
Near-Optimal Pure Exploration in Matrix Games: A Generalization of Stochastic Bandits & Dueling Bandits

Arnab Maiti*

Ross Boczar*

Kevin Jamieson

Lillian J. Ratliff

University of Washington

{arnabm2,rjboczar,ratliff1}@uw.edu, jamieson@cs.washington.edu

Abstract

We study the sample complexity of identifying the pure strategy Nash equilibrium (PSNE) in a two-player zero-sum matrix game with noise. Formally, we are given a stochastic model where any learner can sample an entry (i, j) of the input matrix $A \in [-1, 1]^{n \times m}$ and observe $A_{i,j} + \eta$ where η is a zero-mean 1-sub-Gaussian noise. The aim of the learner is to identify the PSNE of A , whenever it exists, with high probability while taking as few samples as possible. Zhou et al. (2017) presents an instance-dependent sample complexity lower bound that depends only on the entries in the row and column in which the PSNE lies. We design a near-optimal algorithm whose sample complexity matches the lower bound, up to log factors. The problem of identifying the PSNE also generalizes the problem of pure exploration in stochastic multi-armed bandits and dueling bandits, and our result matches the optimal bounds, up to log factors, in both the settings.

1 INTRODUCTION

Pure exploration is a well-studied topic in stochastic multi-armed bandits. Given a set of arms, the aim is to identify the arm with the highest mean with high probability while sequentially sampling the arms. A practically motivated version of this problem is *dueling bandits* where the learner pulls two arms and observes the winner among them. The objective is to identify the Condorcet winner: the arm that has probability

greater than $1/2$ of winning against any other arm. The dueling bandit problem has also been studied from the perspective of pure exploration. A generalization of dueling bandits are matrix games which model scenarios where multiple strategic agents are involved and are competing against each other in a stochastic environment.

Consider a hypothetical situation in which firms F_1 and F_2 both offer the same product. Typically, when both firms set an identical price, the customer base is evenly distributed between them. However, if one firm sets its price higher than the other, it is likely to experience a reduced share of the customer base. This presents a strategic challenge for the firm: should it opt for a lower price p_1 or a higher price p_2 ? If the firms are non-cooperative, the optimal thing for each firm to do is to set the lower price p_1 otherwise they will most probably lose customers. Such a strategic scenario in a stochastic environment can be formally captured by a two-player zero-sum game defined by a matrix A as follows. Consider $A = [0, 0.25; -0.25, 0]$ where $A_{i,j}$ denotes the expected increase in the fraction of customers visiting firm F_1 if first decides to charge p_i and F_2 decides to charge p_j . Both firms setting a low price p_1 corresponds to the entry $(1, 1)$ and such an entry is referred to as pure strategy Nash equilibrium (PSNE) in game theory.

Formally, an entry is a PSNE of A if it is the highest entry in its column and lowest entry in its row. Although, a PSNE need not exist in every matrix A , there are some conditions under which they exist (see Duersch et al. (2012); Radzik (1991); Shapley (1964)). For instance, Duersch et al. (2012) showed that a symmetric two-player zero-sum game has a pure strategy equilibrium if and only if it is not a generalized rock-paper-scissors matrix. In this paper we generalize both stochastic multi-armed bandits and dueling bandits by studying the problem of identifying a PSNE, whenever it exists, in matrix games with noise.

Proceedings of the 27th International Conference on Artificial Intelligence and Statistics (AISTATS) 2024, Valencia, Spain. PMLR: Volume 238. Copyright 2024 by the author(s).

1.1 Problem Setting

We consider two problems settings.

Matrix Games. In this problem setting, there is an arbitrary input matrix $A \in [-1, 1]^{n \times m}$ which is unknown to the learner. The learner can sample an entry (i, j) of A and observe the random variable $X_{i,j} = A_{i,j} + \eta$ where η is a zero-mean 1-sub-Gaussian noise. The aim is to design a δ -PAC learner, that is, a learner which has the following property: if the learner stops after a finite time τ and returns an entry (\hat{i}, \hat{j}) then (\hat{i}, \hat{j}) is a PSNE of A with probability at least $1 - \delta$. The objective is to minimize the number of samples drawn from A .

Dueling Bandits. In dueling bandits setting, there are K arms and a matrix $\mathbf{P} \in [0, 1]^{K \times K}$. The matrix \mathbf{P} has the property that for all i, j we have $\mathbf{P}_{i,j} + \mathbf{P}_{j,i} = 1$. In each round t , a pair of arms i, j are pulled (also known as a duel) and a winner X_t is declared where $\mathbb{P}(X_t = i) = 1 - \mathbb{P}(X_t = j) = \mathbf{P}_{i,j}$. An arm i_* is called a Condorcet winner if for all $j \in [K] \setminus \{i_*\}$ we have $\mathbf{P}_{i_*,j} > 1/2$. In pure exploration, one aims to design a learner that identifies the Condorcet winner, whenever it exists, with probability at least $1 - \delta$ by performing as few duels as possible. Note that dueling bandits is a special case of a matrix game with $A = \mathbf{P}$. Here the PSNE of A is equal to (i_*, i_*) where i_* is the Condorcet winner.

1.2 Contributions

Zhou et al. (2017) initiated the study of identifying the PSNE where they showed that any δ -PAC learner requires $\Omega(\mathbf{H}_1 \log(1/\delta))$ samples to find the PSNE (i_*, j_*) , whenever it exists, where

$$\mathbf{H}_1 = \prod_{i \neq i_\tau} \frac{1}{(A_{i_\tau, j_\tau} - A_{i, j_\tau})^2} + \prod_{j \neq j_\tau} \frac{1}{(A_{i_\tau, j} - A_{i_\tau, j_\tau})^2}.$$

In this paper, we design a δ -PAC learner (Algorithm 3) that achieves the sample complexity of $\mathbf{H}_1 \log(1/\delta)$, up to log factors.

Recall that dueling bandits is a special case of matrix games. Haddenhorst et al. (2021) showed that any Condorcet δ -PAC learner requires $\prod_{i \neq i_\tau} \Delta_{i, i_\tau}^{-2} \log(1/\delta)$ where $\Delta_{i, i_\tau} = \mathbf{P}_{i_\tau, i} - 1/2$. In the notation of matrix games, this sample complexity is equal to $\frac{1}{2} \mathbf{H}_1 \log(1/\delta)$. As our algorithm achieves $\mathbf{H}_1 \log(1/\delta)$, up to log factors, it is a near-optimal Condorcet learner. To the best of our knowledge, this is the first algorithm to achieve this for Condorcet winner identification, without making additional strong assumptions (see the discussion on Dueling Bandits in Section 1.3).

Beyond the theoretical contributions, we also benchmark our algorithm against strong baselines in Sec-

tion 3 and demonstrate that our algorithm is also superior empirically.

1.3 Related Work

Matrix games. As just described Zhou et al. (2017) showed that identifying the PSNE of a zero-sum matrix game requires a sample complexity of $\Omega(\mathbf{H}_1 \cdot \log(1/\delta))$. They also provided an algorithm based on LUCB that achieves an upper bound of $O(\mathbf{H}_1 \log(\mathbf{H}_1/\delta) + \frac{nm}{\Delta})$ where Δ is a matrix-dependent parameter. This upper bound is sub-optimal, and can be far worse than the lower bound. For instance, when δ and the gaps $|A_{i,j} - A_{i_\tau, j_\tau}|$ are constants, the upper bound scales as $O(nm)$ where as the lower bound is $\Omega(n + m)$. Recently, Maiti et al. (2023a) studied the problem of finding ε -Nash Equilibrium in $n \times 2$ games and provided near-optimal instance-dependent bounds. The instance-dependent parameters include the gaps between the entries of the matrix, and difference between the value of the game and reward received from playing a sub-optimal row. Later, Maiti et al. (2023b) extended the techniques of Maiti et al. (2023a) to identify the support of the Nash equilibrium in arbitrary $n \times m$ games. However, the bounds in their paper are sub-optimal.

An important class of learners for matrix games utilizes no-regret algorithms for adversarial bandits (cf. Freund and Schapire (1999)). Specifically, initialize two independent copies of Exp3-IX where the first selects rows to maximize reward, and the other selects columns to minimize loss, where both are fed $A_{i,j} + \eta$. Due to the fact that both algorithms enjoy an external regret bound of \sqrt{KT} for a game played for T time steps with K arms, it can be shown that with constant probability, the average of the plays converges to an ε -approximate Nash equilibrium as soon as T exceeds $\frac{n+m}{\varepsilon^2}$, ignoring log factors. Consequently, this guarantees that the PSNE can be identified with constant probability once T exceeds $\frac{n+m}{\Delta_{\min}^2}$ where

$$\Delta_{\min} = \min \left\{ \min_{i \neq i} A_{i, j} - A_{i, j}, \min_{j \neq j} A_{i, j} - A_{i, j} \right\}.$$

The Tsallis-Inf algorithm of Zimmert and Seldin (2021) is a no-regret algorithm for adversarial bandits that also enjoys some instance-dependent guarantees under certain favorable conditions. While we are not aware of any analysis of this algorithm in the matrix game setting, our experiments will show that while it is a stronger baseline than Exp3-IX, its empirical performance is far inferior to our algorithm.

Stochastic Multi-Armed Bandits. Best arm identification in the fixed confidence setting is a well-studied topic (see Appendix B for the problem formulation). One of the earlier works in Even-

Dar et al. (2002) introduces the Successive Elimination algorithm that achieved an upper bound of $\sum_{i \neq i_*} \Delta_i^{-2} \log(n \Delta_i^{-2} / \delta)$ where i_* is the best arm and $\Delta_i = \mu_{i_*} - \mu_i$. A lower bound of $\sum_{i \neq i_*} \Delta_i^{-2} \log(1/\delta)$ was then established by Mannor and Tsitsiklis (2004) (later it was refined and simplified by Kaufmann et al. (2016)). The lower upper confidence bound (LUCB) algorithm designed by Kalyanakrishnan et al. (2012) achieved an upper bound of $\sum_{i \neq i_*} \Delta_i^{-2} \log(\sum_{j \neq i_*} \Delta_j^{-2} / \delta)$. Eventually, an upper bound of $\sum_{i \neq i_*} \Delta_i^{-2} \log(\log(\Delta_i^{-2}) / \delta)$ was achieved by the Exponential-gap Elimination algorithm of Karnin et al. (2013) and the lil'UCB algorithm of Jamieson et al. (2014). Jamieson et al. (2014) also showed that this upper bound is optimal for two-armed bandits. Later Chen et al. (2017) provided even tighter upper bounds in the general case based on entropy-like terms determined by the gaps.

