
Instance-Dependent Regret Bounds for Nonstochastic Linear Partial Monitoring

Federico Di Gennaro^{*,†}
EPFL, Lausanne, Switzerland
digennarof.00@gmail.com

Khaled Eldowa^{*,‡}
Università degli Studi di Milano, Milan, Italy
& Politecnico di Milano, Milan, Italy
khaled.eldowa3@gmail.com

Nicolò Cesa-Bianchi
Università degli Studi di Milano, Milan, Italy
& Politecnico di Milano, Milan, Italy
nicolo.cesa-bianchi@unimi.it

Abstract

In contrast to the classic formulation of partial monitoring, linear partial monitoring can model infinite outcome spaces, while imposing a linear structure on both the losses and the observations. This setting can be viewed as a generalization of linear bandits where loss and feedback are decoupled in a flexible manner. In this work, we address a nonstochastic (adversarial), finite-actions version of the problem through a simple instance of the exploration-by-optimization method that is amenable to efficient implementation. We derive regret bounds that depend on the game structure in a more transparent manner than previous theoretical guarantees for this paradigm. Our bounds feature instance-specific quantities that reflect the degree of alignment between observations and losses, and resemble known guarantees in the stochastic setting. Notably, they achieve the standard $\sqrt{T}$ rate in easy (locally observable) games and $T^{2/3}$ in hard (globally observable) games, where T is the time horizon. We instantiate these bounds in a selection of old and new partial information settings subsumed by this model, and illustrate that the achieved dependence on the game structure can be tight in interesting cases.

1 Introduction

Partial monitoring models sequential decision-making problems with general feedback structures. A generic nonstochastic partial monitoring problem is characterized by an action space $\mathcal{A}$, an outcome space $\mathcal{Z}$, an observation space Σ , a loss function $Y: \mathcal{A} \times \mathcal{Z} \rightarrow \mathbb{R}$, and an observation function $\Phi: \mathcal{A} \times \mathcal{Z} \rightarrow \Sigma$. All of these components are known to the learner, who interacts with an unknown environment in a series of T rounds. In each round t , the learner picks an action $A_t \in \mathcal{A}$, observes $\Phi(A_t, Z_t)$, and incurs the (unobserved) loss $Y(A_t, Z_t)$, where $Z_t \in \mathcal{Z}$ is a latent outcome. The learner aims to minimize the regret, defined as the difference between the cumulative losses of played actions and the cumulative losses incurred by the best action in hindsight. This general decoupling between losses and observations results in an extraordinary modeling power. Not only does this framework encompass full-information scenarios [23, 13], where the latent outcome can be inferred from the observation, and bandit scenarios [4, 11], where the loss and the observation coincide;

^{*}Equal contribution.

[†]Now at ETH Zürich, Zürich, Switzerland.

[‡]Now at Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK; Grenoble, France.

it can model problems that interpolate between these two extremes, like feedback graphs [2], and problems that go beyond bandit feedback, like dynamic pricing [14]. The primary object of study in this framework is how the structure of the game affects its difficulty, measured via the notion of minimax regret (the best achievable worst-case regret).

When $\mathcal{A}$ and $\mathcal{Z}$ are both finite, one recovers the well-studied *finite* partial monitoring problem. An extensive line of work over the last few decades [46, 14, 3, 20, 6, 36, 37] has shown that all finite games can be classified into one of four classes, each characterized by the achievable minimax rate in terms of T , which can only be $\Omega(T)$, $\Theta(T^{2/3})$, $\Theta(T^{1/2})$, or 0. These classes, respectively, consist of: hopeless games, hard—or (strictly) globally observable—games, easy—or locally observable—games, and trivial games. What delineates these categories are precise ‘observability’ conditions on the loss and observation functions. An important work in our context is that of Lattimore and Szepesvári [38], who introduced the *exploration-by-optimization* (EXO) method (see Section 3 for a detailed description) and achieved—via an intuitive and efficient algorithm—the best known rates for the problem, free from arbitrary and redundant game-dependent constants that appear in prior results.

Nevertheless, the restriction to finite outcome spaces remains a limitation of this model. On the other hand, the *linear* partial monitoring framework [40, 15, 29, 30] allows infinite outcomes, but restricts that, for a given outcome, the loss and observation functions are linear in some representation of the actions. This framework still encompasses a broad spectrum of problems (see Section 2), and arguably, provides a more intuitive observation structure, free from the combinatorial complications of the finite model (which can nonetheless be made to fit within the linear framework [30]). A simpler point of view sees this model as a generalization of the standard linear bandit problem [17, 10], which corresponds to the case when the loss and observation functions coincide. Kirschner et al. [29, 30] studied an instance of this model in the stochastic setting (where the hidden outcome is fixed over the rounds and the learner receives noisy observations), adopting a version of the information-directed sampling [47] algorithm. Remarkably, they showed that the classification theorem continues to hold in the linear setting with analogous observability criteria. Also of note is that their bounds feature certain *alignment* constants that reflect the difficulty of the game within its respective class in an arguably more interpretable manner compared to finite partial monitoring results.

In the nonstochastic (adversarial) setting, a linear partial monitoring problem was studied by Lattimore and Gyorgy [35], who also adopted the EXO approach. In fact, their setup is more general than what we described above since their observation function is taken as arbitrary. Nevertheless, they established an elegant result bounding the regret of a general EXO policy in terms of a generalized version of the information ratio [47], which characterizes the Bayesian regret of information-directed sampling. As such, however, their bounds are still of a largely opaque nature in terms of their dependence on problem-specific parameters. Moreover, in its general form, implementing their EXO algorithm involves solving a daunting optimization problem (over a possibly infinite dimensional space) at every step in order to compute a loss estimation function and a sampling distribution. Hence, though constructive, their approach serves more as a template than a concrete implementable policy. An expanded discussion on related works is provided in Appendix A.

1.1 Contributions

We address a finite-actions version of adversarial linear partial monitoring through a simple instance of the EXO method. Outlined in Section 3, this policy builds upon an exponential weights update and imposes a simple structure on the adopted loss estimator, which is made possible thanks to the linearity of the observation function. As a result, the aforementioned optimization problem is reduced to one of minimizing a convex objective over probability distributions over the actions, subject to a certain convex constraint. This problem is efficiently solvable against many loss (outcome) spaces of interest. Under the observability conditions identified by Kirschner et al. [29] in the stochastic setting, we prove new regret bounds for this policy in Sections 4 and 5 for locally and globally observable games respectively; the former of order $\sqrt{T}$ and the latter of order $T^{2/3}$, as one would expect.

The proofs rely on bounding the optimal value of the optimization program alluded to above via an application of the minimax theorem, following the analysis of EXO by Lattimore and Szepesvári [38] in the finite setting. However, to better exploit the linear structure of our setting, we employ new arguments, including a more natural way of constructing exploration distributions in locally observable games. The resulting bounds feature interpretable instance-dependent quantities, similar to the alignment constants of Kirschner et al. [29, 30], that gauge the difficulty of the game within its

observability class. We then illustrate the versatility of the bounds by instantiating them in a selection of both locally and globally observable games, proving their near-optimality in some cases.

2 Preliminaries

Notation. For a positive integer d , $[d]$ denotes the set $\{1, \dots, d\}$. We use the Iverson bracket notation $\llbracket \cdot \rrbracket$ to denote the indicator function. For a set $S \subseteq \mathbb{R}^d$, $\text{co}(S)$ denotes its convex hull, and $\text{ext}(S)$ the set of its extreme points. Moreover, its (absolute) polar set is defined as $S^\circ := \{x \in \mathbb{R}^d \mid \sup_{s \in S} |\langle x, s \rangle| \leq 1\}$. For a real matrix X , $\text{tr}(X)$, $\text{rk}(X)$, $\text{col}(X)$, $X^\dagger$, and $\|X\|$ denote, respectively, its trace, rank, column space, Moore–Penrose inverse, and operator norm (induced by the Euclidean L_2 -norm). For $p \in [0, \infty]$, let $\mathcal{B}_p(r)$ be the centered L_p ball of radius $r > 0$. We use $\mathbf{1}_d$ and I_d to denote, respectively, the d -dimensional vector with all entries equal to 1 and the $d \times d$ identity matrix. The probability simplex in $\mathbb{R}^d$ is denoted by $\Delta_d := \{x \in [0, 1]^d \mid \|x\|_1 = 1\}$. Given a finite index set Z and a collection of matrices $\{X_z \mid z \in Z\}$, all with the same number of rows, we denote by $\text{span}\{X_z \mid z \in Z\}$ the span of all their columns.

2.1 Problem setting

Let $\mathcal{A} := [k]$ denote the action set, where k is a positive integer. We associate every action $a \in \mathcal{A}$ with a feature vector $\psi_a \in \mathbb{R}^d$ and an observation matrix $M_a \in \mathbb{R}^{d \times n(a)}$, where $n(\cdot)$ maps every action to a positive integer. We will use $\mathcal{X}$ to denote the set $\{\psi_a \mid a \in \mathcal{A}\}$. Let $\mathcal{L} \subset \mathbb{R}^d$ be a convex and compact loss space, which, in our setting, is synonymous with the outcome space $\mathcal{Z}$ as used in the introduction. We assume here that $\mathcal{L}$ contains a centered L_2 ball $\mathcal{B}_2(r)$ of some radius $r > 0$. This simplifies the observation structure of the game and suffices to subsume typical cases such as when $\mathcal{L}$ represents the absolute polar set of $\mathcal{X}$ or a ball of bounded norm. We study the following game between a learner and an oblivious adversary over T rounds. Before the game starts, the adversary secretly chooses a sequence of loss vectors $(\ell_t)_{t \in [T]}$, where $\ell_t \in \mathcal{L}$. Then, at every round $t \in [T]$, the learner: i) selects (possibly at random) an action $A_t \in \mathcal{A}$; ii) observes the signal $\phi_t := M_{A_t}^\top \ell_t \in \mathbb{R}^{n(A_t)}$, which, in general, is distinct from the incurred (but unobserved) loss $\psi_{A_t}^\top \ell_t \in \mathbb{R}$. The feature mapping and the observation matrices are known to the learner. For brevity, let $y_t(a) := \psi_a^\top \ell_t$ for every $a \in \mathcal{A}$. In the notation used in the introduction, one can see that for $a \in \mathcal{A}$ and $\ell \in \mathcal{L}$, the loss map in the setting becomes $Y(a, \ell) = \psi_a^\top \ell$, while the observation map becomes $\Phi(a, \ell) = M_a^\top \ell$.

For round $t \in [T]$, let $\mathcal{H}_t := (A_s, \phi_s)_{s=1}^t$ denote the interaction history up to the end of the round, and let $\mathcal{F}_t := \sigma(\mathcal{H}_t)$ denote the σ -algebra generated by $\mathcal{H}_t$. The learner’s policy can be represented via a sequence $(\pi_t)_{t \in [T]}$ of probability kernels such that π_t maps the history $\mathcal{H}_{t-1}$ to a distribution over the actions, from which A_t is to be sampled. The learner’s goal is to minimize the following standard notion of regret

$$R_T := \mathbb{E} \left[\sum_{t=1}^T y_t(A_t) \right] - \min_{a \in \mathcal{A}} \sum_{t=1}^T y_t(a), \quad (1)$$

where the expectation is with respect to the internal randomness of the algorithm. The dependence of the regret on the learner’s strategy $(\pi_t)_t$ and the sequence of losses $(\ell_t)_t$ is suppressed from the notation for brevity. The minimax regret is defined as $R_T^* := \inf_{(\pi_t)_t} \sup_{(\ell_t)_t} R_T$, where the inf is over all (possibly randomized) learning algorithms and the sup is over all sequences of T loss vectors from $\mathcal{L}$. Similarly to [29, 30], we impose a standard boundedness assumption throughout this work:

Assumption 1. *It holds that $\max_{a, b \in \mathcal{A}, \ell \in \mathcal{L}} |(\psi_a - \psi_b)^\top \ell| \leq 2$.*

As a notational remark, we define $\mathbb{E}_t[\cdot] := \mathbb{E}[\cdot \mid \mathcal{F}_{t-1}]$, with $\mathcal{F}_0$ being the trivial σ -algebra. Moreover, throughout the paper, we treat functions over $\mathcal{A}$ as vectors in $\mathbb{R}^k$. In particular, for $b \in \mathcal{A} (= [k])$, the indicator vector $e_b \in \mathbb{R}^k$ corresponds to the function $a \mapsto \llbracket a = b \rrbracket$ over $\mathcal{A}$.

2.2 Observability conditions and lower bounds

Similarly to the finite partial monitoring setting, the observation structure of the game characterizes the achievable regret. Before we proceed, we introduce some relevant concepts. An action $a \in \mathcal{A}$

is called *Pareto optimal* if $\psi_a \in \text{ext}(\text{co}(\mathcal{X}))$; i.e., ψ_a is an extreme point of the convex hull of $\mathcal{X}$. Let $\mathcal{A}^* \subseteq \mathcal{A}$ satisfy that $\text{ext}(\text{co}(\mathcal{X})) = \{\psi_a \mid a \in \mathcal{A}^*\}$ and contain no two actions with the same features. Additionally, let $k^* := |\mathcal{A}^*|$. This set is well-defined since $\text{ext}(\text{co}(\mathcal{X})) \subseteq \mathcal{X}$; moreover, it is the smallest subset of $\mathcal{A}$ that satisfies $\min_{a \in \mathcal{A}} \langle \psi_a, \ell \rangle = \min_{a^* \in \mathcal{A}^*} \langle \psi_{a^*}, \ell \rangle$ for every $\ell \in \mathbb{R}^d$. The set of actions that are optimal against a loss vector $\ell \in \mathbb{R}^d$ is denoted by $\mathcal{P}(\ell) := \{a \in \mathcal{A} \mid \langle \psi_a - \psi_b, \ell \rangle \leq 0 \text{ for all } b \in \mathcal{A}\}$, which extends to sets $C \subseteq \mathbb{R}^d$ as $\mathcal{P}(C) := \bigcup_{\ell \in C} \mathcal{P}(\ell)$. We state below the observability conditions that determine the difficulty of the game (in terms of the scaling of the regret with T) as identified by Kirschner et al. [29] in the stochastic setting.

Definition 1 (Observability conditions [29]). *A game is called globally observable if*

$$\psi_a - \psi_b \in \text{span}\{M_c \mid c \in \mathcal{A}\} \quad \text{for all } a, b \in \mathcal{A}. \quad (2)$$

Further, a game is called locally observable if for every convex set $C \subseteq \mathbb{R}^d$,

$$\psi_a - \psi_b \in \text{span}\{M_c \mid c \in \mathcal{P}(C)\} \quad \text{for all } a, b \in \mathcal{P}(C). \quad (3)$$

Note that the second condition implies the first by taking C as $\mathbb{R}^d$ (or simply $\{0\}$). Intuitively, in a globally observable game, one can use the observations to estimate the difference between the losses of any pair of actions. In locally observable games, one can do so efficiently; only using observations from actions that perform at least as good as the pair of interest. (See also the equivalent formulation provided in Lemma 4 below, or Lemma 5 in Appendix B.) Kirschner et al. [30] refined these conditions to account for cases when $\mathcal{L}$ is an arbitrary (and possibly low-dimensional) set. We do not pursue this direction here; as mentioned in the problem setting, we focus on the case when $\mathcal{L}$ has a non-empty interior (in particular, we assumed it contains a centered L_2 ball), rendering the conditions above sufficient to derive the following lower bounds.

Proposition 1. *Assume that the game has at least two non-duplicate Pareto optimal actions, and that $\mathcal{B}_2(r) \subseteq \mathcal{L}$ for some $r > 0$. Then, the minimax regret R_T^* is $\Omega(\sqrt{T})$ if the game is locally observable, $\tilde{\Omega}(T^{2/3})$ if it is globally but not locally observable, and $\Omega(T)$ otherwise.*

These bounds are obtained in the same manner as similar results in the finite outcomes setting and the linear stochastic setting [29], see Appendix J. However, one must deal with some technical nuisances since in our case, the stochastic setting is not immediately subsumed by the adversarial one; besides boundedness issues, the manner in which noise is added to the observations can make the relationship a bit subtle, refer to [33, Chapter 29] for details. Over the coming section, we will derive regret upper bounds with matching rates and interpretable game-dependent constants for an EXO-based policy.

2.3 Examples

Owing to its flexibility, this setting can model many online learning problems, classic and contrived, as we now elaborate. In all these examples, $\mathcal{L}$ is typically taken as $\mathcal{X}^\circ$, which is equivalent to $\mathcal{B}_\infty(1)$ when $\psi_a = e_a$ for all $a \in \mathcal{A}$.

Full information. In this game, also known as prediction with expert advice [13], we have that $d = k$, $\psi_a = e_a$, and $M_a = \mathbf{I}_k$ for all $a \in \mathcal{A}$. Hence, the learner observes the entire loss vector upon playing any action, and the game is clearly locally observable.

Feedback graphs. An instance of this game [42, 1, 2] is characterized by a graph $\mathcal{G} = (\mathcal{A}, E)$ over the actions (we focus on undirected graphs). As in the previous game, $d = k$ and $\psi_a = e_a$; however, M_a only has the vectors $(e_b)_{b \in N_{\mathcal{G}}(a)}$ as columns, where $N_{\mathcal{G}}(a)$ is the neighborhood of a (which need not include a itself). Globally observable games correspond to *weakly* observable graphs, where the loss of every action is observable; while locally observable games correspond to *strongly* observable graphs, where each action either has a self-loop, is connected to all other actions, or both.

Linear bandits. Adversarial linear bandits [17, 10] (with a finite action set) is recovered by simply fixing $M_a = \psi_a$ for all $a \in \mathcal{A}$. This is easily verifiable to be a locally observable game, which reduces to a standard bandit problem when $\psi_a = e_a$.

Linear dueling bandits. This game is a quantitative variation of the dueling bandits problem. It was proposed in [29, 30] in the stochastic linear partial monitoring setting, extending the finite-outcomes formulation in [24]. Here, the action set has the form $\mathcal{A} = [m] \times [m]$ for some positive integer m . Given a feature mapping $a \mapsto \psi_a$ from $[m]$ to $\mathbb{R}^d$, we extend it to $\mathcal{A}$ via $(a, b) \mapsto \psi_{a,b} := \psi_a + \psi_b$.

However, the observations are given by $M_{a,b} = \psi_a - \psi_b$; hence, the learner receives relative feedback. This game can also be shown to be locally observable [29, 30].

Bandits with ill-conditioned observers. Modifying a problem from [18], consider a game where $d = k$, $\psi_a = e_a$, and $M_a = (1 - \varepsilon)\mathbf{1}_k/k + \varepsilon e_a$ for some $\varepsilon \in (0, 1]$, such that $\varepsilon = 1$ recovers standard bandits. Note that the observations become, in a sense, less informative as ε decreases. Nevertheless, the game is always locally observable as $e_a - e_b = 1/\varepsilon(M_a - M_b)$.

Composite graph feedback. We describe now another variant of the graph feedback problem, still characterized by an undirected graph $G = (\mathcal{A}, E)$, which we now assume to contain all self-loops. Here, again, $d = k$ and $\psi_a = e_a$; however, upon playing an action, the learner observes the *average* of the losses of its neighbors (which includes the played action itself). That is, if A is the adjacency matrix of the graph and D is the degree matrix, then M_a is the a -th column of AD^{-1} . This is less informative than standard graph feedback, and draws inspiration from problems studied in [55, 25]. Global observability here corresponds to AD^{-1} being invertible, whereas the only locally observable graph is the one corresponding to standard bandits, see Appendix I for more details.

3 Exploration-by-Optimization, the Linear Case

We describe in this section a simple instantiation of the EXO approach [38, 35] suited to our setting, resulting in a policy amenable to efficient implementation. We show in later sections that this policy suffices to obtain tight bounds, even if its specificity excludes it from enjoying the generic guarantees provided by Lattimore and Gyorgy [35] and Lattimore [34] for general EXO. The predictions made by this policy revolve around the predictions of an instance of the exponential weights algorithm, restricted over the set $\mathcal{A}^*$ of Pareto optimal actions (without duplicates). When executed against a sequence of loss functions $(\hat{y}_t)_{t \in [T]}$ with $\hat{y}_t: \mathcal{A} \rightarrow \mathbb{R}$, the exponential weights algorithm outputs at round t a distribution $q_t \in \Delta_k$ given by

$$q_t(a) = \frac{\mathbb{I}[a \in \mathcal{A}^*] \exp(-\eta \sum_{s=1}^{t-1} \hat{y}_s(a))}{\sum_{a' \in \mathcal{A}^*} \exp(-\eta \sum_{s=1}^{t-1} \hat{y}_s(a'))}, \quad (4)$$

where $\eta > 0$ is a learning rate parameter. This would have been a sound approach on its own were we able to access the loss functions $(y_t)_t$ directly. Instead, we are to estimate these loss functions, which, in general, requires carefully exploring the actions. Incurring minimal loss while doing so, to the extent allowed by the structure of the game, is the essential goal of the EXO method. Towards making this tradeoff more explicit, the following lemma provides a generic regret bound for any policy by pivoting around the predictions $(q_t)_t$ produced against some sequence $(\hat{y}_t)_t$ of surrogate loss functions satisfying certain conditions. (Proofs of this section's results are in Appendix D.)

Lemma 1. *Fix a learning policy and define p_t as the law of A_t conditioned on $\mathcal{F}_{t-1}$. Let q_t be as given in (4) for some learning rate $\eta > 0$ and sequence of surrogate loss functions $(\hat{y}_t)_t$ such that $\hat{y}_t \in \mathbb{R}^k$ is $\mathcal{F}_t$ -measurable and satisfies $\max_{a \in \mathcal{A}} |\eta \hat{y}_t(a)| \leq 1$. Further, let $a^* \in \arg \min_{a \in \mathcal{A}^*} \sum_{t=1}^T y_t(a)$. Then, the regret of the policy satisfies*

$$R_T \leq \frac{\log k}{\eta} + \sum_{t=1}^T \mathbb{E} [\langle p_t - q_t, y_t \rangle + \eta \langle q_t, \mathbb{E}_t \hat{y}_t^2 \rangle + \langle q_t - e_{a^*}, y_t - \mathbb{E}_t \hat{y}_t \rangle].$$

While not particularly insightful, results of this form nonetheless motivate the EXO approach, where one picks the predictions and the surrogate loss functions by solving an optimization problem at every step, seeking to minimize an upper bound on the second term in the regret guarantee of Lemma 1. More precisely, at every round t , reliant on the observed history thus far, one decides on a prediction $p_t \in \Delta_k$ and a function $g_t: \mathcal{A} \times \mathbb{R} \rightarrow \mathbb{R}^k$ so as to set $\hat{y}_t = g_t(A_t, \phi_t)$, assuming $\|g_t(A_t, \phi_t)\|_\infty \leq 1/\eta$. To make the dependence on ℓ_t more explicit, we denote by H the $d \times k$ matrix with $\{\psi_a\}_{a \in \mathcal{A}}$ as columns. In particular, this enables us to write $y_t = H^\top \ell_t$. Considering the t -th summand in the second term in the bound of Lemma 1, the term inside the expectation becomes

$$\langle p_t - q_t, H^\top \ell_t \rangle + \eta \langle q_t, \sum_a p_t(a) g_t(a, M_a^\top \ell_t)^2 \rangle + \langle q_t - e_{a^*}, H^\top \ell_t - \sum_a p_t(a) g_t(a, M_a^\top \ell_t) \rangle, \quad (5)$$

where $g_t(a, M_a^\top \ell_t)^2$ is the entry-wise square of $g_t(a, M_a^\top \ell_t)$. The general strategy is to optimize jointly in p_t and g_t the expression above, maximized over the adversary's choice of ℓ_t . Besides some

small deviations, this still aligns closely with the approach of Lattimore and Gyorgy [35].¹ Now, however, exploiting the structure of our setting, we sidestep the demanding task of optimizing over all possible loss estimators by committing to a simple structure commonly used in (adversarial) linear bandit algorithms [10, 33], modified slightly to suit our partial monitoring setup.

Before proceeding in that direction, we introduce some notation. Let $\mathbf{M}$ denote the $d \times (\sum_a n(a))$ matrix obtained by stacking horizontally the observation matrices of all the actions; such that $\text{span}\{M_a \mid a \in \mathcal{A}\} = \text{col}(\mathbf{M})$. For any $\pi \in \Delta_k$, we define $Q(\pi) := \sum_{a \in \mathcal{A}} \pi(a) M_a M_a^\top$. Moreover, with $\delta \in (0, 1/2]$, let $Q_\delta(\pi) := Q((1 - \delta)\pi + \delta \mathbf{1}_k/k)$. This forces the argument of Q to be bounded away from zero, implying that $\text{col}(Q_\delta(\pi)) = \text{col}(\mathbf{M})$ for any $\pi \in \Delta_k$ (see Lemma 8 in Appendix C), which is a convenient property for the analysis. Still, $Q_\delta(\pi)$ need not be invertible as we allow that $\text{rk}(\mathbf{M}) < d$; hence, the ensuing presentation will feature quantities of the form $Q(\pi)^\dagger$ and $Q_\delta(\pi)^\dagger$. Returning to our task, the following lemma specifies a certain form for the pair (p_t, g_t) and provides the resulting regret bound under global observability.

Lemma 2. *In the same setting as Lemma 1, let the predictions of the policy satisfy $p_t = (1 - \delta)\tilde{p}_t + \delta \mathbf{1}_k/k$ for some $\delta \in (0, 1)$ and $\tilde{p}_t \in \Delta_k$, and let the surrogate loss functions satisfy $\hat{y}_t = g_t(A_t, \phi_t)$ where $g_t(a, \phi) = (\mathbf{I}_k - \mathbf{1}_k q_t^\top) H^\top Q(p_t)^\dagger M_a \phi$. Then, assuming global observability, it holds that*

$$R_T \leq \frac{\log k}{\eta} + 2\delta T + \sum_{t=1}^T \mathbb{E} \left[\langle \tilde{p}_t - q_t, H^\top \ell_t \rangle + \eta \omega^2 \sum_{a,b \in \mathcal{A}} q_t(a) q_t(b) (\psi_a - \psi_b)^\top Q_\delta(\tilde{p}_t)^\dagger (\psi_a - \psi_b) \right]$$

provided that $\max_{a,b \in \mathcal{A}} (\psi_a - \psi_b)^\top Q_\delta(\tilde{p}_t)^\dagger (\psi_a - \psi_b) + \max_{c \in \mathcal{A}} \|M_c^\top Q_\delta(\tilde{p}_t)^\dagger M_c\|_2 \leq \frac{2}{\eta\omega}$, where

$$\omega := \sup_{a,b,c \in \mathcal{A}, p \in \Delta_k, \ell \in \mathcal{L}} \frac{|(\psi_b - \psi_c)^\top Q_\delta(p)^\dagger M_a M_a^\top \ell|}{\|M_a^\top Q_\delta(p)^\dagger (\psi_b - \psi_c)\|} \leq \max_{a \in \mathcal{A}, \ell \in \mathcal{L}} \|M_a^\top \ell\|.$$

Here, $Q(p_t)^\dagger M_{A_t} \phi_t =: \hat{\ell}_t$ is an estimator of ℓ_t that is essentially unbiased.² Instead of taking the loss estimate of action $a \in \mathcal{A}$ as $\psi_a^\top \hat{\ell}_t$ as common in linear bandit algorithms, we use a shifted version $(\psi_a - Hq_t)^\top \hat{\ell}_t$ anchored about $Hq_t = \sum_{b \in \mathcal{A}} q_t(b) \psi_b$. Shifting the loss estimators in this manner (via any anchor in $\text{co}(\mathcal{A})$) does not alter the resulting q_t ; however, its effect is that the last term in the bound of Lemma 2 (known as the variance term) is expressed in terms of feature *differences*. This is crucial if one is to take advantage of either observability condition, since they only relate the observation matrices (in terms of which $Q_\delta(\cdot)$ is defined) to feature differences. Specifically using Hq_t as the anchor is, in a certain sense, an optimal choice, see Lemmas 13 and 14 in Appendix C.

With the structures imposed upon g_t and p_t in Lemma 2, we can now reduce our task of minimizing the regret bound to solving a (constrained) optimization problem at every step t (depending on the history only through q_t) over the distribution $\tilde{p}_t \in \mathbb{R}^k$. In the following, let L be a free parameter satisfying $L \geq \omega$. Define the function $\mathcal{E}: \mathcal{A} \times \mathcal{A} \times \Delta_k \rightarrow \mathbb{R}$ as $\mathcal{E}(a, b; p) := (\psi_a - \psi_b)^\top Q_\delta(p)^\dagger (\psi_a - \psi_b)$, and for every pair $q \in \Delta_k$ and $\eta > 0$, define the function $\Lambda_{\eta,q}: \Delta_k \times \mathcal{L} \rightarrow \mathbb{R}$ as³

$$\Lambda_{\eta,q}(p, \ell) := \frac{1}{\eta} \langle p - q, H^\top \ell \rangle + L^2 \sum_{a,b \in \mathcal{A}} q(a) q(b) \mathcal{E}(a, b; p).$$

Next, define $z: \Delta_k \rightarrow \mathbb{R}$ as $z(p) := \max_{a,b \in \mathcal{A}} \mathcal{E}(a, b; p) + \max_{c \in \mathcal{A}} \|M_c^\top Q_\delta(p)^\dagger M_c\|_2$ and its sub-level set $\Xi_\eta := \{p \in \Delta_k: z(p) \leq 2/\eta L\}$. Finally, let

$$\Lambda_{\eta,q}^* := \min_{p \in \Xi_\eta} \max_{\ell \in \mathcal{L}} \Lambda_{\eta,q}(p, \ell) \quad \text{and} \quad \Lambda_\eta^* := \sup_{q \in \Delta_k} \Lambda_{\eta,q}^*.$$

Algorithm 1 summarizes the policy we have arrived at: at round t , given q_t , we choose a distribution $\tilde{p}_t$ satisfying $z(\tilde{p}_t) \leq 2/\eta L$ and $\max_{\ell \in \mathcal{L}} \Lambda_{\eta,q_t}(\tilde{p}_t, \ell) \leq \Lambda_\eta^* + \varepsilon$, for some tolerance $\varepsilon > 0$. In the

¹The core algorithm here corresponds to the FTRL algorithm run on the probability simplex over the actions, while they use FTRL with a general potential over the convex hull of the action set. The two approaches can be related, see [28, Section 5.2]. Moreover, we use a second order bound on the Bregman divergence (or stability) term, which introduces the constraint over the norm of the loss estimators.

²Its expectation is the projection of ℓ_t onto $\text{col}(\mathbf{M})$.

³As an alternative form, Lemma 13 in Appendix C implies that $\Lambda_{\eta,q}(p, \ell) = (1/\eta) \langle p - q, H^\top \ell \rangle + 2L^2 \sum_{a \in \mathcal{A}} q(a) (\psi_a - Hq)^\top Q_\delta(p)^\dagger (\psi_a - Hq)$.

