

Exponential Weights on the Hypercube in Polynomial Time

Sudeep Raja Putta

SPUTTA@CS.UMASS.EDU

College of Information and Computer Sciences

University of Massachusetts Amherst

140 Governors Dr, Amherst, MA 01002, USA

Abstract

We study a general online linear optimization problem (OLO). At each round, a subset of objects from a fixed universe of n objects is chosen, and a linear cost associated with the chosen subset is incurred. We use regret as a measure of performance of our algorithms. Regret is the difference between the total cost incurred over all iterations and the cost of the best fixed subset in hindsight. We consider Full Information, Semi-Bandit and Bandit feedback for this problem. Using characteristic vectors of the subsets, this problem reduces to OLO on the $\{0, 1\}^n$ hypercube. The Exp2 algorithm and its bandit variants are commonly used strategies for this problem. It was previously unknown if it is possible to run Exp2 on the hypercube in polynomial time.

In this paper, we present a polynomial time algorithm called PolyExp for OLO on the hypercube. We show that our algorithm is equivalent to both Exp2 on $\{0, 1\}^n$ as well as Online Mirror Descent (OMD) with Entropic regularization on $[0, 1]^n$ and Bernoulli Sampling. We consider L_∞ adversarial losses. We show PolyExp achieves expected regret bounds that are a factor of \sqrt{n} better than Exp2 in all the three settings. Because of the equivalence of these algorithms, this implies an improvement on Exp2's regret bounds.

Keywords: Online Learning, Bandits, Exponential Weights, Hedge, Online Mirror Descent, EXP2, Online Stochastic Mirror Descent

1. Introduction

Consider the following abstract game which proceeds as a sequence of T rounds. In each round t , a player has to choose a subset S_t from a universe U of n objects. Without loss of generality, assume $U = \{1, 2, \dots, n\} = [n]$. Each object $i \in U$ has an associated loss $c_{t,i}$, which are unknown to the player and may be chosen by an adversary. On choosing S_t , the player incurs the cost $c_t(S_t) = \sum_{i \in S_t} c_{t,i}$. In addition the player receives some feedback about the costs of this round. The goal of the player is to choose the subsets such that the total cost incurred over a period of rounds is close to the total cost of the best subset in hindsight. This difference in costs is called the *regret* of the player. Formally, regret is defined as:

$$R_T = \sum_{t=1}^T c_t(S_t) - \min_{S \subseteq U} \sum_{t=1}^T c_t(S)$$

We can re-formulate the problem as follows. The 2^n subsets of U can be mapped to the vertices of the $\{0, 1\}^n$ hypercube. The vertex corresponding to the set S is represented by its characteristic vector $X(S) = \sum_{i=1}^n 1\{i \in S\}e_i$. From now on, we will work with the hypercube instead of sets and use losses $l_{t,i}$ instead of costs. In each round, the player chooses $X_t \in \{0, 1\}^n$. The loss vector l_t is chosen by an adversary and is unknown to

the player. The loss of choosing X_t is $X_t^\top l_t$. The player receives some feedback about the loss vector. The goal is to minimize regret, which is now defined as:

$$R_T = \sum_{t=1}^T X_t^\top l_t - \min_{X \in \{0,1\}^n} \sum_{t=1}^T X^\top l_t$$

This is the *Online Linear Optimization (OLO)* problem on the hypercube. As the loss vector l_t can be set by an adversary, the player has to use some randomization in its decision process in order to avoid being foiled by the adversary. At each round $t = 1, 2, \dots, T$, the player chooses an action X_t from the decision set $\{0, 1\}^n$, using some internal randomization. Simultaneously, the adversary chooses a loss vector l_t , without access to the internal randomization of the player. Since the player's strategy is randomized and the adversary could be adaptive, we consider the expected regret of the player as a measure of the player's performance. Here the expectation is with respect to the internal randomization of the player and eventually the adversary's randomization if it is an adaptive adversary.

So far, we haven't specified the feedback that the player receives. There could be three kinds of feedback for this problem:

1. *Full Information setting*: At the end of each round t , the player observes the loss vector l_t .
2. *Semi Bandit setting*: At the end of each round t , the player observes the losses of the objects chosen, ie, $l_{t,i} X_{t,i}$ for $i \in [n]$.
3. *Bandit setting*: At the end of each round t , the player only observes the scalar loss incurred $X_t^\top l_t$.

In order to make make quantifiable statements about the regret of the player, we need to restrict the loss vectors the adversary may choose. The two cases considered in the online learning literature are:

1. L_∞ assumption: here we assume that $\|l_t\|_\infty \leq 1$ for all t .
2. L_2 assumption: here we assume that $|X^\top l_t| \leq 1$ for all t and all $X \in \{0, 1\}^n$.

As our decision set contains the vector $\mathbf{1}_n$, the L_2 and L_∞ assumptions are equivalent upto a scaling factor by n . So, we only consider the L_∞ assumption.

There are three major strategies for online optimization, which can be tailored to the problem structure and type of feedback. Although, these can be shown to be equivalent to each other in some form, not all of them may be efficiently implementable. These strategies are:

1. Exponential Weights(EW)(Freund and Schapire, 1997; Littlestone and Warmuth, 1994)
2. Follow the Leader(FTL)(Kalai and Vempala, 2005)
3. Online Mirror Descent(OMD) (Nemirovsky and Yudin, 1983)

For problems of this nature, a commonly used EW type algorithm is Exp2 (Audibert et al., 2011, 2013; Bubeck et al., 2012). For the specific problem of Online Linear Optimization on the hypercube, it was previously unknown if the Exp2 algorithm can be efficiently implemented (Bubeck et al., 2012). So, previous works have resorted to using OMD algorithms for problems of this kind. The main reason for this is that Exp2 explicitly maintains a probability distribution on the decision set. In our case, the size of the decision set is 2^n . So a straightforward implementation of Exp2 would need exponential time and space.

1.1 Our Contributions

We use a key observation: In the case of linear losses the probability distribution of Exp2 can be factorized as a product of n Bernoulli distributions. Using this fact, we design an efficient polynomial time algorithm called *PolyExp* for sampling sampling from and updating these distributions.

We show that PolyExp is equivalent to Exp2. In addition, we show that PolyExp is equivalent to OMD with entropic regularization and Bernoulli sampling. This allows us to analyze PolyExp’s using powerful analysis techniques of OMD.