Dueling Bandits. The dueling bandits problem is a well-studied variant of stochastic bandits (cf. Bengs et al. (2021) for a survey). The aim here is to identify the best arm (according to some rule) by noisy pairwise comparisons. This paper focuses on the Condorcet winner: the arm (assuming it exists) that has a probability greater than $1/2$ of beating every other arm. Haddenhorst et al. (2021) provided a lower bound of $\sum_{i \neq i_*} \Delta_{i,i_*}^{-2} \equiv \frac{1}{2} \mathbf{H}_1$ to identify the Condorcet winner i_* , whenever it exists, where $\Delta_{i,i_*} = \mathbf{P}_{i_*,i} - 1/2$ is the probability that arm i_* beats arm i , minus $1/2$. They also provided an upper bound of $\frac{K}{\Delta_{\min,all}^2}$, up to log factors, where $\Delta_{\min,all} = \min_{i \neq j} |\mathbf{P}_{i,j} - 1/2|$.

In an effort to hit the lower bound, prior works introduced a number of strong assumptions to make the problem easier. The strong assumptions include total order, strong stochastic transitivity (SST) and stochastic triangle inequality (STI). Under total ordering, Mohajer et al. (2017) achieved an upper bound of $\frac{K \log \log K}{\min_{i \neq i} \Delta_{i,i_*}^2}$. Finally, Ren et al. (2020) achieved an upper bound of $\sum_{i \neq i_*} \Delta_{i,i_*}^{-2} \equiv \frac{1}{2} \mathbf{H}_1$ that matches the lower bound up to log factors, but under the strong assumptions of the total ordering setting, SST and STI. We emphasize that our work does *not* make these strong assumptions, yet still achieves the optimal sample complexity of \mathbf{H}_1 , up to log factors.

Regret minimization for dueling bandits has also been studied. Recently Saha and Gaillet (2022) achieved the optimal regret bound of $\sum_{i \neq i_*} \frac{\log T}{\Delta_{i,i_*}}$ with the help of Tsallis-inf from Zimmert and Seldin (2021). This immediately implies an upper bound of $\frac{1}{\min_{i \neq i} \Delta_{i,i_*}} \sum_{i \neq i_*} \frac{1}{\Delta_{i,i_*}}$ with constant probability, up to log factors. In the notation of matrix games, this sample complexity is at least \mathbf{H}_1 , and potentially Δ_{\min}^{-1} times larger than \mathbf{H}_1 for certain instances (see the du-

eling bandit instance of Section 3). To the best of our knowledge, this is the best known sample complexity bound in this setting prior to our work.

2 FIXED CONFIDENCE NEAR-OPTIMAL ALGORITHM

Consider a matrix $A \in [-1, 1]^{n \times m}$ which has a PSNE (i_*, j_*) . Let $\Delta_{i,j} = |A_{i,j} - A_{i_*,j_*}|$ for any i, j . Let $\mathbf{H}_1 = \sum_{i \neq i_*} \frac{1}{\Delta_{i,j_*}^2} + \sum_{j \neq j_*} \frac{1}{\Delta_{i_*,j}^2}$. Let us assume that $A_{i,j_*} < A_{i_*,j_*} < A_{i_*,j}$ for any $(i, j) \neq (i_*, j_*)$, otherwise \mathbf{H}_1 is not well-defined. Now we state the main result.

Theorem 1. *There is an algorithm (Algorithm 3) that, with probability at least $1 - \delta$, takes at most $c_0 \cdot \mathbf{H}_1 \cdot \log(\frac{nm \log(\mathbf{H}_1)}{\delta}) \cdot \log(nm)$ samples from the input matrix A and returns the PSNE (i_*, j_*) . Here c_0 is an absolute constant.*

For the sake of simplicity of presentation, let us assume that n and m are powers of two. Let $\Delta_g := \frac{1}{n+m-2} \sum_{i \neq i_*} \frac{1}{\Delta_{i,j_*}^2} + \sum_{j \neq j_*} \frac{1}{\Delta_{i_*,j}^2}$. Our goal is to identify (i_*, j_*) with a sample complexity that scales as $(n + m - 2) \Delta_g^{-2} = \mathbf{H}_1$, ignoring log factors.

We divide the analysis into multiple parts. First, in Section 2.1, we design an algorithm using a guess of Δ_g . Next, in Section 2.2, we describe the details of the procedures ColMidVal and RowMidVal that were used in Section 2.1. In Section 2.3, we state the main algorithm which has a sample complexity of $(n + m - 2) \Delta_g^{-2}$, ignoring log factors. Finally, in Section 2.4, we address the special case of dueling bandits. In the interest of space, we refer the reader to Appendix E for a few omitted calculations.

2.1 Algorithm with a guess of Δ_g

In this section, we aim to design an algorithm to identify the PSNE (i_*, j_*) when we are given a guess of the parameter Δ_g . Informally, the algorithm proceeds in a logarithmic number of stages. In each stage it aims to eliminate either half of the sub-optimal rows or sub-optimal columns. Suppose at one particular stage the algorithm aims to eliminate half of the sub-optimal rows. The algorithm takes samples from each column and computes a value near the median of the column entries. Then the algorithm chooses the column \mathfrak{b} which has the lowest such value. Then the algorithm samples all the entries of the column \mathfrak{b} sufficiently and removes those rows whose corresponding entries in column \mathfrak{b} are lower than median of the column entries.

Before formally describing our algorithm, we state the guarantees of the two subroutines ColMidVal and

Algorithm 1 MidSearch(A, Δ, δ)

-
- 1: $\mathcal{X}_1 \leftarrow$ rows of A , $\mathcal{Y}_1 \leftarrow$ columns of A
 - 2: **for** $t = 1, 2, \dots$ **do**
 - 3: **if** $\max\{|\mathcal{X}_t|, |\mathcal{Y}_t|\} = 2$ **then**
 - 4: Sample every element in $\mathcal{X}_t \times \mathcal{Y}_t$ $\frac{n+m-2}{2}$ $\cdot \frac{50 \log(16/\delta)}{\Delta^2}$ times and return the PSNE of the sub-matrix formed by $\mathcal{X}_t \times \mathcal{Y}_t$.
 - 5: **else if** $|\mathcal{X}_t| \geq |\mathcal{Y}_t|$ **then**
 - 6: Let \mathcal{C}_j denote the elements of column j in the sub-matrix formed by $\mathcal{X}_t \times \mathcal{Y}_t$.
 - 7: $\varepsilon_t \leftarrow \frac{1}{9} \frac{|\mathcal{X}_t|}{n+m-2} \cdot \Delta$
 - 8: Set $\mathbf{b}_{t,x}(j) \leftarrow \text{ColMidVal}(\mathcal{C}_j, \varepsilon_t, \frac{\delta}{2m^2n^2})$ for each $j \in \mathcal{Y}_t$.
 - 9: $\mathfrak{b} \leftarrow \arg \min_{j \in \mathcal{Y}_t} \mathbf{b}_{t,x}(j)$ and $\mathcal{Y}_{t+1} \leftarrow \mathcal{Y}_t$
 - 10: Sample every entry in $\{(i, \mathfrak{b}) : i \in \mathcal{X}_t\}$ $\frac{n+m-2}{|\mathcal{X}_t|} \cdot \frac{162 \log(4n^2m^2/\delta)}{\Delta^2}$ times and compute the empirical average $\bar{A}_{i,j}$ of these samples.
 - 11: $\mathcal{X}_{t+1} \leftarrow \{i_1, \dots, i_{|\mathcal{X}_t|/2}\} \subset \mathcal{X}_t = \{i_1, \dots, i_{|\mathcal{X}_t|}\}$ where $\bar{A}_{i_1, \mathfrak{b}} \geq \dots \geq \bar{A}_{i_{|\mathcal{X}_t|/2}, \mathfrak{b}}$.
 - 12: **else if** $|\mathcal{X}_t| < |\mathcal{Y}_t|$ **then**
 - 13: Let \mathcal{R}_i denote the elements of row i in the sub-matrix formed by $\mathcal{X}_t \times \mathcal{Y}_t$.
 - 14: $\varepsilon_t \leftarrow \frac{1}{9} \frac{|\mathcal{Y}_t|}{n+m-2} \cdot \Delta$
 - 15: Set $\mathbf{b}_{t,y}(i) \leftarrow \text{RowMidVal}(\mathcal{R}_i, \varepsilon_t, \frac{\delta}{2m^2n^2})$, for each $i \in \mathcal{X}_t$.
 - 16: $\mathfrak{b} \leftarrow \arg \max_{i \in \mathcal{X}_t} \mathbf{b}_{t,y}(i)$ and $\mathcal{X}_{t+1} \leftarrow \mathcal{X}_t$
 - 17: Sample every entry in $\{(\mathfrak{b}, j) : j \in \mathcal{Y}_t\}$ $\frac{n+m-2}{|\mathcal{Y}_t|} \cdot \frac{162 \log(4n^2m^2/\delta)}{\Delta^2}$ times and compute the empirical average $\bar{A}_{i,j}$ of these samples.
 - 18: $\mathcal{Y}_{t+1} \leftarrow \{j_1, \dots, j_{|\mathcal{Y}_t|/2}\} \subset \mathcal{Y}_t = \{j_1, \dots, j_{|\mathcal{Y}_t|}\}$ where $\bar{A}_{\mathfrak{b}, j_1} \leq \dots \leq \bar{A}_{\mathfrak{b}, j_{|\mathcal{Y}_t|/2}}$.
 - 19: **end if**
 - 20: **end for**
-

RowMidVal from Section 2.2 that will be used in our algorithm. Consider $\varepsilon, \delta > 0$. Given a set of n arms \mathcal{A} with means $\mu_1 \geq \mu_2 \geq \dots \geq \mu_n$, ColMidVal outputs a value $\mathbf{b} \in [\mu_{n/2} - \varepsilon, \mu_{n/4+1} + \varepsilon]$ with probability $1 - \delta$ by using at most $O(\frac{1}{\varepsilon^2} \log \frac{1}{\delta})$ samples. Similarly, RowMidVal outputs a value $\mathbf{b} \in [\mu_{3n/4} - \varepsilon, \mu_{n/2+1} + \varepsilon]$ with probability $1 - \delta$ by using at most $O(\frac{1}{\varepsilon^2} \log \frac{1}{\delta})$ samples. We refer the reader to Algorithm 1 for a formal description of our algorithm.

