameters, so action x_{d+1} needs to be selected approximately $\frac{\log(T)}{\alpha^2 \Delta_{\neq}^2} \gtrsim \frac{\log(T)v}{\Delta_{\neq}^2}$ times to identify the problem at hand with high enough probability. Since selecting action x_{d+1} leads to an simple regret larger than $1/2$ under **Problem 1**, this implies that the regret must be at least of the order $\frac{v \log(T)}{\Delta_{\neq}^2}$.

Bounds on $\kappa(\Delta)$ Finally, the following lemma allows to express $\kappa(\Delta)$ as a function of v and Δ_{\max} .

Lemma 7. *For any $i \in \{1, \dots, \lfloor d/2 \rfloor + v1\}$, the gap vector Δ verifies*

$$\kappa(\Delta) = \frac{(1 + \sqrt{v})^2 \Delta_{\max}}{4}.$$

On the one hand, since $v \geq 1$, we see that $v \leq (1 + \sqrt{v})^2 \leq 4v$. On the other hand, $1/2 \leq \Delta_{\max} \leq 2$, so $\kappa(\Delta) \in [\frac{v}{8}, 2v]$. \square

6.2 Lower bound on the worst-case regret

To obtain a lower bound on the worst-case regret, we can rely on Theorem 4 and build, for a well-chosen set of actions, a sequence of parameters θ_T with $\theta_T \in \Theta_{\rho_T, \rho_T}^A$ for ρ_T of order $(v/T)^{1/3}$ such that the regret on the problem θ_T is at least of order $\kappa_*^{1/3} T^{2/3}$ for all T .

Corollary 2. *Let $v \geq 1$ and $d \geq 4$. There exists an action set $\mathcal{A} \in \mathbf{A}_{v,d}$ such that for all algorithms, and all $T \geq 8^4 v$, there exists a bandit problem with parameter $\theta_T \in \mathbb{R}^{d+1}$ such that the regret of this algorithm satisfies*

$$R_T^{\mathcal{A}, \theta_T} \geq \frac{v^{1/3} T^{2/3}}{32},$$

and such that $\kappa_* \in [v/8, 2v]$. This implies in particular that

$$R_T^{\mathcal{A}, \theta_T} \geq \frac{\kappa_*^{1/3} T^{2/3}}{50}.$$

Remark 2. *Corollary 2 follows from Theorem 4, which holds for $d \geq 4$. For smaller values of d (namely, $d \geq 2$), Theorem 5 proves that for a well-chosen set of actions \mathcal{A} with optimal variance v , the worst case regret of any algorithm satisfies $R_T^{\mathcal{A}, \theta_T} \geq \frac{1}{8e} v^{1/3} T^{2/3}$. The formal statement of this result, all well as the proofs, are postponed to Appendix B.*

Note that the assumption $d \geq 2$ is necessary here : if $d = 1$, there are at most two potentially optimal actions (namely, $\max\{x : x \in \mathcal{X}\}$ and $\min\{x : x \in \mathcal{X}\}$). Then, the problem becomes locally observable, and regret rates of the order $\tilde{O}(T^{1/2})$ can be achieved, as shown by the authors of [Kirschner et al., 2020]. Similar lower bounds have already been established in [Kirschner et al., 2020] : the authors show that when the problem is not locally observable, any algorithm must have a regret at least $c_{\mathcal{A}} T^{2/3}$ for some parameter $\theta_T \in \mathbb{R}^{d+1}$, where $c_{\mathcal{A}} > 0$ is a constant depending (not explicitly) on \mathcal{A} . Those results are complementary in some sense: Corollary 2 provides a bound which is sharp up to a factor equal to $\log(T)^{1/3}$, and depends on an explicit problem-dependent constant $c_{\mathcal{A}}$, namely $c \kappa_*^{1/3}$ for some numerical constant $c > 0$, albeit this bound is only valid for some specific action sets \mathcal{A} .

7 Conclusion

In this paper, we addressed the problem of online decision making under biased bandit feedback. We underlined its relation with partial monitoring with linear feedback and linear reward. We designed a new algorithm based on Δ - and G-optimal design, and obtained worst-case and gap-dependent upper bounds on its regret. We obtained lower bounds on the regret for some problem instances, thus showing that these rates are tight up to sub-logarithmic factors in some settings. These rates highlight two behaviors : on the one hand, the worst case rate $\mathcal{O}(\kappa_*^{1/3} \log(T)^{1/3} T^{2/3})$ highlights the cost induced by the biased feedback, and the need to select sub-optimal actions in order to debias it. On the other hand, the gap-depend bound shows that for some instance, the problem can be locally observable around the best action : then, the difficulty of the problem is dominated by the difficulty of the corresponding linear bandit problem, and is no more difficult than this problem. When this is not the case, the regret scales as Δ_{\neq}^{-2} , where Δ_{\neq} is the gap between the best actions of the two groups. This work paves the way for considering other bandit models with unfair feedback, considering for example continuous, multi-dimensional sensitive attributes.

Acknowledgements.

The authors would like to thank Evgenii Chzhen and Nicolas Verzelen for their valuable discussions and suggestions.

The work of A. Carpentier is partially supported by the Deutsche Forschungsgemeinschaft (DFG) Emmy Noether grant MuSyAD (CA 1488/1-1), by the DFG - 314838170, GRK 2297 MathCoRe, by the FG DFG, by the DFG CRC 1294 'Data Assimilation', Project A03, by the Forschungsgruppe FOR 5381 "Mathematical Statistics in the Information Age - Statistical Efficiency and Computational Tractability", Project TP 02, by the Agence Nationale de la Recherche (ANR) and the DFG on the French-German PRCI ANR ASCAI CA 1488/4-1 "Aktive und Batch-Segmentierung, Clustering und Seriation: Grundlagen der KI" and by the UFA-DFH through the French-German Doktorandenkolleg CDFA 01-18 and by the SFI Sachsen-Anhalt for the project RE-BCI. Christophe Giraud received partial support by grant ANR-19-CHIA-0021-01 ("BiSCottE", Agence Nationale de la Recherche) and by the ANR and the DFG on the French-German PRCI ANR-21-CE23-0035 (ASCAI)

References

- [Abbasi-yadkori et al., 2011] Abbasi-yadkori, Y., Pál, D., and Szepesvári, C. (2011). Improved algorithms for linear stochastic bandits. In Shawe-Taylor, J., Zemel, R., Bartlett, P., Pereira, F., and Weinberger, K. Q., editors, *Advances in Neural Information Processing Systems*, volume 24. Curran Associates, Inc.
- [Barik and Honorio, 2021] Barik, A. and Honorio, J. (2021). Fair sparse regression with clustering: An invex relaxation for a combinatorial problem. In *Thirty-Fifth Conference on Neural Information Processing Systems*.
- [Byanjankar et al., 2015] Byanjankar, A., Heikkilä, M., and Mezei, J. (2015). Predicting credit risk in peer-to-peer lending: A neural network approach. In *2015 IEEE Symposium Series on Computational Intelligence*, pages 719–725.
- [Celis et al., 2018] Celis, L. E., Kapoor, S., Salehi, F., and Vishnoi, N. K. (2018). An algorithmic framework to control bias in bandit-based personalization.
- [Chaudhuri and Tewari, 2016] Chaudhuri, S. and Tewari, A. (2016). Phased exploration with greedy exploitation in stochastic combinatorial partial monitoring games. In Lee, D., Sugiyama, M., Luxburg, U., Guyon, I., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc.
- [Chawla and Jagadeesan, 2021] Chawla, S. and Jagadeesan, M. (2021). Individual fairness in advertising auctions through inverse proportionality.

- [Chzhen and Schreuder, 2020] Chzhen, E. and Schreuder, N. (2020). A minimax framework for quantifying risk-fairness trade-off in regression.
- [Claire et al., 2020] Claire, H., Chen, Y., Modi, J., Jung, M. F., and Nikolaidis, S. (2020). Multi-armed bandits with fairness constraints for distributing resources to human teammates. *2020 15th ACM/IEEE International Conference on Human-Robot Interaction (HRI)*, pages 299–308.
- [Elfving, 1952] Elfving, G. (1952). Optimum Allocation in Linear Regression Theory. *The Annals of Mathematical Statistics*, 23(2):255 – 262.
- [Fauw et al., 2018] Fauw, J., Ledsam, J., Romera-Paredes, B., Nikolov, S., Tomasev, N., Blackwell, S., Askham, H., Glorot, X., O’Donoghue, B., Visentin, D., Driessche, G., Lakshminarayanan, B., Meyer, C., Mackinder, F., Bouton, S., Ayoub, K., Chopra, R., King, D., Karthikesalingam, A., and Ronneberger, O. (2018). Clinically applicable deep learning for diagnosis and referral in retinal disease. *Nature Medicine*, 24.
- [Fellman, 1974] Fellman, J. (1974). *On the Allocation of Linear Observations*. Commentationes physico-mathematicae // Societas Scientiarum Fennica. Societas Scientiarum Fennica.
- [Fuster et al., 2022] Fuster, A., Goldsmith-Pinkham, P., Ramadorai, T., and Walther, A. (2022). Predictably unequal? the effects of machine learning on credit markets. *The Journal of Finance*, 77(1):5–47.
- [Hadiji et al., 2020] Hadiji, H., Gerchinovitz, S., Loubes, J.-M., and Stoltz, G. (2020). Diversity-Preserving K-Armed Bandits, Revisited. working paper or preprint.
- [Harman and Jurík, 2008] Harman, R. and Jurík, T. (2008). Computing c-optimal experimental designs using the simplex method of linear programming. *Computational Statistics & Data Analysis*, 53(2):247–254.
- [Khademi et al., 2019] Khademi, A., Lee, S., Foley, D., and Honavar, V. (2019). Fairness in algorithmic decision making: An excursion through the lens of causality. In *The World Wide Web Conference, WWW ’19*, page 2907–2914, New York, NY, USA. Association for Computing Machinery.
- [Kiefer, 1974] Kiefer, J. (1974). General Equivalence Theory for Optimum Designs (Approximate Theory). *The Annals of Statistics*, 2(5):849 – 879.
- [Kilbertus et al., 2017] Kilbertus, N., Rojas-Carulla, M., Parascandolo, G., Hardt, M., Janzing, D., and Schölkopf, B. (2017). Avoiding discrimination through causal reasoning. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS’17*, page 656–666, Red Hook, NY, USA. Curran Associates Inc.
- [Kirschner et al., 2020] Kirschner, J., Lattimore, T., and Krause, A. (2020). Information directed sampling for linear partial monitoring. In Abernethy, J. D. and Agarwal, S., editors, *Conference on Learning Theory*, volume 125 of *Proceedings of Machine Learning Research*, pages 2328–2369. PMLR.
- [Köchling and Wehner, 2020] Köchling, A. and Wehner, M. C. (2020). Discriminated by an algorithm: a systematic review of discrimination and fairness by algorithmic decision-making in the context of hr recruitment and hr development. *Business Research*, pages 1–54.
- [Lattimore and Szepesvári, 2020] Lattimore, T. and Szepesvári, C. (2020). *Bandit Algorithms*. Cambridge University Press.
- [Lin et al., 2014] Lin, T., Abrahao, B., Kleinberg, R., Lui, J., and Chen, W. (2014). Combinatorial partial monitoring game with linear feedback and its applications. In Xing, E. P. and Jebara, T., editors, *Proceedings of the 31st International Conference on Machine Learning*, number 2, pages 901–909, Beijing, China. PMLR.
- [Mehrabi et al., 2021] Mehrabi, N., Morstatter, F., Saxena, N., Lerman, K., and Galstyan, A. (2021). A survey on bias and fairness in machine learning. *ACM Computing Surveys*, 54(6).
- [Oneto and Chiappa, 2020] Oneto, L. and Chiappa, S. (2020). Fairness in machine learning. *Studies in Computational Intelligence*, page 155–196.

- [Papamitsiou and Economides, 2014] Papamitsiou, Z. and Economides, A. A. (2014). Learning analytics and educational data mining in practice: A systematic literature review of empirical evidence. *Journal of Educational Technology & Society*, 17(4):49–64.
- [Patil et al., 2020] Patil, V., Ghalme, G., Nair, V., and Narahari, Y. (2020). Achieving fairness in the stochastic multi-armed bandit problem. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(04):5379–5386.
- [Pázman, 1986] Pázman, A. (1986). *Foundations of Optimum Experimental Design*. Mathematics and its Applications. Springer Netherlands.
- [Perlich et al., 2014] Perlich, C., Dalessandro, B., Raeder, T., Stitelman, O., and Provost, F. (2014). Machine learning for targeted display advertising: transfer learning in action. *Machine Learning*, 95(1):103–127.
- [Pronzato and Sagnol, 2021] Pronzato, L. and Sagnol, G. (2021). Removing inessential points in c- and A-optimal design. *Journal of Statistical Planning and Inference*, 213:233–252.
- [Pukelsheim, 1980] Pukelsheim, F. (1980). On linear regression designs which maximize information. *Journal of statistical planning and inference*, 4:339–364.
- [Raghavan et al., 2020] Raghavan, M., Barocas, S., Kleinberg, J. M., and Levy, K. E. C. (2020). Mitigating bias in algorithmic hiring: evaluating claims and practices. *Proceedings of the 2020 Conference on Fairness, Accountability, and Transparency*.
- [Sagnol, 2010] Sagnol, G. (2010). Plans d’expériences optimaux et application à l’estimation des matrices de trafic dans les grands réseaux : programmation conique du second ordre et sous-modularité.
- [Sagnol and Harman, 2013] Sagnol, G. and Harman, R. (2013). Computing exact d-optimal designs by mixed integer second order cone programming. *The Annals of Statistics*, 43.
- [Vandenberghe et al., 1998] Vandenberghe, L., Boyd, S., and Wu, S.-P. (1998). Determinant maximization with linear matrix inequality constraints. *SIAM Journal on Matrix Analysis and Applications*, 19(2):499–533.
- [Černý and Hladík, 2012] Černý, M. and Hladík, M. (2012). Two complexity results on c-optimality in experimental design. *Computational Optimization and Applications*, 51(3):1397–1408.
- [Wang et al., 2021] Wang, L., Bai, Y., Sun, W., and Joachims, T. (2021). Fairness of exposure in stochastic bandits. In Meila, M. and Zhang, T., editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 10686–10696. PMLR.

A Detailed Fair Phased Elimination algorithm

We present the notations used in Algorithm 4. The phases are indexed by $l \in \mathbb{N}^*$. The sets $\mathcal{X}_l^{(z)}$ for $z \in \{-1, +1\}$ corresponds to actions in group z that are considered as potentially optimal in phase l . The variable \widehat{z}_l^* encodes the group determined as optimal : it is 0 as long as this group has not been determined. The subscript (z) refer to the group z when $z \in \{-1, +1\}$, and otherwise to the estimation of the bias ω^* : for example, the probability $\pi_l^{(z)}$ for $z \in \{-1, +1\}$ and $l > 1$ corresponds to the approximate G -optimal design on $\mathcal{X}_l^{(z)}$. Then, for $z \in \{-1, +1\}$, allocations $\mu^{(z)}$ (resp. $\mu^{(0)}$) correspond to allocation of samples in the exploration phase $\text{Exp}_l^{(z)}$ (resp. $\text{Exp}_l^{(0)}$). Similarly, $V_l^{(z)}$ (resp $V_l^{(0)}$) denotes the variance matrix of the estimator $\begin{pmatrix} \widehat{\gamma}_l^{(z)} \\ \widehat{\omega}_l^{(z)} \end{pmatrix}$ (resp. $\widehat{\omega}_l^{(0)}$) obtained from observations made during phase $\text{Exp}_l^{(z)}$ (resp. $\text{Exp}_l^{(0)}$). Finally, $\text{Explore}_l^{(z)}$ (resp. $\text{Explore}_l^{(0)}$) is a Boolean variable indicating whether the exploration at phase l for group z (resp. for the bias parameter) has been performed. It is used in the proofs to ensure that the corresponding estimators are well defined.