Proposition 1 *For the Online Linear Optimization problem on the $\{0, 1\}^n$ Hypercube, Exp2, OMD with Entropic regularization and Bernoulli sampling, and PolyExp are equivalent.*

This kind of equivalence is rare. To the best of our knowledge, the only other scenario where this equivalence holds is on the probability simplex for the so called experts problem.

In our paper, we focus on the L_∞ assumption. Directly analyzing Exp2 gives regret bounds different from PolyExp. In fact, PolyExp’s regret bounds are a factor of \sqrt{n} better than Exp2. These results are summarized by the table below.

L_∞			
	Full Information	Semi-Bandit	Bandit
Exp2 (direct analysis)	$O(n^{3/2}\sqrt{T})$	$O(n^{3/2}\sqrt{T})$	$O(n^2\sqrt{T})$
PolyExp	$O(n\sqrt{T})$	$O(n\sqrt{T})$	$O(n^{3/2}\sqrt{T})$

However, since we show that Exp2 and PolyExp are equivalent, they must have the same regret bound. This implies an improvement on Exp2’s regret bound.

Proposition 2 *For the Online Linear Optimization problem on the $\{0, 1\}^n$ Hypercube with L_∞ adversarial losses, Exp2, OMD with Entropic regularization and Bernoulli sampling, and PolyExp have the following regret:*

1. *Full Information:* $O(n\sqrt{T})$
2. *Semi-Bandit:* $O(n\sqrt{T})$
3. *Bandit:* $O(n^{3/2}\sqrt{T})$

Under the L_∞ assumption, the established lower bounds presented in Koolen et al. (2010) and Audibert et al. (2011) match the regret upper bound of PolyExp.

1.2 Relation to Previous Works

In previous works on OLO (Dani et al., 2008; Koolen et al., 2010; Audibert et al., 2011; Cesa-Bianchi and Lugosi, 2012; Bubeck et al., 2012; Audibert et al., 2013) the authors consider arbitrary subsets of $\{0, 1\}^n$ as their decision set. This is also called as Online Combinatorial optimization. In our work, the decision set is the entire $\{0, 1\}^n$ hypercube. Moreover, the assumption on the adversarial losses are different. Most of the previous works use L_2 assumption (Bubeck et al., 2012; Dani et al., 2008; Cesa-Bianchi and Lugosi, 2012) and some use L_∞ assumption (Koolen et al., 2010; Audibert et al., 2011).

The Exp2 algorithm has been studied under various names, each with their own modifications and improvements. In its most basic form, it corresponds to the Hedge algorithm from Freund and Schapire (1997) for full information. For combinatorial decision sets, it has been studied by Koolen et al. (2010) for full information. In the bandit case, several variants of Exp2 exist based on the exploration distribution used. These were studied in Dani et al. (2008); Cesa-Bianchi and Lugosi (2012) and Bubeck et al. (2012). It has been proven in Audibert et al. (2011) that Exp2 is provably sub optimal for some decision sets and losses.

Follow the Leader kind of algorithms were introduced by Kalai and Vempala (2005) for full information setting, which can be extended to the bandit settings as well.

Mirror descent style of algorithms were introduced in Nemirovsky and Yudin (1983). For online learning, several works (Abernethy et al., 2009; Koolen et al., 2010; Bubeck et al., 2012; Audibert et al., 2013) consider OMD style of algorithms. Other algorithms such as Hedge, FTRL etc can be shown to be equivalent to OMD with the right regularization function. In fact, Srebro et al. (2011) show that OMD can always achieve a nearly optimal regret guarantee for a general class of online learning problems.

We refer the readers to the books by Cesa-Bianchi and Lugosi (2006), Bubeck and Cesa-Bianchi (2012), Shalev-Shwartz (2012), Hazan (2016) and lectures by Rakhlin and Tewari (2009), Bubeck (2011) for a comprehensive survey of online learning algorithms.

2. Algorithms, Equivalences and Regret

In this section, we describe and analyze the Exp2, OMD with Entropic regularization and Bernoulli Sampling, and PolyExp algorithms and prove their equivalence.

2.1 Exp2

For all the three settings, the loss vector used to update Exp2 must satisfy the condition that $\mathbb{E}_{X_t}[\tilde{l}_t] = l_t$. We can verify that this is true for the three cases. In the bandit case, the estimator was first proposed by Dani et al. (2008). Here, μ is the exploration distribution and γ is the mixing coefficient. We use uniform exploration over $\{0, 1\}^n$ as proposed in Cesa-Bianchi and Lugosi (2012).

Exp2 has several computational drawbacks. First, it uses 2^n parameters to maintain the distribution p_t . Sampling from this distribution in step 1 and updating it step 3 will require exponential time. For the bandit settings, even computing \tilde{l}_t will require exponential time. We state the following regret bounds by analyzing Exp2 directly. Later, we prove that these can be improved. These regret bounds are under the L_∞ assumption.

Algorithm: Exp2

Parameters: Learning Rate η

Let $w_1(X) = 1$ for all $X \in \{0, 1\}^n$. For each round $t = 1, 2, \dots, T$:

1. Sample X_t as below. Play X_t and incur the loss $X_t^\top l_t$.
 - (a) Full Information: $X_t \sim p_t(X) = \frac{w_t(X)}{\sum_{Y \in \{0,1\}^n} w_t(Y)}$.
 - (b) Semi-Bandit: $X_t \sim p_t(X)$
 - (c) Bandit: $X_t \sim q_t(X) = (1 - \gamma)p_t(X) + \gamma\mu(X)$. Here μ is the exploration distribution.
2. See Feedback and construct \tilde{l}_t .
 - (a) Full Information: $\tilde{l}_t = l_t$.
 - (b) Semi-Bandit: $\tilde{l}_{t,i} = \frac{l_{t,i}X_{t,i}}{m_{t,i}}$, where $m_{t,i} = \sum_{Y \in \{0,1\}^n: Y_i=1} p_t(Y)$
 - (c) Bandit: $\tilde{l}_t = P_t^{-1}X_tX_t^\top l_t$, where $P_t = \mathbb{E}_{X \sim q_t}[XX^\top]$
3. Update for all $X \in \{0, 1\}^n$

$$w_{t+1}(X) = \exp(-\eta X^\top \tilde{l}_t)w_t(X) \quad \text{or equivalently} \quad w_{t+1}(X) = \exp(-\eta \sum_{\tau=1}^t X^\top \tilde{l}_\tau)$$

Exp2 attains the regret bounds:

1. **Theorem 7:** For Full Information, if $\eta = \sqrt{\frac{\log 2}{nT}}$, then $E[R_T] \leq 2n^{3/2}\sqrt{T \log 2}$
2. **Theorem 9:** For Semi Bandit, if $\eta = \sqrt{\frac{\log 2}{nT}}$, then $E[R_T] \leq 2n^{3/2}\sqrt{T \log 2}$
3. **Theorem 11:** For Bandit, using uniform exploration on $\{0, 1\}^n$, if $\eta = \sqrt{\frac{\log 2}{9n^2T}}$ and $\gamma = 4n^2\eta$, then $\mathbb{E}[R_T] \leq 6n^2\sqrt{T \log 2}$

2.2 PolyExp

To get a polynomial time algorithm, we replace the sampling and update steps with polynomial time operations. PolyExp uses n parameters represented by the vector x_t . Each element of x_t corresponds to the mean of a Bernoulli distribution. It uses the product of these Bernoulli distributions to sample X_t and uses the update equation mentioned in step 3 to obtain x_{t+1} .

In the Semi-Bandit setting, it is easy to see that $m_{t,i} = \sum_{Y: Y_i=1} \prod_{j=1}^n \text{Bernoulli}(x_{t,j}) = x_{t,i}$. In the Bandit setting, we can sample X_t by sampling from $\prod_{i=1}^n \text{Bernoulli}(x_{t,i})$ with probability $1 - \gamma$ and sampling from μ with probability γ . As we use the uniform distribution over $\{0, 1\}^n$ for exploration, this is equivalent to sampling from $\prod_{i=1}^n \text{Bernoulli}(1/2)$. So we

Algorithm: PolyExp

Parameters: Learning Rate η

Let $x_{i,1} = 1/2$ for all $i \in [n]$. For each round $t = 1, 2, \dots, T$:

1. Sample X_t as below. Play X_t and incur the loss $X_t^\top l_t$.
 - (a) Full information: $X_{i,t} \sim \text{Bernoulli}(x_{i,t})$
 - (b) Semi-Bandit: $X_{i,t} \sim \text{Bernoulli}(x_{i,t})$
 - (c) Bandit: With probability $1 - \gamma$ sample $X_{i,t} \sim \text{Bernoulli}(x_{i,t})$ and with probability γ sample $X_t \sim \mu$
2. See Feedback and construct \tilde{l}_t
 - (a) Full information: $\tilde{l}_t = l_t$
 - (b) Semi-Bandit: $\tilde{l}_{t,i} = \frac{l_{t,i} X_{t,i}}{x_{t,i}}$
 - (c) Bandit: $\tilde{l}_t = P_t^{-1} X_t X_t^\top l_t$, where $P_t = (1 - \gamma)\Sigma_t + \gamma\mathbb{E}_{X \sim \mu}[X X^\top]$. The matrix Σ_t is $\Sigma_t[i, j] = x_{i,t}x_{j,t}$ if $i \neq j$ and $\Sigma_t[i, i] = x_i$ for all $i, j \in [n]$
3. Update for all $i \in [n]$:

$$x_{i,t+1} = \frac{x_{i,t}}{x_{i,t} + (1 - x_{i,t}) \exp(\eta \tilde{l}_{i,t})} \quad \text{or equivalently} \quad x_{i,t+1} = \frac{1}{1 + \exp(\eta \sum_{\tau=1}^t \tilde{l}_{i,\tau})}$$

can sample from μ in polynomial time. The matrix $P_t = \mathbb{E}_{X \sim \mu}[X X^\top] = (1 - \gamma)\Sigma_t + \gamma\Sigma_\mu$. Here Σ_t and Σ_μ are the covariance matrices when $X \sim \prod_{i=1}^n \text{Bernoulli}(x_{t,i})$ and $X \sim \prod_{i=1}^n \text{Bernoulli}(1/2)$ respectively. It can be verified that $\Sigma_t[i, j] = x_{i,t}x_{j,t}$, $\Sigma_\mu[i, j] = 1/4$ if $i \neq j$ and $\Sigma_t[i, i] = x_i$, $\Sigma_\mu[i, i] = 1/2$ for all $i, j \in [n]$. So P_t^{-1} can be computed in polynomial time.

2.3 Equivalence of Exp2 and PolyExp

We prove that running Exp2 is equivalent to running PolyExp.

Theorem 3 *Under linear losses \tilde{l}_t , Exp2 on $\{0, 1\}^n$ is equivalent to PolyExp. At round t , The probability that PolyExp chooses X is $\prod_{i=1}^n (x_{i,t})^{X_i} (1 - x_{i,t})^{(1-X_i)}$ where $x_{i,t} = (1 + \exp(\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau}))^{-1}$. This is equal to the probability of Exp2 choosing X at round t , ie:*

$$\prod_{i=1}^n (x_{i,t})^{X_i} (1 - x_{i,t})^{(1-X_i)} = \frac{\exp(-\eta \sum_{\tau=1}^{t-1} X^\top \tilde{l}_\tau)}{\sum_{Y \in \{0,1\}^n} \exp(-\eta \sum_{\tau=1}^{t-1} Y^\top \tilde{l}_\tau)}$$

At every round, the probability distribution p_t in Exp2 is the same as the product of Bernoulli distributions in PolyExp. This is because our decision set is the entire $\{0, 1\}^n$ hypercube. The vector \tilde{l}_t computed by Exp2 and PolyExp will be same. Hence, Exp2 and PolyExp are equivalent. Note that this equivalence is true for any sequence of losses as long as they are linear.

2.4 Online Mirror Descent

We present the OMD algorithm for linear losses on general finite decision sets. Our exposition is adapted from Bubeck and Cesa-Bianchi (2012) and Shalev-Shwartz (2012). Let $\mathcal{X} \subset \mathbb{R}^n$ be an open convex set and $\bar{\mathcal{X}}$ be the closure of \mathcal{X} . Let $\mathcal{K} \in \mathbb{R}^d$ be a finite decision set such that $\bar{\mathcal{X}}$ is the convex hull of \mathcal{K} . Since \mathcal{K} could be any finite set, we omit the semi-bandit case in the algorithm as this setting is meaningful only on $\{0, 1\}^n$.