Algorithm 1 Anchored Exploration-by-Optimization

- 1: **input:** learning rate $\eta > 0$, stability parameter $\delta \in (0, 1/2]$, sub-optimality tolerance $\varepsilon \geq 0$, scale parameter $L \geq \omega$
 - 2: **initialize:** $\hat{y}_0 = \mathbf{0}$
 - 3: **for** $t = 1, \dots, T$ **do**
 - 4: $\forall a \in \mathcal{A}$, set $q_t(a) \propto \llbracket a \in \mathcal{A}^* \rrbracket \exp(-\eta \sum_{s=1}^{t-1} \hat{y}_s(a))$
 - 5: choose $\tilde{p}_t \in \Delta_k$ such that $\max_{\ell \in \mathcal{L}} \Lambda_{\eta, q_t}(\tilde{p}_t, \ell) \leq \Lambda_\eta^* + \varepsilon$ and $z(\tilde{p}_t) \leq \frac{2}{\eta L}$
 - 6: set $p_t = (1 - \delta)\tilde{p}_t + \delta \mathbf{1}_k/k$
 - 7: execute $A_t \sim p_t$ and observe $\phi_t = M_{A_t}^\top \ell_t$
 - 8: $\forall a \in \mathcal{A}$, set $\hat{y}_t(a) = (\psi_a - Hq_t)^\top Q(p_t)^\dagger M_{A_t} \phi_t$ ▷ or simply, $\hat{y}_t(a) = \psi_a^\top Q(p_t)^\dagger M_{A_t} \phi_t$
 - 9: **end for**
-

following round, q_{t+1} is computed via the exponential weights update in (4) with $\hat{y}_t$ as given in Lemma 2. As shown in Appendix D (in the proof of Lemma 3 below), the functions $z(\cdot)$ and $\max_{\ell \in \mathcal{L}} \Lambda_{\eta, q}(\cdot, \ell)$ are both convex. Hence, the problem at hand is a (finite-dimensional) convex program, the structure of which depends on that of the loss space $\mathcal{L}$, over which the max in the first (and more troublesome) term of $\Lambda_{\eta, q}$ is taken. Typical examples of $\mathcal{L}$, like balls of bounded norm or the polar set of $\mathcal{X}$, can further simplify the form of the objective function, see Appendix G for details. Note that the parameter δ is introduced to maintain numerical stability. Turning back to the regret, the following bound readily follows from Lemma 2.

Proposition 2. *Under global observability, Algorithm 1 satisfies $R_T \leq \frac{\log k}{\eta} + \eta \Lambda_\eta^* T + (2\delta + \eta\varepsilon)T$.*

What remains now is deriving an upper bound on Λ_η^* depending on the structure of the game, and tuning the algorithm's parameters accordingly. Note that an adequate tuning of η would depend on Λ_η^* (or a satisfactory upper bound), which might be difficult to compute. Following Lattimore and Szepesvári [38], we can use a learning rate sequence $(\eta_t)_t$, which is updated at round $t + 1$ using the attained value at the optimization problem of round t . This still allows recovering essentially the same regret rates, details are included in Appendix H for completeness. Now, for bounding Λ_η^* , the properties of our $\Lambda_{\eta, q}$ allow us to invoke Sion's minimax theorem, as Lattimore and Szepesvári [38] did in the finite setting.

Lemma 3. *For any $\eta > 0$ and $q \in \Delta_k$, it holds that*

$$\Lambda_{\eta, q}^* = \max_{\ell \in \mathcal{L}} \min_{p \in \Xi_\eta} \Lambda_{\eta, q}(p, \ell).$$

The next two sections will start from this result and derive bounds on Λ_η^* for locally and globally observable games, exploiting the linear structure of the problem.

4 Locally Observable Games

As alluded to before, what distinguishes locally observable games is that exploration, roughly speaking, is cheap. The following lemma provides an alternative definition of local observability that appeals to this intuition.

Lemma 4. *A game is locally observable if and only if it holds for all $a \in \mathcal{A}^*$, $\ell \in \mathbb{R}^d$, and $a^* \in \arg \min_{a' \in \mathcal{A}} \psi_{a'}^\top \ell$ that $\psi_a - \psi_{a^*} \in \text{span} \{M_b \mid \psi_b^\top \ell \leq \psi_{a^*}^\top \ell\}$.*

Fixing $q \in \Delta_k$ and $\eta > 0$, we show in this section how to exploit this property to bound Λ_η^* . We sketch here the main arguments, deferring details to Appendix E, where proofs of this section's results can be found. Note that Lemma 3 reduces the task of bounding $\Lambda_{\eta, q}^*$ to finding a uniform upper bound on $\min_{p \in \Xi_\eta} \Lambda_{\eta, q}(p, \ell)$ holding against any fixed loss vector ℓ . With the clairvoyant-like knowledge of ℓ that this affords us, we start by constructing a specific distribution $p \in \Xi_\eta$ that depends on ℓ and derive an upper bound on $\Lambda_{\eta, q}$ when evaluated at this pair. Using this result, we can then bound $\Lambda_{\eta, q}^*$ and Λ_η^* by passing to worst case over $\ell \in \mathcal{L}$ and $q \in \Delta_k$.

Our choice of p will take the form $p = \gamma\pi + (1 - \gamma)\hat{p}$, with $\gamma \in (0, 1/2]$ and $\pi, \hat{p} \in \Delta_k$. For brevity, let $\bar{\gamma} = 1 - \gamma$. The role of π is to ensure that $p \in \Xi_\eta$ and that of $\hat{p}$ is to control the magnitude of

$\Lambda_{\eta,q}(p, \ell)$. As a starting point, it is easy to show that

$$\begin{aligned} & \Lambda_{\eta,q}(\gamma\pi + \bar{\gamma}\hat{p}, \ell) \\ & \leq \frac{2\gamma}{\eta} + \frac{\bar{\gamma}}{\eta} \langle \hat{p} - q, H^\top \ell \rangle + L^2 \sum_{a,b \in \mathcal{A}} q(a)q(b) (\psi_a - \psi_b)^\top Q_\delta(\gamma\pi + \bar{\gamma}\hat{p})^\dagger (\psi_a - \psi_b) \end{aligned} \quad (6)$$

$$\leq \frac{2\gamma}{\eta} + \frac{\bar{\gamma}}{\eta} \langle \hat{p} - q, H^\top \ell \rangle + 2L^2 \sum_{a,b \in \mathcal{A}} q(a)q(b) (\psi_a - \psi_b)^\top Q_\delta(\hat{p})^\dagger (\psi_a - \psi_b). \quad (7)$$

We proceed first with the selection of $\hat{p}$ aiming to minimize the second and third terms, which exhibit a tradeoff between minimizing our loss against ℓ (compared to q) and sufficiently exploring to minimize the variance of the loss estimator. In finite partial monitoring [37, 38], the analogous predicament is resolved via a ‘water transfer’ operator. Leveraging properties of the same spirit as Lemma 4, this technique involves perturbing q to ensure that each action receives comparable mass to actions whose estimators rely on the observations of the former, without incurring more loss in the process. While some adaptation of this approach can be utilized here (more details below), we will describe in what follows an alternative manner of redistributing the mass in q that is arguably more natural (especially considering the structure of (7)), and that, moreover, only relies on the simple characterization of local observability given by Lemma 4, without explicit reference to the neighborhood relation used in finite partial monitoring.

To control the third term, it is helpful to first relate the feature differences $(\psi_a - \psi_b)_{a,b \in \mathcal{A}}$ to the observations matrices $(M_a)_{a \in \mathcal{A}}$ (or their collective form $\mathbf{M}$) seeing that $Q_\delta(\hat{p})$ is a linear combination of the latter. While this feat only requires global observability, the stronger assumption of local observability allows one to do so in a manner that facilitates controlling the second term simultaneously, as will be shortly shown. Towards this goal, we define some relevant objects. Let $n := \sum_{a \in \mathcal{A}} n(a)$ and $k^* := |\mathcal{A}^*|$. Note that any $\mathbf{v} \in \mathbb{R}^n$ can, with a slight notation abuse, be decomposed as $(\mathbf{v}(a))_{a \in \mathcal{A}}$ with $\mathbf{v}(a) \in \mathbb{R}^{n(a)}$ such that $\mathbf{M}\mathbf{v} = \sum_{a \in \mathcal{A}} M_a \mathbf{v}(a)$. In locally observable games, as will be shortly shown, Lemma 4 gives that for any loss vector $\ell \in \mathcal{L}$, action $a \in \mathcal{A}^*$, and optimal action $a^* \in \arg \min_{c \in \mathcal{A}} \psi_c^\top \ell$, there exists a vector $\mathbf{v} \in \mathbb{R}^n$ such that $\psi_a - \psi_{a^*} = \mathbf{M}\mathbf{v}$ and that, at the same time, having $\mathbf{v}(b) \neq \mathbf{0}$ for some $b \in \mathcal{A}$ implies that $\psi_b^\top \ell \leq \psi_{a^*}^\top \ell$. Hence, the following set is not empty under local observability:

$$\begin{aligned} \mathcal{W}_\ell^{\text{loc}} := \left\{ \boldsymbol{\lambda} = (\boldsymbol{\lambda}_a)_{a \in \mathcal{A}^*} \in \mathbb{R}^{k^* \times n} \mid \exists a^* \in \arg \min_{c \in \mathcal{A}} \psi_c^\top \ell \ \forall (a, b) \in \mathcal{A}^* \times \mathcal{A}, \right. \\ \left. \psi_a - \psi_{a^*} = \mathbf{M}\boldsymbol{\lambda}_a \text{ and } \boldsymbol{\lambda}_a(b) \neq \mathbf{0} \Rightarrow \psi_b^\top \ell \leq \psi_{a^*}^\top \ell \right\}. \end{aligned}$$

Each member of this set is a sequence of weight-vectors in $\mathbb{R}^n$, one for each action in $\mathcal{A}^*$, that allows recovering the difference between its feature vector and that of a fixed optimal action by linearly combining columns from the observation matrices of better- (or similarly) performing actions. Finally, for $\boldsymbol{\lambda} \in \mathcal{W}_\ell^{\text{loc}}$ and $a \in \mathcal{A}^*$, let $\text{supp}(\boldsymbol{\lambda}, a) := \{b \in \mathcal{A} \mid \boldsymbol{\lambda}_a(b) \neq \mathbf{0}\}$, and define $\text{supp}(\boldsymbol{\lambda}) := \bigcup_{a \in \mathcal{A}^*} \text{supp}(\boldsymbol{\lambda}, a)$; that is, the set of all actions whose observations are relied on.

Now, fixing some $\boldsymbol{\lambda} \in \mathcal{W}_\ell^{\text{loc}}$ and $a^* \in \arg \min_{a' \in \mathcal{A}} \psi_{a'}^\top \ell$ (such that $\psi_a - \psi_{a^*} = \mathbf{M}\boldsymbol{\lambda}_a$ for all $a \in \mathcal{A}$), one can show that

$$\begin{aligned} \sum_{a,b \in \mathcal{A}} q(a)q(b) (\psi_a - \psi_b)^\top Q_\delta(\hat{p})^\dagger (\psi_a - \psi_b) & \leq 2 \sum_{a \in \mathcal{A}} q(a) (\psi_a - \psi_{a^*})^\top Q_\delta(\hat{p})^\dagger (\psi_a - \psi_{a^*}) \\ & \leq 2 \max_{s \in \mathcal{A}^*} (\sum_{c \in \mathcal{A}} \|\boldsymbol{\lambda}_s(c)\|) \sum_{c \in \mathcal{A}} \sum_{a \in \mathcal{A}^*} q(a) \|\boldsymbol{\lambda}_a(c)\| \|M_c^\top Q_\delta(\hat{p})^\dagger M_c\|. \end{aligned}$$

An easily demonstrated fact is that for any $r \in \Delta_k$, $\sum_{c \in \mathcal{A}} r(c) \|M_c^\top Q_\delta(r)^\dagger M_c\| \leq 2 \text{rk}(\mathbf{M}) \leq 2d$. In view of the bound reported above, one can exploit this property by setting $\hat{p} = \sum_{a \in \mathcal{A}^*} q(a) \nu_a$, where $\nu_a \in \Delta_k$ is such that $\nu_a(c) \propto \|\boldsymbol{\lambda}_a(c)\|$. That is, we choose $\hat{p}$ as a mixture (weighted according to q) of k^* distributions, each associated with a Pareto optimal action $a \in \mathcal{A}^*$ and supported on $\text{supp}(\boldsymbol{\lambda}, a)$. In particular, action a transfers a fraction of its total mass $q(a)$ to action c proportionally to $\|\boldsymbol{\lambda}_a(c)\|$, representing the ‘importance’ of c to a . (For a^* , simply set $\nu_{a^*} = e_{a^*}$.) As an aside, though arrived at through a different manner, the exploration distribution used by Lattimore and Szepesvári [38] in the analysis of EXO in finite partial monitoring assumes a form similar to $\hat{p}$, only

that, roughly speaking, ν_a is taken there as uniform over $\text{supp}(\lambda, a)$. We show in Appendix E.1 a way to incorporate this alternative choice (among others) into the analysis scheme showcased in this section, though resulting in a generally worse bound than what we report below for our choice of $\hat{p}$.

Proceeding with the analysis, note that by the definition of $\mathcal{W}^{\text{loc}}$, $\psi_b^\top \ell \leq \psi_a^\top \ell$ for all $b \in \text{supp}(\lambda, a)$; hence, it holds that $\langle \nu_a - e_a, H^\top \ell \rangle \leq 0$ for all $a \in \mathcal{A}^*$. And since $\hat{p} - q = \sum_{a \in \mathcal{A}^*} (\nu_a - e_a)$, we see now that the second term in Equation (7) becomes non-positive. While for the third term, a refinement of the arguments laid above gives that

$$\sum_{a,b \in \mathcal{A}} q(a)q(b) (\psi_a - \psi_b)^\top Q_\delta(\hat{p})^\dagger (\psi_a - \psi_b) \leq 4\beta_\lambda^2 \min\{\text{rk}(\mathbf{M}), |\text{supp}(\lambda)|\},$$

where $\beta_\lambda := \max_{a \in \mathcal{A}^*} \sum_{b \in \mathcal{A}} \|\lambda_a(b)\|$ represents the ‘difficulty’ of recovering the features from the observations via λ . The alternative bound featuring $|\text{supp}(\lambda)|$ brings an improvement in problems with abundant feedback per played action, like the full-information and graph feedback problems. We then proceed by taking the infimum of the bound above over $\lambda \in \mathcal{W}_\ell^{\text{loc}}$, before passing to the worst case over $\ell \in \mathcal{L}$. What remains now is the less challenging task of bounding the first term in (7). This is done by choosing $\gamma \propto \eta$, provided η is small enough, with the leading constant chosen to ensure that $z(p) \leq 2/\eta L$; specifically, it is an upper bound on the value of an optimal-design-like criterion that π is chosen to minimize. Ultimately, we arrive at the following result:

Theorem 1. *In locally observable games, it holds that*

$$\Lambda_\eta^* \leq \max_{\ell \in \mathcal{L}} \inf_{\lambda \in \mathcal{W}_\ell^{\text{loc}}} 8L^2 \beta_\lambda^2 \min\{\text{rk}(\mathbf{M}), |\text{supp}(\lambda)|\} + 2L(1 + \beta_{\text{glo}}^2) \min\{\text{rk}(\mathbf{M}), w^*\}$$

provided $\frac{1}{\eta} \geq 2L(1 + \beta_{\text{glo}}^2) \min\{\text{rk}(\mathbf{M}), w^*\}$, where $\beta_{\text{glo}} := \max_{a,b \in \mathcal{A}} \min_{v \in \mathbb{R}^n : \psi_a - \psi_b = Mv} \sum_{c \in \mathcal{A}} \|v(c)\|$

and $w^* := \min_{S \subseteq \mathcal{A}} |S| \max_{b \in \mathcal{A}} \|M_b^\top U (\sum_{s \in S} U^\top M_s M_s^\top U)^{-1} U^\top M_b\| \leq k$.

Here, $U \in \mathbb{R}^{d \times \text{rk}(\mathbf{M})}$ denotes a matrix whose columns $(u_i)_{i \in [\text{rk}(\mathbf{M})]}$ form an orthonormal basis for $\text{col}(\mathbf{M})$.⁴ Let $\beta_{\text{loc}} := \max_{\ell \in \mathcal{L}} \min_{\lambda \in \mathcal{W}_\ell^{\text{loc}}} \beta_\lambda$, which satisfies $\beta_{\text{loc}} \geq \beta_{\text{glo}}$ (see Lemma 17 in Appendix E). Assuming for simplicity that $L, \beta_{\text{loc}} \geq 1$; the following simpler bound is obtained immediately:

$$\Lambda_\eta^* \leq 12L^2 \beta_{\text{loc}}^2 d.$$

We see that under local observability, similarly to finite partial monitoring [38], there exist constants α, B such that $\Lambda_\eta^* \leq \alpha$ given that $\eta \leq 1/B$. Assuming δ and ε are sufficiently small, setting $\eta = \min\{1/B, \sqrt{\log k / \alpha T}\}$ yields via Proposition 2 that R_T is $\mathcal{O}(\sqrt{\alpha T \log k})$. We refer again to Appendix H for a discussion on adaptive learning rates following [38]. We also note here that B can be chosen conservatively as it only affects the regret additively. We now use the bound of Theorem 1 to obtain regret bounds for the locally observable games discussed in Section 2. Omitted details and derivations are provided in Appendix I.

Full information. Here, playing any action is sufficiently informative; for any $\ell \in \mathcal{L}$, we can trivially pick $\lambda \in \mathcal{W}_\ell^{\text{loc}}$ supported only on an optimal action and satisfying $\beta_\lambda \leq \sqrt{2}$. Moreover, $w^* \leq 1$ taking S as a singleton. Hence, we recover the order-optimal $\mathcal{O}(\sqrt{T \log k})$ regret bound [13].

Strongly observable feedback graphs. Generalizing the example above, we can, against any $\ell \in \mathcal{L}$, find $\lambda \in \mathcal{W}_\ell^{\text{loc}}$ supported on an independent set (a set of actions no two distinct members of which are neighbors) and satisfying $\beta_\lambda \leq 2$. This set can be iteratively constructed by first selecting an optimal action, removing its neighbors from the graph, and repeating this (proceeding with the best remaining action) until the graph is empty. This gives that $\beta_\lambda^2 |\text{supp}(\lambda)|$ is bounded up to a constant by $\alpha(\mathcal{G})$, the independence number of the graph (size of a maximal independent set). The same can be shown to hold for w^* , allowing us to recover the near-optimal $\mathcal{O}(\sqrt{\alpha(\mathcal{G}) T \log k})$ regret bound for this problem [42, 2].

Linear bandits. Since the features and the observation vectors coincide, it is immediate that $\beta_{\text{loc}} \leq 2$. The regret bound is then of order $\sqrt{dT \log(k)}$ (as achieved in [10] and [33, Chp. 27]), which reduces to the near-optimal $\sqrt{kT \log(k)}$ rate for standard bandits.

⁴The inverse of $\sum_{s \in S} U^\top M_s M_s^\top U$ exists when $\text{span}(\{M_s\}_{s \in S}) = \text{col}(\mathbf{M})$, see Lemma 9 in Appendix C.

Linear dueling bandits. It also holds in this example that $\beta_{\text{loc}} \leq 2$, implying again that the regret bound is of order $\sqrt{dT \log(k)}$.

Bandits with ill-conditioned observers. Since $e_a - e_b = 1/\varepsilon(M_a - M_b)$, we have that $\beta_{\text{loc}} \leq 2/\varepsilon$. Thus, the regret bound is of order $1/\varepsilon \sqrt{kT \log k}$. We show in Appendix J that this bound is tight up to logarithmic factors.

5 Globally Observable Games

On the other hand, exploration is generally costly in globally observable games. Here, starting from the decomposition in (6), we simply pick $\hat{p} = q$ and place the burden of controlling the variance term entirely on π . Consequently, γ must be kept sufficiently large; precisely, of order $\sqrt{\eta}$. This causes the final bound to scale with $1/\sqrt{\eta}$, hence becoming η -dependent, which aligns with the finite outcomes setting [38]. Building on this, the next theorem provides two bounds on Λ_η^* . The first resembles the square root of the second term in the bound of Theorem 1, while the second bound features a refined alignment constant but has a generally worse dependence on the structure of the observation matrices.

Theorem 2. *In globally observable games, it holds that*

$$\Lambda_\eta^* \leq 4\sqrt{1/\eta}L \min\left\{(1 + \beta_{\text{glo}})\sqrt{\min\{\text{rk}(\mathbf{M}), w^*\}}, (1 + \beta_{2,\text{glo}})\sqrt{u^*}\right\}$$

provided that $\frac{1}{\eta} \geq (1 + L^2) \min\left\{(1 + \beta_{\text{glo}}^2)\min\{\text{rk}(\mathbf{M}), w^*\}, (1 + \beta_{2,\text{glo}}^2)u^*\right\}$, where

$$\beta_{2,\text{glo}} := \max_{a,b \in \mathcal{A}} \|\mathbf{M}^\dagger(\psi_a - \psi_b)\|, \text{ and } u^* := \min_{S \subseteq \mathcal{A}} |S| \|\mathbf{M}^\top U (\sum_{s \in S} U^\top M_s M_s^\top U)^{-1} U^\top \mathbf{M}\| \leq k.$$

Note that $w^* \leq u^*$, and that $\beta_{2,\text{glo}} = \max_{a,b \in \mathcal{A}} \min_{\mathbf{v} \in \mathbb{R}^n: \psi_a - \psi_b = \mathbf{M}\mathbf{v}} \|\mathbf{v}\|$; hence, $\beta_{2,\text{glo}} \leq \beta_{\text{glo}} \leq \sqrt{k}\beta_{2,\text{glo}}$. This theorem shows that in the present regime, there exist constants α, B such that $\Lambda_\eta^* \leq \alpha/\sqrt{\eta}$ given that $\eta \leq 1/B$. Assuming again that δ and ε are sufficiently small, setting $\eta = \min\{1/B, (\log k/\alpha T)^{2/3}\}$ yields via Proposition 2 that R_T is $\mathcal{O}((\alpha T)^{2/3}(\log k)^{1/3})$. The adaptive tuning of η is discussed in Appendix H. We now revisit the globally observable examples from Section 2.

Weakly observable feedback graphs. Here, it is immediate to verify that $\beta_{\text{glo}} \leq 2$. A subset of vertices (actions) is said to be a total dominating set if the union of their neighborhoods covers all vertices. Moreover, the total domination number of the graph, $\delta(\mathcal{G})$, is defined as the size of a minimal total dominating set. It can be easily shown (see Appendix I) that $w^* \leq \delta(\mathcal{G})$, leading to a $\mathcal{O}((\delta(\mathcal{G}) \log k)^{1/3} T^{2/3})$ regret bound. Further refining this bound to scale with the *weak* domination number (which only considers weakly observable nodes) used in [1] requires a careful combination of the analyses of this section and the previous one to eliminate unnecessary forced exploration over the strongly observable portion of the graph.

Composite graph feedback. It holds here that $w^* = u^* = \text{rk}(\mathbf{M}) = k$; hence, the second bound is generally better. The regret bound is then $\mathcal{O}((k \log k)^{1/3} (\beta_{2,\text{glo}} T)^{2/3})$, the tightness of which remains to be studied. In Appendix I, we provide a more detailed discussion about the interpretation of $\beta_{2,\text{glo}}$ in this problem, as well as two instantiations of which in simple families of graphs.

6 Limitations and Future Work

Besides enjoying the same rates in T as the bounds derived in this work, the regret bounds derived by Kirschner et al. [29, 30] in the stochastic setting feature analogous alignment constants to β_{loc} and β_{glo} , see [29, Lemma 13] or [30, Lemma 5]. Comparatively, our bounds enjoy some extra versatility as they easily yield nearly tight guarantees for feedback-rich problems like full-information and feedback graphs, which were not addressed in [29, 30]. However, our exposition is limited to full-dimensional loss spaces (like [29] but unlike [30]) and finite actions. The former aspect precludes us from recovering tight guarantees for finite partial monitoring as the observation conditions do not match in general, see [30, Example 8]. Addressing this aspect in the adversarial setting is left for future work. Likewise for studying generic compact action sets, for which a full classification theorem is lacking also in the stochastic setting. Lastly, characterizing how the minimax regret depends on instance-based constants remains a general challenging direction.

Acknowledgments and Disclosure of Funding

KE and NCB acknowledge the financial support from the FAIR project, funded by the NextGenerationEU program within the PNRR-PE-AI scheme (M4C2, investment 1.3, line on Artificial Intelligence), the MUR PRIN grant 2022EKNE5K (Learning in Markets and Society), funded by the NextGenerationEU program within the PNRR scheme (M4C2, investment 1.1), the EU Horizon CL4-2022-HUMAN-02 research and innovation action under grant agreement 101120237, project ELIAS (European Lighthouse of AI for Sustainability), and the One Health Action Hub, University Task Force for the resilience of territorial ecosystems, funded by Università degli Studi di Milano (PSR 2021-GSA-Linea 6).

References

- [1] Noga Alon, Nicolò Cesa-Bianchi, Ofer Dekel, and Tomer Koren. Online learning with feedback graphs: Beyond bandits. In *Proceedings of The 28th Conference on Learning Theory*, volume 40 of *Proceedings of Machine Learning Research*, pages 23–35. PMLR, 2015.
- [2] Noga Alon, Nicolò Cesa-Bianchi, Claudio Gentile, Shie Mannor, Yishay Mansour, and Ohad Shamir. Nonstochastic multi-armed bandits with graph-structured feedback. *SIAM Journal on Computing*, 46(6):1785–1826, 2017.
- [3] András Antos, Gábor Bartók, Dávid Pál, and Csaba Szepesvári. Toward a classification of finite partial-monitoring games. *Theoretical Computer Science*, 473:77–99, 2013. ISSN 0304-3975.
- [4] Peter Auer, Nicolò Cesa-Bianchi, Yoav Freund, and Robert E Schapire. The nonstochastic multiarmed bandit problem. *SIAM journal on computing*, 32(1):48–77, 2002.
- [5] Michel Balinski. On the graph structure of convex polyhedra in n -space. *Pacific Journal of Mathematics*, 11, 06 1961. doi: 10.2140/pjm.1961.11.431.
- [6] Gábor Bartók, Dean P Foster, Dávid Pál, Alexander Rakhlin, and Csaba Szepesvári. Partial monitoring—classification, regret bounds, and algorithms. *Mathematics of Operations Research*, 39(4):967–997, 2014.
- [7] Gábor Bartók, Dávid Pál, and Csaba Szepesvári. Minimax regret of finite partial-monitoring games in stochastic environments. In *Proceedings of the 24th Annual Conference on Learning Theory*, volume 19 of *Proceedings of Machine Learning Research*, pages 133–154. PMLR, 2011.
- [8] Gábor Bartók, Navid Zolghadr, and Csaba Szepesvári. An adaptive algorithm for finite stochastic partial monitoring. In *Proceedings of the 29th International Conference on Machine Learning (ICML-12)*, ICML ’12, pages 1727–1734. Omnipress, 2012.
- [9] Aharon Ben-Tal and Arkadi Nemirovski. *Lectures on Modern Convex Optimization*. MPS/SIAM Series on Optimization. SIAM, Philadelphia, PA, 2001. ISBN 978-0-89871-491-3.
- [10] Sébastien Bubeck, Nicolò Cesa-Bianchi, and Sham M Kakade. Towards minimax policies for online linear optimization with bandit feedback. In *Proceedings of the 25th Annual Conference on Learning Theory*, volume 23 of *Proceedings of Machine Learning Research*, pages 41.1–41.14. PMLR, 2012.
- [11] Sébastien Bubeck, Nicolo Cesa-Bianchi, et al. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1):1–122, 2012.
- [12] David Carlson, Emile Haynsworth, and Thomas Markham. A generalization of the Schur complement by means of the Moore-Penrose inverse. *SIAM Journal on Applied Mathematics*, 26(1):169–175, 1974.
- [13] Nicolo Cesa-Bianchi, Yoav Freund, David Haussler, David P Helmbold, Robert E Schapire, and Manfred K Warmuth. How to use expert advice. *Journal of the ACM (JACM)*, 44(3):427–485, 1997.

- [14] Nicolo Cesa-Bianchi, Gábor Lugosi, and Gilles Stoltz. Regret minimization under partial monitoring. *Mathematics of Operations Research*, 31(3):562–580, 2006.
- [15] Sougata Chaudhuri and Ambuj Tewari. Phased exploration with greedy exploitation in stochastic combinatorial partial monitoring games. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- [16] Alon Cohen, Tamir Hazan, and Tomer Koren. Tight bounds for bandit combinatorial optimization. In *Conference on Learning Theory*, pages 629–642. PMLR, 2017.
- [17] Varsha Dani, Sham M Kakade, and Thomas Hayes. The price of bandit information for online optimization. In *Advances in Neural Information Processing Systems*, volume 20. Curran Associates, Inc., 2007.
- [18] Khaled Eldowa, Nicolò Cesa-Bianchi, Alberto Maria Metelli, and Marcello Restelli. Information capacity regret bounds for bandits with mediator feedback. *Journal of Machine Learning Research*, 25(353):1–36, 2024.
- [19] Erhan Baki Ermis and Venkatesh Saligrama. Distributed detection in sensor networks with limited range multimodal sensors. *IEEE Transactions on Signal Processing*, 58(2):843–858, 2010. doi: 10.1109/TSP.2009.2033300.
- [20] Dean Foster and Alexander Rakhlin. No internal regret via neighborhood watch. In *Proceedings of the Fifteenth International Conference on Artificial Intelligence and Statistics*, volume 22 of *Proceedings of Machine Learning Research*, pages 382–390. PMLR, 2012.
- [21] Dylan J Foster, Sham M Kakade, Jian Qian, and Alexander Rakhlin. The statistical complexity of interactive decision making. *arXiv preprint arXiv:2112.13487*, 2021.
- [22] Dylan J Foster, Alexander Rakhlin, Ayush Sekhari, and Karthik Sridharan. On the complexity of adversarial decision making. In *Advances in Neural Information Processing Systems*, volume 35, pages 35404–35417. Curran Associates, Inc., 2022.
- [23] Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997.
- [24] Pratik Gajane and Tanguy Urvoy. Utility-based dueling bandits as a partial monitoring game. *arXiv preprint arXiv:1507.02750*, 2015.
- [25] Manjesh Hanawal, Venkatesh Saligrama, Michal Valko, and Rémi Munos. Cheap bandits. In *International Conference on Machine Learning*, pages 2133–2142. PMLR, 2015.
- [26] Elad Hazan. *Introduction to Online Convex Optimization*. MIT Press, 2022.
- [27] Andrew Holbrook. Differentiating the pseudo determinant. *Linear Algebra and its Applications*, 548:293–304, 2018. ISSN 0024-3795.
- [28] Shinji Ito and Kei Takemura. An exploration-by-optimization approach to best of both worlds in linear bandits. In *Advances in Neural Information Processing Systems*, volume 36, pages 71582–71602. Curran Associates, Inc., 2023.
- [29] Johannes Kirschner, Tor Lattimore, and Andreas Krause. Information directed sampling for linear partial monitoring. In *Conference on Learning Theory*, pages 2328–2369. PMLR, 2020.
- [30] Johannes Kirschner, Tor Lattimore, and Andreas Krause. Linear partial monitoring for sequential decision making: Algorithms, regret bounds and applications. *The Journal of Machine Learning Research*, 24(1):16606–16650, 2023.
- [31] Junpei Komiyama, Junya Honda, and Hiroshi Nakagawa. Regret lower bound and optimal algorithm in finite stochastic partial monitoring. In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015.