Algorithm 4 Fair Phased Elimination (detailed version)

```

1: Input:  $\delta, T, k = |\mathcal{X}|$ 
2: Initialize:  $\Delta_{\max} \leftarrow \max_{\gamma \in \mathcal{C}(\mathcal{X}), (x, x') \in \mathcal{X}^2} (x - x')^\top \gamma$ ,  $\text{Recovery} \leftarrow \emptyset$ ,  $t \leftarrow 0$ ,  $l \leftarrow 1$ 
3:  $\hat{z}_1^* \leftarrow 0$ ,  $\mathcal{X}_1^{(+1)} \leftarrow \{x : z_x = 1\}$ ,  $\mathcal{X}_1^{(-1)} \leftarrow \{x : z_x = -1\}$ ,  $\hat{\Delta}_x^1 \leftarrow \Delta_{\max}$  for  $x \in \mathcal{X}$ 
4: while  $t < T$  do
5:   Initialize:  $\epsilon_l \leftarrow 2^{1-l} \Delta_{\max}$ ,  $\hat{z}_{l+1}^* \leftarrow \hat{z}_l^*$ ,  $\hat{\Delta}^{l+1} \leftarrow \hat{\Delta}^l$ ,  $\text{Explore}_l^{(z)} \leftarrow \text{False}$  for  $z \in \{-1, 0, +1\}$ 
6:   for  $z \in \{-1, +1\}$  such that  $z \neq -\hat{z}_l^*$  do ▷ G-optimal Exploration and Elimination
7:      $\pi_l^{(z)} \leftarrow \underset{\pi}{\operatorname{argmin}} \left\{ \max_{x \in \mathcal{X}_l^{(z)}} a_x^\top V(\pi) + a_x : \pi \in \mathcal{P}_{\mathcal{X}_l^{(z)}}, |\operatorname{supp}(\pi)| \leq \frac{(d+1)(d+2)}{2} \right\}$ 
8:      $\mu_l^{(z)}(x) \leftarrow \left\lceil \frac{2(d+1)\pi_l^{(z)}(x)}{\epsilon_l^2} \log \left( \frac{kl(l+1)}{\delta} \right) \right\rceil$  for all  $x \in \mathcal{X}_l^{(z)}$ 
9:      $n_l^{(z)} \leftarrow \sum_{x \in \mathcal{X}_l^{(z)}} \mu_l^{(z)}(x)$ ,  $\text{Exp}_l^{(z)} \leftarrow \{t+1, \dots, T \wedge (t+n_l^{(z)})\}$ 
10:    if  $t+n_l^{(z)} \leq T$  then
11:       $\text{Explore}_l^{(z)} \leftarrow \text{True}$ , choose each action  $x \in \mathcal{X}_l^{(z)}$  exactly  $\mu_l^{(z)}(x)$  times
12:       $V_l^{(z)} \leftarrow \sum_{t \in \text{Exp}_l^{(z)}} a_{x_t} a_{x_t}^\top$ ,  $\hat{\theta}_l^{(z)} \leftarrow \left( V_l^{(z)} \right)^+ \sum_{t \in \text{Exp}_l^{(z)}} y_t a_{x_t}$ 
13:       $\mathcal{X}_{l+1}^{(z)} \leftarrow \left\{ x \in \mathcal{X}_l^{(z)} : \max_{x' \in \mathcal{X}_l^{(z)}} (a_{x'} - a_x)^\top \hat{\theta}_l^{(z)} \leq 3\epsilon_l \right\}$ 
14:    else for  $t \in \text{Exp}_l^{(z)}$ , sample empirical best action in  $\mathcal{X}_l^{(z)}$ 
15:    end if
16:     $t \leftarrow t + n_l^{(z)}$ 
17:  end for
18:  if  $\hat{z}_l^* = 0$  then
19:    compute the  $\hat{\Delta}^l$ -optimal design  $\hat{\mu}_l$  and the corresponding regret  $\kappa(\hat{\Delta}^l)$ 
20:    if  $\epsilon_l \leq \left( \kappa(\hat{\Delta}^l) \log(T)/T \right)^{1/3}$  then ▷ Recovery phase
21:       $\text{Recovery} \leftarrow \{t, \dots, T\}$ 
22:      sample empirical best action in  $\mathcal{X}_{l+1}^{(-1)} \cup \mathcal{X}_{l+1}^{(1)}$  until the end of the budget,  $t \leftarrow T$ 
23:    else ▷  $\hat{\Delta}^l$ -optimal Exploration and Elimination
24:       $\mu_l^{(0)}(x) \leftarrow \left\lceil \frac{2\hat{\mu}_l(x)}{\epsilon_l^2} \log \left( \frac{l(l+1)}{\delta} \right) \right\rceil$  for all  $x \in \mathcal{X}$ 
25:       $n_l^{(0)} \leftarrow \sum_{x \in \mathcal{X}} \mu_l^{(0)}(x)$ ,  $\text{Exp}_l^{(0)} \leftarrow \{t, \dots, T \wedge (t+n_l^{(0)})\}$ 
26:      if  $t+n_l^{(0)} \leq T$  then
27:         $\text{Explore}_l^{(0)} \leftarrow \text{True}$ , choose each action  $x \in \mathcal{X}$  exactly  $\mu_l^{(0)}(x)$  times
28:         $V_l^{(0)} \leftarrow \sum_{t \in \text{Exp}_l^{(0)}} a_{x_t} a_{x_t}^\top$ ,  $\hat{\omega}_l^{(0)} \leftarrow e_{d+1}^\top \left( V_l^{(0)} \right)^+ \sum_{t \in \text{Exp}_l^{(0)}} y_t a_{x_t}$ 
29:        for  $x \in \mathcal{X}_{l+1}^{(-1)} \cup \mathcal{X}_{l+1}^{(1)}$  do
30:           $\hat{m}_{l,x} \leftarrow a_x^\top \hat{\theta}_l^{(z_x)} - z_x \hat{\omega}_l^{(0)}$ ,  $\hat{\Delta}_x^{l+1} \leftarrow \left( \max_{x' \in \mathcal{X}_{l+1}^{(-1)} \cup \mathcal{X}_{l+1}^{(1)}} \hat{m}_{l,x'} - \hat{m}_{l,x} + 4\epsilon_l \right) \wedge \Delta_{\max}$ 
31:        end for
32:        for  $z \in \{-1, +1\}$  do
33:          if  $\max_{x \in \mathcal{X}_{l+1}^{(z)}} \hat{m}_{l,x} - 2\epsilon_l \geq \max_{x \in \mathcal{X}_{l+1}^{(-z)}} \hat{m}_{l,x} + 2\epsilon_l$  then  $\hat{z}_{l+1}^* \leftarrow z$ 
34:          end if
35:        end for
36:        else sample empirical best action in  $\mathcal{X}_{l+1}^{(-1)} \cup \mathcal{X}_{l+1}^{(1)}$  until the end of the budget,  $t \leftarrow T$ 
37:        end if
38:         $t \leftarrow t + n_l^{(0)}$ 
39:      end if
40:    end if
41:     $l \leftarrow l + 1$ 
42: end while

```

B Lower bound on the worst-case regret for $d \geq 2$

Theorem 5. Let $v \geq 1$, $d \geq 2$ and $T \geq 4^3 v$. There exists an action set $\mathcal{A} \in \mathbf{A}_{v,d}$ such that for any algorithm, there exists a bandit problem with parameter $\theta_T \in \mathbb{R}^{d+1}$ such that the regret of this algorithm on the problem characterized by θ_T satisfies

$$R_T^{\mathcal{A}, \theta_T} \geq \frac{1}{8e} v^{1/3} T^{2/3}.$$

Lemma 8. Let the action set be given by $\mathcal{A} = \left\{ \begin{pmatrix} x_1 \\ z_{x_1} \end{pmatrix}, \dots, \begin{pmatrix} x_{d+1} \\ z_{x_{d+1}} \end{pmatrix} \right\}$, where $\begin{pmatrix} x_1 \\ z_{x_1} \end{pmatrix} = e_1 + e_{d+1}$, $\begin{pmatrix} x_i \\ z_{x_i} \end{pmatrix} = e_i - e_{d+1}$ for $i \in \{2, \dots, d\}$, and $\begin{pmatrix} x_{d+1} \\ z_{x_{d+1}} \end{pmatrix} = -\left(1 - \frac{2}{\sqrt{v}+1}\right) e_1 - e_{d+1}$. It holds that

$$\min_{\pi \in \mathcal{P}^{\mathcal{A}}} \left\{ e_{d+1}^\top \left(\sum_{\begin{pmatrix} x \\ z \end{pmatrix} \in \mathcal{A}} \pi_x \begin{pmatrix} x \\ z \end{pmatrix} \begin{pmatrix} x \\ z \end{pmatrix}^\top \right)^+ e_{d+1} \right\} = \kappa.$$

Proof. Consider the actions \mathcal{A} defined in the sketch of proof of Lemma 8. By Lemma 8, $\mathcal{A} \in \mathbf{A}_{v,d}$. We will introduce two bandit problems characterized by two parameters $\theta_T^{(1)}$ and $\theta_T^{(2)}$ - assuming that the noise ξ_t is Gaussian and i.i.d. - and we prove that for any algorithm, the regret for one of those two problems must be of larger order than $v^{1/3} T^{2/3}$.

We also consider the two alternative problems defined in the sketch of proof of Lemma 8 as follows. For a small $1/4 > \rho_T > 0$ where $\rho_T = T^{-1/3} v^{1/3}$ (satisfied since $T > 4^3 v$), the two alternative action parameters are defined as:

$$\begin{aligned} \gamma_T^{(1)} &= \frac{1 + \rho_T}{2} e_1 + \frac{1 - \rho_T}{2} e_2 - \frac{\rho_T}{2} \left(\sum_{3 \leq j \leq d} e_j \right) \\ \gamma_T^{(2)} &= \frac{1 - \rho_T}{2} e_1 + \frac{1 + \rho_T}{2} e_2 + \frac{\rho_T}{2} \left(\sum_{3 \leq j \leq d} e_j \right). \end{aligned}$$

On top of this, two bias parameters are defined as $\omega_T^{(1)} = -\frac{\rho_T}{2}$ and $\omega_T^{(2)} = \frac{\rho_T}{2}$. Through this, we define the two bandit problems of the sketch of proof of Lemma 8 characterized by $\theta_T^{(1)} = \begin{pmatrix} \gamma_T^{(1)} \\ \omega_T^{(1)} \end{pmatrix}$ and $\theta_T^{(2)} = \begin{pmatrix} \gamma_T^{(2)} \\ \omega_T^{(2)} \end{pmatrix}$ - and where the distribution of the noise ξ_t is supposed to be Gaussian and i.i.d. We refer to these two problems respectively as **Problem 1** and **Problem 2**. We write $R_T^{(1)}$, $\mathbb{P}^{(1)}$ and $\mathbb{E}^{(1)}$ (respectively $R_T^{(2)}$, $\mathbb{P}^{(2)}$ and $\mathbb{E}^{(2)}$) for the regret, probability and expectation for the first bandit problem, when the parameter is $\theta_T^{(1)}$ (respectively the second bandit problem with $\theta_T^{(2)}$). We also write $\mathbb{P}_j^{(i)}$ for the distribution of a sample received in **Problem i** when sampling action x_j at any given time t - note that by definition of the bandit problems, this distribution does not depend on t and on the past samples given that action x_j is sampled.

The three following facts hold on these two bandit problems:

Fact 1 The parameters $\gamma_T^{(1)}$ and $\gamma_T^{(2)}$ are chosen so that x_1 is the unique best action for **Problem 1**, and x_2 is the unique best action for **Problem 2**. Choosing any sub-optimal action induces an instantaneous regret of at least ρ_T , and choosing the very sub-optimal action x_{d+1} induces an instantaneous regret of at least $1/2$.

Fact 2 Because of the chosen bias parameters, the distributions of the evaluations of all actions but x_{d+1} are exactly the same under the two bandit problems characterized by $\theta_T^{(1)}$ and $\theta_T^{(2)}$ - i.e. exactly the same data is observed under the two alternative bandit problems defined by the two alternative parameters for all actions but x_{d+1} . More precisely, for $i \in \{1, 2\}$, in **Problem i** and at any time t , when sampling action x_i where $i \leq 2$, we observe a sample distributed according to $\mathcal{N}(1/2, 1)$ - i.e. $\mathbb{P}_j^{(i)}$ is $\mathcal{N}(1/2, 1)$ - and when sampling action x_i where $2 < i \leq d+1$, we observe a sample distributed according to $\mathcal{N}(0, 1)$ - i.e. $\mathbb{P}_j^{(i)}$ is $\mathcal{N}(0, 1)$.

Fact 3 The distributions of the outcomes of the evaluation of action x_{d+1} differs in the two bandit problems. Set $\alpha = 2/(\sqrt{v} + 1)$. In **Problem 1**, $\mathbb{P}_{d+1}^{(1)}$ is $\mathcal{N}(-\frac{1-\alpha-\rho_T\alpha}{2}, 1)$. In **Problem 2**, $\mathbb{P}_{d+1}^{(2)}$ is $\mathcal{N}(-\frac{1-\alpha+\rho_T\alpha}{2}, 1)$. So that the difference between the means of the evaluations of action x_{d+1} in the two bandit problems is $\bar{\Delta} = \rho_T\alpha = \frac{2\rho_T}{\sqrt{v}+1} \leq \frac{2\rho_T}{\sqrt{v}}$.

For $i \leq d+1$, we write $N_i(T)$ for the number of times that action x_i has been selected before time T . In **Problem 1**, choosing the action x_{d+1} leads to an instantaneous regret larger than $\frac{1}{2}$ (**Fact 1**), so that

$$R_T^{(1)} \geq \frac{\mathbb{E}^{(1)}[N_{x_{d+1}}(T)]}{2}.$$

If $\mathbb{E}^{(1)}[N_{d+1}(T)] \geq \frac{T^{2/3}v^{1/3}}{2}$, then Theorem 2 follows immediately; we therefore consider from now on the case when

$$\mathbb{E}^{(1)}[N_{d+1}(T)] \leq \frac{T^{2/3}v^{1/3}}{2}. \quad (11)$$

Now, let us define the event

$$F = \left\{ N_1(T) \geq \frac{T}{2}v^{1/3} \right\}.$$

Note that action x_1 is optimal for **Problem 1** and that action x_2 is optimal for **Problem 2** (**Fact 1**). Since choosing an action that is sub-optimal leads to an instantaneous regret larger than ρ_T (**Fact 1**), we also have

$$R_T^{(1)} \geq \frac{T\rho_T}{2} \mathbb{P}^{(1)}(\bar{F})$$

and

$$R_T^{(2)} \geq \frac{T\rho_T}{2} \mathbb{P}^{(2)}(F).$$

Then, Bretagnolle-Huber inequality (see, e.g., Theorem 14.2 in [Lattimore and Szepesvári, 2020]) implies that

$$R_T^{(1)} + R_T^{(2)} \geq \frac{T\rho_T}{4} \exp\left(-KL\left(\mathbb{P}^{(1)}, \mathbb{P}^{(2)}\right)\right).$$

For the choice $\rho_T = T^{-1/3}v^{1/3}$, this implies that

$$R_T^{(1)} + R_T^{(2)} \geq \frac{T^{2/3}v^{1/3}}{4} \exp\left(-KL\left(\mathbb{P}^{(1)}, \mathbb{P}^{(2)}\right)\right). \quad (12)$$

Now, the Kullback-Leibler divergence between $\mathbb{P}^{(1)}$ and $\mathbb{P}^{(2)}$ can be rewritten as follows (see, e.g., Lemma 15.1 in [Lattimore and Szepesvári, 2020]) :

$$KL(\mathbb{P}^{(1)}, \mathbb{P}^{(2)}) = \frac{1}{2} \sum_{j \leq d+1} \mathbb{E}^{(1)}[N_j(T)] KL(\mathbb{P}_j^{(1)}, \mathbb{P}_j^{(2)}).$$

By **Fact 2**, we have that for any $j \leq d$, $\mathbb{P}_j^{(1)} = \mathbb{P}_j^{(2)}$. So that

$$KL(\mathbb{P}^{(1)}, \mathbb{P}^{(2)}) = \frac{1}{2} \mathbb{E}^{(1)}[N_{d+1}(T)] KL(\mathbb{P}_{d+1}^{(1)}, \mathbb{P}_{d+1}^{(2)}).$$

By the characterization of $\mathbb{P}_{d+1}^{(1)}, \mathbb{P}_{d+1}^{(2)}$ in **Fact 3**, and recalling that the Kullback-Leibler divergence between two normalized Gaussian distributions is given by the squared distance between their means, we find that

$$KL(\mathbb{P}^{(1)}, \mathbb{P}^{(2)}) = \frac{1}{2} \mathbb{E}^{(1)}[N_{d+1}(T)] \bar{\Delta}^2.$$

Thus, by the definition of $\bar{\Delta}$ in **Fact 3** and by Equation (11)

$$KL(\mathbb{P}^{(1)}, \mathbb{P}^{(2)}) = \frac{1}{2} \mathbb{E}^{(1)}[N_{d+1}(T)] \left(\frac{2\rho_T}{\sqrt{v}+1} \right)^2 \leq \frac{T^{2/3}v^{1/3}}{4} \times \frac{4\rho_T^2}{v} = 1, \quad (13)$$

reminding that $\rho_T = T^{-1/3}v^{1/3}$.

Combining Equations (12) and (13) implies that

$$\max \left\{ R_T^{(1)}, R_T^{(2)} \right\} \geq \frac{T^{2/3}\kappa^{1/3}}{8} \exp(-1),$$

which concludes the proof of Theorem 5. \square

C Proofs

C.1 Proof of Theorem 2

We begin by defining for $z \in \{-1, 0, +1\}$

$$L^{(z)} = \max \left\{ l \geq 1 : \text{Explore}_l^{(z)} = \text{True} \right\}$$

the largest integer l such that $\text{Explore}_l^{(z)} = \text{True}$. Recall that $\kappa_* = \Delta_{\max} v$, where v is the e_{d+1} -optimal variance. By definition of the algorithm, for all $l \leq L^{(0)} + 1$, $\hat{\Delta}^l \leq \Delta_{\max}$, so $\kappa(\hat{\Delta}^l) \leq \kappa_*$. Now, let us also define

$$L_T = \max \left\{ l \geq 1 : \epsilon_l > \left(\frac{\kappa_* \log(T)}{T} \right)^{1/3} \right\}.$$

Then, if $\text{Recovery} \neq \emptyset$, we must have $L^{(0)} \geq L_T$. Moreover, we see that since $\epsilon_{L_T} = 2^{1-L_T} \Delta_{\max}$, we have $L_T \leq 1 + 3 \log_2 \left(T \Delta_{\max}^{1/3} / (\kappa_* \log(T)) \right) \leq 4 \log_2(T)$ when $T > 1$.

We define a "bad" event \mathcal{F} , such that, on $\overline{\mathcal{F}}$, our estimators $\hat{\gamma}_l^{(z)}$ and $\hat{\omega}_l^{(z)}$ are close to the true parameters γ^* and ω^* for all rounds l . More precisely, let

$$\mathcal{F} = \bigcup_{l \geq 1} \mathcal{F}_l, \tag{14}$$

where for $l \geq 1$

$$\begin{aligned} \mathcal{F}_l = & \left\{ \exists z \in \{-1, 1\} \text{ such that } \text{Explore}_l^{(z)} = \text{True}, \text{ and } x \in \mathcal{X}_l^{(z)} \text{ such that } \left| \begin{pmatrix} \hat{\gamma}_l^{(z)} - \gamma^* \\ \hat{\omega}_l^{(z)} - \omega^* \end{pmatrix}^\top \begin{pmatrix} x \\ z_x \end{pmatrix} \right| \geq \epsilon_l \right\} \\ & \bigcup \left\{ \text{Explore}_l^{(0)} = \text{True} \text{ and } \left| \hat{\omega}_l^{(0)} - \omega^* \right| \geq \epsilon_l \right\}. \end{aligned}$$

Then, the regret decomposes as

$$R_T \leq \sum_{t \leq T} \mathbb{E}_{|\overline{\mathcal{F}}} [(x^* - x_t)^\top \gamma^*] + 2T\mathbb{P}[\mathcal{F}]. \tag{15}$$

The following lemma relies on concentration of Gaussian variables to bound the probability of the event \mathcal{F} .

Lemma 9. $\mathbb{P}(\mathcal{F}) \leq 2\delta$.