A continuous function $F : \mathcal{X} \rightarrow \mathbb{R}$ is *Legendre* if 1) F is strictly convex and has continuous partial derivatives on \mathcal{X} , 2) $\lim_{x \rightarrow \bar{\mathcal{X}}/\mathcal{X}} \|\nabla F(x)\| = +\infty$. The *Legendre-Fenchel conjugate* of F is:

$$F^*(\theta) = \sup_{x \in \bar{\mathcal{X}}} (x^\top \theta - F(x))$$

The *Bregman divergence* $D_F : \bar{\mathcal{X}} \times \mathcal{X} \rightarrow \mathbb{R}$ is:

$$D_F(x||y) = F(x) - F(y) - \nabla F(y)^\top (x - y)$$

Algorithm: Online Mirror Descent with Regularization $F(x)$

Parameters: Learning Rate η

Pick $x_1 = \arg \min_{x \in \bar{\mathcal{X}}} F(x)$. For each round $t = 1, 2, \dots, T$:

1. Let p_t be a distribution on \mathcal{K} such that $\mathbb{E}_{X \sim p_t}[X] = x_t$. Sample X_t as below and incur the loss $X_t^\top l_t$
 - (a) Full information: $X_t \sim p_t$
 - (b) Bandit: With probability $1 - \gamma$ sample $X_t \sim p_t$ and with probability γ sample $X_t \sim \mu$.
2. See Feedback and construct \tilde{l}_t
 - (a) Full information: $\tilde{l}_t = l_t$
 - (b) Bandit: $\tilde{l}_t = P_t^{-1} X_t X_t^\top l_t$, where $P_t = (1 - \gamma)\mathbb{E}_{X \sim p_t}[X X^\top] + \gamma\mathbb{E}_{X \sim \mu}[X X^\top]$.
3. Let y_{t+1} satisfy: $y_{t+1} = \nabla F^*(\nabla F(x_t) - \eta \tilde{l}_t)$
4. Update $x_{t+1} = \arg \min_{x \in \bar{\mathcal{X}}} D_F(x||y_{t+1})$

2.5 Equivalence of PolyExp and Online Mirror Descent

For our problem, $\mathcal{K} = \{0, 1\}^n$, $\bar{\mathcal{X}} = [0, 1]^n$ and $\mathcal{X} = (0, 1)^n$. We use entropic regularization:

$$F(x) = \sum_{i=1}^n x_i \log x_i + (1 - x_i) \log(1 - x_i)$$

This function is Legendre. The OMD algorithm does not specify the probability distribution p_t that should be used for sampling. The only condition that needs to be met is $\mathbb{E}_{X \sim p_t}[X] =$

x_t , i.e. x_t should be expressed as a convex combination of $\{0, 1\}^n$ and probability of picking X is its coefficient in the linear decomposition of x_t . An easy way to achieve this is by using Bernoulli sampling like in PolyExp. Hence, we have the following equivalence theorem:

Theorem 4 *Under linear losses \tilde{l}_t , OMD on $[0, 1]^n$ with Entropic Regularization and Bernoulli Sampling is equivalent to PolyExp. The sampling procedure of PolyExp satisfies $\mathbb{E}[X_t] = x_t$. The update of OMD with Entropic Regularization is the same as PolyExp.*

In the semi-bandit setting, OMD can be used on $[0, 1]^n$ as follows: Sample $X_t \sim p_t$ and construct $\tilde{l}_{t,i} = l_{t,i}X_{t,i}/m_{t,i}$ where $m_{t,i} = \sum_{Y:Y_i=1} p(Y)$. If we use Bernoulli sampling, $m_{t,i} = x_{t,i}$. This is the same as PolyExp. Similarly, in the bandit case, if we use Bernoulli sampling, $\mathbb{E}_{X \sim p_t}[XX^\top] = \Sigma_t$.

2.6 Regret of PolyExp via OMD analysis

Since OMD and PolyExp are equivalent, we can use the standard analysis tools of OMD to derive a regret bound for PolyExp. These regret bounds are under the L_∞ assumption. PolyExp attains the regret bounds:

1. **Theorem 16:** For Full Information, if $\eta = \sqrt{\frac{\log 2}{T}}$, then $E[R_T] \leq 2n\sqrt{T \log 2}$
2. **Theorem 18:** For Semi Bandit, if $\eta = \sqrt{\frac{\log 2}{T}}$, then $E[R_T] \leq 2n\sqrt{T \log 2}$
3. **Theorem 20:** For Bandit, using uniform exploration on $\{0, 1\}^n$, $\eta = \sqrt{\frac{3 \log 2}{8nT}}$ and $\gamma = 4n\eta$, then $\mathbb{E}[R_T] \leq 4n^{3/2}\sqrt{6T \log 2}$

We have shown that Exp2 on $\{0, 1\}^n$ with linear losses is equivalent to PolyExp. We have also shown that PolyExp's regret bounds are tighter than the regret bounds that we were able to derive for Exp2. This naturally implies an improvement for Exp2's regret bounds as it is equivalent to PolyExp and must attain the same regret.

3. Comparison of Exp2's and PolyExp's regret proofs

Consider the results we have shown so far. We proved that PolyExp and Exp2 on the hypercube are equivalent. So logically, they should have the same regret bounds. But, our proofs say that PolyExp's regret is $O(\sqrt{n})$ better than Exp2's regret. What is the reason for this apparent discrepancy?

The answer lies in the choice of learning rate η and the application of the inequality $e^{-x} \leq 1 + x - x^2$ in our proofs. This inequality is valid when $x \geq -1$. When analyzing Exp2, x is $\eta X^\top l_t = \eta L_t(X)$. So, to satisfy the constraints $x \geq -1$ we enforce that $|\eta L_t(X)| \leq 1$. Since $|L_t(X)| \leq n$, $\eta \leq 1/n$. This governs the optimal η parameter that we are able to get using Exp2's proof technique. When analyzing PolyExp, x is $\eta l_{t,i}$ and we enforce that $|\eta l_{t,i}| \leq 1$. Since we already assume $|l_{t,i}| \leq 1$, we get that $\eta \leq 1$. PolyExp's proof technique allows us to find a better η and achieve a better regret bound.

