

Sampling Equilibria: Fast No-Regret Learning in Structured Games

Daniel Beaglehole* Max Hopkins† Daniel Kane‡ Sihan Liu§ Shachar Lovett¶

Abstract

Learning and equilibrium computation in games are fundamental problems across computer science and economics, with applications ranging from politics to machine learning. Much of the work in this area revolves around a simple algorithm termed *randomized weighted majority* (RWM), also known as “Hedge” or “Multiplicative Weights Update,” which is well known to achieve statistically optimal rates in adversarial settings (Littlestone and Warmuth ’94, Freund and Schapire ’99). Unfortunately, RWM comes with an inherent computational barrier: it requires maintaining and sampling from a distribution over all possible actions. In typical settings of interest the action space is exponentially large, seemingly rendering RWM useless in practice.

In this work, we refute this notion for a broad variety of *structured* games, showing it is possible to efficiently (approximately) sample the action space in RWM in *polylogarithmic* time. This gives the first efficient no-regret algorithms for problems such as the (*discrete*) *Colonel Blotto game*, *matroid congestion*, *matroid security*, and basic *dueling games*. As an immediate corollary, we give a polylogarithmic time meta-algorithm to compute approximate Nash Equilibria for these games that is exponentially faster than prior methods in several important settings. Further, our algorithm is the first to efficiently compute equilibria for more involved variants of these games with general sums, more than two players, and, for Colonel Blotto, multiple resource types. Our results also greatly generalize earlier work on efficient RWM-based techniques for exponential strategy sets from (Cesa-Bianchi and Lugosi ’09).

*Department of Computer Science and Engineering, UCSD, CA 92092. Email: dbeaglehole@ucsd.edu

†Department of Computer Science and Engineering, UCSD, CA 92092. Email: nmhopkin@eng.ucsd.edu. Supported by NSF Award DGE-1650112.

‡Department of Computer Science and Engineering / Department of Mathematics, UCSD, California, CA 92092. Email: dakane@eng.ucsd.edu. Supported by NSF Award CCF-1553288 (CAREER) and a Sloan Research Fellowship.

§Department of Computer Science and Engineering, UCSD, CA 92092. Email: sil046@ucsd.edu

¶Department of Computer Science and Engineering, UCSD, CA 92092. Email: slovett@cs.ucsd.edu. Supported by NSF Award CCF-1953928.

Contents

1	Introduction	3
1.1	Results	4
1.1.1	The Colonel Blotto Game	5
1.1.2	Congestion Games	5
1.1.3	Security Games	6
1.1.4	Dueling Games	7
1.2	Techniques	7
1.2.1	Sampling via Dynamic Programming	8
1.2.2	MCMC-methods	9
1.3	Discussion	10
1.3.1	Dynamic Programming vs MCMC-sampling	10
1.3.2	Further Open Problems	10
1.4	Further Related Work	11
1.4.1	No-regret learning with structured loss	11
1.4.2	Computing equilibria for Colonel Blotto	11
2	Preliminaries	12
2.1	Game Theory	12
2.2	Linear Hypergraph Game	13
2.3	No-Regret Learning in Games	13
2.4	Randomized Weighted Majority Algorithm	14
3	Playing Games via MCMC-Sampling	15
3.1	Glauber Dynamics and Fractionally Log-Concave Games	17
3.2	Dueling Games and the JSV Chain	20
4	Playing Games via DP-Sampling	22
4.1	Sampling via estimation of partition function	24
4.2	Computing the Partition Function	26
4.3	Applications of the meta algorithm	29
4.4	Multi-resource allocation games	31
A	No-regret Learning and Equilibrium Computation	36
B	Bit-complexity and Stability of Numeric Operations	36
C	Implementing Glauber Dynamics	38

1 Introduction

Online learning and equilibrium computation in games has long played a major role in our understanding of human behavior and general multi-agent systems, with applications ranging all the way from politics [BBD⁺19, KR12] and national defense [Tam11] to complexity theory [DGP09, Rub16] and machine learning [FS96, FS99, AB21]. Perhaps the most celebrated line of work in this area is the introduction and analysis of randomized weighted-majority (RWM) and its ‘mixed’ variant (Optimistic) Hedge [LW94, FS96, RS13, DFG21]. These powerful algorithms allow players to engage in repeated gameplay *without regret*, in the sense that the overall loss experienced by any player is not much more than that of the best fixed strategy, even against an arbitrary, adaptive adversary. Such a guarantee is not only powerful in its own right, but is also known to converge quickly to equilibria when performed by all players in repeated rounds of play [CBL06].

Randomized weighted majority is a surprisingly simple algorithm given its powerful guarantees. In each round of a repeated game, a player following RWM samples a strategy s with probability proportional to its (exponentiated) historical loss $\ell(s)$:

$$(1.1) \quad \Pr[\text{Player chooses } s] \propto \beta^{\ell(s)}$$

for some specified ‘learning rate’ $\beta \in (0, 1)$. RWM is also well studied in the setting where the player ‘plays’ the distribution itself (typically called a mixed strategy), and experiences its expected loss. This variant, called *Hedge*, is perhaps the best studied algorithm in all of learning in games [FS96, CBL06].

Unfortunately, while RWM and Hedge are *statistically* optimal [LW94, FS96], they come with an inherent computational barrier: both techniques crucially rely on tracking a distribution over all possible actions. Since the number of actions is typically exponential in the relevant parameters of the game (e.g. in the famous Colonel Blotto problem), this seems to render both Hedge and RWM completely infeasible.

It turns out, however, that this intuition is not entirely correct. In many important settings the distributions that arise from RWM are *highly structured*, and while it still may not be possible to efficiently *output* the distribution itself as in Hedge, it is sometimes possible to efficiently *sample* from it. It is known, for instance, that RWM can be implemented in *polylogarithmic* time when actions are given by the k -edges of a complete hypergraph and rewards decompose linearly over vertices [WK08]. This raises an important question:

When is it possible to efficiently sample in Randomized Weighted Majority?

Toward this end, we introduce a natural generalization of the complete hypergraph setting we call *linear hypergraph games*, where actions are given by k -edges of an *arbitrary* hypergraph, and the reward of any edge similarly decomposes as a sum over individual reward functions on its vertices (see Section 2.2 for more detail). This simple definition captures a surprising number of settings studied in the literature including resource allocation problems like Colonel Blotto [Bor53], along with other widely-studied settings such as congestion [Ros73], security [Tam11, ADH⁺19, Sze17, BB19], and basic dueling games [IKL⁺11, ADH⁺19].

In this work, we show it is indeed possible to efficiently (approximately) sample from RWM over several important subclasses of linear hypergraph games including Colonel Blotto and its variants, matroid congestion [VA06], matroid security [Tam11, ADH⁺19, Sze17, BB19], and basic dueling games [IKL⁺11, ADH⁺19]. This leads to the first algorithms for no-regret learning in these settings that are *polylogarithmic* in the size of the state space, and thereby the first polylogarithmic time algorithms for (approximate) equilibrium computation. On top of giving an exponential improvement over prior results, this also constitutes the first efficient algorithm for equilibrium computation whatsoever in several more involved settings such as dice games, Colonel Blotto with multiple resources, and for multiplayer and general-sum variants of all games we consider. We informally summarize the main results in Table 1.

Our techniques are largely based on two main paradigms: dynamic programming, and Monte Carlo Markov Chains (MCMC). Generalizing seminal work on learning k -sets and other structured concepts [WK08, KWK⁺10], we show that the distributions arising from RWM on linear hypergraph games correspond to well-studied structure in approximate sampling and statistical physics called *external fields*. In resource allocation games like Colonel Blotto that are played over (ordered) fixed-size partitions of n , we exploit this structure to build a dynamic program that approximately computes the normalization factor of Equation (1.1) (often called the partition function). On the other hand, in settings like matroid congestion and security, we rely on deep results from the MCMC-sampling literature showing that any hypergraph that is a sufficiently good high dimensional expander can be sampled under arbitrary

external fields [ALOV19, AASV21, AJK⁺21]. To the authors’ knowledge, these are the first applications of approximate sampling techniques to game theory.

Game	Known runtime (exact)	Our runtime (apx.)
(Discrete) Colonel Blotto	$n^{13}k^{14}$ [ADH ⁺ 19]	$k^4 \log^4(n)L_{\max}^7$
Dice	NA	$k^4 \log^4(n)$
Matroid Security	n^4k^2 [BB19]	$k^3 \log^3(n)L_{\max}^4$
Matroid Congestion	nk [ARV08]	$k^4 \log^3(n)L_{\max}^4$

Table 1: Rough asymptotic runtimes for equilibrium computation in contrast to previous methods with two players. The approximation factor is set to $O(1)$ along with some other game specific parameters (See Definition 3.5 for details).

1.1 Results We briefly review the theory of games, equilibria, and no-regret learning before discussing our results in more formality. Games are mathematical objects that model (possibly non-cooperative) interaction between rational agents. A (simultaneous) game consists of a set of actions A_i for each player, and reward functions R_i mapping action tuples to rewards (real numbers). Players seek to maximize their own reward, and optimal play is typically characterized by Nash equilibria: randomized strategies such that no player can improve by deviating. By the historic result of [Nas51], every finite game has at least one NE. As they are not always efficiently computable [DGP06], one often instead hopes to understand weaker notions such as Coarse Correlated Equilibria (CCE), where the strategies of different players are chosen in coordination with one another (see Section 2.1).

There is a deep connection between equilibrium computation and no-regret learning in games. We consider the typical adaptive online setting in which, in each round, a learner chooses an action and receives an adversarially selected loss that may depend on the learner’s previous actions (see [CBL06, Chapter 4]). An algorithm is said to have “no-regret” when the expected loss suffered by the learner in sequential rounds grows sublinearly compared to the loss of the best fixed action in hindsight. No-regret learning is itself a powerful tool, as it allows for optimal play against sub-optimal opponents (unlike equilibria which only model the setting where all agents play optimally). Furthermore, it is well known that any no-regret algorithm¹ leads to approximate equilibrium computation with similar runtime simply by simulating the algorithm for all players for sufficiently many rounds. RWM, for instance, is well-known to satisfy the following (optimal) regret guarantee.

LEMMA 1.1. (RWM IS NO-REGRET [CBL06, LEMMA 4.1]) *The regret of RWM over T rounds, N actions, and with rewards in $[-L_{\max}, L_{\max}]$, satisfies*

$$\text{Reg}_T \leq O\left(L_{\max} \sqrt{T \log(N)}\right),$$

against any adaptive adversary (with high probability).

In fact, it is important to note in our setting that essentially all guarantees of RWM also hold in the approximate regime, where the learner only δ -approximately samples from the distribution in Equation (1.1) in each round (in Total Variation distance). We call such algorithms δ -approximate RWM (δ -RWM). It is not hard to show that δ -RWM also satisfies the above regret guarantees for small enough δ (see Lemma 2.1). We now cover four of the main settings in which we give new algorithms for no-regret learning and equilibrium computation through efficient implementatin of δ -RWM: Colonel Blotto, Matroid Security, Matroid Congestion, and Dueling games. We note that all results are given in the algebraic computation model for simplicity (where algebraic operations such as addition and subtraction are considered to be unit time), but can easily be moved to the standard bit model with no substantial loss in running time (see Appendix B).

¹Formally we may need to require that a players strategy depends only on the opponents history and not their own [CBL06]. This is satisfied by all algorithms considered in this work.

1.1.1 The Colonel Blotto Game The Colonel Blotto game was originally described by Borel in 1921 [Bor53] and formalizes how warring colonels should distribute soldiers over different battlefields. In the most general version of this game, two colonels have n_1 and n_2 soldiers that they must assign to k different battlefields, each with a non-negative integer weight. A colonel wins a battle (receiving its weight in reward) if they assign more armies to that battle than their opponent. Each colonel seeks to maximize the total weight of battles won in a single assignment.

Despite its breadth of applications and simplicity to state, the first polynomial time algorithm to compute optimal strategies for this game was only developed recently in [ADH⁺19]. This breakthrough result deservedly received significant media attention [Ins16, Mew16], but struggled to see any practical use due to an infeasible $O(n^{13}k^{14})$ running time (where $n = \max\{n_1, n_2\}$). To this day, this is the only known algorithm to provably compute exact optimal strategies for the (discrete) Colonel Blotto game with arbitrary parameters in polynomial time. Though some progress has been made towards more practical algorithms in different settings [BDD⁺17], even these methods cannot handle parameters beyond a few hundred troops [VLS18].

Indeed, solving the Colonel Blotto problem is now only more relevant than it was in 1921, with practical applications in a large swath of market competitions including advertising and auctions [Rob06], budget allocation [Kva07], elections [LP02], and even ecological modeling [GP09]. We give the first no-regret learning algorithm for the Colonel Blotto games under the most general setting [KR21], where rewards are heterogeneous across battles and players and different players are allowed different troop capacities. Moreover, our algorithm runs in time *polylogarithmic* in the state space, making it extremely efficient in the regime where $n \gg k$ (i.e. there are many more troops than battlefields).

THEOREM 1.1. (BLOTTO WITHOUT REGRET (INFORMAL COROLLARY 4.1)) *In a Colonel Blotto game, for a player with n soldiers, k battlefields, and maximum reward bounded by L_{\max} , δ -RWM can be implemented over T rounds of play in time:*

$$\tilde{O}\left(T^3 L_{\max} k \log(n) \delta^{-1}\right),$$

and is no-regret. In the regime where $n = O(k^2)$, we give a faster algorithm running in time $\tilde{O}(Tnk)$.

Theorem 1.1 is the first no-regret algorithm for Colonel Blotto in online adaptive settings, and also gives the fastest known algorithms to compute (approximate) Nash equilibria provided the game is zero-sum, and approximate coarse correlated equilibrium in general sum settings with many players. We state the theorem for the two-player zero-sum setting here.

COROLLARY 1.1. (EQUILIBRIUM COMPUTATION FOR BLOTTO (INFORMAL COROLLARY 4.2)) *Let $n = \max(n_1, n_2)$, where n_1, n_2 are the soldier counts for the two player Colonel Blotto game. Let L_{\max} be maximum reward of the game. There exists an algorithm to compute an ε -approximate Nash equilibrium for the two-player Colonel Blotto Game in time*

$$\tilde{O}\left(L_{\max}^7 k^4 \log^4(n) \varepsilon^{-6}\right)$$

with high probability. When $n = O(k^2)$, we give a faster algorithm running in time $\tilde{O}(nk^2 L_{\max}^2 \varepsilon^{-2})$.

Not only is this algorithm exponentially faster than any prior work in most relevant scenarios (namely when $n \gg k$), it is also the first known method for computing CCE for *multiplayer* Blotto at all. Even more generally, our algorithm extends to a number of other variants of Blotto (or ‘resource allocation’ problems) such as Dice games and settings with multiple types of troops known as the *Multi-resource Colonel Blotto* problem [BDD⁺17] (though in this latter setting we lose the logarithmic dependence on n). We cover these further applications in Section 4.3 and Section 4.4.

1.1.2 Congestion Games Another natural example is a *congestion game*, a class introduced by Rosenthal [Ros73] to model resource competition among greedy players. In a congestion game, m players compete to select from a set of n resources and receive rewards depending on how many players chose a particular resource. Classical examples of congestion games include routing traffic (pick the least congested route) and variants of the famed El Farol Bar Problem [Art94] (players aim to choose a bar that is neither too under nor over-crowded).

Unlike Blotto, equilibrium computation is known to be hard for congestion games, namely (PPAD \cap PLS)-complete [BR21]. However, this can be circumvented when the underlying strategy spaces are sufficiently combinatorially

structured. It has long been known, for instance, that a Nash equilibrium can be found in time $\tilde{O}(m^2 n q k)$ via iterated best-response when all strategies are given by the bases of a rank- k matroid² over n resources of q types [ARV08]. We show matroid congestion games are similarly well-behaved under RWM, and provide a near-optimal no-regret algorithm in both a computational and statistical sense.

THEOREM 1.2. (CONGESTION WITHOUT REGRET (INFORMAL COROLLARY 3.1)) *Let $\mathcal{I} = \{\{A_i\}_{i=1}^m, c\}$ be a congestion game over a size- n ground set Ω with q resource types where each A_i is the set of bases of a rank- k matroid. Then δ -RWM can be implemented for T rounds in time*

$$\tilde{O}\left(kT \log(n) \log(\delta^{-1})(q + kmT)\right),$$

and is no-regret.

To our knowledge, this is the first efficient no-regret algorithm for matroid congestion. Moreover, in the setting where there are polylog(n) resource types, the algorithm leads to exponentially faster (approximate) equilibrium computation than the typical best response strategy (albeit for CCE rather than Nash).

COROLLARY 1.2. (EQUILIBRIUM COMPUTATION FOR CONGESTION (INFORMAL COROLLARY 3.2)) *Let $\mathcal{I} = \{\{A_i, R_i\}_{i=1}^m\}$ be a congestion game over a size- n ground set Ω with q resource types where each A_i is the set of bases of a rank- k matroid. Then there exists an algorithm to compute an ε -CCE in time*

$$\tilde{O}\left(m^2 L_{\max}^4 k^4 \log^3(n) \varepsilon^{-4} + qm L_{\max}^2 k^2 \log^2(n) \varepsilon^{-2}\right)$$

with high probability.

1.1.3 Security Games While slightly less intuitive, games modeling security also fit within the resource allocation paradigm. *Security games* are a basic two-player setting modeling the behavior of a limited-resource player defending n targets, and an adversarial attacker. Each target in the game has a cost to defend, and a “ k -resource” defender may choose a (possibly restricted) k -set to defend. Similarly, each target has a cost to attack, and the attacker chooses a single element, receiving a reward depending on whether or not the selected target was defended by the opponent. Depending on the cost/reward structure, security games model several real-world scenarios, ranging from allocating defensive resources at military checkpoints to choosing a path to transmit critical resources (in the latter the attacker actually *wins* if they attack a ‘defended’ node). Indeed, security games have actually seen significant use in critical real-life infrastructure such as checkpoint placements at LAX and US Coast Guard and Federal Air Marshal Service patrol schedules [Tam11].

Given their practical importance, it is no surprise equilibrium computation is well-studied in the security game setting [Tam11, ADH⁺19, Sze17, BB19], and polynomial time algorithms are known in several settings, notably including when allocation constraints are given by matroid bases [Sze17, BB19]. Unfortunately, as is the case in previous work on Blotto, known algorithms are not practically useful and have large polynomial factors in the number of targets. We take a major step toward resolving this issue by showing δ -RMW can be implemented in time *polylogarithmic* in n , an exponential improvement over prior techniques [Tam11, Sze17, BB19].

THEOREM 1.3. (SECURITY WITHOUT REGRET (INFORMAL COROLLARY 3.3)) *Let \mathcal{I} a security game over the bases of a rank- k matroid over n targets with q distinct defender costs. Then δ -RMW can be implemented for T rounds in time*

$$\tilde{O}(kT \log(n) \log(\delta^{-1})(q + T)),$$

and is no-regret.

COROLLARY 1.3. (EQUILIBRIUM COMPUTATION FOR SECURITY (INFORMAL COROLLARY 3.4)) *Let \mathcal{I} be a security game over the bases of a rank- k matroid over n targets with q distinct attacker and defender costs. Then it is possible to compute an ε -CCE in time*

$$\tilde{O}\left(L_{\max}^4 k^3 \log^3(n) \varepsilon^{-4} + qL_{\max}^2 k^2 \log^2(n) \varepsilon^{-2}\right).$$

If the game is zero-sum, the resulting strategy is ε -Nash.

²Matroid bases can be thought of as a generalization of the combinatorial properties enjoyed by spanning trees, see Section 3.1 for details.

1.1.4 Dueling Games Finally, we make a slight departure from the resource allocation framework to consider the popular class of *dueling games* studied in [IKL⁺11, ADH⁺19]. Dueling games model competitive optimization between two players over a randomized set of events. We will focus our attention on one of the simplest dueling games called *ranking duel* (also known as the ‘Search Engine game’) where two players compete to rank n elements over a known distribution μ , and win a round if they rank $x \sim \mu$ higher than the opponent. This classically models the problem of search engines competing to optimize a page ranking given a known distribution over searches.