Now, the first term of (15) can be decomposed as

$$\sum_{t \leq T} (x^* - x_t)^\top \gamma^* \leq \sum_{z \in \{-1, 0, +1\}} \sum_{l=1}^{L^{(z)}+1} \sum_{t \in \text{Exp}_l^{(z)}} (x^* - x_t)^\top \gamma^* + \sum_{t \in \text{Recovery}} (x^* - x_t)^\top \gamma^*,$$

where we use as convention that the sum over an empty set is null. Note that for $z \in \{-1, +1\}$, during the phase $\text{Exp}_l^{(z)}$ the algorithm only samples actions from $\mathcal{X}_l^{(z)}$. By contrast, during the phase $\text{Exp}_l^{(0)}$, even

actions eliminated from the sets $\mathcal{X}_l^{(z)}$ can be sampled. Finally, if the algorithm stops during phase $\text{Exp}_{L^{(0)}+1}^{(0)}$, but does not have enough budget to complete the last $\hat{\Delta}^l$ -optimal Exploration and Elimination Phase, it samples the remaining actions in the set $\mathcal{X}_{L^{(0)}+2}^{(-1)} \cup \mathcal{X}_{L^{(0)}+2}^{(+1)}$. Hence, the first term of (15) can be upper-bounded by

$$\begin{aligned} \sum_{t \leq T} (x^* - x_t)^\top \gamma^* &\leq \sum_{z \in \{-1, +1\}} \sum_{l=1}^{L_T} \left(\sum_{x \in \mathcal{X}_l^{(z)}} \mu_l^{(z)}(x) \right) \max_{x \in \mathcal{X}_l^{(z)}} (x^* - x)^\top \gamma^* \\ &+ \sum_{z \in \{-1, +1\}} \sum_{l=L_T+1}^{L^{(z)}+1} \sum_{t \in \text{Exp}_l^{(z)}} (x^* - x_t)^\top \gamma^* + \sum_{t \in \text{Recovery}} (x^* - x_t)^\top \gamma^* \\ &+ \sum_{l=1}^{L^{(0)}} \sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \Delta_x + \mathbb{1} \left\{ \text{Explore}_{L^{(0)}+1}^{(0)} = \text{False} \right\} \sum_{t \in \text{Exp}_{L^{(0)}+1}^{(0)}} \max_{x \in \mathcal{X}_{L^{(0)}+2}^{(-1)} \cup \mathcal{X}_{L^{(0)}+2}^{(+1)}} (x^* - x)^\top \gamma^*. \end{aligned} \quad (16)$$

We begin by bounding the sum of the regret corresponding to the Recovery phase and to the phases $\text{Exp}_L^{(z)}$ for $z \in \{-1, +1\}$ and $l > L_T$ on the event $\overline{\mathcal{F}}$.

Bound on
$$\sum_{z \in \{-1, +1\}} \sum_{l=L_T+1}^{L^{(z)}+1} \sum_{t \in \text{Exp}_l^{(z)}} (x^* - x_t)^\top \gamma^* + \sum_{t \in \text{Recovery}} (x^* - x_t)^\top \gamma^*.$$

Lemma 10. *Let $x^* \in \arg\max_{x \in \mathcal{X}} x^\top \gamma^*$ be an optimal action. Then, on the event $\overline{\mathcal{F}}$ defined in Equation (14), for $l \geq 1$ such that $\text{Explore}_l^{(z_{x^*})} = \text{True}$,*

$$\mathcal{X}_{l+1}^{(z_{x^*})} \subset \left\{ x \in \mathcal{X}_1^{(z_{x^*})} : (x^* - x)^\top \gamma^* < 10\epsilon_{l+1} \right\}. \quad (17)$$

Moreover, for $l \geq 1$ such that $\text{Explore}_l^{(-z_{x^*})} = \text{True}$,

$$\mathcal{X}_{l+1}^{(-z_{x^*})} \subset \left\{ x \in \mathcal{X}_1^{(-z_{x^*})} : (x^* - x)^\top \gamma^* < 42\epsilon_{l+1} \right\}.$$

Recall that if $\text{Recovery} \neq \emptyset$, $L^{(0)} \geq L_T$. Then, all actions sampled during the Recovery phase belong to $\mathcal{X}_{l+1}^{(-1)} \cup \mathcal{X}_{l+1}^{(+1)}$ for some $l \geq L_T$. Lemma 10 shows that, on $\overline{\mathcal{F}}$, for $l \geq L_T$, the actions in $\mathcal{X}_{l+1}^{(z)}$ are sub-optimal by at most $42\epsilon_{L_T+1}$. Then, we get that on the event $\overline{\mathcal{F}}$,

$$\begin{aligned} \sum_{z \in \{-1, +1\}} \sum_{l=L_T+1}^{L^{(z)}+1} \sum_{t \in \text{Exp}_l^{(z)}} (x^* - x_t)^\top \gamma^* + \sum_{t \in \text{Recovery}} (x^* - x_t)^\top \gamma^* &\leq T \times 42\epsilon_{L_T+1} \\ &\leq 42\kappa_*^{1/3} T^{2/3} \log(T)^{1/3}. \end{aligned} \quad (18)$$

Bound on
$$\sum_{l=1}^{L^{(0)}} \sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \Delta_x + \mathbb{1} \left\{ \text{Explore}_{L^{(0)}+1}^{(0)} = \text{False} \right\} \sum_{t \in \text{Exp}_{L^{(0)}+1}^{(0)}} \max_{x \in \mathcal{X}_{L^{(0)}+2}^{(-1)} \cup \mathcal{X}_{L^{(0)}+2}^{(+1)}} (x^* - x)^\top \gamma^*.$$

We begin by bounding the term $\mathbb{1} \left\{ \text{Explore}_{L^{(0)}+1}^{(0)} = \text{False} \right\} \sum_{t \in \text{Exp}_{L^{(0)}+1}^{(0)}} \max_{x \in \mathcal{X}_{L^{(0)}+2}^{(-1)} \cup \mathcal{X}_{L^{(0)}+2}^{(+1)}} (x^* - x)^\top \gamma^*$. Recall

that $n_{L^{(0)}+1}^{(0)} = \sum_{x \in \mathcal{X}} \mu_{L^{(0)}+1}^{(0)}(x)$ is the budget that would be necessary to complete the $\widehat{\Delta}^l$ -optimal Exploration and Elimination phase at phase $L^{(0)} + 1$. On the one hand, Lemma 10 implies that on the event $\overline{\mathcal{F}}$,

$$\mathbb{1} \left\{ \text{Explore}_{L^{(0)}+1}^{(0)} = \text{False} \right\} \sum_{t \in \text{Exp}_{L^{(0)}+1}^{(0)}} \max_{x \in \mathcal{X}_{L^{(0)}+2}^{(-1)} \cup \mathcal{X}_{L^{(0)}+2}^{(+1)}} (x^* - x)^\top \gamma^* \leq 42n_{L^{(0)}+1}^{(0)} \epsilon_{L^{(0)}+2} \leq 21n_{L^{(0)}+1}^{(0)} \epsilon_{L^{(0)}+1}.$$

On the other hand, for all $l \leq L^{(0)} + 1$, the definition of $\widehat{\Delta}^l$ implies that $\widehat{\Delta}_x^l \geq \epsilon_l$ for all $x \in \mathcal{X}$. Therefore, $21n_{L^{(0)}+1}^{(0)} \epsilon_{L^{(0)}+1} \leq 21n_{L^{(0)}+1}^{(0)} \min_x \widehat{\Delta}_x^{L^{(0)}+1}$. This implies that on $\overline{\mathcal{F}}$,

$$\mathbb{1} \left\{ \text{Explore}_{L^{(0)}+1}^{(0)} = \text{False} \right\} \sum_{t \in \text{Exp}_{L^{(0)}+1}^{(0)}} \max_{x \in \mathcal{X}_{L^{(0)}+2}^{(-1)} \cup \mathcal{X}_{L^{(0)}+2}^{(+1)}} (x^* - x)^\top \gamma^* \leq 21 \sum_{x \in \mathcal{X}} \mu_{L^{(0)}+1}^{(0)}(x) \widehat{\Delta}_x^{L^{(0)}+1}. \quad (19)$$

Next, to bound the remaining terms of Equation (16), we bound the regret $\sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \Delta_x$ of exploration phase $\text{Exp}_l^{(0)}$ using the following lemma.

Lemma 11. *For all $l > 0$, and $z \in \{-1, +1\}$, we have*

$$\sum_{x \in \mathcal{X}_l^{(z)}} \mu_l^{(z)}(x) \leq \frac{2(d+1)}{\epsilon_l^2} \log \left(\frac{kl(l+1)}{\delta} \right) + \frac{(d+1)(d+2)}{2}.$$

and on $\overline{\mathcal{F}}$, we have

$$\sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \Delta_x \leq \sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \widehat{\Delta}_x^l \leq \frac{2\kappa(\widehat{\Delta}^l)}{\epsilon_l^2} \log \left(\frac{l(l+1)}{\delta} \right) + 2(d+1).$$

Then, Equation (19) and Lemma 11 imply that on $\overline{\mathcal{F}}$

$$\begin{aligned} \sum_{l=1}^{L^{(0)}} \sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \Delta_x + \mathbb{1} \left\{ \text{Explore}_{L^{(0)}+1}^{(0)} = \text{False} \right\} \sum_{t \in \text{Exp}_{L^{(0)}+1}^{(0)}} \max_{x \in \mathcal{X}_{L^{(0)}+2}^{(-1)} \cup \mathcal{X}_{L^{(0)}+2}^{(+1)}} (x^* - x)^\top \gamma^* \\ \leq 21 \sum_{l=1}^{L^{(0)}+1} \sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \widehat{\Delta}_x^l \\ \leq 42 \sum_{l=1}^{L^{(0)}+1} \frac{\kappa(\widehat{\Delta}^l)}{\epsilon_l^2} \log \left(\frac{l(l+1)}{\delta} \right) + 42(d+1)(L^{(0)}+1) \end{aligned} \quad (20)$$

We rely on the following Lemma to bound $\kappa(\widehat{\Delta}^l)$.

Lemma 12. *On $\overline{\mathcal{F}}$, we have for any $l \geq 1$ and any $\tau > 0$*

$$\kappa(\widehat{\Delta}^l) \leq 513 \left(1 + \frac{\epsilon_l}{\tau} \right) \kappa(\Delta \vee \tau).$$

and

$$\kappa(\widehat{\Delta}^l) \geq \kappa(\Delta \vee \epsilon_l).$$

Lemma 11 and Lemma 12 with $\tau = \epsilon_{L^{(0)}}$ imply that on $\overline{\mathcal{F}}$,

$$\begin{aligned} \sum_{l=1}^{L^{(0)}+1} \frac{\kappa(\widehat{\Delta}^l)}{\epsilon_l^2} \log \left(\frac{l(l+1)}{\delta} \right) &\leq 513 \kappa(\Delta \vee \epsilon_{L^{(0)}}) \log \left(\frac{(L^{(0)}+1)(L^{(0)}+2)}{\delta} \right) \left(\sum_{l=1}^{L^{(0)}+1} \frac{1}{\epsilon_l^2} + \sum_{l=1}^{L^{(0)}+1} \frac{1}{\epsilon_l \epsilon_{L^{(0)}}} \right) \\ &\leq 1026 \kappa(\Delta \vee \epsilon_{L^{(0)}}) \log \left(\frac{6L^{(0)}}{\delta} \right) \left(\frac{16}{\epsilon_{L^{(0)}}^2} + \frac{4}{\epsilon_{L^{(0)}}^2} \right) \\ &\leq 20520 \log \left(\frac{6L^{(0)}}{\delta} \right) \frac{\kappa(\widehat{\Delta}^{L^{(0)}})}{\epsilon_{L^{(0)}}^2} \end{aligned} \quad (21)$$

where the last line follows from the second claim of Lemma 12. Now, by definition of $L^{(0)}$, $\epsilon_{L^{(0)}} \geq \left(\kappa(\widehat{\Delta}^{L^{(0)}}) \log(T)/T\right)^{1/3}$. Then, Equation (21) implies that

$$\sum_{l=1}^{L^{(0)}+1} \frac{\kappa(\widehat{\Delta}^l)}{\epsilon_l^2} \log\left(\frac{l(l+1)}{\delta}\right) \leq 20520 \log\left(\frac{6L^{(0)}}{\delta}\right) \kappa(\widehat{\Delta}^{L^{(0)}})^{1/3} \log(T)^{-2/3} T^{2/3}. \quad (22)$$

Moreover, we observe that during each phase l , but the last one, we sample at least

$$\max_{z \in \{-1, 1\}} \sum_{x \in \mathcal{X}_l^{(z)}} \tau_{l,x}^{(z)} \geq \frac{2(d+1)}{\delta_l^2} \log(kl(l+1)/\delta)$$

actions during the G-optimal explorations, so the number of phases $L^{(0)}$ is never larger than

$$\ell_T = 1 \vee \log_4(T).$$

Using this remark, together with Equations (20) and (22), we find that on $\overline{\mathcal{F}}$

$$\begin{aligned} \sum_{l=1}^{L^{(0)}} \sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \widehat{\Delta}_x^l + \mathbb{1}\left\{\text{Explore}_{L^{(0)}+1}^{(0)} = \text{False}\right\} \sum_{t \in \text{Exp}_{L^{(0)}+1}^{(0)}} \max_{x \in \mathcal{X}_{L^{(0)}+2}^{(-1)} \cup \mathcal{X}_{L^{(0)}+2}^{(+1)}} (x^* - x)^\top \gamma^* \\ \leq 2^{20} \log\left(\frac{6L^{(0)}}{\delta}\right) \kappa(\widehat{\Delta}^{L^{(0)}}) T^{2/3} \log(T)^{-2/3} + 42\ell_T. \end{aligned} \quad (23)$$

Bound on $\sum_{z \in \{-1, +1\}} \sum_{l=1}^{L_T} \left(\sum_{x \in \mathcal{X}_l^{(z)}} \mu_l^{(z)}(x) \right) \max_{x \in \mathcal{X}_l^{(z)}} (x^* - x)^\top \gamma^*$. We bound the remaining term in Equation (16) using the first claim in Lemma 11 and Lemma 10. On $\overline{\mathcal{F}}$,

$$\begin{aligned} \sum_{z \in \{-1, +1\}} \sum_{l=1}^{L_T} \left(\sum_{x \in \mathcal{X}_l^{(z)}} \mu_l^{(z)}(x) \right) \max_{x \in \mathcal{X}_l^{(z)}} (x^* - x)^\top \gamma^* &\leq 2 \sum_{l=1}^{L_T} \left(\frac{2(d+1)}{\epsilon_l^2} \log\left(\frac{kl(l+1)}{\delta}\right) + \frac{(d+1)(d+2)}{2} \right) 42\epsilon_l \\ &\leq \frac{336(d+1)}{\epsilon_{L_T}} \log\left(\frac{kL_T(1+L_T)}{\delta}\right) + 84(d+1)(d+2)\Delta_{\max} \\ &\leq 208(d+1)\kappa_*^{-1/3} T^{1/3} \log(T)^{-1/3} \log\left(\frac{kL_T(1+L_T)}{\delta}\right) \\ &\quad + 168(d+1)(d+2). \end{aligned} \quad (24)$$

Combing Equations (15), (16), (18), (23), and (24), and using $\delta = T^{-1}$, $\kappa(\widehat{\Delta}^{L^{(0)}}) \leq \kappa_*$ and $L_T \leq 4T/\log(2)$, we get for all $T \geq 1$

$$R_T \leq C \left(\kappa_*^{1/3} T^{2/3} \log(T)^{1/3} + (d \vee \kappa_*) \log(T) + d^2 + d\kappa_*^{-1/3} T^{1/3} \log(kT) \log(T)^{-1/3} \right)$$

for some absolute constant $C > 0$. Finally, for

$$T \geq \frac{((d \vee \kappa_*)^{3/2} \log(T)) \vee d^3}{\sqrt{\kappa_*}} \vee \frac{(d \log(kT))^3}{(\kappa_* \log(T))^2},$$

we get

$$R_T \leq C' \kappa_*^{1/3} T^{2/3} \log(T)^{1/3}.$$

C.2 Proof of Theorem 3

The beginning of the proof of Theorem 3 follows the same lines as the proof of Theorem 2. We begin by decomposing the regret as

$$R_T \leq \sum_{t \leq T} \mathbb{E}_{|\mathcal{F}} [(x^* - x_t)^\top \gamma^*] + 2T\mathbb{P}[\mathcal{F}]. \quad (25)$$

where \mathcal{F} is defined in Equation (14). On the one hand, Lemma 9 implies $T\mathbb{P}[\mathcal{F}] \leq 2\delta T$. Then, Equation (25) implies

$$\begin{aligned} R_T \leq & 4\delta T + \mathbb{E}_{|\mathcal{F}} \left[\sum_{z \in \{-1, +1\}} \sum_{l \geq 1}^{L^{(z)}+1} \sum_{t \in \text{Exp}_l^{(z)}} (x^* - x_t)^\top \gamma^* \right] + \mathbb{E}_{|\mathcal{F}} \left[\sum_{t \in \text{Recovery}} (x^* - x_t)^\top \gamma^* \right] \\ & + \mathbb{E}_{|\mathcal{F}} \left[\sum_{l=1}^{L^{(0)}} \sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \Delta_x \right] + \mathbb{E}_{|\mathcal{F}} \left[\mathbb{1} \left\{ \text{Explore}_{L^{(0)}+1}^{(0)} = \text{False} \right\} \sum_{t \in \text{Exp}_{L^{(0)}+1}^{(0)}} \max_{x \in \mathcal{X}_{L^{(0)}+2}^{(-1)} \cup \mathcal{X}_{L^{(0)}+2}^{(+1)}} (x^* - x)^\top \gamma^* \right] \end{aligned} \quad (26)$$

where \mathcal{F} is defined in Equation (14), and where we used the convention that the sum over an empty set is null.

Bound on $\mathbb{1} \left\{ \text{Explore}_{L^{(0)}+1}^{(0)} = \text{False} \right\} \sum_{t \in \text{Exp}_{L^{(0)}+1}^{(0)}} \max_{x \in \mathcal{X}_{L^{(0)}+1}^{(z)}} (x^* - x)^\top \gamma^*.$

Similarly to the proof of Theorem 2, we use Lemma 10 and Lemma 12 to show that on $\overline{\mathcal{F}}$

$$\mathbb{1} \left\{ \text{Explore}_{L^{(0)}+1}^{(0)} = \text{False} \right\} \sum_{t \in \text{Exp}_{L^{(0)}+1}^{(0)}} \max_{x \in \mathcal{X}_{L^{(0)}+1}^{(z)}} (x^* - x)^\top \gamma^* \leq 21 \sum_{x \in \mathcal{X}} \mu_{L^{(0)}+1}^{(0)}(x) \widehat{\Delta}_x^{L^{(0)}+1}. \quad (27)$$

Bound on $\sum_{z \in \{-1, +1\}} \sum_{l \geq 1}^{L^{(z)}+1} \sum_{t \in \text{Exp}_l^{(z)}} (x^* - x_t)^\top \gamma^*.$

Lemma 10 shows that for $l \leq L^{(z)}$, the actions in $\mathcal{X}_{l+1}^{(z)}$ are sub-optimal by at most an additional factor at most $21\epsilon_l$. Let us set $l_{\Delta_{\min}} = \lceil -\log_2(\Delta_{\min}/21) \rceil$, so that

$$\frac{\Delta_{\min}}{42} \leq \epsilon_{l_{\Delta_{\min}}} \leq \frac{\Delta_{\min}}{21}.$$

For $l \geq l_{\Delta_{\min}}$, we have $\mathcal{X}_{l+1}^{(-1)} \cup \mathcal{X}_{l+1}^{(+1)} = \{x_{z^*}\}$. Thus, $l^{(-z_{x^*})} \leq l_{\Delta_{\min}}$, and for $l \geq l_{\Delta_{\min}}$, the algorithm selects only x^* during the phase $\text{Exp}_l^{(z^*)}$. Then, combining Lemmas 11 and 10, and the fact that $L^{(z)} + 1 \leq \ell_T$, we find that, on $\overline{\mathcal{F}}$,

$$\begin{aligned} \sum_{z \in \{-1, +1\}} \sum_{l=1}^{L^{(z)}+1} \sum_{t \in \text{Exp}_l^{(z)}} (x^* - x_t)^\top \gamma^* & \leq \sum_{z \in \{-1, +1\}} \sum_{l=1}^{l_{\Delta_{\min}}+1 \wedge \ell_T} \left(\sum_{x \in \mathcal{X}_l^{(z)}} \mu_l^{(z)}(x) \right) \max_{x \in \mathcal{X}_l^{(z)}} (x^* - x)^\top \gamma^* \\ & \leq 2 \sum_{l=1}^{l_{\Delta_{\min}}+1 \wedge \ell_T} \left(\frac{2(d+1)}{\epsilon_l^2} \log \left(\frac{kl(l+1)}{\delta} \right) + \frac{(d+1)(d+2)}{2} \right) 42\epsilon_l \\ & \leq 84(d+1)(d+2) + 2^{l_{\Delta_{\min}}} \times 336(d+1) \log \left(\frac{k(1+\ell_T)(2+\ell_T)}{\delta} \right) \\ & \leq 84(d+1)(d+2) + \frac{14112(d+1)}{\Delta_{\min}} \log \left(\frac{k(1+\ell_T)(2+\ell_T)}{\delta} \right). \end{aligned} \quad (28)$$

Bound on
$$\sum_{t \in \text{Recovery}} (x^* - x_t)^\top \gamma^* + \sum_{l=1}^{L^{(0)}} \sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \Delta_x + \sum_{x \in \mathcal{X}} \mu_{L^{(0)}+1}^{(0)}(x) \widehat{\Delta}_x^{L^{(0)}+1}.$$

We use the following lemma to bound the number of phases necessary to eliminate the sub-optimal group.