References

- Jacob D Abernethy, Elad Hazan, and Alexander Rakhlin. Competing in the dark: An efficient algorithm for bandit linear optimization. 2009.
- Jean-Yves Audibert, Sébastien Bubeck, and Gábor Lugosi. Minimax policies for combinatorial prediction games. In *Proceedings of the 24th Annual Conference on Learning Theory*, pages 107–132, 2011.
- Jean-Yves Audibert, Sébastien Bubeck, and Gábor Lugosi. Regret in online combinatorial optimization. *Mathematics of Operations Research*, 39(1):31–45, 2013.
- Sébastien Bubeck. Introduction to online optimization. *Lecture Notes*, pages 1–86, 2011.
- Sébastien Bubeck and Nicolo Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandit problems. *Foundations and Trends® in Machine Learning*, 5(1): 1–122, 2012.
- Sébastien Bubeck, Nicolo Cesa-Bianchi, and Sham Kakade. Towards minimax policies for online linear optimization with bandit feedback. In *Annual Conference on Learning Theory*, volume 23, pages 41–1. Microtome, 2012.
- Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, learning, and games*. Cambridge university press, 2006.
- Nicolo Cesa-Bianchi and Gábor Lugosi. Combinatorial bandits. *Journal of Computer and System Sciences*, 78(5):1404–1422, 2012.
- Varsha Dani, Sham M Kakade, and Thomas P Hayes. The price of bandit information for online optimization. In *Advances in Neural Information Processing Systems*, pages 345–352, 2008.
- Yoav Freund and Robert E Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of computer and system sciences*, 55(1):119–139, 1997.
- Elad Hazan. Introduction to online convex optimization. *Foundations and Trends® in Optimization*, 2(3-4):157–325, 2016.
- Adam Kalai and Santosh Vempala. Efficient algorithms for online decision problems. *Journal of Computer and System Sciences*, 71(3):291–307, 2005.
- Wouter M Koolen, Manfred K Warmuth, and Jyrki Kivinen. Hedging structured concepts. In *COLT*, pages 93–105. Citeseer, 2010.
- Nick Littlestone and Manfred K Warmuth. The weighted majority algorithm. *Information and computation*, 108(2):212–261, 1994.
- Arkadii Semenovich Nemirovsky and David Borisovich Yudin. Problem complexity and method efficiency in optimization. 1983.

Alexander Rakhlin and A Tewari. Lecture notes on online learning. *Draft, April, 2009.*

Shai Shalev-Shwartz. Online learning and online convex optimization. *Foundations and Trends® in Machine Learning*, 4(2):107–194, 2012.

Nati Srebro, Karthik Sridharan, and Ambuj Tewari. On the universality of online mirror descent. In *Advances in neural information processing systems*, pages 2645–2653, 2011.

Appendix A. A note on Lower bounds

We consider the lower bounds presented in Audibert et al. (2011) under the L_∞ assumption. These bounds address the maximal minimax regret on arbitrary subsets of the hypercube and not the minimax regret on the hypercube. They prove that there exists a subset $S \subset \{0, 1\}^n$ and a sequence of loss vectors such that any OLO algorithm on S has to have regret lower bounded by:

1. $\Omega(n\sqrt{T})$ in the Full Information.
2. $\Omega(n\sqrt{T})$ in the Semi-Bandit.
3. $\Omega(n^{3/2}\sqrt{T})$ in the Bandit.

Since we consider the entire hypercube, these lower bounds do not directly apply to our setting. However, the regret upper bounds of PolyExp match these lower bounds. Finding lower bounds on the hypercube under L_∞ assumption is an open question.

Appendix B. $\{-1, +1\}^n$ Hypercube Case

In Bubeck et al. (2012), the authors state that it is not known if it is possible to sample from the exponential weights distribution in polynomial time for $\{-1, +1\}^n$ hypercube. We show how to use PolyExp on $\{0, 1\}^n$ for $\{-1, +1\}^n$. We show that the regret of such an algorithm on $\{-1, +1\}^n$ will be a constant factor away from the regret of the algorithm on $\{0, 1\}^n$. Thus, we can use PolyExp to obtain a polynomial time algorithm for $\{-1, +1\}^n$ hypercube.

Full information and bandit algorithms which work on $\{0, 1\}^n$ can be modified to work on $\{-1, +1\}^n$. The semi-bandit case is meaningful only on $\{0, 1\}^n$, so we omit it. The general strategy is as follows:

Theorem 5 *Exp2 on $\{-1, +1\}^n$ using the sequence of losses l_t is equivalent to PolyExp on $\{0, 1\}^n$ using the sequence of losses $2\tilde{l}_t$. Moreover, the regret of Exp2 on $\{-1, 1\}^n$ will equal the regret of PolyExp using the losses $2\tilde{l}_t$.*

Hence, using the above strategy, PolyExp can be run in polynomial time on $\{-1, 1\}^n$ and since the losses are doubled its regret only changes by a constant factor.

Appendix C. Proofs

C.1 Exp2 Regret Proofs

First, we directly analyze Exp2’s regret for the three kinds of feedback.

1. Sample $X_t \in \{0, 1\}^n$, play $Z_t = 2X_t - \mathbf{1}$ and incur loss $Z_t^\top l_t$.
 - (a) Full information: $X_t \sim p_t$
 - (b) Bandit: $X_t \sim q_t = (1 - \gamma)p_t + \gamma\mu$
2. See feedback and construct \tilde{l}_t
 - (a) Full information: $\tilde{l}_t = l_t$
 - (b) Bandit: $\tilde{l}_t = P_t^{-1} Z_t Z_t^\top l_t$ where $P_t = \mathbb{E}_{X \sim q_t} [(2X - \mathbf{1})(2X - \mathbf{1})^\top]$
3. Update algorithm using $2\tilde{l}_t$

C.1.1 FULL INFORMATION

Lemma 6 *Let $L_t(X) = X^\top l_t$. If $|\eta L_t(X)| \leq 1$ for all $t \in [T]$ and $X \in \{0, 1\}^n$, the Exp2 algorithm satisfies for any X :*

$$\sum_{t=1}^T p_t^\top L_t - \sum_{t=1}^T L_t(X) \leq \eta \sum_{t=1}^T p_t^\top L_t^2 + \frac{n \log 2}{\eta}$$

Proof (Adapted from Hazan (2016) Theorem 1.5) Let $Z_t = \sum_{Y \in \{0,1\}^n} w_t(Y)$. We have:

$$\begin{aligned} Z_{t+1} &= \sum_{Y \in \{0,1\}^n} \exp(-\eta L_t(Y)) w_t(Y) \\ &= Z_t \sum_{Y \in \{0,1\}^n} \exp(-\eta L_t(Y)) p_t(Y) \end{aligned}$$