- [32] Joon Kwon and Vianney Perchet. Online learning and Blackwell approachability with partial monitoring: Optimal convergence rates. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics*, volume 54 of *Proceedings of Machine Learning Research*, pages 604–613. PMLR, 2017.
- [33] T. Lattimore and Cs. Szepesvári. *Bandit Algorithms*. Cambridge University Press, 2020.
- [34] Tor Lattimore. Minimax regret for partial monitoring: Infinite outcomes and Rustichini’s regret. In *Conference on Learning Theory*, pages 1547–1575. PMLR, 2022.
- [35] Tor Lattimore and Andras Gyorgy. Mirror descent and the information ratio. In *Conference on Learning Theory*, pages 2965–2992. PMLR, 2021.
- [36] Tor Lattimore and Csaba Szepesvári. Cleaning up the neighborhood: A full classification for adversarial partial monitoring. In *Algorithmic Learning Theory*, pages 529–556. PMLR, 2019.
- [37] Tor Lattimore and Csaba Szepesvári. An information-theoretic approach to minimax regret in partial monitoring. In *Proceedings of the Thirty-Second Conference on Learning Theory*, volume 99 of *Proceedings of Machine Learning Research*, pages 2111–2139. PMLR, 2019.
- [38] Tor Lattimore and Csaba Szepesvári. Exploration by optimisation in partial monitoring. In *Conference on Learning Theory*, pages 2488–2515. PMLR, 2020.
- [39] B. Laurent and Pascal Massart. Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, 28, 10 2000.
- [40] Tian Lin, Bruno Abrahao, Robert Kleinberg, John Lui, and Wei Chen. Combinatorial partial monitoring game with linear feedback and its applications. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 901–909. PMLR, 2014.
- [41] Gábor Lugosi, Shie Mannor, and Gilles Stoltz. Strategies for prediction under imperfect monitoring. *Mathematics of Operations Research*, 33(3):513–528, 2008.
- [42] Shie Mannor and Ohad Shamir. From bandits to experts: On the value of side-observations. In *Advances in neural information processing systems*, volume 24, 2011.
- [43] Shie Mannor and Nahum Shimkin. On-line learning with imperfect monitoring. In *Learning Theory and Kernel Machines*, pages 552–566, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg. ISBN 978-3-540-45167-9.
- [44] Shie Mannor, Vianney Perchet, and Gilles Stoltz. Set-valued approachability and online learning with partial monitoring. *Journal of Machine Learning Research*, 15(94):3247–3295, 2014.
- [45] Kenneth Nordström. Convexity of the inverse and Moore–Penrose inverse. *Linear Algebra and its Applications*, 434(6):1489–1512, 2011.
- [46] Fabio Piccolboni and Christian Schindelhauer. Discrete prediction games with arbitrary feedback and loss. *COLT*, 2001.
- [47] Daniel Russo and Benjamin Van Roy. Learning to optimize via information-directed sampling. In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.
- [48] Aldo Rustichini. Minimizing regret: The general case. *Games and Economic Behavior*, 29 (1-2):224–243, 1999.
- [49] Maurice Sion. On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171 – 176, 1958.
- [50] G. W. Stewart. On the perturbation of pseudo-inverses, projections and linear least squares problems. *SIAM Review*, 19(4):634–662, 1977.

- [51] Taira Tsuchiya, Junya Honda, and Masashi Sugiyama. Analysis and design of Thompson sampling for stochastic partial monitoring. In *Advances in Neural Information Processing Systems*, volume 33, pages 8861–8871. Curran Associates, Inc., 2020.
- [52] Taira Tsuchiya, Shinji Ito, and Junya Honda. Best-of-both-worlds algorithms for partial monitoring. In *Proceedings of The 34th International Conference on Algorithmic Learning Theory*, volume 201 of *Proceedings of Machine Learning Research*, pages 1484–1515. PMLR, 2023.
- [53] Taira Tsuchiya, Shinji Ito, and Junya Honda. Stability-penalty-adaptive follow-the-regularized-leader: Sparsity, game-dependency, and best-of-both-worlds. In *Advances in Neural Information Processing Systems*, volume 36, pages 47406–47437. Curran Associates, Inc., 2023.
- [54] Taira Tsuchiya, Shinji Ito, and Junya Honda. Exploration by optimization with hybrid regularizers: Logarithmic regret with adversarial robustness in partial monitoring. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pages 48768–48790. PMLR, 21–27 Jul 2024.
- [55] Michal Valko, Remi Munos, Branislav Kveton, and Tomáš Kocák. Spectral bandits for smooth graph functions. In Eric P. Xing and Tony Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 46–54, Beijing, China, 22–24 Jun 2014. PMLR.
- [56] Hastagiri P Vanchinathan, Gábor Bartók, and Andreas Krause. Efficient partial monitoring with prior information. In *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014.
- [57] Fuzhen Zhang. *The Schur complement and its applications*, volume 4. Springer Science & Business Media, 2006.

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: All the claims made in the abstract and introduction are formalized via theorems stated in the main body and formally proven in the appendix.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: We explicitly added a section where limitations are discussed.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory assumptions and proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [\[Yes\]](#)

Justification:

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental result reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental setting/details

Question: Does the paper specify all the training and test details (e.g., data splits, hyper-parameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment statistical significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).
- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.

- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments compute resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code of ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: The research conducted in this paper fully conforms to the NeurIPS Code of Ethics. There are no ethical concerns related to data collection, experiments, or other aspects of the work, as it focuses purely on theoretical analysis.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: This paper focuses on theoretical advancements in partial monitoring and does not have direct societal applications. As such, it does not explicitly discuss potential positive or negative societal impacts. The work is foundational in nature and does not involve technologies that could be misused or present ethical concerns in its current form.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.
- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.

- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.
- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.

- If this information is not available online, the authors are encouraged to reach out to the asset’s creators.

13. **New assets**

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. **Crowdsourcing and research with human subjects**

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. **Institutional review board (IRB) approvals or equivalent for research with human subjects**

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [NA]

Justification:

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.

16. **Declaration of LLM usage**

Question: Does the paper describe the usage of LLMs if it is an important, original, or non-standard component of the core methods in this research? Note that if the LLM is used only for writing, editing, or formatting purposes and does not impact the core methodology, scientific rigorousness, or originality of the research, declaration is not required.

Answer: [NA]

Justification: The core method development in this research does not involve LLMs as any important, original, or non-standard components.

Guidelines:

- The answer NA means that the core method development in this research does not involve LLMs as any important, original, or non-standard components.
- Please refer to our LLM policy (<https://neurips.cc/Conferences/2025/LLM>) for what should or should not be described.

A Additional Related Works

The work of Rustichini [48] is generally recognized as the first work on partial monitoring. Notably, he studied a modified notion of regret under which non-trivial guarantees can be derived for hopeless games, see also [43, 41, 44, 32, 34]. Finite partial monitoring has also been extensively studied in the stochastic setting [7, 8, 56, 31, 51]. Among the works cited earlier in the adversarial (finite) setting, of most relevance here are those by Lattimore and Szepesvári [37, 38]. The first proved a minimax theorem for all games with finite actions, linking the minimax (adversarial) regret and the worst-case Bayesian regret. For finite partial monitoring, they bounded the latter by presenting an algorithm (Mario sampling) enjoying a bounded information ratio in the Bayesian setting, which in turn, implies a non-constructive guarantee on the former. The second work, as mentioned already, proposed the EXO method, and showed that its regret matches the non-constructive bounds of the first work. The subsequent work of Lattimore and Gyorgy [35] does, in some sense, link these two results in a general linear partial monitoring setup by showing that the regret of EXO can be bounded in terms of a general notion of information ratio. We note here that the results by Kirschner et al. [29, 30] in the *stationary* stochastic linear setting cannot be combined with these findings to yield bounds on the adversarial minimax regret or the regret of EXO, as that requires bounding (some form of) the information ratio for arbitrary (finitely supported) priors over *sequences* of latent outcomes.

In the same line of work, Lattimore [34] expanded on the results of Lattimore and Gyorgy [35] considering all games with infinite outcomes and finite actions. He showed that the optimal rate in T is indeed determined by the behavior of the information ratio, and that for every $p \in [1/2, 1]$, there exist games where the minimax regret is T^p up to subpolynomial factors. Further, Foster et al. [22] showed that the regret of EXO can also be bounded in terms of the Decision-Estimation Coefficient, a general complexity measure for interactive decision making [21], which too is rather opaque in terms of its dependence on game-specific parameters.

The first work on linear partial monitoring is by Lin et al. [40], followed by the closely related work of Chaudhuri and Tewari [15]. Both study the stochastic setting and only prove $T^{2/3}$ bounds under a global observability condition. On the other hand, they focus on designing efficient algorithms in the face of combinatorial action spaces that can be exponentially large. On a different thread, Ito and Takemura [28] employ a version of EXO in the standard linear bandit setting (hence our algorithm and theirs naturally follow a similar general template) and obtain best-of-both-worlds (BOBW) bounds (near optimal bounds in both the adversarial and the stochastic settings). Their techniques and contributions remain distinct from ours as they do not address partial monitoring. In a related line of work [52, 53, 54], EXO is used to obtain BOBW bounds in partial monitoring, but only for finite outcomes.

B The Geometry of a Linear Partial Monitoring Game

This section provides a brief overview of some partial monitoring concepts and terminology, see [3, 6, 29, 33], which will be used in the proofs, particularly those of the lower bounds in Appendix J. We remark that the definitions used here mostly concern the case when $\mathcal{B}_2(r) \subseteq \mathcal{L}$ for some $r > 0$, which is our focus in this work. Recall that $\mathcal{X} := \{\psi_a \mid a \in \mathcal{A}\}$, and define $\mathcal{V} := \text{co}(\mathcal{X})$, which is a convex polytope. An action $a \in \mathcal{A}$ is called *Pareto optimal* if $\psi_a \in \text{ext}(\mathcal{X})$; that is, ψ_a is a vertex of $\mathcal{V}$. The convex cone

$$C_a := \{\ell \in \mathbb{R}^d \mid \forall b \in \mathcal{A}, \langle \psi_a - \psi_b, \ell \rangle \leq 0\} = \{\ell \in \mathbb{R}^d \mid \forall x \in \mathcal{V}, \langle \psi_a - x, \ell \rangle \leq 0\}$$

is called the *cell* of a . A Pareto optimal action a satisfies that $\{\psi_a\} \cap \text{co}(\mathcal{X} \setminus \{\psi_a\}) = \emptyset$ since ψ_a is a vertex of $\mathcal{V}$. Hence, by the separating hyperplane theorem, there exists a vector $u \in \mathbb{R}^d$ such that $\psi_a^\top u < x^\top u$ for all $x \in \mathcal{X} \setminus \{\psi_a\}$, yielding that $\psi_a^\top u < \psi_b^\top u$ for all $b \in \mathcal{A}$ with $\psi_b \neq \psi_a$. It follows that $\dim(C_a) = d$, where $\dim(S)$ is the dimension of the affine hull of S , since we can construct a small enough Euclidean ball around u where the same condition holds (recall that $\mathcal{A}$ is finite). Conversely, $\dim(C_a) = d$ implies that a is Pareto optimal. This is because if ψ_a can be written as convex combination of points in $\mathcal{X} \setminus \{\psi_a\}$, then for every such point ψ_b (with positive weight), $\langle \psi_a - \psi_b, \ell \rangle = 0$ for all $\ell \in C_a$; hence, C_a cannot be full-dimensional.

The polytope $\mathcal{V}$ induces a graph (the 1-skeleton of $\mathcal{V}$) whose vertices are the vertices of $\mathcal{V}$, two of which are adjacent in this graph if they are the endpoints of an edge of $\mathcal{V}$. Two non-duplicate (i.e., having different features) Pareto optimal actions a and b are said to be *neighbors* if ψ_a and ψ_b are

adjacent in the 1-skeleton of $\mathcal{V}$. It is known that this graph is connected [5]. For any two neighbors a and b , it holds that $\dim(C_a \cap C_b) = d - 1$. This is true since

$$C_a \cap C_b = \{\ell \in \mathbb{R}^d \mid \langle \psi_a - \psi_b, \ell \rangle = 0 \text{ and } \forall c \in \mathcal{A}, \langle \psi_a - \psi_c, \ell \rangle \leq 0\},$$

$\mathcal{A}$ is finite, and the fact that there exists (by virtue of ψ_a and ψ_b being endpoints of an edge of $\mathcal{V}$) a vector $u \in \mathbb{R}^d$ such that $\psi_a^\top u < x^\top u$ for all $x \in \mathcal{V} \setminus [\psi_a, \psi_b]$ (which also excludes duplicates of a and b) and $\psi_a^\top u = \psi_b^\top u$. Conversely, if $\dim(C_a \cap C_b) = d - 1$ for two Pareto optimal actions a and b , then they must be neighbors. To see this, note that each $\ell \in C_a \cap C_b$ can be mapped to a face of $\mathcal{V}$ containing both ψ_a and ψ_b in addition to points $x \in \mathcal{V}$ satisfying $\langle x, \ell \rangle = \langle \psi_a, \ell \rangle$. Since the intersection of two faces is also a face, the intersection of all faces containing ψ_a and ψ_b is still a face of $\mathcal{V}$. If a and b are not neighbors, this face must contain some vertex ψ_c different from ψ_a and ψ_b . Therefore, for any $\ell \in C_a \cap C_b$, $\langle \psi_c, \ell \rangle = \langle \psi_a, \ell \rangle = \langle \psi_b, \ell \rangle$, implying that $\dim(C_a \cap C_b) \leq d - 2$ since ψ_a, ψ_b , and ψ_c are affinely independent being all distinct vertices of $\mathcal{V}$. Lastly, for neighboring actions a and b , define their neighborhood as

$$\mathcal{N}_{a,b} := \{c \in \mathcal{A} \mid \psi_c \in [\psi_a, \psi_b]\} = \{c \in \mathcal{A} \mid \psi_c = \alpha\psi_a + (1 - \alpha)\psi_b, \alpha \in [0, 1]\},$$

which includes the duplicates of a and b . (Note that the neighborhood relation is defined only for non-duplicate Pareto optimal actions.)

Recall the definition $\mathcal{P}(\ell) := \{a \in \mathcal{A} \mid \langle \psi_a - \psi_b, \ell \rangle \leq 0 \text{ for all } b \in \mathcal{A}\}$ for $\ell \in \mathbb{R}^d$, extending to sets $D \subseteq \mathbb{R}^d$ via $\mathcal{P}(D) := \cup_{\ell \in D} \mathcal{P}(\ell)$. Lemma 26 in [29] shows that local observability as defined via Definition 1 is equivalent to the more usual local observability condition used in finite partial monitoring. The following lemma expands on that adding two more equivalent conditions, one of which was featured in Lemma 4.

Lemma 5. *The following conditions are equivalent.*

i) For every convex set $D \subseteq \mathbb{R}^d$,

$$\psi_a - \psi_b \in \text{span}\{M_c \mid c \in \mathcal{P}(D)\} \text{ for all } a, b \in \mathcal{P}(D).$$

ii) For any pair of neighboring actions a and b ,

$$\psi_a - \psi_b \in \text{span}\{M_c \mid c \in \mathcal{N}_{a,b}\}.$$

iii) For any Pareto optimal action a , $\ell \in \mathbb{R}^d$, and $a^* \in \arg \min_{a' \in \mathcal{A}} \psi_{a'}^\top \ell$,

$$\psi_a - \psi_{a^*} \in \text{span}\{M_b \mid \psi_b^\top \ell \leq \psi_a^\top \ell\}.$$

iv) For any pair of Pareto optimal actions a and b , and any $\ell \in \mathbb{R}^d$,

$$\psi_a - \psi_b \in \text{span}\{M_c \mid \psi_c^\top \ell \leq \max\{\psi_a^\top \ell, \psi_b^\top \ell\}\}.$$

Proof. $i) \Leftrightarrow ii)$ This is given by [29, Lemma 26]. (Note that, unlike [29], our setting allows duplicate actions; however, their proof extends immediately to the case when duplicates are present.)

$i) \Rightarrow iii)$ Since a is Pareto optimal, $\psi_a \in \text{ext}(\mathcal{V})$, implying that $\{\psi_a\} \cap \text{co}(\mathcal{X} \setminus \{\psi_a\}) = \emptyset$. Hence, by the separating hyperplane theorem, there exists a vector $v \in \mathbb{R}^d$ such that $\psi_a^\top v < x^\top v$ for all $x \in \mathcal{X} \setminus \{\psi_a\}$, yielding that $\psi_a^\top v < \psi_b^\top v$ for all $b \in \mathcal{A}$ with $\psi_b \neq \psi_a$. Define E as the line segment between v and ℓ , i.e.,

$$E := \{\ell' \in \mathbb{R}^d \mid \ell' = (1 - \lambda)\ell + \lambda v, \lambda \in [0, 1]\}.$$

Now, considering that $a, a^* \in \mathcal{P}(E)$, Condition $i)$ implies that

$$\psi_a - \psi_{a^*} \in \text{span}\{M_b \mid b \in \mathcal{P}(E)\}.$$

We now show that for every $b \in \mathcal{P}(E)$, $\psi_b^\top \ell \leq \psi_a^\top \ell$, from which the lemma follows directly. If $\psi_b = \psi_a$, the inequality holds trivially. Suppose now that $\psi_b \neq \psi_a$. By the definition of v , $\psi_a^\top v < \psi_b^\top v$. Hence, since $b \in \mathcal{P}(E)$, we have that for some $\lambda \in [0, 1]$,

$$(1 - \lambda)\psi_b^\top \ell - (1 - \lambda)\psi_a^\top \ell \leq \lambda\psi_a^\top v - \lambda\psi_b^\top v \leq 0.$$

Dividing by $1 - \lambda$ (note that $\lambda < 1$) gives that $\psi_b^\top \ell \leq \psi_a^\top \ell$ as required.

$iii) \Rightarrow iv)$ This can be seen immediately upon writing $\psi_a - \psi_b = \psi_a - \psi_{a^*} - (\psi_b - \psi_{a^*})$.

$iv) \Rightarrow ii)$ For neighbors a and b , $[\psi_a, \psi_b]$ is a face (in particular, an edge) of the polytope $\mathcal{V}$; hence, there exists $r \in \mathbb{R}$ and $\ell \in \mathbb{R}^d$ such that $x^\top \ell = r$ for $x \in [\psi_a, \psi_b]$ and $y^\top \ell > r$ for $y \in \mathcal{V} \setminus [\psi_a, \psi_b]$. This gives, via the definition of $\mathcal{N}_{a,b}$, that for any $c \in \mathcal{A}$, $\psi_c^\top \ell \leq \max\{\psi_a^\top \ell, \psi_b^\top \ell\} \Rightarrow c \in \mathcal{N}_{a,b}$, which established the sought implication. $\square$

C Auxiliary Results

We will use $\mathcal{S}_{\succeq}^n$ in the following to denote the set of (symmetric) positive semi-definite matrices in $\mathbb{R}^{n \times n}$, while $\succeq$ will refer to the Loewner order: for any $F, G \in \mathcal{S}_{\succeq}^n$, $F \succeq G$ is equivalent to $F - G \in \mathcal{S}_{\succeq}^n$.

Lemma 6. *Let n and m be two positive integers, and let $\mathcal{V}$ be a subspace of $\mathbb{R}^n$. Then, the mapping $(X, G) \mapsto X^\top G^\dagger X$ is convex on $\mathbb{R}^{n \times m} \times \{G \in \mathcal{S}_{\succeq}^n \mid \text{col}(G) = \mathcal{V}\}$ with respect to the Loewner order.*

Proof. This is a direct consequence of Theorem 8 in [45]. $\square$

Lemma 7. *Let n and m be two positive integers. For $F, G \in \mathcal{S}_{\succeq}^n$ and $X \in \mathbb{R}^{n \times m}$ such that $\text{col}(X) \subseteq \text{col}(G)$, it holds that*

$$X^\top (F + G)^\dagger X \preceq X^\top G^\dagger X.$$

Proof. This is a direct consequence of Theorem 5 in [12]. $\square$

Recall that for $\pi \in \Delta_k$ and $\delta \in (0, 1)$,

$$Q(\pi) := \sum_{a \in \mathcal{A}} \pi(a) M_a M_a^\top \quad \text{and} \quad Q_\delta(\pi) := Q((1 - \delta)\pi + \delta \mathbf{1}_k/k).$$

Recall also that $U \in \mathbb{R}^{d \times \text{rk}(\mathbf{M})}$ is a matrix whose columns $(u_i)_{i \in [\text{rk}(\mathbf{M})]}$ form an orthonormal basis for $\text{col}(\mathbf{M})$, with $\mathbf{M}$ being the $d \times (\sum_a n(a))$ matrix obtained by horizontally stacking all the observations matrices $(M_a)_{a \in \mathcal{A}}$. Now, for $\pi \in \Delta_k$ and $\delta \in (0, 1)$, define

$$B(\pi) := \sum_{a \in \mathcal{A}} \pi(a) U^\top M_a M_a^\top U = U^\top Q(\pi) U \quad \text{and} \quad B_\delta(\pi) := B((1 - \delta)\pi + \delta \mathbf{1}_k/k). \quad (8)$$

Finally, recall that given a finite index set Z and a collection of matrices $\{X_z \mid z \in Z\}$, all with the same number of rows, we denote by $\text{span}\{X_z \mid z \in Z\}$ the span of all their columns.

Lemma 8. *For any $\pi \in \Delta_k$, $\text{col}(Q(\pi)) = \text{span}\{M_a \mid \pi(a) > 0\}$.*

Proof. Let $n := \sum_{a \in \mathcal{A}} n_a$, v_π be the n -dimensional vector constructed by concatenating the vectors $\sqrt{\pi(a)} \mathbf{1}_{n(a)}$ in order, and D_π be the $n \times n$ diagonal matrix with v_π on the diagonal. Note then that $Q(\pi)$ can be written as $\mathbf{M} D_\pi^2 \mathbf{M}^\top$. Hence,

$$\text{col}(Q(\pi)) = \text{col}(\mathbf{M} D_\pi^2 \mathbf{M}^\top) = \text{col}(\mathbf{M} D_\pi) = \text{span}\{M_a \mid \pi(a) > 0\}.$$

$\square$

Lemma 9. *For any $\pi \in \Delta_k$, the following conditions are equivalent.*

- i) $\text{span}\{M_a \mid \pi(a) > 0\} = \text{col}(\mathbf{M})$;
- ii) $\text{col}(Q(\pi)) = \text{col}(\mathbf{M})$;
- iii) $B(\pi)$ is positive definite;
- iv) $\text{span}\{U^\top M_a \mid \pi(a) > 0\} = \mathbb{R}^{\text{rk}(\mathbf{M})}$.

Proof.

i) $\Leftrightarrow$ ii) Follows directly from Lemma 8.

i) $\Rightarrow$ iii) For any non-zero $v \in \mathbb{R}^{\text{rk}(\mathbf{M})}$,

$$v^\top B(\pi) v = v^\top U^\top Q(\pi) U v = (Uv)^\top Q(\pi) (Uv) = \sum_{a \in \mathcal{A}} \pi(a) \|M_a^\top U v\|^2.$$

Note that $Uv \neq 0$ since U has full rank, and $Uv \in \text{col}(\mathbf{M})$ by the definition of U . Hence, thanks to Condition *i*), $M_a^\top Uv \neq 0$ for some $a \in \mathcal{A}$ with $\pi(a) > 0$, yielding that $\sum_{a \in \mathcal{A}} \pi(a) \|M_a^\top Uv\|^2 > 0$.
iii) $\Rightarrow$ *iv*) Applying an analogous argument to the one used in the proof of Lemma 8, one can show that

$$\text{col}(B(\pi)) = \text{span}\{U^\top M_a \mid \pi(a) > 0\}.$$

The required property then follows since $\text{col}(B(\pi)) = \mathbb{R}^{\text{rk}(\mathbf{M})}$ by the assumption that $B(\pi)$ is positive definite.

iv) $\Rightarrow$ *i*) For any $x \in \mathbb{R}^d$, Condition *iv*) implies that $U^\top x \in \mathbb{R}^{\text{rk}(\mathbf{M})}$ can be written as $U^\top x = \sum_{a \in \mathcal{A}} U^\top M_a v_a$ for some collection of vectors $(v_a)_{a \in \mathcal{A}}$ such that $v_a \in \mathbb{R}^{n(a)}$ and $v_a = \mathbf{0}$ whenever $\pi(a) = 0$. Now, take x to be in $\text{col}(\mathbf{M})$. Since $\text{col}(M_a) \in \text{col}(\mathbf{M})$ and $UU^\top$ is the projection matrix onto $\text{col}(\mathbf{M})$, we obtain that

$$x = UU^\top x = \sum_{a \in \mathcal{A}} UU^\top M_a v_a = \sum_{a \in \mathcal{A}} M_a v_a.$$

This shows that $\text{col}(\mathbf{M}) \subseteq \text{span}\{M_a \mid \pi(a) > 0\}$. The required property then readily follows as the converse statement is trivial. $\square$

Lemma 10. Fix some $\pi \in \Delta_k$ such that $B(\pi)$ is positive definite. Then, $U B(\pi)^{-1} U^\top = Q(\pi)^\dagger$.

Proof. For brevity, let $Q := Q(\pi)$, $B := B(\pi)$, and $X := UB^{-1}U^\top$. Also, recall that $B = U^\top QU$. Since U has orthonormal columns, $U^\top U = \mathbf{I}_{\text{rk}(\mathbf{M})}$ and $UU^\top$ is the orthogonal projection matrix onto $\text{col}(U)$, which coincides with $\text{col}(\mathbf{M})$. Additionally, note that the row and column spaces of Q are the same thanks to symmetry, both of which also coincide with $\text{col}(\mathbf{M})$ via Lemma 9. Hence,

$$UBU^\top = UU^\top QUU^\top = Q.$$

We now prove the sought identity, that $X = Q^\dagger$, by verifying the Moore–Penrose conditions:

$$\begin{aligned} QXQ &= UBU^\top UB^{-1}U^\top UBU^\top = UBU^\top = Q, \\ XQX &= UB^{-1}U^\top UBU^\top UB^{-1}U^\top = UB^{-1}U^\top = X, \\ (QX)^\top &= (UU^\top)^\top = UU^\top = QX, \\ (XQ)^\top &= (UU^\top)^\top = UU^\top = XQ. \end{aligned}$$

$\square$

Lemma 11. For any $\pi \in \Delta_k$ and $\delta \in (0, 1)$, it holds that $\text{col}(Q_\delta(\pi)) = \text{col}(\mathbf{M})$, $B_\delta(\pi)$ is positive definite, and $U B_\delta(\pi)^{-1} U^\top = Q_\delta(\pi)^\dagger$.

Proof. As $(1 - \delta)\pi + \delta \mathbf{1}_k/k$ has full support, the requirements follow from Lemma 9, Lemma 10, and the definitions of $Q_\delta(\pi)$ and $B_\delta(\pi)$. $\square$

Lemma 12 (Extension of Theorem 21.1 in [33]). Define the functions $f, g: \Delta_k \rightarrow [-\infty, +\infty]$ as

$$f(\pi) := \log \det B(\pi) \quad \text{and} \quad g(\pi) := \max_{a \in \mathcal{A}} \text{tr}(M_a^\top UB(\pi)^{-1} U^\top M_a).$$

Then, for any $\pi^* \in \Delta_k$, the following are equivalent:

- i) π^* is a minimizer of g ;
- ii) π^* is a maximizer of f ;
- iii) $g(\pi^*) = \text{rk}(\mathbf{M})$.

Proof. For brevity, let $r := \text{rk}(\mathbf{M})$ and $g_a(\pi) := \text{tr}(M_a^\top U B(\pi)^{-1} U^\top M_a)$ for $a \in \mathcal{A}$. When $B(\pi)$ is ill-conditioned, f and g map to $-\infty$ and $+\infty$ respectively. It can be verified via standard arguments that f and g indeed attain their maximum and minimum respectively on Δ_k .

Via the definition of B (see (8)) and the fact that $d(\log \det X) = \text{tr}(X^{-1} dX)$, we have that

$$(\nabla f(\pi))_a = \text{tr}(B(\pi)^{-1} U^\top M_a M_a^\top U) = \text{tr}(M_a^\top U B(\pi)^{-1} U^\top M_a) = g_a(\pi).$$

Using the linearity of the trace, we get that

$$\sum_{a \in \mathcal{A}} \pi(a) (\nabla f(\pi))_a = \sum_{a \in \mathcal{A}} \pi(a) g_a(\pi) = \text{tr}\left(B(\pi)^{-1} \sum_{a \in \mathcal{A}} \pi(a) U^\top M_a M_a^\top U\right) = \text{tr}(\mathbf{I}_r) = r.$$

ii) $\Rightarrow$ i) By the concavity of f , the first-order optimality condition gives that for any $\pi \in \Delta_k$,

$$0 \geq \langle \nabla f(\pi^*), \pi - \pi^* \rangle = \sum_a \pi(a) g_a(\pi^*) - r.$$

Taking π as a Dirac over some $a \in \mathcal{A}$ gives that $g_a(\pi^*) \leq r$; hence, $g(\pi^*) \leq r$. However, for any $\pi \in \Delta_k$, $g(\pi) \geq \sum_{a \in \mathcal{A}} \pi(a) g_a(\pi) = r$, so π^* minimizes g and $\min_{\pi \in \Delta_k} g(\pi) = r$.

iii) $\Rightarrow$ ii) If $g(\pi^*) = r$, then $g_a(\pi^*) \leq r$ for every $a \in \mathcal{A}$; hence,

$$\langle \nabla f(\pi^*), \pi - \pi^* \rangle = \sum_{a \in \mathcal{A}} \pi(a) g_a(\pi^*) - r \leq 0$$

for all $\pi \in \Delta_k$. The concavity of f therefore implies that π^* maximizes f .

i) $\Rightarrow$ iii) From the first part we know that $\min_{\pi \in \Delta_k} g(\pi) = r$; hence, π^* , being a minimizer of g , satisfies $g(\pi^*) = r$. $\square$

Lemma 13. For any $p, q \in \Delta_k$, it holds that

$$\sum_{a, b \in \mathcal{A}} q(a) q(b) (\psi_a - \psi_b)^\top Q_\delta(p)^\dagger (\psi_a - \psi_b) = 2 \sum_{a \in \mathcal{A}} q(a) (\psi_a - Hq)^\top Q_\delta(p)^\dagger (\psi_a - Hq).$$

Proof. Note that via the definition of H , we have that $Hq = \sum_{b \in \mathcal{A}} q(b) \psi_b$. For brevity, let $X := Q_\delta(p)^\dagger$. It is also noteworthy that the only relevant property of $Q_\delta(p)^\dagger$ here is its symmetry. Starting with the left-hand-side, we have that

$$\begin{aligned} \sum_{a, b \in \mathcal{A}} q(a) q(b) (\psi_a - \psi_b)^\top X (\psi_a - \psi_b) &= 2 \sum_{a \in \mathcal{A}} q(a) \psi_a^\top X \psi_a - 2 \sum_{a, b \in \mathcal{A}} q(a) q(b) \psi_a^\top X \psi_b \\ &= 2 \sum_{a \in \mathcal{A}} q(a) \psi_a^\top X \psi_a - 2(Hq)^\top X Hq. \end{aligned}$$

where the equivalence of the two cross terms in the initial expansion follows from the symmetry of X . Next, expanding the right-hand-side of the sought equality yields that

$$\begin{aligned} 2 \sum_{a \in \mathcal{A}} q(a) (\psi_a - Hq)^\top X (\psi_a - Hq) &= 2 \sum_{a \in \mathcal{A}} q(a) \psi_a^\top X \psi_a - 4 \sum_{a \in \mathcal{A}} q(a) \psi_a^\top X Hq + 2(Hq)^\top X Hq \\ &= 2 \sum_{a \in \mathcal{A}} q(a) \psi_a^\top X \psi_a - 2(Hq)^\top X Hq, \end{aligned}$$

which concludes the proof. $\square$

Lemma 14. For any $p, q \in \Delta_k$, it holds that

$$\sum_{a \in \mathcal{A}} q(a) (\psi_a - Hq)^\top Q_\delta(p)^\dagger (\psi_a - Hq) = \min_{r \in \Delta_k} \sum_{a \in \mathcal{A}} q(a) (\psi_a - Hr)^\top Q_\delta(p)^\dagger (\psi_a - Hr).$$

Proof. Define $F: \Delta_k \rightarrow \mathbb{R}$ such that $F(r) = \sum_{a \in \mathcal{A}} q(a) (\psi_a - Hr)^\top Q_\delta(p)^\dagger (\psi_a - Hr)$, which is easily seen to be convex ($Q_\delta(p)^\dagger$ is positive semi-definite). Its gradient is given by

$$\nabla F(r) = 2H^\top Q_\delta(p)^\dagger Hr - 2H^\top Q_\delta(p)^\dagger Hq.$$

This clearly vanishes when $r = q$; hence, via the convexity of F , it is minimized at q . $\square$

D Proofs of Section 3

As a starting point for the analysis, the following lemma states a standard result for exponential weights. A proof can be easily extracted, for example, from the proof of Theorem 1.5 in [26].