Equilibrium computation is well-studied in dueling games, and algorithms are known in a few settings via a mix of bilinear embedding techniques and reduction to non-competitive optimization [IKL⁺11, ADH⁺19]. As in previous settings, however, the algorithms are too slow to be of practical use. In contrast, we focus our attention only on the basic ranking duel, but give both a faster algorithm and a novel no-regret guarantee over the original space.

THEOREM 1.4. (DUELING WITHOUT REGRET (INFORMAL COROLLARY 3.5)) *Let \mathcal{I} be an instance of ranking duel. Then δ -RMW can be implemented for T rounds in time*

$$\tilde{O}(T^2 n^7 \log(\delta^{-1})),$$

and is no-regret.

COROLLARY 1.4. (EQUILIBRIUM COMPUTATION FOR RANKING DUEL (INFORMAL COROLLARY 3.6)) *Let \mathcal{I} be an instance of ranking duel. Then it is possible to compute an ε -CCE in time*

$$\tilde{O}(n^9 \varepsilon^{-4}).$$

If the game is zero-sum, the resulting strategy is ε -Nash.

While a running time of $O(n^9)$ can hardly be claimed as practical, the broader technique used in this result has the eventual possibility of running in near-linear time. We discuss this further in Section 1.3.

1.2 Techniques At their core, our results all stem from the ability to approximately sample distributions arising from randomized weighted majority on various linear hypergraph games. Recall that RWM maintains a mixed strategy, which we denote as the RWM distribution, whose probabilities are proportional to their (exponentiated) total historical loss (negative reward):

$$(1.2) \quad \forall x \in A_i : \Pr(x) \propto \beta^{\ell^T(x)}$$

where $\ell^T(x)$ is the total loss experienced by pure strategy x up to round T . As discussed earlier in the section, we typically cannot hope to maintain this distribution explicitly, but it may still be possible to sample from it in polylogarithmic time. Furthermore, while sampling such a distribution exactly is a challenging task (and very few such algorithms are known), *approximate* sampling is perfectly sufficient in our setting. Indeed, our approximate variant δ -RWM satisfies essentially the same guarantees as RWM itself.

LEMMA 1.2. (δ -RWM IS NO-REGRET (LEMMA 2.1)) *δ -RWM over N actions has*

$$\text{Reg}(T) \leq O\left(L_{\max} \sqrt{T \log N} + \delta L_{\max} T\right)$$

expected regret, where L_{\max} is the maximum loss experienced by any action.

As discussed at the start of Section 1, no-regret algorithms like RWM are classically used to compute equilibria of the base game by simulating repeated play across all players. While much of the current literature centers around the Hedge algorithm that ‘plays’ an entire mixed strategy in each round, classical (and therefore approximate) RWM still leads to equilibrium computation with high probability, by the classic result that no-regret implies equilibrium computation [FS99].

LEMMA 1.3. (APPROXIMATE RWM \rightarrow EQUILIBRIA (INFORMAL COROLLARY 2.1)) *Let \mathcal{I} be an m -player game where each player has at most N strategies. Let $\{(x_1^{(t)}, \dots, x_m^{(t)})\}_{t=1}^T$ be the strategies arising from T rounds of δ -approximate RWM. There exist universal constants $C > 0$ such that for $T = C \cdot L_{\max}^2 \varepsilon^{-2} \cdot \log(N)$ rounds, and approximation parameter $\delta \leq \varepsilon / (CL_{\max})$, these strategies constitute an ε -CCE with high probability (Nash if the game is two-player zero-sum).*

As a result, efficient no-regret learning and equilibrium computation truly reduces to the existence of an efficient approximate sampling scheme for distributions arising in the execution of (δ -approximate) RWM. Of course, this is easier said than done. While approximate sampling is easier than its exact variant, it is still a challenging problem, even over structured domains. Using a mixture of novel sampling techniques and reductions to known methods in the literature, we show it is indeed possible to efficiently sample from RWM across a wide variety of structured games. Our strategies fall into two main paradigms: dynamic programming (DP), and Monte Carlo Markov Chains (MCMC).

1.2.1 Sampling via Dynamic Programming We start with the former: sampling in basic resource allocation settings via dynamic programming. At its most general, resource allocation problems are played over (possibly constrained) fixed size partitions of n . The discrete Colonel Blotto game on n troops and k battlefields is the simplest example of this problem, where the strategy space corresponds to the set of all k -size (ordered) partitions of n (i.e. assignments x_1, \dots, x_k such that $\sum x_i = n$). In this section, we will focus only on the Colonel Blotto problem—general resource allocation follows from very similar arguments (see Section 4 for details).

Our goal is now to design an algorithm for approximately sampling distributions over strategies of the Colonel Blotto game that arise from RWM. In this setting, it will actually be easier to solve an equivalent problem, computing the normalizing factor of Equation (1.2), otherwise known as the *partition function*:

$$f_k(n) = \sum_{x_1 + \dots + x_k = n} \beta^{\ell(x)} = \sum_{x_1 + \dots + x_k = n} \prod_{h=1}^k \beta^{\ell_h(x_h)},$$

where $\ell_h(x_h)$ is the historical losses from the h -th battlefield over previous rounds of play if one were to place x_h soldiers on that battlefield. Notice that once we know the value of $f_{k'}(n')$ for all $k' \leq k$ and $n' \leq n$, it is actually possible to exactly sample from Equation (1.2) (and therefore implement RWM). In particular, one does this simply by sequentially sampling the number of troops to put in each battlefield conditional on prior choices in the following manner:

$$\begin{aligned} \Pr[x_1 = y] &\propto \beta^{\ell_1(y)} \cdot f_{k-1}(n - y) \text{ for the first battlefield,} \\ \Pr[x_{h+1} = y | x_{1..h}] &\propto \beta^{\ell_{h+1}(y)} \cdot f_{k-h-1} \left(n - \left(\sum_{j=1}^h x_j \right) - y \right) \text{ for the remaining battlefields.} \end{aligned}$$

One can easily check the joint distribution arising from this procedure is exactly the RWM distribution.

Thus we have reduced our problem to computing the partition functions $f_{k'}(n')$. This can be done by a simple dynamic programming argument, and in particular by noticing that:

$$(1.3) \quad f_{k'}(n') = \sum_{i=0}^{n'} \beta^{\ell_{k'}(i)} \cdot f_{k'-1}(n' - i).$$

Since filling each entry $f_{k'}(n')$ takes time at most $O(n)$ given that f_{h-1} is pre-computed, we can fill the entire DP table in time $O(n^2k)$.³

While this procedure already gives the first no-regret learning algorithm for Blotto in the adversarial setting and by far the fastest known equilibrium computation, one can still hope to do much better. Indeed, it is known that there exist ε -Nash Equilibrium with support that is *logarithmic* in the size of the state-space [LY94], so there is hope in building a *polylogarithmic* time algorithm (equivalently, a polynomial time algorithm in the description complexity of the equilibria). We show this is indeed possible by building an approximation scheme for the above DP. The key is to observe that the partition functions $f_{k'}(n')$ are bounded and monotonic. Roughly speaking, this means f can be approximated within multiplicative $(1 \pm \varepsilon)$ factors by a piece-wise function with only $\text{poly}(k \log(n)/\varepsilon)$ pieces (which is polylogarithmic in the size of the state space).

By carefully computing and maintaining approximate versions of the partition function, we can run a modified variant of the same dynamic program that computes approximations for all nk partition functions $f_{k'}(n')$ (despite their

³We note that this can actually be improved to near-linear in n using the Fast Fourier Transform.

sizes, these can indeed be presented in only $\text{poly}(k \log(n))$ bits due to being piece-wise). Once we have approximately computed the partition functions, it is easy to show that a similar sampling scheme as discussed for the exact case gives an efficient approximate sampling scheme for RWM running in $\text{poly}(k \log(n)/\varepsilon)$ time. Combined with Lemma 1.2 and Lemma 1.3, this results in the first polylogarithmic time algorithm for no-regret learning and (approximate) equilibrium computation for Colonel Blotto (Theorem 1.1 and Corollary 1.1), as well as for a number of related resource allocation variants discussed later in the paper (e.g. multi-resource Blotto and Dice games).

1.2.2 MCMC-methods While dynamic programming is a powerful algorithmic method for structured computation, there are many combinatorial settings common to games we cannot hope to handle via such techniques. Building an analogous exact-counting based DP for games over bipartite matchings or matroids, for instance, would give efficient algorithms for classical #P-hard problems such as the permanent and counting matroid bases [CPV95]. On the other hand, we do actually know of approximation algorithms for these problems based on a powerful tool called MCMC-sampling [JSV04, ALOV19].

MCMC-sampling is an elegant method for approximately sampling from a distribution μ with exponential size support usually traced back to Ulam and Von Neumann in the 1940s (see e.g. [Eck87]). The idea is simple. Imagine we can construct a Markov chain (random process) M satisfying the following three conditions:

1. The stationary distribution of M is μ
2. A single step of M can be implemented efficiently
3. M converges quickly to its stationary distribution.

Approximate sampling would then simply boil down to finding a starting configuration and running the chain until it is within δ of stationary (this typically takes around $O(\log(N/\delta))$ samples for a good chain).

Unsurprisingly, while MCMC-sampling itself is a simple technique, the design and analysis of Markov chains is a difficult task, and general recipes for their construction are known in very few scenarios. One particularly well-studied setting in the literature that arises from simulation problems in statistical physics are *external fields*. Given a hypergraph $\Omega \subset \binom{[n]}{k}$, the distribution arising from external field $w \in \mathbb{R}_+^n$ simply assigns each k -set a probability proportional to the product of its fields:

$$\Omega^w(s) \propto \prod_{v \in s} w(v).$$

External fields often correspond to particularly natural problems, and are well-studied in the literature. In a recent breakthrough series of works, for instance, it was shown that approximate sampling under external fields is possible whenever the underlying state-space is a good enough high dimensional expander [KO20, AL20, ALOV19, AJK⁺21].⁴

This is particularly relevant to our setting since it is a simple observation that the distributions arising from RWM on a linear hypergraph game are exactly given by the application of an external field over the action space.

OBSERVATION 1.1. (RWM \rightarrow EXTERNAL FIELDS (INFORMAL OBSERVATION 3.1)) *Let $\mathcal{I} = \{(A_i, R_i)_{i=1}^m\}$ be an m -player linear hypergraph game. Then for any A_i and any round of play, Player i 's RWM distribution can be written as the application of an external field w to A_i .*

As a result, no-regret learning and equilibrium computation are possible in any linear hypergraph game whose state space can be sampled under arbitrary external fields. As a short detour, it is worth noting that the result in the previous section can be phrased as a slight refinement of this statement. At a technical level, our resource allocation algorithm simply corresponds to an efficient approximate sampling scheme for fixed-size partitions of n under *monotonic* external fields (corresponding to the fact that assigning more troops to a battlefield always results in at least as many victories).

Many well-studied games in the literature have state spaces where efficient approximate sampling schemes under external fields exist. In this work we focus mostly on games played on matroids (e.g. matroid congestion, security), and dueling games arising from bipartite matchings such as ranking duel. Both settings have well-known sampling schemes over external fields [ALOV19, CGM19, JSV04], which leads to our results for Congestion, Security, and Dueling games (Theorems 1.2, 1.3, 1.4 and Corollaries 1.2, 1.3, 1.4 respectively).

⁴More formally, when the space satisfies a property known as ‘fractional log-concavity.’

1.3 Discussion In this work, we present two potential paradigms for learning in games via approximate sampling. In this section we touch on the pros and cons of each method, their likelihood to generalize beyond the settings considered in this work, and natural open problems.

1.3.1 Dynamic Programming vs MCMC-sampling Broadly speaking, the DP and MCMC approaches we develop in this work seem to be largely incomparable. Dynamic programming works well in relatively unconstrained resource allocation problems, where recursive structure allows for inductive computation of the partition function. On the other hand, typical MCMC methods (which are usually *local* in nature) actually fail drastically in this sort of setting due to the need for global coordination. One natural example of this issue appears in the Colonel Blotto game. Imagine a scenario where Colonel A has k more troops than Colonel B , then there always exists a configuration where A wins every battle by assigning one more troop in each battlefield than B . Finding this sort of optimum, however, requires coordinated planning across battlefields. Typical MCMC methods like Glauber dynamics (see Section 3) only look at a few battlefields at a time, and therefore struggle to converge to such solutions. Simulations confirm this intuition—even for small n and k Glauber dynamics seem to exhibit very poor mixing on distributions arising in RWM.

On the other hand, as we mentioned in the previous section, our dynamic programming approach has a significant issue in any setting with non-trivial combinatorial structure. In particular, because the underlying method relies on exactly computing the partition function, constructing any such method for a problem like matroid games is #P-hard. On the other hand, local chains such as the Glauber Dynamics mix extremely fast in these settings, providing near-optimal algorithms.

Of course, neither of these arguments rules out either approach. It is possible there exist successful MCMC methods for Blotto that are more *global* in nature—indeed the main insight leading to the resolution of approximate permanent was exactly such a Markov chain that avoided these issues [JSV04]. On the other hand, there may exist DP-based approaches that do not go through computing the partition function. Understanding in which scenarios these two or other potential sampling methods may apply remains an interesting and important open problem if we wish to extend efficient learnability in games beyond the few structured settings considered in this work.

1.3.2 Further Open Problems Our work gives the first no-regret learning guarantees and polylogarithmic equilibrium computation for several well-studied settings in game theory, but there is still much to be done. Perhaps the most obvious open directions involve improving the computational efficiency (and therefore practicality) of our algorithms. The polynomial dependencies of our algorithm would be universally improved if we can show that the δ -approximate *optimistic* variant of RWM achieves $\tilde{O}(1)$ regret in games like its deterministic counter-part, Optimistic Hedge [DFG21], even for polynomial approximation δ .

QUESTION 1.1. (OPTIMISTIC-RWM) *Does δ -RWM with weights $\{w_i^{(t)}\}_{i,t}$ and optimistic updates $w_i^{(t+1)} \leftarrow w_i^{(t)} \cdot \beta^{2\ell_i^{(t)} - \ell_i^{(t-1)}}$ achieve $\text{polylog}(T)$ regret in games (even for $\delta = (\text{polylog } N)^{-1}$)?*

These techniques are well known to give a substantial improvement in the exact setting [DFG21], but their analysis is subtle and may be nontrivial to adapt to the δ -approximate sampling variant we need for efficient computation.

Similarly, our algorithm for dueling games (while faster than prior work in the worst-case), is not practical at $O(n^9)$ running time. One interesting question is whether the MCMC-sampling technique can be improved in this setting using the fact that the weights arising from RWM are not arbitrary, but exhibit *monotonic* structure (in the sense that ranking a page higher is always better). This actually corresponds to a well-studied problem in the sampling and geometry of polynomials literature (monotone permanent [BHVW11]), but giving an improved sampling algorithm over the JSV-chain [JSV04] remains an interesting open problem.

QUESTION 1.2. (SAMPLING WITH MONOTONE WEIGHTS) *Can a perfect matching in a complete bipartite graph with n nodes under monotone external fields be sampled in faster than $O(n^7)$ time? In near-linear time?*

There is certainly hope in this direction, as recent years have seen many breakthroughs towards near-optimal MCMC methods, including similar linear time guarantees on problems that once seemed infeasible [ALOV19, CGM19, CLV21].

Another natural direction is to try to strengthen the type of equilibria we compute in multiplayer and general-sum games. Foremost in this direction are the so-called Correlated Equilibria (CE), a substantially stronger notion than CCE

which allows a player to switch strategies even after they receive instructions from the coordinator. It was recently shown that a variant of Optimistic Hedge converges quickly to CE in multiplayer, general-sum games [ADF⁺22]. It is an open question whether an approximate, sampled variant could do the same. We pose this as the following,

QUESTION 1.3. (CORRELATED EQUILIBRIA WITH RWM) *Is there a variant of δ -approximate RWM that converges to CE and remains efficiently samplable? A variant that achieves $\tilde{O}(T^{-1})$ convergence rate?*

Finally, we end with a concrete direction toward answering our original question: *when can one efficiently sample from RWM?* We rely in part of this work on a series of breakthroughs in the approximate sampling and high dimensional expansion literatures [KO20, AL20, ALOV19, AJK⁺21] leading to a sufficient condition called fractional log-concavity [AASV21, AJK⁺21] for sampling hypergraphs under arbitrary external fields (generalizing an earlier result for matroids [ALOV19]). This is in fact a stronger guarantee than we actually need to ensure efficient sampling for RWM. Not only are the fields we study typically additionally structured (e.g. monotonic), but we are also okay with some amount of decay in the mixing time depending (logarithmically) on the field size. Is there a characterization of such objects in terms of geometry of polynomials or high dimensional expansion?

QUESTION 1.4. *Is there a general condition on hypergraphs (e.g. in terms of high dimensional expansion, geometry of polynomials) that allows for approximate sampling under external fields with polylogarithmic dependence on the worst field size? What about under structural constraints (e.g. monotonicity)?*

1.4 Further Related Work

1.4.1 No-regret learning with structured loss Online learning over exponentially large classes with structured losses has been considered previously in other contexts (e.g. [Kle05, CHK17, HKW09, KWK⁺10, ABL14, VLS19]). Much of this work considers the combinatorial bandit setting [CBL12], which typically competes against a non-adaptive adversary, but has restricted information. This work introduces the notion of sampling from the MWU distribution in structured games by constructing a linear embedding and performing MWU explicitly on each dimension (their ComBand algorithm). We note two distinctions with our current work. First, since the ComBand algorithm focuses on the partial information setting, the regret guarantee of ComBand has worse dependence on dimension than MWU. Furthermore, the natural embedding of a linear hypergraph game into binary vectors (which maps a k -set to its corresponding weight k vector in $\{0, 1\}^\Omega$) typically has dimension linear in n , leading to exponentially worse runtime compared to our techniques in all of our settings except ranking duel.

There are two additional works which also consider efficient implementation of RWM [WK08, HKW09], but only for the very special settings of k -sets and permutations (which are generalized by our framework). Also of note is the later work of [KWK⁺10], who built a new hedge-based algorithm for these settings called component-hedge that also gives efficient online learning in a few additional cases (e.g. for spanning trees).

One may also consider a kernelized approach to sampling in RWM [FLLK22]. We note that the kernel approach (which in particular implies the ability to compute the partition function) cannot be applied efficient in many settings (e.g. matroids, where this problem is known to be #P-hard). Indeed, it is worth noting that the ability to exactly compute the partition function implies the ability to exactly sample efficiently (at least in self-reducible settings), so in this sense requiring efficient kernel computation is strictly stronger than our approach. Further, since the standard embedding of our setting into a 0/1-polytope has dimension that is linear in n (or more accurately the number of vertices in the hypergraph), a naive application of a kernelized method leads to bounds that are linear rather than polylogarithmic in n .

1.4.2 Computing equilibria for Colonel Blotto The Colonel Blotto game is one of the most well studied problems in algorithmic game theory—we restrict our attention here to some of the most notable and relevant results. As mentioned previously, the first known algorithm to compute exact Nash equilibria strategies for discrete Colonel Blotto was introduced in [ADH⁺19], who consider games that are asymmetric across battles and across players (allowing different troop capacity and rewards across battles and players). This work remains the only known algorithm for exact equilibrium computation that is polynomial in the number of troops and battlefields, though follow-up work gave a more practical (but potentially exponential time) algorithm [BDD⁺17].