Since $e^{-x} \leq 1 - x + x^2$ for $x \geq -1$, we have that $\exp(-\eta L_t(Y)) \leq 1 - \eta L_t(Y) + \eta^2 L_t(Y)^2$ (Because we assume $|\eta L_t(X)| \leq 1$). So,

$$\begin{aligned} Z_{t+1} &\leq Z_t \sum_{Y \in \{0,1\}^n} (1 - \eta L_t(Y) + \eta^2 L_t(Y)^2) p_t(Y) \\ &= Z_t (1 - \eta p_t^\top L_t + \eta^2 p_t^\top L_t^2) \end{aligned}$$

Using the inequality $1 + x \leq e^x$,

$$Z_{t+1} \leq Z_t \exp(-\eta p_t^\top L_t + \eta^2 p_t^\top L_t^2)$$

Hence, we have:

$$Z_{T+1} \leq Z_1 \exp\left(-\sum_{t=1}^T \eta p_t^\top L_t + \sum_{t=1}^T \eta^2 p_t^\top L_t^2\right)$$

For any $X \in \{0, 1\}^n$, $w_{T+1}(X) = \exp(-\sum_{t=1}^T \eta L_t(X))$. Since $w_{T+1}(X) \leq Z_{T+1}$ and $Z_1 = 2^n$, we have:

$$\exp\left(-\sum_{t=1}^T \eta L_t(X)\right) \leq 2^n \exp\left(-\sum_{t=1}^T \eta p_t^\top L_t + \sum_{t=1}^T \eta^2 p_t^\top L_t^2\right)$$

Taking the logarithm on both sides manipulating this inequality, we get:

$$\sum_{t=1}^T p_t^\top L_t - \sum_{t=1}^T L_t(X) \leq \eta \sum_{t=1}^T p_t^\top L_t^2 + \frac{n \log 2}{\eta}$$

■

Theorem 7 *In the full information setting, if $\eta = \sqrt{\frac{\log 2}{nT}}$, Exp2 attains the regret bound:*

$$E[R_T] \leq 2n^{3/2} \sqrt{T \log 2}$$

Proof Using $L_t(X) = X^\top l_t$ and applying expectation with respect to the randomness of the player to definition of regret, we get:

$$E[R_T] = \sum_{t=1}^T \sum_{X \in \{0,1\}^n} p_t(X) L_t(X) - \min_{X^* \in \{0,1\}^n} \sum_{t=1}^T L_t(X^*) = \sum_{t=1}^T p_t^\top L_t - \min_{X^* \in \{0,1\}^n} \sum_{t=1}^T L_t(X^*)$$

Applying Lemma 6, we get $E[R_T] \leq \eta \sum_{t=1}^T p_t^\top L_t^2 + n \log 2 / \eta$. Since $|L_t(X)| \leq n$ for all $X \in \{0,1\}^n$, we get $\sum_{t=1}^T p_t^\top L_t^2 \leq Tn^2$.

$$E[R_T] \leq \eta T n^2 + \frac{n \log 2}{\eta}$$

Optimizing over the choice of η , we get the regret is bounded by $2n^{3/2} \sqrt{T \log 2}$ if we choose $\eta = \sqrt{\frac{\log 2}{nT}}$.

To apply Lemma 6, $|\eta L_t(X)| \leq 1$ for all $t \in [T]$ and $X \in \{0,1\}^n$. Since $|L_t(X)| \leq n$, we have $\eta \leq 1/n$. ■

C.1.2 SEMI BANDIT

Lemma 8 *Let $\tilde{L}_t(X) = X^\top \tilde{l}_t$, where $\tilde{l}_{t,i} = l_{t,i} X_{t,i} / m_{t,i}$. If $|\eta \tilde{L}_t(X)| \leq 1$ for all $t \in [T]$ and $X \in \{0,1\}^n$, the Exp2 algorithm satisfies for any X*

$$\sum_{t=1}^T p_t^\top L_t - \sum_{t=1}^T L_t(X) \leq \eta \mathbb{E} \left[\sum_{t=1}^T p_t^\top \tilde{L}_t^2 \right] + \frac{n \log 2}{\eta}$$

Proof Since the algorithm essentially runs Exp2 using the losses $\tilde{L}_t(X)$ and $|\eta \tilde{L}_t(X)| \leq 1$, we can apply Lemma 6:

$$\sum_{t=1}^T p_t^\top \tilde{L}_t - \sum_{t=1}^T \tilde{L}_t(X) \leq \eta \sum_{t=1}^T p_t^\top \tilde{L}_t^2 + \frac{n \log 2}{\eta}$$

Apply expectation with respect to X_t . Using the fact that $\mathbb{E}[\tilde{l}_t] = l_t$:

$$\sum_{t=1}^T p_t^\top L_t - \sum_{t=1}^T L_t(X) \leq \eta \mathbb{E}[\sum_{t=1}^T p_t^\top \tilde{L}_t^2] + \frac{n \log 2}{\eta}$$

■

Theorem 9 *In the semi-bandit setting, if $\eta = \sqrt{\frac{\log 2}{nT}}$, Exp2 attains the regret bound:*

$$E[R_T] \leq 2n^{3/2} \sqrt{T \log 2}$$

Proof Applying expectation with respect to the randomness of the player to the definition of regret, we get:

$$\mathbb{E}[R_T] = \mathbb{E}[\sum_{t=1}^T L_t(X_t) - \min_{X^* \in \{0,1\}^n} L_t(X^*)] = \sum_{t=1}^T p_t^\top L_t - \min_{X^* \in \{0,1\}^n} \sum_{t=1}^T L_t(X^*)$$

Applying Lemma 8, we get:

$$\mathbb{E}[R_T] \leq \eta \mathbb{E}[\sum_{t=1}^T p_t^\top \tilde{L}_t^2] + \frac{n \log 2}{\eta}$$

We follow the proof technique of Audibert et al. (2011) Theorem 17. We have that:

$$\begin{aligned} \mathbb{E}[p_t^\top \tilde{L}_t^2] &= \mathbb{E}[\sum_{X \in \{0,1\}^n} p_t(X) (X^\top \tilde{l}_t)^2] = \mathbb{E}_{X_t} \mathbb{E}_X [(X^\top \tilde{l}_t)^2] \\ &= \mathbb{E}_{X_t} \mathbb{E}_X [(\sum_{i=1}^n X_i \tilde{l}_{t,i})^2] = \mathbb{E}_{X_t} \mathbb{E}_X [\sum_{i=1}^n \sum_{j=1}^n X_i \tilde{l}_{t,i} X_j \tilde{l}_{t,j}] \\ &= \mathbb{E}_{X_t} \mathbb{E}_X [\sum_{i=1}^n \sum_{j=1}^n \frac{X_{t,j} l_{t,i}}{m_{i,t}} \frac{X_{t,i} l_{t,j}}{m_{j,t}} X_i X_j] \\ &\leq \mathbb{E}_{X_t} \mathbb{E}_X [\sum_{i=1}^n \sum_{j=1}^n l_{t,i} l_{t,j} \frac{X_{t,j}}{m_{i,t}} \frac{X_j}{m_{j,t}}] = (\sum_{i=1}^n l_{t,i})^2 \leq n^2 \end{aligned}$$