Lemma 15. *Assuming that at every round t , $\eta \hat{y}_t(a) \geq -1$ holds for every action a ; then, for any $a^* \in \mathcal{A}^*$, the sequence of predictions $(q_t)_t$ defined by (4) satisfies*

$$\sum_{t=1}^T \langle q_t - e_{a^*}, \hat{y}_t \rangle \leq \frac{\log k}{\eta} + \eta \sum_{t=1}^T \langle q_t, \hat{y}_t^2 \rangle.$$

Lemma 1. *Fix a learning policy and define p_t as the law of A_t conditioned on $\mathcal{F}_{t-1}$. Let q_t be as given in (4) for some learning rate $\eta > 0$ and sequence of surrogate loss functions $(\hat{y}_t)_t$ such that $\hat{y}_t \in \mathbb{R}^k$ is $\mathcal{F}_t$ -measurable and satisfies $\max_{a \in \mathcal{A}} |\eta \hat{y}_t(a)| \leq 1$. Further, let $a^* \in \arg \min_{a \in \mathcal{A}^*} \sum_{t=1}^T y_t(a)$. Then, the regret of the policy satisfies*

$$R_T \leq \frac{\log k}{\eta} + \sum_{t=1}^T \mathbb{E} [\langle p_t - q_t, y_t \rangle + \eta \langle q_t, \mathbb{E}_t \hat{y}_t^2 \rangle + \langle q_t - e_{a^*}, y_t - \mathbb{E}_t \hat{y}_t \rangle].$$

Proof. Firstly, we have that

$$\begin{aligned} R_T &= \mathbb{E} \left[\sum_{t=1}^T y_t(A_t) - y_t(a^*) \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T \langle p_t - e_{a^*}, y_t \rangle \right] \\ &= \mathbb{E} \left[\sum_{t=1}^T \langle p_t - q_t, y_t \rangle + \sum_{t=1}^T \langle q_t - e_{a^*}, \hat{y}_t \rangle + \sum_{t=1}^T \langle q_t - e_{a^*}, y_t - \hat{y}_t \rangle \right], \end{aligned}$$

where the second equality follows from the definition of p_t , the linearity of expectation, and the tower rule. Next, we apply Lemma 15 to the middle term ($\eta \hat{y}_t(a) \geq -1$ holds by assumption) to get that

$$\begin{aligned} R_T &\leq \frac{\log k}{\eta} + \mathbb{E} \left[\sum_{t=1}^T \langle p_t - q_t, y_t \rangle + \eta \sum_{t=1}^T \langle q_t, \hat{y}_t^2 \rangle + \sum_{t=1}^T \langle q_t - e_{a^*}, y_t - \hat{y}_t \rangle \right] \\ &= \frac{\log k}{\eta} + \sum_{t=1}^T \mathbb{E} [\langle p_t - q_t, y_t \rangle + \eta \langle q_t, \mathbb{E}_t \hat{y}_t^2 \rangle + \langle q_t - e_{a^*}, y_t - \mathbb{E}_t \hat{y}_t \rangle], \end{aligned}$$

where the equality is another application of the linearity of expectation and the tower rule, using the fact that q_t is $\mathcal{F}_{t-1}$ measurable as it only depends on $(\hat{y}_s)_{s=1}^{t-1}$. $\square$

Lemma 2. *In the same setting as Lemma 1, let the predictions of the policy satisfy $p_t = (1 - \delta) \tilde{p}_t + \delta \mathbf{1}_k / k$ for some $\delta \in (0, 1)$ and $\tilde{p}_t \in \Delta_k$, and let the surrogate loss functions satisfy $\hat{y}_t = g_t(A_t, \phi_t)$ where $g_t(a, \phi) = (\mathbf{I}_k - \mathbf{1}_k q_t^\top) H^\top Q(p_t)^\dagger M_a \phi$. Then, assuming global observability, it holds that*

$$R_T \leq \frac{\log k}{\eta} + 2\delta T + \sum_{t=1}^T \mathbb{E} \left[\langle \tilde{p}_t - q_t, H^\top \ell_t \rangle + \eta \omega^2 \sum_{a,b \in \mathcal{A}} q_t(a) q_t(b) (\psi_a - \psi_b)^\top Q_\delta(\tilde{p}_t)^\dagger (\psi_a - \psi_b) \right]$$

provided that $\max_{a,b \in \mathcal{A}} (\psi_a - \psi_b)^\top Q_\delta(\tilde{p}_t)^\dagger (\psi_a - \psi_b) + \max_{c \in \mathcal{A}} \|M_c^\top Q_\delta(\tilde{p}_t)^\dagger M_c\|_2 \leq \frac{2}{\eta \omega}$, where

$$\omega := \sup_{a,b,c \in \mathcal{A}, p \in \Delta_k, \ell \in \mathcal{L}} \frac{|(\psi_b - \psi_c)^\top Q_\delta(p)^\dagger M_a M_a^\top \ell|}{\|M_a^\top Q_\delta(p)^\dagger (\psi_b - \psi_c)\|} \leq \max_{a \in \mathcal{A}, \ell \in \mathcal{L}} \|M_a^\top \ell\|.$$

Proof. Fix some $a^* \in \mathcal{A}^*$. If $\|g_t(A_t, \phi_t)\|_\infty \leq \eta^{-1}$ holds in all rounds, Lemma 1 gives that $R_T \leq \eta^{-1} \log k + \sum_{t=1}^T \mathbb{E} \Gamma_t$, where

$$\Gamma_t := \langle p_t - q_t, H^\top \ell_t \rangle + \eta \langle q_t, \sum_a p_t(a) g_t(a, M_a^\top \ell_t)^2 \rangle + \langle q_t - e_{a^*}, H^\top \ell_t - \sum_a p_t(a) g_t(a, M_a^\top \ell_t) \rangle. \quad (9)$$

Concerning the third term in (9), we have that

$$\begin{aligned}
\sum_a p_t(a) g_t(a, M_a^\top \ell_t) &= (\mathbf{I}_k - \mathbf{1}_k q_t^\top) H^\top Q(p_t)^\dagger \sum_a p_t(a) M_a M_a^\top \ell_t \\
&= (\mathbf{I}_k - \mathbf{1}_k q_t^\top) H^\top Q(p_t)^\dagger Q(p_t) \ell_t \\
&= (\mathbf{I}_k - \mathbf{1}_k q_t^\top) H^\top U B_\delta(\tilde{p}_t)^{-1} U^\top U B_\delta(\tilde{p}_t) U^\top \ell_t \\
&= (\mathbf{I}_k - \mathbf{1}_k q_t^\top) H^\top U U^\top \ell_t \\
&= (\mathbf{I}_k - \mathbf{1}_k q_t^\top) H^\top \ell_t,
\end{aligned}$$

where U and $B_\delta(\cdot)$ are defined in Appendix C (see (8)), and the last equality holds since each row in $(\mathbf{I}_k - \mathbf{1}_k q_t^\top) H^\top$ is a convex combination of vectors of the form $\psi_a - \psi_b$ with $a, b \in \mathcal{A}$, which, via global observability, belong to $\text{col}(M)$, onto which $U U^\top$ projects. Hence, we get that

$$\begin{aligned}
\langle q_t - e_{a^*}, H^\top \ell_t - \sum_a p_t(a) g_t(a, M_a^\top \ell_t) \rangle &= \langle q_t - e_{a^*}, H^\top \ell_t - (\mathbf{I}_k - \mathbf{1}_k q_t^\top) H^\top \ell_t \rangle \\
&= \langle q_t - e_{a^*}, \mathbf{1}_k q_t^\top H^\top \ell_t \rangle \\
&= 0
\end{aligned}$$

using in the last equality that $q_t, e_{a^*} \in \Delta_k$ and all coordinates of $\mathbf{1}_k q_t^\top H^\top \ell_t$ are identical. In what follows, we will refer to the individual coordinates of $g_t(\cdot, \cdot)$ as $g_t(b; \cdot, \cdot)$ for $b \in \mathcal{A}$. Shifting to the second term, fix $b \in \mathcal{A}$ and observe that

$$\begin{aligned}
\sum_a p_t(a) g_t(b; a, M_a^\top \ell_t)^2 &= \sum_a p_t(a) ((\psi_b - H q_t)^\top Q_\delta(\tilde{p}_t)^\dagger M_a M_a^\top \ell_t)^2 \\
&= \sum_a p_t(a) (q_t^\top (\psi_b \mathbf{1}_k^\top - H)^\top Q_\delta(\tilde{p}_t)^\dagger M_a M_a^\top \ell_t)^2 \\
&\leq \sum_a p_t(a) \sum_c q_t(c) ((\psi_b - \psi_c)^\top Q_\delta(\tilde{p}_t)^\dagger M_a M_a^\top \ell_t)^2 \\
&\leq \omega^2 \sum_a p_t(a) \sum_c q_t(c) \|M_a^\top Q_\delta(\tilde{p}_t)^\dagger (\psi_b - \psi_c)\|^2 \\
&= \omega^2 \sum_a p_t(a) \sum_c q_t(c) (\psi_b - \psi_c)^\top Q_\delta(\tilde{p}_t)^\dagger M_a M_a^\top Q_\delta(\tilde{p}_t)^\dagger (\psi_b - \psi_c) \\
&= \omega^2 \sum_c q_t(c) (\psi_b - \psi_c)^\top Q_\delta(\tilde{p}_t)^\dagger Q_\delta(\tilde{p}_t) Q_\delta(\tilde{p}_t)^\dagger (\psi_b - \psi_c) \\
&= \omega^2 \sum_c q_t(c) (\psi_b - \psi_c)^\top Q_\delta(\tilde{p}_t)^\dagger (\psi_b - \psi_c),
\end{aligned}$$

where the first inequality is an application of Jensen's inequality, while the second inequality follows from the definition of ω .

As for the first term in (9), it is immediate that

$$\langle p_t - q_t, H^\top \ell_t \rangle = \delta \langle \mathbf{1}_k/k - q_t, H^\top \ell_t \rangle + (1 - \delta) \langle \tilde{p}_t - q_t, H^\top \ell_t \rangle \leq 2\delta + \langle \tilde{p}_t - q_t, H^\top \ell_t \rangle,$$

where we used the definition of p_t and the assumptions that $\max_{a, b \in \mathcal{A}, \ell \in \mathcal{L}} |(\psi_a - \psi_b)^\top \ell| \leq 2$ and $\delta \in (0, 1)$.

Finally, we examine the boundedness condition. Fixing $a, b \in \mathcal{A}$, we have that

$$\begin{aligned}
|g_t(b; a, M_a^\top \ell_t)| &= |(\psi_b - H q_t)^\top Q_\delta(\tilde{p}_t)^\dagger M_a M_a^\top \ell_t| \\
&= |q_t^\top (\psi_b \mathbf{1}_k^\top - H)^\top Q_\delta(\tilde{p}_t)^\dagger M_a M_a^\top \ell_t| \\
&\leq \sum_c q_t(c) |(\psi_b - \psi_c)^\top Q_\delta(\tilde{p}_t)^\dagger M_a M_a^\top \ell_t| \\
&\leq \omega \sum_c q_t(c) \|M_a^\top (Q_\delta(\tilde{p}_t)^\dagger)^{1/2} (Q_\delta(\tilde{p}_t)^\dagger)^{1/2} (\psi_b - \psi_c)\| \\
&\leq \omega \sum_c q_t(c) \|M_a^\top (Q_\delta(\tilde{p}_t)^\dagger)^{1/2}\| \| (Q_\delta(\tilde{p}_t)^\dagger)^{1/2} (\psi_b - \psi_c) \| \\
&\leq \frac{\omega}{2} \sum_c q_t(c) \left(\|M_a^\top (Q_\delta(\tilde{p}_t)^\dagger)^{1/2}\|^2 + \| (Q_\delta(\tilde{p}_t)^\dagger)^{1/2} (\psi_b - \psi_c) \|^2 \right) \\
&\leq \frac{\omega}{2} \left(\|M_a^\top Q_\delta(\tilde{p}_t)^\dagger M_a\| + \max_{c \in \mathcal{A}} (\psi_b - \psi_c)^\top Q_\delta(\tilde{p}_t)^\dagger (\psi_b - \psi_c) \right),
\end{aligned}$$

where the first and second inequalities again follow from Jensen's inequality and the definition of ω respectively. The required result now follows from Lemma 1. $\square$

Lemma 3. *For any $\eta > 0$ and $q \in \Delta_k$, it holds that*

$$\Lambda_{\eta, q}^* = \max_{\ell \in \mathcal{L}} \min_{p \in \Xi_\eta} \Lambda_{\eta, q}(p, \ell).$$

Proof. Recall that

$$\Lambda_{\eta,q}(p, \ell) := \frac{1}{\eta} \langle p - q, H^\top \ell \rangle + L^2 \sum_{a,b \in \mathcal{A}} q(a)q(b) \mathcal{E}(a, b; p)$$

for $p \in \Delta_k$, and $\ell \in \mathcal{L}$; and that for any $a, b \in \mathcal{A}$,

$$\mathcal{E}(a, b; p) := (\psi_a - \psi_b)^\top Q_\delta(p)^\dagger (\psi_a - \psi_b).$$

Also, recall that $\Xi_\eta := \{p \in \Delta_k : z(p) \leq 2/\eta L\}$, where

$$z(p) := \max_{a,b \in \mathcal{A}} \mathcal{E}(a, b; p) + \max_{c \in \mathcal{A}} \|M_c^\top Q_\delta(p)^\dagger M_c\|_2,$$

and that $\Lambda_{\eta,q}^* := \min_{p \in \Xi_\eta} \max_{\ell \in \mathcal{L}} \Lambda_{\eta,q}(p, \ell)$.

We start by showing that for a fixed pair $(a, b) \in \mathcal{A}^2$, $\mathcal{E}(a, b; p)$ is convex in p . Via Lemma 11 in Appendix C, $\text{col}(Q_\delta(p)) = \text{col}(\mathbf{M})$ for any $p \in \Delta_k$. Hence, the convexity of $\mathcal{E}$ in p follows from Lemma 6 in Appendix C and the fact that $Q_\delta(p)$ is affine in p . Moreover, for a fixed $c \in \mathcal{A}$, the function $p \mapsto \|M_c^\top Q_\delta(p)^\dagger M_c\|_2$ can also be easily shown to be convex by invoking once again Lemma 6 and using that $Q_\delta(p)$ is affine in p and that the spectral norm is convex and non-decreasing in the Loewner order. Also note that both functions are continuous as the rank of $Q_\delta(p)$ is constant for any $p \in \Delta_k$ (see Corollary 3.5 in [50]), or simply via the continuity of matrix inversion since $Q_\delta(p)^\dagger = U B_\delta(p)^{-1} U^\top$ and $B_\delta(p)$ (defined in (8)) is positive definite for any $p \in \Delta_k$ as asserts Lemma 11 in Appendix C. Hence, the function z is continuous and convex in $p \in \Delta_k$, and its sub-level set $\Xi_\eta := \{p \in \Delta_k : z(p) \leq 2/\eta L\}$ is convex and compact.

The arguments above imply that for a fixed $\ell \in \mathcal{L}$, $p \mapsto \Lambda_{\eta,q}(p, \ell)$ is continuous and convex over the compact set Ξ_η . Moreover, for a fixed $p \in \Delta_k$, $\ell \mapsto \Lambda_{\eta,q}(p, \ell)$ is trivially continuous and concave over the compact set $\mathcal{L}$. Therefore, Sion's minimax theorem [49] asserts that

$$\min_{p \in \Xi_\eta} \max_{\ell \in \mathcal{L}} \Lambda_{\eta,q}(p, \ell) = \max_{\ell \in \mathcal{L}} \min_{p \in \Xi_\eta} \Lambda_{\eta,q}(p, \ell).$$

The theorem then follows from the definition of $\Lambda_{\eta,q}^*$. $\square$

E Proofs of Section 4

Lemma 4. *A game is locally observable if and only if it holds for all $a \in \mathcal{A}^*$, $\ell \in \mathbb{R}^d$, and $a^* \in \arg \min_{a' \in \mathcal{A}} \psi_{a'}^\top \ell$ that $\psi_a - \psi_{a^*} \in \text{span} \{M_b \mid \psi_b^\top \ell \leq \psi_{a^*}^\top \ell\}$.*

Proof. This follows from Lemma 5 in Appendix B. $\square$

Before moving on to the proof of Theorem 1, we define an analogous set to $\mathcal{W}^{\text{loc}}$ that does not take the losses into account. Recall that global observability implies that for any pair of actions (a, b) , there exists a (not necessarily unique) weight vector $\mathbf{v} \in \mathbb{R}^n$ such that $\psi_a - \psi_b = \mathbf{M}\mathbf{v}$. The following is the set of all possible action-pair-wise assignments of such weight vectors:

$$\mathcal{W}^{\text{glo}} := \{\boldsymbol{\xi} = (\boldsymbol{\xi}_{a,b})_{a,b \in \mathcal{A}} \in \mathbb{R}^{k^2 \times n} \mid \forall a, b \in \mathcal{A}, \psi_a - \psi_b = \mathbf{M}\boldsymbol{\xi}_{a,b}\}.$$

Recall that

$$\beta_{\text{glo}} := \max_{a,b \in \mathcal{A}} \min_{\mathbf{v} \in \mathbb{R}^n : \psi_a - \psi_b = \mathbf{M}\mathbf{v}} \sum_{c \in \mathcal{A}} \|\mathbf{v}(c)\|.$$

The following lemma provides a slightly different characterization of β_{glo} featuring $\mathcal{W}^{\text{glo}}$.

Lemma 16. *It holds that*

$$\beta_{\text{glo}} = \min_{\boldsymbol{\xi} \in \mathcal{W}^{\text{glo}}} \max_{a,b \in \mathcal{A}} \sum_{c \in \mathcal{A}} \|\boldsymbol{\xi}_{a,b}(c)\|.$$

Proof. For $a, b \in \mathcal{A}$, let $\mathbf{v}_{a,b}^*$ refer to an arbitrary member of $\arg \min_{\mathbf{v} \in \mathbb{R}^n : \psi_a - \psi_b = \mathbf{M}\mathbf{v}} \sum_{c \in \mathcal{A}} \|\mathbf{v}(c)\|$. Via the definition of $\mathcal{W}^{\text{glo}}$, $(\mathbf{v}_{a,b}^*)_{a,b \in \mathcal{A}} \in \mathcal{W}^{\text{glo}}$; moreover, it holds that $(\mathbf{v}_{a,b}^*)_{a,b \in \mathcal{A}} \in \arg \min_{\boldsymbol{\xi} \in \mathcal{W}^{\text{glo}}} \sum_{c \in \mathcal{A}} \|\boldsymbol{\xi}_{a,b}(c)\|$ for all $a, b \in \mathcal{A}$ simultaneously. This implies that $(\mathbf{v}_{a,b}^*)_{a,b \in \mathcal{A}} \in \arg \min_{\boldsymbol{\xi} \in \mathcal{W}^{\text{glo}}} \max_{a,b \in \mathcal{A}} \sum_{c \in \mathcal{A}} \|\boldsymbol{\xi}_{a,b}(c)\|$, concluding the proof. $\square$

The following lemma relates β_{glo} and β_{loc} .

Lemma 17. *It holds that $\beta_{\text{glo}} \leq \beta_{\text{loc}}$.*

Proof. Recall that

$$\beta_{\text{loc}} := \max_{\ell \in \mathcal{L}} \min_{\lambda \in \mathcal{W}_{\ell}^{\text{loc}}} \max_{a \in \mathcal{A}^*} \sum_{c \in \mathcal{A}} \|\lambda_a(c)\|,$$

where

$$\mathcal{W}_{\ell}^{\text{loc}} := \left\{ \lambda = (\lambda_a)_{a \in \mathcal{A}^*} \in \mathbb{R}^{k^* \times n} \mid \exists a^* \in \arg \min_{c \in \mathcal{A}} \psi_c^{\top} \ell \ \forall (a, b) \in \mathcal{A}^* \times \mathcal{A}, \right. \\ \left. \psi_a - \psi_{a^*} = M\lambda_a \text{ and } \lambda_a(b) \neq \mathbf{0} \Rightarrow \psi_b^{\top} \ell \leq \psi_{a^*}^{\top} \ell \right\}.$$

Additionally, define

$$\mathcal{W}_{\ell}^{\text{glo}} := \left\{ \lambda = (\lambda_a)_{a \in \mathcal{A}^*} \in \mathbb{R}^{k^* \times n} \mid \exists a^* \in \arg \min_{c \in \mathcal{A}} \psi_c^{\top} \ell \ \forall a \in \mathcal{A}^*, \psi_a - \psi_{a^*} = M\lambda_a \right\}$$

for $\ell \in \mathcal{L}$ and, with some notation abuse,

$$\mathcal{W}_{a'}^{\text{glo}} := \left\{ \lambda = (\lambda_a)_{a \in \mathcal{A}^*} \in \mathbb{R}^{k^* \times n} \mid \forall a \in \mathcal{A}^*, \psi_a - \psi_{a'} = M\lambda_a \right\}$$

for $a' \in \mathcal{A}^*$. Then,

$$\begin{aligned} \beta_{\text{loc}} &= \max_{\ell \in \mathcal{L}} \min_{\lambda \in \mathcal{W}_{\ell}^{\text{loc}}} \max_{a \in \mathcal{A}^*} \sum_{c \in \mathcal{A}} \|\lambda_a(c)\| \\ &\geq \max_{\ell \in \mathcal{L}} \min_{\lambda \in \mathcal{W}_{\ell}^{\text{glo}}} \max_{a \in \mathcal{A}^*} \sum_{c \in \mathcal{A}} \|\lambda_a(c)\| \\ &\geq \max_{b \in \mathcal{A}^*} \min_{\lambda \in \mathcal{W}_b^{\text{glo}}} \max_{a \in \mathcal{A}^*} \sum_{c \in \mathcal{A}} \|\lambda_a(c)\| \\ &\geq \max_{b \in \mathcal{A}^*} \max_{a \in \mathcal{A}^*} \min_{\lambda \in \mathcal{W}_b^{\text{glo}}} \sum_{c \in \mathcal{A}} \|\lambda_a(c)\| \\ &= \max_{a, b \in \mathcal{A}^*} \min_{v \in \mathbb{R}^n : \psi_a - \psi_b = Mv} \sum_{c \in \mathcal{A}} \|v(c)\|, \end{aligned}$$

where the first inequality holds since $\mathcal{W}_{\ell}^{\text{loc}} \subseteq \mathcal{W}_{\ell}^{\text{glo}}$. Concerning the second inequality, as argued in the proof of Lemma 5, it holds for any Pareto optimal action b that $\{\psi_a\} \cap \text{co}(\mathcal{X} \setminus \{\psi_a\}) = \emptyset$ (since $\psi_b \in \text{ext}(\mathcal{V})$); hence, by the separating hyperplane theorem, there exists a loss vector $\ell_b \in \mathbb{R}^d$ such that $\psi_b^{\top} \ell_b < x^{\top} \ell_b$ for all $x \in \mathcal{X} \setminus \{\psi_b\}$, yielding that only b and its duplicates can minimize $c \mapsto \psi_c^{\top} \ell_b$ over $\mathcal{A}$. The inequality then follows by restricting the maximization to any selection of loss vectors of the form $(\ell_b)_{b \in \mathcal{A}^*}$ and using the fact that $\mathcal{W}_{\ell_b}^{\text{glo}} = \mathcal{W}_b^{\text{glo}}$.

On the other hand, we have that

$$\beta_{\text{glo}} = \max_{a, b \in \mathcal{A}} \min_{v \in \mathbb{R}^n : \psi_a - \psi_b = Mv} \sum_{c \in \mathcal{A}} \|v(c)\|;$$

hence, showing that the maximum on the right-hand side of the equation above is attained by a pair $a, b \in \mathcal{A}^*$ suffices to conclude the proof. Let $a, b \in \mathcal{A}$ be two arbitrary actions. By the definition of $\mathcal{A}^*$, there exist two probability distributions $p_1, p_2 \in \Delta_{k^*}$ such that $\psi_a = \sum_{h \in \mathcal{A}^*} p_1(h) \psi_h$ and $\psi_b = \sum_{h \in \mathcal{A}^*} p_2(h) \psi_h$. Then,

$$\psi_a - \psi_b = \sum_{h \in \mathcal{A}^*} p_1(h) \psi_h - \sum_{j \in \mathcal{A}^*} p_2(j) \psi_j = \sum_{h, j \in \mathcal{A}^*} p_1(h) p_2(j) (\psi_h - \psi_j),$$

and $\psi_a - \psi_b = \mathbf{M} \sum_{h,j \in \mathcal{A}^*} p_1(h)p_2(j)\mathbf{v}_{h,j}^*$, with $\mathbf{v}_{h,j}^* \in \arg \min_{\mathbf{v} \in \mathbb{R}^n : \psi_h - \psi_j = \mathbf{M}\mathbf{v}} \sum_{c \in \mathcal{A}} \|\mathbf{v}(c)\|$. Hence,

$$\begin{aligned} \min_{\mathbf{v} \in \mathbb{R}^n : \psi_a - \psi_b = \mathbf{M}\mathbf{v}} \sum_{c \in \mathcal{A}} \|\mathbf{v}(c)\| &\leq \sum_{c \in \mathcal{A}} \left\| \sum_{h,j \in \mathcal{A}^*} p_1(h)p_2(j)\mathbf{v}_{h,j}^*(c) \right\| \\ &\leq \sum_{h,j \in \mathcal{A}^*} p_1(h)p_2(j) \sum_{c \in \mathcal{A}} \|\mathbf{v}_{h,j}^*(c)\| \\ &\leq \max_{h,j \in \mathcal{A}^*} \sum_{c \in \mathcal{A}} \|\mathbf{v}_{h,j}^*(c)\| \\ &= \max_{h,j \in \mathcal{A}^*} \min_{\mathbf{v} \in \mathbb{R}^n : \psi_h - \psi_j = \mathbf{M}\mathbf{v}} \sum_{c \in \mathcal{A}} \|\mathbf{v}(c)\|, \end{aligned}$$

where the second inequality is an application of Jensen's inequality. We have thus shown that

$$\beta_{\text{glo}} = \max_{a,b \in \mathcal{A}^*} \min_{\mathbf{v} \in \mathbb{R}^n : \psi_a - \psi_b = \mathbf{M}\mathbf{v}} \sum_{c \in \mathcal{A}} \|\mathbf{v}(c)\| \leq \beta_{\text{loc}}.$$

□

The proof of Theorem 1 relies on the following Proposition. (The definition of $B(\cdot)$ is given in (8).)