Due to the difficulty of understanding the discrete version, a number of works have also considered Colonel Blotto's continuous relaxation. It should be noted that the equilibria in the continuous version do not apply to the discrete

case [PRG22]. The works by [Gro50, GW50, Las02, Tho18, KR21] consider the case that troop counts are identical (symmetric) for both players. Later, the symmetric case was also studied when Colonels have different values for battles [KR21]. On the other hand, when the troop counts of the two players differ, constructing/computing equilibria becomes more complicated. In [Rob06], the author constructs equilibrium strategies explicitly in the case that the rewards for each battlefield are the same (homogeneous). The authors of [MM15] consider Blotto with heterogeneous rewards and asymmetric troops counts (but with only two battles). In [SLS14], they consider more than two battles but with strict assumptions on the battle weights. More recently, the authors in [PRG22] present an efficient algorithm to compute approximate Nash equilibria in the two-player continuous Colonel Blotto game with asymmetric troop and battle values.

There is also a breadth of work that constructs strategies for approximate and exact equilibrium under constrained parameter settings of Borel's two-player discrete version. Beginning with [Har08], the author constructs optimal strategies explicitly when the troop counts and battle rewards are identical. In [VLS18], the authors give an algorithm to compute equilibria with fixed approximation (decaying with the number of battles). They also give an algorithm to compute the best-response strategy to a given distribution over soldiers in each battlefield using dynamic programming. In [BAEJ20, Tho17], the authors describe equilibria in the symmetric case where the number of soldiers is the same for both players. Moreover, [BAEJ20] introduces the multiplayer variant. In [VL21], the authors construct equilibria under particular conditions for an extension of the Colonel Blotto game that accounts for pre-allocations and resource effectiveness.

2 Preliminaries

All throughout the paper, for integers $a \leq b$ we denote by $[a, b]$ the set $\{a, \dots, b\}$ and shorthand $[n] = [1, n]$. We use the notations $\widetilde{O}(f)$ to hide polylogarithmic dependencies on the argument. Given a finite set Ω , we denote by $\Delta(\Omega)$ the (convex) polytope of all distributions defined on Ω . We denote by 2^Ω as the power set of Ω , i.e. the set of all subsets of Ω . Given two finite sets Ω_1, Ω_2 , we denote by $\Omega_1 \times \Omega_2$ as the Cartesian product of the two sets, i.e. $(x, y) \in \Omega_1 \times \Omega_2$ if $x \in \Omega_1$ and $y \in \Omega_2$. We will use $\binom{\Omega}{k}$ to denote all size- k subsets of the ground set Ω . Given two integers $k, n \in \mathbb{Z}^+$, we will use $P_k(n)$ to represent the set of ordered size- k partition of n .

2.1 Game Theory

DEFINITION 2.1. (MULTIPLAYER SIMULTANEOUS GAME) An m -player **Simultaneous Game** is a tuple $\{\{A_i\}_{i=1}^m, \{R_i\}_{i=1}^m\}$ where A_i denotes the finite set of actions available for the i -th player and $R_i : A_1 \times \dots \times A_m \mapsto \mathbb{R}$ denotes the reward function for the i -th player.

Given a set of actions a_1, \dots, a_m , we often write a_{-i} to represent the combined action tuples without a_i , i.e. $(a_1, \dots, a_{i-1}, a_{i+1}, \dots, a_m)$, and $R_i(a_i, a_{-i}) = R_i(a_1, \dots, a_m)$ where we have abused notation in the input ordering to R_i for simplicity of notation.

In a game, a player can choose to play an action, often called a pure strategy, or to draw randomly from a *mixed strategy* given by a probability distribution over the set of available actions.

DEFINITION 2.2. (MIXED STRATEGY) Let $\{\{A_i\}_{i=1}^m, \{R_i\}_{i=1}^m\}$ be an m -player simultaneous game. For the i -th player, the set of mixed strategies are all possible probability distributions over the actions A_i . Let $\mathbf{s}_i \in \Delta(A_i)$ be the mixed strategy chosen by the i -th player. Then, the expected reward received by the i -th player is given by $\mathbb{E}_{a_1 \sim \mathbf{s}_1, \dots, a_m \sim \mathbf{s}_m} [R_i(a_1, \dots, a_i, \dots, a_m)]$.

We will also make use of the following notion of a joint strategy.

DEFINITION 2.3. (JOINT STRATEGY) A joint strategy is a distribution $\sigma \in \Delta(A_1 \times \dots \times A_m)$. If players were to participate in a joint strategy, then a central coordinator samples an action tuple $a = (a_1, \dots, a_m) \sim \sigma$, and each player then plays the action a_i correspondingly. As a result, the expected reward of the i -th player is given by $\mathbb{E}_{a \sim \sigma} R_i(a)$.

For a set of actions $a_i^{(t)}$ for $i \in [m]$, and $t \in [T]$, we will often write $\frac{1}{T} \sum_{t=1}^T a_i^{(t)}$ as the mixed strategy of player i such that action $a_i^{(t)}$ is played with probability $1/T$, and $\frac{1}{T} \sum_{t=1}^T a_1^{(t)} \otimes \dots \otimes a_m^{(t)}$ as the joint mixed strategy such that the action tuple $(a_1^{(t)}, \dots, a_m^{(t)})$ is played with probability $1/T$.

It is well known that if all players play a game optimally, the resulting strategy tuples compose of a *Nash Equilibrium* of the game.

DEFINITION 2.4. (NASH EQUILIBRIUM) *In an m -player game $\{\{A_i\}_{i=1}^m, \{R_i\}_{i=1}^m\}$, a tuple of mixed strategies (s_1, \dots, s_m) composes an ε -Nash Equilibria (ε -NE) if for all $i \in [m]$ it satisfies:*

$$\mathbb{E}_{s_1, \dots, s_m} [R_i(s_1, \dots, s_i, \dots, s_m)] \geq \sup_{s' \in \Delta(A_i)} \mathbb{E}_{s_1, \dots, s_m} [R_i(s_1, \dots, s', \dots, s_m)] - \varepsilon,$$

where the mixed strategies s_i for $i \in [m]$ are mutually independent.

In multiplayer and general-sum games, computing Nash equilibria is challenging. In fact, this problem is known to be complete for PPD [DGP06], a complexity class containing many other computationally hard problems. A standard and arguably more realistic goal is to find the so-called Coarse Correlated Equilibria (CCEs) of the multi-player game, a relaxation of Nash Equilibrium introduced by Aumann [Aum74].

In a CCE, all players together sample from a joint mixed strategy (in contrast to NE where players *independently* sample from their own mixed strategy). Although a player i cannot benefit from switching to any single action s'_i before the joint strategy is sampled, once a strategy s_i is sampled from a CCE distribution (becoming known to each player), a player may improve their outcome by deviating (using the fact that her strategy is correlated with other players'). Thus, CCE apply to situations where a player must commit to their strategy up front and are unable to deviate after sampling.

DEFINITION 2.5. *Let $\mathcal{I} = \{(A_i, R_i)_{i=1}^m\}$ be an m -player game. An ε -approximate coarse correlated equilibrium (ε -CCE) is a joint mixed strategy $\sigma \in \Delta(A_1 \times \dots \times A_m)$ that satisfies:*

$$\forall i \in [m], \text{ and actions } a'_i \in A_i : \mathbb{E}_{a \sim \sigma} R_i(a) \geq \mathbb{E}_{a \sim \sigma} R_i(a'_i, a_{-i}) - \varepsilon.$$

2.2 Linear Hypergraph Game Given a ground set of vertices Ω , a hypergraph H is a collection of subsets of Ω called hyperedges. If all hyperedges of H are in $\binom{\Omega}{k}$, the graph is called k -uniform. In this work we study a special class of games whose reward functions can be 'decomposed' based on the underlying structure of the game's action space. More formally, we consider games played over k -uniform hypergraphs whose rewards *decompose linearly* over vertices. We denote this class of games as Linear Hypergraph Games.

DEFINITION 2.6. (LINEAR HYPERGRAPH GAMES) *Let $\mathcal{I} = \{(A_i, R_i)_{i=1}^m\}$ be an m -player game. We call \mathcal{I} a **linear hypergraph game** if for all $i \in [m]$ there is a groundset Ω_i and parameter $k_i \in \mathbb{N}$ such that $A_i \subset \binom{\Omega_i}{k_i}$ and a 'vertex-wise' reward function $R_i^{\Omega_i} : \Omega_i \times A_{-i}$ such that for all $t_1, \dots, t_n \in A_1 \times \dots \times A_n$*

$$R_i(t_i, t_{-i}) = \sum_{v \in t_i} R_i^{\Omega_i}(v, t_{-i}).$$

In other words, each element v in the ground set Ω_i has a certain reward with respect to any choice of the opponents, and the reward of a k -set is simply the sum of its individual rewards. Many important games that are well-studied in the game theory literature falls under this category, e.g. Colonel Blotto Games, Security Games, Congestion Games, Dueling Games, etc. In fact, it should be noted a similar notion has been studied in the online learning setting in [KWK⁺10], who develop an efficient no-regret algorithm called Component Hedge for linear losses over basic structures such as the complete hypergraph, truncated permutations, and spanning trees.

2.3 No-Regret Learning in Games We consider the framework of No-Regret Learning in Games (see [CBL06, DFG21] and references therein). In this framework, a game is iterated with one or more players implementing a no-regret learning algorithm to adaptively choose strategies. At the t -th round of the game, each player selects a mixed strategy $\mathbf{s}_i^{(t)}$, and samples the action $a_i^{(t)} \sim \mathbf{s}_i^{(t)}$, where the choice of $\mathbf{s}_i^{(t)}$ depends only on $a_j^{(t')}$ for $j \in [m]$ and $t' < t$.

The goal for each player is to optimize her regret, defined as the following.

DEFINITION 2.7. (REGRET) *At the T -th round of the game, the regret for the i -th player is defined as*

$$\text{Reg}_{T,i} := \max_{a^* \in A_i} \sum_{t=1}^T R_i(a_1^{(t)}, \dots, a_{i-1}^{(t)}, a^*, a_{i+1}^{(t)}, \dots, a_m^{(t)}) - \sum_{t=1}^T R_i(a_1^{(t)}, \dots, a_m^{(t)}).$$

It is classical result that if all players follow no-regret learning strategies, the overall dynamics quickly converge to a Nash or Coarse-Correlated Equilibria (CCEs) of the game (see e.g. [FS99, CBL06, OvS13]).

THEOREM 2.1. (NO-REGRET IMPLIES EQUILIBRIUM COMPUTATION [FS99]) *Suppose m players are playing under the No-Regret Learning in Games framework. Let $\sigma^* := \frac{1}{T} \sum_{t=1}^T a_1^{(t)} \otimes \cdots \otimes a_m^{(t)}$ be the average mixed joint strategies played by the i -th player over T rounds. Then, σ^* forms an $T^{-1} \max(\text{Reg}_{T,1}, \dots, \text{Reg}_{T,m})$ -approximate CCE of the game, where $R_i^{(T)}$ is the regret for the i -th player at the T -th round. When $m = 2$ and the game is zero-sum, the mixed strategies $(\frac{1}{T} \sum_{t=1}^T a_1^{(t)}, \frac{1}{T} \sum_{t=1}^T a_2^{(t)})$ constitute a $T^{-1} \max(\text{Reg}_{T,1}, \text{Reg}_{T,2})$ -approximate Nash Equilibrium.*

2.4 Randomized Weighted Majority Algorithm As in online learning, no-regret learning in games studies the regret of a player against the opponents' strategies in repeated play with respect to the best single strategy in hindsight.

One of the most frequently used tools in no-regret learning is the randomized weighted majority (RWM) algorithm. For player $i \in [m]$, RWM maintains the mixed strategies from $\Delta(A_i)$ as follows: at the first round, it chooses uniformly among the actions.

At the $(T + 1)$ -st round, a cumulative reward is computed for each action $x \in A_i$

$$r^{(T+1)}(x) = \sum_{t=1}^T R_i(s_1^{(t)}, \dots, s_{i-1}^{(t)}, x, s_{i+1}^{(t)}, \dots, s_m^{(t)}),$$

and RWM chooses the mixed strategy $RM^{(T+1)}(\beta)$ (which we refer to as the RWM distribution) such that

$$(2.4) \quad \Pr[RM^{(T+1)}(\beta) = x] \propto \beta^{-r^{(T+1)}(x)}.$$

It is well known that if any player samples according to $RM^{(T+1)}(\beta)$ in each round, her expected regret will be bounded in the worst case by $O_{\beta,N}(\sqrt{T})$.

In games whose action sets are exponentially large, exactly sampling from the the RWM distribution may be intractable in relevant cases. Nonetheless, we show that similar regret bounds hold even when one *approximately* samples the RWM distributions in each round (the proof is given in Appendix).

DEFINITION 2.8. (APPROXIMATE SAMPLING) *We say a randomized algorithm \mathcal{A} with output space Ω δ -approximately samples a distribution μ over Ω if the output of \mathcal{A} is δ -close to μ in TV-distance.*

We call any strategy that δ -approximately samples from $RM^{(t)}(\beta)$ in each of T rounds of repeated play δ -approximate RWM, and denote this class of algorithms by $\delta\text{-RWM}_{\beta}^T$. It is not hard to show that δ -RWM has near-optimal regret in the adversarial setting (see Appendix A).

LEMMA 2.1. (δ -RWM IS NO-REGRET) *Let \mathcal{I} be an m -player game with at most N actions and L_{\max} reward. If the i -th player follows $\delta\text{-RWM}_{\beta}^T$ in T rounds of play with learning rate $\beta = 1 - \sqrt{\log(N)}/T$ and approximation factor $\delta \leq \sqrt{\log(N)/T}$,⁵ then for any $\eta > 0$ they experience regret at most*

$$\text{Reg}_{T,i} \leq O\left(L_{\max} \sqrt{T} \left(\sqrt{\log N} + \sqrt{\log(1/\eta)} \right)\right)$$

with probability at least $1 - \eta$.

As an immediate corollary, for any game, if we can approximate sample from the RWM distribution efficiently, we immediately get an efficient no-regret learner. In addition, connecting it with Theorem 2.1, we also obtain the following corollary for equilibrium computation with RWM.

COROLLARY 2.1. (EQUILIBRIUM COMPUTATION WITH δ -RWM) *Let $\mathcal{I} = \{(A_i, R_i)_{i=1}^m\}$ be an m -player where $|A_i| \leq N$ for each $i \in [m]$ and reward bounded by L_{\max} . Suppose the game is played repeatedly for T rounds.*

⁵We are assuming $T \gg \log N$. Otherwise, the regret bound becomes $L_{\max}T$, which can be achieved by any arbitrary sequence of choices.

If there is an algorithm which can perform δ -RWM $_{\beta}^T$ for each player i in time $f_I(T, \delta)$, then there exists an algorithm which computes an ε -CCE of \mathcal{I} (Nash if the game is 2-player zero-sum) with probability at least $1 - \eta$. Moreover, the algorithm runs in time

$$O\left(m \cdot f_I\left(C \cdot L_{\max} \varepsilon^{-2} \log(Nm/\eta), \frac{\varepsilon}{C \cdot L_{\max}}\right)\right).$$

for some universal constant $C > 0$.

In Sections 3 and 4, we develop two different types of methods of approximate sampling from the RWM distribution in many well-studied games and discuss their implications. Before moving on, however, it is convenient to briefly discuss one computational consideration that frequently occurs in efficient implementation of RWM. In particular, it will often be the case that our algorithm needs to deal with piece-wise constant functions that map from $[0, n]$ to \mathbb{R} (e.g. reward functions in Blotto for each battlefield are 2-piecewise in this sense). To represent such functions, we will use the following data structure that we refer as a *succinct representation*.

DEFINITION 2.9. (SUCCINCT DESCRIPTIONS OF PIECEWISE CONSTANT FUNCTIONS) Let $f : [0, n] \mapsto \mathbb{R}$ be a q -piecewise constant function. The succinct description of f , denoted as D_f , consists of q tuples of the form $(a_i, b_i, y_i) \in (\mathbb{Z}^+, \mathbb{Z}^+, \mathbb{R})$ such that for all $x \in [a_i, b_i]$, $f(x) = y_i$ and the intervals $\{[a_1, b_1] \cdots [a_q, b_q]\}$ partition $[0, n]$.

We will often write $|D_f|$ to denote the number of intervals contained in the succinct description. Finally, note that assuming access to succinct descriptions does not lose much generality, as given query access to a standard representation for the monotonic piece-wise function in question (e.g. in the RAM model), it is easy to build a succinct description in time $q \log(n)$ by binary search.

3 Playing Games via MCMC-Sampling

In this section, we develop the connections between linear hypergraph games, the RWM distribution, and efficient sampling. In doing so, we unlock access to powerful tools from the sampling literature for the first time in the context of games. This allows for a number of immediate applications including the first no-regret algorithms for well-studied settings such as matroids.

With this in mind, let's first recall the basic framework of Monte Carlo Markov Chains: a powerful tool for approximately sampling from large spaces like the RWM distribution. More formally, consider the following problem: given a distribution π over a large state space A , we'd like to *approximately* sample a state from π in $\text{polylog}(|A|)$ time. MCMC-sampling is an elegant approach to this problem in which one defines a Markov chain M on A satisfying the following three conditions:

1. The stationary distribution of M is π
2. A single step of M can be implemented efficiently
3. M converges quickly to its stationary distribution.

As long as these three conditions hold, it is possible to efficiently sample from π up to any desired accuracy simply by running the Markov chain from any starting position a few steps and outputting the resulting state. More formally, recall that the *mixing time* of a Markov chain M is the number of steps until the resulting distribution is close in TV-distance to π :

DEFINITION 3.1. (MIXING TIME) The mixing time of a Markov chain M is the worst-case number of steps until the total variation distance of M is close to its stationary measure:

$$T(M, \delta) := \min_{t \in \mathbb{N}} : \forall \pi_s, TV(M^t \pi_s, \pi) \leq \delta.$$

Thus one only needs to run the chain $T(M, \delta)$ times (from any starting position) in order to δ -approximately sample from π .

While MCMC-sampling is a promising approach, designing efficient Markov chains is typically a challenging task. However, in structured settings such as linear hypergraph games, the distributions arising from RWM seem to be more conducive to the approach. In fact, they correspond to well-studied structure in the approximate sampling literature called *external fields*.

DEFINITION 3.2. (EXTERNAL FIELD) *Let π be a distribution over a k -uniform hypergraph $H \subset \binom{\Omega}{k}$. The distribution given by π ‘under external field w ’ for $w \in \mathbb{R}_+^\Omega$ has measure proportional to the product of w across each k -set:*

$$\pi^w(s) \propto \pi(s) \prod_{v \in s} w(v).$$

When π is uniform over H , we often just write H^w for π^w .

It is a simple observation that the distribution arising from RWM (or indeed any reasonable variant) is exactly given by the application of an external field to the state space.

OBSERVATION 3.1. (RWM \rightarrow EXTERNAL FIELDS) *Let $\mathcal{I} = \{(A_i, R_i)_{i=1}^m\}$ be an m -player linear hypergraph game. Then for any $i \in [m]$, $RM^{(T)}(\beta)$ can be written as the application of an external field w to A_i such that A_i^w has minimum probability at most $\frac{\beta^{-2L_{\max}T}}{|A_i|}$.*

Proof. Assume $i = 1$ without loss of generality (for simplicity of notation). Recall that RWM operates at round $T + 1$ by exponentiating the total loss over the previous T rounds:

$$\Pr\left[RM^{(T)}(\beta) = a\right] \propto \beta^{-\sum_{j=1}^T R_1(a, s^{(j)})}$$

where $s^{(1)}, \dots, s^{(T)} \in A_{-1}$ are the historical strategies played by players $\{2, \dots, n\}$ in rounds one through T . For simplicity of notation, let $\ell^{(T)}(a) := -\sum_{t=1}^T R(a, s^{(t)})$ be the total loss. We can similarly define this quantity for any element of the ground set $v \in \Omega$ as:

$$\ell^{(T)}(v) = -\sum_{t=1}^T R_1^\Omega(v, s^{(t)}).$$

Switching the summations, linearity promises we can express $\ell_T(s)$ as a sum over $\ell_T(v)$:

$$\ell_T(a) = \sum_{v \in a} \ell_T(v).$$

As a result, the RWM distribution is proportional to the product of (exponentiated) total loss for each vertex:

$$\Pr(a) \propto \prod_{v \in a} \beta^{\ell^{(T)}(v)}.$$

This is exactly A_1 under the external field $w \in \mathbb{R}_+^\Omega$ where $w(v) = \beta^{\ell^{(T)}(v)}$. Since the distribution started uniformly over A_1 and is update by at most $\beta^{L_{\max}}$ in each step, the minimum probability is at worst $\frac{\beta^{-2L_{\max}}}{|A_1|}$. \square

As an immediate corollary, we get an efficient no-regret algorithm for any linear hypergraph game whose state space can be approximately sampled under arbitrary external fields, a well-studied problem in approximate sampling. In the remainder of the section, we’ll show how such results lead to new efficient no-regret learning algorithms for many well-studied settings in game theory.