Now we begin the analysis of Algorithm 1. First, we state the following proposition that provides a lower bound on the gaps Δ_{i,j^*} . The lower bound later allows us to argue that a constant fraction of the entries in the optimal column j_* is well below A_{i^*,j^*} , and this fact plays a crucial role in lowering the number of samples required in each stage.

Proposition 1. Consider a subset $S = \{i_1, \dots, i_\ell\} \subset [n]$ such that $\Delta_{i_1, j^*} \leq \dots \leq \Delta_{i_\ell, j^*}$. Let $\ell \geq 2$. Then for any $s \geq \lceil \ell/4 \rceil + 1$, $\Delta_{i_s, j^*} \geq \frac{1}{2} \frac{|S|}{n+m-2} \Delta_g$.

Proof. Consider an index $s \geq \lceil \ell/4 \rceil + 1$. Now we have the following:

$$\begin{aligned} \frac{|S|}{n+m-2} \Delta_g^{-2} &= \frac{1}{|S|} \overset{\circ}{@} \overset{\times}{\times}_{i \neq i^*} \frac{1}{\Delta_{i, j^*}^2} + \overset{\times}{\times}_{j \neq j^*} \frac{1}{\Delta_{i^*, j}^2} A \\ &\geq \frac{1}{|S|} \overset{\times}{\times}_{i \in S \setminus \{i^*\}} (\Delta_{i, j^*})^{-2} \\ &\geq \frac{1}{|S|} \overset{[\ell/4]+1}{\times}_{k=2} (\Delta_{i_k, j^*})^{-2} \\ &\geq \frac{[\ell/4]}{|S|} \cdot \Delta_{i_s, j^*}^{-2} \\ &\geq (2\Delta_{i_s, j^*})^{-2}, \end{aligned}$$

where the second to last inequality holds since $\Delta_{i_s, j^*} \geq \Delta_{i_k, j^*}$ for all $k \leq \lceil \ell/4 \rceil + 1$, and the last inequality is due to the fact that $|S| = \ell$. \square

Similarly, we have the following proposition.

Proposition 2. Consider a subset $S = \{j_1, \dots, j_\ell\} \subset [m]$ such that $\Delta_{i^*, j_1} \leq \dots \leq \Delta_{i^*, j_\ell}$. Let $\ell \geq 2$. Then for any $s \geq \lceil \ell/4 \rceil + 1$, $\Delta_{i^*, j_s} \geq \frac{1}{2} \frac{|S|}{n+m-2} \Delta_g$.

Next, we establish the sample complexity of our algorithm in the following lemma.

Lemma 1. The sample complexity of the procedure MidSearch(A, Δ, δ) is $c \cdot \frac{n+m-2}{\Delta^2} \cdot \log \frac{nm}{\delta} \cdot \log(nm)$ where c is an absolute constant.

Proof. Let $\mathcal{T}_1 := \{t : |\mathcal{X}_{t+1}| = \frac{|\mathcal{X}_t|}{2}\}$ and $\mathcal{T}_2 := \{t : |\mathcal{Y}_{t+1}| = \frac{|\mathcal{Y}_t|}{2}\}$, and consider an iteration $t \in \mathcal{T}_1$. Recall that $\varepsilon_t = \frac{1}{9} \frac{|\mathcal{X}_t|}{n+m-2} \cdot \Delta$. Due to Lemma 3, ColMidVal($\mathcal{C}_j, \varepsilon_t, \frac{\delta}{2m^2n^2}$) takes $\frac{c \cdot \log(2n^2m^2/\delta)}{\varepsilon_t^2}$ samples where c is an absolute constant. Hence, the total number of samples taken in the iteration t is upper bounded by $|\mathcal{Y}_t| \cdot \frac{n+m-2}{|\mathcal{X}_t|} \cdot \frac{81c \cdot \log(2n^2m^2/\delta)}{\Delta^2} + |\mathcal{X}_t| \cdot \frac{n+m-2}{|\mathcal{X}_t|} \cdot \frac{162 \log(4n^2m^2/\delta)}{\Delta^2} \leq (c+1) \cdot (n+m-2) \cdot \frac{162 \log(4n^2m^2/\delta)}{\Delta^2}$, since $|\mathcal{Y}_t| \leq |\mathcal{X}_t|$. By an identical argument, this bound applies to each iteration $t \in \mathcal{T}_2$.

Finally consider iteration t such that $\max\{|\mathcal{X}_t|, |\mathcal{Y}_t|\} = 2$. The total number of samples taken in this iteration t is upper bounded by $4 \cdot \frac{n+m-2}{2} \cdot \frac{50 \log(16/\delta)}{\Delta^2}$.

We now get the desired sample complexity as $|\mathcal{T}_1| = \log_2(n) - 1$ and $|\mathcal{T}_2| = \log_2(m) - 1$. \square

Finally, we establish the correctness of our algorithm in the following lemma. The proof of the lemma makes use of the fact that the entries around the median of column j_* are far below A_{i_γ, j_γ} . The sampling done from each column ensures that the entries around the median of column \mathfrak{b} are also far below A_{i_γ, j_γ} . This along with the fact that $A_{i_\gamma, \mathfrak{b}} \geq A_{i_\gamma, j_\gamma}$ ensures that we don't eliminate the row i_* .

Lemma 2. *If $\Delta \leq \Delta_g$, then the procedure $\text{MidSearch}(A, \Delta, \delta)$ returns (i_*, j_*) with probability at least $1 - \delta$.*

Proof. Let us consider the case when $\min\{n, m\} \geq 4$. The other cases can be proved analogously. Let

$$\mathcal{T}_1 := \{t : |\mathcal{X}_{t+1}| = \frac{|\mathcal{X}_t|}{2}\} \text{ and } \mathcal{T}_2 := \{t : |\mathcal{Y}_{t+1}| = \frac{|\mathcal{Y}_t|}{2}\}.$$

Consider an iteration $t \in \mathcal{T}_1$ such that $|\mathcal{X}_t| \geq 4$ and let us assume that $i_* \in \mathcal{X}_t$ and $j_* \in \mathcal{Y}_t$. Observe that $\mathcal{Y}_{t+1} = \mathcal{Y}_t$. We now show that with probability at least $1 - \frac{\delta}{nm}$, we have $i_* \in \mathcal{X}_{t+1}$. Let $\Delta_t = \frac{|\mathcal{X}_t|}{n+m-2} \Delta_g$. Recall the definition of ε_t and observe that $\varepsilon_t \leq \Delta_t/9$.

Consider an index $j \in \mathcal{Y}_t$. Let us relabel the indices in \mathcal{X}_t as $\{i_1^{(j)}, \dots, i_{|\mathcal{X}_t|}^{(j)}\}$ such that $A_{i_1^{(j)}, j} \geq \dots \geq A_{i_{|\mathcal{X}_t|}^{(j)}, j}$. Let G_j be the event that $\mathfrak{b}_{t,x}(j) \in [A_{i_{j \times t_j = 2}, j}^{(j)} - \varepsilon_t, A_{i_{j \times t_j = 4+1}, j}^{(j)} + \varepsilon_t]$. Due to Lemma 4, event G_j holds with probability at least $1 - \frac{\delta}{2m^2n^2}$.

Let us assume that G_j holds for all $j \in \mathcal{Y}_t$. This happens with probability at least $1 - \frac{\delta}{2nm^2}$ due to a union bound. As event G_{j_γ} holds, we have the following due to Proposition 1:

$$\begin{aligned} \mathfrak{b}_{t,x}(j_*) &\leq A_{i_{j \times t_j = 4+1}, j_\gamma}^{(j_\gamma)} + \frac{\Delta_t}{9}, \\ &\leq A_{i_\gamma, j_\gamma} - \frac{\Delta_t}{2} + \frac{\Delta_t}{9} = A_{i_\gamma, j_\gamma} - \frac{7\Delta_t}{18}. \end{aligned}$$

Recall the definition of \mathfrak{b} . Since event $G_{\mathfrak{b}}$ holds, we have the following for all $\frac{|\mathcal{X}_t|}{2} \leq s \leq |\mathcal{X}_t|$:

$$A_{i_s^{(\mathfrak{b})}, \mathfrak{b}} \leq \mathfrak{b}_{t,x}(\mathfrak{b}) + \frac{\Delta_t}{9} \leq \mathfrak{b}_{t,x}(j_*) + \frac{\Delta_t}{9} \leq A_{i_\gamma, j_\gamma} - \frac{5\Delta_t}{18}.$$

Let us assume that for all $i \in \mathcal{X}_t$, $|\bar{A}_{i, \mathfrak{b}} - A_{i, \mathfrak{b}}| \leq \frac{\Delta_t}{9}$. This happens with probability at least $1 - \frac{\delta}{2nm^2}$ due to the sub-Gaussian tail bound and union bound. We now have the following for all $\frac{|\mathcal{X}_t|}{2} \leq s \leq |\mathcal{X}_t|$:

$$\bar{A}_{i_s^{(\mathfrak{b})}, \mathfrak{b}} \leq A_{i_s^{(\mathfrak{b})}, \mathfrak{b}} + \frac{\Delta_t}{9} \leq A_{i_\gamma, j_\gamma} - \frac{3\Delta_t}{18}.$$

This implies that $|\{i \in \mathcal{X}_t : \bar{A}_{i, \mathfrak{b}} \leq A_{i_\gamma, j_\gamma} - \frac{3\Delta_t}{18}\}| \geq \frac{|\mathcal{X}_t|}{2}$. Recall that $\mathcal{X}_{t+1} = \{i_1, \dots, i_{|\mathcal{X}_t|/2}\} \subset \mathcal{X}_t =$

$\{i_1, \dots, i_{|\mathcal{X}_t|}\}$ where $\bar{A}_{i_1, \mathfrak{b}} \geq \dots \geq \bar{A}_{i_{|\mathcal{X}_t|}, \mathfrak{b}}$. Now observe that $\bar{A}_{i_\gamma, \mathfrak{b}} \geq A_{i_\gamma, \mathfrak{b}} - \frac{\Delta_t}{9} \geq A_{i_\gamma, j_\gamma} - \frac{\Delta_t}{9}$. Hence, $i_* \in \mathcal{X}_{t+1}$.