Proposition 3. *Suppose that the game is locally observable. Then, for any $\ell \in \mathcal{L}$, $q \in \Delta_k$, $\lambda \in \mathcal{W}_\ell^{\text{loc}}$, and $\xi \in \mathcal{W}^{\text{glo}}$, it holds that*

$$\begin{aligned} \min_{p \in \Xi_\eta} \Lambda_{\eta,q}(p, \ell) &\leq 8L^2 \beta_\lambda^2 \sum_{b \in \mathcal{A}} q^\lambda(b) \|M_b^\top Q(q^\lambda)^\dagger M_b\| \\ &\quad + 2L(1 + \beta_\xi^2) \min_{\pi \in \Delta_k} \max_{c \in \mathcal{A}} \|M_c^\top UB(\pi)^{-1} U^\top M_c\| \end{aligned}$$

provided that $1/\eta \geq 2L(1 + \beta_\xi^2) \min_{\pi \in \Delta_k} \max_{c \in \mathcal{A}} \|M_c^\top UB(\pi)^{-1} U^\top M_c\|$, where

$$\beta_\lambda := \max_{a \in \mathcal{A}^*} \sum_{b \in \mathcal{A}} \|\lambda_a(b)\|, \quad \beta_\xi := \max_{a,b \in \mathcal{A}} \sum_{c \in \mathcal{A}} \|\xi_{a,b}(c)\|, \quad \text{and}$$

$$q^\lambda(b) := q(b) \mathbb{I}[b \in \mathcal{A}^*, \lambda_b = \mathbf{0}] + \underbrace{\sum_{a \in \mathcal{A}^* : \lambda_a \neq \mathbf{0}} q(a) \frac{\|\lambda_a(b)\|}{\sum_{c \in \mathcal{A}} \|\lambda_a(c)\|}}_{=: q_-^\lambda(b)}.$$

Proof. We show this by constructing a certain distribution $p \in \Xi_\eta$ and bounding $\Lambda_{\eta,q}(p, \ell)$. Our choice of p will take the form $p = \gamma\pi + (1 - \gamma)\hat{p}$, with $\gamma \in (0, 1/2]$ and $\pi, \hat{p} \in \Delta_k$. For brevity, let $\bar{\gamma} = 1 - \gamma$. The role of π is to ensure that $\gamma\pi + \bar{\gamma}\hat{p} \in \Xi_\eta$ and that of $\hat{p}$ is to control the magnitude of $\Lambda_{\eta,q}(\gamma\pi + \bar{\gamma}\hat{p}, \ell)$. Let $a^* \in \arg \min_{a' \in \mathcal{A}^*} \langle \psi_{a'}, \ell \rangle$. Now,

$$\begin{aligned} \Lambda_{\eta,q}(\gamma\pi + \bar{\gamma}\hat{p}, \ell) &= \frac{\gamma}{\eta} \langle \pi - q, H^\top \ell \rangle + \frac{\bar{\gamma}}{\eta} \langle \hat{p} - q, H^\top \ell \rangle + L^2 \sum_{a,b \in \mathcal{A}} q(a)q(b) \mathcal{E}(a, b; \gamma\pi + \bar{\gamma}\hat{p}) \\ &\leq \frac{2\gamma}{\eta} + \frac{\bar{\gamma}}{\eta} \langle \hat{p} - q, H^\top \ell \rangle + L^2 \sum_{a,b \in \mathcal{A}} q(a)q(b) \mathcal{E}(a, b; \gamma\pi + \bar{\gamma}\hat{p}) \\ &\leq \frac{2\gamma}{\eta} + \frac{\bar{\gamma}}{\eta} \langle \hat{p} - q, H^\top \ell \rangle + 2L^2 \sum_{a \in \mathcal{A}} q(a) \mathcal{E}(a, a^*; \gamma\pi + \bar{\gamma}\hat{p}), \end{aligned} \tag{10}$$

where the first inequality uses the assumption that $\max_{a,b \in \mathcal{A}, \ell \in \mathcal{L}} |(\psi_a - \psi_b)^\top \ell| \leq 2$, and the second inequality follows from the definition of $\mathcal{E}$, Lemma 13, and Lemma 14 (noting that $\mathcal{E}(a, a^*; p) = (\psi_a - He_{a^*})^\top Q_\delta(p)^\dagger (\psi_a - He_{a^*})$). We now proceed with the selection of $\hat{p}$ with the purpose of minimizing the second and third terms. A convenient choice is to select $\hat{p}$ so as to perform better than q against the loss vector ℓ . For that, it suffices to consider $\hat{p}$ of the form

$$\hat{p} = \sum_{a \in \mathcal{A}^*} q(a) \nu_a \quad \text{such that} \quad \forall a \in \mathcal{A}^* \quad \langle \nu_a - e_a, H^\top \ell \rangle \leq 0, \tag{11}$$

where $\nu_a \in \Delta_k$. That is, $\hat{p}$ is a mixture (with weights proportional to q) of $|\mathcal{A}^*|$ distributions, each associated with a Pareto optimal action and is chosen to outperform this action against ℓ . This choice allows us to simultaneously eliminate the second term in (10) and bound the third term independently of tunable parameters. The former fact is easily demonstrated:

$$\frac{\bar{\gamma}}{\eta} \langle \hat{p} - q, H^\top \ell \rangle = \frac{\bar{\gamma}}{\eta} \sum_{a \in \mathcal{A}^*} q(a) \langle \nu_a - e_a, H^\top \ell \rangle \leq 0. \quad (12)$$

Moving over to the third term, fix $\lambda \in \mathcal{W}_\ell^{\text{loc}}$ such that $\psi_a - \psi_{a^*} = M\lambda_a$ for all $a \in \mathcal{A}^*$, which exists via local observability. Then, for any $a \in \mathcal{A}^*$ and any $p \in \Delta_k$, we have that

$$\begin{aligned} \mathcal{E}(a, a^*; p) &= (\psi_a - \psi_{a^*})^\top Q_\delta(p)^\dagger (\psi_a - \psi_{a^*}) \\ &= \|(Q_\delta(p)^\dagger)^{1/2} (\psi_a - \psi_{a^*})\|^2 \\ &= \|(Q_\delta(p)^\dagger)^{1/2} M\lambda_a\|^2 \\ &= \|\sum_{c \in \mathcal{A}} (Q_\delta(p)^\dagger)^{1/2} M_c \lambda_a(c)\|^2 \\ &\leq (\sum_{c \in \mathcal{A}} \|(Q_\delta(p)^\dagger)^{1/2} M_c \lambda_a(c)\|)^2 \\ &\leq (\sum_{c \in \mathcal{A}} \|(Q_\delta(p)^\dagger)^{1/2} M_c\| \|\lambda_a(c)\|)^2 \\ &\leq (\sum_{c \in \mathcal{A}} \|\lambda_a(c)\|) \sum_{c \in \mathcal{A}} \|(Q_\delta(p)^\dagger)^{1/2} M_c\|^2, \end{aligned} \quad (13)$$

where the first inequality is an application of the triangle inequality, the second a consequence of the definition of the spectral norm, and the third an application of the Cauchy-Schwarz inequality. Hence,

$$\begin{aligned} \sum_{a \in \mathcal{A}} q(a) \mathcal{E}(a, a^*; p) &= \sum_{a \in \mathcal{A}^*} q(a) \mathcal{E}(a, a^*; p) \\ &\leq \sum_{a \in \mathcal{A}^*} q(a) (\sum_{c \in \mathcal{A}} \|\lambda_a(c)\|) \sum_{c \in \mathcal{A}} \|\lambda_a(c)\| \|(Q_\delta(p)^\dagger)^{1/2} M_c\|^2 \\ &\leq \max_{s \in \mathcal{A}^*} (\sum_{c \in \mathcal{A}} \|\lambda_s(c)\|) \sum_{a \in \mathcal{A}^*} q(a) \sum_{c \in \mathcal{A}} \|\lambda_a(c)\| \|(Q_\delta(p)^\dagger)^{1/2} M_c\|^2 \\ &= \max_{s \in \mathcal{A}^*} (\sum_{c \in \mathcal{A}} \|\lambda_s(c)\|) \sum_{c \in \mathcal{A}} \sum_{a \in \mathcal{A}^*} q(a) \|\lambda_a(c)\| \|(Q_\delta(p)^\dagger)^{1/2} M_c\|^2 \\ &= \max_{s \in \mathcal{A}^*} (\sum_{c \in \mathcal{A}} \|\lambda_s(c)\|)^2 \sum_{c \in \mathcal{A}} \sum_{a \in \mathcal{A}^*} q(a) \frac{\|\lambda_a(c)\|}{\max_{s \in \mathcal{A}^*} \sum_{b \in \mathcal{A}} \|\lambda_s(b)\|} \|(Q_\delta(p)^\dagger)^{1/2} M_c\|^2 \\ &\leq \max_{s \in \mathcal{A}^*} (\sum_{c \in \mathcal{A}} \|\lambda_s(c)\|)^2 \sum_{c \in \mathcal{A}} \sum_{a \in \mathcal{A}^* : \lambda_a \neq \mathbf{0}} q(a) \frac{\|\lambda_a(c)\|}{\sum_{b \in \mathcal{A}} \|\lambda_a(b)\|} \|(Q_\delta(p)^\dagger)^{1/2} M_c\|^2, \end{aligned} \quad (15)$$

where the first equality holds since q is only supported on $\mathcal{A}^*$. Note that the right-hand-side of (15) now offers a natural way to choose $\hat{p}$ as to respect the constraint in (11) and simultaneously simplify the bound. Simply, let $\nu_a \in \Delta_k$ (as referred to in (11)) be such that $\nu_a(b) \propto \|\lambda_a(b)\|$ for $b \in \mathcal{A}$ if $\lambda_a \neq \mathbf{0}$; otherwise, if $\lambda_a = \mathbf{0}$ (which is the case for a^* and its duplicates), simply set $\nu_a = e_a$. This choice satisfies the condition in (11) since the fact that $\lambda \in \mathcal{W}_\ell^{\text{loc}}$ implies that any action in the support of ν_a has loss smaller than or equal to that of a . To emphasize its dependence on λ and q , we denote this choice of $\hat{p}$ by q^λ . For clarity, we state its definition explicitly below: for $b \in \mathcal{A}$, let

$$\begin{aligned} q^\lambda(b) &:= \sum_{a \in \mathcal{A}^*} q(a) \left(\mathbb{I}[\lambda_a = \mathbf{0}, a = b] + \mathbb{I}[\lambda_a \neq \mathbf{0}] \frac{\|\lambda_a(b)\|}{\sum_{c \in \mathcal{A}} \|\lambda_a(c)\|} \right) \\ &= q(b) \mathbb{I}[b \in \mathcal{A}^*, \lambda_b = \mathbf{0}] + \underbrace{\sum_{a \in \mathcal{A}^* : \lambda_a \neq \mathbf{0}} q(a) \frac{\|\lambda_a(b)\|}{\sum_{c \in \mathcal{A}} \|\lambda_a(c)\|}}_{=: q_-^\lambda(b)}. \end{aligned}$$

For convenience, we have also defined above an abridged mixture q^λ , which excludes actions $a \in \mathcal{A}^*$ for which $\lambda_a = \mathbf{0}$. Now, with $\beta_\lambda := \max_{s \in \mathcal{A}^*} \sum_{c \in \mathcal{A}} \|\lambda_s(c)\|$, (15) implies that

$$\begin{aligned} \sum_{a \in \mathcal{A}} q(a) \mathcal{E}(a, a^*; \gamma\pi + \bar{\gamma}q^\lambda) &\leq \beta_\lambda^2 \sum_{b \in \mathcal{A}} q_-^\lambda(b) \|(Q_\delta(\gamma\pi + \bar{\gamma}q^\lambda)^\dagger)^{1/2} M_b\|^2 \\ &= \beta_\lambda^2 \sum_{b \in \mathcal{A}} q_-^\lambda(b) \|M_b^\top Q_\delta(\gamma\pi + \bar{\gamma}q^\lambda)^\dagger M_b\| \\ &\leq \frac{\beta_\lambda^2}{(1-\delta)(1-\gamma)} \sum_{b \in \mathcal{A}} q_-^\lambda(b) \|M_b^\top Q(q^\lambda)^\dagger M_b\| \\ &\leq 4\beta_\lambda^2 \sum_{b \in \mathcal{A}} q_-^\lambda(b) \|M_b^\top Q(q^\lambda)^\dagger M_b\|, \end{aligned}$$

where the penultimate step is an application of Lemma 7 using the definition of Q_δ and the fact that $\text{col}(M_b) \subseteq \text{col}(Q(q^\lambda))$ whenever $q^\lambda(b) > 0$ (see Lemma 8), and the last step uses that $\delta, \gamma \leq 1/2$. Hence, combined with (10) and (12), this entails that

$$\Lambda_{\eta,q}(\gamma\pi + \bar{\gamma}q^\lambda, \ell) \leq \frac{2\gamma}{\eta} + 8L^2\beta_\lambda^2 \sum_{b \in \mathcal{A}} q_-^\lambda(b) \|M_b^\top Q(q^\lambda)^\dagger M_b\|. \quad (16)$$

The smallest admissible value for γ depends on the choice of π , which should be chosen so as to ensure that $(\gamma\pi + \bar{\gamma}q^\lambda) \in \Xi_\eta$; that is, $z(\gamma\pi + \bar{\gamma}q^\lambda) \leq 2/\eta L$. Now, fix $\xi \in \mathcal{W}^{\text{glo}}$ and define $\beta_\xi := \max_{a,b \in \mathcal{A}} \sum_{c \in \mathcal{A}} \|\xi_{a,b}(c)\|$. Analogously to (13), we have that for any $a, b \in \mathcal{A}$,

$$\begin{aligned} \mathcal{E}(a, b; \gamma\pi + \bar{\gamma}q^\lambda) &\leq \left(\sum_{c \in \mathcal{A}} \|\xi_{a,b}(c)\| \right) \sum_{c \in \mathcal{A}} \|\xi_{a,b}(c)\| \|(Q_\delta(\gamma\pi + \bar{\gamma}q^\lambda)^\dagger)^{1/2} M_c\|^2 \\ &\leq \left(\sum_{c \in \mathcal{A}} \|\xi_{a,b}(c)\| \right)^2 \max_{c \in \mathcal{A}} \|(Q_\delta(\gamma\pi + \bar{\gamma}q^\lambda)^\dagger)^{1/2} M_c\|^2 \\ &\leq \beta_\xi^2 \max_{c \in \mathcal{A}} \|(Q_\delta(\gamma\pi + \bar{\gamma}q^\lambda)^\dagger)^{1/2} M_c\|^2. \end{aligned} \quad (17)$$

Then, enforcing that $\text{col}(Q(\pi)) = \text{col}(\mathbf{M})$, we get that

$$\begin{aligned} z(\gamma\pi + \bar{\gamma}q^\lambda) &= \max_{a,b \in \mathcal{A}} \mathcal{E}(a, b; \gamma\pi + \bar{\gamma}q^\lambda) + \max_{c \in \mathcal{A}} \|(Q_\delta(\gamma\pi + \bar{\gamma}q^\lambda)^\dagger)^{1/2} M_c\|^2 \\ &\leq (1 + \beta_\xi^2) \max_{c \in \mathcal{A}} \|(Q_\delta(\gamma\pi + \bar{\gamma}q^\lambda)^\dagger)^{1/2} M_c\|^2 \\ &\leq \frac{1}{(1-\delta)\gamma} (1 + \beta_\xi^2) \max_{c \in \mathcal{A}} \|(Q(\pi)^\dagger)^{1/2} M_c\|^2 \\ &\leq \frac{2}{\gamma} (1 + \beta_\xi^2) \max_{c \in \mathcal{A}} \|(Q(\pi)^\dagger)^{1/2} M_c\|^2 \\ &= \frac{2}{\gamma} (1 + \beta_\xi^2) \max_{c \in \mathcal{A}} \|(UB(\pi)^{-1}U^\top)^{1/2} M_c\|^2, \end{aligned} \quad (18)$$

where the second inequality follows from Lemma 7 as $\text{col}(M_b) \subseteq \text{col}(Q(\pi))$ by the assumption on π , the third inequality holds since $\delta \leq 1/2$, and the last equality also follows from the assumption on π (see Lemmas 9 and 10). At this juncton, we pick

$$\pi \in \arg \min_{\pi' \in \Delta_k} \max_{c \in \mathcal{A}} \|(UB(\pi')^{-1}U^\top)^{1/2} M_c\|^2,$$

which satisfies that $\text{col}(Q(\pi)) = \text{col}(\mathbf{M})$ as this is equivalent to $B(\pi)$ being non-singular, see Lemma 9. With this choice of π , equating the right-hand side of (18) to $2/\eta L$ dictates setting

$$\gamma = \eta L (1 + \beta_\xi^2) \min_{\pi' \in \Delta_k} \max_{c \in \mathcal{A}} \|(UB(\pi')^{-1}U^\top)^{1/2} M_c\|^2.$$

To ensure that this choice of γ is valid (i.e., $\gamma \leq 1/2$), we enforce that

$$1/\eta \geq 2L(1 + \beta_\xi^2) \min_{\pi' \in \Delta_k} \max_{c \in \mathcal{A}} \|(UB(\pi')^{-1}U^\top)^{1/2} M_c\|^2.$$

Plugging this value of γ back in (16) finally yields that

$$\begin{aligned} \min_{(p,r) \in \Xi_\eta} \Lambda_{\eta,q}(p,r,\ell) &\leq 8L^2\beta_\lambda^2 \sum_{b \in \mathcal{A}} q_-^\lambda(b) \|M_b^\top Q(q^\lambda)^\dagger M_b\| \\ &\quad + 2L(1 + \beta_\xi^2) \min_{\pi' \in \Delta_k} \max_{c \in \mathcal{A}} \|(UB(\pi')^{-1}U^\top)^{1/2} M_c\|^2. \end{aligned}$$

□

Theorem 1. *In locally observable games, it holds that*

$$\begin{aligned} \Lambda_\eta^* &\leq \max_{\ell \in \mathcal{L}} \inf_{\lambda \in \mathcal{W}_\ell^{\text{loc}}} 8L^2\beta_\lambda^2 \min\{\text{rk}(\mathbf{M}), |\text{supp}(\lambda)|\} + 2L(1 + \beta_{\text{glo}}^2) \min\{\text{rk}(\mathbf{M}), w^*\} \\ \text{provided } \frac{1}{\eta} &\geq 2L(1 + \beta_{\text{glo}}^2) \min\{\text{rk}(\mathbf{M}), w^*\}, \text{ where } \beta_{\text{glo}} := \max_{a,b \in \mathcal{A}} \min_{\mathbf{v} \in \mathbb{R}^n: \psi_a - \psi_b = \mathbf{M}\mathbf{v}} \sum_{c \in \mathcal{A}} \|\mathbf{v}(c)\| \\ \text{and } w^* &:= \min_{S \subseteq \mathcal{A}} |S| \max_{b \in \mathcal{A}} \|M_b^\top U (\sum_{s \in S} U^\top M_s M_s^\top U)^{-1} U^\top M_b\| \leq k. \end{aligned}$$

Proof. Fix $\ell \in \mathcal{L}$, $q \in \Delta_k$, and $\lambda \in \mathcal{W}_\ell^{\text{loc}}$. Concerning the first term in the bound of Proposition 3, we have that

$$\begin{aligned} \sum_{b \in \mathcal{A}} q_-^\lambda(b) \|M_b^\top Q(q^\lambda)^\dagger M_b\| &\leq \sum_{b \in \mathcal{A}} q^\lambda(b) \|M_b^\top Q(q^\lambda)^\dagger M_b\| \\ &\leq \sum_{b \in \mathcal{A}} q^\lambda(b) \text{tr}(M_b^\top Q(q^\lambda)^\dagger M_b) \\ &= \sum_{b \in \mathcal{A}} q^\lambda(b) \text{tr}(Q(q^\lambda)^\dagger M_b M_b^\top) \\ &= \text{tr}(Q(q^\lambda)^\dagger \sum_{b \in \mathcal{A}} q^\lambda(b) M_b M_b^\top) \\ &= \text{tr}(Q(q^\lambda)^\dagger Q(q^\lambda)) = \text{rk}(Q(q^\lambda)) \leq \text{rk}(\mathbf{M}), \end{aligned}$$

where the last inequality follows from Lemma 8. At the same time, whenever $q^\lambda(b) > 0$, an application of Lemma 7 gives that

$$\|M_b^\top Q(q^\lambda)^\dagger M_b\| \leq \|M_b^\top (q^\lambda(b) M_b M_b^\top)^\dagger M_b\| = \frac{\|M_b^\top (M_b M_b^\top)^\dagger M_b\|}{q^\lambda(b)} = \frac{\|M_b^\dagger M_b\|}{q^\lambda(b)} \leq \frac{1}{q^\lambda(b)},$$

where the last inequality uses that $M_b^\dagger M_b$ is an orthogonal projector. Then,

$$\sum_{b \in \mathcal{A}} q_-^\lambda(b) \|M_b^\top Q(q^\lambda)^\dagger M_b\| \leq \sum_{b \in \mathcal{A}: q_-^\lambda(b) > 0} \frac{q_-^\lambda(b)}{q^\lambda(b)} \leq |\{b \in \mathcal{A}: q_-^\lambda(b) > 0\}| = |\text{supp}(\lambda)|.$$

For the second term, setting π as the uniform distribution over some $S \subseteq \mathcal{A}$ (assuming $\sum_{s \in S} U^\top M_s M_s^\top U$ is invertible) gives that

$$\max_{b \in \mathcal{A}} \|M_b^\top UB(\pi)^{-1}U^\top M_b\| = |S| \max_{b \in \mathcal{A}} \|M_b^\top U (\sum_{s \in S} U^\top M_s M_s^\top U)^{-1} U^\top M_b\|.$$

Hence,

$$\min_{\pi \in \Delta_k} \max_{b \in \mathcal{A}} \|M_b^\top UB(\pi)^{-1}U^\top M_b\| \leq \min_{S \subseteq \mathcal{A}} |S| \max_{b \in \mathcal{A}} \|M_b^\top U (\sum_{s \in S} U^\top M_s M_s^\top U)^{-1} U^\top M_b\| =: w^*$$

Note that for any $b \in \mathcal{A}$, the norm on the right-hand side of the inequality is bounded by 1 whenever $b \in S$ (again via an application of Lemma 7 as above). Hence, choosing S as $\mathcal{A}$ yields k as an upper bound on w^* . At the same time,

$$\min_{\pi \in \Delta_k} \max_{b \in \mathcal{A}} \|M_b^\top UB(\pi)^{-1}U^\top M_b\| \leq \min_{\pi \in \Delta_k} \max_{b \in \mathcal{A}} \text{tr}(M_b^\top UB(\pi)^{-1}U^\top M_b).$$

The right-hand side is a form of the G -optimal design criterion (with matrices as elements of the design space), for which one can show that the minimizer coincides with the maximizer of $\log \det(\sum_{a \in \mathcal{A}} \pi(a) U^\top M_a M_a^\top U)$ and attains a value of $\text{rk}(\mathbf{M})$, see Lemma 12.

Let $\beta_{\text{glo}} := \min_{\xi \in \mathcal{W}_{\text{glo}}} \beta_{\xi}$. Combining the four bounds derived above with Proposition 3 (which holds for any $\lambda \in \mathcal{W}_{\ell}^{\text{loc}}$ and $\xi \in \mathcal{W}_{\text{glo}}$) and Lemma 16 gives that

$$\min_{p \in \Xi_{\eta}} \Lambda_{\eta,q}(p, \ell) \leq \inf_{\lambda \in \mathcal{W}_{\ell}^{\text{loc}}} 8L^2 \beta_{\lambda}^2 \min\{\text{rk}(\mathbf{M}), |\text{supp}(\lambda)|\} + 2L(1 + \beta_{\text{glo}}^2) \min\{\text{rk}(\mathbf{M}), w^*\},$$

as long as $1/\eta \geq 2L(1 + \beta_{\text{glo}}^2) \min\{\text{rk}(\mathbf{M}), w^*\}$. Taking the maximum over $\ell \in \mathcal{L}$ in this bound yields an upper bound for $\Lambda_{\eta,q}^*$ via Lemma 3, which is also an upper bound for Λ_{η}^* as it does not depend on q . Finally, note that the image of the map $\ell \mapsto \mathcal{W}_{\ell}^{\text{loc}}$ is finite as $\mathcal{W}_{\ell}^{\text{loc}}$ depends on ℓ only via the ordering induced by the latter over the actions; hence, the maximum value over $\ell \in \mathcal{L}$ of the bound above is attainable. $\square$

E.1 Alternative exploration distributions

We revisit here the key argument in the proof of Proposition 3; that of choosing $\hat{p}$ to control $\sum_{a \in \mathcal{A}} q(a) \mathcal{E}(a, a^*; \hat{p})$ while keeping $\langle \hat{p} - q, H^{\top} \ell \rangle$ non-positive, having fixed $\ell \in \mathcal{L}$ and $a^* \in \arg \min_{a \in \mathcal{A}} \psi_a^{\top} \ell$. As argued in that proof, for any choice of $\lambda \in \mathcal{W}_{\ell}^{\text{loc}}$ (such that $\psi_a - \psi_{a^*} = \mathbf{M} \lambda_a$ for all $a \in \mathcal{A}$), choosing $\hat{p} = \sum_{a \in \mathcal{A}^*} q(a) \nu_a$ with ν_a supported on $\text{supp}(\lambda, a)$ suffices to ensure that $\langle \hat{p} - q, H^{\top} \ell \rangle \leq 0$. On the other hand, establishing an upper bound on $\mathcal{E}(a, a^*; \hat{p})$ that is proportional to $\sum_{c \in \text{supp}(\lambda, a)} \nu_a(c) \|M_c^{\top} Q_{\delta}(\hat{p})^{\dagger} M_c\|$ allows using the fact that $\sum_{a \in \mathcal{A}} r(a) \|M_a^{\top} Q_{\delta}(r)^{\dagger} M_a\|$ is at most $2 \text{rk}(\mathbf{M})$ for any $r \in \Delta_k$ (also, at most twice the size of the support of r) to bound $\sum_{a \in \mathcal{A}} q(a) \mathcal{E}(a, a^*; \hat{p})$ as such (see the proof of Theorem 1 above). A slight generalization of an argument in the proof of Proposition 3 gives that for any $\alpha \in \mathbb{R}$ and $a \in \mathcal{A}^*$ (with $\lambda_a \neq 0$)

$$\begin{aligned} \mathcal{E}(a, a^*; \hat{p}) &= (\psi_a - \psi_{a^*})^{\top} Q_{\delta}(\hat{p})^{\dagger} (\psi_a - \psi_{a^*}) \\ &= \|(Q_{\delta}(\hat{p})^{\dagger})^{1/2} (\psi_a - \psi_{a^*})\|^2 \\ &= \|(Q_{\delta}(\hat{p})^{\dagger})^{1/2} \mathbf{M} \lambda_a\|^2 \\ &= \|\sum_{c \in \text{supp}(\lambda, a)} (Q_{\delta}(\hat{p})^{\dagger})^{1/2} M_c \lambda_a(c)\|^2 \\ &\leq (\sum_{c \in \text{supp}(\lambda, a)} \|(Q_{\delta}(\hat{p})^{\dagger})^{1/2} M_c \lambda_a(c)\|)^2 \\ &\leq (\sum_{c \in \text{supp}(\lambda, a)} \|(Q_{\delta}(\hat{p})^{\dagger})^{1/2} M_c\| \|\lambda_a(c)\|)^2 \\ &\leq (\sum_{c \in \text{supp}(\lambda, a)} \|\lambda_a(c)\|^{2-2\alpha}) (\sum_{c \in \text{supp}(\lambda, a)} \|\lambda_a(c)\|^{2\alpha} \|M_c^{\top} Q_{\delta}(\hat{p})^{\dagger} M_c\|) \\ &= (\sum_{c \in \text{supp}(\lambda, a)} \|\lambda_a(c)\|^{2-2\alpha}) (\sum_{c \in \text{supp}(\lambda, a)} \|\lambda_a(c)\|^{2\alpha}) \\ &\quad \cdot \left(\sum_{c \in \text{supp}(\lambda, a)} \frac{\|\lambda_a(c)\|^{2\alpha}}{\sum_{b \in \text{supp}(\lambda, a)} \|\lambda_a(b)\|^{2\alpha}} \|M_c^{\top} Q_{\delta}(\hat{p})^{\dagger} M_c\| \right), \end{aligned}$$

where the last inequality is an application of the Cauchy-Schwarz inequality. One can then set $\nu_a(c) \propto \|\lambda_a(c)\|^{2\alpha}$, where $\alpha = 1/2$ corresponds to the choice of ν_a used in the proof of Proposition 3 and $\alpha = 0$ yields $\nu_a(c) \propto \mathbb{I}[c \in \text{supp}(\lambda, a)]$, which is essentially the analogous choice of $\hat{p}$ to that used in the analysis of Lattimore and Szepesvári [38] in finite partial monitoring. In any case, the resulting bound is

$$\begin{aligned} \sum_{a \in \mathcal{A}} q(a) \mathcal{E}(a, a^*; \hat{p}) &\leq \sum_{a \in \mathcal{A}} q(a) (\sum_{c \in \text{supp}(\lambda, a)} \|\lambda_a(c)\|^{2-2\alpha}) (\sum_{c \in \text{supp}(\lambda, a)} \|\lambda_a(c)\|^{2\alpha}) \\ &\quad \cdot (\sum_{c \in \mathcal{A}} \nu_a(c) \|M_c^{\top} Q_{\delta}(\hat{p})^{\dagger} M_c\|) \\ &\leq \max_{a \in \mathcal{A}^*} (\sum_{c \in \text{supp}(\lambda, a)} \|\lambda_a(c)\|^{2-2\alpha}) (\sum_{c \in \text{supp}(\lambda, a)} \|\lambda_a(c)\|^{2\alpha}) \\ &\quad \cdot \sum_{b \in \mathcal{A}} \hat{p}(b) \|M_b^{\top} Q_{\delta}(\hat{p})^{\dagger} M_b\|. \end{aligned}$$

Since for any $\alpha \in \mathbb{R}$ the Cauchy-Schwarz inequality gives that

$$(\sum_{c \in \text{supp}(\lambda, a)} \|\lambda_a(c)\|^{2-2\alpha}) (\sum_{c \in \text{supp}(\lambda, a)} \|\lambda_a(c)\|^{2\alpha}) \geq (\sum_{c \in \text{supp}(\lambda, a)} \|\lambda_a(c)\|)^2,$$

choosing $\alpha = 1/2$ minimizes the leading factor in the bound above. To be clear, it is not claimed here that this choice of $\hat{p}$ minimizes $\sum_{a \in \mathcal{A}} q(a) \mathcal{E}(a, a^*; \hat{p})$, neither under the structural constraint that $\hat{p} = \sum_{a \in \mathcal{A}^*} q(a) \nu_a$ with ν_a supported on $\text{supp}(\lambda, a)$ nor the general constraint that $\langle \hat{p} - q, H^\top \ell \rangle \leq 0$, neither for a fixed q nor after passing to the sup over $q \in \Delta_k$. It is not generally clear if such an optimal distribution can be expressed in closed form. The aim, rather, was to show that our choice of $\hat{p}$ is fairly natural and justified in pursuit of a simple and interpretable bound (one that isolates in a standalone leading factor the ‘cost’ of using the weights $(\lambda_a)_a$) through our linear-bandit-inspired analysis.

To expand on this discussion, we consider now a special case where a tighter analysis leads to a different value of α being the most convenient. Assume that $\text{col}(M_b)$ and $\text{col}(M_c)$ are orthogonal for any $b \neq c$. In such a case, it holds that

$$\begin{aligned} \mathcal{E}(a, a^*; \hat{p}) &= (\psi_a - \psi_{a^*})^\top Q_\delta(\hat{p})^\dagger (\psi_a - \psi_{a^*}) \\ &= (M\lambda_a)^\top Q_\delta(\hat{p})^\dagger (M\lambda_a) \\ &= \left(\sum_{c \in \text{supp}(\lambda, a)} M_c \lambda_a(c) \right)^\top Q_\delta(\hat{p})^\dagger \left(\sum_{c' \in \text{supp}(\lambda, a)} M_{c'} \lambda_a(c') \right) \\ &= \sum_{c \in \text{supp}(\lambda, a)} (M_c \lambda_a(c))^\top Q_\delta(\hat{p})^\dagger (M_c \lambda_a(c)) \\ &\leq \sum_{c \in \text{supp}(\lambda, a)} \|\lambda_a(c)\|^2 \|M_c^\top Q_\delta(\hat{p})^\dagger M_c\| \\ &= \left(\sum_{c' \in \text{supp}(\lambda, a)} \|\lambda_a(c')\|^2 \right) \sum_{c \in \text{supp}(\lambda, a)} \frac{\|\lambda_a(c)\|^2}{\sum_{b \in \text{supp}(\lambda, a)} \|\lambda_a(b)\|^2} \|M_c^\top Q_\delta(\hat{p})^\dagger M_c\| \end{aligned}$$

where the fourth equality follows from the assumed structure of the observation matrices. This immediately suggests setting $\nu_a(c) \propto \|\lambda_a(c)\|^2$ (or picking $\alpha = 1$), obtaining as result that

$$\sum_{a \in \mathcal{A}} q(a) \mathcal{E}(a, a^*; \hat{p}) \leq \max_{a \in \mathcal{A}^*} \left(\sum_{c \in \text{supp}(\lambda, a)} \|\lambda_a(c)\|^2 \right) \sum_{b \in \mathcal{A}} \hat{p}(b) \|M_b^\top Q_\delta(\hat{p})^\dagger M_b\|,$$

which features a better leading factor than what we obtained with $\alpha = 1/2$ in the general case. Again, this choice is not necessarily optimal in any precise sense, but it allows exploiting the structure of this special case in a simple and natural manner. Devising a general choice of $\hat{p}$ that can recover improved bounds in special cases as such is a worthy direction for refining the bound of Theorem 1.

F Proofs of Section 5

The proof of Theorem 2 relies on the following proposition. (The definition of $B(\cdot)$ is given in (8).)