Lemma 13. *On the event $\overline{\mathcal{F}}$ defined in Equation (14), for $l \geq 1$ such that $\epsilon_l \leq \frac{\Delta_{\neq}}{8}$ and $\text{Explore}_L^{(0)} = \text{True}$, $\widehat{z}_{l+1}^* = z_{x^*}$.*

Let $l_{\Delta_{\neq}} = \lceil -\log(\Delta_{\neq}/8)/\log(2) \rceil$ be such that

$$\frac{\Delta_{\neq}}{16} \leq \epsilon_{l_{\Delta_{\neq}}} \leq \frac{\Delta_{\neq}}{8}. \quad (29)$$

Lemma 13 implies that on $\overline{\mathcal{F}}$, $L^{(0)} \leq l_{\Delta_{\neq}}$.

To bound the remaining terms, we consider two cases, corresponding to $\text{Recovery} = \emptyset$ and $\text{Recovery} \neq \emptyset$.

Case 1: $\text{Recovery} = \emptyset$. Our case assumption implies that

$$\sum_{t \in \text{Recovery}} (x^* - x_t)^\top \gamma^* = 0. \quad (30)$$

Lemma 12 implies that

$$\sum_{l=1}^{L^{(0)}} \sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \Delta_x + \sum_{x \in \mathcal{X}} \mu_{L^{(0)}+1}^{(0)}(x) \widehat{\Delta}_x^{L^{(0)}+1} \leq \sum_{l=1}^{L^{(0)}+1} \sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \widehat{\Delta}_x^l.$$

Moreover, $L^{(0)} \leq l_{\Delta_{\neq}} \wedge \ell_T$, so on $\overline{\mathcal{F}}$

$$\sum_{l=1}^{L^{(0)}+1} \sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \widehat{\Delta}_x^l \leq \sum_{l=1}^{(l_{\Delta_{\neq}} \wedge \ell_T)+1} \sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \widehat{\Delta}_x^l.$$

Using Lemma 11, we find that on $\overline{\mathcal{F}}$

$$\begin{aligned} \sum_{l=1}^{(l_{\Delta_{\neq}} \wedge \ell_T)+1} \sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \widehat{\Delta}_x^l &\leq \sum_{l=1}^{(l_{\Delta_{\neq}} \wedge \ell_T)+1} \frac{2\kappa(\widehat{\Delta}^l)}{\epsilon_l^2} \log \left(\frac{l(l+1)}{\delta} \right) + 2(d+1)(\ell_T+1) \\ &\leq 2 \log \left(\frac{(\ell_T+1)(\ell_T+2)}{\delta} \right) \sum_{l=1}^{l_{\Delta_{\neq}}+1} \frac{\kappa(\widehat{\Delta}^l)}{\epsilon_l^2} + 2(d+1)(\ell_T+1). \end{aligned}$$

Using Lemma 12 with $\tau = \Delta_{\neq}$ and (29), we have on $\overline{\mathcal{F}}$

$$\begin{aligned} \sum_{l=1}^{l_{\Delta_{\neq}}+1} \frac{\kappa(\widehat{\Delta}^l)}{\epsilon_l^2} &\leq 513\kappa(\Delta \vee \Delta_{\neq}) \sum_{l=1}^{l_{\Delta_{\neq}}+1} (\epsilon_l^{-2} + \epsilon_l^{-1}/\Delta_{\neq}) \\ &\leq \frac{2^{20}\kappa(\Delta \vee \Delta_{\neq})}{\Delta_{\neq}^2}. \end{aligned}$$

We obtain on $\overline{\mathcal{F}}$

$$\sum_{l=1}^{L^{(0)}+1} \sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \widehat{\Delta}_x^l \leq 2^{23} \log \left(\frac{(\ell_T+1)(\ell_T+2)}{\delta} \right) \frac{\kappa(\Delta \vee \Delta_{\neq})}{\Delta_{\neq}^2} + 2(d+1)(\ell_T+1). \quad (31)$$

Combining Equations (28), (27), (30), and (31), we find that on $\overline{\mathcal{F}}$, when $\text{Recovery} = \emptyset$, there exists an absolute constant $c > 0$ such that for $\delta = T^{-1}$,

$$\begin{aligned} & \sum_{z \in \{-1, +1\}} \sum_{l \geq 1}^{L^{(z)}+1} \sum_{t \in \text{Exp}_l^{(z)}} (x^* - x_t)^\top \gamma^* + \sum_{t \in \text{Recovery}} (x^* - x_t)^\top \gamma^* + \sum_{l=1}^{L^{(0)}} \sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \Delta_x \\ & + \mathbb{1}\{\text{Explore}_{L^{(0)}+1}^{(0)} = \text{False}\} \sum_{t \in \text{Exp}_{L^{(0)}+1}^{(0)}} \max_{x \in \mathcal{X}_{L^{(0)}+2}^{(-1)} \cup \mathcal{X}_{L^{(0)}+2}^{(+1)}} (x^* - x)^\top \gamma^* \\ & \leq c \left(d^2 + \left(\frac{d}{\Delta_{\min}} \vee \frac{\kappa(\Delta \vee \Delta_{\neq})}{\Delta_{\neq}^2} \right) \log(T) + \frac{d}{\Delta_{\min}} \log(k) \right). \end{aligned} \quad (32)$$

Case 2: $\text{Recovery} \neq \emptyset$. In this case, the algorithm enters Recovery at phase $L^{(0)}$, so $\text{Explore}_{L^{(0)}+1}^{(0)} = \text{False}$ and $\text{Exp}_{L^{(0)}+1}^{(0)} = \emptyset$, and

$$\mathbb{1}\{\text{Explore}_{L^{(0)}+1}^{(0)} = \text{False}\} \sum_{t \in \text{Exp}_{L^{(0)}+1}^{(0)}} \max_{x \in \mathcal{X}_{L^{(0)}+2}^{(-1)} \cup \mathcal{X}_{L^{(0)}+2}^{(+1)}} (x^* - x)^\top \gamma^* = 0. \quad (33)$$

Using Lemma 10, we see that

$$\sum_{t \in \text{Recovery}} (x^* - x_t)^\top \gamma^* \leq 21T\epsilon_{L^{(0)}+1}.$$

On the other hand, in the Recovery phase, $\epsilon_{L^{(0)}+1} \leq \left(\kappa(\widehat{\Delta}^{L^{(0)}+1}) \log(T)/T \right)^{1/3}$. Thus,

$$\sum_{t \in \text{Recovery}} (x^* - x_t)^\top \gamma^* \leq \frac{21\kappa(\widehat{\Delta}^{L^{(0)}+1}) \log(T)}{\epsilon_{L^{(0)}+1}^2}.$$

Now, Lemma 11 show that

$$\sum_{l=1}^{L^{(0)}} \sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \Delta_x \leq 4 \log(2L^{(0)}\delta^{-1}) \sum_{l=1}^{L^{(0)}} \frac{\kappa(\widehat{\Delta}^l)}{\epsilon_l^2} + 4dL^{(0)}.$$

Combining these results, and using $L^{(0)} \leq \ell_T$, we see that

$$\sum_{t \in \text{Recovery}} (x^* - x_t)^\top \gamma^* + \sum_{l=1}^{L^{(0)}} \sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \Delta_x \leq 4dL^{(0)} + (4 \log(2\ell_T\delta^{-1}) \vee 21 \log(T)) \sum_{l=1}^{L^{(0)}+1} \frac{\kappa(\widehat{\Delta}^l)}{\epsilon_l^2}. \quad (34)$$

Using Lemma 12 with $\tau = \epsilon_{L^{(0)}}$, we see that

$$\begin{aligned} \sum_{l=1}^{L^{(0)}+1} \frac{\kappa(\widehat{\Delta}^l)}{\epsilon_l^2} & \leq 513 \sum_{l=1}^{L^{(0)}+1} \frac{\kappa(\Delta \vee \epsilon_{L^{(0)}})}{\epsilon_l^2} + 513 \sum_{l=1}^{L^{(0)}+1} \frac{\kappa(\Delta \vee \epsilon_{L^{(0)}})}{\epsilon_{L^{(0)}} \epsilon_l} \\ & \leq 8202 \frac{\kappa(\Delta \vee \epsilon_{L^{(0)}})}{\epsilon_{L^{(0)}}^2} + 2052 \frac{\kappa(\Delta \vee \epsilon_{L^{(0)}})}{\epsilon_{L^{(0)}}^2}. \end{aligned}$$

Now, the algorithm enters the Recovery phase before finding the best group, so we must have $L^{(0)} \leq l_{\Delta_{\neq}}$. This implies that

$$\sum_{l=1}^{L^{(0)}+1} \frac{\kappa(\widehat{\Delta}^l)}{\epsilon_l^2} \leq 2^{22} \frac{\kappa(\Delta \vee \epsilon_{L^{(0)}})}{\Delta_{\neq}^2}.$$

Finally, note that $L^{(0)} \geq L_T$, so $\epsilon_{L^{(0)}} \leq \epsilon_{L_T} = \varepsilon_T$, and

$$\sum_{l=1}^{L^{(0)}+1} \frac{\kappa(\widehat{\Delta}^l)}{\epsilon_l^2} \leq 2^{22} \frac{\kappa(\Delta \vee \varepsilon_T)}{\Delta_{\neq}^2}. \quad (35)$$

Combining Equations (28), (33), (34), and (35), we find that on $\overline{\mathcal{F}}$, when $\text{Recovery} \neq \emptyset$, there exists an absolute constant $c > 0$ such that for $\delta = T^{-1}$,

$$\begin{aligned} & \sum_{z \in \{-1, +1\}} \sum_{l \geq 1}^{L^{(z)}+1} \sum_{t \in \text{Exp}_l^{(z)}} (x^* - x_t)^\top \gamma^* + \sum_{t \in \text{Recovery}} (x^* - x_t)^\top \gamma^* + \sum_{l=1}^{L^{(0)}} \sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \Delta_x \\ & + \mathbf{1}\{\text{Explore}_{L^{(0)}+1}^{(0)} = \text{False}\} \sum_{t \in \text{Exp}_{L^{(0)}+1}^{(0)}} \max_{x \in \mathcal{X}_{L^{(0)}+2}^{(-1)} \cup \mathcal{X}_{L^{(0)}+2}^{(+1)}} (x^* - x)^\top \gamma^* \\ & \leq c \left(d^2 + \left(\frac{d}{\Delta_{\min}} \vee \frac{\kappa(\Delta \vee \varepsilon_T)}{\Delta_{\neq}^2} \right) \log(T) + \frac{d \log(k)}{\Delta_{\min}} \right). \end{aligned} \quad (36)$$

Conclusion We conclude the proof of Theorem 3 by combining Equations (26), (32) and (36).

C.3 Proof of Theorem 4

We prove Theorem 4 for the set of actions \mathcal{A} defined in the sketch of proof of Theorem 4: $\mathcal{A} = \left\{ \begin{pmatrix} x_1 \\ z_{x_1} \end{pmatrix}, \dots, \begin{pmatrix} x_{d+1} \\ z_{x_{d+1}} \end{pmatrix} \right\}$, where $\begin{pmatrix} x_i \\ z_{x_i} \end{pmatrix} = e_i + e_{d+1}$, for $i \in \{2, \dots, \lfloor d/2 \rfloor\}$, $\begin{pmatrix} x_i \\ z_{x_i} \end{pmatrix} = e_i - e_{d+1}$ for $i \in \{\lfloor d/2 \rfloor + 1, \dots, d\}$, and $\begin{pmatrix} x_{d+1} \\ z_{x_{d+1}} \end{pmatrix} = -\left(1 - \frac{2}{\sqrt{v}+1}\right) e_1 - e_{d+1}$. Then, by Lemma 6, for this choice of action set, we have $\mathcal{A} \in \mathbf{A}_{v,d}$.

We consider the set of bandit problems defined in the sketch of proof of Theorem 4: for $i \in \{1, \dots, \lfloor d/2 \rfloor + 1\}$

Problem i is characterized by the parameter $\theta^{(i)}$, where $\theta^{(i)} = \begin{pmatrix} \gamma^{(i)} \\ \omega^{(i)} \end{pmatrix}$ is defined as:

$$\begin{aligned} \gamma^{(1)} &= \frac{1 + \Delta_{\neq} - \Delta_{\min}}{2} \left(\sum_{1 \leq j \leq \lfloor d/2 \rfloor} e_j \right) + \frac{1 - \Delta_{\neq} - \Delta_{\min}}{2} \left(\sum_{\lfloor d/2 \rfloor + 1 \leq j \leq d} e_j \right) + \Delta_{\min} e_1 + \Delta_{\min} e_{\lfloor d/2 \rfloor + 1} \\ \gamma^{(i)} &= \gamma^{(1)} + 2\Delta_{\min} e_i + 2\Delta_{\min} e_{\lfloor d/2 \rfloor + i} \quad \forall i \in \{2, \dots, \lfloor d/2 \rfloor\} \\ \gamma^{(\lfloor d/2 \rfloor + 1)} &= \frac{1 - \Delta_{\neq} - \Delta_{\min}}{2} \left(\sum_{1 \leq j \leq \lfloor d/2 \rfloor} e_j \right) + \frac{1 + \Delta_{\neq} - \Delta_{\min}}{2} \left(\sum_{\lfloor d/2 \rfloor + 1 \leq j \leq d} e_j \right) + \Delta_{\min} e_1 + \Delta_{\min} e_{\lfloor d/2 \rfloor + 1}, \end{aligned}$$

and the bias parameters are defined as $\omega^{(i)} = -\frac{\Delta_{\neq}}{2} \quad \forall i \in \{1, \dots, \lfloor d/2 \rfloor\}$, and otherwise $\omega^{(\lfloor d/2 \rfloor + 1)} = \frac{\Delta_{\neq}}{2}$. We write $\mathbb{E}^{(i)}, \mathbb{P}^{(i)}, R_T^{(i)}$ for resp. the probability, expectation, and regret, in **Problem i**. Note that this choice of parameters ensures that $\forall i \in \{1, \dots, \lfloor d/2 \rfloor + 1\}$, $\theta^{(i)} \in \Theta_{\Delta_{\min}, \Delta_{\neq}}^{\mathcal{A}}$.

The following facts hold:

Fact 1 For any $i \in \{1, \dots, \lfloor d/2 \rfloor + 1\}$, action x_i is the unique optimal action in **Problem i**. Since $1/2 \geq \Delta_{\neq} \geq \Delta_{\min}$, sampling any other (sub-optimal) action leads to an instantaneous regret of at least Δ_{\min} . Moreover, choosing an action in the group $-z_i$ leads to an instantaneous regret of at least Δ_{\neq} .

Fact 2 In **Problem i** for any $i \in \{1, \dots, \lfloor d/2 \rfloor + 1\}$, action $d+1$ is very sub-optimal and sampling it leads to an instantaneous regret higher than $(1 - 2/(\sqrt{v}+1))(1 - \Delta_{\neq} + \Delta_{\min}) + (1 + \Delta_{\neq} + \Delta_{\min})/2 \geq 1/2$, since $v \geq 1$ and $1/2 \geq \Delta_{\neq} \geq \Delta_{\min}$.

Fact 3 In **Problem i**, for $i \in \{1, \dots, \lfloor d/2 \rfloor + 1\}$, when sampling action x_j at time, t the distribution of the observation does not depend on t or on the past (except through the choice of x_j) and is $\mathbb{P}_j^{(i)}$. It is characterized as:

$$\begin{aligned} \forall i \in \{1, \dots, \lfloor d/2 \rfloor + 1\}, \mathbb{P}_1^{(i)}, \mathbb{P}_{\lfloor d/2 \rfloor + 1}^{(i)} & \text{ are } \mathcal{N}((1 + \Delta_{\min})/2, 1) \\ \forall i \in \{1, \dots, \lfloor d/2 \rfloor + 1\}, \forall j \in \{2, \dots, d\} \setminus \{\lfloor d/2 \rfloor + 1, i, \lfloor d/2 \rfloor + i\}, \mathbb{P}_j^{(i)} & \text{ is } \mathcal{N}((1 - \Delta_{\min})/2, 1), \\ \forall i \in \{2, \lfloor d/2 \rfloor\}, \mathbb{P}_i^{(i)} & \text{ is } \mathcal{N}((1 + 3\Delta_{\min})/2, 1) \quad \mathbb{P}_{\lfloor d/2 \rfloor + i}^{(i)} \text{ is } \mathcal{N}((1 + 3\Delta_{\min})/2, 1) \\ \forall i \in \{1, \lfloor d/2 \rfloor\}, \mathbb{P}_{d+1}^{(i)} & \text{ is } \mathcal{N}(-(1 - \alpha)(1 + \Delta_{\neq} + \Delta_{\min})/2 + \Delta_{\neq}/2, 1), \\ \mathbb{P}_{d+1}^{(\lfloor d/2 \rfloor + 1)} & \text{ is } \mathcal{N}(-(1 - \alpha)(1 - \Delta_{\neq} + \Delta_{\min})/2 - \Delta_{\neq}/2, 1) \quad \text{where } \alpha = 2/(\sqrt{v} + 1). \end{aligned}$$

So that:

Fact 3.1 For any $i \in \{2, \dots, \lfloor d/2 \rfloor\}$, between **Problem 1** and **Problem i**, the only actions that provide different evaluations when sampled are action i and action $\lfloor d/2 \rfloor + i$, and the mean gaps in both cases is $2\Delta_{\min}$.