3.1 Glauber Dynamics and Fractionally Log-Concave Games The past few years have seen major advances in approximate sampling various combinatorial objects under external fields [ALOV19, ALO20, AASV21, AJK⁺21]. The recent breakthroughs have largely been driven by new analysis techniques for a simple local Markov chain arising from the study of Ising models in statistical physics called the *Glauber Dynamics*. Starting from a state $\sigma \in H \subset \binom{\Omega}{k}$, the (single-site) Glauber Dynamics for a distribution π over A_i are given by the following two-stage procedure:

1. “Down-Step:” Remove a vertex v uniformly at random from σ .
2. “Up-Step:” Sample from π conditional on $\sigma \setminus v$.

It is not hard to show that π is the stationary distribution of this process. The first major breakthrough towards rapid mixing of Glauber Dynamics was due to Anari, Liu, Oveis-Gharan, and Vinzant [ALOV19], who used tools developed in the high dimensional expansion literature [KO20] to prove rapid-mixing of Glauber Dynamics on a broad class of combinatorial objects called *matroids*.

DEFINITION 3.3. (MATROIDS) *Let Ω be a ground set and \mathcal{J} a family of subsets of Ω . (Ω, \mathcal{J}) is called a matroid if it satisfies*

1. **Non-emptiness:** \mathcal{J} contains at least one subset.
2. **Downward-closure:** For all $S \in \mathcal{J}$ and $S' \subset S$, $S' \in \mathcal{J}$
3. **Exchange-property:** For all $S, S' \in \mathcal{J}$ s.t. $|S| > |S'|$, there exists $x \in S$ s.t. $S' \cup x \in \mathcal{J}$.

An element $S \in \mathcal{J}$ is a **basis** if it is maximal, and the **rank** of the matroid (denoted $r(\mathcal{J})$) is the size of its largest basis.

Note that the bases of a matroid make up an $r(\mathcal{J})$ -uniform hypergraph over vertex set Ω . These objects are perhaps best thought of as generalizing the combinatorial structure seen in spanning trees (which form the bases of a ‘graphic’ matroid). ALOV’s [ALOV19] major breakthrough was to prove rapid mixing of Glauber Dynamics on matroid bases, a problem known as the Mihail-Vazirani Conjecture (this result was later optimized by Cryan, Guo, and Mousa [CGM19]). Since matroids maintain their structure under external fields (see e.g. [AJK⁺21]), this leads to the following MCMC-algorithm for sampling matroid bases under arbitrary external fields.

THEOREM 3.1. (GLAUBER DYNAMICS ON MATROIDS [AJK⁺21, THEOREM 5]) *Let H be the set of bases of a rank- k matroid (Ω, \mathcal{J}) . Let $w \in \mathbb{R}_+^\Omega$ be any external field. Then, the single-step Glauber Dynamics on H^w has mixing time*

$$T(GD, \delta) \leq O\left(k \log\left(\frac{\log(|\Omega|/w_*)}{\delta}\right)\right).$$

A substantial amount of progress has been made since ALOV and CGM’s works. In fact, recently Anari, Jain, Koehler, Pham, and Vuong [AJK⁺21] introduced an even more general class of hypergraphs that can be sampled under arbitrary external fields called fractionally log-concave hypergraphs.⁶ All of our results extend to this setting, but to our knowledge matroids already capture most settings of interest in game theory so we focus just on this case for concreteness.

Since matroids are typically exponential size in their rank, we will need implicit access in order to build efficient algorithms. This is typically done through various types of oracle access to the independent sets. For simplicity of presentation, we will assume access to a *contraction oracle*, a standard operation on matroids that restricts the object to independent sets containing some fixed $S \in \mathcal{J}$.

DEFINITION 3.4. (CONTRACTION ORACLE) *Let (Ω, \mathcal{J}) be a rank- k matroid. A rank- r contraction oracle inputs an independent set $S \in \mathcal{J}$ of size r , and outputs (query access to) the contracted matroid (Ω, \mathcal{J}_S) , where*

$$\mathcal{J}_S = \{T : T \cup S \in \mathcal{J}\}.$$

⁶Formally this requires a slight generalization known as the q -step Glauber Dynamics.

Crucially we will only use rank- $(k-1)$ contraction oracles on matroids with rank k (and thus drop rank from the notation below). This can always be implemented in $|\Omega|$ applications of a standard independence oracle (which decides given $S \subseteq \Omega$ whether $S \in \mathcal{J}$), but can often be implemented much more efficiently. For instance, it is easy to see in the case of uniform (unconstrained) matroids, this can be implemented in $O(k)$ time simply by removing each element of S from the list.

Before stating the main guarantees for no-regret learning and equilibrium computation, it will be useful to discuss a few finer-grained properties typical to games on matroids that help parameterize related computation complexities. First, while not strictly necessary, it will be convenient to restrict our attention to games where the action sets of all players are given by k_i -uniform hypergraphs on some shared groundset Ω , i.e. $A_i \subseteq \binom{\Omega}{k_i}$ for all i . Given such a game $\mathcal{I} = \{(A_i, R_i)_{i=1}^m\}$, we will typically write $k_{-i} = \max_{j \neq i} k_j$ to denote the maximum support of any viable opponent strategy.

Second, it will be useful to introduce an important property of congestion and security games we call *collision-sensitivity*: the vertex-wise reward of an element $v \in \Omega$ only changes if v is also selected by another opponent.

DEFINITION 3.5. (COLLISION-SENSITIVE GAMES) *Let $\mathcal{I} = \{(A_i, R_i)_{i=1}^m\}$ be an m -player linear hypergraph game where $A_i \subseteq 2^\Omega$. We call the rewards of player i ‘collision-sensitive’ if for all $v \in \Omega$, the vertex-wise reward of v only changes if another opponent also selects v :*

$$\forall v \in \Omega \text{ and } s, s' \in A_{-i} \text{ s.t. } v \notin s, s' : R_i^\Omega(v, s) = R_i^\Omega(v, s').$$

We will write $\text{NC}_i : \Omega \mapsto \mathbb{R}$ as the function specifying the i -th player’s no-collision reward function for each vertex, i.e. $\text{NC}_i(x) = R_i^\Omega(x, s)$ for $x \notin s$. We say a collision-sensitive reward has support q if the no-collision reward function NC_i for each player takes on at most q values across all vertices $v \in \Omega$.

In a sense, collision-sensitivity can be thought of as an independence criterion on the vertices: roughly speaking, actions taken on v do not effect actions taken on w for $w \neq v$. With these definitions out of the way, we can now state our main guarantees for no-regret learning on matroids.

THEOREM 3.2. (RWM ON MATROIDS) *Let $\mathcal{I} = \{(A_i, R_i)_{i=1}^m\}$ be an m -player linear hypergraph game on a size- n ground set Ω . If A_i is collision-sensitive with support q , then it is possible to implement δ -RWM $_\beta^T$ in time*

$$O\left(k_i T(\text{CO} + q \log(n) + m k_{-i} T \log(n)) \log\left(\frac{k_i \log(n) + L_{\max} T \log(\beta^{-1})}{\delta}\right)\right),$$

assuming access to a q -piecewise succinct description of NC_i encoded under an ordering of Ω and a contraction oracle matching the same ordering.

The proof of Theorem 3.2 is not particularly interesting beyond combining Observation 3.1 and Theorem 3.1 and involves mostly tedious implementation details of Glauber Dynamics on matroids. We give these details in Appendix C for completeness.

Before moving on, we briefly note this result is nearly tight in many of the main parameters. For instance, the dependence on k, m , and n is $\tilde{O}(mk \log(n))$,⁷ which in many cases (e.g. uniform matroid) is the number of bits required even to express a set of pure strategies for each player. The bound is also linear in q , which is easily seen to be necessary since one needs to know the q distinct values in order to sample.

Many games in the literature satisfy the conditions of Theorem 3.2. We’ll end this subsection by giving a few concrete examples. Perhaps the most well-studied variant of these games is a popular setting called *congestion games*. Congestion games are a natural model for resource competition where m -players compete to share n resources and receive rewards dependent on the number of players sharing the same resource.

⁷This may increase when the support of other players is non-constant. E.g. if all players are playing bases of k -uniform matroids, we require $mk^2 \log(n)$

DEFINITION 3.6. (CONGESTION GAME) Given a ground set Ω , an m -player congestion game on Ω consists of a collection $\{A_i\}_{i=1}^m$ and a reward function $c : \Omega \times [m] \rightarrow \mathbb{R}$ where each $A_i \subseteq 2^\Omega$ are the strategies of player i , and the reward on the actions (s_1, \dots, s_m) is given by:

$$R_i(s_i, s_{-i}) = \sum_{e \in s_i} c(e, |e(s)|)$$

where $e(s) = \{s_j : e \in s_j\}$.

Congestion games are particularly well-studied on matroid bases, which are the only structure on which best response is known to converge to Nash in polynomial time. However, to our knowledge Theorem 3.2 provides the first no-regret algorithm for congestion games.

COROLLARY 3.1. (MATROID CONGESTION WITHOUT REGRET) Let $\mathcal{I} = \{\{A_i\}_{i=1}^m, c\}$ be a congestion game where each A_i is the set of bases of a rank- k_i matroid on a ground set Ω of size n satisfying $\max_{i \in [m]} k_i = k$. Suppose $\text{NC}(e) = c(e, 1)$, the no collision reward function, is q -piecewise under some ordering of Ω . Then there is a no-regret learning algorithm for \mathcal{I} with regret:

$$\text{Reg}_T \leq O\left(L_{\max} \sqrt{T} \cdot \left(\sqrt{k \log(n)} + \sqrt{\log(1/\eta)}\right)\right)$$

with probability at least $1 - \eta$ that runs in time

$$O(kT(\text{CO} + q \log(n) + Tmk \log(n)) \log(L_{\max}Tk \log(n)))$$

assuming access to a q -piecewise succinct description of NC encoded in some ordering of Ω and a contraction oracle matching this ordering.

Proof. It is enough to argue the game is linear, as the result then follows immediately from Theorem 3.2. Denote by s be the strategy tuples chosen by the players. Recall that the reward of any strategy $s_i \in A_i \subset \binom{\Omega}{k_i}$ in the congestion game is given by:

$$R(s_i, s_{-i}) = \sum_{e \in s_i} c(e, |e(s)|).$$

c can easily be extended into the desired vertex-wise reward function, so we are done. \square

COROLLARY 3.2. (EQUILIBRIUM COMPUTATION FOR MATROID CONGESTION) Let $\mathcal{I} = \{\{A_i\}_{i=1}^m, c\}$ be a congestion game where each A_i is the set of bases of a rank- k_i matroid on a shared ground set Ω of size n satisfying $\max_{i \in [m]} k_i = k$. Suppose $\text{NC}(e) = c(e, 1)$ is q -piecewise under some ordering of Ω . Then it is possible to compute an ε -CCE with probability at least $1 - \eta$ in time

$$O\left(mL_{\max}^2 k^2 \log(mn/\eta) \varepsilon^{-2} (\text{CO} + q \log(n) + mL_{\max}^2 k^2 \log^2(mn/\eta) \varepsilon^{-2}) \log(L_{\max}k \log(mn/\eta)/\varepsilon)\right).$$

We note that these results also easily generalizes matroid congestion over any FLC, unlike the best response strategy for computing Nash. Furthermore, we note that Hedge is actually known to converge to better equilibria [KPT09] than original techniques based on best response, which gives this approach an additional potential advantage.

Another setting particularly well-suited to matroids are *security games*, which model a variety of attack/defense scenarios.

DEFINITION 3.7. (SECURITY GAME) A security game $\mathcal{I} = (A_d, A_a, \{r, \zeta, c, \rho\})$ over ground set Ω consists of defender actions $A_d \subseteq 2^\Omega$, attacker actions $A_a = \Omega$, and reward/cost functions $r, \zeta, c, \rho : \Omega \rightarrow \mathbb{R}$. Let $S \in A_d$, $i \in A_a$ be the actions taken by the defender and the attacker respectively. The reward matrices are given by:

$$R_d(S, i) = \begin{cases} r(i) & \text{if } i \in S \\ c(i) & \text{else,} \end{cases} \quad \text{and} \quad R_a(S, i) = \begin{cases} \zeta(i) & \text{if } i \in S \\ \rho(i) & \text{else.} \end{cases}$$

Security games can model a couple natural settings dependent on the choice of parameters. One basic setting is where the defender has k security resources to defend a set of n targets, and ‘wins’ if the attacker chooses a defended target. On the other hand, the model also captures the complement of this game where the defender chooses k targets to distribute key resources, and the attacker wins if they intercept this distribution (pick one of the k marked targets). Security games have broad applicability in practice, and indeed have been used in cases such as assigning security checkpoints at LAX [Tam11].

Security games are inherently linear in their natural representation and thus admit efficient no-regret algorithms when the defender’s state-space is a matroid (simulating RWM for the attacker is trivial as it corresponds to a size- $|\Omega|$ multinomial distribution).

COROLLARY 3.3. (SECURITY WITHOUT REGRET) *Let $\mathcal{I} = (A_d, A_a, \{r, \zeta, c, \rho\})$ be a security game where A_d are the bases of a rank- k matroid on the ground set Ω and $A_a = \Omega$. Suppose $c, \rho : \Omega \mapsto \mathbb{R}$ are q -piecewise under some ordering of Ω . Then there exists a no-regret learning algorithm for \mathcal{I} with regret:*

$$\text{Reg}_T \leq O\left(L_{\max} \sqrt{T} \cdot \left(\sqrt{k \log(n)} + \sqrt{\log(1/\eta)}\right)\right)$$

with probability at least $1 - \eta$ that runs in time

$$O(kT(\text{CO} + q \log(n) + T \log(n)) \log(k \log(n) T L_{\max})),$$

assuming access to q -piecewise succinct descriptions for c and ρ encoded in some ordering of Ω and a contraction oracle matching this ordering.

Proof. Note that the attacker’s strategy consists of bases of just rank-1 matroid so implementing RWM for her is trivial. We focus on the implementation for the defender side. Again it is enough to show the game is linear. By definition, we have

$$R_D(S, i) = \sum_{j \in S} R_D^\Omega(j, i)$$

where $R_D^\Omega : [n] \times [n] \rightarrow \mathbb{R}$ is

$$R_D^\Omega(i, j) = \begin{cases} r(i) & \text{if } i = j \\ c(i) & \text{else.} \end{cases}$$

□

As an immediate corollary, we also get fast equilibrium computation.

COROLLARY 3.4. (EQUILIBRIUM COMPUTATION FOR SECURITY) *Let $\mathcal{I} = (A_d, A_a, \{r, \zeta, c, \rho\})$ be a security game where A_d are the bases of a rank- k matroid on the ground set Ω and $A_a = \Omega$. Suppose $c, \rho : \Omega \mapsto \mathbb{R}$ are q -piecewise under some ordering of Ω . Then it is possible to compute an ε -CCE (Nash if the game is zero-sum) with probability at least $1 - \eta$ in time*

$$O\left(L_{\max}^2 k^2 \log(n/\eta) \varepsilon^{-2} (\text{CO} + q \log(n) + L_{\max}^2 k \varepsilon^{-2} \log^2(n/\eta)) \log\left(k \log(n/\eta) \varepsilon^{-1} L_{\max}\right)\right),$$

assuming access to an q -piecewise succinct description for c and ρ encoded under some ordering of Ω and a contraction oracle matching this ordering.

We note that this result easily generalizes to settings with multiple attackers or an attacker who chooses targets corresponding to a matroid basis.

3.2 Dueling Games and the JSV Chain Matroids (or more generally FLC’s) are not the only type of constrained state space that can be sampled under arbitrary external fields. Indeed, long before these results Jerrum, Sinclair, and Vigoda [JSV04] famously proved (in work on approximating the permanent) that bipartite matchings have this property as well. We give an improved version of their result due to Bezáková, Štefankovič, Vazirani, and Vigoda [BŠVV04].

THEOREM 3.3. (JSV CHAIN) *Let $(K_{n,n}, w)$ be an edge-weighted complete bipartite graph, and consider the distribution over perfect matchings given by:*

$$\Pr(M) \propto \prod_{e \in M} w_e.$$

It is possible to δ -approximately sample from this distribution in $\tilde{O}(n^7 \log \frac{1}{\delta w_{\min}})$ time, where w_{\min} is the minimum weight.

Note one can phrase this result as a sampling algorithm for permutations over external fields, where the state-space is viewed as a subset of $[n]^n$. Like matroids, bipartite matchings are very natural objects and underlie a fair number of well-studied games. In this section, we focus on the setting of *dueling games*. Dueling games model two player competitive optimization over a shared ground set.

DEFINITION 3.8. (DUELING GAMES) *A dueling game $\mathcal{I} = (\Omega, \mu, A_1, A_2)$ consists of a set Ω , a distribution μ over Ω , and strategy spaces $A_1, A_2 \subset \mathbb{R}_+^\Omega$. The reward matrices are given by the probability of ranking $x \sim \mu$ higher than the opponent:*

$$R_1(s, t) = \Pr_{x \sim \mu}[s(x) > t(x)] - \Pr_{x \sim \mu}[t(x) > s(x)],$$

and likewise:

$$R_2(s, t) = \Pr_{x \sim \mu}[t(x) > s(x)] - \Pr_{x \sim \mu}[s(x) > t(x)].$$

There is no known polynomial time algorithm for computing equilibria of general dueling games. We will give a general algorithm for a class of dueling games we call *unrestricted*.

DEFINITION 3.9. (UNRESTRICTED DUELING GAMES) *A dueling game $\mathcal{I} = (\Omega, \mu, A_1, A_2)$ is called un-restricted if there exist subsets $S_1, S_2 \subset \mathbb{R}$ with $|S_1| = |S_2| = |\Omega|$ such that A_1 (respectively A_2) consists of all possible assignments of Ω to S_1 (respectively S_2).*

It is not hard to see that unrestricted dueling games are linear over perfect matchings in a complete bipartite graph. As a result, we can use the JSV-chain to simulate optimistic hedge in polynomial time.

THEOREM 3.4. (SAMPLING UNRESTRICTED DUELING GAMES) *Let $\mathcal{I} = (\Omega, \mu, A_1, A_2)$ be an unrestricted dueling game where $|\Omega| = n$. Then it is possible to implement δ -RMW $_\beta^T$ in time $\tilde{O}(T^2 n^7 \log(1/\delta))$.*

Proof. We focus on player 1. The result is analogous for player 2. Strategies in an unrestricted dueling game correspond to perfect matchings in the complete bipartite graph $K_{n,n}$, where the LHS corresponds to elements of Ω , and the RHS corresponds to elements in S_1 . To fit into our prior framework of linearity and external fields, one may view these perfect matchings as elements of E^n (where E is the edge set of $K_{n,n}$). Recall that the reward is given by:

$$R_1(s, t) = \Pr_{x \sim \mu}[s(x) > t(x)] - \Pr_{x \sim \mu}[t(x) > s(x)]$$

It is not hard to see this is linear over the edges of matching:

$$R_1(s, t) = \sum_{e \in S} R_1^\Omega(e, t)$$

where $R_1^\Omega : E \times A_2 \rightarrow \mathbb{R}$ is given by:

$$R_1^\Omega(\{v, w\}, t) = \begin{cases} \mu(v) & \text{if } w > t(v) \\ -\mu(v) & \text{if } w < t(v) \\ 0 & \text{else} \end{cases}$$

As a result, Observation 3.1 implies that RWM is given by the application of an external field over the edges of perfect bipartite matchings, which is exactly the distribution considered in Theorem 3.3. All that is left is to efficiently build access to the weights of the underlying bipartite graph, which is a small onetime cost that is asymptotically dominated by even the mixing time of the JSV chain. As a result, it is enough to run Theorem 3.3 T times, which gives the resulting runtime bound. \square

We note that this result can easily be generalized to a slightly larger class of games where $|S_1|$ and $|S_2|$ may be larger than Ω , and specific edges in the bipartite representation may be disallowed (i.e. we might add the constraint that $x \in \Omega$ can never be given rank 1). Such strategies correspond to sampling matchings on a generic bipartite graph (rather than $K_{n,n}$), and no-regret learning can also be performed by the JSV-chain.