Proposition 4. *In globally observable games, it holds for any $\ell \in \mathcal{L}$ and $q \in \Delta_k$ that*

$$\Lambda_\eta^* \leq 4\sqrt{\frac{1}{\eta}} L \min \left\{ \sqrt{(1 + \beta_{\text{glo}}^2) \min_{\pi \in \Delta_k} \max_{b \in \mathcal{A}} \|M_b^\top U B(\pi)^{-1} U^\top M_b\|}, \sqrt{(1 + \beta_{2, \text{glo}}^2) \min_{\pi \in \Delta_k} \|\mathbf{M}^\top U B(\pi)^{-1} U^\top \mathbf{M}\|} \right\}$$

provided that

$$1/\eta \geq (1 + L^2) \min \left\{ (1 + \beta_{\text{glo}}^2) \min_{\pi \in \Delta_k} \max_{b \in \mathcal{A}} \|M_b^\top U B(\pi)^{-1} U^\top M_b\|, (1 + \beta_{2, \text{glo}}^2) \min_{\pi \in \Delta_k} \|\mathbf{M}^\top U B(\pi)^{-1} U^\top \mathbf{M}\| \right\}$$

where

$$\beta_{2, \text{glo}} := \min_{\xi \in \mathcal{W}^{\text{glo}}} \max_{a, b \in \mathcal{A}} \|\xi_{a, b}\| = \max_{a, b \in \mathcal{A}} \min_{\mathbf{v} \in \mathbb{R}^n : \psi_a - \psi_b = \mathbf{M}\mathbf{v}} \|\mathbf{v}\| = \max_{a, b \in \mathcal{A}} \|\mathbf{M}^\dagger (\psi_a - \psi_b)\|.$$

Proof. As in the proof of Proposition 3, we choose a certain distribution $p \in \Xi_\eta$ and bound $\Lambda_{\eta, q}(p, \ell)$. Here, we simply set $p = \gamma\pi + (1 - \gamma)q$, with $\gamma \in (0, 1]$ and $\pi \in \Delta_k$ such that $\text{col}(Q(\pi)) = \text{col}(\mathbf{M})$. For brevity, let $\bar{\gamma} = 1 - \gamma$. With these choices we obtain that

$$\begin{aligned} \Lambda_{\eta, q}(\gamma\pi + \bar{\gamma}q, \ell) &= \frac{\gamma}{\eta} \langle \pi - q, H^\top \ell \rangle + \frac{\bar{\gamma}}{\eta} \langle q - q, H^\top \ell \rangle + L^2 \sum_{a, b \in \mathcal{A}} q(a) q(b) \mathcal{E}(a, b; \gamma\pi + \bar{\gamma}q) \\ &\leq \frac{2\gamma}{\eta} + L^2 \sum_{a, b \in \mathcal{A}} q(a) q(b) \mathcal{E}(a, b; \gamma\pi + \bar{\gamma}q), \end{aligned} \tag{19}$$

where the inequality uses the assumption that $\max_{a,b \in \mathcal{A}, \ell \in \mathcal{L}} |(\psi_a - \psi_b)^\top \ell| \leq 2$. Fix some $\xi \in \mathcal{W}^{\text{glo}}$, which exists via global observability. Then, similarly to (17) in the proof of Proposition 3, for any $a, b \in \mathcal{A}$,

$$\begin{aligned} \mathcal{E}(a, b; \gamma\pi + \bar{\gamma}q) &\leq \beta_\xi^2 \max_{c \in \mathcal{A}} \|(Q_\delta(\gamma\pi + \bar{\gamma}q)^\dagger)^{1/2} M_c\|^2 \\ &\leq \frac{1}{(1-\delta)\gamma} \beta_\xi^2 \max_{c \in \mathcal{A}} \|M_c^\top Q(\pi)^\dagger M_c\| \\ &\leq \frac{2}{\gamma} \beta_\xi^2 \max_{c \in \mathcal{A}} \|M_c^\top Q(\pi)^\dagger M_c\|, \end{aligned}$$

where the second step follows from Lemma 7 using the assumption that $\text{col}(Q(\pi)) = \text{col}(\mathbf{M})$, and the last step uses that $\delta \leq 1/2$. Then, choosing $\xi \in \arg \min_{\xi' \in \mathcal{W}^{\text{glo}}} \beta_{\xi'}$, we get via Lemma 16 that

$$\begin{aligned} \sum_{a,b \in \mathcal{A}} q(a)q(b)\mathcal{E}(a, b; \gamma\pi + \bar{\gamma}q) &\leq \frac{2}{\gamma} \beta_{\text{glo}}^2 \max_{b \in \mathcal{A}} \|M_b^\top Q(\pi)^\dagger M_b\| \\ &\leq \frac{2}{\gamma} (1 + \beta_{\text{glo}}^2) \max_{b \in \mathcal{A}} \|M_b^\top Q(\pi)^\dagger M_b\|. \end{aligned}$$

Alternatively, for any $\xi \in \mathcal{W}^{\text{glo}}$, we also have that

$$\begin{aligned} \mathcal{E}(a, b; \gamma\pi + \bar{\gamma}q) &= (\psi_a - \psi_b)^\top Q_\delta(\gamma\pi + \bar{\gamma}q)^\dagger (\psi_a - \psi_b) \\ &= \|(Q_\delta(\gamma\pi + \bar{\gamma}q)^\dagger)^{1/2} (\psi_a - \psi_b)\|^2 \\ &= \|(Q_\delta(\gamma\pi + \bar{\gamma}q)^\dagger)^{1/2} \mathbf{M} \xi_{a,b}\|^2 \\ &\leq \|\xi_{a,b}\|^2 \|(Q_\delta(\gamma\pi + \bar{\gamma}q)^\dagger)^{1/2} \mathbf{M}\|^2 \\ &\leq \frac{2}{\gamma} \|\xi_{a,b}\|^2 \|\mathbf{M}^\top Q(\pi)^\dagger \mathbf{M}\|, \end{aligned}$$

again using the assumption that $\text{col}(Q(\pi)) = \text{col}(\mathbf{M})$. With $\xi \in \arg \min_{\xi' \in \mathcal{W}^{\text{glo}}} \max_{a,b \in \mathcal{A}} \|\xi'_{a,b}\|$ we obtain that

$$\sum_{a,b \in \mathcal{A}} q(a)q(b)\mathcal{E}(a, b; \gamma\pi + \bar{\gamma}q) \leq \frac{2}{\gamma} \beta_{2,\text{glo}}^2 \|\mathbf{M}^\top Q(\pi)^\dagger \mathbf{M}\| \leq \frac{2}{\gamma} (1 + \beta_{2,\text{glo}}^2) \|\mathbf{M}^\top Q(\pi)^\dagger \mathbf{M}\|,$$

where

$$\beta_{2,\text{glo}} := \max_{a,b \in \mathcal{A}} \|\mathbf{M}^\dagger(\psi_a - \psi_b)\| = \max_{a,b \in \mathcal{A}} \min_{v \in \mathbb{R}^n : \psi_a - \psi_b = \mathbf{M}v} \|v\| = \min_{\xi \in \mathcal{W}^{\text{glo}}} \max_{a,b \in \mathcal{A}} \|\xi_{a,b}\|.$$

The first equality follows via the minimum norm property of the Moore-Penrose inverse, and the second is due to the structure of $\mathcal{W}^{\text{glo}}$ (see also the proof of Lemma 16 in Appendix E). Let

$$v_1(\pi) := (1 + \beta_{\text{glo}}^2) \max_{b \in \mathcal{A}} \|M_b^\top U B(\pi)^{-1} U^\top M_b\|$$

and

$$v_2(\pi) := (1 + \beta_{2,\text{glo}}^2) \|\mathbf{M}^\top U B(\pi)^{-1} U^\top \mathbf{M}\|.$$

Returning back to (19), we have shown that

$$\Lambda_{\eta,q}(\gamma\pi + \bar{\gamma}q, \ell) \leq \frac{2\gamma}{\eta} + \frac{2}{\gamma} L^2 \min\{v_1(\pi), v_2(\pi)\}.$$

This expression is minimized at $\gamma = \sqrt{\eta}L \min\{\sqrt{v_1(\pi)}, \sqrt{v_2(\pi)}\}$, yielding that

$$\Lambda_{\eta,q}(\gamma\pi + \bar{\gamma}q, \ell) \leq 4\sqrt{\frac{1}{\eta}} L \min\{\sqrt{v_1(\pi)}, \sqrt{v_2(\pi)}\}.$$

For $\gamma \leq 1$ to hold, it must hold that

$$1/\eta \geq L^2 \min\{v_1(\pi), v_2(\pi)\}.$$

Consider also that

$$\begin{aligned} z(\gamma\pi + \bar{\gamma}q) &= \max_{a,b \in \mathcal{A}} \mathcal{E}(a,b; \gamma\pi + \bar{\gamma}q) + \max_{c \in \mathcal{A}} \|M_c^\top Q_\delta(\gamma\pi + \bar{\gamma}q)^\dagger M_c\| \\ &\leq \frac{2}{\gamma} (1 + \beta_{\text{glo}}^2) \max_{b \in \mathcal{A}} \|M_b^\top Q(\pi)^\dagger M_b\|. \end{aligned}$$

And that at the same time,

$$\begin{aligned} z(\gamma\pi + \bar{\gamma}q) &= \max_{a,b \in \mathcal{A}} \mathcal{E}(a,b; \gamma\pi + \bar{\gamma}q) + \max_{c \in \mathcal{A}} \|M_c^\top Q_\delta(\gamma\pi + \bar{\gamma}q)^\dagger M_c\| \\ &\leq \frac{2}{\gamma} \beta_{2,\text{glo}}^2 \|\mathbf{M}^\top Q(\pi)^\dagger \mathbf{M}\| + \frac{2}{\gamma} \max_{b \in \mathcal{A}} \|M_b^\top Q(\pi)^\dagger M_b\| \\ &\leq \frac{2}{\gamma} (1 + \beta_{2,\text{glo}}^2) \|\mathbf{M}^\top Q(\pi)^\dagger \mathbf{M}\|, \end{aligned}$$

where the second inequality holds since the columns of every M_b are also columns of $\mathbf{M}$ via its definition; hence, we can construct a matrix C_b with $\|C_b\| \leq 1$ such that $\mathbf{M} = M_b C_b$. We have then shown that,

$$\begin{aligned} z(\gamma\pi + \bar{\gamma}q) &\leq \frac{2}{\gamma} \min\{v_1(\pi), v_2(\pi)\} \\ &= \frac{2}{\sqrt{\eta}L} \min\{\sqrt{v_1(\pi)}, \sqrt{v_2(\pi)}\}. \end{aligned}$$

Therefore, for $\gamma\pi + \bar{\gamma}q \in \Xi_\eta$ to hold, it suffices that the last expression is no larger than $2/\eta L$; meaning that η must also satisfy that

$$1/\eta \geq \min\{v_1(\pi), v_2(\pi)\}.$$

Thus, if we pick $\pi \in \arg \min_{\pi' \in \Delta_k} \min\{v_1(\pi'), v_2(\pi')\}$,⁵ then enforcing that

$$1/\eta \geq (1 + L^2) \min_{\pi \in \Delta_k} \min\{v_1(\pi), v_2(\pi)\} = (1 + L^2) \min\left\{\min_{\pi \in \Delta_k} v_1(\pi), \min_{\pi \in \Delta_k} v_2(\pi)\right\}$$

suffices to have

$$\begin{aligned} \min_{p \in \Xi_\eta} \Lambda_{\eta,q}(p, \ell) &\leq 4\sqrt{\frac{1}{\eta}} L \min_{\pi \in \Delta_k} \min\{\sqrt{v_1(\pi)}, \sqrt{v_2(\pi)}\} \\ &= 4\sqrt{\frac{1}{\eta}} L \min\left\{\min_{\pi \in \Delta_k} \sqrt{v_1(\pi)}, \min_{\pi \in \Delta_k} \sqrt{v_2(\pi)}\right\} \\ &= 4\sqrt{\frac{1}{\eta}} L \min\left\{\sqrt{\min_{\pi \in \Delta_k} v_1(\pi)}, \sqrt{\min_{\pi \in \Delta_k} v_2(\pi)}\right\}. \end{aligned}$$

Since this bound does not depend on ℓ or q , it is also an upper bound for $\Lambda_{\eta,q}^*$ (via Lemma 3) and Λ_η^* . $\square$

Theorem 2. *In globally observable games, it holds that*

$$\Lambda_\eta^* \leq 4\sqrt{1/\eta} L \min\left\{(1 + \beta_{\text{glo}})\sqrt{\min\{\text{rk}(\mathbf{M}), w^*\}}, (1 + \beta_{2,\text{glo}})\sqrt{u^*}\right\}$$

provided that $\frac{1}{\eta} \geq (1 + L^2) \min\left\{(1 + \beta_{\text{glo}}^2)\min\{\text{rk}(\mathbf{M}), w^*\}, (1 + \beta_{2,\text{glo}}^2)u^*\right\}$, where

$$\beta_{2,\text{glo}} := \max_{a,b \in \mathcal{A}} \|\mathbf{M}^\dagger(\psi_a - \psi_b)\|, \text{ and } u^* := \min_{S \subseteq \mathcal{A}} |S| \|\mathbf{M}^\top U (\sum_{s \in S} U^\top M_s M_s^\top U)^{-1} U^\top \mathbf{M}\| \leq k.$$

Proof. It follows from the proof of Theorem 1 that

$$\sqrt{\min_{\pi \in \Delta_k} \max_{b \in \mathcal{A}} \|M_b^\top U B(\pi)^{-1} U^\top M_b\|} \leq \sqrt{\min\{\text{rk}(\mathbf{M}), w^*\}}.$$

Similarly, restricting π to be uniform over a subset of actions gives that

$$\min_{\pi \in \Delta_k} \|\mathbf{M}^\top U B(\pi)^{-1} U^\top \mathbf{M}\| \leq \min_{S \subseteq \mathcal{A}} |S| \|\mathbf{M}^\top U (\sum_{s \in S} U^\top M_s M_s^\top U)^{-1} U^\top \mathbf{M}\| =: u^*.$$

This is again bounded by k , which is attained with $S = \mathcal{A}$. The theorem then follows from Proposition 4. $\square$

⁵Note that the functions $v_1(\pi)$ and $v_2(\pi)$ do attain their minimum in π over Δ_k . A minimizer π^* of either function must satisfy our requirement that $\text{col}(Q(\pi^*)) = \text{col}(\mathbf{M})$ as this is equivalent to $B(\pi)$ being positive definite, see Lemma 9, which is necessary for $B(\pi)$ to be non-singular.

G Implementation

Define

$$\begin{aligned}\Phi_{\eta,q}(p) &:= \max_{\ell \in \mathcal{L}} \Lambda_{\eta,q}(p, \ell) \\ &= \frac{1}{\eta} \max_{\ell \in \mathcal{L}} \underbrace{\langle p - q, H^\top \ell \rangle + L^2}_{=:\sigma_{\mathcal{L}}(H(p-q))} \sum_{a,b \in \mathcal{A}} q(a)q(b) \underbrace{(\psi_a - \psi_b)^\top Q_\delta(p)^\dagger (\psi_a - \psi_b)}_{=:\mathcal{E}(a,b;p)},\end{aligned}$$

where $\sigma_{\mathcal{L}}$ is the support function of $\mathcal{L}$, which we assumed to be compact. In Algorithm 1, we are asked to solve an optimization problem of the following form at every round:

$$\min_{p \in \Xi_\eta} \Phi_{\eta,q}(p), \quad \Xi_\eta := \{p \in \Delta_k : z(p) \leq 2/\eta L\},$$

where $\eta > 0$ and $q \in \Delta_k$ are given and $z(p) := \max_{a,b \in \mathcal{A}} \mathcal{E}(a,b;p) + \max_{c \in \mathcal{A}} \|M_c^\top Q_\delta(p)^\dagger M_c\|$. This was shown to be a convex problem in the proof of Lemma 3. Going further, we point out here that it can in many interesting cases be represented as a semidefinite program, a well-structured and widely-studied class of convex programs that can be solved in polynomial time via interior-point methods [9]. In terms of semidefinite representability, the crucial components to study here are the functions $\mathcal{E}(a,b;\cdot)$, $\|M_c^\top Q_\delta(\cdot)^\dagger M_c\|$, and $\sigma_{\mathcal{L}}(\cdot)$; that the program is semidefinite representable (SDR) would then follow via the combination rules outlined in [9, Chapters 3 and 4], which include taking the intersection of SDR sets, taking the maximum of finitely many SDR functions, and summing (with non-negative weights) SDR functions.

Fix some positive integer r and a $d \times r$ matrix C that satisfies $\text{col}(C) \subseteq \text{col}(\mathbf{M})$, and consider the mapping $p \mapsto \|C^\top Q_\delta(p)^\dagger C\|$ for $p \in \Delta_k$. This mapping, assuming global observability, subsumes the first two functions specified above; namely, $p \mapsto \mathcal{E}(a,b;p)$ and $p \mapsto \|M_c^\top Q_\delta(p)^\dagger M_c\|$ for any $a, b, c \in \mathcal{A}$. We show now that the mapping in question is SDR by showing that the set

$$\Gamma := \{(p, G, t) \in \Delta_k \times \mathcal{S}^r \times \mathbb{R} \mid C^\top Q_\delta(p)^\dagger C \preceq G \text{ and } \|G\| \leq t\}$$

coincides with the solution set of a number of linear matrix inequalities, where $\mathcal{S}^r$ is the set of symmetric $r \times r$ matrices. This is sufficient since the epigraph of the mapping in question coincides with the projection of Γ onto the (p, t) space. We firstly note that the semidefinite representability Δ_k is immediate; it can be characterized by a set of simple linear inequalities. Consider now the following symmetric matrix:

$$J(p, G) := \begin{bmatrix} Q_\delta(p) & C \\ C^\top & G \end{bmatrix}.$$

By the properties of the (generalized) Schur complement (see, e.g., [57, Theorem 1.20]), we have that

$$J(p, G) \succeq 0 \iff Q_\delta(p) \succeq 0, \text{ col}(C) \subseteq \text{col}(Q_\delta(p)), \text{ and } C^\top Q_\delta(p)^\dagger C \preceq G.$$

For any $p \in \Delta_k$, the first two conditions are always satisfied thanks to the structure of $Q_\delta(p)$ (convex combination of positive semidefinite matrices), the assumption that $\text{col}(C) \subseteq \text{col}(\mathbf{M})$, and the fact that $\text{col}(Q_\delta(p)) = \text{col}(\mathbf{M})$ as asserts Lemma 11. Moreover, it holds that

$$\|G\| \leq t \iff G \preceq t\mathbf{I}_r \text{ and } G \succeq -t\mathbf{I}_r.$$

Hence, for $p \in \Delta_k$ and $G \in \mathcal{S}^r$,

$$(p, G, t) \in \Gamma \iff J(p, G) \succeq 0, \ G \preceq t\mathbf{I}_r, \text{ and } G \succeq -t\mathbf{I}_r;$$

that is, membership in Γ can be characterized by linear matrix inequalities (recalling that $Q_\delta(p)$ is affine in p), implying it is SDR.

As for $\sigma_{\mathcal{L}}$, consider now the common scenario when $\mathcal{L} = \mathcal{B}_p(c) := \{\ell \in \mathbb{R}^d \mid \|\ell\|_p \leq c\}$ for some $p \geq 1$ and $c > 0$; that is, $\mathcal{L}$ is a centered L_p ball. The support function then becomes $\sigma_{\mathcal{L}}(x) = c\|x\|_q$ with $1/p + 1/q = 1$. For rational q , the L_q norm is cone quadratic representable [9, Chapter 3], hence SDR (note that the input to $\sigma_{\mathcal{L}}$ is $H(p - q)$, which is affine in our variable p and thus preserves semidefinite representability). This also holds when $p = 1$ ($q = \infty$) and $p = \infty$ ($q = 1$).

Another common case is when $\mathcal{L}$ is $\mathcal{X}^\circ$, the absolute polar set of $\mathcal{X}$.⁶ Here, $\sigma_{\mathcal{L}}(H(p - q))$ is the optimal value of the following program:

$$\max_{\ell \in \mathbb{R}^d} \sum_{a \in \mathcal{A}} (p(a) - q(a)) \psi_a^\top \ell \quad \text{s.t.} \quad |\psi_a^\top \ell| \leq 1 \quad \forall a \in \mathcal{A}.$$

Via strong duality, it is straightforward to show that

$$\sigma_{\mathcal{L}}(H(p - q)) = \min_{\mathbf{w} \in \mathbb{R}^k: H(p - q) = H\mathbf{w}} \|\mathbf{w}\|_1 = \min_{\mathbf{u}_1, \mathbf{u}_2 \in \mathbb{R}_{\geq 0}^k: H(p - q) = H(\mathbf{u}_1 - \mathbf{u}_2)} (\mathbf{u}_1 + \mathbf{u}_2)^\top \mathbf{1}_k.$$

This is the minimum value of a linear function under linear constraints, which can be easily incorporated into the full program adding $\mathbf{u}_1$ and $\mathbf{u}_2$ as auxiliary variables.

H Adaptive Learning Rate

Algorithm 2 Adaptive Anchored Exploration-by-Optimization

- 1: **input:** stability parameter $\delta \in (0, 1/2]$, sub-optimality tolerance $\varepsilon \geq 0$, scale parameter $L \geq \omega$, $B > 0$
 - 2: **initialize:** $\eta_1 = \min\{B^{-1}, \sqrt{\log k}\}$, $\hat{y}_0(a) = 0 \quad \forall a \in \mathcal{A}$
 - 3: **for** $t = 1, \dots, T$ **do**
 - 4: $\forall a \in \mathcal{A}$, set $q_t(a) \propto \llbracket a \in \mathcal{A}^* \rrbracket \exp(-\eta_t \sum_{s=1}^{t-1} \hat{y}_s(a))$
 - 5: choose $\tilde{p}_t \in \Delta_k$ such that $\max_{\ell \in \mathcal{L}} \Lambda_{\eta_t, q_t}(\tilde{p}_t, \ell) \leq \Lambda_{\eta_t}^* + \varepsilon$ and $z(\tilde{p}_t) \leq \frac{2}{\eta_t L}$
 - 6: set $p_t = (1 - \delta)\tilde{p}_t + \delta \mathbf{1}_k/k$
 - 7: execute $A_t \sim p_t$ and observe $\phi_t = M_t^\top \ell_t$
 - 8: $\forall a \in \mathcal{A}$, set $\hat{y}_t(a) = (\psi_a - Hq_t)^\top Q(p_t)^\dagger M_{A_t} \phi_t$
 - 9: set $V_t = \max\{0, \max_{\ell \in \mathcal{L}} \Lambda_{\eta_t, q_t}(\tilde{p}_t, \ell)\}$ and $\eta_{t+1} = \min\left\{\frac{1}{B}, \sqrt{\frac{\log k}{1 + \sum_{s=1}^t V_s}}\right\}$
 - 10: **end for**
-

Following Lattimore and Szepesvári [38], we briefly describe here an ‘adaptive’ variant of Algorithm 1 that tunes the learning rate in an online fashion, sidestepping (for the most part) the need to have prior knowledge concerning the value of Λ_η^* . This variant is described in Algorithm 2, with the main difference being the use of a learning rate schedule $(\eta_t)_t$, where η_t is set based on the attained values at the optimization problems of past rounds. Note that the algorithm is not completely parameter-free, besides the stability (δ) and tolerance (ε) parameters (which can in principle be reduced at will), it requires a sufficiently large constant $B > 0$ as input to guarantee that the learning rates are sufficiently small to meet the requirements of Theorems 1 and 2. However, B can be picked conservatively as it only affects the regret through an additive term as shown below.

Adapting our analysis along the same lines as the proof of Theorem 6 in [38], one can show (analogously to Proposition 2) that Algorithm 2 satisfies (assuming global observability)

$$\begin{aligned} R_T &\leq 2\delta T + \frac{\log k}{\eta_T} + \mathbb{E} \left[\sum_{t=1}^T \eta_t V_t \right] \\ &\leq 2\delta T + 5\mathbb{E} \left[\sqrt{\left(1 + \sum_{t=1}^T V_t\right) \log k} \right] + \mathbb{E} \left[\max_{t \in [T]} V_t \right] \sqrt{\log k} + B \log k. \end{aligned}$$

Locally observable games. Via Theorem 1, there exist constants α, B' such that $\Lambda_\eta^* \leq \alpha$ given that $\eta \leq 1/B'$. Hence, as long as $B \geq B'$, we have that $V_t \leq \Lambda_{\eta_t}^* + \varepsilon \leq \alpha + \varepsilon$. Thus,

$$R_T \leq 2\delta T + 5\sqrt{(1 + (\alpha + \varepsilon)T) \log k} + (\alpha + \varepsilon)\sqrt{\log k} + B \log k,$$

where, assuming δ and ε are picked sufficiently small (e.g., $o(T^{-1/2})$), the dominant term is of order $\sqrt{\alpha T \log k}$.

⁶Note that for $\mathcal{X}^\circ$ to be compact, $\mathcal{X}$ must span $\mathbb{R}^d$.

Globally observable games. Theorem 2 asserts that there exist constants α, B' such that $\Lambda_\eta^* \leq \alpha/\sqrt{\eta}$ given that $\eta \leq 1/B'$. Hence, assuming that $B \geq \max\{1, B'\}$, we get that

$$V_t \leq \Lambda_{\eta_t}^* + \varepsilon \leq \alpha/\sqrt{\eta_t} + \varepsilon \leq \alpha/\sqrt{\eta_t} + \varepsilon/\sqrt{\eta_t}.$$

Then, from the proof of Proposition 7 in [38] we obtain that

$$\begin{aligned} R_T \leq 2\delta T + 5 \left(\frac{\alpha + \varepsilon}{(\log k)^{1/4}} T + \max\{1, B^2 \log k\}^{3/4} \right)^{2/3} \sqrt{\log k} \\ + \alpha \left(\frac{\alpha + \varepsilon}{(\log k)^{1/4}} T + \max\{1, B^2 \log k\}^{3/4} \right)^{1/3} (\log k)^{1/4} + \varepsilon + B \log k, \end{aligned}$$

where, assuming δ and ε are picked sufficiently small (e.g., $o(T^{-1/3})$), the dominant term is of order $(\alpha T)^{2/3} (\log k)^{1/3}$.

As a final remark, note that ε is our error tolerance when minimizing $\max_{\ell \in \mathcal{L}} \Lambda_{\eta_t, q_t}(p, \ell)$ in $p \in \Xi_{\eta_t}$, where the optimal value is Λ_{η_t, q_t}^* . The properties of this optimization program are discussed in Appendix G. On the other hand, neither algorithm undertakes the (generally) more challenging task of computing $\Lambda_{\eta_t}^*$ up to some tolerance. Still, when tuning a fixed learning rate, this task is somewhat unavoidable if tight regret guarantees are sought, and therein lies the utility of the online learning rate schedule described here.

I Examples - Continued

I.1 Linear bandits

Since $\psi_a = M_a$ for all $a \in \mathcal{A}$ and $\mathcal{L} = \mathcal{X}^\circ$, $\max_{a \in \mathcal{A}, \ell \in \mathcal{L}} |M_a^\top \ell| \leq 1$. Hence, $\omega \leq 1$ (see its definition in Lemma 2) and L can be chosen as 1. Moreover, trivially, $\psi_a - \psi_b = M_a - M_b = \mathbf{M}(e_a - e_b)$, which shows that the game is locally observable. Finally, it is immediate that for any $\ell \in \mathcal{L}$ and $a^* \in \arg \min_{a \in \mathcal{A}} \psi_a^\top \ell$, $(e_a - e_{a^*})_{a \in \mathcal{A}^*} \in \mathcal{W}_\ell^{\text{loc}}$, and consequently, $\beta_{\text{loc}} := \max_{\ell \in \mathcal{L}} \min_{\lambda \in \mathcal{W}_\ell^{\text{loc}}} \max_{a \in \mathcal{A}^*} \|\lambda_a\|_1 \leq 2$. We then obtain via Theorem 1 (and Lemma 17) that $\Lambda_\eta^* \lesssim d$ (i.e., Λ_η^* is no larger than d up to a universal constant).

I.2 Linear dueling bandits

Recall that here, the action set has the form $\mathcal{A} = [m] \times [m]$ for some positive integer m . With a fixed feature mapping $a \mapsto \psi_a$ from $[m]$ to $\mathbb{R}^d$, we have that for $(a, b) \in \mathcal{A}$, $\psi_{a,b} := \psi_a + \psi_b$ and $M_{a,b} = \psi_a - \psi_b$. If $\mathcal{L} = ((\psi_a)_{a \in \mathcal{A}})^\circ$, then $\omega \leq 2$ and L can be taken as 2. As pointed out in [30], the Pareto optimal actions are those of the form (a, a) for $a \in [m]$. Fix $\ell \in \mathcal{L}$, and let $(a, b) \in \arg \min_{(a', b') \in \mathcal{A}} \langle \psi_{a', b'}, \ell \rangle$. (Here, a and b are not necessarily distinct and, in any case, satisfy $\psi_a = \psi_b = \min_{c \in [m]} \langle \psi_c, \ell \rangle$.) Now, for any $(c, c) \in \mathcal{A}$, we can write

$$\psi_{c,c} - \psi_{a,b} = \psi_c + \psi_c - \psi_a - \psi_b = M_{c,a} + M_{c,b}.$$

At the same time, $\langle \psi_{c,a} - \psi_{c,c}, \ell \rangle = \langle \psi_a - \psi_c, \ell \rangle \leq 0$ and $\langle \psi_{c,b} - \psi_{c,c}, \ell \rangle = \langle \psi_b - \psi_c, \ell \rangle \leq 0$. This shows that the game is locally observable (see Lemma 4) and that $\beta_{\text{loc}} \leq 2$ as in the linear bandit case. Hence, it also holds here that $\Lambda_\eta^* \lesssim d$.

I.3 Learning with graph feedback

The game here, to recall, is characterized by an undirected graph $\mathcal{G} = (\mathcal{A}, E)$ and we have that $d = k$, $\psi_a = e_a$, and $M_a \in \mathbb{R}^{k \times |N_{\mathcal{G}}(a)|}$ has the vectors $(e_b)_{b \in N_{\mathcal{G}}(a)}$ as columns, where $N_{\mathcal{G}}(a)$ is the neighborhood of a in $\mathcal{G}$ (which need not include a). In the terminology of [1], a graph is said to be weakly observable if $\forall a \in \mathcal{A}, \exists b \in \mathcal{A}: a \in N_{\mathcal{G}}(b)$, while in a strongly observable graph, it holds that $\forall a \in \mathcal{A}$, either $a \in N_{\mathcal{G}}(a)$, $a \in N_{\mathcal{G}}(b) \forall b \in \mathcal{A}$, or both. Here, and in all the examples to follow in this section, all actions are Pareto optimal. We consider $\mathcal{L}$ in this example to be $\mathcal{B}_\infty(1)$, which coincides with $((e_a)_{a \in [d]})^\circ$. Note that $\max_{a \in \mathcal{A}, \ell \in \mathcal{L}} \|M_a^\top \ell\| = \max_{a \in \mathcal{A}} \sqrt{|N_{\mathcal{G}}(a)|}$, which can be as large as $\sqrt{k}$, so we derive in the following a tighter bound on ω .

Recall from Lemma 2 that

$$\omega := \sup_{a,b,c \in \mathcal{A}, p \in \Delta_k, \ell \in \mathcal{L}} \frac{|(\psi_b - \psi_c)^\top Q_\delta(p)^\dagger M_a M_a^\top \ell|}{\|M_a^\top Q_\delta(p)^\dagger (\psi_b - \psi_c)\|}.$$

Now fix $a, b, c \in \mathcal{A}$, $p \in \Delta_k$, and $\ell \in \mathcal{L}$. Firstly, the Cauchy-Schwarz inequality gives that

$$|(\psi_b - \psi_c)^\top Q_\delta(p)^\dagger M_a M_a^\top \ell| \leq \|M_a^\top Q_\delta(p)^\dagger (\psi_b - \psi_c)\|_1 \|M_a^\top \ell\|_\infty.$$

Notice now that $\|M_a^\top \ell\|_\infty \leq \|\ell\|_\infty \leq 1$. By the structure of the observation matrices, $Q_\delta(p)$ is a diagonal matrix with $(P(a))_{a \in \mathcal{A}}$ as the diagonal entries, where

$$P(a) := \delta + (1 - \delta) \sum_{b \in N_G(a)} p(b).$$

Hence, $Q_\delta(p)^\dagger (\psi_b - \psi_c) = P(b)^{-1} e_b - P(c)^{-1} e_c$, and consequently, $M_a^\top Q_\delta(p)^\dagger (\psi_b - \psi_c)$ has at most two possibly non-zero entries:

$$\alpha := P(b)^{-1} \llbracket b \in N_G(a) \rrbracket \quad \text{and} \quad \beta := -P(c)^{-1} \llbracket c \in N_G(a) \rrbracket.$$

Then, using that $|\alpha| + |\beta| = \sqrt{(|\alpha| + |\beta|)^2} \leq \sqrt{2} \sqrt{\alpha^2 + \beta^2}$, we get that

$$\|M_a^\top Q_\delta(p)^\dagger (\psi_b - \psi_c)\|_1 \leq \sqrt{2} \|M_a^\top Q_\delta(p)^\dagger (\psi_b - \psi_c)\|,$$

which implies via the arguments above that

$$|(\psi_b - \psi_c)^\top Q_\delta(p)^\dagger M_a M_a^\top \ell| \leq \sqrt{2} \|M_a^\top Q_\delta(p)^\dagger (\psi_b - \psi_c)\|,$$

and consequently, $\omega \leq \sqrt{2}$. Alternatively, if $\mathcal{L} = [0, 1]^k$, the non-negativity of the losses and the structure of $M_a^\top Q_\delta(p)^\dagger (\psi_b - \psi_c)$ described above yield that

$$\begin{aligned} |(\psi_b - \psi_c)^\top Q_\delta(p)^\dagger M_a M_a^\top \ell| &\leq \max\{|\alpha|, |\beta|\} \\ &= \|M_a^\top Q_\delta(p)^\dagger (\psi_b - \psi_c)\|_\infty \leq \|M_a^\top Q_\delta(p)^\dagger (\psi_b - \psi_c)\|. \end{aligned}$$

Hence, it simply holds in this variant that $\omega \leq 1$.