Fact 3.2 Between **Problem 1** and **Problem** $\lfloor d/2 \rfloor + 1$, the only action that provide different evaluation when sampled is action $d + 1$, and the mean gap in this case is $\alpha\Delta_{\neq}$.

For $j \leq d + 1$, we write $N_j(T)$ for the total number of times action x_j has been selected before time T . Then, for $j \in \{1, \dots, \lfloor d/2 \rfloor\}$, let $E^{(j)} = \{N_i(T) \leq T/2\}$. Note that for $i \in \{1, \dots, \lfloor d/2 \rfloor\}$, in **Problem i** the action x_i is the optimal action. Therefore, for any efficient algorithm, for all $i \in \{1, \dots, \lfloor d/2 \rfloor\}$ the event $E^{(i)}$ should have a low probability under $\mathbb{P}^{(i)}$. Indeed, for $i \in \{1, \dots, \lfloor d/2 \rfloor\}$, the regret of the algorithm under **Problem i** can be lower-bounded as follows - see **Facts 1 and 2**:

$$R_T^{(i)} \geq \sum_{j \leq \lfloor d/2 \rfloor, j \neq i} \mathbb{E}^{(i)}[N_j(T)] \Delta_{\min} + \sum_{\lfloor d/2 \rfloor + 1 \leq j \leq d} \mathbb{E}^{(i)}[N_j(T)] \Delta_{\neq} + \frac{\mathbb{E}^{(i)}[N_{d+1}(T)]}{2}. \quad (37)$$

Since $\sum_j \mathbb{E}^{(i)}[N_j(T)] = T$ and $\Delta_{\min} \leq \Delta_{\neq} \leq \frac{1}{2}$, this implies together with **Facts 1**:

$$R_T^{(i)} \geq \left(T - \mathbb{E}^{(i)}[N_i(T)]\right) \Delta_{\min}$$

Using the definition of $E^{(i)}$, we find that

$$R_T^{(i)} \geq \frac{T\Delta_{\min}}{2} \mathbb{P}^{(i)}(E^{(i)}). \quad (38)$$

In particular for **Problem 1**, for any $i \in \{1, \dots, \lfloor d/2 \rfloor\}$,

$$R_T^{(1)} \geq \frac{T\Delta_{\min}}{2} \mathbb{P}^{(1)}(\overline{E^{(i)}}). \quad (39)$$

since $E^{(1)} \supset \overline{E^{(i)}}$.

Similarly, let us also define the event $F = \left\{ \sum_{i \leq \lfloor d/2 \rfloor} N_i(T) \geq T/2 \right\}$. Then, in **Problem 1**, the group 1 contains the optimal action, and so for any efficient algorithm, the event F should have a low probability under $\mathbb{P}^{(1)}$. Indeed, Equation (37) also implies

$$R_T^{(1)} \geq \left(T - \mathbb{E}^{(1)} \left[\sum_{i \leq \lfloor d/2 \rfloor} N_i(T) \right]\right) \Delta_{\neq} \geq \frac{T\Delta_{\neq}}{2} \mathbb{P}^{(1)}(\overline{F}). \quad (40)$$

On the other hand, for any efficient algorithm, the event F should have high probability under $\mathbb{P}^{(\lfloor d/2 \rfloor + 1)}$. Indeed, under problem **Problem** $\lfloor d/2 \rfloor + 1$, the regret can be lower-bounded as follows - see **Facts 1 and 2**:

$$R_T^{(\lfloor d/2 \rfloor + 1)} \geq \sum_{j \leq \lfloor d/2 \rfloor} \mathbb{E}^{(\lfloor d/2 \rfloor + 1)}[N_j(T)] \Delta_{\neq} + \sum_{\lfloor d/2 \rfloor + 2 \leq j \leq d} \mathbb{E}^{(\lfloor d/2 \rfloor + 1)}[N_j(T)] \Delta_{\min} + \frac{\mathbb{E}^{(\lfloor d/2 \rfloor + 1)}[N_{d+1}(T)]}{2}.$$

which implies that

$$R_T^{(\lfloor d/2 \rfloor + 1)} \geq \sum_{j \leq \lfloor d/2 \rfloor} \mathbb{E}^{(\lfloor d/2 \rfloor + 1)} [N_j(T)] \Delta_{\neq} \geq \frac{T \Delta_{\neq}}{2} \mathbb{P}^{(\lfloor d/2 \rfloor + 1)}(F). \quad (41)$$

Now, Bretagnolle-Huber inequality (see, e.g., Theorem 14.2 in [Lattimore and Szepesvári, 2020]) implies that for all $i \in \{2, \dots, \lfloor d/2 \rfloor\}$,

$$\frac{1}{2} \exp \left(-KL \left(\mathbb{P}^{(1)}, \mathbb{P}^{(i)} \right) \right) \leq \mathbb{P}^{(i)} \left(E^{(i)} \right) + \mathbb{P}^{(1)} \left(\overline{E^{(i)}} \right) \quad (42)$$

and that

$$\frac{1}{2} \exp \left(-KL \left(\mathbb{P}^{(1)}, \mathbb{P}^{(\lfloor d/2 \rfloor + 1)} \right) \right) \leq \mathbb{P}^{(\lfloor d/2 \rfloor + 1)}(F) + \mathbb{P}^{(1)}(\overline{F}). \quad (43)$$

On the one hand, Equation (42) implies that for any $i \in \{2, \dots, \lfloor d/2 \rfloor\}$,

$$\begin{aligned} KL \left(\mathbb{P}^{(1)}, \mathbb{P}^{(i)} \right) &\geq -\log \left(2\mathbb{P}^{(i)} \left(E^{(i)} \right) + 2\mathbb{P}^{(1)} \left(\overline{E^{(i)}} \right) \right) \\ &\geq \log(T) - \log \left(2T\mathbb{P}^{(i)} \left(E^{(i)} \right) + 2T\mathbb{P}^{(1)} \left(\overline{E^{(i)}} \right) \right). \end{aligned} \quad (44)$$

Combining Equations (38), (39), and (44), we find that

$$KL \left(\mathbb{P}^{(1)}, \mathbb{P}^{(i)} \right) \geq \log(T) - \log \left(\frac{4(R_T^{(i)} + R_T^{(1)})}{\Delta_{\min}} \right). \quad (45)$$

On the other hand, Equation (43) implies that

$$\begin{aligned} KL \left(\mathbb{P}^{(1)}, \mathbb{P}^{(\lfloor d/2 \rfloor + 1)} \right) &\geq -\log \left(2\mathbb{P}^{(\lfloor d/2 \rfloor + 1)}(F) + 2\mathbb{P}^{(1)}(\overline{F}) \right) \\ &\geq \log(T) - \log \left(2T\mathbb{P}^{(\lfloor d/2 \rfloor + 1)}(F) + 2T\mathbb{P}^{(1)}(\overline{F}) \right). \end{aligned} \quad (46)$$

Combining Equations (38), (39), and (46), we find that

$$KL \left(\mathbb{P}^{(1)}, \mathbb{P}^{(\lfloor d/2 \rfloor + 1)} \right) \geq \log(T) - \log \left(\frac{4(R_T^{(\lfloor d/2 \rfloor + 1)} + R_T^{(1)})}{\Delta_{\neq}} \right). \quad (47)$$

Also, note that for all $i \in \{2, \dots, \lfloor d/2 \rfloor + 1\}$, the Kullback-Leibler divergence between $\mathbb{P}^{(1)}$ and $\mathbb{P}^{(i)}$ can be decomposed as follows (see, e.g., Lemma 15.1 in [Lattimore and Szepesvári, 2020]) :

$$KL(\mathbb{P}^{(1)}, \mathbb{P}^{(i)}) = \sum_{j \leq d+1} \mathbb{E}^{(1)} [N_j(T)] KL(\mathbb{P}_j^{(1)}, \mathbb{P}_j^{(i)}). \quad (48)$$

Lower bound in $d\Delta_{\min}^{-1} \log T$. By design, for $i \in \{2, \dots, \lfloor d/2 \rfloor\}$, all actions but x_i and $x_{\lfloor d \rfloor + i}$ have the same distribution under $\mathbb{P}^{(1)}$ and $\mathbb{P}^{(i)}$ - see **Fact 3.1**. Then, Equation (48) becomes from **Fact 3.1** and from the expression of KL divergence between standard Gaussian distributions:

$$KL(\mathbb{P}^{(1)}, \mathbb{P}^{(i)}) = \frac{4\Delta_{\min}^2}{2} \mathbb{E}^{(1)} [N_i(T)] + \frac{4\Delta_{\min}^2}{2} \mathbb{E}^{(1)} [N_{\lfloor d \rfloor + i}(T)].$$

So that, summing over $i \in \{2, \dots, \lfloor d/2 \rfloor\}$, and by **Fact 1**:

$$\sum_{i \in \{2, \dots, \lfloor d/2 \rfloor\}} KL(\mathbb{P}^{(1)}, \mathbb{P}^{(i)}) \leq 2\Delta_{\min} R_T^{(1)}.$$

So that by Equation (45) (summing over $i \in \{2, \dots, \lfloor d/2 \rfloor\}$):

$$\begin{aligned} 2\Delta_{\min} R_T^{(1)} &\geq \sum_{i \in \{2, \dots, \lfloor d/2 \rfloor\}} \left[\log(T) - \log \left(\frac{4(R_T^{(i)} + R_T^{(1)})}{\Delta_{\min}} \right) \right] \\ &= (\lfloor d/2 \rfloor - 1) \log(T) - \sum_{i \in \{2, \dots, \lfloor d/2 \rfloor\}} \log \left(\frac{4(R_T^{(i)} + R_T^{(1)})}{\Delta_{\min}} \right). \end{aligned}$$

Let us assume that our algorithm satisfies $\max_{i \leq \lfloor d/2 \rfloor} R_T^{(i)} \leq \frac{d \log(T)}{\Delta_{\min}}$ - otherwise the bound immediately follows for this algorithm. Then

$$\begin{aligned} R_T^{(1)} &\geq \frac{1}{2\Delta_{\min}} (\lfloor d/2 \rfloor - 1) \log(T) - \frac{1}{2\Delta_{\min}} \sum_{i \in \{2, \dots, \lfloor d/2 \rfloor\}} \log \left(\frac{8d \log T}{\Delta_{\min}^2} \right) \\ &\geq \frac{1}{2\Delta_{\min}} (\lfloor d/2 \rfloor - 1) \left[\log(T) - \log \left(\frac{8d \log(T)}{\Delta_{\min}^2} \right) \right]. \end{aligned} \quad (49)$$

Sine $d \geq 4$, we note that $\lfloor d/2 \rfloor - 1 \geq d/5$. This concludes the proof for this part of the bound.

Lower bound in $v\Delta_{\neq}^{-2} \log T$. By design, all actions but x_{d+1} have the same evaluation under **Problem 1** and **Problem $\lfloor d/2 \rfloor + 1$** - see **Fact 3.2**. Then, by **Fact 3.2** and the expression between the KL divergence of standard Gaussians, Equation (48) becomes

$$KL(\mathbb{P}^{(1)}, \mathbb{P}^{(\lfloor d/2 \rfloor + 1)}) = \mathbb{E}^{(1)} [N_{d+1}(T)] \frac{(\alpha \Delta_{\neq})^2}{2} = \frac{1}{2} \mathbb{E}^{(1)} [N_{d+1}(T)] \left(\frac{2\Delta_{\neq}}{\sqrt{v} + 1} \right)^2.$$

Combined with equation (47), this implies that

$$\frac{1}{2} \mathbb{E}^{(1)} [N_{d+1}(T)] \left(\frac{2\Delta_{\neq}}{\sqrt{v} + 1} \right)^2 \geq \log(T) - \log \left(\frac{4(R_T^{(\lfloor d/2 \rfloor + 1)} + R_T^{(1)})}{\Delta_{\neq}} \right). \quad (50)$$

Let us assume that our algorithm satisfies $\max_{i \leq \lfloor d/2 \rfloor + 1} R_T^{(i)} \leq \frac{v \log(T)}{\Delta_{\neq}^2}$ - otherwise the bound immediately follows for this algorithm. We then have

$$\frac{1}{2} \mathbb{E}^{(1)} [N_{d+1}(T)] \left(\frac{2\Delta_{\neq}}{\sqrt{v} + 1} \right)^2 \geq \log(T) - \log \left(\frac{8v \log(T)}{\Delta_{\neq}^3} \right).$$

Using Equation (37), we find that

$$R_T^{(1)} \geq \frac{v+1}{4\Delta_{\neq}^2} \left[\log(T) - \log \left(\frac{8v \log(T)}{\Delta_{\neq}^3} \right) \right]. \quad (51)$$

Lower bound in $v\Delta_{\neq}^{-2}$. Let us assume that our algorithm satisfies $\max_{i \leq \lfloor d/2 \rfloor + 1} R_T^{(i)} \leq \frac{v}{\Delta_{\neq}^2}$ - otherwise the bound immediately follows for this algorithm. Then, Equation (50) implies

$$\frac{1}{2} \mathbb{E}^{(1)} [N_{d+1}(T)] \left(\frac{2\Delta_{\neq}}{\sqrt{v}} \right)^2 \geq \log(T) - \log \left(\frac{8v}{\Delta_{\neq}^3} \right).$$

Using again Equation (37), we find that

$$R_T^{(1)} \geq \frac{v+1}{4\Delta_{\neq}^2} \log \left(\frac{T\Delta_{\neq}^3}{8v} \right). \quad (52)$$

We conclude the proof of Theorem 4 by combining Equations (49), (51) and (52).

C.4 Proof of Corollary 1

The proof is immediate if we apply Theorem 4 for $v = \frac{\kappa}{2}$. Then, Theorem 4 ensures that there exists an set of actions $\mathcal{A} \in \mathbf{A}_{\frac{\kappa}{2}, d}^{\kappa}$ such that for all budget $T \geq 1$, there exists a parameter $\theta \in \Theta_{\Delta_{\min}, \Delta_{\neq}}^{\mathcal{A}}$ such that

$$R_T^{\mathcal{A}, \theta} \geq \left[\frac{1}{10\Delta_{\min}} \log(T) \left[1 - \frac{\log\left(\frac{8d \log(T)}{\Delta_{\min}^2}\right)}{\log(T)} \right] \right] \vee \left[\frac{\frac{\kappa}{2} + 1}{4\Delta_{\neq}^2} \log(T) \left[1 - \frac{\log\left(\frac{16v \log(T)}{\Delta_{\neq}^3}\right)}{\log(T)} \right] \right].$$

Moreover, for this parameter, $\kappa(\Delta) \leq 2 \times \frac{\kappa}{2}$, which shows that $\theta \in \Theta_{\Delta_{\min}, \Delta_{\neq}, \kappa}^{\mathcal{A}}$.

C.5 Proof of Corollary 2

To prove Corollary 2, we rely on the set of actions \mathcal{A} defined as $\mathcal{A} = \left\{ \begin{pmatrix} x_1 \\ z_{x_1} \end{pmatrix}, \dots, \begin{pmatrix} x_{d+1} \\ z_{x_{d+1}} \end{pmatrix} \right\}$, where $\begin{pmatrix} x_i \\ z_{x_i} \end{pmatrix} = e_i + e_{d+1}$, for $i \in \{2, \dots, \lfloor d/2 \rfloor\}$, $\begin{pmatrix} x_i \\ z_{x_i} \end{pmatrix} = e_i - e_{d+1}$ for $i \in \{\lfloor d/2 \rfloor + 1, \dots, d\}$, and $\begin{pmatrix} x_{d+1} \\ z_{x_{d+1}} \end{pmatrix} = -\left(1 - \frac{2}{\sqrt{v}+1}\right) e_1 - e_{d+1}$. Then, by Lemma 6, for this choice of actions \mathcal{A} , $\mathcal{A} \in \mathbf{A}_{v, d}$.

We consider the set of bandit problems defined in the sketch of proof of Theorem 4 for $\Delta_{\min} = \Delta_{\neq} = \rho_T = (8ev/T)^{1/3}$. Note that when $T \geq 4 \times 8^3 v$, $\rho_T \leq \frac{1}{8}$. Then, for $i \in \{1, \dots, \lfloor d/2 \rfloor + 1\}$ **Problem i** is characterized by the parameter $\theta_T^{(i)}$, where $\theta_T^{(i)} = \begin{pmatrix} \gamma^{(i)} \\ \omega^{(i)} \end{pmatrix}$ is defined as:

$$\begin{aligned} \gamma^{(1)} &= \frac{1}{2} \left(\sum_{1 \leq j \leq \lfloor d/2 \rfloor} e_j \right) + \frac{1 - 2\rho_T}{2} \left(\sum_{\lfloor d/2 \rfloor + 1 \leq j \leq d} e_j \right) + \rho_T e_1 + \rho_T e_{\lfloor d/2 \rfloor + 1} \\ \gamma^{(i)} &= \gamma^{(1)} + 2\rho_T e_i + 2\rho_T e_{\lfloor d/2 \rfloor + i} \quad \forall i \in \{2, \dots, \lfloor d/2 \rfloor\} \\ \gamma^{(\lfloor d/2 \rfloor + 1)} &= \frac{1 - 2\rho_T}{2} \left(\sum_{1 \leq j \leq \lfloor d/2 \rfloor} e_j \right) + \frac{1}{2} \left(\sum_{\lfloor d/2 \rfloor + 1 \leq j \leq d} e_j \right) + \rho_T e_1 + \rho_T e_{\lfloor d/2 \rfloor + 1}, \end{aligned}$$

and the bias parameters are defined as $\omega^{(i)} = -\frac{\rho_T}{2} \forall i \in \{1, \dots, \lfloor d/2 \rfloor\}$, and otherwise $\omega^{(\lfloor d/2 \rfloor + 1)} = \frac{\rho_T}{2}$. We write $\mathbb{E}^{(i)}, \mathbb{P}^{(i)}, R_T^{(i)}$ for resp. the probability, expectation, and regret, in **Problem i**. Note that this choice of parameters ensures that $\forall i \in \{1, \dots, \lfloor d/2 \rfloor + 1\}$, $\theta_T^{(i)} \in \Theta_{\rho_T, \rho_T}^{\mathcal{A}}$.