Finally, we'll look at a classic dueling game that fit into the unrestricted framework: ranking duel. Ranking duel (or the 'search engine game') is a game where two players compete to choose the best ranking of n items. One of these items is pulled from a known distribution, and the player who ranked it higher wins.

Ranking duel is an unrestricted game where $S_1, S_2 = [n]$ and the action spaces $A_1, A_2 = \mathcal{S}_n$, i.e. permutations of n . As a result Theorem 3.4 immediately implies an efficient algorithm for sampling in δ -RWM.

COROLLARY 3.5. (RANKING DUEL WITHOUT REGRET) *Let $\mathcal{I} = ([n], \mu, \mathcal{S}_n, \mathcal{S}_n)$ be an instance of ranking duel. Then there exists an algorithm with regret:*

$$\text{Reg}_T \leq O\left(\sqrt{T} \cdot \left(\sqrt{k \log(n)} + \sqrt{\log(1/\eta)}\right)\right)$$

with probability at least $1 - \eta$ that runs in time $\tilde{O}(T^2 n^7)$.

As a corollary we get the fastest known equilibrium computation for ranking duel,

COROLLARY 3.6. (EQUILIBRIUM COMPUTATION FOR RANKING DUEL) *Let $\mathcal{I} = ([n], \mu, \mathcal{S}_n, \mathcal{S}_n)$ be an instance of ranking duel. Then there exists an algorithm computing an ε -CCE (Nash if the game is zero-sum) with probability at least $1 - \eta$ in time $\tilde{O}(n^9 \log(1/\eta)/\varepsilon^4)$*

Unfortunately, while the JSV-chain is an improvement over previous extended linear programming approaches to dueling games [IKL⁺11, ADH⁺19], n^9 can hardly be called a practical running time. In fact, it should be noted there is a faster known no-regret algorithm for perfect bipartite matchings called PERMELEARN that runs in $O(Tn^4)$ time.

Thus Theorem 3.4 is perhaps more interesting from the perspective of the method than the result itself. Designing faster and simpler Markov chains for sampling bipartite matchings has long been a favorite open problem in the sampling community. Our setting gives a nice intermediate version of this problem, as the matching problems arising from unrestricted dueling games have particularly nice structure. In particular, they correspond to *monotonic weightings*, in the sense that for every fixed vertex v on the LHS on the graph, the edge weight of $w(v, i) \geq w(v, j)$ if $i \geq j$. Matchings with monotonic weights are actually well-studied in the literature, including the resolution of the monotone column permanent conjecture [BHVW11] and rapid mixing of the switch chain⁸ for binary monotonic weights [DJM17]. However despite these related results, a fast algorithm for sampling general bipartite matchings with monotonic weights remains an interesting open problem, and the application to practical no-regret algorithms for dueling games gives yet another motivation for its study.

4 Playing Games via DP-Sampling

While MCMC-sampling is a powerful tool, standard techniques like Glauber Dynamics may not perform well in settings that require global coordination across coordinates. In this section, we develop a new sampling technique toward this end based on Dynamic Programming, taking advantage of the fact that many settings of interest, such as Colonel Blotto, additionally exhibit certain recursive structure. In particular, we consider a large class of problems called *Resource Allocation Games* that broadly generalize the Colonel Blotto game.

Resource Allocation Game. In a resource allocation game, each player assigns fungible items to some number of battlefields. Namely, for the i -th player, the action space A_i is the set of ordered size k_i partition of n_i for $k_i, n_i \in \mathbb{Z}^+$.⁹ One can see the action space is indeed a hypergraph where the vertices correspond to pairs (h, x) interpreted as "assigning x items to the h -th battlefield", and a strategy is simply a subset of vertices $(1, x_1), \dots, (k, x_k)$ satisfying $\sum_{h=1}^k x_h = n$. Let the variables $x_{i,h}$ denote the number of items assigned by the i -th player to the h -th battlefields and

⁸The switch chain is in essence the 2-step Glauber Dynamics on the view of matchings as a subset of E^n .

⁹While we define actions as all k -partitions of $[n]$, our algorithm also applies to action spaces that are subsets of these partitions with arbitrary assignment constraints on each battlefield (i.e., of the form, at most m items can be assigned to battle j).

$A_{-i} := A_1 \times \cdots \times A_{i-1} \times A_{i+1} \times \cdots \times A_m$ denote the set of action tuples from players other than player i . The reward structure of a resource allocation game is defined by a set of battlefield reward functions $r_{i,h} : [0, n] \times A_{-i} \mapsto \mathbb{R}$ for each player $i \in [m]$ and battlefield $h \in [k_i]$. Let $\mathbf{a}_{-i} \in A_{-i}$ be the actions picked by players other than player i . The total reward received by the i -th player is given by the sum of rewards over individual battles:

$$R_i(x_{i,1}, \dots, x_{i,h}, \mathbf{a}_{-i}) = \sum_{h=1}^{k_i} r_{i,h}(x_{i,h}, \mathbf{a}_{-i}).$$

Additionally, let $r_{i,h,\mathbf{a}_{-i}} : [0, n] \mapsto \mathbb{R}$ be the restriction of $r_{i,h}$ after fixing the strategies of the other players i.e. $r_{i,h,\mathbf{a}_{-i}}(x) = r_{i,h}(x, \mathbf{a}_{-i})$. We say the resource allocation game is q -piecewise and monotonic if for every $i \in [m]$, $h \in [k_i]$, and $\mathbf{a}_{-i} \in A_{-i}$, $r_{i,h,\mathbf{a}_{-i}}$ is a q -piecewise constant and monotonically increasing function.

In the remainder of this section, we discuss how one achieves no-regret learning for the first player (the algorithm for the other players is analogous). For this purpose, we will drop the subscript indicating the player number and let $n = n_1, k = k_1, r_h = r_{1,h}, \mathbf{a} = \mathbf{a}_{-1}$.

RWM Distributions. To achieve no-regret learning, we will need to approximately sample from the distributions arising from running the RWM algorithm on Resource Allocation Games. Assume that the players have played the game for T rounds. Let $a^{(t)} \in A_{-1}$ be the action tuples observed from all but the first player at the t -th round. For an action $x \in A_1$ that assigns x_h items to the h -th battlefield, RWM will set its weight $w_T(x)$ as

$$(4.5) \quad w_T(x) = \beta^{\sum_{h=1}^k \sum_{t=1}^T r_h(x_h, a^{(t)})},$$

where $\beta \in [1/2, 1)$ is the learning rate. For simplicity, we will define the cumulative battlefield reward function (negated to simplify the syntax appearing after)

$$(4.6) \quad \ell_h^{(T+1)}(x_h) = - \sum_{t=1}^T r_j(x_h; a^{(t)}),$$

so the weight for strategy x can alternatively be written as $\prod_{h=1}^k \beta^{\ell_h^{(T+1)}(x_h)}$. We remark that, if the resource allocation game is monotonic and q piecewise, $\ell_j^{(T+1)}(\cdot)$ is also monotonically increasing and a $((q-1) \cdot t + 1)$ -piecewise constant function. Though the domain of ℓ_j is of size $n+1$, the property allows us to represent it succinctly in space $O(qt)$, which is critical when we try to design algorithms whose runtime depends on n polylogarithmically. In this section, we focus on how one could design algorithms to efficiently sample from $RM^{(T)}(\beta)$ approximately given succinct descriptions of the functions $r_{h,\mathbf{a}^{(t)}}$ for all $h \in [k]$ and $t \in T$ (recall that $r_{h,\mathbf{a}^{(t)}}$ is the restriction of the battlefield reward function r_h after fixing the other players' actions). This allows us to implement δ -RWM in polylogarithmic time.

THEOREM 4.1. (RWM IN RESOURCE ALLOCATION GAME) *Let $\mathcal{I} = \{(A_i, R_i)_{i=1}^m\}$ be an m -player monotonic, q -piecewise resource allocation game where $A_1 = P_k(n)$ for $n, k \in \mathbb{Z}^+$. Suppose the reward of the first player is bounded by L_{\max} . Then it is possible to implement δ -RWM $_{\beta}^T$ in time*

$$O(Tk) \cdot \left(\min \left(T^2 L_{\max} \log(1/\beta) / \delta \cdot \zeta_1 (\log T + \log \zeta_2) \cdot \log n, n \log n \right) + \min(Tq, n) \right).$$

where $\zeta_1 := \min(L_{\max} \log(1/\beta) / \delta, q)$, $\zeta_2 := \max(L_{\max} \log(1/\beta) / \delta, q)$, assuming access to a Tq -piecewise succinct description of $\ell_h^{(t)}$ defined in Equation (4.6) for all $h \in [k], t \in [T]$.

As corollaries of the above theorem, we obtain efficient no-regret learners and algorithms for computing CCEs in Resource Allocation Games.

In the remainder of this section, we present the sampling algorithm whose analysis leads to the proof of Theorem 4.1, and discuss a number of applications to games in the literature including Colonel Blotto and its variants.

4.1 Sampling via estimation of partition function We will focus for the moment on how one could sample from the RWM distribution just in round T . For that purpose, we often omit the superscript T for the function $\ell_h^{(T)}$ (Equation (4.6)) for simplicity. To sample with a dynamic program, we define the partition function for an RWM distribution in resource allocation.

Partition Function The partition function $f_h : [0] \cup \mathbb{Z}^+ \mapsto \mathbb{R}^+$ for $h \in [k]$ is defined as the sum of partial weights of all strategies that allocate y soldiers in the subgame induced on the first h battles. Namely,

$$(4.7) \quad f_h(y) = \sum_{x_1 + \dots + x_h = y} \prod_{i=1}^h \beta^{\ell_i(x_i)}.$$

It has long been known that efficient algorithms for computing the partition function of a self-reducible problem imply efficient (approximate) samplers for the problem's solution space [JVV86]. As one can see, computing the partition function f_h in our setting simply corresponds to counting the number of (weighted) size h partitions of y , which is exactly such a self-reducible problem. Consequently, if one has the values for the partition functions precomputed, one can use them to sample from the RWM distribution efficiently. We provide the detailed sampling procedure below for completeness.

In particular, this is done by sampling the number of items to put in each battlefield sequentially, conditioned appropriately on prior choices. One puts $x_1 \in \{0, \dots, n\}$ soldier to the first battlefield with probability

$$(4.8) \quad \Pr[x_1 = y] \propto \beta^{\ell_1(y)} \cdot f_{k-1}(n - y).$$

To sample from the $(h + 1)$ -th battlefield conditioned on the fact that one has put x_1, \dots, x_h soldiers in battles $1 \dots h$, it is enough to sample according to the distribution

$$(4.9) \quad \Pr[x_{h+1} = y | x_{1 \dots h}] \propto \beta^{\ell_{h+1}(y)} \cdot f_{k-h-1} \left(n - \left(\sum_{j=1}^h x_j \right) - y \right).$$

The probabilities defined according to Equations (4.8), (4.9) yields exactly the RWM distribution, but computing the partition function exactly can be quite costly. For this purpose, we consider the notion of δ -approximations (multiplicative) of functions.

DEFINITION 4.1. (δ -APPROXIMATION) Given $f : [0, n] \mapsto \mathbb{R}^+$ and $\hat{f} : [0, n] \mapsto \mathbb{R}^+$, we say \hat{f} is a δ -approximation of f if for all $x \in [0, n]$ we have

$$(1 - \delta)f(x) \leq \hat{f}(x) \leq (1 + \delta)f(x).$$

Fortunately, with $\delta/(Ck)$ -approximations of the partition functions for some sufficiently large constant C , one can still perform δ -approximate sampling from the RWM distribution (See proof of Lemma 4.1).

Another important observation for achieving efficient (approximate) sampling is that ℓ_h , the reward function for each battlefield, is a piece-wise constant function. Hence, further optimization is possible if the approximations used for each partition function is also piece-wise constant (and this is indeed the case as we will see shortly). The pseudo-code for the sampling algorithm that takes as input the succinct descriptions of the (approximate) partition functions and the reward functions is given below.

Require: Succinct descriptions $D_{\hat{f}_h}$ for $h \in [k]$; succinct description D_{g_h} of the function $g_h(x) = \beta^{\ell_h(x)}$.

- 1: Initialize the number of unused soldiers $u = n$.
- 2: Initialize an empty assignment description S .
- 3: **for** $h = 1 \dots k$ **do**
- 4: Compute the succinct description D_{κ_h} of the function $\kappa_h : [0, u] \mapsto \mathbb{R}^+$ defined as

$$(4.10) \quad \kappa_h(i) := g_h(i) \cdot \hat{f}_{k-h}(u - i).$$

▷ $|D_{\kappa_h}| = |D_{g_h}| + |D_{\hat{f}_h}|$ and D_{κ_h} can be computed in time linear with respect to the description length.

- 5: {Compute intervals}
- 6: **for** $(a_i, b_i, y_i) \in D_{\kappa_h}$ **do**
- 7: Compute the cumulative weight of constant intervals.

$$(4.11) \quad v_i := y_i \cdot (b_i - a_i + 1).$$

- 8: {Sample an interval}
 - 9: Sample $j \in 1 \dots |D_{\kappa_h}|$ according to the weight vector v .
 - 10: {Sample soldiers used in battle h }
 - 11: Sample z_h uniformly for $\{a_j, a_j + 1, \dots, b_j - 1, b_j\}$ where $[a_j, b_j]$ is the j -th constant interval in D_{κ_h} .
 - 12: Add z_h to the strategy S description.
 - 13: $u \leftarrow u - z_h$.
 - 14: **return** S .
-

LEMMA 4.1. Let $\mathcal{I} = \{(A_i, R_i)_{i=1}^m\}$ be an m -player monotonic, q -piecewise resource allocation game where $A_i = P_{k_i}(n_i)$ for $n_i, k_i \in \mathbb{Z}^+$. At the t -th round, for each $h \in [k]$, let \hat{f}_h be $\delta/(2k)$ -approximations of the partition function defined in Equation (4.7), and let $g_h(x) = \beta^{\ell_h(x)}$. Assume one is given the succinct descriptions $D_{\hat{f}_h}$ and D_{g_h} . Then, there exists an algorithm **Partition-Sampling** which performs δ -approximate sampling from $RM^{(T)}(\beta)$ in time

$$k \cdot O(p + \min(Tq, n) + \log n)$$

where $p := \max |D_{\hat{f}_h}|$.

Proof. If we were to perform the sampling process with f_h instead of \hat{f}_h , we would get exactly the distribution $RM^{(t)}(\beta)$. This follows from repeated applications of Bayes' rule and that we are sampling the correct conditional distribution. Namely,

$$\Pr[x_1 = y_1, \dots, x_k = y_k] = \prod_{h=1}^k \Pr[x_h = y_h \mid x_{h'} = y_{h'}, \forall h' < h]$$

To make sure that we are performing δ -approximate sampling overall, it suffices if we perform δ/k approximate sampling for each conditional distribution (each battlefield). To perform exact sampling, one needs to compute the weight

$$\kappa_h^*(i) := g_h(i) \cdot f_{k-h-1}(u - i).$$

Since for every $h \in [k]$ we have \hat{f}_h are $\delta/(2k)$ approximations of f_h , we also have κ_h is $\delta/(2k)$ approximations of κ_h^* , which implies that

$$(1 - \delta/(2k)) \sum_{i=0}^n \kappa_h^*(i) \leq \sum_{i=0}^n \kappa_h(i) \leq (1 + \delta/(2k)) \sum_{i=0}^n \kappa_h^*(i).$$

It then follows that the distribution defined by κ_h^* and κ_h differs by at most δ in total variation distance.

To analyze the runtime, we note that for each battlefield, we first compute the succinct description D_{κ_h} defined in Equation (4.10). Since it is the point-wise multiplication between g_h , which is $\min(Tq, n)$ -piecewise constant, and \hat{f}_{k-h-1} , which is p -piecewise constant, κ_h will be $p + \min(Tq, n)$ piecewise constant. To construct the succinct description of D_{κ_h} , one maintains two pointers $a = 0, b = u$ and keeps track of the interval from D_{g_h} that a is in and the interval from $D_{\hat{f}_h}$ that b is in. Then, one shifts a forward and b backward to seek for constant intervals of D_{κ_h} . It is easy to see the runtime of the construction is linear with respect to $|D_{\kappa_h}|$. Then, computing v_i requires scanning through the succinct description of D_{κ_h} once. After that, we first sample from a multinomial distribution with support at most p , which takes time $O(p)$. Then, we sample from a uniform distribution with support at most n , which takes time $O(\log n)$. Adding everything together then gives our final runtime. \square

4.2 Computing the Partition Function We now move to showing how to (approximately) compute the partition function. As a warmup, we will first show how this can be done exactly via dynamic programming. In particular, we want to fill a $k \times n$ table such that the (h, y) entry corresponds to the value $f_h(y)$.

PROPOSITION 4.1. *The values $f_h(y)$ for all $h \in [k]$ and $y \in [0, n]$ can be computed in time $O(nk \log n)$.*

Proof. Notice that we have the following recursion

$$(4.12) \quad f_h(y) = \sum_{x=0}^y \beta^{\ell_h(x)} \cdot f_{h-1}(y-x).$$

f_h is exactly the convolution of $\beta^{\ell_h(\cdot)}$ and f_{h-1} . Using Discrete Fast Fourier Transform, f_h can be evaluated in time $O(n \log n)$ ([Bri88]). Hence, in total, the entire DP table can be filled in time $O(nk \log n)$. \square

Next, we will discuss how one can develop faster algorithms when n is substantially larger than k, T , and L_{\max} . Our main technical result is an algorithm to pre-compute the partition functions “approximately” whose runtime depends on n *polylogarithmically*.

PROPOSITION 4.2. *There exists an algorithm **Approx-DP** which constructs $\hat{f}_1, \dots, \hat{f}_k$ such that \hat{f}_i is a δ -approximation of the partition function f_i pointwisely, and runs in time*

$$O(kT^2 L_{\max} \log(1/\beta)/\delta \cdot \zeta_1 (\log T + \log \zeta_2) \cdot \log n),$$

where $\zeta_1 := \min(L_{\max} \log(1/\beta)/\delta, q), \zeta_2 := \max(L_{\max} \log(1/\beta)/\delta, q)$.

This seems a bit surprising as there are in total $k \cdot (n + 1)$ values that we need to pre-compute ($f_h(y)$ for all $h \in [k]$ and $y \in [0, n]$). However, notice that we are only interested in computing approximations to these values. And, as each f_h is itself a monotonically increasing function, we can approximate it with a sufficiently simple piece-wise function.

FACT 4.1. *Given a monotonically increasing function $f : [0, n] \mapsto \mathbb{R}^+$, it can be δ -approximated by a function that is d -piecewise constant where $d = \Theta(\log(\max_x f(x)/\min_x f(x))/\delta)$.*

Hence, the algorithm **Approx-DP** does not need to output the entire $k \times (n + 1)$ tables specifying the partition functions. Rather it can just construct the succinct descriptions of a series of functions $\hat{f}_1, \dots, \hat{f}_k$ such that \hat{f}_i is a δ -approximation of f_i . By Claim 4.1, we can indeed find such \hat{f}_i that are $\Theta(\log(\beta^{TL_{\max}})/\delta) = \Theta(TL_{\max} \log(1/\beta)/\delta)$ piecewise constant. As the first building block of the **Approx-DP**, we demonstrate the routine which, given query access to an unknown monotonically increasing function, constructs a piecewise constant approximation of the function.