Next, we study the problem-dependent terms appearing in the bounds of Theorem 1 and Theorem 2 in weakly and strongly observable graphs respectively, leading to a bound of order $\alpha(\mathcal{G})$ on Λ_η^* in the former and a bound of order $\sqrt{\delta(\mathcal{G})/\eta}$ in the latter.

Strongly observable graphs. Fix $\ell \in \mathcal{L}$, as mentioned in Section 4, we aim to find $\lambda \in \mathcal{W}_\ell^{\text{loc}}$ that is (essentially) supported on an independent set and satisfying $\beta_\lambda \leq 2$. Let $a^* \in \arg \min_{a \in \mathcal{A}} \langle \psi_a, \ell \rangle$. Notice here that $\mathcal{A}^* = \mathcal{A}$ and $\langle \psi_a, \ell \rangle = \ell(a)$. We construct such a λ by iteratively constructing two functions $g_1, g_2: \mathcal{A} \setminus \{a^*\} \rightarrow \mathcal{A}$ such that for all $a \in \mathcal{A} \setminus \{a^*\}$,

- i) $a \in N_G(g_1(a))$
- ii) $a^* \in N_G(g_2(a))$
- iii) $\ell(g_1(a)) \leq \ell(a)$
- iv) $\ell(g_2(a)) \leq \ell(a)$
- v) $|\text{img}(g_1) \cup \text{img}(g_2)| \leq \alpha(\mathcal{G}) + 1$.

Fix $a \in \mathcal{A} \setminus \{a^*\}$. Note that for $\lambda \in \mathcal{W}_\ell^{\text{loc}}$ and $b \in \mathcal{A}$, $\lambda_a(b) \in \mathbb{R}^{|N_G(b)|}$ with each coordinate corresponding to a neighbor of b . Since, by assumption, $a \in N_G(g_1(a))$, we set $\lambda_a(g_1(a)) = e_a$ where we abuse notation and use e_a in this expression as the indicator vector in $\mathbb{R}^{|N_G(b)|}$ for the coordinate corresponding to a . Similarly, set $\lambda_a(g_2(a)) = -e_{a^*}$ using that $a^* \in N_G(g_2(a))$. While for $b \notin \{g_1(a), g_2(a)\}$, set $\lambda_a(b) = \mathbf{0}$. Finally, we naturally set $\lambda_{a^*}(b) = \mathbf{0}$ for all $b \in \mathcal{A}$. Using the assumed properties above, we get that indeed $\psi_a - \psi_{a^*} = M\lambda_a$, $\lambda \in \mathcal{W}_\ell^{\text{loc}}$ (since $\ell(g_1(a)), \ell(g_2(a)) \leq \ell(a)$), $\beta_\lambda \leq 2$, and $\text{supp}(\lambda) \leq \alpha(\mathcal{G}) + 1$, as desired.

The two functions can be constructed using the iterative node elimination procedure detailed in Algorithm 3. This construction satisfies the required properties for g_1 and g_2 as listed above. In particular, $\text{img}(g_1)$ is an independent set, and $|\text{img}(g_2) \setminus \text{img}(g_1)| \leq 1$. The arguments laid so far

illustrate that the game is locally observable (see Lemma 4), and, moreover, that the first addend in the bound of Theorem 1 is of order $\alpha(\mathcal{G})$. What remains now is bounding w^* (note that $\beta_{\text{glo}} \leq 2$ using Lemma 17 and the fact that $\beta_{\text{loc}} \leq 2$). Let $S^* = \text{img}(g_1) \cup \text{img}(g_2)$, which is a total dominating set (that is, every node in the graph, including the ones in S^* , is connected to a node in S^*). Then,

$$\begin{aligned} w^* &:= \min_{S \subseteq \mathcal{A}} |S| \max_{b \in \mathcal{A}} \|M_b^\top (\sum_{s \in S} M_s M_s^\top)^{-1} M_b\| \\ &\leq |S^*| \max_{b \in \mathcal{A}} \|M_b^\top (\sum_{s \in S^*} M_s M_s^\top)^{-1} M_b\| \leq |S^*| \max_{b \in \mathcal{A}} \|M_b^\top M_b\| \leq |S^*| \leq \alpha(\mathcal{G}) + 1, \end{aligned}$$

where we have used that $\sum_{s \in S^*} M_s M_s^\top$ is a diagonal matrix where every entry on the diagonal is lower bounded by 1 (thanks to S^* being a total dominating set).

Algorithm 3 Constructing g_1 and g_2

input: undirected graph $\mathcal{G} = (\mathcal{A}, E)$, $\ell \in \mathbb{R}^d$, $a^* \in \arg \min_{a \in \mathcal{A}} \ell(a)$
output: $g_1, g_2: \mathcal{A} \setminus \{a^*\} \rightarrow \mathcal{A}$
initialize: $\mathcal{G}' \leftarrow \mathcal{G}$
while $\mathcal{G}'$ contains at least one vertex **do**
 if $V_{\mathcal{G}'} = \mathcal{A}$ **then** $b \leftarrow a^*$ **else** select $b \in \arg \min_{a \in V_{\mathcal{G}'}} \ell(a)$
 $\triangleright V_{\mathcal{G}'}$ denotes the set of all vertices in $\mathcal{G}'$
 set $g_1(a) = b$ for all $a \in N_{\mathcal{G}'}(b) \setminus \{a^*\}$
 $\triangleright$ note that $b \in N_{\mathcal{G}'}(b)$ if $b \neq a^*$, otherwise $b \in N_{\mathcal{G}}(a^*)$ implying it has already been eliminated
 $\mathcal{G}' \leftarrow \mathcal{G}'[V_{\mathcal{G}'} \setminus (N_{\mathcal{G}'}(b) \cup \{b\})]$ $\triangleright$ for $S \subseteq V_{\mathcal{G}'}$, $\mathcal{G}'[S]$ denotes the sub-graph induced by S
end while
if $a^* \in N_{\mathcal{G}}(a^*)$ **then**
 set $g_2(a) = a^*$ for all $a \in \mathcal{A} \setminus a^*$
else $\triangleright$ in this case, $N_{\mathcal{G}}(a^*) = \mathcal{A} \setminus \{a^*\}$ and $g_1(a) = a^*$ for all $a \in \mathcal{A} \setminus \{a^*\}$
 select $c \in \arg \min_{a \in \mathcal{A} \setminus a^*} \ell(a)$
 set $g_2(a) = c$ for all $a \in \mathcal{A} \setminus a^*$
end if
return g_1 and g_2

Weakly observable graphs. It is immediate to verify that weakly observable graphs are globally observable with $\beta_{\text{glo}} \leq 2$. Moreover, w^* is easily seen to be bounded by the total domination number $\delta(\mathcal{G})$ (size of a minimal total dominating set) by following the same argument laid above in the strongly observable case.

I.4 Full information

This is a special case of strongly observable graphs, with the graph being a clique (fully connected graph) with self loops. For this graph, $\alpha(\mathcal{G}) = 1$; hence, the arguments laid out in the previous example yield that Λ_η^* is simply upper bounded by a universal constant (independent of the number of actions). In fact, since M_a is the identity matrix for all a , it is easily seen that $w^* = 1$. Moreover, the construction we described for strongly observable graphs (Algorithm 3) trivially yields, for every $\ell \in \mathcal{L}$, weights $\lambda \in \mathcal{W}_\ell^{\text{loc}}$ that satisfy $\text{supp}(\lambda) = \{a^*\}$ (with $a^* \in \arg \min_{a \in \mathcal{A}} \psi_a^\top \ell$) and $\beta_\lambda = \max_{a \in \mathcal{A}} \|\lambda_a(a^*)\| = \max_{a \in \mathcal{A}} \|e_a - e_{a^*}\| = \sqrt{2}$.

I.5 Bandits with ill-conditioned observers

In this game: $d = k$, $\psi_a = e_a$, and $M_a = (1 - \varepsilon)\mathbf{1}_k/k + \varepsilon e_a$ for some $\varepsilon \in (0, 1]$. Moreover, we take $\mathcal{L} = \mathcal{B}_\infty(1)$. Since $M_a \in \Delta_k$ for all a , we have that $\|M_a\|_1 = 1$. Hence, the Cauchy-Schwarz inequality gives that $\max_{a \in \mathcal{A}, \ell \in \mathcal{L}} |M_a^\top \ell| \leq 1$, implying that $\omega \leq 1$. As noted in Section 4, $e_a - e_b = 1/\varepsilon(M_a - M_b)$ for any $a, b \in \mathcal{A}$. Hence, it holds for any $\ell \in \mathcal{L}$ and $a^* \in \arg \min_{a \in \mathcal{A}} \psi_a^\top \ell$ that $(e_a/\varepsilon - e_{a^*}/\varepsilon)_{a \in \mathcal{A}} \in \mathcal{W}_\ell^{\text{loc}}$, and consequently, that $\beta_{\text{loc}} \leq \max_{a, b \in \mathcal{A}} \|e_a/\varepsilon - e_b/\varepsilon\|_1 = 2/\varepsilon$. Theorem 1 then gives that $\Lambda_\eta^* \lesssim d/\varepsilon^2$.

I.6 Learning with composite graph feedback

This game is a variant of the graph feedback problem (on an undirected graph $G = (\mathcal{A}, E)$ with all self-loops) where, again, $d = k$ and $\psi_a = e_a$; however, the feedback vector M_a is the a -th row of the $k \times k$ matrix $D^{-1}A$; hence $\mathbf{M} = (D^{-1}A)^\top = AD^{-1}$. Here, D is the degree matrix and A is the adjacency matrix. Hence, after playing A_t , the learner observes the average of the losses $\{\ell_t(i) : i \in N_G(A_t)\}$, which include $\ell_t(A_t)$. As mentioned in Section 2, this feedback is less informative than standard graph feedback, and draws inspiration from problems studied in [55, 25] motivated by applications in signal processing (SNETs [19]). Once again we take $\mathcal{L} = \mathcal{B}_\infty(1)$. Note that $\|M_a\|_1 = 1$; hence, $\max_{a \in \mathcal{A}, \ell \in \mathcal{L}} |M_a^\top \ell| \leq 1$, implying again that $\omega \leq 1$. Depending on the structure of the graph, the game could be hopeless, globally or locally observable (or even trivial if there is just one action).

If the graph is empty of edges except self-loops (i.e., $\mathbf{M} = \mathbf{I}_k$), we recover the standard bandit game, which is locally observable. In fact, this is a necessary condition for locally observability. To see this, fix $\ell \in \mathcal{L}$ such that $\ell(a) < \ell(b) < \min_{c \in \mathcal{A} \setminus \{a, b\}} \ell(c)$. Then, $e_b - e_a \notin \text{span}\{M_a, M_b\}$ if a and b are connected in the graph, which rules out local observability via Lemma 4 (note that all actions are Pareto optimal). As for global observability, one needs that $e_b - e_a \in \text{span}\{M_c : c \in \mathcal{A}\}$ for all $a, b \in \mathcal{A}$. If we define $\mathcal{S} := \text{span}\{e_a - e_b : \text{for all } a, b \in \mathcal{A}\}$, global observability then translates to $\mathcal{S} \subseteq \text{col}(\mathbf{M})$. Now, suppose that $\det(\mathbf{M}) = 0$. This would imply that $\dim(\text{col}(\mathbf{M}))$ is at most $k - 1$. Then, since $\dim(\mathcal{S}) = k - 1$, global observability would require that $\dim(\text{col}(\mathbf{M})) = k - 1$ and $\mathcal{S} = \text{col}(\mathbf{M})$. This can never hold, however, because the vector $\mathbf{1}_k$ is orthogonal to $\mathcal{S}$, while any column of $\mathbf{M}$ has a non-zero component along $\mathbf{1}_k$, implying hence that $\text{col}(\mathbf{M}) \not\subseteq \mathcal{S}$. Therefore, in order to be globally observable, the game must satisfy that $\det(\mathbf{M}) \neq 0$, which is also clearly a sufficient condition.

Assume in what follows that the game is globally observable, meaning that $\mathbf{M}$ has full rank. As mentioned in Section 5, it holds here that $u^* = k$. To see this, note that Lemma 9 implies in this case that for any orthonormal basis U , $\sum_{s \in S} U^\top M_s M_s^\top U$ is invertible only when $S = \mathcal{A}$. Hence,

$$u^* = k \|\mathbf{M}^\top U (U^\top \mathbf{M} \mathbf{M}^\top U)^{-1} U^\top \mathbf{M}\| = k.$$

As for w^* , we know from Theorem 1 that $w^* \leq k$. Moreover, recalling that $(M_a)_{a \in \mathcal{A}}$ are vectors in this example, we have that

$$w^* \geq \min_{\pi \in \Delta_k} \max_{b \in \mathcal{A}} M_b^\top U B(\pi)^{-1} U^\top M_b = k,$$

where the equality follows from Lemma 12 using that $\text{rk}(\mathbf{M}) = k$ by assumption. Hence, it holds too that $w^* = k$. Theorem 2 then yields a bound of order $\beta_{2, \text{glo}} \sqrt{k/\eta}$ on Λ_η^* .

We now examine the two quantities β_{glo} and $\beta_{2, \text{glo}}$ in this problem. The assumed Invertibility of $\mathbf{M}$ implies that

$$\beta_{\text{glo}} = \max_{a, b \in \mathcal{A}} \|\mathbf{M}^{-1}(e_a - e_b)\|_1 \quad \text{and} \quad \beta_{2, \text{glo}} = \max_{a, b \in \mathcal{A}} \|\mathbf{M}^{-1}(e_a - e_b)\|_2.$$

It holds here, as it does in general, that $\beta_{2, \text{glo}} \leq \beta_{\text{glo}} \leq \sqrt{k} \beta_{2, \text{glo}}$. Recall that $\mathbf{M} = AD^{-1}$ and let $S := D^{-1/2}AD^{-1/2}$. Note that S is symmetric and $\mathbf{M}$ could be rewritten in terms of S as $\mathbf{M} = D^{1/2}SD^{-1/2}$. Further, $\mathbf{M}$ and S are similar, thus they have the same eigenvalues. For a square matrix V , let $\sigma(V)$ denote its smallest singular value, and let $\{w_i(V)\}_{i=1, \dots, k}$ denote its eigenvalues. Lastly, let $\deg(a)$ denote the degree of action a , and let $\deg_{\min}$ and $\deg_{\max}$ denote, respectively, the smallest and the largest degree of any action in the graph. Now,

$$\begin{aligned} \beta_{2, \text{glo}} &= \max_{a, b \in \mathcal{A}} \|\mathbf{M}^{-1}(e_a - e_b)\|_2 = \max_{a, b \in \mathcal{A}} \|D^{1/2}S^{-1}D^{-1/2}(e_a - e_b)\|_2 \\ &\leq \sqrt{\frac{\deg_{\max}}{\deg(a)} + \frac{\deg_{\max}}{\deg(b)}} \|S^{-1}\|_2 \leq \frac{\sqrt{2\deg_{\max}/\deg_{\min}}}{\sigma(S)} \\ &= \frac{\sqrt{2\deg_{\max}/\deg_{\min}}}{\min_{i=1, \dots, k} |w_i(S)|} = \frac{\sqrt{2\deg_{\max}/\deg_{\min}}}{\min_{i=1, \dots, k} |w_i(\mathbf{M})|}, \end{aligned}$$

where the penultimate step follows from the symmetry of S , and the last step uses that $\mathbf{M}$ and S are similar. Note also that $w_i(S) = 1 - w_i(L^{\text{sym}})$ where $L^{\text{sym}} = \mathbf{I}_k - S$ is the symmetrically

normalized Laplacian of the graph $\mathcal{G}$. Hence, it also holds that

$$\beta_{2,\text{glo}} \leq \frac{\sqrt{2\deg_{\max}/\deg_{\min}}}{\min_{i=1,\dots,k} |1 - w_i(L^{\text{sym}})|}.$$

I.6.1 Cycle graphs

As an example, take $\mathcal{G}$ to be a k -nodes cycle graph (with self-loops), where $k \geq 3$. For convenience, we index actions starting from 0 instead of 1 in the following discussion; that is we take $\mathcal{A} = \{0, \dots, k-1\}$. We also extend this notation to vectors and matrices indexed by the actions. For this family of graphs, we have that

$$A_{a,b} = \mathbb{I}[b \in \{(a-1) \bmod k, a, (a+1) \bmod k\}] \quad \text{and} \quad \mathbf{M} = \frac{1}{3}A.$$

If $k \bmod 3 = 0$, then $\mathbf{M}$ is singular, which is synonymous with a hopeless (i.e., not globally observable) game in this problem as argued before. A simple non-zero vector v that belongs to $\ker(\mathbf{M})$ in this case is given by

$$v(a) = \begin{cases} 2, & \text{if } a \bmod 3 = 0 \\ -1, & \text{otherwise} \end{cases}$$

for $a \in \mathcal{A}$, see Graph (i) in Figure 1 for a schematic representation on a 9-node cycle. When $k \bmod 3 \in \{1, 2\}$, $\mathbf{M}$ is invertible and the game is globally observable. In particular, for $a, b \in \mathcal{A}$, one can verify that

$$\mathbf{M}_{a,b}^{-1} = \frac{2}{\sqrt{3}} \sin\left(\frac{2\pi k}{3}\right) \cos\left(\frac{2\pi h(a,b)}{3}\right) + \sqrt{3} \sin\left(\frac{2\pi h(a,b)}{3}\right),$$

where $h(a,b) = (a-b) \bmod k$. We now characterize β_{glo} and $\beta_{2,\text{glo}}$ up to small constants, focusing on the case when $k \bmod 3 = 1$. (The remaining case, when $k \bmod 3 = 2$, can be treated in a similar manner.) Firstly, via the triangle's inequality,

$$\beta_{2,\text{glo}} = \max_{a,b \in \mathcal{A}} \|\mathbf{M}^{-1}(e_a - e_b)\|_2 \leq 2 \max_{a \in \mathcal{A}} \|\mathbf{M}^{-1}e_a\|_2.$$

And similarly, $\beta_{\text{glo}} \leq 2 \max_{a \in \mathcal{A}} \|\mathbf{M}^{-1}e_a\|_1$. As the entries of $\mathbf{M}^{-1}$ are determined by $h(a,b)$, all its columns are identical up to cyclical shifts. Hence, the bound above is attained by any action. From the formula stated above for $\mathbf{M}^{-1}$, one can check that each of its columns contains $1 + 2(k-1)/3$ entries with value 1 and $(k-1)/3$ entries with value -2 , Graph (ii) in Figure 1 provides a schematic representation of one column of $\mathbf{M}^{-1}$ on a 10-node cycle. Evaluating the L_1 and L_2 norms we obtain

$$\beta_{2,\text{glo}} \leq 2\sqrt{2k-1} \quad \text{and} \quad \beta_{\text{glo}} \leq 2 + 8(k-1)/3.$$

On the other hand, we can obtain a lower bound by taking a and b as neighbors. In which case, $\mathbf{M}^{-1}(e_a - e_b)$ contains $(k-1)/3$ entries with value 3, $(k-1)/3$ entries with value -3 , and $1 + (k-1)/3$ entries with value 0, a visual representation is provided by Graph (iii) in Figure 1. Then, again evaluating the L_1 and L_2 norms gives that

$$\beta_{2,\text{glo}} \geq \sqrt{6(k-1)} \quad \text{and} \quad \beta_{\text{glo}} \geq 2(k-1).$$

This shows that β_{glo} is $\Theta(k)$, while $\beta_{2,\text{glo}}$ is $\Theta(\sqrt{k})$, providing thus an appreciable improvement relative to the former in the bound of Theorem 2 (recalling that here, $w^* = u^* = k$).

I.6.2 Complete bipartite graphs

As a second example, we consider a complete balanced bipartite graph $G_{k/2,k/2}$ (with self loops), assuming k is even. The vertices of this graph are partitioned into two subsets of size $k/2$; $L = \{1, \dots, k/2\}$ and $R = \{k/2 + 1, \dots, k\}$ such that for $a, b \in \mathcal{A}$ with $a < b$,

$$A_{a,b} = A_{b,a} = \mathbb{I}[a \in L \text{ and } b \in R],$$

while $A_{a,a} = 1$ for all $a \in \mathcal{A}$. Hence, all nodes in this graph have degree $k/2 + 1$ and

$$\mathbf{M} = \frac{1}{k/2 + 1}A = \frac{1}{k/2 + 1} \begin{pmatrix} \mathbf{I}_{k/2} & \mathbf{J}_{k/2} \\ \mathbf{J}_{k/2} & \mathbf{I}_{k/2} \end{pmatrix}.$$




Figure 1: The first figure represents a (non-zero) assignment of weights on a 9-node cycle such that the sum at any three consecutive nodes is 0, illustrating that M is singular. On a 10-node cycle, the second figure provides a representation of $M^{-1}e_a$; i.e., the a -th column of M^{-1} , where a is an arbitrary action in $\mathcal{A}$. The last figure is a representation of $M^{-1}(e_a - e_b)$ on the same graph, where b is a neighbor of a .

where $I_{k/2}$ is the $(k/2) \times (k/2)$ identity matrix and $J_{k/2}$ is the $(k/2) \times (k/2)$ matrix with all entries equal to 1. In a similar manner, we can write any vector $x \in \mathbb{R}^k$ in the block form

$$x = \begin{pmatrix} x_L \\ x_R \end{pmatrix}$$

with $x_L, x_R \in \mathbb{R}^{k/2}$. If $k = 2$, then M is singular and the game is hopeless. Otherwise, when $k \geq 4$, M is full rank and its inverse is given by

$$M^{-1} = \frac{1}{k/2 - 1} \begin{pmatrix} ((k^2/4 - 1)I_{k/2} - (k/2)J_{k/2}) & J_{k/2} \\ J_{k/2} & (k^2/4 - 1)I_{k/2} - (k/2)J_{k/2} \end{pmatrix}.$$

If two actions a and b are on the same side; i.e., $a, b \in L$ or $a, b \in R$, then

$$M^{-1}(e_a - e_b) = (k/2 + 1)(e_a - e_b).$$

Otherwise, if $a \in L$ and $b \in R$, then

$$M^{-1}(e_a - e_b) = (k/2 + 1) \left(e_a - e_b + \frac{1}{k/2 - 1} \begin{pmatrix} -\mathbf{1}_{k/2} \\ \mathbf{1}_{k/2} \end{pmatrix} \right).$$

In either case, both $\|M^{-1}(e_a - e_b)\|_2$ and $\|M^{-1}(e_a - e_b)\|_1$ are $\Theta(k)$. Hence, both $\beta_{2,\text{glo}}$ and β_{glo} are $\Theta(k)$. Compared to the previous example, $\beta_{2,\text{glo}}$ grows faster with the number of actions, while β_{glo} remains of the same order.

J Lower Bounds

In this section, we provide a proof for Proposition 1. This is an extension of the lower bounds in [29, Appendix G], with a construction that is compatible with the adversarial setting. The proof still relies on designing hard instances in a stochastic linear partial monitoring framework. However, the noise is added to the loss vector itself. This noisy loss vector, moreover, is designed to respect a boundedness condition. The latter is achieved via truncation, following [16], at the cost of a factor logarithmic in the time horizon. We also provide a tailored lower bound for the bandit game with ill-conditioned observers that matches the upper bound we achieved.

J.1 Stochastic Linear Partial Monitoring with Parameter Noise

We will assume henceforth that, for some positive integer m , $n(a) = m$ for all $a \in \mathcal{A}$; this can be taken as $\max_a n(a)$ while padding with zeros the matrices with fewer columns. Moreover, as

mentioned in Section 2, we assume that $\mathcal{B}_2(r) \subseteq \mathcal{L}$ for some $r > 0$. Let $\theta \in \mathbb{R}^d$. We will use θ to parameterize a stochastic linear partial monitoring problem as follows. At round t , the player select a (possibly random) action $A_t \in \mathcal{A}$ and observes the signal $\phi_t = M_{A_t}^\top \ell_t \in \mathbb{R}^m$, where $\ell_t \in \mathbb{R}^d$ is given by

$$\ell_t = \theta + X_t,$$

with X_t drawn i.i.d. across rounds from a Gaussian distribution $\mathcal{N}(\mathbf{0}, \Sigma)$, where $\Sigma \in \mathbb{R}^{d \times d}$ is some covariance matrix. Extending [16] (see also [18, Theorem 9]), we describe in the following a truncation scheme that allows adapting to the adversarial setting lower bound constructions that employ unbounded noise. Let $\text{clip}: \mathbb{R}^d \rightarrow \mathcal{L}$ denote some clipping operation. Consider the following two notions of (random) regret:

$$\begin{aligned}\widehat{R}_T(\theta) &:= \max_{a \in \mathcal{A}} \sum_{t=1}^T (\psi_{A_t} - \psi_a)^\top \ell_t \\ \widetilde{R}_T(\theta) &:= \max_{a \in \mathcal{A}} \sum_{t=1}^T (\psi_{A_t} - \psi_a)^\top \text{clip}(\ell_t),\end{aligned}$$

where $(\ell_t)_t$ are drawn as described above and the actions $(A_t)_t$ are chosen via some policy $(\pi_t)_t$. Moreover, define the stochastic regret $\bar{R}_T(\theta)$ as:

$$\bar{R}_T(\theta) := \max_{a \in \mathcal{A}} \mathbb{E}_\theta \sum_{t=1}^T (\psi_{A_t} - \psi_a)^\top \ell_t = \max_{a \in \mathcal{A}} \mathbb{E}_\theta \sum_{t=1}^T (\psi_{A_t} - \psi_a)^\top \theta,$$

where the expectation is taken over the aforementioned sequences $(\ell_t)_t$ and $(A_t)_t$. (The dependence on $(\pi_t)_t$ is not explicit in the notation since it will be fixed throughout.) The equality in the above display follows from the fact that X_t is centered and independent from A_t .

Lemma 18. *Let Θ be a finite subset of $\mathbb{R}^d$ with $|\Theta| = N$, and let $\alpha^* := \max_{a,b \in \mathcal{A}} \|\psi_a - \psi_b\|$. Then,*

$$R_T^* \geq \frac{1}{N} \sum_{\theta \in \Theta} \bar{R}_T(\theta) - \frac{\sqrt{2}\alpha^* T^{3/2}}{N} \sum_{\theta \in \Theta} (\|\theta\| + \sqrt{\mathbb{E}\|X_1\|^2}) \sqrt{\mathbb{P}_\theta(\text{clip}(\ell_1) \neq \ell_1)}.$$

While if there exists some $\mathcal{R} > 0$ such that $\max_{\theta \in \Theta} \widehat{R}_T(\theta) \leq \mathcal{R}$ uniformly, then

$$R_T^* \geq \frac{1}{N} \sum_{\theta \in \Theta} \bar{R}_T(\theta) - \frac{\mathcal{R}T}{N} \sum_{\theta \in \Theta} \mathbb{P}_\theta(\text{clip}(\ell_1) \neq \ell_1).$$

Proof. Firstly, note that

$$R_T^* \geq \sup_{\theta} \mathbb{E}_\theta \widetilde{R}_T(\theta) \geq \frac{1}{N} \sum_{\theta \in \Theta} \mathbb{E}_\theta \widetilde{R}_T(\theta) \quad \text{and} \quad \mathbb{E} \widehat{R}_T(\theta) \geq \bar{R}_T(\theta),$$

the latter holding via Jensen's inequality. Define the event $B := \{\forall t \in [T], \text{clip}(\ell_t) = \ell_t\}$. Clearly, $\widetilde{R}_T(\theta) = \widehat{R}_T(\theta)$ whenever B occurs. Hence,

$$\mathbb{E}_\theta \widehat{R}_T(\theta) = \mathbb{E}_\theta [\mathbb{I}[B] \widehat{R}_T(\theta)] + \mathbb{E}_\theta [\mathbb{I}[\bar{B}] \widehat{R}_T(\theta)] \leq \mathbb{E}_\theta \widetilde{R}_T(\theta) + \mathbb{E}_\theta [\mathbb{I}[\bar{B}] \widehat{R}_T(\theta)].$$

Rather crudely, we have that

$$\begin{aligned}\widehat{R}_T(\theta) &= \max_{a \in \mathcal{A}} \sum_{t=1}^T (\psi_{A_t} - \psi_a)^\top (X_t + \theta) \\ &\leq \max_{a \in \mathcal{A}} \sum_{t=1}^T \|\psi_{A_t} - \psi_a\| \|X_t + \theta\| \\ &\leq \max_{a,b \in \mathcal{A}} \|\psi_a - \psi_b\| \sum_{t=1}^T \|X_t + \theta\| \\ &\leq \max_{a,b \in \mathcal{A}} \|\psi_a - \psi_b\| \left(\|\theta\|T + \sum_{t=1}^T \|X_t\| \right).\end{aligned}$$

And,

$$\begin{aligned}
\widehat{R}_T^2(\theta) &\leq \max_{a,b \in \mathcal{A}} \|\psi_a - \psi_b\|^2 \left(\|\theta\|T + \sum_{t=1}^T \|X_t\| \right)^2 \\
&\leq 2 \max_{a,b \in \mathcal{A}} \|\psi_a - \psi_b\|^2 \left(\|\theta\|^2 T^2 + \left(\sum_{t=1}^T \|X_t\| \right)^2 \right) \\
&\leq 2 \max_{a,b \in \mathcal{A}} \|\psi_a - \psi_b\|^2 \left(\|\theta\|^2 T^2 + T \sum_{t=1}^T \|X_t\|^2 \right).
\end{aligned}$$

Now, via the Cauchy-Schwarz inequality,

$$\begin{aligned}
\mathbb{E}_\theta [\llbracket \bar{B} \rrbracket \widehat{R}_T(\theta)] &\leq \sqrt{\mathbb{E}_\theta [\llbracket \bar{B} \rrbracket^2] \mathbb{E}_\theta [\widehat{R}_T^2(\theta)]} \\
&= \sqrt{\mathbb{E}_\theta [\llbracket \bar{B} \rrbracket] \mathbb{E}_\theta [\widehat{R}_T^2(\theta)]} \\
&= \sqrt{\mathbb{P}_\theta(\bar{B}) \mathbb{E}_\theta [\widehat{R}_T^2(\theta)]} \\
&\leq \sqrt{2} \max_{a,b \in \mathcal{A}} \|\psi_a - \psi_b\| \sqrt{\mathbb{P}_\theta(\bar{B})} \left(\|\theta\|T + \sqrt{T} \sqrt{\sum_{t=1}^T \mathbb{E} \|X_t\|^2} \right) \\
&\leq \sqrt{2} \max_{a,b \in \mathcal{A}} \|\psi_a - \psi_b\| (\|\theta\| + \sqrt{\mathbb{E} \|X_1\|^2}) \sqrt{\mathbb{P}_\theta(\bar{B})} T,
\end{aligned}$$

Where we have used the fact that $(X_t)_t$ is an i.i.d. sequence. Moreover, a union bound gives that

$$\mathbb{P}_\theta(\bar{B}) \leq \sum_{t=1}^T \mathbb{P}_\theta(\text{clip}(\ell_t) \neq \ell_t) = T \mathbb{P}_\theta(\text{clip}(\ell_1) \neq \ell_1).$$

Let $\alpha^* := \max_{a,b \in \mathcal{A}} \|\psi_a - \psi_b\|$. We have shown that

$$\mathbb{E}_\theta [\llbracket \bar{B} \rrbracket \widehat{R}_T(\theta)] \leq \sqrt{2} \alpha^* (\|\theta\| + \sqrt{\mathbb{E} \|X_1\|^2}) \sqrt{\mathbb{P}_\theta(\text{clip}(\ell_1) \neq \ell_1)} T^{3/2}.$$

In the simpler case when $\max_{\theta \in \Theta} \widehat{R}_T(\theta) \leq \mathcal{R}$ uniformly, we can similarly show that

$$\mathbb{E}_\theta [\llbracket \bar{B} \rrbracket \widehat{R}_T(\theta)] \leq \mathbb{P}_\theta(\text{clip}(\ell_1) \neq \ell_1) \mathcal{R} T.$$

The theorem then follows using that

$$\mathbb{E}_\theta \widetilde{R}_T(\theta) \geq \mathbb{E}_\theta \widehat{R}_T(\theta) - \mathbb{E}_\theta [\llbracket \bar{B} \rrbracket \widehat{R}_T(\theta)] \geq \bar{R}_T(\theta) - \mathbb{E}_\theta [\llbracket \bar{B} \rrbracket \widehat{R}_T(\theta)].$$

□

In the following, we will generally use $\mathbb{P}_\theta$ to denote the probability measure over the history of interaction $\mathcal{H}_T := (A_s, \phi_s)_{s=1}^T$ induced by our fixed policy $(\pi_t)_t$ and the environment parameterized by θ . Moreover, $\mathbb{P}_{\theta,a}$ will denote the distribution over observations when action a is played. This is stationary (does not depend on t) since $(\ell_t)_t$ is an i.i.d. sequence. Let $D_{\text{KL}}(P \parallel Q)$ denote the KL-divergence between two measure P and Q . The following is a standard result based on the chain rule for the KL-divergence; see, e.g., [33, Exercise 15.8].