Hence, we get:

$$\mathbb{E}[R_T] \leq \eta n^2 T + \frac{n \log 2}{\eta}$$

Rest of the proof is similar to the proof of Theorem 7. ■

C.1.3 BANDIT

Lemma 10 *Let $\tilde{L}_t(X) = X^\top \tilde{l}_t$, where $\tilde{l}_t = P_t^{-1} X_t X_t^\top l_t$. If $|\eta \tilde{L}_t(X)| \leq 1$ for all $t \in [T]$ and $X \in \{0, 1\}^n$, the Exp2 algorithm with uniform exploration satisfies for any X*

$$\sum_{t=1}^T q_t^\top L_t - \sum_{t=1}^T L_t(X) \leq \eta \mathbb{E}[\sum_{t=1}^T q_t^\top \tilde{L}_t^2] + \frac{n \log 2}{\eta} + 2\gamma nT$$

Proof We have that:

$$\sum_{t=1}^T q_t^\top \tilde{L}_t - \sum_{t=1}^T \tilde{L}_t(X) = (1 - \gamma) \left(\sum_{t=1}^T p_t^\top \tilde{L}_t - \sum_{t=1}^T \tilde{L}_t(X) \right) + \gamma \left(\sum_{t=1}^T \mu^\top \tilde{L}_t - \sum_{t=1}^T \tilde{L}_t(X) \right)$$

Since the algorithm essentially runs Exp2 using the losses $\tilde{L}_t(X)$ and $|\eta \tilde{L}_t(X)| \leq 1$, we can apply Lemma 6:

$$\sum_{t=1}^T q_t^\top \tilde{L}_t - \sum_{t=1}^T \tilde{L}_t(X) \leq (1 - \gamma) \left(\frac{n \log 2}{\eta} + \eta \sum_{t=1}^T p_t^\top \tilde{L}_t^2 \right) + \gamma \left(\sum_{t=1}^T \mu^\top \tilde{L}_t - \sum_{t=1}^T \tilde{L}_t(X) \right)$$

Apply expectation with respect to X_t . Using the fact that $\mathbb{E}[\tilde{l}_t] = l_t$ and $\mu^\top L_t - L_t(X) \leq 2n$:

$$\begin{aligned} \sum_{t=1}^T q_t^\top L_t - \sum_{t=1}^T L_t(X) &\leq (1 - \gamma) \left(\frac{n \log 2}{\eta} + \eta \mathbb{E}[\sum_{t=1}^T p_t^\top \tilde{L}_t^2] \right) + \gamma \left(\sum_{t=1}^T \mu^\top L_t - \sum_{t=1}^T L_t(X) \right) \\ &\leq \eta \mathbb{E}[\sum_{t=1}^T q_t^\top \tilde{L}_t^2] + \frac{n \log 2}{\eta} + 2\gamma nT \end{aligned}$$

■

Theorem 11 *In the bandit setting, if $\eta = \sqrt{\frac{\log 2}{9n^2 T}}$ and $\gamma = 4n^2 \eta$, Exp2 with uniform exploration on $\{0, 1\}^n$ attains the regret bound:*

$$\mathbb{E}[R_T] \leq 6n^2 \sqrt{T \log 2}$$

Proof Applying expectation with respect to the randomness of the player to the definition of regret, we get:

$$\mathbb{E}[R_T] = \mathbb{E} \left[\sum_{t=1}^T L_t(X_t) - \min_{X^* \in \{0, 1\}^n} L_t(X^*) \right] = \sum_{t=1}^T q_t^\top L_t - \min_{X^* \in \{0, 1\}^n} \sum_{t=1}^T L_t(X^*)$$

Applying Lemma 10

$$\mathbb{E}[R_T] \leq \eta \mathbb{E}[\sum_{t=1}^T q_t^\top \tilde{L}_t^2] + \frac{n \log 2}{\eta} + 2\gamma nT$$

We follow the proof technique of Bubeck et al. (2012) Theorem 4. We have that:

$$\begin{aligned}
 q_t^\top \tilde{L}_t^2 &= \sum_{X \in \{0,1\}^n} q_t(X) (X^\top \tilde{l}_t)^2 = \sum_{X \in \{0,1\}^n} q_t(X) (\tilde{l}_t^\top X X^\top \tilde{l}_t) \\
 &= \tilde{l}_t^\top P_t \tilde{l}_t = \tilde{l}_t^\top X_t X_t^\top P_t^{-1} P_t P_t^{-1} X_t X_t^\top \tilde{l}_t = (X_t^\top \tilde{l}_t)^2 X_t^\top P_t^{-1} X_t \\
 &\leq n^2 X_t^\top P_t^{-1} X_t = n^2 \text{Tr}(P_t^{-1} X_t X_t^\top)
 \end{aligned}$$

Taking expectation, we get $E[q_t^\top \tilde{L}_t^2] \leq n^2 \text{Tr}(P_t^{-1} \mathbb{E}[X_t X_t^\top]) = n^2 \text{Tr}(P_t^{-1} P_t) = n^3$. Hence,

$$\mathbb{E}[R_T] \leq \eta n^3 T + \frac{n \log 2}{\eta} + 2\gamma n T$$

However, in order to apply Lemma 10, we need that $|\eta X^\top \tilde{l}_t| \leq 1$. We have that

$$|\eta X^\top \tilde{l}_t| = \eta |(X_t^\top l_t) X^\top P_t^{-1} X_t| \leq 1$$

As $|X_t^\top l_t| \leq n$ and $|X_t^\top X_t| \leq n$, we get $\eta n |X_t^\top P_t^{-1} X_t| \leq \eta n |X_t^\top X_t| \|P_t^{-1}\| \leq \eta n^2 \|P_t^{-1}\| \leq 1$. The matrix $P_t = (1 - \gamma)\Sigma_t + \gamma\Sigma_\mu$. The smallest eigenvalue of Σ_μ is $1/4$ (Cesa-Bianchi and Lugosi, 2012). So $P_t \succeq \frac{\gamma}{4} I_n$ and $P_t^{-1} \preceq \frac{4}{\gamma} I_n$. We should have that $\frac{4n^2\eta}{\gamma} \leq 1$. Substituting $\gamma = 4n^2\eta$ in the regret inequality, we get:

$$\begin{aligned}
 \mathbb{E}[R_T] &\leq \eta n^3 T + 8\eta n^3 T + \frac{n \log 2}{\eta} \\
 &\leq 9\eta n^3 T + \frac{n \log 2}{\eta}
 \end{aligned}$$