LEMMA 4.2. *Given a monotonically increasing function $f : [0, n] \mapsto \mathbb{R}^+$ and query access to f , there exists an algorithm **Piecewise-Approximate** which outputs a piecewise function \hat{f} satisfying that*

- \hat{f} is d -piecewise constant for $d = \Theta(\log(\max_x f(x)/\min_x f(x))/\delta)$.

- $(1 - \delta) \cdot f(x) \leq \hat{f}(x) \leq f(x)$.
- the algorithm runs in time $O(\log n \cdot d \cdot Q)$.

where Q is the cost of making a single query to f .

Proof. The algorithm proceeds by iteratively finding the longest interval from a starting point such that the function values at the endpoints are within a $(1 + \delta)$ -factor of each other, and then letting \hat{f} on this interval be the constant function given by the value of f on the right endpoint. By the monotonicity of f , it is easy to see that \hat{f} is indeed a δ -approximation of f on this interval. The algorithm then repeats this process starting from the next value on which \hat{f} is not yet defined and repeats until \hat{f} has been defined on the entire domain.

Since the function values of f increase by at least a factor of $(1 + \delta)$ between each interval and the last, the total number of intervals is at most $d = \Theta(\log(\max_x f(x)/\min_x f(x))/\delta)$. \square

Algorithm 1 Piecewise-Approximate

Require: Query access to $f : [0, n] \mapsto \mathbb{R}^+$; Approximation accuracy δ .

- 1: $D \leftarrow \{\}, a \leftarrow 0$.
 - 2: **while** $a \leq n$ **do**
 - 3: Binary Search to find the largest b such that $f(b) \leq f(a) \cdot (1 + \delta)$.
 - 4: Add $(a, b, f(a))$ to D . (setting $\hat{f}(x)$ to $f(a)$ for all $a \leq x \leq b$.)
 - 5: $a \leftarrow b + 1$.
 - 6: **return** D .
-

With this in mind the construction of the series of piecewise constant approximation functions $\hat{f}_1, \dots, \hat{f}_k$ becomes clear: one initializes $\hat{f}_1 = \mathbf{Piecewise-Approximate}(f_1)$ and then defines \hat{f}_h recursively as $\hat{f}_h = \mathbf{Piecewise-Approximate}(f'_h)$ where $f'_h(y) = \sum_{x=0}^y \hat{f}_{h-1}(x) \cdot g_h(y-x)$. Recall that the routine **Piecewise-Approximate** requires query access to the input function. Hence, we need to show how one could implement query access to f_1 and f'_h for $h \in [k]$ efficiently (independent of n). The former is easy given the succinct description of f_1 since $f_1 = \beta^{\ell_1(x)}$ is a $T \cdot q$ piecewise constant function. To realize the latter, we use the fact that f'_h is the convolution of \hat{f}_h and g_h which are both piecewise constant functions. In particular, we give the following routine which efficiently implements query access to convolutions of piecewise constant functions.

Algorithm 2 Convolution-Query

Require: Succinct descriptions D_f, D_g of two piece-wise constant functions f, g ; Query point x .

- 1: Preprocess D_f to extend each tuple into the form (a_i, b_i, y_i, s_i) where $s_i := \sum_{j < i} (b_j - a_j + 1) \cdot y_j$.
- 2: Let

$$F(z) = \begin{cases} \sum_{i=0}^z f(i) & \text{for } i \geq 0, \\ 0 & \text{otherwise.} \end{cases}$$

- 3: $res \leftarrow 0$.
 - 4: **for** each interval $(a_i, b_i, y_i) \in D_g$ **do**
 - 5: $res += (F(x - a_i) - F(x - b_i - 1)) \cdot y_i$.
 - 6: **return** res .
-

We note that the pre-processing incurs a one-time cost for each new function f and does not need to be performed for different queries with respect to the same function f . We also note $F(z)$ can be evaluated by first binary search the first index j such that for $(a_j, b_j, y_j, s_j) \in D_f$ we have $b_j \geq z$. Then use the equality $F(z) = s_j + (z - a_j) \cdot y_j$.

LEMMA 4.3. Assume one is given the succinct description D_f, D_g of two functions $f, g : [0, n] \mapsto \mathbb{R}^+$. Let $(f \star g)(x) = \sum_{i=0}^x f(i) \cdot g(x-i)$. Suppose $|D_f| = p_f$ and $|D_g| = p_g$ with $p_f < p_g$. There exists an algorithm **Convolution-Query** that takes $O(p_f + p_g)$ time to preprocess D_f, D_g and then takes $O(p_f \cdot \log p_g)$ time to return query access $(f \star g)(x)$ for each $x \in [0, n]$.

Proof. In the pre-processing step, for each tuple $(a_i, b_i, y_i) \in D_f$, we add an extra number s_i which denotes the prefix-sum of all elements before the interval $[a_i, b_i]$. This can be done easily by scanning through the tuples of succinct description in order in one pass. Then, if one want to query the prefix sum $F(z)$, one can just find out which interval z falls into by binary search and then computes in constant time with s_i .

Then, using the fact that g is piecewise constant, we can rewrite the convolution query $(f \star g)(x)$ as

$$\sum_{i=0}^x f(x-i) \cdot g(i) = \sum_{i=1}^{|D_g|} y_i \cdot \left(\sum_{j=x-b_i-1}^{x-a_i} f(j) \right) = \sum_{i=1}^{|D_g|} y_i \cdot (F(x-a_i) - F(x-b_i-1)).$$

where (a_i, b_i, y_i) are tuples in D_g . Since evaluating each query to F takes at most $O(\log |D_f|)$ time (for binary search), the above expression can be evaluated in time $O(|D_g| \log |D_f|)$. \square

We are now ready to present the pseudocode and analysis of **Approx-DP**, whose analysis then lead to the proof of Proposition 4.2.

Algorithm 3 Approx-DP

Require: For each $h \in [k]$, succinct description D_{g_h} of the function $g_h(x) = \beta^{-\ell_h(x)}$; number of goods n ; number of battlefields k ; approximation error δ .

1: Initialize \hat{f}_1 as g_1 .

$$D_{\hat{f}_1} \leftarrow \mathbf{Piecewise-Approximate}(g_1, \delta/(4k)).$$

2: **for** $h = 2 \cdots k$ **do**

3: Consider the function

$$(4.13) \quad f'_h(y) = \sum_{x=0}^y \hat{f}_{h-1}(x) \cdot g_h(y-x) = \sum_{x=0}^y \hat{f}_{h-1}(x) \cdot \beta^{\ell_h(y-x)}.$$

4: Set \hat{f}_h to be the piecewise approximation of f'_h .

$$D_{\hat{f}_h} \leftarrow \mathbf{Piecewise-Approximate}(f'_h, \delta/(4k)).$$

5: **return** $D_{\hat{f}_h}$ for $h \in [k]$.

We note that the query access to g_1 is by simply reading from the succinct description D_{g_1} . Further, query access to $f'_h(y)$ is implemented by routine **Convolution-Query** $(D_{\hat{f}_{h-1}}, D_{g_h}, y)$ (see the Appendix for its pseudocode and runtime).

Proof. [Proof of Proposition 4.2] We will show via induction that each \hat{f}_h is monotonically increasing and \hat{f}_h an $(h \cdot \delta/k)$ -approximation of the original function f_h . Consider the function f'_h defined in Equation (4.13). Since both $\hat{f}_{h-1}(x)$ and $\beta^{-\ell_h(x)}$ are monotonically increasing, their convolution f'_h is also monotonically increasing. Besides, by our inductive hypothesis, \hat{f}_{h-1} is an $((h-1) \cdot \delta/k)$ -approximation of f_{h-1} , implying that f'_h is an $((h-1) \cdot \delta/k)$ -approximation of f_h . By Lemma 4.2, \hat{f}_h is a (δ/k) -approximation of f'_h , and consequently a $(h \cdot \delta/k)$ -approximation of f_h .

The runtime of the algorithm is dominated by the $(k-1)$ times we call the routine **Piecewise-Approximate** in Line 4. Notice that $\max_x f_h(x) / \min_x f_h(x)$ can be at most $\beta^{-\ell(x)}$ where $\ell(x)$ is at most TL_{\max} . As we have argued, each

f'_h is always a δ -approximation of f_h , which implies that $\max_x f'_h(x) / \min_x f'_h(y)$ is at most $\exp(2 \cdot \log(1/\beta)TL_{\max})$. Therefore, each \hat{f}_h is a $O(\log(1/\beta)TL_{\max})$ constant function.

Piecewise-Approximate uses the routine **Convolution-Query** as its query access to the input function f'_h . By Lemma 4.3, **Convolution-Query** incurs a one-time cost of $O(TL_{\max} \log(1/\beta)/\delta + Tq)$ to preprocess the succinct descriptions of $g_h(x) = \beta^{\ell_h(x)}$ and \hat{f}_{h-1} . Then, each query takes time $O(T \cdot \zeta_1 \cdot (\log T + \log \zeta_2))$, where $\zeta_1 := \min(L_{\max} \log(1/\beta)/\delta, q)$, $\zeta_2 := \max(L_{\max} \log(1/\beta)/\delta, q)$. By Lemma 4.2, each call to **Piecewise-Approximate** then takes time

$$O\left(T^2 L_{\max} \log(1/\beta)/\delta \cdot \zeta_1 (\log T + \log \zeta_2) \cdot \log n\right).$$

The overall runtime just multiplies the entire expression by k . \square

Combining Lemmas 4.1, 4.1, and Lemma 4.2, we then obtain an efficient algorithm for δ -RWM in Resource Allocation Games, which concludes the proof of Theorem 4.1.

4.3 Applications of the meta algorithm

In this section, we describe the main applications of our sampling algorithm.

Colonel Blotto Game A well-studied example of the resource allocation game is the Colonel Blotto Game. In the game, m players try to assign $\{n_i\}_{i=1}^m$ troops to k different battlefields. For the i -th battlefield, the player who places more soldiers wins the battle and earn a reward of $w_i \in \mathbb{Z}^+$ (ties are broken e.g. lexicographically).¹⁰ This can be viewed as a resource allocation game where the reward function r_j is simply the threshold function $r_j(x) = w_j \cdot \mathbb{1}(x > y)$ where y is the maximum number of soldiers placed by the other players. It is easy to see that r_j is monotonically increasing and 2-piecewise constant. Hence, Theorem 4.1 immediately gives an efficient no-regret learning algorithm for the Colonel Blotto Game.

COROLLARY 4.1. (COLONEL BLOTTO WITHOUT REGRET) *Let \mathcal{I} be an m -player Colonel Blotto Game where the i -th player tries to assign n_i soldiers to k battlefields satisfying $n_i \leq n$. Then there is a no-regret learning algorithm for \mathcal{I} with regret:*

$$\text{Reg}_T \leq O\left(L_{\max} \sqrt{T} \cdot \left(\sqrt{k \log(n)} + \sqrt{\log(1/\eta)}\right)\right)$$

with probability at least $1 - \eta$ that runs in time

$$Tk \cdot O\left(\min\left(n \log n, T^2 L_{\max} \log(TL_{\max}) \log n\right) + m\right).$$

Proof. By Lemma 2.1, to achieve no-regret learning in T rounds, we simply need to perform δ -approximate sampling from the RWM distributions with learning rate $\log(1/\beta)$, where $\delta = \sqrt{k \log n / T}$ and $\beta = 1 - \sqrt{k \log n / T}$. For $T \geq C \cdot k \log n$ for a sufficiently large constant, we have $\log(1/\beta) = O(\sqrt{k \log n / T})$. Hence, $\log(1/\beta)/\delta = O(1)$. To show that we can perform the sampling process efficiently, we will apply Theorem 4.1 with $q = 2$ since Colonel Blotto is a resource allocation game whose reward function is always 2-piecewise. To do so, we need to construct and maintain the succinct descriptions of the cumulative battlefield reward function $\ell_h^{(t)}$ required by the sampling algorithm. Let $a_{j,h}^{(t)}$ be the number of soldiers that the j -th player assigns to the h -th battlefield at the t -th round. Then, for the first player, we essentially have

$$\ell_h^{(t)}(x) = w_h \cdot \sum_{t'=1}^{t-1} \mathbb{1}\left\{x > \max_{j=2 \dots m} a_{j,h}^{(t')}\right\}.$$

After observing the assignments from other players at round t , one first computes the maximum $v_h^{(t)} = \max_{j=2 \dots m} a_{j,h}^{(t)}$, which can be done in $O(m)$ time. After that, one essentially adds $\ell_h^{(t-1)}$ with the threshold function $\mathbb{1}\{x > v_h^{(t)}\}$, which takes time $O(\min(T, n))$, which is strictly dominated by the sampling time. Overall, the update just adds an additive factor of $O(Tkm)$ in total. \square

¹⁰In the zero-sum variant the other player also loses w_i .

This immediately gives an algorithm for computing ε -approximate CCEs in the Colonel Blotto Game (or Nash Equilibria when $m = 2$). Namely, we simulate no-regret playing for all m players simultaneously for $T = C \cdot kL_{\max}^2 \varepsilon^{-2} \log(mn/\eta)$ many rounds.

COROLLARY 4.2. (EQUILIBRIUM COMPUTATION FOR COLONEL BLOTTO GAMES) *Let \mathcal{I} be an m -player Colonel Blotto Game where the i -th player tries to assign n_i soldiers to k battlefields satisfying $n_i \leq n$. There exists an algorithm to compute an ε -approximate CCE (Nash if the game is two-player zero-sum) with probability at least $1 - \eta$ in time*

$$m \cdot O\left(\min\left(nk^2 L_{\max}^2 \varepsilon^{-2} \log(m/\eta) \cdot \log^2 n, k^4 L_{\max}^7 \varepsilon^{-6} \log^4(n) \cdot \log^3(m/\eta) \cdot \log\left(kL_{\max} \varepsilon^{-1} \log(mn/\eta)\right) \right) + mk^2 \varepsilon^{-2} L_{\max}^2 \log(mn/\eta) \right).$$

In the regime where n is far greater than k , ε , and L_{\max} , this gives an exponential improvement over prior algorithms [ADH⁺19] (at the cost of being approximate rather than exact).

Dice Game Dice Games are a randomized variant of Blotto first proposed in [DSDMDB06] where m -players¹¹ construct and roll dice, with the highest roller winning the game. These games are another natural instance of resource allocation. More formally, in the dice game each of m -players has a k_i -sided die and n_i points to distribute. The i -th player builds their die by assigning a number of dots to each face such that the sum is exactly n_i (where $k_i, n_i \in \mathbb{Z}^+$). Note that the action space A_i is then exactly the set of ordered partitions $P_{k_i}(n_i)$. The rewards are determined by “rolling” the m dice simultaneously; the player with the highest roll wins. In other words, for each player a face is selected uniformly at random and independently, and the player with more dots on the chosen faces wins and earns a reward of 1. The reward function is given by the expected reward of this process.

Let $x_{i,h}$ be the number of dots that the i -th player placed on the h -th face of her die for $i \in [m]$ and $h \in k_i$, and denote by X_i the random variable representing the number of dots obtained by player i after rolling. The expected utility for the first player is given by

$$\begin{aligned} \Pr[X_1 > \max(X_2, \dots, X_m)] &= \frac{1}{k_1} \sum_{h=1}^{k_1} \Pr[x_{1,h} > \max(X_2, \dots, X_m)] \\ &= \frac{1}{k_1} \sum_{h=1}^{k_1} \prod_{i=2}^m \Pr[x_{1,h} > X_i]. \\ &= \frac{1}{k_1} \sum_{h=1}^{k_1} \prod_{i=2}^m \frac{1}{k_i} \sum_{h'=1}^{k_i} \mathbb{1}(x_{1,h} > x_{i,h'}). \end{aligned}$$

Let $k = \max_i k_i$. It is not hard to see that $\frac{1}{k_i} \sum_{h'=1}^{k_i} \mathbb{1}(x_{1,h} > x_{i,h'})$ is an $O(k)$ -piecewise monotonic function and $\prod_{i=2}^m \frac{1}{k_i} \sum_{h'=1}^{k_i} \mathbb{1}(x_{1,h} > x_{i,h'})$ is an $O(mk)$ piecewise monotonic function. Hence, this is indeed a $O(mk)$ -piecewise monotonic resource allocation game. Applying our meta-algorithm immediately gives the following results on no-regret learning and equilibrium computation in dice games.

COROLLARY 4.3. (NO-REGRET LEARNING IN DICE GAMES) *Let $\mathcal{I} = \{(A_i, R_i)_{i=1}^m\}$ be an m player dice game such that $\max_i k_i \leq k$ and $\max_i n_i \leq n$. Then there is a no-regret learning algorithm for \mathcal{I} with regret*

$$\text{Reg}_T \leq O\left(L_{\max} \sqrt{T} \cdot \left(\sqrt{k \log(n)} + \sqrt{\log(1/\eta)} \right) \right)$$

with probability at least $1 - \eta$, and runs in time

$$Tk \cdot O\left(\min\left(T^2 \cdot \log(Tm) \cdot \log n, n \log n \right) + Tkm \right).$$

¹¹The original dice game is defined for two players. This can be nonetheless generalized naturally to an m -player setting.

Proof. Similar to the proof of Theorem 4.1, we will apply Theorem 4.1 with $q := mk$ and $L_{\max} = 1$. The main step is to maintain the succinct description of $\ell_h^{(t)}$. In dice games, the cumulative reward functions $\ell_h^{(t)}$ for all faces are identical and take the form

$$\ell^{(t)}(x) = \sum_{t'=1}^{t-1} \frac{1}{k_1} \prod_{i=2}^m \frac{1}{k_i} \sum_{h=1}^{k_i} \mathbb{1}(x > x_{i,h}^{(t')}).$$

After observing the actions $x_{i,h}^{(t)}$ for $i \in [m]$, $h \in [k_i]$ at the t -th round, we first compute the functions $v_i^{(t)}(x) = \frac{1}{k_i} \sum_{h=1}^{k_i} \mathbb{1}(x > x_{i,h}^{(t)})$, which are all at most k -piecewise. This requires sorting $x_{i,1}^{(t)}, \dots, x_{i,k_i}^{(t)}$, which takes time at most $O(k \log k)$. As a result, constructing all $v_2^{(t)}, \dots, v_m^{(t)}$ takes time in total $O(mk \log k)$. Then, we will point-wisely multiply all $v_i^{(t)}$ together. One can proceed in a divide and conquer manner: in the first pass, multiply together $v_i^{(t)}$ in groups of two, in the second pass, multiply together $v_i^{(t)}$ in groups of four and continues until all $v_i^{(t)}$ are multiplied together. There will be $\log m$ passes and the computation cost for each pass is at most $O(mk)$. Hence, the process takes time $O(mk \log m)$ in total. Then, we add the resulting function to $\ell^{(t-1)}$, which incurs another cost of $O(Tkm)$. Hence, in total, (assuming $\log m < T$), it takes time $O(Tkm)$ to update $\ell^{(t)}$ at one round. \square

COROLLARY 4.4. (EQUILIBRIUM FOR DICE GAMES) *Let $\mathcal{I} = \{(A_i, R_i)_{i=1}^m\}$ be a dice game with $n = \max(n_1, \dots, n_m)$, and $k = \max(k_1, \dots, k_m)$. There exists an algorithm to compute an ε -CCE (Nash if the game is two-player zero-sum) with probability at least $1 - \eta$ in time*

$$m \cdot O\left(\min\left(nk^2 \varepsilon^{-2} \log(m/\eta) \cdot \log^2 n, k^4 \varepsilon^{-6} \log^4(n) \cdot \log^3(m/\eta) \cdot \log(k\varepsilon^{-1} m \log(n/\eta)) \right) + mk^4 \varepsilon^{-4} \log^2(mn/\eta) \right).$$

4.4 Multi-resource allocation games A natural generalization of the resource allocation game is when each player has multiple resource types. This occurs naturally in many settings: a Colonel in Blotto, for instance, might have access to multiple unit types including troops, tanks, and planes (this variant was introduced in [BDD⁺17]). One would expect that the reward functions should vary depending on which types of units the Colonel chooses.