Following the lines of the proof of Theorem 4, we find that for some $i \in \{1, \dots, \lfloor d/2 \rfloor + 1\}$, there exists a **Problem i** such that for $\theta_T = \theta_T^{(i)}$,

$$R_T^{\mathcal{A}, \theta_T} \geq \frac{v+1}{4(8ev/T)^{2/3}} \left[1 \wedge \log\left(\frac{T(8ev/T)}{8v}\right) \right]. \quad (53)$$

To conclude the proof of Corollary 2, note that on the one hand, for all **Problems i**, $\Delta_{\max} \geq 1/2$, so $\kappa_* = \Delta_{\max} v \geq \frac{v}{2}$. On the other hand, we always have $\Delta_{\max} \leq 2$, so $\kappa_* \leq 2v$.

C.6 Auxiliary Lemmas

C.6.1 Proof of Lemmas 1 and 2

We begin by proving Lemma 1. Recall that π is a G-optimal design for the set $\{a_x : x \in \mathcal{X}\}$, and that μ is defined as $\mu(x) = \lceil m\pi(x) \rceil$ for all $x \in \mathcal{X}$.

We first observe that $V(\pi) = A_\pi^\top A_\pi$, where A_π is the matrix with lines given by $[\sqrt{\pi(x)} a_x^\top]_{x \in \mathcal{X}}$. Since the supports of μ and π are the same, we get that $\text{Range}(A_\pi^\top) = \text{Range}(A_\mu^\top)$. As a consequence

$$\text{Range}(V(\pi)) = \text{Range}(A_\pi^\top) = \text{Range}(A_\mu^\top) = \text{Range}(V(\mu)),$$

and $x \in \text{Range}(V(\mu))$ for all $x \in \mathcal{X}$. This ensures that $a_x^\top \hat{\theta}_\mu$ is an unbiased estimator of $a_x^\top \theta^*$.

Furthermore $V(\mu) \succcurlyeq mV(\pi)$, so the variance $a_x^\top V(\mu)^+ a_x$ of $a_x^\top \hat{\theta}_\mu$ is upper-bounded by $a_x^\top V(\mu)^+ a_x \leq m^{-1} a_x^\top V(\pi)^+ a_x$. Now, the General Equivalence Theorem of Kiefer and Pukelshein shows that $\max_{x \in \mathcal{X}} a_x^\top V(\pi)^+ a_x \leq d+1$. Thus, $a_x^\top V(\mu)^+ a_x \leq m^{-1}(d+1)$.

We now prove Lemma 2. Recall that $\pi \in \mathcal{M}_{e_{d+1}}^\mathcal{X}$ is such that $e_{d+1} \in \text{Range } V(\pi)$, and that μ is defined as $\mu(x) = \lceil m\pi(x) \rceil$ for all $x \in \mathcal{X}$. Using similar arguments, we can show that $e_{d+1} \in \text{Range}(V(\mu))$, which ensures that $e_{d+1}^\top \hat{\theta}_\mu$ is an unbiased estimator of $e_{d+1}^\top \theta^*$. The second part of the Lemma follows directly using that $V(\mu) \succcurlyeq mV(\pi)$.

C.6.2 Proof of Lemma 3

Setting $\mu \cdot \Delta = (\mu(x)\Delta_x)_{x \in \mathcal{X}}$ and

$$V_\Delta(\lambda) = \sum_{x \in \mathcal{X}} \lambda_x \begin{pmatrix} \Delta_x^{-1/2} x \\ \Delta_x^{-1/2} z_x \end{pmatrix} \begin{pmatrix} \Delta_x^{-1/2} x \\ \Delta_x^{-1/2} z_x \end{pmatrix}^\top,$$

we observe that $V_\Delta(\mu \cdot \Delta) = V(\mu)$. Hence,

$$\kappa(\Delta) = \min_{\substack{\mu \in \mathcal{M}^+ \\ e_{d+1}^\top V_\Delta(\mu \cdot \Delta)^+ e_{d+1} \leq 1}} \sum_{x \in \mathcal{X}} (\mu \cdot \Delta)_x.$$

We observe that $e_{d+1} \in \text{Range}(V(\mu))$ is equivalent to $e_{d+1} \in \text{Range}(V_\Delta(\mu \cdot \Delta))$. Hence, $\mu^\Delta \cdot \Delta = \lambda^\Delta$ where

$$\lambda^\Delta \in \underset{\substack{\lambda \in \mathbb{R}_+^\mathcal{X} \\ e_{d+1} \in \text{Range}(V_\Delta(\lambda)) \\ e_{d+1}^\top V_\Delta(\lambda)^+ e_{d+1} \leq 1}}{\text{argmin}} \sum_{x \in \mathcal{X}} \lambda_x.$$

The conclusion then follows by noticing that by homogeneity, $\lambda^\Delta = \kappa^\Delta \pi^\Delta$.

C.6.3 Proof of Lemma 4

Proof of Claim i) The proof of the first claim is immediate by definition of κ . Indeed, let $\widetilde{\mathcal{M}} = \left\{ \mu \in \mathcal{M}_{e_{d+1}}^\mathcal{X} : e_{d+1}^\top V(\mu)^+ e_{d+1} \leq 1 \right\}$ be the set of measures μ admissible for estimating ω^* with a precision level 1. Then,

$$\kappa(c\Delta) = \min_{\mu \in \widetilde{\mathcal{M}}} \sum_x \mu(x) c\Delta_x = c \min_{\mu \in \widetilde{\mathcal{M}}} \sum_x \mu(x) \Delta_x = c\kappa(\Delta).$$

Proof of Claim ii) The proof of the second claim is also straightforward. If $\Delta \leq \Delta'$, then for all $\mu \in \widetilde{\mathcal{M}}$, $\sum_x \mu(x) \Delta_x \leq \sum_x \mu(x) \Delta'_x$. Recall that $\mu^{\Delta'} = \text{argmin}_{\mu \in \widetilde{\mathcal{M}}} \sum_x \mu(x) \Delta'_x$. Then,

$$\kappa(\Delta') = \sum_x \mu^{\Delta'}(x) \Delta'_x \geq \sum_x \mu^{\Delta'}(x) \Delta_x \geq \min_{\mu \in \widetilde{\mathcal{M}}} \sum_x \mu(x) \Delta_x = \kappa(\Delta).$$

Proof of Claim iii) To prove the third claim, note that

$$\begin{aligned} \kappa(\Delta \vee \Delta') &= \min_{\mu \in \widetilde{\mathcal{M}}} \sum_x \mu(x) (\Delta_x \vee \Delta'_x) \\ &\geq \min_{\mu \in \widetilde{\mathcal{M}}} \left(\sum_x \mu(x) \Delta_x \vee \sum_x \mu(x) \Delta'_x \right) \\ &\geq \left(\min_{\mu \in \widetilde{\mathcal{M}}} \sum_x \mu(x) \Delta_x \right) \vee \left(\min_{\mu \in \widetilde{\mathcal{M}}} \sum_x \mu(x) \Delta'_x \right) \\ &\geq \kappa(\Delta) \vee \kappa(\Delta'). \end{aligned}$$

Proof of Claim iv) Recall that

$$\kappa(\Delta) = \min_{\mu \in \tilde{\mathcal{M}}} \sum_x \mu(x) \Delta_x.$$

Let us define a sequence $(\mu_n)_{n \in \mathbb{N}} \in \tilde{\mathcal{M}}^{\mathbb{N}}$ such that $\sum_x \mu_n(x) \Delta_x \xrightarrow{n \rightarrow \infty} \kappa(\Delta)$, and let us denote $\kappa_n = \sum_x \mu_n(x) \Delta_x$. According to Claim ii), we have

$$\kappa(\Delta) \leq \kappa(\Delta \vee \epsilon) = \min_{\mu \in \tilde{\mathcal{M}}} \sum_x \mu(x) (\Delta_x \vee \epsilon) \leq \sum_x \mu_n(x) \Delta_x + \epsilon \sum_x \mu_n(x).$$

It follows that for all n ,

$$\kappa(\Delta) \leq \liminf_{\epsilon \rightarrow 0^+} \kappa(\Delta \vee \epsilon) \leq \limsup_{\epsilon \rightarrow 0^+} \kappa(\Delta \vee \epsilon) \leq \kappa_n.$$

Letting n go to infinity, we get that $\lim_{\epsilon \rightarrow 0^+} \kappa(\Delta \vee \epsilon) = \kappa(\Delta)$.

C.6.4 Proof of Lemma 5

Let $\tilde{\mu}$ be defined by (8). Let us define

$$\begin{aligned} (x^{(1)}, x^{(2)}, \tilde{\gamma}) &\in \operatorname{argmax}_{(x, x') \in \mathcal{X}, \gamma \in \mathcal{C}(\mathcal{X})} (x - x')^\top \gamma \\ \bar{x} &= \frac{1}{2}(x^{(1)} + x^{(2)}) \\ \tilde{n} &= \sum_{x \in \mathcal{X}} \tilde{\mu}(x) \\ \text{and } \tilde{x} &= \frac{1}{\tilde{n}} \sum_{x \in \mathcal{X}} \tilde{\mu}(x) x. \end{aligned}$$

Recall that v can equivalently be defined as the budget necessary to estimate the bias with a variance smaller than 1. Therefore, we have

$$\tilde{n} \geq v. \tag{54}$$

On the one hand, the definition of Δ_{\max} implies that $(x^{(1)} - x^{(2)})^\top \tilde{\gamma} = \Delta_{\max}$. On the other hand, by definition of $\tilde{\kappa}$ and $\tilde{\mu}$,

$$\begin{aligned} \tilde{\kappa} &\geq \sum_{x \in \mathcal{X}} \tilde{\mu}(x) (x^{(1)} - x)^\top \tilde{\gamma} \\ &= \tilde{n} (x^{(1)} - \tilde{x})^\top \tilde{\gamma}. \end{aligned}$$

Using Equation (54), we find that

$$\begin{aligned} \frac{\tilde{\kappa}}{v} &\geq (x^{(1)} - \bar{x})^\top \tilde{\gamma} + (\bar{x} - \tilde{x})^\top \tilde{\gamma} \\ &= \frac{\Delta_{\max}}{2} + (\bar{x} - \tilde{x})^\top \tilde{\gamma}. \end{aligned} \tag{55}$$

Now, since $\tilde{\gamma} \in \mathcal{C}(\mathcal{X})$, we also have $-\tilde{\gamma} \in \mathcal{C}(\mathcal{X})$, and therefore

$$\begin{aligned} \tilde{\kappa} &\geq \sum_{x \in \mathcal{X}} \tilde{\mu}(x) (x^{(2)} - x)^\top (-\tilde{\gamma}) \\ &= \tilde{n} (\tilde{x} - x^{(2)})^\top \tilde{\gamma} \end{aligned}$$

Using again Equation (54), we find that

$$\begin{aligned} \frac{\tilde{\kappa}}{v} &\geq (\tilde{x} - \bar{x})^\top \tilde{\gamma} + (\bar{x} - x^{(2)})^\top \tilde{\gamma} \\ &= (\tilde{x} - \bar{x})^\top \tilde{\gamma} + \frac{\Delta_{\max}}{2}. \end{aligned} \tag{56}$$

Combining Equations (55) and (56), we find that

$$\frac{\tilde{\kappa}}{v} \geq \frac{\Delta_{\max}}{2} + |(\bar{x} - \tilde{x})^\top \tilde{\gamma}|.$$

This implies in particular that $\tilde{\kappa} \geq \frac{\Delta_{\max} v}{2} = \frac{\kappa_*}{2}$, which concludes the proof of the lemma.

C.6.5 Proof of Lemmas 6 and 8

We prove Lemma 6. The proof of Lemma 8 follows by noticing that the two actions sets are equal up to a permutation of the direction of some basis vectors. To prove Lemma 8, we rely on Elfving's characterization of c -optimal design, given in Theorem 1. Theorem 1 shows that for $\pi \in \mathcal{P}^{\{1, \dots, d+1\}}$ to be e_{d+1} -optimal, there must exist $t > 0$ and $\zeta \in \{-1, +1\}^{d+1}$ such that

$$\begin{aligned} \sum_{1 \leq i \leq d+1} \pi_i &= 1 \\ 0 &= \pi_1 \zeta_1 - \left(1 - \frac{2}{\sqrt{v} + 1}\right) \pi_{d+1} \zeta_{d+1} \\ \forall i \in \{2, \dots, d\}, 0 &= \pi_i \zeta_i \\ t &= \sum_{1 \leq i \leq \lfloor d/2 \rfloor} \pi_i \zeta_i - \sum_{\lfloor d/2 \rfloor + 1 \leq i \leq d+1} \pi_i \zeta_i. \end{aligned}$$

Solving this system, we find that $t^{-2} = v$. Note that the unicity of the solution for the corresponding probability measure π guarantees that te_{d+1} belongs to the boundary of \mathcal{S} .

C.6.6 Proof of Lemma 7

For a given parameter γ^* , let us denote by Δ_i the gap corresponding to the action i . To compute $\kappa(\Delta)$, we could want to rely on Lemma 3 to find the Δ -optimal design, corresponding to the e_{d+1} -optimal design on the rescaled features $\Delta_x^{-1/2} \begin{pmatrix} x \\ z_x \end{pmatrix}$. Theorem 1 indeed allows us to compute such a design, as seen in the proof of Lemma 6. Unfortunately, we cannot rescale the features using the true gaps, since $\Delta_{x^*} = 0$. To circumvent this problem, we rely on the following reasoning :

1. We use Lemma 3 and Theorem 1 to compute the design $\mu^{\Delta \vee \epsilon}$ for $\epsilon \in (0, \Delta_{\min})$; and the corresponding regret $\kappa(\Delta \vee \epsilon)$;
2. We find the value of $\kappa(\Delta)$ by noticing that $\epsilon \mapsto \kappa(\Delta \vee \epsilon)$ is continuous at 0.

For $\epsilon \in (0, \Delta_{\min})$, define $\bar{\Delta} = \Delta \vee \epsilon$, and $\bar{x} = \bar{\Delta}_x^{-1/2} x$. Let $\bar{\pi}$ denote the e_{d+1} -optimal design for the rescaled features \bar{x} , and let \bar{v} denote its variance. Then, Lemma 3 ensures that $\kappa(\bar{\Delta}) = \bar{v}$.

Now, Theorem 1 shows that there exists $\zeta \in \{-1, +1\}^{d+1}$ such that

$$\begin{aligned} \sum_{1 \leq i \leq d+1} \bar{\pi}_i &= 1 \\ 0 &= \bar{\pi}_1 \zeta_1 \bar{\Delta}_1^{-1/2} - \left(1 - \frac{2}{\sqrt{\bar{v}} + 1}\right) \bar{\pi}_{d+1} \zeta_{d+1} \bar{\Delta}_{d+1}^{-1/2} \\ \forall i \in \{2, \dots, d\}, 0 &= \bar{\pi}_i \zeta_i \bar{\Delta}_i^{-1/2} \\ \bar{v}^{-1/2} &= \sum_{1 \leq i \leq \lfloor d/2 \rfloor} \bar{\pi}_i \zeta_i \bar{\Delta}_i^{-1/2} - \sum_{\lfloor d/2 \rfloor + 1 \leq i \leq d+1} \bar{\pi}_i \zeta_i \bar{\Delta}_i^{-1/2} \end{aligned}$$

and $\bar{v}^{-1/2} e_{d+1}$ belongs to the boundary of \mathcal{S} . Solving this system, we find that

$$\kappa(\bar{\Delta})^{-1/2} = \bar{v}^{-1/2} = \frac{\left(\frac{2}{\sqrt{\bar{v}} + 1}\right) \bar{\Delta}_{d+1}^{-1/2}}{1 + \left(1 - \frac{2}{\sqrt{\bar{v}} + 1}\right) \bar{\Delta}_{d+1}^{-1/2} \bar{\Delta}_1^{-1/2}}.$$

As in Lemma 6, the unicity of the solution for the corresponding probability measure $\bar{\pi}$ guarantees that $\bar{v}^{-1/2}e_{d+1}$ belongs to the boundary of the Elfving's set. Now, $\epsilon \leq \Delta_{\min}$, so

$$\kappa(\bar{\Delta})^{-1/2} = \kappa(\Delta \vee \epsilon)^{-1/2} = \frac{\left(\frac{2}{\sqrt{v}+1}\right) \Delta_{d+1}^{-1/2}}{1 + \left(1 - \frac{2}{\sqrt{v}+1}\right) \Delta_{d+1}^{-1/2} \epsilon^{1/2}}.$$

The fourth claim of Lemma 4 ensures that $\kappa(\Delta \vee \epsilon) \xrightarrow{\epsilon \rightarrow 0} \kappa(\Delta)$. Therefore,

$$\kappa(\Delta) = \lim_{\epsilon \rightarrow 0} \left(\frac{\left(\frac{2}{\sqrt{v}+1}\right) \Delta_{d+1}^{-1/2}}{1 + \left(1 - \frac{2}{\sqrt{v}+1}\right) \Delta_{d+1}^{-1/2} \epsilon^{1/2}} \right)^{-2} = \frac{(\sqrt{v}+1)^2 \Delta_{d+1}}{4}.$$

C.6.7 Proof of Lemma 9

Lemma 9 follows directly from Lemmas 14 and 15.

Lemma 14.

$$\mathbb{P} \left(\exists l \geq 1, z \in \{-1, 1\} \text{ such that } \text{Explore}_l^{(z)} = \text{True}, \text{ and } x \in \mathcal{X}_l^{(z)} \text{ such that } \left| \begin{pmatrix} \hat{\gamma}_l^{(z)} - \gamma^* \\ \hat{\omega}_l^{(z)} - \omega^* \end{pmatrix}^\top \begin{pmatrix} x \\ z_x \end{pmatrix} \right| \geq \epsilon_l \right) \leq \delta.$$

Lemma 15.

$$\mathbb{P} \left(\exists l \geq 1 \text{ such that } \text{Explore}_l^{(0)} = \text{True} \text{ and } \left| \hat{\omega}_l^{(0)} - \omega^* \right| \geq \epsilon_l \right) \leq \delta.$$

C.6.8 Proof of Lemma 10

To prove Lemma 10, we rely on the following key lemma. This lemma proves that on $\bar{\mathcal{F}}$, i.e. when the error bounds hold, the algorithm never eliminates the best action or the best group.