Lemma 19. *For any $\theta, \theta' \in \mathbb{R}^d$, it holds that*

$$D_{\text{KL}}(\mathbb{P}_\theta \parallel \mathbb{P}_{\theta'}) = \sum_{t=1}^T \mathbb{E}_\theta [D_{\text{KL}}(\mathbb{P}_{\theta,A_t} \parallel \mathbb{P}_{\theta',A_t})] = \sum_{a \in \mathcal{A}} T_\theta(a) D_{\text{KL}}(\mathbb{P}_{\theta,A_t} \parallel \mathbb{P}_{\theta',A_t}),$$

where $T_\theta(a) := \mathbb{E}_\theta \sum_{t=1}^T \llbracket A_t = a \rrbracket$.

As an extra piece of helpful notation, for $S \subseteq \mathcal{A}$, let $T_\theta(S) := \mathbb{E}_\theta \sum_{t=1}^T \llbracket A_t \in S \rrbracket$.

J.2 Regime-based lower bounds

Besides the use of truncation and parameter noise, the constructions used in the following three lower bounds are adapted from [29, Appendix G] and [33, Chapter 37]. Analogous constructions appear also in earlier works on finite partial monitoring like [3, 6].

Theorem 3. *Assume that $\mathcal{B}_2(r) \subseteq \mathcal{L}$ for some $r > 0$. If the game is globally but not locally observable then R_T^* is $\tilde{\Omega}(T^{2/3})$.*

Proof. By Definition 1 and Lemma 5, there exist at least three non-duplicate Pareto optimal actions, two of which, say a and b , are neighbors and satisfy

$$\psi_a - \psi_b \notin \text{span}\{M_c : c \in \mathcal{N}_{a,b}\}.$$

Hence, we can write $\psi_a - \psi_b = u + v$ with u belonging to $\text{span}\{M_c : c \in \mathcal{N}_{a,b}\}$ and $v \neq 0$ belonging to its orthogonal complement. Let $\theta \in \text{relint}(C_a \cap C_b)$. Then $\mathcal{P}(\theta)$ coincides with $\mathcal{N}_{a,b}$; in particular, there exists a game dependent $\varepsilon > 0$ such that any $c \notin \mathcal{N}_{a,b}$ satisfies $\langle \psi_c - \psi_a, \theta \rangle \geq \varepsilon$. Moreover, we can assume that $\|\theta\| = r/4$ ($C_a \cap C_b$ is a cone). Let $0 < \Delta \leq 1$ be a constant to be chosen to later, and let $q := \frac{r}{4\|v\|}v$. Now, define

$$\theta_a = \theta - \Delta q \quad \text{and} \quad \theta_b = \theta + \Delta q,$$

and notice that

$$\langle \psi_a - \psi_b, \theta - \Delta q \rangle = -\frac{\Delta r}{4\|v\|} \langle \psi_a - \psi_b, v \rangle = -\frac{\Delta r\|v\|^2}{4\|v\|} = -\frac{\Delta r\|v\|}{4} < 0.$$

Similarly, $\langle \psi_b - \psi_a, \theta + \Delta q \rangle = -\frac{\Delta r\|v\|}{4} < 0$. For brevity, let $\tilde{\Delta} := \frac{\Delta r\|v\|}{4}$. Note that any $c \in \mathcal{N}_{a,b}$ satisfies that $c \in \mathcal{P}(\theta)$. Hence,

$$\langle \psi_c - \psi_a, \theta - \Delta q \rangle + \langle \psi_c - \psi_b, \theta + \Delta q \rangle = \langle \psi_c - \psi_b - \psi_c + \psi_a, \Delta q \rangle = \tilde{\Delta}.$$

Therefore, $\max\{\langle \psi_c - \psi_a, \theta_a \rangle, \langle \psi_c - \psi_b, \theta_b \rangle\} \geq \tilde{\Delta}/2$. In particular, define

$$S_a := \{c \in \mathcal{N}_{a,b} \mid \langle \psi_c - \psi_a, \theta_a \rangle \geq \tilde{\Delta}/2\}, \quad S_b := \{c \in \mathcal{N}_{a,b} \mid \langle \psi_c - \psi_a, \theta_a \rangle < \tilde{\Delta}/2\} = \mathcal{N}_{a,b} \setminus S_a.$$

Also note that $\min\{\langle \psi_c - \psi_a, \theta_a \rangle, \langle \psi_c - \psi_b, \theta_b \rangle\} \geq 0$ for any $c \in \mathcal{N}_{a,b}$ since $\psi_c \in [\psi_a, \psi_b]$. Moving over to actions $c \notin \mathcal{N}_{a,b}$, we have that

$$\langle \psi_c - \psi_a, \theta - \Delta q \rangle \geq \varepsilon - \langle \psi_c - \psi_a, \Delta q \rangle \geq \varepsilon - \Delta \|\psi_c - \psi_a\| \|q\| \geq \varepsilon - \frac{r\Delta\alpha^*}{4},$$

where $\alpha^* := \max_{c,d \in \mathcal{A}} \|\psi_c - \psi_d\| > 0$. Assume that $\Delta \leq \frac{2\varepsilon}{r\alpha^*}$; thus, $\langle \psi_c - \psi_a, \theta_a \rangle \geq \varepsilon/2$, and similarly, $\langle \psi_c - \psi_b, \theta_b \rangle \geq \varepsilon/2$. Moreover, since $\|v\| \leq \alpha^*$, we have that $\Delta \leq \frac{2\varepsilon}{r\|v\|}$, implying that $\tilde{\Delta} \leq \varepsilon/2$. For brevity, let $\bar{\mathcal{N}}_{a,b} = \mathcal{A} \setminus \mathcal{N}_{a,b}$. Overall, we have that $\theta_a \in C_a$ and $\theta_b \in C_b$. In particular,

$$\bar{R}_T(\theta_a) = \mathbb{E}_{\theta_a} \sum_{t=1}^T (\psi_{A_t} - \psi_a)^\top \theta_a \geq \frac{\varepsilon}{2} T_{\theta_a}(\bar{\mathcal{N}}_{a,b}) + \frac{\tilde{\Delta}}{2} T_{\theta_a}(S_a),$$

and,

$$\begin{aligned} \bar{R}_T(\theta_b) &= \mathbb{E}_{\theta_b} \sum_{t=1}^T (\psi_{A_t} - \psi_b)^\top \theta_b \geq \frac{\varepsilon}{2} T_{\theta_b}(\bar{\mathcal{N}}_{a,b}) + \frac{\tilde{\Delta}}{2} T_{\theta_b}(S_b) \\ &\geq \frac{\tilde{\Delta}}{2} T_{\theta_b}(\bar{\mathcal{N}}_{a,b} \cup S_b) \\ &= \frac{\tilde{\Delta}}{2} (T - T_{\theta_b}(S_a)) \\ &= \frac{\tilde{\Delta}}{2} (T - T_{\theta_a}(S_a)) - \frac{\tilde{\Delta}}{2} (T_{\theta_b}(S_a) - T_{\theta_a}(S_a)) \\ &\geq \frac{\tilde{\Delta}}{2} (T - T_{\theta_a}(S_a)) - \frac{\tilde{\Delta}}{2} T \sqrt{(1/2) D_{\text{KL}}(\mathbb{P}_{\theta_a} \parallel \mathbb{P}_{\theta_b})}, \end{aligned}$$

where the last inequality follows from [33, Exercise 14.4] and Pinsker's inequality. Hence,

$$\bar{R}_T(\theta_a) + \bar{R}_T(\theta_b) \geq \frac{\varepsilon}{2} T_{\theta_a}(\bar{\mathcal{N}}_{a,b}) + \frac{\tilde{\Delta}}{2} T \left(1 - \sqrt{(1/2) D_{\text{KL}}(\mathbb{P}_{\theta_a} \parallel \mathbb{P}_{\theta_b})} \right).$$

Now, let $X_t \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_d)$. Recall that, $\ell_t = \theta_a + X_t$ under environment θ_a , and $\ell_t = \theta_b + X_t$ under environment θ_b . For any action c , it then holds that $\mathbb{P}_{\theta_a, c}$ is $\mathcal{N}(M_c^\top \theta_a, \sigma^2 M_c^\top M_c)$, and that

$$D_{\text{KL}}(\mathbb{P}_{\theta_a, c} \parallel \mathbb{P}_{\theta_b, c}) = D_{\text{KL}}\left(\mathcal{N}(M_c^\top \theta_a, \sigma^2 M_c^\top M_c) \parallel \mathcal{N}(M_c^\top \theta_b, \sigma^2 M_c^\top M_c)\right).$$

Note that the first distribution is supported on $M_c^\top \theta_a + \text{col}(M_c^\top M_c) = \text{col}(M_c^\top)$, which coincides with the support of the second. Hence, the KL-divergence is finite and can be written as (using, e.g., Equation 3.1 by Holbrook [27]):

$$\begin{aligned} D_{\text{KL}}(\mathbb{P}_{\theta_a, c} \parallel \mathbb{P}_{\theta_b, c}) &= \frac{1}{2\sigma^2} (M_c^\top \theta_a - M_c^\top \theta_b)^\top (M_c^\top M_c)^\dagger (M_c^\top \theta_a - M_c^\top \theta_b) \\ &= \frac{1}{2\sigma^2} (\theta_a - \theta_b)^\top M_c (M_c^\top M_c)^\dagger M_c^\top (\theta_a - \theta_b) \\ &= \frac{1}{2\sigma^2} (\theta_a - \theta_b)^\top M_c M_c^\dagger (\theta_a - \theta_b) \\ &= \frac{1}{2\sigma^2} (-2\Delta q)^\top M_c M_c^\dagger (-2\Delta q) \\ &= \frac{2\Delta^2}{\sigma^2} q^\top M_c M_c^\dagger q = \frac{2\Delta^2}{\sigma^2} \|M_c M_c^\dagger q\|^2, \end{aligned}$$

where we have used that $M_c M_c^\dagger$ is the orthogonal projection onto $\text{col}(M_c)$. This also means that if $c \in \mathcal{N}_{a,b}$, then $M_c M_c^\dagger q = 0$ by the construction of q . Otherwise, we simply have that $\|M_c M_c^\dagger q\| \leq \|q\| = r/4$. Hence, via Lemma 19, we have that

$$D_{\text{KL}}(\mathbb{P}_{\theta_a} \parallel \mathbb{P}_{\theta_b}) \leq \frac{r^2 \Delta^2}{8\sigma^2} T_{\theta_a}(\bar{\mathcal{N}}_{a,b}) = \frac{2\tilde{\Delta}^2}{\sigma^2 \|v\|^2} T_{\theta_a}(\bar{\mathcal{N}}_{a,b}).$$

Which gives that

$$\bar{R}_T(\theta_a) + \bar{R}_T(\theta_b) \geq \frac{\varepsilon}{2} T_{\theta_a}(\bar{\mathcal{N}}_{a,b}) + \frac{\tilde{\Delta}}{2} T \left(1 - \frac{\tilde{\Delta}}{\sigma \|v\|} \sqrt{T_{\theta_a}(\bar{\mathcal{N}}_{a,b})} \right).$$

Now, set $\Delta = 2\sigma T^{-1/3}/r$, assuming T is large enough so that $\Delta \leq \min\{1, \frac{2\varepsilon}{r\alpha^*}\}$. Precisely, assuming that $T^{1/3} \geq \sigma \max\{2/r, \alpha^*/\varepsilon\}$. We then have that $\tilde{\Delta} = \sigma \|v\| T^{-1/3}/2$. Now, if $T_{\theta_a}(\bar{\mathcal{N}}_{a,b}) \leq T^{2/3}$, then,

$$\bar{R}_T(\theta_a) + \bar{R}_T(\theta_b) \geq \frac{\tilde{\Delta}}{2} T \left(1 - \frac{\tilde{\Delta}}{\sigma \|v\|} T^{1/3} \right) = \frac{\tilde{\Delta}}{4} T = \frac{\sigma \|v\|}{8} T^{2/3}.$$

Otherwise, we directly get that $\bar{R}_T(\theta_a) + \bar{R}_T(\theta_b) \geq \frac{\varepsilon}{2} T^{2/3}$. Hence, in total, we have that

$$\frac{1}{2} (\bar{R}_T(\theta_a) + \bar{R}_T(\theta_b)) \geq \min\left\{ \frac{\sigma \|v\|}{16}, \frac{\varepsilon}{4} \right\} T^{2/3}.$$

We now prepare for an application of Lemma 18. First, notice that $\mathbb{E}\|X_1\|^2 = \sigma^2 d$. Second, pick the clipping operator as

$$\text{clip}(\ell_t) := \mathbb{I}[\|\ell_t\| \leq r] \ell_t + \mathbb{I}[\|\ell_t\| > r] \frac{r \ell_t}{\|\ell_t\|}.$$

Then, for $\theta' \in \{\theta_a, \theta_b\}$,

$$\begin{aligned} \mathbb{P}_{\theta'}(\text{clip}(\ell_1) \neq \ell_1) &= \mathbb{P}_{\theta'}(\|\ell_t\| > r) \leq \mathbb{P}_{\theta'}(\|\theta'\| + \|X_1\| > r) \\ &\leq \mathbb{P}_{\theta'}(\|X_1\| > r/2) \\ &= \mathbb{P}_{\theta'}(\|X_1\|^2 > r^2/4), \end{aligned}$$

which uses that $\|\theta_a\| \leq \|\theta\| + \Delta\|q\| \leq r/2$, with the same holding for θ_b . Lemma 1 by Laurent and Massart [39] gives that for $x > 0$,

$$\mathbb{P}(\|X_1\|^2 > \sigma^2 d + 2\sigma^2 \sqrt{dx} + 2\sigma^2 x) \leq e^{-x}.$$

Assuming $\sigma^2 \leq r^2/(8d)$, setting $x = r^2/(80\sigma^2)$ gives that

$$\mathbb{P}(\|X_1\|^2 > r^2/4) \leq e^{-r^2/(80\sigma^2)}.$$

Then, Lemma 18 yields that

$$\begin{aligned} R_T^* &\geq \frac{1}{2} \sum_{\theta' \in \{\theta_a, \theta_b\}} \bar{R}_T(\theta') - \frac{\sqrt{2}\alpha^* T^{3/2}}{2} \sum_{\theta' \in \{\theta_a, \theta_b\}} (\|\theta'\| + \sqrt{\mathbb{E}\|X_1\|^2}) \sqrt{\mathbb{P}_{\theta'}(\text{clip}(\ell_1) \neq \ell_1)} \\ &\geq \min\left\{\frac{\sigma\|v\|}{16}, \frac{\varepsilon}{4}\right\} T^{2/3} - \sqrt{2}\alpha^* T^{3/2} (r/2 + \sigma\sqrt{d}) e^{-r^2/(160\sigma^2)} \\ &\geq \min\left\{\frac{\sigma\|v\|}{16}, \frac{\varepsilon}{4}\right\} T^{2/3} - \sqrt{2}\alpha^* r T^{3/2} e^{-r^2/(160\sigma^2)}, \end{aligned}$$

where the last step uses that $\sigma \leq r/\sqrt{8d}$. Now, we can pick $\sigma = \frac{r}{\sqrt{240 \log(T)}}$, which is a valid choice when T is large enough so that $\sigma^2 \leq r^2/(8d)$, meaning that $T \geq e^{d/30}$. With this choice of σ we obtain that

$$R_T^* \geq \min\left\{\frac{r\|v\|}{32\sqrt{60 \log(T)}}, \frac{\varepsilon}{4}\right\} T^{2/3} - \sqrt{2}\alpha^* r,$$

which concludes the proof. $\square$

Theorem 4. Assume that $\mathcal{B}_2(r) \subseteq \mathcal{L}$ for some $r > 0$. If the game is locally observable and there is at least one pair of non-duplicate Pareto optimal actions, then R_T^* is $\tilde{\Omega}(\sqrt{T})$.

Proof. By the connectedness of the neighborhood graph, there must be at least two pair of neighboring actions, which we again denote by a and b . Most of the proof proceeds almost identically to that of Theorem 3. The main difference is that $\psi_a - \psi_b \in \text{span}\{M_c : c \in \mathcal{N}_{a,b}\}$ by local observability. Moreover, $\bar{\mathcal{N}}_{a,b}$ could be empty. Nevertheless, we can directly take

$$v := \psi_a - \psi_b,$$

which is strictly positive, and again define $q := \frac{r}{4\|v\|} v$. Then, for some $\theta \in \text{relint}(C_a \cap C_b)$, we again choose

$$\theta_a = \theta - \Delta q \quad \text{and} \quad \theta_b = \theta + \Delta q.$$

Following the proof of Theorem 3 verbatim, we reach that

$$\begin{aligned} \bar{R}_T(\theta_a) + \bar{R}_T(\theta_b) &\geq \frac{\varepsilon}{2} T_{\theta_a}(\bar{\mathcal{N}}_{a,b}) + \frac{\tilde{\Delta}}{2} T(1 - \sqrt{(1/2) D_{\text{KL}}(\mathbb{P}_{\theta_a} \parallel \mathbb{P}_{\theta_b})}) \\ &\geq \frac{\tilde{\Delta}}{2} T(1 - \sqrt{(1/2) D_{\text{KL}}(\mathbb{P}_{\theta_a} \parallel \mathbb{P}_{\theta_b})}). \end{aligned}$$

We choose, again, that $X_t \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_d)$. Now, however, we use the bound $D_{\text{KL}}(\mathbb{P}_{\theta_a, c} \parallel \mathbb{P}_{\theta_b, c}) \leq \frac{2\tilde{\Delta}^2}{\sigma^2\|v\|^2}$ for all actions c , even the ones in $\mathcal{N}_{a,b}$. This yields via Lemma 19 that

$$\bar{R}_T(\theta_a) + \bar{R}_T(\theta_b) \geq \frac{\tilde{\Delta}}{2} T \left(1 - \frac{\tilde{\Delta}}{\sigma\|v\|} \sqrt{T}\right).$$

Taking $\Delta = 2\sigma T^{-1/2}/r$, (assuming again that T is large enough to satisfy the condition set on Δ in the proof of Theorem 3),⁷ we have that $\tilde{\Delta} = \sigma\|v\|T^{-1/2}/2$, which implies that

$$\bar{R}_T(\theta_a) + \bar{R}_T(\theta_b) \geq \frac{\sigma\|v\|}{8} \sqrt{T}.$$

⁷If $\bar{\mathcal{N}}_{a,b}$ is empty (also meaning that $T_{\theta}(\bar{\mathcal{N}}_{a,b}) = 0$), one can satisfy the conditions on Δ in the proof of Theorem 3 taking ε as an arbitrarily large number.

The rest of the proof again follows that of Theorem 3, with an identical use of Lemma 18 and tuning of σ , finally yielding that

$$R_T^* \geq \frac{r\|v\|}{32\sqrt{60\log(T)}}\sqrt{T} - \sqrt{2}\alpha^*r$$

for sufficiently large T . $\square$

Theorem 5. Assume that $\mathcal{B}_2(r) \subseteq \mathcal{L}$ for some $r > 0$. If the game is not globally observable, then R_T^* is $\Omega(T)$.

Proof. The game not being globally observable means that there exists $a, b \in \mathcal{A}$ such that

$$\psi_a - \psi_b \notin \text{span}\{M_c : c \in \mathcal{A}\}.$$

This implies that they have distinct features, which in turn implies that there are at least two non-duplicate Pareto optimal actions (vertices). Now, if the property above holds for some pair of actions (a, b) , then it must hold for at least one pair of non-duplicate Pareto optimal actions. This is because the features of any action can be written as a convex combination (with positive weights) of the feature of one or more Pareto optimal actions. Hence, $\psi_a - \psi_b$ can be written as a linear combination of feature differences for Pareto optimal actions. This implies that we can take a and b to be (non-duplicate) Pareto optimal actions. In fact, we can take them to be neighbors too; since the neighborhood graph is connected, we can find a path $(a_i)_{i=1}^s$ where $a_1 = a$, $a_s = b$, and (a_i, a_{i+1}) are neighbors (endpoint of an edge in the polytope $\mathcal{V} = \text{co}(\mathcal{X})$). And since $\psi_a - \psi_b = \sum_{i=1}^{s-1} (\psi_{a_i} - \psi_{a_{i+1}})$, at least one pair of neighbors must satisfy the non-observability condition above. Thus, without loss of generality, a and b are assumed to be neighbors.

Now, define v as the component of $\psi_a - \psi_b$ in the orthogonal complement of $\text{span}\{M_c : c \in \mathcal{A}\}$, and set $q := \frac{r}{4\|v\|}v$. Neither of which is zero by assumption. With this altered choice of v (and the corresponding one for q), we can follow the proof of Theorem 3 verbatim till we reach that

$$\begin{aligned} \bar{R}_T(\theta_a) + \bar{R}_T(\theta_b) &\geq \frac{\varepsilon}{2}T_{\theta_a}(\bar{\mathcal{N}}_{a,b}) + \frac{\tilde{\Delta}}{2}T(1 - \sqrt{(1/2) D_{\text{KL}}(\mathbb{P}_{\theta_a} \parallel \mathbb{P}_{\theta_b})}) \\ &\geq \frac{\tilde{\Delta}}{2}T(1 - \sqrt{(1/2) D_{\text{KL}}(\mathbb{P}_{\theta_a} \parallel \mathbb{P}_{\theta_b})}). \end{aligned}$$

Observe now that q belongs to the orthogonal complement of $\text{col}(M_c)$ for all $c \in \mathcal{A}$, meaning that, via the derivation in the proof of Theorem 3 (again with $X_t \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}_d)$), $D_{\text{KL}}(\mathbb{P}_{\theta_{a,c}} \parallel \mathbb{P}_{\theta_{b,c}}) = 0$. Hence, $D_{\text{KL}}(\mathbb{P}_{\theta_a} \parallel \mathbb{P}_{\theta_b}) = 0$ via Lemma 19. Then, choosing $\Delta = \min\{1, \frac{2\varepsilon}{r\alpha^*}\}$ yields that⁸

$$\bar{R}_T(\theta_a) + \bar{R}_T(\theta_b) \geq \min\left\{\frac{r\|v\|}{8}, \frac{\varepsilon\|v\|}{4\alpha^*}\right\} T.$$

Next, a similar application of Lemma 18 to that in the proof of Theorem 3 gives (with a sufficiently small σ) that

$$R_T^* \geq \min\left\{\frac{r\|v\|}{16}, \frac{\varepsilon\|v\|}{8\alpha^*}\right\} T - \sqrt{2}\alpha^*rT^{3/2}e^{-r^2/(160\sigma^2)}.$$

Since σ does not appear in the first term (in fact, adding noise here is redundant), it can be chosen sufficiently small so that

$$R_T^* \geq \min\left\{\frac{r\|v\|}{32}, \frac{\varepsilon\|v\|}{16\alpha^*}\right\} T,$$

proving the theorem. $\square$

J.3 A lower bound for bandits with ill-conditioned observers

As defined in Section 2, we have in this game that $d = k$, $\psi_a = e_a$, and $M_a = (1 - \varepsilon)\mathbf{1}_k/k + \varepsilon e_a$ for some $\varepsilon \in (0, 1]$. Here, we take $\mathcal{L} = \mathcal{B}_\infty(1)$ (see Appendix I). In fact, we clip the losses in this lower bound construction to belong to $[0, 1]^k$. This example is adapted from [18, Section 6], where they consider the case when ψ_a too is $(1 - \varepsilon)\mathbf{1}_k/k + \varepsilon e_a$, and prove a $\tilde{\Omega}(\sqrt{kT})$ lower bound in their Theorem 9. The game we consider here is harder as asserts the following theorem.

⁸As remarked in the proof of Theorem 4, we can take ε to be an arbitrarily large number if $\bar{\mathcal{N}}_{a,b}$ is empty.

Theorem 6. *In the bandit game with ill-conditioned observers, it holds that R_T^* is $\tilde{\Omega}(1/\varepsilon\sqrt{kT})$.*

Proof. Besides the altered definition of ψ_a , the proof is essentially identical to that of [18, Proposition 8]. Nonetheless, we report here the full proof in our notation for completeness. Consider k environments $\{\theta_a\}_{a \in \mathcal{A}}$ such that

$$\theta_a := \left(\frac{1}{2} + \Delta \frac{1-\varepsilon}{k}\right) \mathbf{1}_k - \Delta e_a,$$

besides an auxiliary environment $\theta_0 := (1/2)\mathbf{1}_k$, where $0 < \Delta \leq 1/4$. Notice then that for any $b \in \mathcal{A} \setminus \{a\}$,

$$\langle \psi_b - \psi_a, \theta_a \rangle = \langle e_b - e_a, \theta_a \rangle = \Delta.$$

Hence, $\theta_a \in C_a$, and in a similar fashion to the proof of Theorem 3, we have that

$$\begin{aligned} \bar{R}_T(\theta_a) &= \mathbb{E}_{\theta_a} \sum_{t=1}^T \langle \psi_{A_t} - \psi_a, \theta_a \rangle = \Delta(T - T_{\theta_a}(a)) \\ &= \Delta(T - T_{\theta_0}(a)) - \Delta(T_{\theta_a}(a) - T_{\theta_0}(a)) \\ &\geq \Delta(T - T_{\theta_0}(a)) - \Delta T \sqrt{(1/2) D_{\text{KL}}(\mathbb{P}_{\theta_0} \parallel \mathbb{P}_{\theta_a})}. \end{aligned}$$

In this game, we will take X_t to be distributed according to $\mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{1}_d \mathbf{1}_d^\top)$, for some $\sigma > 0$. This means that in environment θ_a , $\ell_t = \theta_a + Z_t \mathbf{1}_k$, where $(Z_t)_t$ is an i.i.d. sequence of random variables, each sampled from $\mathcal{N}(0, \sigma^2)$. Notice that, for all $c \in \mathcal{A}$, $M_c \in \Delta_k$; hence, $M_c^\top \ell_t = M_c^\top \theta_a + Z_t$. So we get that

$$\begin{aligned} D_{\text{KL}}(\mathbb{P}_{\theta_0, c} \parallel \mathbb{P}_{\theta_a, c}) &= D_{\text{KL}}\left(\mathcal{N}(M_c^\top \theta_0, \sigma^2) \parallel \mathcal{N}(M_c^\top \theta_a, \sigma^2)\right) \\ &= \frac{1}{2\sigma^2} (\theta_0 - \theta_a)^\top M_c M_c^\top (\theta_0 - \theta_a). \end{aligned}$$

Notice, moreover that

$$(\theta_0 - \theta_a)^\top M_c = \Delta \left(e_a - \frac{1-\varepsilon}{k} \mathbf{1}_k \right)^\top \left(\frac{1-\varepsilon}{k} \mathbf{1}_k + \varepsilon e_c \right) = \varepsilon \Delta \mathbb{I}[a = c].$$

Thus, $D_{\text{KL}}(\mathbb{P}_{\theta_0, c} \parallel \mathbb{P}_{\theta_a, c}) = \frac{\varepsilon^2 \Delta^2}{2\sigma^2} \mathbb{I}[a = c]$ Lemma 19 then implies that

$$\bar{R}_T(\theta_a) \geq \Delta(T - T_{\theta_0}(a)) - \Delta T \sqrt{\frac{\varepsilon^2 \Delta^2}{4\sigma^2} T_{\theta_0}(a)} = \Delta(T - T_{\theta_0}(a)) - \frac{\varepsilon \Delta^2}{2\sigma} T \sqrt{T_{\theta_0}(a)}.$$

Hence, via Jensen's inequality,

$$\begin{aligned} \frac{1}{k} \sum_{a \in \mathcal{A}} \bar{R}_T(\theta_a) &\geq \Delta(T - T/K) - \frac{\varepsilon \Delta^2}{2\sigma} T \sqrt{T/k} \\ &\geq \frac{\Delta}{2} T \left(1 - \frac{\varepsilon \Delta}{\sigma} \sqrt{T/k} \right) \\ &= \frac{\sigma}{8\varepsilon} \sqrt{kT} \end{aligned}$$

where in the last step we have picked $\Delta = \frac{\sigma\sqrt{k}}{2\varepsilon\sqrt{T}}$, assuming that $T \geq 4\sigma^2 k/\varepsilon^2$ so that $\Delta \leq 1/4$. To apply Lemma 18, we choose the clipping operator as follows (where $\ell_t(c) = e_c^\top \ell_t$):

$$[\text{clip}(\ell_t)](c) := \max\{\min\{1, \ell_t(c)\}, 0\}.$$

Note that under any environment θ_a , $1/4 \leq \theta_a(c) \leq 3/4$ for all $c \in \mathcal{A}$ by the definition of θ_a and the fact that $0 \leq \Delta \leq 1/4$. Hence, since $\ell_t(a) = \theta_a(c) + Z_t$, if $\text{clip}(\ell_t) \neq \ell_t$, then it must be that $|Z_t| > 1/4$. And since $Z_t \sim \mathcal{N}(0, \sigma^2)$, we have that $\mathbb{P}(|Z_t| > 1/4) \leq 2 \exp(-1/(32\sigma^2))$. Moreover, under environment θ_a , it holds for any pair of actions c and d that

$$\langle \psi_c - \psi_d, \ell_t \rangle = \langle e_c - e_d, \theta_a + Z_t \mathbf{1}_k \rangle = \langle e_c - e_d, \theta_a \rangle \leq \Delta,$$

which implies that $\hat{R}_T(\theta_a) \leq \Delta T$ uniformly. Therefore, we can apply the second part of Lemma 18 to obtain that

$$R_T^* \geq \frac{\sigma}{8\varepsilon} \sqrt{kT} - 2\Delta T^2 e^{-1/32\sigma^2} = \frac{\sigma}{\varepsilon} \sqrt{kT} \left(\frac{1}{8} - T e^{-1/32\sigma^2} \right) = \frac{\varepsilon^{-1}}{64\sqrt{2\log(16T)}} \sqrt{kT},$$

where the first equality uses our choice of Δ , and to obtain the second we set $\sigma = \frac{1}{4\sqrt{2\log(16T)}}$. $\square$