Analogously we can show that for an iteration $t \in \mathcal{T}_2$ such that $|\mathcal{Y}_t| \geq 4$, $i_* \in \mathcal{X}_t$ and $y_* \in \mathcal{Y}_t$, with probability at least $1 - \frac{\delta}{nm}$ we have $i_* \in \mathcal{X}_{t+1}$ and $y_* \in \mathcal{Y}_{t+1}$.

Observe that the algorithm terminates at a fixed iteration $t_* = \log_2(nm) - 1$. Let $p_t = \mathbb{P}(i_* \in \mathcal{X}_t, j_* \in \mathcal{Y}_t | i_* \in \mathcal{X}_{t-1}, j_* \in \mathcal{Y}_{t-1})$. Now we have the following due to the chain rule and Bernoulli's inequality,

$$\mathbb{P}(i_* \in \mathcal{X}_{t_\gamma}, j_* \in \mathcal{Y}_{t_\gamma}) = \prod_{t=2}^{t_\gamma} p_t \geq 1 - \frac{\delta}{nm} t_\gamma \geq 1 - \delta/2.$$

Hence, with probability at least $1 - \delta/2$ we have an iteration t_* such that $\mathcal{X}_{t_\gamma} = \{i_*, \mathfrak{b}\}$ and $\mathcal{Y}_{t_\gamma} = \{j_*, \mathfrak{b}\}$. Under the assumption of the existence of such an iteration t_* , we now show that with probability at least $1 - \delta/2$ we return (i_*, j_*) . Let $\Delta_{t_\gamma} = \frac{|\mathcal{X}_{t_\gamma}|}{n+m-2} \Delta_g$. For all $(i, j) \in \mathcal{X}_{t_\gamma} \times \mathcal{Y}_{t_\gamma}$, let $\bar{A}_{i,j}$ denote the empirical average of the $\frac{n+m-2}{|\mathcal{X}_{t_\gamma}|} \cdot \frac{50 \log(16/\delta)}{\Delta_g^2}$ samples. Let us assume that for all $(i, j) \in \mathcal{X}_{t_\gamma} \times \mathcal{Y}_{t_\gamma}$, $|\bar{A}_{i,j} - A_{i,j}| \leq \frac{\Delta_{t_\gamma}}{5}$. This happens with probability at least $1 - \delta/2$ due to sub-gaussian tail bound and union bound. Now observe that $A_{i_\gamma, j_\gamma} - \frac{\Delta_{t_\gamma}}{5} \leq \bar{A}_{i_\gamma, j_\gamma} \leq A_{i_\gamma, j_\gamma} + \frac{\Delta_{t_\gamma}}{5}$. Due to Proposition 2, we have $\bar{A}_{i_\gamma, \mathfrak{b}} \geq A_{i_\gamma, \mathfrak{b}} - \frac{\Delta_{t_\gamma}}{5} \geq A_{i_\gamma, j_\gamma} + \frac{\Delta_{t_\gamma}}{2} - \frac{\Delta_{t_\gamma}}{5} > \bar{A}_{i_\gamma, j_\gamma}$. Similarly due to Proposition 1, we have $\bar{A}_{\mathfrak{b}, j_\gamma} \leq A_{\mathfrak{b}, j_\gamma} + \frac{\Delta_{t_\gamma}}{5} \leq A_{i_\gamma, j_\gamma} - \frac{\Delta_{t_\gamma}}{2} + \frac{\Delta_{t_\gamma}}{5} < \bar{A}_{i_\gamma, j_\gamma}$. Hence, we return (i_*, j_*) . \square

2.2 COLMIDVAL & ROWMIDVAL Subroutine

Here, we define ColMidVal (Algorithm 2) which solves the subproblem of finding a value near the median of means of n arms; RowMidVal is defined analogously in Appendix C along with its guarantees.

Consider $\varepsilon, \delta > 0$. Given a set of n arms \mathcal{A} with means $\mu_1 \geq \mu_2 \geq \dots \geq \mu_n$, we show that ColMidVal outputs a value $\mathfrak{b} \in [\mu_{n/2} - \varepsilon, \mu_{n/4+1} + \varepsilon]$ with probability $1 - \delta$ by using at most $O(\frac{1}{\varepsilon^2} \log \frac{1}{\delta})$ samples. We assume that n is a multiple of 4 for simplicity.

We first establish the sample complexity of ColMidVal in the following lemma.

Lemma 3. *ColMidVal($\mathcal{A}, \varepsilon, \delta$) requires at most $\frac{c \log(1/\delta)}{\varepsilon^2}$ samples.*

Proof. As k, δ_1 are constants, the total number of sam-

Algorithm 2 Col MidVal ($\mathcal{A}, \varepsilon, \delta$)

-
- 1: $\ell \leftarrow \lceil 14 \log(\frac{1}{\delta}) \rceil$, $\delta_1 \leftarrow 0.05$, $\delta_2 \leftarrow 0.05$, $k \leftarrow 108 \log(\frac{4}{\delta_2})$ and $z \leftarrow \frac{k}{3} + 1$.
 - 2: **for** $i = 1$ to ℓ **do**
 - 3: $\mathcal{B}_i \leftarrow \{k \text{ arms sampled independently and uniformly at random from } \mathcal{A}\}$.
 - 4: Sample each arm $j \in \mathcal{B}_i$ $\frac{2 \log(\frac{2k}{\varepsilon^2})}{\varepsilon^2}$ times and compute the empirical mean $\hat{\mu}_j$.
 - 5: $v_i \leftarrow z$ -th highest value in $\{\hat{\mu}_j : j \in \mathcal{B}_i\}$
 - 6: **end for**
 - 7: **return** the median of $\{v_i : i \in [\ell]\}$
-

ples is

$$14 \log \frac{1}{\delta} \cdot k \cdot \frac{2 \log(\frac{2k}{\varepsilon^2})}{\varepsilon^2} \leq \frac{c \log(1/\delta)}{\varepsilon^2}$$

for an absolute constant c . \square

Next, we establish the correctness of Col MidVal. The proof follows from the standard median of means argument.

Lemma 4. *Given n arms \mathcal{A} with means $\mu_1 \geq \mu_2 \geq \dots \geq \mu_n$, Col MidVal ($\mathcal{A}, \varepsilon, \delta$) outputs a value $\mathfrak{b} \in [\mu_{n/2} - \varepsilon, \mu_{n/4+1} + \varepsilon]$ with probability at least $1 - \delta$.*

Proof. Consider an iteration i . Let $n_{1,i}$ and $n_{2,i}$ be the number of times arms 1 through $n/4$ and arms $n/4 + 1$ through $n/2$ were drawn for \mathcal{B}_i , respectively; that is, the counts of occurrences corresponding to the first and second quantiles of highest arm means. We then say the event E_i holds if

- $\frac{7k}{40} \leq n_{1,i} \leq \frac{k}{3}$ and $\frac{7k}{40} \leq n_{2,i} \leq \frac{k}{3}$, and
- for all $j \in \mathcal{B}_i$, $|\hat{\mu}_j - \mu_j| \leq \varepsilon$.

By a standard sub-Gaussian tail bound and union bound, we have that $|\hat{\mu}_j - \mu_j| \leq \varepsilon$ simultaneously for all $j \in \mathcal{B}_i$ with probability at least $1 - \delta_1$. This implies that any sampled arm not in the first quantile has an empirical mean of at most $\mu_{n/4+1} + \varepsilon$. Similarly any sampled arm from the first and second quantile has an empirical mean of at least $\mu_{n/2} - \varepsilon$. Next, observe that $\mathbb{E}[n_{1,i}] = \mathbb{E}[n_{2,i}] = k/4$. Using the Chernoff bound and union bound, we have $\frac{7k}{40} \leq |n_{1,i}| \leq \frac{k}{3}$ and $\frac{7k}{40} \leq |n_{2,i}| \leq \frac{k}{3}$ with probability at least $1 - \delta_2$. Thus, by a union bound, $\mathbb{P}[E_i] \geq 1 - \delta_1 - \delta_2 = \frac{9}{10}$.

We first claim that if E_i holds, then $v_i \in [\mu_{n/2} - \varepsilon, \mu_{n/4+1} + \varepsilon]$. We see that the z -th highest empirical mean is in the interval $[\mu_{n/2} - \varepsilon, \mu_{n/4+1} + \varepsilon]$ if there are *at most* $z - 1$ empirical means greater than $\mu_{n/4+1} + \varepsilon$ and there are *at least* z empirical means greater than

or equal to $\mu_{n/2} - \varepsilon$. Noting that $z = \frac{k}{3} + 1$, on E_i , the first condition occurs as $n_{1,i} \leq k/3 = z - 1$, and the second condition occurs as $n_{1,i} + n_{2,i} \geq \frac{7k}{20} > z$.