Lemma 16. *On the event $\bar{\mathcal{F}}$, for all $x^* \in \arg\max_{x \in \mathcal{X}} x^\top \gamma^*$ and all l such that $\text{Explore}_l^{(z_{x^*})} = \text{True}$, $x^* \in \mathcal{X}_{l+1}^{(z_{x^*})}$. Moreover, on the event $\bar{\mathcal{F}}$, for all l such that $\text{Explore}_l^{(0)} = \text{True}$, there exists $x^* \in \arg\max_{x \in \mathcal{X}} x^\top \gamma^*$ such that $\hat{z}_{l+1}^* \neq -z_{x^*}$.*

Let $l \geq 1$ be such that $\text{Explore}_l^{(z_{x^*})} = \text{True}$. Then, on $\bar{\mathcal{F}}$, $x^* \in \mathcal{X}_{l+1}^{(z_{x^*})}$ by Lemma 16. Moreover, for all $x \in \mathcal{X}_{l+1}^{(z_{x^*})}$, by definition of $\mathcal{X}_{l+1}^{(z_{x^*})}$, we have that on $\bar{\mathcal{F}}$

$$\left(\begin{pmatrix} x^* \\ z_{x^*} \end{pmatrix} - \begin{pmatrix} x \\ z_x \end{pmatrix} \right)^\top \begin{pmatrix} \hat{\gamma}_l^{(z)} \\ \hat{\omega}_l^{(z)} \end{pmatrix} \leq 3\epsilon_l.$$

which implies that

$$\left(\begin{pmatrix} x^* \\ z_{x^*} \end{pmatrix} - \begin{pmatrix} x \\ z_x \end{pmatrix} \right)^\top \begin{pmatrix} \gamma^* \\ \omega^* \end{pmatrix} \leq 3\epsilon_l + \left| \begin{pmatrix} x^* \\ z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \hat{\gamma}_l^{(z)} - \gamma^* \\ \hat{\omega}_l^{(z)} - \omega^* \end{pmatrix} \right| + \left| \begin{pmatrix} x \\ z_x \end{pmatrix}^\top \begin{pmatrix} \hat{\gamma}_l^{(z)} - \gamma^* \\ \hat{\omega}_l^{(z)} - \omega^* \end{pmatrix} \right|.$$

Thus, on the event $\bar{\mathcal{F}}$, for all $x \in \mathcal{X}_{l+1}^{(z_{x^*})}$

$$(x^* - x)^\top \gamma^* < 5\epsilon_l,$$

which proves Equation (17). To prove the second claim of Lemma 10, assume that for all $x' \in \arg\max_{x \in \mathcal{X}} x^\top \gamma^*$, $z_{x'} = z_{x^*}$ (when this does not hold, the second claim follows from Equation (17)). Now, let $l \geq 1$ be such that $\text{Explore}_l^{(-z_{x^*})} = \text{True}$. By Lemma 16, on $\bar{\mathcal{F}}$, $x^* \in \mathcal{X}_l^{(z_{x^*})}$ and $\hat{z}_l^* = 0$. Then, the algorithm is unable to

determine the group containing the best set during the phase $\text{Exp}_{l-1}^{(0)}$, so there must exist $x' \in \mathcal{X}_l^{(-z_{x^*})}$ such that

$$\begin{pmatrix} x^* \\ z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \hat{\gamma}_{l-1}^{(z_{x^*})} \\ \hat{\omega}_{l-1}^{(z_{x^*})} \end{pmatrix} \leq \begin{pmatrix} x' \\ -z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \hat{\gamma}_{l-1}^{(-z_{x^*})} \\ \hat{\omega}_{l-1}^{(-z_{x^*})} \end{pmatrix} + 2z_{x^*}\hat{\omega}_{l-1}^{(0)} + 4\epsilon_{l-1}.$$

It follows that

$$\begin{pmatrix} x^* - x' \\ 2z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \gamma^* \\ \omega^* \end{pmatrix} \leq \begin{pmatrix} x^* \\ z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \gamma^* - \hat{\gamma}_{l-1}^{(z_{x^*})} \\ \omega^* - \hat{\omega}_{l-1}^{(z_{x^*})} \end{pmatrix} + \begin{pmatrix} x' \\ -z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \hat{\gamma}_{l-1}^{(-z_{x^*})} - \gamma^* \\ \hat{\omega}_{l-1}^{(-z_{x^*})} - \omega^* \end{pmatrix} + 2z_{x^*}\hat{\omega}_{l-1}^{(0)} + 4\epsilon_{l-1}.$$

On $\overline{\mathcal{F}}$, this implies that

$$\begin{pmatrix} x^* - x' \\ 2z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \gamma^* \\ \omega^* \end{pmatrix} < 2z_{x^*}\hat{\omega}_{l-1}^{(0)} + 6\epsilon_{l-1}$$

so

$$(x^* - x')^\top \gamma^* \leq 2z_{x^*} \left(\hat{\omega}_{l-1}^{(0)} - \omega^* \right) + 6\epsilon_{l-1} < 8\epsilon_{l-1} = 16\epsilon_l. \quad (57)$$

Moreover, for all $x \in \mathcal{X}_{l+1}^{(-z_{x^*})}$ we have $(a_{x'} - a_x)^\top \hat{\theta}_l^{(-z_{x^*})} \leq 3\epsilon_l$, so following the same lines as for the first claim, we get $(x' - x)^\top \gamma^* < 5\epsilon_l$. Combining this bound with (57), we get

$$\max_{x \in \mathcal{X}_{l+1}^{(-z_{x^*})}} (x^* - x)^\top \gamma^* < 21\epsilon_l.$$

This concludes the proof of Lemma 10.

C.6.9 Proof of Lemma 11

For $z \in \{-1, +1\}$ and $l > 0$,

$$\sum_x \mu_l^{(z)}(x) \leq \sum_x \frac{2(d+1)\pi_l^{(z)}(x)}{\epsilon_l^2} \log \left(\frac{kl(l+1)}{\delta} \right) + |\text{supp}(\pi_l^{(z)})|.$$

Now, $\text{supp}(\pi_l^{(z)}) \leq \frac{(d+1)(d+2)}{2}$ and $\sum_x \pi_l^{(z)}(x) = 1$, so

$$\sum_x \mu_l^{(z)}(x) \leq \frac{2(d+1)}{\epsilon_l^2} \log \left(\frac{kl(l+1)}{\delta} \right) + \frac{(d+1)(d+2)}{2}$$

which proves the first claim of Lemma 11.

To prove the second claim, we bound the regret for bias estimation at stage l as follows. On $\overline{\mathcal{F}}$, we have $\Delta_x \leq \hat{\Delta}_x^l$ for all $x \in \mathcal{X}$ and $l \geq 1$, so

$$\sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \Delta_x \leq \sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \hat{\Delta}_x^l.$$

Recall that $\hat{\mu}_l$ is the $\hat{\Delta}^l$ -optimal design, and that for all $x \in \mathcal{X}$, $\mu_l^{(0)}(x) = \lceil \frac{2\hat{\mu}_l(x)}{\epsilon_l^2} \log \left(\frac{l(l+1)}{\delta} \right) \rceil$. Since $\hat{\Delta}_x^l \leq \Delta_{\max} \leq 2$ for all $x \in \mathcal{X}$, we have

$$\sum_{x \in \mathcal{X}} \mu_l^{(0)}(x) \hat{\Delta}_x^l \leq \sum_{x \in \mathcal{X}} \frac{2\hat{\mu}_l(x)}{\epsilon_l^2} \log \left(\frac{l(l+1)}{\delta} \right) \hat{\Delta}_x^l + 2|\text{supp}(\mu_l^{(0)})|$$

and $|\text{supp}(\mu_l^{(0)})| \leq d+1$, so

$$\sum_x \mu_l^{(0)}(x) \Delta_x \leq \frac{2}{\epsilon_l^2} \log \left(\frac{l(l+1)}{\delta} \right) \sum_{x \in \mathcal{X}} \hat{\mu}_l(x) \hat{\Delta}_x^l + 2(d+1).$$

By definition of $\hat{\mu}_l(x)$, we have that

$$\sum_{x \in \mathcal{X}} \hat{\mu}_l(x) \hat{\Delta}_x^l = \kappa(\hat{\Delta}^l).$$

It follows that, on $\overline{\mathcal{F}}$,

$$\sum_x \mu_l^{(0)}(x) \Delta_x \leq \sum_x \mu_l^{(0)}(x) \hat{\Delta}_x^l \leq \frac{2}{\epsilon_l^2} \log \left(\frac{l(l+1)}{\delta} \right) \kappa(\hat{\Delta}^l) + 2(d+1).$$

C.6.10 Proof of Lemma 12

For the first claim, we rely on the next lemma.

Lemma 17. *Let us set $\ell_x = \max \{l \geq 1 : x \in \mathcal{X}_l^{(-1)} \cup \mathcal{X}_l^{(1)}\}$. On $\overline{\mathcal{F}}$, we have for any $l \geq 1$*

1. $\hat{\Delta}_x^l \leq \Delta_x + 16\epsilon_l$ for all $x \in \mathcal{X}_l^{(-1)} \cup \mathcal{X}_l^{(1)}$ (i.e. for all x such that $l \leq \ell_x$);
2. if $\Delta_x \geq 21\epsilon_l$ then $\ell_x \leq l$;
3. $\epsilon_{\ell_x} < \Delta_x$ for all $x \in \mathcal{X}$.

Lemma 12 relies on the following remarks : if Δ, Δ' are such that $\Delta_x \leq \Delta'_x$ for all $x \in \mathcal{X}$, then by Lemma 4 (ii), $\kappa(\Delta) \leq \kappa(\Delta')$. Let us now prove that for all $l \geq 1$ and all $x \in \mathcal{X}$, $\hat{\Delta}_x^l \leq 513(\Delta \vee \epsilon_l)$.

Case $\epsilon_l \geq \Delta_x$. On $\overline{\mathcal{F}}$, we have $l \leq \ell_x - 1$ according to the third claim of Lemma 17. So, on $\overline{\mathcal{F}}$,

$$\hat{\Delta}_x^l \leq \Delta_x + 16\epsilon_l \leq 17(\Delta_x \vee \epsilon_l).$$

Case $\epsilon_l < \Delta_x$. Then, on $\overline{\mathcal{F}}$, we have $32\epsilon_{l+5} < \Delta_x$ and so $l+5 \geq \ell_x$ according to the second claim of Lemma 17. Hence, on $\overline{\mathcal{F}}$, according to Lemma 17, we have

$$\begin{aligned} \hat{\Delta}_x^l &\leq \max_{k=0, \dots, 5} \hat{\Delta}_x^{\ell_x - k} \leq \Delta_x + 16\epsilon_{\ell_x - 5} \\ &\leq \Delta_x + 512\epsilon_{\ell_x} \leq 513\Delta_x. \end{aligned}$$

Thus, for all $l \geq 1$ and all $x \in \mathcal{X}$,

$$\hat{\Delta}_x^l \leq 513(\Delta \vee \epsilon_l).$$

Now, let $\widetilde{\mathcal{M}} = \left\{ \mu \in \mathcal{M}_{e_{d+1}}^{\mathcal{X}} : e_{d+1}^\top V(\mu)^+ e_{d+1} \geq 1 \right\}$ the measures μ admissible for estimating ω^* with a precision level 1. Note that for all $a, b, c > 0$,

$$(1 + ab^{-1})(c \vee b) = (c + cab^{-1}) \vee (a + b) \geq c \vee (a + b) \geq c \vee a. \quad (58)$$

Using Equation (58) with $a = \Delta_x$, $b = \tau$ and $c = \epsilon$, we see that

$$\kappa(\Delta \vee \epsilon) = \min_{\mu \in \widetilde{\mathcal{M}}} \sum_x \mu(x) (\Delta_x \vee \epsilon) \leq (1 + \epsilon/\tau) \min_{\mu \in \widetilde{\mathcal{M}}} \sum_x \mu(x) (\Delta_x \vee \tau) = (1 + \epsilon/\tau) \kappa(\Delta \vee \tau).$$

Using Lemma 4 together with $\hat{\Delta}_x^l \leq 513(\Delta \vee \epsilon_l)$, we find that

$$\kappa(\hat{\Delta}_x^l) \leq 513\kappa(\Delta \vee \epsilon_l) \leq 513(1 + \epsilon_l/\tau) \kappa(\Delta \vee \tau).$$

This proves the first claim of Lemma 12.

To prove the second claim, we use Lemma 4 and the fact that for all x , $\hat{\Delta}_x^l \geq \epsilon_l$. Moreover, on $\overline{\mathcal{F}}$, $\hat{\Delta}_x^l \geq \Delta_x$ for all $x \in \mathcal{X}$. Then, $\kappa(\hat{\Delta}) \geq \kappa(\epsilon_l \vee \Delta)$ by Lemma 4 (iii).

C.6.11 Proof of Lemmas 13

To prove Lemma 13, let us consider l such that $\epsilon_l \leq \frac{\Delta_{\neq}}{8}$. According to Lemma 16, on $\overline{\mathcal{F}}$ we know that $\hat{z}_l^* \neq -z_{x^*}$. When $\hat{z}_l^* = z_{x^*}$, then we also have $\hat{z}_{l+1}^* = z_{x^*}$ and the conclusion follows immediately. Let us consider now the case where $\hat{z}_l^* = 0$. By definition of Δ_{\neq} , for all $x' \in \mathcal{X}_{l+1}^{(-z_{x^*})}$,

$$(x^* - x')^\top \gamma^* \geq \Delta_{\neq}.$$

This implies that

$$\begin{aligned} \begin{pmatrix} x^* \\ z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \hat{\gamma}_l^{(z_{x^*})} \\ \hat{\omega}_l^{(z_{x^*})} \end{pmatrix} - z_{x^*} \hat{\omega}_l^{(0)} &\geq \max_{x \in \mathcal{X}_{l+1}^{(-z_{x^*})}} \begin{pmatrix} x \\ -z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \hat{\gamma}_l^{(-z_{x^*})} \\ \hat{\omega}_l^{(-z_{x^*})} \end{pmatrix} + z_{x^*} \hat{\omega}_l^{(0)} \\ &+ \begin{pmatrix} x^* \\ z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \hat{\gamma}_l^{(z_{x^*})} - \gamma^* \\ \hat{\omega}_l^{(z_{x^*})} - \omega^* \end{pmatrix} + \min_{x \in \mathcal{X}_{l+1}^{(-z_{x^*})}} \begin{pmatrix} x \\ -z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \gamma^* - \hat{\gamma}_l^{(-z_{x^*})} \\ \omega^* - \hat{\omega}_l^{(-z_{x^*})} \end{pmatrix} \\ &+ \Delta_{\neq} + 2z_{x^*} (\omega^* - \hat{\omega}_l^{(0)}). \end{aligned}$$

On $\overline{\mathcal{F}}$, it follows that

$$\begin{pmatrix} x^* \\ z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \hat{\gamma}_l^{(z_{x^*})} \\ \hat{\omega}_l^{(z_{x^*})} \end{pmatrix} - z_{x^*} \hat{\omega}_l^{(0)} - 2\epsilon_l \geq \max_{x \in \mathcal{X}_{l+1}^{(-z_{x^*})}} \begin{pmatrix} x \\ -z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \hat{\gamma}_l^{(-z_{x^*})} \\ \hat{\omega}_l^{(-z_{x^*})} \end{pmatrix} + z_{x^*} \hat{\omega}_l^{(0)} - 6\epsilon_l + \Delta_{\neq}.$$

When $\Delta_{\neq} \geq 8\epsilon_l$, this implies that $\widehat{z_{l+1}^*} = z_{x^*}$.

C.6.12 Proof of Lemma 14

Recall that $\xi_t = y_t - x_t^\top \gamma^* - z_{x_t} \omega^*$. For $l \geq 0$ and $z \in \{-1, +1\}$, when $\text{Explore}_l^{(z)} = \text{True}$, the least square estimator $\begin{pmatrix} \hat{\gamma}_l^{(z)} \\ \hat{\omega}_l^{(z)} \end{pmatrix}$ is given by

$$\begin{aligned} \begin{pmatrix} \hat{\gamma}_l^{(z)} \\ \hat{\omega}_l^{(z)} \end{pmatrix} &= \left(V_l^{(z)}\right)^+ \sum_{t \in \text{Exp}_l^{(z)}} \left(\begin{pmatrix} x_t \\ z_{x_t} \end{pmatrix}^\top \begin{pmatrix} \gamma^* \\ \omega^* \end{pmatrix} + \xi_t \right) \begin{pmatrix} x_t \\ z_{x_t} \end{pmatrix} \\ &= \left(V_l^{(z)}\right)^+ \left(V_l^{(z)}\right) \begin{pmatrix} \gamma^* \\ \omega^* \end{pmatrix} + \left(V_l^{(z)}\right)^+ \sum_{t \in \text{Exp}_l^{(z)}} \xi_t \begin{pmatrix} x_t \\ z_{x_t} \end{pmatrix}, \end{aligned}$$

where $\left(V_l^{(z)}\right)^+$ is a generalized inverse of $V_l^{(z)}$. Since $V_l^{(z)} \left(V_l^{(z)}\right)^+ V_l^{(z)} = V_l^{(z)}$, multiplying the left and right hand side of the last equation by $V_l^{(z)}$, we find that

$$V_l^{(z)} \begin{pmatrix} \hat{\gamma}_l^{(z)} - \gamma^* \\ \hat{\omega}_l^{(z)} - \omega^* \end{pmatrix} = V_l^{(z)} \left(V_l^{(z)}\right)^+ \sum_{t \in \text{Exp}_l^{(z)}} \xi_t \begin{pmatrix} x_t \\ z_{x_t} \end{pmatrix}. \quad (59)$$

By Lemma 1, for all $x \in \mathcal{X}_l^{(z)}$, $\begin{pmatrix} x \\ z_x \end{pmatrix} \in \text{Range} \left(V_l^{(z)}\right)$, so

$$V_l^{(z)} \left(V_l^{(z)}\right)^+ \begin{pmatrix} x \\ z_x \end{pmatrix} = \begin{pmatrix} x \\ z_x \end{pmatrix}. \quad (60)$$

Then,

$$\begin{aligned}
\begin{pmatrix} \hat{\gamma}_l^{(z)} - \gamma^* \\ \hat{\omega}_l^{(z)} - \omega^* \end{pmatrix}^\top \begin{pmatrix} x \\ z_x \end{pmatrix} &= \begin{pmatrix} \hat{\gamma}_l^{(z)} - \gamma^* \\ \hat{\omega}_l^{(z)} - \omega^* \end{pmatrix}^\top V_l^{(z)} (V_l^{(z)})^+ \begin{pmatrix} x \\ z_x \end{pmatrix} \\
&= \sum_{t \in \text{Exp}_l^{(z)}} \begin{pmatrix} x_t \\ z_{x_t} \end{pmatrix}^\top (V_l^{(z)})^+ V_l^{(z)} (V_l^{(z)})^+ \begin{pmatrix} x \\ z_x \end{pmatrix} \xi_t \\
&= \sum_{t \in \text{Exp}_l^{(z)}} \begin{pmatrix} x_t \\ z_{x_t} \end{pmatrix}^\top (V_l^{(z)})^+ \begin{pmatrix} x \\ z_x \end{pmatrix} \xi_t,
\end{aligned}$$

where the first and third lines follow from Equation (60), and the second line follows from Equation (59). By definition of our algorithm, conditionally on $\mathcal{X}_l^{(z)}$ and $\text{Explore}_l^{(z)} = \text{True}$, the variables $(\xi_t)_{t \in \text{Exp}_l^{(z)}}$ are independent centered normal gaussian variables. Then,

$$\mathbb{P}_{|\mathcal{X}_l^{(z)}, \text{Explore}_l^{(z)} = \text{True}} \left(\left| \begin{pmatrix} \hat{\gamma}_l^{(z)} - \gamma^* \\ \hat{\omega}_l^{(z)} - \omega^* \end{pmatrix}^\top \begin{pmatrix} x \\ z_x \end{pmatrix} \right| \geq \sqrt{2 \sum_{t \in \text{Exp}_l^{(z)}} \left(\begin{pmatrix} x_t \\ z_{x_t} \end{pmatrix}^\top (V_l^{(z)})^+ \begin{pmatrix} x \\ z_x \end{pmatrix} \right)^2 \log \left(\frac{kl(l+1)}{\delta} \right)} \right) \leq \frac{\delta}{kl(l+1)}.$$