More formally, in the multi-resource allocation game, the i -th player has B_i types of fungible items. We denote by $n_{i,b}$ the number of type- b items that the i -th player possesses. Her strategy is an allocation of these items to k_i battlefields. We denote $X_{i,b,h}$ as the number of type b items that the i -th player assigns to the h -th battlefield. Similar to the single-resource allocation game, for each player i and each battlefield h , there is a battlefield reward function

$$r_{i,h} : ([0, n_{i,1}]) \times \dots \times ([0, n_{i,B}]) \times A_{-i} \mapsto \mathbb{R},$$

where we recall A_{-i} is the set of strategy tuples from the players other than i . Let $S \in A_{-i}$ be the strategies used by the other players, the total reward for the i -th player on strategy $X = \{X_{i,b,h}\}$ is given by summing over the rewards on each individual battlefield:

$$R_i(X, S) = \sum_{h=1}^{k_i} r_{i,h}(X_{i,1,h}, \dots, X_{i,B,h}, S).$$

We now prove a variant of Theorem 4.1 for the multi-resource setting.

THEOREM 4.2. (RWM IN MULTI-RESOURCE ALLOCATION GAME) *Let $\mathcal{I} = \{(A_i, R_i)_{i=1}^m\}$ be an m -player multi-resource allocation game where $A_i = P_{k_i}(n_{i,1}) \times \dots \times P_{k_i}(n_{i,B})$ for $n_1, \dots, n_B, k \in \mathbb{Z}^+$. Suppose the reward of the first player is bounded by L_{\max} , $\max_{i \in [m], b \in [B]} n_{i,b} = n$, $\max_{i \in [m]} k_i = k$. Then it is possible to implement (exact) RMW_β^T in time*

$$O\left(TkB(n+1)^{2B} + mTk^2B \cdot (n+1)^B \right).$$

assuming query access to $r_{i,h} : ([0, n_{i,1}]) \times \dots \times ([0, n_{i,B}]) \times A_{-i} \mapsto \mathbb{R}$.

Proof. We proceed to analyze the partition function of the RWM distribution. As usual, we will conduct the analysis from the first player’s perspective and drop the subscript used to index the player. Suppose the game is played for T rounds and the observed actions from the other players’ are $S^{(1)}, \dots, S^{(T)}$. Then, similar to Equation (4.6), for an assignment $x \in A_i$, we will define the cumulative reward

$$\ell_h^{(T+1)}(X_{1,h}, \dots, X_{b,h}) = \sum_{t=1}^T r_h(X_{1,h}, \dots, X_{B,h}, S^{(t)}).$$

For simplicity, we will abbreviate $(X_{1,h}, \dots, X_{b,h})$ as $X_{*,h}$. Then, we have the weight for the action x is simply

$$w_T(X) = \prod_{h=1}^k \beta_h^{\ell_h^{(T+1)}(X_{*,h})}.$$

After dropping the superscript marking the rounds, accordingly, the partition function is now $f_h : ([0, n])^B \mapsto \mathbb{R}^+$ for $h \in [k]$ defined as

$$f_h(y) = \sum_{X \in \mathcal{X}_y} \prod_{h=1}^k \beta_h^{\ell_h(X_{*,h})},$$

where $\mathcal{X}_y = \{X \in (\mathbb{Z}^+)^{B \times k} \mid \forall b \in [B], \sum_{h=1}^k X_{b,h} = y_b\}$. We still have the recursion

$$(4.14) \quad f_h(y) = \sum_{z \in \mathcal{Z}_y} f_h(y - z) \cdot \beta_h^{\ell_h(z)},$$

where $\mathcal{Z}_y = \{z \in ([0, n])^B \mid \forall b, z_b \leq y_b\}$.

Compared to the single-resource allocation game, there are $O(k \cdot (n + 1)^B)$ partition function values we need to compute. Using dynamic programming and the recursion stated in Equation (4.14), each of them now takes time at most $O(B(n + 1)^B)$. Hence, filling the entire DP table takes time $O(kB \cdot (n + 1)^{2B})$.

After that, we likewise sample the assignment for each battlefield sequentially. We will write $\mathbf{n} = (n_1, \dots, n_B)$. For the first battlefield, we sample

$$\Pr[X_{*,1} = y] \propto \beta_1^{\ell_1(y)} \cdot f_{k-1}(\mathbf{n} - y).$$

To sample from the $(h + 1)$ -st battlefield, one sample according to the distribution

$$\Pr[X_{*,h+1} = y \mid X_{*,1..h}] \propto \beta_{h+1}^{\ell_{h+1}(y)} \cdot f_{k-h-1}\left(\mathbf{n} - \left(\sum_{j=1}^h X_{*,j}\right) - y\right).$$

The domain size of the distributions we sample from is $O((n + 1)^B)$. To compute the probabilities of each element takes $O(B)$ times. Hence, the sampling time for one battlefield is $O(B \cdot (n + 1)^B)$. Hence, the runtime of the sampling process is dominated by that of computing the partition functions.

Finally, we discuss how we maintain the function $\ell_h^{(t)}$ at round t . To do that, we gather the strategy $S^{(t)}$ observed at the t -th round from other players and then query $r_h(z, S^{(t)})$ for each $z \in ([0, n])^B$, and add that to $\ell_h^{(t-1)}(z)$. Each query takes time $O(mBk)$ (to write down the input). In total, maintaining $\ell_h^{(t)}$ in T rounds takes time $O(mTk^2B(n + 1)^B)$. Adding this together with the time for computing the partition function then gives our final runtime. \square

As an immediate application, we get no-regret learning and equilibrium computation for the multi-resource allocation games such as the multi-resource Colonel Blotto problem.

COROLLARY 4.5. (MULTI-RESOURCE WITHOUT REGRET) Let $\mathcal{I} = \{(A_i, R_i)_{i=1}^m\}$ be an m -player multi-resource allocation game where $A_i = P_{k_i}(n_{i,1}) \times \cdots \times P_{k_i}(n_{i,B})$ for $n_1, \dots, n_B, k \in \mathbb{Z}^+$. Suppose the reward of the first player is bounded by L_{\max} , $\max_{i \in [m], b \in [B]} n_{i,b} = n$, $\max_{i \in [m]} k_i = k$. Then there is a no-regret learning algorithm for \mathcal{I} with regret:

$$\text{Reg}_T \leq O \left(L_{\max} \sqrt{T} \cdot \left(\sqrt{\sum_{b=1}^B \log(n_B)} + \sqrt{\log(1/\eta)} \right) \right)$$

with probability at least $1 - \eta$ that runs in time

$$O(TkB(n+1)^{2B} + mTk^2B \cdot (n+1)^B)$$

assuming query access to $r_{i,h} : ([0, n_{i,1}]) \times \cdots \times ([0, n_{i,B}]) \times A_{-i} \mapsto \mathbb{R}$.

COROLLARY 4.6. (EQUILIBRIUM COMPUTATION FOR MULTI-RESOURCE ALLOCATION GAMES) Let $\mathcal{I} = \{(A_i, R_i)_{i=1}^m\}$ be an m -player multi-resource allocation game where $A_i = P_{k_i}(n_{i,1}) \times \cdots \times P_{k_i}(n_{i,B})$ for $n_1, \dots, n_B, k \in \mathbb{Z}^+$. Suppose the reward of the first player is bounded by L_{\max} , $\max_{i \in [m], b \in [B]} n_{i,b} = n$, $\max_{i \in [m]} k_i = k$. There exists an algorithm to compute an ε -approximate CCE (Nash if two-player zero-sum) with probability at least $1 - \eta$ in time

$$O \left(m(n+1)^{2B} k^2 B^2 L_{\max}^2 \varepsilon^{-2} \log(mn/\eta) + m^2 (n+1)^B k^3 B^2 L_{\max}^2 \varepsilon^{-2} \cdot \log(mn/\eta) \right)$$

assuming query access to $r_{i,h} : ([0, n_{i,1}]) \times \cdots \times ([0, n_{i,B}]) \times A_{-i} \mapsto \mathbb{R}$.

Acknowledgements

We thank Thuy Duong Vuong, Nima Anari, and Kuikui Liu for helpful discussions on MCMC-based approaches for Dueling games and Colonel Blotto, and Saeed Seddighin for discussions on algorithms for Colonel Blotto and related problems. We also thank Maxwell Fishelson for enlightening discussion of optimistic hedge and its variants.

References

- [AASV21] Yeganeh Alimohammadi, Nima Anari, Kirankumar Shiragur, and Thuy-Duong Vuong, *Fractionally log-concave and sector-stable polynomials: counting planar matchings and more*, Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, 2021, pp. 433–446.
- [AB21] Alexandr Andoni and Daniel Beaglehole, *Learning to hash robustly, guaranteed*, arXiv preprint arXiv:2108.05433 (2021).
- [ABL14] Jean-Yves Audibert, Sébastien Bubeck, and Gábor Lugosi, *Regret in online combinatorial optimization*, Mathematics of Operations Research **39** (2014), no. 1, 31–45.
- [ADF⁺22] Ioannis Anagnostides, Constantinos Daskalakis, Gabriele Farina, Maxwell Fishelson, Noah Golowich, and Tuomas Sandholm, *Near-optimal no-regret learning for correlated equilibria in multi-player general-sum games*, Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing, 2022, pp. 736–749.
- [ADH⁺19] Amirmahdi Ahmadinejad, Sina Dehghani, Mohammadtaghi Hajiaghayi, Brendan Lucier, Hamid Mahini, and Saeed Seddighin, *From duels to battlefields: Computing equilibria of blotto and other games*, Mathematics of Operations Research **44** (2019), no. 4, 1304–1325.
- [AJK⁺21] Nima Anari, Vishesh Jain, Frederic Koehler, Huy Tuan Pham, and Thuy-Duong Vuong, *Entropic independence i: modified log-sobolev inequalities for fractionally log-concave distributions and high-temperature ising models*, arXiv preprint arXiv:2106.04105 (2021).
- [AL20] Vedat Levi Alev and Lap Chi Lau, *Improved analysis of higher order random walks and applications*, arXiv preprint arXiv:2001.02827 (2020).
- [ALO20] Nima Anari, Kuikui Liu, and Shayan Oveis Gharan, *Spectral independence in high-dimensional expanders and applications to the hardcore model*, arXiv preprint arXiv:2001.00303 (2020).
- [ALOV19] Nima Anari, Kuikui Liu, Shayan Oveis Gharan, and Cynthia Vinzant, *Log-concave polynomials ii: high-dimensional walks and an fpras for counting bases of a matroid*, Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, 2019, pp. 1–12.
- [Art94] W Brian Arthur, *Inductive reasoning and bounded rationality*, The American economic review **84** (1994), no. 2, 406–411.

- [ARV08] Heiner Ackermann, Heiko Röglin, and Berthold Vöcking, *On the impact of combinatorial structure on congestion games*, Journal of the ACM (JACM) **55** (2008), no. 6, 1–22.
- [Aum74] Robert J. Aumann, *Subjectivity and correlation in randomized strategies*, Journal of Mathematical Economics **1** (1974), no. 1, 67–96.
- [BAEJ20] Enric Boix-Adserà, Benjamin L. Edelman, and Siddhartha Jayanti, *The multiplayer colonel blotto game*, Proceedings of the 21st ACM Conference on Economics and Computation (2020).
- [BB19] Mourad Baïou and Francisco Barahona, *Faster algorithms for security games on matroids*, Algorithmica **81** (2019), no. 3, 1232–1246.
- [BBD⁺19] Soheil Behnezhad, Avrim Blum, Mahsa Derakhshan, MohammadTaghi Hajiaghayi, Christos H. Papadimitriou, and Saeed Seddighin, *Optimal strategies of blotto games: Beyond convexity*, 2019.
- [BDD⁺17] Soheil Behnezhad, Sina Dehghani, Mahsa Derakhshan, MohammadTaghi Hajiaghayi, and Saeed Seddighin, *Faster and simpler algorithm for optimal strategies of blotto game*, Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (2017), 369–375.
- [BHVW11] Petter Brändén, James Haglund, Mirkó Visontai, and David G Wagner, *Proof of the monotone column permanent conjecture*, Notions of Positivity and the Geometry of Polynomials, Springer, 2011, pp. 63–78.
- [Bor53] Emile Borel, *The theory of play and integral equations with skew symmetric kernels*, Econometrica **21** (1953), no. 1, 97.
- [BR21] Yakov Babichenko and Aviad Rubinfeld, *Settling the complexity of nash equilibrium in congestion games*, Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, 2021, pp. 1426–1437.
- [Bri88] E Oran Brigham, *The fast fourier transform and its applications*, Prentice-Hall, Inc., 1988.
- [BŠVV04] Ivona Bezáková, Daniel Štefankovic, Vijay V Vazirani, and Eric Vigoda, *Approximating the permanent in $o(n^7)$ time*.
- [CBL06] Nicolo Cesa-Bianchi and Gábor Lugosi, *Prediction, learning, and games*, Cambridge university press, 2006.
- [CBL12] ———, *Combinatorial bandits*, Journal of Computer and System Sciences **78** (2012), no. 5, 1404–1422.
- [CGM19] Mary Cryan, Heng Guo, and Giorgos Mousa, *Modified log-sobolev inequalities for strongly log-concave distributions*, 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2019, pp. 1358–1370.
- [CHK17] Alon Cohen, Tamir Hazan, and Tomer Koren, *Tight bounds for bandit combinatorial optimization*, Conference on Learning Theory, PMLR, 2017, pp. 629–642.
- [CLV21] Zongchen Chen, Kuikui Liu, and Eric Vigoda, *Optimal mixing of glauber dynamics: Entropy factorization via high-dimensional expansion*, Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, 2021, pp. 1537–1550.
- [CPV95] Charles J Colbourn, J Scott Provan, and Dirk Vertigan, *The complexity of computing the tutte polynomial on transversal matroids*, Combinatorica **15** (1995), no. 1, 1–10.
- [DFG21] Constantinos Daskalakis, Maxwell Fishelson, and Noah Golowich, *Near-optimal no-regret learning in general games*, Advances in Neural Information Processing Systems **34** (2021), 27604–27616.
- [DGP06] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou, *The complexity of computing a nash equilibrium*, Proceedings of the thirty-eighth annual ACM symposium on Theory of computing - STOC 06 (2006).
- [DGP09] Constantinos Daskalakis, Paul W Goldberg, and Christos H Papadimitriou, *The complexity of computing a nash equilibrium*, SIAM Journal on Computing **39** (2009), no. 1, 195–259.
- [DJM17] Martin Dyer, Mark Jerrum, and Haiko Müller, *On the switch markov chain for perfect matchings*, Journal of the ACM (JACM) **64** (2017), no. 2, 1–33.
- [DSDMDB06] Bart De Schuymer, Hans De Meyer, and Bernard De Baets, *Optimal strategies for equal-sum dice games*, Discrete applied mathematics **154** (2006), no. 18, 2565–2576.
- [Eck87] Roger Eckhardt, *Stan ulam, john von neumann, and the monte carlo method*, Los Alamos Science **15** (1987), no. 131-136, 30.
- [FLLK22] Gabriele Farina, Chung-Wei Lee, Haipeng Luo, and Christian Kroer, *Kernelized multiplicative weights for 0/1-polyhedral games: Bridging the gap between learning in extensive-form and normal-form games*, arXiv preprint arXiv:2202.00237 (2022).
- [FS96] Yoav Freund and Robert E. Schapire, *Game theory, on-line prediction and boosting*, Proceedings of the ninth annual conference on Computational learning theory - COLT '96 (1996).
- [FS99] ———, *Adaptive game playing using multiplicative weights*, Games and Economic Behavior **29** (1999), no. 1-2, 79–103.
- [GP09] Russell Golman and Scott E. Page, *General blotto: Games of allocative strategic mismatch*, Public Choice **138** (2009), no. 3/4, 279–299.
- [Gro50] Oliver Gross, *The symmetric blotto game [1]*, Rand Corporation, 1950.
- [GW50] Oliver Gross and Robert Wagner, *A continuous colonel blotto game*, Tech. report, Rand Project Air Force Santa Monica Ca, 1950.
- [Har08] Sergiu Hart, *Discrete colonel blotto and general lotto games*, International Journal of Game Theory **36** (2008), no. 3, 441–460.
- [HKW09] David P Helmbold and Manfred K Warmuth, *Learning permutations with exponential weights.*, Journal of Machine

- Learning Research **10** (2009), no. 7.
- [IKL⁺11] Nicole Immerlica, Adam Tauman Kalai, Brendan Lucier, Ankur Moitra, Andrew Postlewaite, and Moshe Tennenholtz, *Dueling algorithms*, Proceedings of the forty-third annual ACM symposium on Theory of computing, 2011, pp. 215–224.
- [Ins16] Insider, *scientists say they can predict two-party outcomes after solving a 95-year-old game theory problem.*, 2016.
- [JSV04] Mark Jerrum, Alistair Sinclair, and Eric Vigoda, *A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries*, Journal of the ACM (JACM) **51** (2004), no. 4, 671–697.
- [JV86] Mark R Jerrum, Leslie G Valiant, and Vijay V Vazirani, *Random generation of combinatorial structures from a uniform distribution*, Theoretical computer science **43** (1986), 169–188.
- [Kle05] Robert David Kleinberg, *Online decision problems with large strategy sets*, Ph.D. thesis, Massachusetts Institute of Technology, 2005.
- [KO20] Tali Kaufman and Izhar Oppenheim, *High order random walks: Beyond spectral gap*, Combinatorica (2020), 1–37.
- [KPT09] Robert Kleinberg, Georgios Piliouras, and Éva Tardos, *Multiplicative updates outperform generic no-regret learning in congestion games*, Proceedings of the forty-first annual ACM symposium on Theory of computing, 2009, pp. 533–542.
- [KR12] Dan Kovenock and Brian Roberson, *Coalitional colonel blotto games with application to the economics of alliances*, Journal of Public Economic Theory **14** (2012), no. 4, 653–676.
- [KR21] ———, *Generalizations of the general lotto and colonel blotto games*, Economic Theory **71** (2021), no. 3, 997–1032.
- [Kva07] Dmitriy Kvasov, *Contests with limited resources*, Journal of Economic Theory **136** (2007), no. 1, 738–748.
- [KWK⁺10] Wouter M Koolen, Manfred K Warmuth, Jyrki Kivinen, et al., *Hedging structured concepts.*, COLT, Citeseer, 2010, pp. 93–105.
- [Las02] Jean-François Laslier, *How two-party competition treats minorities*, Review of Economic Design **7** (2002), no. 3, 297–307.
- [LP02] Jean-François Laslier and Nathalie Picard, *Distributive politics and electoral competition*, Journal of Economic Theory **103** (2002), no. 1, 106–130.
- [LW94] Nick Littlestone and Manfred K Warmuth, *The weighted majority algorithm*, Information and computation **108** (1994), no. 2, 212–261.
- [LY94] Richard J Lipton and Neal E Young, *Simple strategies for large zero-sum games with applications to complexity theory*, Proceedings of the twenty-sixth annual ACM symposium on Theory of computing, 1994, pp. 734–740.
- [Mew16] Mewright, *Umd-led team first to solve well-known game theory scenario*, Feb 2016.
- [MM15] Scott T Macdonell and Nick Mastronardi, *Waging simple wars: a complete characterization of two-battlefield blotto equilibria*, Economic Theory **58** (2015), no. 1, 183–216.
- [Nas51] John Nash, *Non-cooperative games*, Annals of mathematics (1951), 286–295.
- [OvS13] Georg Ostrovski and Sebastian van Strien, *Payoff performance of fictitious play*, arXiv preprint arXiv:1308.4049 (2013).
- [PRG22] Vianney Perchet, Philippe Rigollet, and Thibaut Le Gouic, *An algorithmic solution to the blotto game using multi-marginal couplings*, 2022.
- [Rob06] Brian Roberson, *The colonel blotto game*, Economic Theory **29** (2006), no. 1, 1–24.
- [Ros73] Robert W Rosenthal, *A class of games possessing pure-strategy nash equilibria*, International Journal of Game Theory **2** (1973), no. 1, 65–67.
- [RS13] Sasha Rakhlin and Karthik Sridharan, *Optimization, learning, and games with predictable sequences*, Advances in Neural Information Processing Systems **26** (2013).
- [Rub16] Aviad Rubinfeld, *Settling the complexity of computing approximate two-player nash equilibria*, 2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS), IEEE, 2016, pp. 258–265.
- [Sch96] James C Schatzman, *Accuracy of the discrete fourier transform and the fast fourier transform*, SIAM Journal on Scientific Computing **17** (1996), no. 5, 1150–1166.
- [SLS14] Galina Schwartz, Patrick Loiseau, and Shankar S Sastry, *The heterogeneous colonel blotto game*, 2014 7th international conference on NETwork Games, COntrol and OPTimization (NetGCoop), IEEE, 2014, pp. 232–238.
- [Sze17] Dávid Szeszlér, *Security games on matroids*, Mathematical Programming **161** (2017), no. 1, 347–364.
- [Tam11] Milind Tambe, *Security and game theory: algorithms, deployed systems, lessons learned*, Cambridge university press, 2011.
- [Tho17] Caroline Thomas, *N-dimensional blotto game with heterogeneous battlefield values*, Economic Theory **65** (2017), no. 3, 509–544.
- [Tho18] ———, *N-dimensional blotto game with heterogeneous battlefield values*, Economic Theory **65** (2018), no. 3, 509–544.
- [VA06] Berthold Vöcking and R Aachen, *Congestion games: Optimization in competition.*, ACiD, 2006, pp. 9–20.
- [VL21] Dong Quan Vu and Patrick Loiseau, *Colonel blotto games with favoritism: Competitions with pre-allocations and asymmetric effectiveness*, 2021.
- [VLS18] Dong Quan Vu, Patrick Loiseau, and Alonso Silva, *Efficient computation of approximate equilibria in discrete colonel blotto games*, Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence (2018).
- [VLS19] ———, *Combinatorial bandits for sequential learning in colonel blotto games*, 2019 IEEE 58th Conference on Decision

and Control (CDC), IEEE, 2019, pp. 867–872.