Finally, let X_i be the indicator variable $\mathbb{1}_{E_i}$, and let $Z = \sum_{i=1}^{\ell} X_i$ be the total number of times E_i occurs. If $Z > \ell/2$, then the median of $\{v_i : i \in [\ell]\}$ lies in $[\mu_{n/2} - \varepsilon, \mu_{n/4+1} + \varepsilon]$. Observing that $\mathbb{E}[Z] \geq \frac{9\ell}{10}$, where $\ell = \lceil 14 \log(\frac{1}{\delta}) \rceil$, by the Chernoff bound we have $\mathbb{P}[Z \geq 0.54\ell] \geq 1 - \delta$. Hence, Col MidVal ($\mathcal{A}, \varepsilon, \delta$) returns a value $\mathfrak{b} \in [\mu_{n/2} - \varepsilon, \mu_{n/4+1} + \varepsilon]$ with probability at least $1 - \delta$. \square

2.3 Algorithm without the knowledge of Δ_g

First we describe the guarantees of a procedure based on stochastic bandit algorithms that is used to verify whether an entry $(\mathfrak{i}, \mathfrak{j})$ is a PSNE of A or not.

Lemma 5. *Given an entry $(\mathfrak{i}, \mathfrak{j})$ and parameters Δ and δ , there is a procedure Verify possessing the following properties with probability $1 - \delta$:*

- If $\Delta \leq \Delta_g$ and $(\mathfrak{i}, \mathfrak{j}) = (i_*, j_*)$, then $(\mathfrak{i}, \mathfrak{j})$ is “accepted” as the PSNE of A by Verify.
- If $(\mathfrak{i}, \mathfrak{j}) \neq (i_*, j_*)$, then $(\mathfrak{i}, \mathfrak{j})$ is “not accepted” as the PSNE of A by Verify.

Moreover, Verify terminates after taking $c \cdot \mathbf{H}_* \cdot \log(\frac{\log(\mathbf{H}_*)}{\delta})$ samples from A where c is an absolute constant and $\mathbf{H}_* = \frac{n+m-2}{\Delta^2}$.

We refer the reader to Appendix D for more details.

Algorithm 3 is a meta-algorithm for finding (i_*, j_*) with high probability without the knowledge of Δ_g . Theorem 1 establishes the guarantees for Algorithm 3.

Algorithm 3

-
- 1: **Given:** Probability error term δ .
 - 2: **for** $t = 1, 2, \dots$ **do**
 - 3: $\Delta_t \leftarrow \frac{1}{2^{t-1}}$ and $\delta_t \leftarrow \frac{\delta}{4t^2}$
 - 4: $(i_t, j_t) \leftarrow \text{MidSearch}(A, \Delta_t, \delta_t)$
 - 5: **if** Verify($A, i_t, j_t, \delta_t, \Delta_t$) = “accepted” **then**
 - 6: **Return** (i_t, j_t) as the PSNE
 - 7: **end if**
 - 8: **end for**
-

Proof of Theorem 1. Recall that $\Delta_g = \frac{(n+m-2)}{\mathbf{H}_1}^{1/2}$. Consider an iteration $t \leq \lceil \log_2(1/\Delta_g) \rceil + 1$. If $(i_t, j_t) \neq (i_*, j_*)$ then, due to Lemma 5, (i_t, j_t) is not returned as the PSNE with probability at least $1 - \frac{\delta}{4t^2}$. Consider the iteration $t_* = \lceil \log_2(1/\Delta_g) \rceil + 2$. Observe that $\Delta_{t_*} \leq \frac{1}{2^{\lceil \log_2(1/\Delta_g) \rceil}} = \Delta_g$. Due to Lemma 2, with probability at least $1 - \frac{\delta}{4t_*^2}$, we have $(i_{t_*}, j_{t_*}) = (i_*, j_*)$. Due

to Lemma 5, the pair (i_*, j_*) is accepted and returned as a PSNE with probability at least $1 - \frac{\delta}{4t^2}$.

Hence, with probability at least $1 - \frac{\delta}{4t^2} - \prod_{t=1}^{\infty} \frac{\delta}{4t^2} > 1 - \delta$ the algorithm terminates by the iteration $t_* = \lceil \log_2(1/\Delta_g) \rceil + 2$ and returns (i_*, j_*) as the PSNE.

Let us assume that the algorithm terminates by the iteration $t_* = \lceil \log_2(1/\Delta_g) \rceil + 2$. Due to Lemma 1 and Lemma 5, the number of samples in any iteration t is upper bounded by $c \cdot \frac{n+m-2}{\Delta_t^2} \cdot \log \frac{nm \log(\frac{1}{\Delta_t})}{\delta_t} \cdot \log(nm)$ where c is an absolute constant. The sample complexity τ of Algorithm 3 is upper bounded as follows:

$$\begin{aligned} \tau &\leq \sum_{t=1}^{\infty} c \cdot \frac{n+m-2}{\Delta_t^2} \cdot \log \frac{nm \log(\frac{1}{\Delta_t})}{\delta_t} \cdot \log(nm), \\ &\leq \sum_{t=1}^{\infty} c \cdot 4^t (n+m-2) \cdot \log \frac{4t^3 nm}{\delta} \cdot \log(nm), \\ &\leq c \cdot (n+m-2) \cdot \log \frac{4t_*^3 \cdot nm}{\delta} \cdot \log(nm) \cdot 4^{t_*}, \end{aligned}$$

where the second inequality holds since $\Delta_t = \frac{1}{2^{t-1}}$ and $\delta_t = \frac{\delta}{4t^2}$. Using the fact that $\sum_{t=1}^{\infty} 4^t \leq \frac{4^{t_*+1}}{3}$, we further deduce that

$$\begin{aligned} \tau &\leq \frac{4c}{3} \cdot (n+m-2) \cdot \log \frac{4t_*^3 \cdot nm}{\delta} \cdot \log(nm) \cdot 2^{2t_*}, \\ &\leq \frac{256c}{3} \cdot \frac{n+m-2}{\Delta_g^2} \cdot \log \frac{4t_*^3 \cdot nm}{\delta} \cdot \log(nm), \end{aligned}$$

where the last inequality is due to $t_* \leq \log_2(1/\Delta_g) + 3$. The claim in the theorem follows from the fact that $\Delta_g = \left(\frac{n+m-2}{H_1}\right)^{1/2}$. \square

2.4 Dueling Bandits

Recall that in the context of dueling bandits defined by the matrix \mathbf{P} , we have $\Delta_{i,j} = |1/2 - \mathbf{P}_{i,j}|$. The lower bound to identify the Condorcet winner i_* , whenever it exists, scales roughly as $\sum_{i \neq i_*} \Delta_{i,i_*}^{-2}$. Now, an instance of dueling can be interpreted as a two-player zero-game on a matrix A where $A = \mathbf{P}$. It is straightforward to see that the PSNE of A is (i_*, i_*) . Algorithm 3 achieves the bound of $\sum_{i \neq i_*} \Delta_{i,i_*}^{-2}$ up to log factors.

3 EXPERIMENTS

To illustrate the performance of MidSearch, we introduce the following class of $d \times d$ reward matrices, where we use bold to denote a vector of repeated en-

tries:

$$A_{\text{hard}}(\Delta_{\min}, \beta) := \begin{array}{cc} & \begin{array}{c} 3 \\ 2 \end{array} \\ \begin{array}{c} 6 \\ 6 \\ 4 \end{array} & \begin{array}{c} 0.5 \\ \text{---} \\ 0.5 - \Delta_{\min} \\ \text{---} \\ 0.5 - \beta \end{array} & \begin{array}{c} | \\ \text{---} \\ | \\ \text{---} \\ | \end{array} & \begin{array}{c} 0.5 + \Delta_{\min} \\ \text{---} \\ 0.5I + U \end{array} & \begin{array}{c} | \\ \text{---} \\ | \\ \text{---} \\ | \end{array} & \begin{array}{c} (\mathbf{0.5} + \beta)^\top \\ \text{---} \\ 5 \end{array} & \begin{array}{c} 3 \\ 7 \\ 5 \end{array} \end{array}$$

where U is the matrix with all ones on the (strictly) upper triangular entries and $0 < \Delta_{\min} < \beta$. These instances are difficult for the class of learners that play two no-regret learners against themselves, as the second row and second column have high and low sums, respectively, causing $A_{2,2}$ to potentially appear optimal. Note, this is also a dueling bandit instance, which is a special case of a matrix game.

As baselines, we compare MidSearch¹ to Tsallis-INF (Zimmert and Seldin, 2021), Exp3-IX (Neu, 2015), LUCB-G (Zhou et al., 2017), and uniform random sampling for different matrices in A_{hard} . Each algorithm is given a fixed budget T of samples that it can draw from the input matrix and we measure its percentage of identifying the PSNE for the given value of T . Code to replicate these experiments is available at <https://github.com/aistats2024-noisy-psne/midsearch>; supplementary experimental results can be found in Appendix F.

Let $H_1 = \frac{d-2}{\beta^2} + \frac{1}{\Delta_{\min}^2}$. Figure 1 shows that even with a budget of $50H_1$ samples, the Exp3-IX, LUCB-G², and Uniform algorithms already begin to degrade in performance at $d = 32$, returning the correct PSNE on 282, 121, and 99 of the 300 trials, respectively (Exp3-IX degrades slightly slower, worsening to 150/300 at $d = 128$).

Thus, for larger matrices, we only compare MidSearch and Tsallis-INF. In Figure 2 we plot the performance of MidSearch vs. Tsallis-INF for instances of $A_{\text{hard}}(\Delta_{\min}, 0.1)$ with $d = 1024$, varying only Δ_{\min} . We see that after $T = 50H_1$ samples, MidSearch remains close to perfect for the $d = 1024$ matrix, while the accuracy of Tsallis-INF decreases as Δ_{\min} decreases from 0.1 to 0.0125.

Note that the sample complexity for Uniform is $d^2 \Delta_{\min}^{-2}$, LUCB-G is no worse than $d \Delta_{\min}^{-2} + d^2$, Exp3-IX is no worse than $d \Delta_{\min}^{-2}$, Tsallis-INF is no worse

¹The subroutines COLMIDVAL and ROWMIDVAL are wasteful. We optimize them for empirical performance by using an aggressive procedure similar to successive halving. Furthermore, a fixed budget of T is achieved for MIDSEARCH by setting $\Delta \approx \left(\frac{n+m}{T}\right)^{1/2}$ and $\delta \approx \text{constant}$.

²For LUCB-G, we set $\delta \approx \text{constant}$ and return the PSNE it identifies by using at most T samples. We take the PSNE of the empirical mean matrix (if it exists) if the budget has been reached.