Expanding $\left(\begin{pmatrix} x_t \\ z_{x_t} \end{pmatrix}^\top (V_l^{(z)})^+ \begin{pmatrix} x \\ z_x \end{pmatrix} \right)^2 = \begin{pmatrix} x \\ z_x \end{pmatrix}^\top (V_l^{(z)})^+ \begin{pmatrix} x_t \\ z_{x_t} \end{pmatrix} \begin{pmatrix} x_t \\ z_{x_t} \end{pmatrix}^\top (V_l^{(z)})^+ \begin{pmatrix} x \\ z_x \end{pmatrix}$, and using the definition of $V_l^{(z)}$, we find that

$$\mathbb{P}_{|\mathcal{X}_l^{(z)}, \text{Explore}_l^{(z)} = \text{True}} \left(\left| \begin{pmatrix} \hat{\gamma}_l^{(z)} - \gamma^* \\ \hat{\omega}_l^{(z)} - \omega^* \end{pmatrix}^\top \begin{pmatrix} x \\ z_x \end{pmatrix} \right| \geq \sqrt{2 \begin{pmatrix} x \\ z_x \end{pmatrix}^\top (V_l^{(z)})^+ V_l^{(z)} (V_l^{(z)})^+ \begin{pmatrix} x \\ z_x \end{pmatrix} \log \left(\frac{kl(l+1)}{\delta} \right)} \right) \leq \frac{\delta}{kl(l+1)}$$

which in turn implies (using Equation (60))

$$\mathbb{P}_{|\mathcal{X}_l^{(z)}, \text{Explore}_l^{(z)} = \text{True}} \left(\left| \begin{pmatrix} \hat{\gamma}_l^{(z)} - \gamma^* \\ \hat{\omega}_l^{(z)} - \omega^* \end{pmatrix}^\top \begin{pmatrix} x \\ z_x \end{pmatrix} \right| \geq \sqrt{2 \left\| \begin{pmatrix} x \\ z_x \end{pmatrix} \right\|_{(V_l^{(z)})^+}^2 \log \left(\frac{kl(l+1)}{\delta} \right)} \right) \leq \frac{\delta}{kl(l+1)}$$

Now, using Lemma 1 and the definition of μ_l^z , we see that for all $x \in \mathcal{X}_l^{(z)}$,

$$\begin{pmatrix} x \\ z_x \end{pmatrix}^\top (V_l^{(z)})^+ \begin{pmatrix} x \\ z_x \end{pmatrix} \leq \frac{\epsilon_l^2}{2 \log(kl(l+1)/\delta)}.$$

Finally, for all $x \in \mathcal{X}_l^{(z)}$,

$$\begin{aligned}
&\mathbb{P}_{|\mathcal{X}_l^{(z)}, \text{Explore}_l^{(z)} = \text{True}} \left(\left| \begin{pmatrix} \hat{\gamma}_l^{(z)} - \gamma^* \\ \hat{\omega}_l^{(z)} - \omega^* \end{pmatrix}^\top \begin{pmatrix} x \\ z_x \end{pmatrix} \right| \geq \epsilon_l \right) \\
&\leq \mathbb{P}_{|\mathcal{X}_l^{(z)}, \text{Explore}_l^{(z)} = \text{True}} \left(\left| \begin{pmatrix} \hat{\gamma}_l^{(z)} - \gamma^* \\ \hat{\omega}_l^{(z)} - \omega^* \end{pmatrix}^\top \begin{pmatrix} x \\ z_x \end{pmatrix} \right| \geq \sqrt{2 \left\| \begin{pmatrix} x \\ z_x \end{pmatrix} \right\|_{(V_l^{(z)})^+}^2 \log \left(\frac{kl(l+1)}{\delta} \right)} \right) \leq \frac{\delta}{kl(l+1)}.
\end{aligned}$$

Integrating out the conditioning on the value of $\mathcal{X}_l^{(z)}$ and $\text{Explore}_l^{(z)}$ and using a union bound yields the desired result.

C.6.13 Proof of Lemma 15

The proof is similar to that of Lemma 14. If $\text{Explore}_l^{(0)} = \text{True}$, then $\hat{\omega}_l$ is defined as

$$\hat{\omega}_l^{(0)} = e_{d+1}^\top \left(V_l^{(0)} \right)^+ \sum_{t \in \text{Exp}_l^{(0)}} \left(\begin{pmatrix} x_t \\ z_{x_t} \end{pmatrix}^\top \begin{pmatrix} \gamma^* \\ \omega^* \end{pmatrix} + \xi_t \right) \begin{pmatrix} x_t \\ z_{x_t} \end{pmatrix}.$$

Since $\begin{pmatrix} x \\ z_x \end{pmatrix}_{x \in \mathcal{X}}$ spans \mathbb{R}^{d+1} , μ is finite and $e_{d+1} \in \text{Range}(V(\hat{\mu}_l))$. Then, according to Lemma 2, for every round l , we have $e_{d+1} \in \text{Range}(V_l^{(0)})$, so $V_l^{(0)} \left(V_l^{(0)} \right)^+ e_{d+1} = e_{d+1}$. This implies that

$$\hat{\omega}_l^{(0)} - \omega^* = \sum_{t \in \text{Exp}_l^{(0)}} e_{d+1}^\top \left(V_l^{(0)} \right)^+ \begin{pmatrix} x_t \\ z_{x_t} \end{pmatrix} \xi_t.$$

By definition of our algorithm, conditionally on $\text{Explore}_l^{(0)} = \text{True}$, the variables $(\xi_t)_{t \in \text{Exp}_l^{(0)}}$ are independent centered normal gaussian variables. Then,

$$\mathbb{P}_{|\text{Explore}_l^{(0)} = \text{True}} \left(\left| \hat{\omega}_l^{(0)} - \omega^* \right| \geq \sqrt{2 \sum_{t \in \text{Exp}_l^{(z)}} \left(e_{d+1}^\top \left(V_l^{(0)} \right)^+ \begin{pmatrix} x_t \\ z_{x_t} \end{pmatrix} \right)^2 \log \left(\frac{l(l+1)}{\delta} \right)} \right) \leq \frac{\delta}{l(l+1)}.$$

Using again $V_l^{(0)} \left(V_l^{(0)} \right)^+ e_{d+1} = e_{d+1}$ and the definition of $V_l^{(0)}$, we find that

$$\mathbb{P}_{|\text{Explore}_l^{(0)} = \text{True}} \left(\left| \hat{\omega}_l^{(0)} - \omega^* \right| \geq \sqrt{2 e_{d+1}^\top \left(V_l^{(0)} \right)^+ e_{d+1} \log \left(\frac{l(l+1)}{\delta} \right)} \right) \leq \frac{\delta}{l(l+1)}. \quad (61)$$

Now, Lemma 2 and the definition of $\mu_l^{(0)}$ imply that

$$e_{d+1}^\top \left(V_l^{(0)} \right)^+ e_{d+1} \leq \frac{\epsilon_l^2}{2 \log(l(l+1)/\delta)}.$$

Finally, Equation (61) implies that

$$\mathbb{P}_{|\text{Explore}_l^{(0)} = \text{True}} \left(\left| \hat{\omega}_l^{(0)} - \omega^* \right| \geq \epsilon_l \right) \leq \frac{\delta}{l(l+1)}.$$

Using a union bound over the phases $\text{Exp}_l^{(0)}$ yields the result.

C.6.14 Proof of Lemma 16

To prove Lemma 16, we begin by showing that it is enough to prove that for $l \geq 1$,

$$\begin{aligned} \mathcal{F}_l \supset & \left\{ \exists x^* \in \underset{x \in \mathcal{X}}{\text{argmax}} x^\top \gamma^* : \text{Explore}_l^{(z_{x^*})} = \text{True and } x^* \notin \mathcal{X}_{l+1}^{(z_{x^*})} \right\} \\ \cup & \left\{ \bigcap_{l' \leq l} \overline{\left\{ \exists x^* \in \underset{x \in \mathcal{X}}{\text{argmax}} x^\top \gamma^* : \text{Explore}_{l'}^{(z_{x^*})} = \text{True and } x^* \notin \mathcal{X}_{l'+1}^{(z_{x^*})} \right\}} \right. \\ & \left. \bigcap \left\{ \text{Explore}_l^{(0)} = \text{True and } \forall x^* \in \underset{x \in \mathcal{X}}{\text{argmax}} x^\top \gamma^*, \hat{z}_{l+1}^* = -z_{x^*} \right\} \right\}. \end{aligned} \quad (62)$$

Indeed, denoting $\mathcal{F}_l^{(1)} = \left\{ \exists x^* \in \underset{x \in \mathcal{X}}{\text{argmax}} x^\top \gamma^* : \text{Explore}_l^{(z_{x^*})} = \text{True and } x^* \notin \mathcal{X}_{l+1}^{(z_{x^*})} \right\}$ and $\mathcal{F}_l^{(2)} = \left\{ \text{Explore}_l^{(0)} = \text{True and } \forall x^* \in \underset{x \in \mathcal{X}}{\text{argmax}} x^\top \gamma^*, \hat{z}_{l+1}^* = -z_{x^*} \right\}$, we see that Equation (62) would

then be rewritten as

$$\mathcal{F}_l \supset \mathcal{F}_l^{(1)} \cup \left\{ \bigcap_{l' \leq l} \overline{\mathcal{F}_{l'}^{(1)}} \cap \mathcal{F}_l^{(2)} \right\}$$

which implies

$$\bigcup_{l \geq 1} \mathcal{F}_l \supset \bigcup_{l \geq 1} \left\{ \mathcal{F}_l^{(1)} \cup \left\{ \bigcap_{l' \leq l} \overline{\mathcal{F}_{l'}^{(1)}} \cap \mathcal{F}_l^{(2)} \right\} \cup \mathcal{F}_{l'}^{(1)} \right\} \supset \bigcup_{l \geq 1} \left\{ \mathcal{F}_l^{(1)} \cup \mathcal{F}_l^{(2)} \right\}.$$

Then, Equation (62) would imply that

$$\overline{\mathcal{F}} = \overline{\bigcup_{l \geq 1} \mathcal{F}_l} \subset \overline{\bigcup_{l \geq 1} \left\{ \mathcal{F}_l^{(1)} \cup \mathcal{F}_l^{(2)} \right\}} = \bigcap_{l \geq 1} \left\{ \overline{\mathcal{F}_l^{(1)}} \cap \overline{\mathcal{F}_l^{(2)}} \right\},$$

thus proving Lemma 16. To prove Equation (62), we show that both $\mathcal{F}_l^{(1)}$ and $\bigcap_{l' \leq l} \overline{\mathcal{F}_{l'}^{(1)}} \cap \mathcal{F}_l^{(2)}$ imply \mathcal{F}_l .

If $\mathcal{F}_l^{(1)}$ is true: then $\exists x^* \in \arg\max_{x \in \mathcal{X}} : \text{Explore}_l^{(z_{x^*})} = \text{True}$ and $x^* \notin \mathcal{X}_{l+1}^{(z_{x^*})}$.

Without loss of generality, assume that $l > 1$ is the smallest integer such that $\text{Explore}_l^{(z_{x^*})} = \text{True}$ and $x^* \notin \mathcal{X}_{l+1}^{(z_{x^*})}$. Then, necessarily $x^* \in \mathcal{X}_l^{(z_{x^*})}$ (because either $l = 1$, or $\text{Explore}_{l-1}^{(z_{x^*})} = \text{True}$). Now, because $x^* \in \mathcal{X}_l^{(z_{x^*})} \setminus \mathcal{X}_{l+1}^{(z_{x^*})}$, there exists $x \in \mathcal{X}_l^{(z_{x^*})}$ such that

$$(x - x^*)^\top \hat{\gamma}_l^{(z_{x^*})} \geq 3\epsilon_l$$

and in particular

$$x^\top \hat{\gamma}_l^{(z_{x^*})} - \epsilon_l > (x^*)^\top \hat{\gamma}_l^{(z_{x^*})} + \epsilon_l.$$

Recall that by definition of x^* , $(\gamma^*)^\top (x^* - x) \geq 0$. This in turn implies that

$$\begin{pmatrix} x \\ z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \hat{\gamma}_l^{(z_{x^*})} - \gamma^* \\ \hat{\omega}_l^{(z_{x^*})} - \omega^* \end{pmatrix} - \epsilon_l > \begin{pmatrix} x^* \\ z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \hat{\gamma}_l^{(z_{x^*})} - \gamma^* \\ \hat{\omega}_l^{(z_{x^*})} - \omega^* \end{pmatrix} + \epsilon_l.$$

The last equation implies that either $\begin{pmatrix} x \\ z_x \end{pmatrix}^\top \begin{pmatrix} \hat{\gamma}_l^{(z)} - \gamma^* \\ \hat{\omega}_l^{(z)} - \omega^* \end{pmatrix} > \epsilon_l$ or $\begin{pmatrix} x^* \\ z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \hat{\gamma}_l^{(z)} - \gamma^* \\ \hat{\omega}_l^{(z)} - \omega^* \end{pmatrix} < -\epsilon_l$, which in turn implies \mathcal{F}_l .

If $\bigcap_{l' \leq l} \overline{\mathcal{F}_{l'}^{(1)}} \cap \mathcal{F}_l^{(2)}$ is true: then $\text{Explore}_l^{(0)} = \text{True}$ and $\forall x^* \in \arg\max_{x \in \mathcal{X}} x^\top \gamma^*, \widehat{z}_{l+1}^* = -z_{x^*}$. Moreover, for all $l' \leq l$, $\text{Explore}_{l'}^{(z_{x^*})} = \text{False}$ or $x^* \in \mathcal{X}_{l'+1}^{(z_{x^*})}$.

Note that this case can only hold if all optimal actions x^* belong to the same group z_{x^*} . Without loss of generality, assume that $l > 1$ is the smallest integer such that $\text{Explore}_l^{(0)} = \text{True}$ and $\widehat{z}_{l+1}^* = -z_{x^*}$, and for all $l' \leq l$, $\text{Explore}_{l'}^{(z_{x^*})} = \text{False}$ or $x^* \in \mathcal{X}_{l'+1}^{(z_{x^*})}$. Note that because $\text{Explore}_l^{(0)} = \text{True}$, necessarily $\text{Explore}_{l'}^{(0)} = \text{True}$ for all $l' \leq l$, and in particular $x^* \in \mathcal{X}_{l+1}^{(z_{x^*})}$.

Then, there exists $x \in \mathcal{X}_{l+1}^{(-z_{x^*})}$ such that

$$\begin{pmatrix} x \\ -z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \hat{\gamma}_l^{(-z_{x^*})} \\ \hat{\omega}_l^{(-z_{x^*})} \end{pmatrix} - \begin{pmatrix} x^* \\ z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \hat{\gamma}_l^{(z_{x^*})} \\ \hat{\omega}_l^{(z_{x^*})} \end{pmatrix} + 2z_{x^*} \hat{\omega}_l^{(0)} \geq 4\epsilon_l.$$

Recall that all optimal actions x^* are in the same group z_{x^*} , so $(\gamma^*)^\top (x^* - x) > 0$. This in turn implies that

$$\begin{pmatrix} x \\ -z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \hat{\gamma}_l^{(-z_{x^*})} - \gamma^* \\ \hat{\omega}_l^{(-z_{x^*})} - \omega^* \end{pmatrix} - \begin{pmatrix} x^* \\ z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \hat{\gamma}_l^{(z_{x^*})} - \gamma^* \\ \hat{\omega}_l^{(z_{x^*})} - \omega^* \end{pmatrix} + 2z_{x^*} (\hat{\omega}_l^{(0)} - \omega^*) \geq 4\epsilon_l.$$

The last equation implies that either $\begin{pmatrix} x \\ -z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \hat{\gamma}_l^{(-z_{x^*})} - \gamma^* \\ \hat{\omega}_l^{(-z_{x^*})} - \omega^* \end{pmatrix} \geq \epsilon_l$, or $\begin{pmatrix} x^* \\ z_{x^*} \end{pmatrix}^\top \begin{pmatrix} \hat{\gamma}_l^{(z_{x^*})} - \gamma^* \\ \hat{\omega}_l^{(z_{x^*})} - \omega^* \end{pmatrix} \leq -\epsilon_l$, or $z_{x^*} (\hat{\omega}_l^{(0)} - \omega^*) \geq \epsilon_l$, which in turn implies \mathcal{F}_l .

C.6.15 Proof of Lemma 17

The first claim holds for $l = 1$. For $l \geq 1$, for any $x \in \mathcal{X}_{l+1}^{(-1)} \cup \mathcal{X}_{l+1}^{(1)}$, we have $\widehat{\Delta}_x^{l+1} \leq \Delta_x + 8\epsilon_l$ on $\overline{\mathcal{F}}$ according to the definition of $\widehat{\Delta}^{l+1}$ and \mathcal{F} . The first claim then follows.

For the second claim, Lemma 10 gives that, on $\overline{\mathcal{F}}$, $\Delta_x < 21\epsilon_l$ for any $x \in \mathcal{X}_{l+1}^{(-1)} \cup \mathcal{X}_{l+1}^{(1)}$. So $\Delta_x \geq 21\epsilon_l$ implies $x \notin \mathcal{X}_{l+1}^{(-1)} \cup \mathcal{X}_{l+1}^{(1)}$ and hence $l \geq \ell_x$ on $\overline{\mathcal{F}}$.

For the third claim, we notice that

$$\max_{x' \in \mathcal{X}_{\ell_x}^{(z_x)}} (a_{x'} - a_x)^\top \widehat{\theta}_{\ell_x}^{(z_x)} > 3\epsilon_{\ell_x},$$

since $x \notin \mathcal{X}_{\ell_x+1}$. Since the left-hand side is smaller than $\Delta_x + 2\epsilon_{\ell_x}$ on $\overline{\mathcal{F}}$, we get $\Delta_x > \epsilon_{\ell_x}$.