[WK08] Manfred K Warmuth and Dima Kuzmin, *Randomized online pca algorithms with regret bounds that are logarithmic in the dimension*, Journal of Machine Learning Research **9** (2008), no. Oct, 2287–2320.

A No-regret Learning and Equilibrium Computation

LEMMA A.1. (LEMMA 4.1 FROM [CBL06] REPHRASED) Let $\mathcal{I} = \{(A_i, R_i)_{i=1}^m\}$ be an m -player game played repeatedly for T rounds. Denote $\mathbf{s}_i^{(t)}$ as the mixed strategy chosen by the i -th player, and $a_i^{(t)} \sim \mathbf{s}_i^{(t)}$ as the action sampled. Assume the i -th player follows an algorithm which computes $\mathbf{s}_i^{(t)}$ solely based on $a_j^{(t')}$ for $t' < t$ and $j \in [m] \setminus \{i\}$. Furthermore, suppose the following is true

$$\sup_{b_j^{(t)} \in A_j \text{ for } t \in [T], j \neq i} \mathbb{E} \left[\max_{e \in A_i} \sum_{t=1}^T R_i(e, b_{-i}^{(t)}) - \sum_{t=1}^T R_i(a_i^{(t)}, b_{-i}^{(t)}) \right] \leq B.$$

Then, for all $\delta \in (0, 1)$, with probability at least $1 - \delta$, it holds

$$\max_{e \in A_i} \sum_{t=1}^T R_i(e, a_{-i}^{(t)}) - \sum_{t=1}^T R_i(a_i^{(t)}, a_{-i}^{(t)}) \leq B + L_{\max} \sqrt{T/2 \log(1/\delta)}.$$

at the end of the repeated play for the i -th player.

Proof. [Proof of Lemma 2.1] If one follows exactly from the Randomized Weighted Majority algorithm, one has the guarantee that

$$\sup_{b_j^{(t)} \in A_j \text{ for } t \in [T], j \neq i} \mathbb{E}_{a_i^{(t)} \sim RM_i^{(t)}(\beta)} \left[\max_{e \in A_i} \sum_{t=1}^T R_i(e, b_{-i}^{(t)}) - \sum_{t=1}^T R_i(a_i^{(t)}, b_{-i}^{(t)}) \right] \leq L_{\max} \sqrt{T \log N}.$$

In reality, since we are performing κ -approximate sampling, we have $\text{TV}(RM_i^{(t)}(\beta), \mathbf{s}_i^{(t)}) \leq \kappa$. Since the reward of the game is bounded by L_{\max} , we have

$$\sup_{b_j^{(t)} \in A_j \text{ for } t \in [T], j \neq i} \mathbb{E}_{a_i^{(t)} \sim \mathbf{s}_i^{(t)}} \left[\max_{e \in A_i} \sum_{t=1}^T R_i(e, b_{-i}^{(t)}) - \sum_{t=1}^T R_i(a_i^{(t)}, b_{-i}^{(t)}) \right] \leq L_{\max} \sqrt{T \log N} + \kappa \cdot L_{\max} \cdot T.$$

Then, by Lemma A.1, it then holds

$$\frac{1}{T} \max_{e \in A_i} \sum_{t=1}^T R_i(e, a_{-i}^{(t)}) - \frac{1}{T} \sum_{t=1}^T R_i(a_i^{(t)}, a_{-i}^{(t)}) \leq L_{\max} \sqrt{\log N/T} + \kappa \cdot L_{\max} + L_{\max} \sqrt{\log(1/\delta)/(2T)}.$$

with probability at least $1 - \delta$. Setting $T = L_{\max}^2 \varepsilon^{-2} \log(N/\delta)$ and $\kappa = \min(1/2, \varepsilon)$ then gives the average regret is bounded by $O(\varepsilon)$. \square

Proof. [Proof of Corollary 2.1] By Lemma 2.1, if we set $T = C \cdot (L_{\max}^2 \varepsilon^{-2} \log(Nm/\eta))$, $\delta = \varepsilon/(CL_{\max})$ for a sufficient large constant and simulate the repeated game playing for T rounds where each player makes her decision based on δ -RWM, the regret of the i -th player is bounded by ε with probability at least $1 - \delta/m$. By union bound, this holds for all players simultaneously. This results then follows from Theorem 2.1. \square

B Bit-complexity and Stability of Numeric Operations

To avoid un-necessary technical details of the bit-complexities of numbers and time complexities of algebraic operations, the algorithmic results in the main body of this work are stated in the *Algebraic Computation Model*. In particular, we assume additions, subtractions, multiplications, divisions, exponentiation and comparisons can be carried between real

numbers in constant time, and that the computing device has query access to the digits of the real numbers. We remark that, for all the games we study, if the rewards are rational numbers with bounded bit complexities, our algorithms can all be implemented exactly in the RAM computation model with their runtime increased by at most polynomial factors. Unsurprisingly, if one is more careful with the numeric precision needed and maintains only multiplicative approximations of each algebraic operations, our algorithms can be implemented in the RAM model losing only poly-logarithmic factors. In this section, we discuss some standard techniques to this end for the reader interested in any implementation of our algorithms.

We first discuss how we represent and perform algebraic operations on numbers whose absolute values are exponentially large or small. While writing these numbers down exactly is costly, for the purpose of δ -approximate sampling, it is actually sufficient to keep $\text{poly}(\delta)$ multiplicative approximations of these numbers. Fortunately, these approximations can indeed be represented much more succinctly using the scientific notations. For convenience, for a number $a \in \mathbb{R}^+$, we will call \tilde{a} a δ -approximation of a if we have $(1 - \delta)a \leq \tilde{a} \leq (1 + \delta)a$ and a one-sided δ -approximation if we have $(1 - \delta)a \leq \tilde{a} \leq a$.

FACT B.1. *Given $a \in \mathbb{R}^+$ satisfying $\exp(-q) \leq a \leq \exp(q)$ for $q \in \mathbb{Z}^+$, let \tilde{a} be written in scientific form keeping $\Theta(\log(1/\delta))$ many significant figures. Then, \tilde{a} is a one-sided δ -approximation of a and can be represented using $\Theta(\log(\delta^{-1}) + \log q)$ many bits.*

Instead of performing exact arithmetic computations, we can perform ‘approximate’ arithmetic operations on real numbers in all our algorithms.

CLAIM B.1. *Let $a, b \in \mathbb{R}^+$ be two numbers in scientific notations with s significant figures.*

- *one-sided δ -approximation of additions and multiplications can be performed in time $O(s)$ and $O(s \log s)$ respectively.*
- *δ -approximation of division can be performed in time $O(s + \log(1/\delta))$.*
- *Given $1 \leq \alpha \leq 2$ that has s significant figures and $i \in \mathbb{Z}^+ \cup \{0\}$, a one-side δ -approximation of α^i can be computed in time $O(s \cdot \log^2 i \cdot \log(\delta^{-1}))$.*

Proof. The first two claims follow from the definition of (one-sided) δ -approximation. We proceed to show that one can perform approximate exponentiation efficiently. In particular, we argue that α^i can be computed fairly accurately via fast exponentiation while keeping $C \cdot \log(i/\delta)$ significant figures throughout the computation for some large enough constant C . By doing so, we can make sure the approximation to α is a one-sided ξ -approximation where $\xi = \delta/i^c$ for some large enough constant c . Consequently, the approximation of β^j for any $j \in [i]$ that is a power of 2 is within $(1 \pm 4^{\log_2(i)} \cdot \xi)$. It then follows α^i can be approximated within $(1 \pm \xi \cdot O(\log i) \cdot 4^{\log_2(i)}) = (1 \pm \delta)$ when c is sufficiently large. \square

Unsurprisingly, the output of applying a series of arithmetic operations will be within multiplicative factors of the result obtained by replacing each operation with its approximate counterpart.

FACT B.2. *Given a variable y that is the result of V arithmetic operations including Addition, Multiplication and Division on the inputs $x_1, \dots, x_n \in \mathbb{R}^+$ in scientific notations with s significant figures, let \hat{y} be the variable obtained by replacing all the arithmetic operations with their $\delta/(10V)$ -approximate counterparts for small enough δ . Then, \hat{y} will be a δ -approximation of y . Moreover, if only additions and multiplications are used, the approximation is one-sided i.e. $\hat{y} \leq y$.*

Careful readers may find that subtraction is excluded when we discuss approximate algebraic operations in Claim B.1. For two numbers $a, b \in \mathbb{R}^+$ and \hat{a}, \hat{b} be their δ approximation counterparts, $\hat{a} - \hat{b}$ may be wildly different from $a - b$ when a is substantially larger than b . Yet, subtraction between real numbers is indeed used in two different places. Firstly, subtractions occur in Discrete Fourier Transform, which is used in Proposition 4.2 to compute the convolution between functions. The numeric stability of DFT varies among different implementations and depends on a number of subtle factors (See [Sch96]). If numeric stability indeed becomes an issue in the actual implementation, one can fallback to evaluate the convolution in the brute-force manner, which increases the complexity from $O(nk \log n)$ to $O(n^2k)$.

Another place where subtractions are used is in Algorithm 2 to compute range sum of piece-wise function efficiently. As such, we need a numerically more stable technique for performing range sum query in place of the prefix sum technique. In particular, given the succinct description of a q -piecewise constant function $f : \{0\} \cup [n] \mapsto \mathbb{R}^+$ and $\delta \in (0, 1)$ beforehand, we want to perform some preprocessing in time $q \cdot \text{polylog}(n, q, \delta^{-1})$ and then answer a series of queries of the form $\sum_{i=a}^b f(i)$ within $(1 \pm \delta)$ multiplicative factors in time $\text{polylog}(n, q, \delta^{-1})$.

CLAIM B.2. *Given the succinct description D_f of a q -piecewise constant function $f : \{0\} \cup [n] \mapsto \mathbb{R}$ where the function values contain at most s significant figures, there exists an algorithm **Range-Sum-Query** which performs some preprocess in time $O(qs) + \text{polylog}(q, \delta^{-1}, n)$, and can compute one-sided δ -approximation to query of the form $\sum_{i=a}^b f(i)$ in time $\text{polylog}(n, q, \delta^{-1})$.*

Proof. Let $D_f = \{(a_1, b_1, y_1), \dots, (a_q, b_q, y_q)\}$. In the preprocessing step, we first compute the range sum of all intervals $[a_i, b_i]$ (approximately). Denote the results as an array $[s_1, \dots, s_q]$. Then, we build a segment tree with the array, where nodes store the approximate range sum of intervals of lengths that are powers of 2. This takes $O(q \log(q))$ arithmetic operations. The data structure then allows us to answer $\sum_{i=a}^b s_i$ with $O(\log(q))$ arithmetic operations with enough accuracy. Then, when we receives a query $\sum_{i=a}^b f(i)$. We first binary search for the intervals of f that a, b fall into respectively. Next, one uses the pre-built segment tree to answer the range sum of any intervals that are strictly contained in $[a, b]$ and then adds the sum of remaining elements. It is then not hard to see that the pre-processing step takes time at most $O(qs) + \text{polylog}(q, \delta^{-1}, n)$, and answering each query takes time $\text{polylog}(q, \delta^{-1}, n)$. \square

Lastly, we discuss the building block of sampling: sampling from multinomial distributions. Typically, our algorithm computes a vector w_1, \dots, w_n and then samples from the multinomial distribution X where $\Pr[X = i] \propto w_i$. This is simple in the Algebraic computation model as one can easily reduce this to sampling from uniform distributions over real intervals (which can be done at assumed unit cost). In particular, one first computes the prefix sum W_1, \dots, W_n . Then, one samples z from the uniform distribution over the interval $[0, W_n]$ and returns j for $W_{j-1} < z < W_j$. This clearly takes at most time $O(n)$. In the bit-complexity model, we can nonetheless achieve approximate sampling from arbitrary multinomial distributions with similar runtime.

CLAIM B.3. *Given a weight vector (w_1, \dots, w_n) in scientific notations with s significant figures, δ -approximate sampling from the multinomial distribution X such that $\Pr[X = i] \propto w_i$ can be done in time $\tilde{O}(n \cdot (\log(\delta^{-1}) + s))$.*

Proof. After reading the input, one first truncates to make sure each w_i has at most $O(\log(n/\delta))$ many significant figures as that is already enough for the specified ‘sampling accuracy’. After that, all arithmetic operations will be carried out with their $c \cdot \delta/n$ approximations for some sufficiently small constant c . In the next step, one normalizes the weight vector and rounds each w_i to their nearest multiple of $c \cdot \delta/n$. Doing so changes the distribution by at most $c \cdot \delta$ in total variation distance. One can then multiply all w_i by a factor of $n/(c \cdot \delta)$ to make everything an integer. Finally, one can do the same thing as sampling in the Algebraic computation model: computing the prefix sums and reducing the problem to sampling from uniform distributions, now over integer intervals. The integers in the interval can be at most $n^2/(c \cdot \delta)$ so the runtime is dominated by the preliminary computations performed. \square

We note that in many cases we actually require a slightly more complicated sampling procedure where we wish to sample from a q -piecewise support- n multinomial with $q \ll n$. This can be done similarly in time $\tilde{O}((q + \log(n))(\log(\delta^{-1}) + s))$ by first sampling one of the q piecewise intervals by the above technique, then sampling uniformly within the interval.

C Implementing Glauber Dynamics

This section is devoted to proving Theorem 3.2, which we repeat here for convenience.

THEOREM C.1. (RWM ON MATROIDS) *Let $\mathcal{I} = \{(A_i, R_i)_{i=1}^m\}$ be an m -player game on a size- n ground set Ω . If A_i consists of the bases of a rank- k matroid, is linear, and is collision-sensitive with support q , then it is possible to implement δ -RWM $_{\beta}^T$ in time*

$$O\left(k_i T(\text{CO} + q \log(n) + mk_{-i} T \log(n)) \log\left(\frac{k_i \log(n) + L_{\max} T \log(\beta^{-1})}{\delta}\right)\right),$$

assuming access to a q -piecewise succinct description of NC_i encoded under an ordering of Ω and a contraction oracle matching the same ordering.

Proof. Since A_i is given by the bases of a rank- k matroid, the single-step Glauber Dynamics on Ω^w mix in time

$$T(\text{GD}(\Omega^w), \delta) \leq O\left(k \log\left(\frac{\log(|A_i|/w_*)}{\delta}\right)\right)$$

for any external field $w \in \mathbb{R}_+^\Omega$. In our setting, we have that $|A_i| \leq n^k$, and $w_* \geq \beta^{L_{\max}T}$, and the process needs to be repeated once per round bringing the complexity to:

$$O\left(kT \log\left(\frac{k \log(n) + L_{\max}T \log(\beta^{-1})}{\delta}\right)\right)$$

times the implementation time of a single step of Glauber dynamics.

It is left to bound this cost. To implement a step of GD in the t -th round of optimistic hedge, we first remove a uniformly random element from our current basis, then re-sample from the conditional distribution. The first step can easily be implemented in $O(\log(k))$ time. The latter step requires more care. Let \hat{e} denote the $(k-1)$ -size set resulting from the down-step of the walk. Query the contraction oracle on \hat{e} and call the resulting set $S_e \subset E$. Notice that by definition, the conditional measure of any $x \in S_e$ is proportional to $\beta^{-r_t(x)}$ where

$$r_t(x) = \sum_{j=1}^t R_j^\Omega(x, s^{(j)}).$$

Thus to perform the conditional sampling efficiently, it is sufficient to compute the external field for each element in S_e and sample from the corresponding multinomial distribution.

While implementing this naively would require time at least $|S_e|$ to check the weight of each element in the conditional distribution, this can be circumvented via our assumption that our game is collision-sensitive with bounded support. In particular, assume for the moment we have access to a succinct description for the vertex-wise total rewards $r_t(v)$ that is $(q + tmk_{-i})$ -piecewise, and that the output of the contraction oracle respects the order of the description (we will argue this can be constructed efficiently shortly). As a result, the total rewards in S_e are $(q + tmk_{-i})$ -piecewise as well. This means that using query access to CO,¹² we can build a succinct description for total rewards on the elements in $|S_e|$ (labeled by their index in CO). Sampling from the corresponding multinomial distribution in the algebraic computation model then takes $O((q + tmk_{-i}))$ time, and one can then feed the sampled index into CO to receive the correct vertex. Altogether, a single step of GD can therefore be implemented in $O(\text{CO} + (q + tmk_{-i}) \log(n))$ time assuming access to the appropriate description of total rewards.

Finally, we argue we can construct and maintain the succinct descriptions of the vertex-wise reward functions over T rounds efficiently. Recall we start with an q -piecewise succinct description for the no-collision vertex-wise reward values. In each round, at most mk_{-i} new elements of Ω are introduced into the history, and since the game is collision-sensitive the resulting succinct description of rewards is at most $(q + tmk_{-i})$ -piecewise in the t -th round as desired. The computational cost stems from noting that it is actually sufficient just to update the rewards for vertices which have appeared in the opponent history (and the number of rounds in which it has appeared). During look-up, computing the total reward for any vertex v that has appeared t times can be computed in $O(1)$ time by simply adding the stored value $(T-t)R_i(v, s)$ for any $s \not\ni v$. The cost of building the succinct description is therefore asymptotically dominated by the sampling procedure above, which gives the final complexity. \square

¹²Formally we are also assuming here one has query access to the size of the output of the contraction oracle. Note this can be easily implemented in $\text{polylog}(n)$ time even if one does not assume such access.