Figure 1: Plot for each algorithm on the $A_{\text{hard}}(0:05;0:1)$ instance with $N = 300$ trials and with a budget of $50H_1$ samples; they-axis shows what percentage of the time each algorithm's output was correct at different checkpoints during each run. We also overlay the corresponding Wilson binomial confidence interval ($\alpha = 95\%$).

Figure 2: Plot of MidSearch vs. Tsallis-INF for varying ϵ_{\min} parameter. Each algorithm was run for $N = 300$ trials on the $A_{\text{hard}}(\epsilon_{\min};0:1)$ instance ($d = 1024$) with a budget of $50H_1$ samples.

than $d \frac{1}{\epsilon_{\min}} + \frac{1}{\epsilon_{\min}^2}$; MidSearch is no worse than $H_1 = d \frac{1}{\epsilon_{\min}^2} + \frac{1}{\epsilon_{\min}}$. Based on these experiments, we conjecture that these bounds are indeed tight.

4 CONCLUSION

In this paper we generalized stochastic bandits and dueling bandits by studying the problem of identifying the PSNE in two-player zero-sum games, whenever it exists. We designed an algorithm that matches the lower bound from Zhou et al. (2017), up to log factors. To the best of our knowledge, this is also the first algorithm that achieves near-optimal guarantees

for identifying Condorcet winner in dueling bandits. Our work leads to the following interesting problems. First, can one design a ϵ -PAC learner that has optimal bounds even in the log factors, similar to the extensive work in multi-armed bandits? Although PSNE is not guaranteed to exist in every matrix, a mixed strategy Nash Equilibrium is always guaranteed to exist. The optimal sample complexity to identify an ϵ -Nash Equilibrium in an interesting open problem (see Maiti et al. (2023a) for more discussion). Last-iterate convergence in matrix games has also gained significant interest in recent years (see Wei et al. (2021)). An important open problem here is to study last-iterate convergence in matrix games with noisy bandit feed-

back. Finally, can one design an algorithm for dueling bandits that simultaneously has optimal sample complexity for identifying the Condorcet winner, when it exists, and optimal regret guarantees with respect to the Condorcet winner (see Saha and Gaillard (2022) for the definition and near-optimal result for regret), up to log factors.

ACKNOWLEDGEMENTS

This work was supported in part by NSF TRIPODS CCF Award #2023166 and a Northrop Grumman University Research Award.

References

- Viktor Bengs, Róbert Busa-Fekete, Adil El Mesaoudi-Paul, and Eyke Hüllermeier. Preference-based online learning with dueling bandits: A survey. *The Journal of Machine Learning Research* 22(1):278–385, 2021.
- Lijie Chen, Jian Li, and Mingda Qiao. Towards instance optimal bounds for best arm identification. In *Conference on Learning Theory*, pages 535–592. PMLR, 2017.
- Peter Duersch, Jörg Oechssler, and Burkhard C Schipper. Pure strategy equilibria in symmetric two-player zero-sum games. *International Journal of Game Theory*, 41:553–564, 2012.
- Eyal Even-Dar, Shie Mannor, and Yishay Mansour. Pac bounds for multi-armed bandit and markov decision processes. In *Computational Learning Theory: 15th Annual Conference on Computational Learning Theory, COLT 2002 Sydney, Australia, July 8–10, 2002 Proceedings*, pages 255–270. Springer, 2002.
- Yoav Freund and Robert E Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1-2):79–103, 1999.
- Björn Haddendorst, Viktor Bengs, Jasmin Brandt, and Eyke Hüllermeier. Testification of condorcet winners in dueling bandits. In *Uncertainty in Artificial Intelligence*, pages 1195–1205. PMLR, 2021.
- Kevin Jamieson, Matthew Malloy, Robert Nowak, and Sébastien Bubeck. lil'ucb: An optimal exploration algorithm for multi-armed bandits. In *Conference on Learning Theory*, pages 423–439. PMLR, 2014.
- Shivaram Kalyan Krishnan, Ambuj Tewari, Peter Auer, and Peter Stone. Pac subset selection in stochastic multi-armed bandits. In *ICML*, volume 12, pages 655–662, 2012.
- Zohar Karnin, Tomer Koren, and Oren Somekh. Almost optimal exploration in multi-armed bandits. In *International conference on machine learning*, pages 1238–1246. PMLR, 2013.
- Emilie Kaufmann, Olivier Cappé, and Aurélien Garivier. On the complexity of best arm identification in multi-armed bandit models. *Journal of Machine Learning Research* 17:1–42, 2016.
- Arnab Maiti, Kevin Jamieson, and Lillian Ratli. Instance-dependent sample complexity bounds for zero-sum matrix games. In *International Conference on Artificial Intelligence and Statistics*, pages 9429–9469. PMLR, 2023a.
- Arnab Maiti, Kevin Jamieson, and Lillian J Ratli. Logarithmic regret for matrix games against an adversary with noisy bandit feedback. *arXiv preprint arXiv:2306.13233*, 2023b.
- Shie Mannor and John N Tsitsiklis. The sample complexity of exploration in the multi-armed bandit problem. *Journal of Machine Learning Research* 5(Jun):623–648, 2004.
- Soheil Mohajer, Changho Suh, and Adel Elmahdy. Active learning for top-k rank aggregation from noisy comparisons. In *International Conference on Machine Learning*, pages 2488–2497. PMLR, 2017.
- Gergely Neu. Explore no more: Improved high-probability regret bounds for non-stochastic bandits. *Advances in Neural Information Processing Systems* 28, 2015.
- Tadeusz Radzik. Saddle point theorems. *International Journal of Game Theory*, 20:23–32, 1991.
- Wenbo Ren, Jia Liu, and Ness Shro. The sample complexity of best-k items selection from pairwise comparisons. In *International Conference on Machine Learning*, pages 8051–8072. PMLR, 2020.
- Aadirupa Saha and Pierre Gaillard. Versatile dueling bandits: Best-of-both world analyses for learning from relative preferences. In *International Conference on Machine Learning*, pages 19011–19026. PMLR, 2022.
- Lloyd Shapley. Some topics in two-person games. *Advances in game theory* 52:1–29, 1964.
- Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, and Haipeng Luo. Linear last-iterate convergence in constrained saddle-point optimization. In *International Conference on Learning Representations*, 2021.
- Yichi Zhou, Jialian Li, and Jun Zhu. Identify the nash equilibrium in static games with random payoffs. In *International Conference on Machine Learning*, pages 4160–4169. PMLR, 2017.
- Julian Zimmert and Yevgeny Seldin. Tsallis-inf: An optimal algorithm for stochastic and adversarial bandits. *The Journal of Machine Learning Research* 22(1):1310–1358, 2021.

Checklist

1. For all models and algorithms presented, check if you include:
 - (a) A clear description of the mathematical setting, assumptions, algorithm, and/or model. [Yes]
 - (b) An analysis of the properties and complexity (time, space, sample size) of any algorithm. [Yes]
 - (c) (Optional) Anonymized source code, with specification of all dependencies, including external libraries. [Not Applicable]
2. For any theoretical claim, check if you include:
 - (a) Statements of the full set of assumptions of all theoretical results. [Yes]
 - (b) Complete proofs of all theoretical results. [Yes]
 - (c) Clear explanations of any assumptions. [Yes]
3. For all figures and tables that present empirical results, check if you include:
 - (a) The code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL). [Yes]
 - (b) All the training details (e.g., data splits, hyperparameters, how they were chosen). [Not Applicable]
 - (c) A clear definition of the specific measure or statistics and error bars (e.g., with respect to the random seed after running experiments multiple times). [Yes]
 - (d) A description of the computing infrastructure used. (e.g., type of GPUs, internal cluster, or cloud provider). [Not Applicable]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets, check if you include:
 - (a) Citations of the creator If your work uses existing assets. [Not Applicable]
 - (b) The license information of the assets, if applicable. [Not Applicable]
 - (c) New assets either in the supplemental material or as a URL, if applicable. [Not Applicable]
 - (d) Information about consent from data providers/curators. [Not Applicable]
 - (e) Discussion of sensible content if applicable, e.g., personally identifiable information or offensive content. [Not Applicable]
5. If you used crowdsourcing or conducted research with human subjects, check if you include:
 - (a) The full text of instructions given to participants and screenshots. [Not Applicable]
 - (b) Descriptions of potential participant risks, with links to Institutional Review Board (IRB) approvals if applicable. [Not Applicable]
 - (c) The estimated hourly wage paid to participants and the total amount spent on participant compensation. [Not Applicable]

A Technical Lemmas

Lemma 6 (sub-Gaussian tail bound). Let $X_1; X_2; \dots; X_n$ be i.i.d samples from a σ -sub-Gaussian distribution with mean μ . Then we have the following for any $0 < \epsilon < 1$:

$$P \left(\frac{1}{n} \sum_{i=1}^n X_i \geq \mu + \epsilon \right) \leq \frac{\sigma^2}{n \epsilon^2} e^{-\frac{n \epsilon^2}{2 \sigma^2}}$$

Lemma 7 (Chernoff Bound). Let $X_1; X_2; \dots; X_n$ be i.i.d samples from a Bernoulli distribution with mean μ . Then we have the following for any $0 < \epsilon < 1$:

$$P \left(\frac{1}{n} \sum_{i=1}^n X_i \geq (1 + \epsilon) \mu \right) \leq e^{-\frac{n \mu \epsilon^2}{3}} \text{ and } P \left(\frac{1}{n} \sum_{i=1}^n X_i \leq (1 - \epsilon) \mu \right) \leq e^{-\frac{n \mu \epsilon^2}{3}}$$

B Stochastic Multi-Armed Bandit

In stochastic Multi-Armed Bandit (MAB), there are n arms where each arm i is associated with a sub-distribution D_i whose mean is μ_i . When an arm i is pulled (or sampled) by a learner, it observes the random variable $X_i \sim D_i$. The arms with the highest mean is referred as the best-arm. In pure exploration, one aims to design a learner that aims to find the best arm with probability at least $1 - \delta$ by pulling as few arms as possible.

Let $i^* = \arg \max_i \mu_i$ and $\Delta_i = \mu_{i^*} - \mu_i$. Optimal bounds to identify the best-arm scales roughly as $\frac{1}{\Delta_i^2} \log(1/\delta)$. Now, an instance of MAB can be interpreted as a two-player zero-game on a matrix $A \in \mathbb{R}^{n \times n}$ where $A_{i,i^*} = \mu_i$. It is easy to observe that the PSNE of A is $(i^*, 1)$. It is easy to observe that our algorithm essentially becomes a variant of Sequential Halving for this special case of MAB and achieves the bound of $\frac{1}{\Delta_i^2} \log(1/\delta)$, up to log factors.

C RowMidVal Subroutine

Here, we define RowMidVal which solves the subproblem of finding a value near the median of means of k arms.

Consider $\delta > 0$. Given a set of n arms A with means $\mu_1, \mu_2, \dots, \mu_n$, RowMidVal outputs a value $\hat{\mu} \in [\mu_{(n-k+1)}, \mu_{(n-k+2)}]$ with probability $1 - \delta$ by using at most $O(\frac{1}{\delta} \log \frac{1}{\delta})$ samples. The proof of this claim follows from an analogous argument as that of ColMidVal.

Algorithm 4 RowMidVal ($A; \delta; \epsilon$)

- 1: $\ell = \lceil 14 \log(\frac{1}{\delta}) \rceil, \alpha_1 = 0.05, \alpha_2 = 0.05, k = \lceil 108 \log(\frac{4}{\delta}) \rceil$ and $z = \frac{k}{3} + 1$.
- 2: for $i = 1$ to ℓ do
- 3: B_i of k arms sampled independently and uniformly at random from A .
- 4: Sample each arm $j \in B_i$ for $\frac{2 \log(\frac{2k}{\delta})}{\alpha_2}$ times and compute the empirical mean b_j .
- 5: v_i = z -th lowest value in $\{b_j : j \in B_i\}$
- 6: end for
- 7: return the median of $\{v_i : i \in [\ell]\}$

D Verification algorithm

In this section, we focus on the following subproblem:

Subproblem: Given an entry (μ, β) and a parameter ϵ , do the following with probability $1 - \delta$:

- If $\mu \geq \mu_{i^*} - \epsilon$ and $(\mu, \beta) = (\mu_{i^*}; j^*)$, then (μ, β) is accepted as the PSNE of A

^ If $(p, p) \notin (i_?, j_?)$, then (p, p) is not accepted as the PSNE of A .

Consider a stochastic multi-armed bandit instance with means $\mu_1, \mu_2, \dots, \mu_n$. Let $H_0 = \sum_{i=2}^n \frac{1}{(\mu_1 - \mu_i)^2}$. Let Alg be a best arm identification algorithm that does the following:

- ^ If the best arm is unique, Alg finds the best arm by using at most $c_1 H_0 \log(\frac{\log H_0}{\epsilon})$ samples with probability at least $1 - \epsilon/2$ where c_1 is an absolute constant.
- ^ If the best arm is not unique, Alg does not terminate with probability at least $1 - \epsilon/2$.

See Karnin et al. (2013) for one such algorithm. We use Alg to design a procedure called Verify that achieves the desired guarantees for the above the sub-problem.

Algorithm 5 Verify $(A; p, p; \epsilon)$

- 1: Let \mathcal{B}_1 be a multi-armed bandit instance with means $\mu_i = A_{i,p}$ for all $i \in [n]$
- 2: Let \mathcal{B}_2 be a multi-armed bandit instance with means $\mu_j = A_{p,j}$ for all $j \in [m]$
- 3: $H_? = \frac{n+m-2}{2}$
- 4: Run Alg on the bandit instance \mathcal{B}_1 and terminate it after it takes $c_1 H_? \log(\frac{\log H_?}{\epsilon})$ samples.
- 5: If Alg terminates without returning arm p as the best arm, then return not accepted.
- 6: Run Alg on the bandit instance \mathcal{B}_2 and terminate it after it takes $c_1 H_? \log(\frac{\log H_?}{\epsilon})$ samples.
- 7: If Alg terminates without returning arm p as the best arm, then return not accepted.
- 8: Return accepted.

We first establish the sample complexity of Verify in the following lemma.

Lemma 8. Verify takes at most $2c_1 H_? \log(\frac{\log H_?}{\epsilon})$ samples before terminating

Proof. The lemma trivially follows from lines 4 and 6 of Verify. □

We prove the following proposition that is useful to establish the correctness of our algorithm

Proposition 3. If $(p, p) \notin (i_?, j_?)$, then either $A_{p,p} < A_{i_?,p}$ or $A_{p,p} > A_{p,j_?}$.

Proof. Let us assume that $A_{p,p} \geq A_{i_?,p}$ and $A_{p,p} \leq A_{p,j_?}$. This implies that $A_{i_?,p} \leq A_{p,j_?}$ which is a contradiction as $A_{p,j_?} < A_{i_?,j_?} < A_{i_?,p}$. □

The following lemma establishes the correctness of Verify.

Lemma 9. Verify does the following with probability $1 - \epsilon$:

- ^ If $\mu_p \geq \mu_g$ and $(p, p) = (i_?, j_?)$, then (p, p) is accepted as the PSNE of A .
- ^ If $(p, p) \notin (i_?, j_?)$, then (p, p) is not accepted as the PSNE of A .

Proof. Let us first consider the case when $\mu_p \geq \mu_g$ and $(p, p) = (i_?, j_?)$. Recall the bandit instances \mathcal{B}_1 and \mathcal{B}_2 . Let $H_x = \sum_{i \in [n], i \neq i_?} \frac{1}{(A_{i_?,i} - A_{i_?,p})^2}$ and $H_y = \sum_{j \in [m], j \neq j_?} \frac{1}{(A_{i_?,j} - A_{i_?,p})^2}$. Observe that $H_x + H_y = \frac{n+m-2}{2} = H_?$. Now observe that $i_?$ is the best arm of the instance \mathcal{B}_1 and $j_?$ is the best arm of the instance \mathcal{B}_2 . Hence, Alg does the following with probability $1 - \epsilon/2$:

- ^ For the instance \mathcal{B}_1 , Alg returns $i_?$ as the best arm after taking at most $c_1 H_x \log(\frac{\log H_x}{\epsilon/2})$ samples.
- ^ For the instance \mathcal{B}_2 , Alg returns $j_?$ as the best arm after taking at most $c_1 H_y \log(\frac{\log H_y}{\epsilon/2})$ samples.

Hence, we accept (p, p) with probability at least $1 - \epsilon/2$.

Next, we consider the case when $(p, p) \notin (i_?, j_?)$. Due to Proposition 3, either p is not the best arm of the instance \mathcal{B}_1 or p is not the best of arm of the instance \mathcal{B}_2 . Hence, with probability $1 - \epsilon/2$, (p, p) is not accepted. □

E Omitted Calculations

Here we state the calculations omitted from various proofs in the main body.

Omitted Calculations from the proof of Lemma 4. Recall that we generate the set B_i by sampling k arms independently and uniformly at random from A where $k = d108 \log(\frac{4}{2})e = 474$. Hence, $\frac{k}{3} + 1$ is an integer equal to 159 and $\frac{7k}{20} > 165$. Also recall that $E[n_{1,i}] = E[n_{2,i}] = k/4$. First we have the following due to the Chernoff bound:

$$\begin{aligned} P[n_{1,i} \geq \frac{7k}{40}] &= P[n_{1,i} \geq 1 + \frac{3}{10} \frac{k}{4}] \\ &= P[n_{1,i} \geq 1 + \frac{3}{10} E[n_{1,i}]] \\ &\leq \exp\left(-\frac{E[n_{1,i}] (0.3)^2}{2}\right) \\ &= \exp\left(-\frac{k \cdot 0.09}{4 \cdot 2}\right) \\ &\leq \exp\left(-\frac{108 \log(\frac{4}{2}) \cdot 0.09}{4 \cdot 2}\right) \\ &= \exp\left(-1.215 \log \frac{4}{2}\right) \\ &< \exp\left(-\log \frac{4}{2}\right) = \frac{2}{4} \end{aligned}$$

Next we have the following due to the Chernoff bound

$$\begin{aligned} P[n_{1,i} \leq \frac{k}{3}] &= P[n_{1,i} \leq 1 + \frac{1}{3} \frac{k}{4}] \\ &= P[n_{1,i} \leq 1 + \frac{1}{3} E[n_{1,i}]] \\ &\leq \exp\left(-\frac{E[n_{1,i}] (1/3)^2}{3}\right) \\ &= \exp\left(-\frac{k}{4 \cdot 3 \cdot 9}\right) \\ &\leq \exp\left(-\frac{108 \log(\frac{4}{2})}{108}\right) \\ &= \exp\left(-\log \frac{4}{2}\right) = \frac{2}{4} \end{aligned}$$

Similarly, we can show using Chernoff bound that $\frac{7k}{40} \leq n_{2,i} \leq \frac{k}{3}$ occurs with probability at least $1 - \frac{2}{4}$.

Recall the definition of Z . Also recall that $\delta = d14 \log(1.1)e$ and $E[Z] = \frac{9}{10}$. Now we have the following due to the

Cherno bound:

$$\begin{aligned}
 P[Z \leq 0.54] &= P\left[n_{1,i} \leq 1 - \frac{2}{5} \frac{9}{10}\right] \\
 &= P\left[n_{1,i} \leq 1 - \frac{2}{5} E[Z]\right] \\
 &\leq \exp\left[-\frac{E[Z] - (0.4)^2}{2}\right] \\
 &\leq \exp\left[-\frac{9 - 0.16}{10 \cdot 2}\right] \\
 &\leq \exp\left[-\frac{9 - 14 \log\left(\frac{1}{t}\right) - 0.16}{10 \cdot 2}\right] \\
 &= \exp\left[-1.008 \log\left(\frac{1}{t}\right)\right] \\
 &< \exp\left[-\log\left(\frac{1}{t}\right)\right] = \frac{1}{t}
 \end{aligned}$$

□

Omitted calculations in the proof of Theorem 1. Using Lemma 1 we get that at any iteration t , MidSearch uses at most $c_1 \frac{n+m}{2} \log\left(\frac{nm}{t}\right) \log(nm)$ samples where c_1 is an absolute constant. Using Lemma 5 we get that at

any iteration t , Verify uses at most $c_2 \frac{n+m}{2} \log\left(\frac{n+m}{t} A\right)$ where c_2 is absolute constant. As $n, m \geq 1$,

we have $\frac{1}{2m-1} \leq \frac{m}{2m-1}$ which implies $n+m \leq 2nm$. Now we have the following:

$$\begin{aligned}
 &\log\left(\frac{n+m}{t} A\right) \leq \log\left(\frac{2nm}{t} A\right) \\
 &\leq \log\left(\frac{\log(2nm) + 2 \log\left(\frac{1}{t}\right)}{t} A\right) \\
 &\leq \log\left(\frac{2nm + 2nm \log\left(\frac{1}{t}\right)}{t} A\right) \quad (\text{as } \log(x) < x) \\
 &= \log\left(\frac{2nm \log\left(\frac{e}{t}\right)}{t} A\right)
 \end{aligned}$$

Now the total number of samples taken at iteration t is upper bounded as follows:

$$\begin{aligned}
 &= c_1 \frac{n+m}{2} \log\left(\frac{nm}{t}\right) \log(nm) + c_2 \frac{n+m}{2} \log\left(\frac{n+m}{t} A\right) \\
 &\leq c_1 \frac{n+m}{2} \log\left(\frac{nm}{t}\right) \log(nm) + c_2 \frac{n+m}{2} \log\left(\frac{2nm \log\left(\frac{e}{t}\right)}{t} A\right) \\
 &\leq c \frac{n+m}{2} \log\left(\frac{nm \log\left(\frac{1}{t}\right)}{t} A\right) \log(nm)
 \end{aligned}$$

where c is an absolute constant.

□

F Additional Plots

In this section, we perform few additional experiments. We consider the hard instance $A_{\text{hard}}(\alpha_{\min}; \cdot)$ described in Section 3 with $d = 128; 256; 512$. As shown in Section 3 both `MidSearch` and `Tsallis-INF` perform far better than `Exp3-IX`, `LUCB-G` and uniform random sampling. Hence, we compare the performance of `MidSearch` and `Tsallis-INF` here by measuring the probability of identifying the PSNE. First in Figure 3, we $\alpha_{\min} = 0:1$ and vary α and compare the performance of `MidSearch` and `Tsallis-INF`. We observe that `MidSearch` has a better performance than `Tsallis-INF`. Next in Figure 4, we $\alpha_{\min} = 0:01$ and vary α and compare the performance of `MidSearch` and `Tsallis-INF`. We observe that `MidSearch` has a better performance than `Tsallis-INF`. Finally in Figure 5, we vary both α and α_{\min} and compare the performance of `MidSearch` and `Tsallis-INF`. We observe that `MidSearch` has a better performance than `Tsallis-INF`.

Figure 3: Plot of `MidSearch` vs. `Tsallis-INF` for varying α_{\min} parameter. Each algorithm was run for $N = 300$ trials on the $A_{\text{hard}}(\alpha_{\min}; 0:1)$ instance with a budget of $50H_1$ samples.

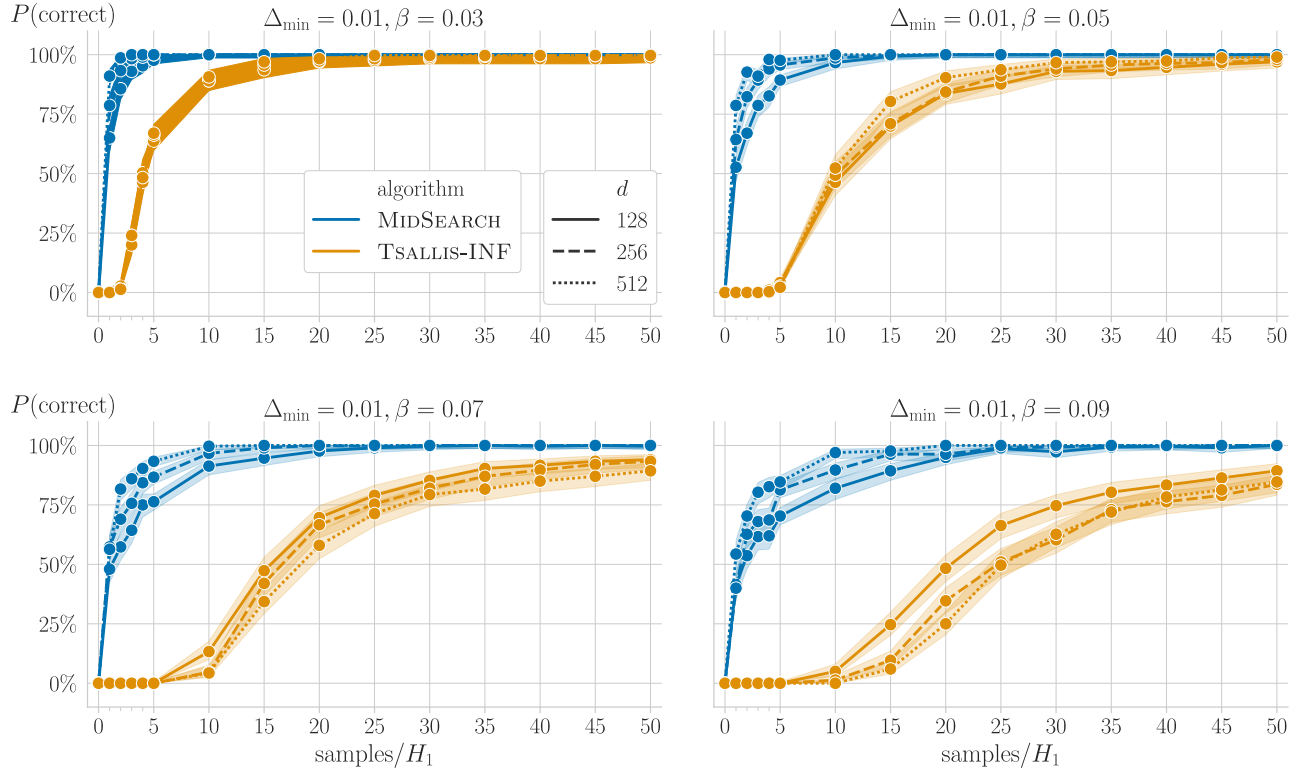


Figure 4: Plot of MidSearch vs. Tsallis-INF for varying β parameter. Each algorithm was run for $N = 300$ trials on the $A_{\text{hard}}(0.01, \beta)$ instance with a budget of $50H_1$ samples.

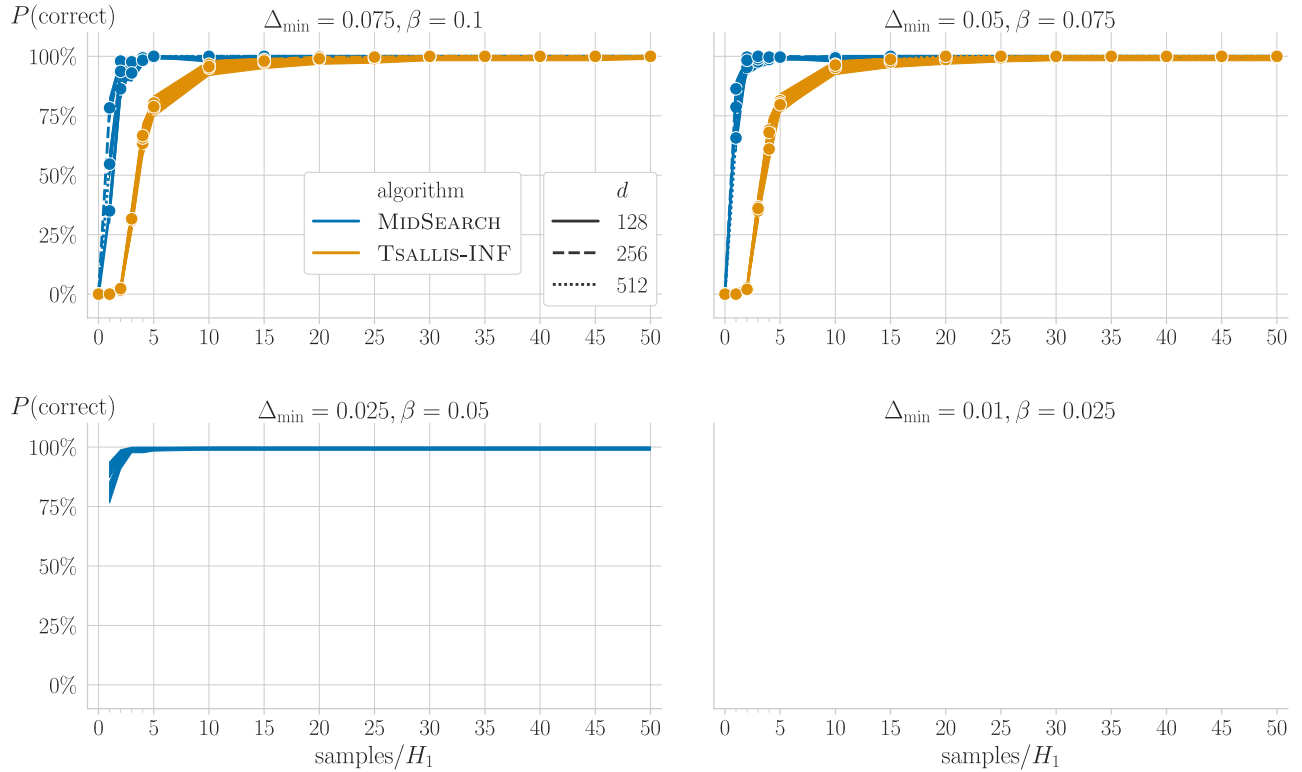


Figure 5: Plot of MidSearch vs. Tsallis-INF for varying both Δ_{\min} and β parameters. Each algorithm was run for $N = 300$ trials on the $A_{\text{hard}}(\Delta_{\min}, \beta)$ instance with a budget of $50H_1$ samples.