Optimizing over the choice of η , we get $\mathbb{E}[R_T] \leq 2n^2 \sqrt{9T \log 2}$ when $\eta = \sqrt{\frac{\log 2}{9n^3 T}}$. ■

C.2 Equivalence Proofs of PolyExp

C.2.1 EQUIVALENCE TO EXP2

Lemma 12 *For any sequence of losses \tilde{l}_t , the following is true for all $t = 1, 2, \dots, T$:*

$$\prod_{i=1}^n \left(1 + \exp\left(-\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau}\right)\right) = \sum_{Y \in \{0,1\}^n} \exp\left(-\eta \sum_{\tau=1}^{t-1} Y^\top \tilde{l}_\tau\right)$$

Proof Consider $\prod_{i=1}^n (1 + \exp(-\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau}))$. It is a product of n terms, each consisting of 2 terms, 1 and $\exp(-\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau})$. On expanding the product, we get a sum of 2^n terms. Each of these terms is a product of n terms, either a 1 or $\exp(-\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau})$. If it is 1, then

$Y_i = 0$ and if it is $\exp(-\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau})$, then $Y_i = 1$. So,

$$\begin{aligned}
 \prod_{i=1}^n (1 + \exp(-\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau})) &= \sum_{Y \in \{0,1\}^n} \prod_{i=1}^n \exp(-\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau})^{Y_i} \\
 &= \sum_{Y \in \{0,1\}^n} \prod_{i=1}^n \exp(-\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau} Y_i) \\
 &= \sum_{Y \in \{0,1\}^n} \exp(-\eta \sum_{i=1}^n \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau} Y_i) \\
 &= \sum_{Y \in \{0,1\}^n} \exp(-\eta \sum_{\tau=1}^{t-1} Y^\top \tilde{l}_\tau)
 \end{aligned}$$

■

Theorem 3 Under linear losses \tilde{l}_t , *Exp2* on $\{0,1\}^n$ is equivalent to *PolyExp*. At round t , The probability that *PolyExp* chooses X is $\prod_{i=1}^n (x_{i,t})^{X_i} (1 - x_{i,t})^{(1-X_i)}$ where $x_{i,t} = (1 + \exp(\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau}))^{-1}$. This is equal to the probability of *Exp2* choosing X at round t , ie:

$$\prod_{i=1}^n (x_{i,t})^{X_i} (1 - x_{i,t})^{(1-X_i)} = \frac{\exp(-\eta \sum_{\tau=1}^{t-1} X^\top \tilde{l}_\tau)}{\sum_{Y \in \{0,1\}^n} \exp(-\eta \sum_{\tau=1}^{t-1} Y^\top \tilde{l}_\tau)}$$

Proof The proof is via straightforward substitution of the expression for $x_{i,t}$ and applying Lemma 12.

$$\begin{aligned}
 \prod_{i=1}^n (x_{i,t})^{X_i} (1 - x_{i,t})^{(1-X_i)} &= \prod_{i=1}^n \left(\frac{1}{1 + \exp(\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau})} \right)^{X_i} \left(\frac{\exp(\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau})}{1 + \exp(\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau})} \right)^{1-X_i} \\
 &= \prod_{i=1}^n \left(\frac{\exp(-\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau})}{1 + \exp(-\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau})} \right)^{X_i} \left(\frac{1}{1 + \exp(-\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau})} \right)^{1-X_i} \\
 &= \prod_{i=1}^n \frac{\exp(-\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau})^{X_i}}{1 + \exp(-\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau})} \\
 &= \frac{\prod_{i=1}^n \exp(-\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau})^{X_i}}{\prod_{i=1}^n (1 + \exp(-\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau}))} \\
 &= \frac{\prod_{i=1}^n \exp(-\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau} X_i)}{\prod_{i=1}^n (1 + \exp(-\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau}))} \\
 &= \frac{\exp(-\eta \sum_{i=1}^n \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau} X_i)}{\prod_{i=1}^n (1 + \exp(-\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau}))} \\
 &= \frac{\exp(-\eta \sum_{\tau=1}^{t-1} X^\top \tilde{l}_\tau)}{\prod_{i=1}^n (1 + \exp(-\eta \sum_{\tau=1}^{t-1} \tilde{l}_{i,\tau}))}
 \end{aligned}$$

$$= \frac{\exp(-\eta \sum_{\tau=1}^{t-1} X^\top \tilde{l}_{i,\tau})}{\sum_{Y \in \{0,1\}^n} \exp(-\eta \sum_{\tau=1}^{t-1} Y^\top \tilde{l}_\tau)}$$

■

C.2.2 EQUIVALENCE TO OMD

Lemma 13 *The Fenchel Conjugate of $F(x) = \sum_{i=1}^n x_i \log x_i + (1 - x_i) \log(1 - x_i)$ is:*

$$F^*(\theta) = \sum_{i=1}^n \log(1 + \exp(\theta_i))$$

Proof Differentiating $x^\top \theta - F(x)$ wrt x_i and equating to 0:

$$\begin{aligned} \theta_i - \log x_i + \log(1 - x_i) &= 0 \\ \frac{x_i}{1 - x_i} &= \exp(\theta_i) \\ x_i &= \frac{1}{1 + \exp(-\theta_i)} \end{aligned}$$

Substituting this back in $x^\top \theta - F(x)$, we get $F^*(\theta) = \sum_{i=1}^n \log(1 + \exp(\theta_i))$. It is also straightforward to see that $\nabla F^*(\theta)_i = (1 + \exp(-\theta_i))^{-1}$ ■

Theorem 4 *Under linear losses \tilde{l}_t , OMD on $[0, 1]^n$ with Entropic Regularization and Bernoulli Sampling is equivalent to PolyExp. The sampling procedure of PolyExp satisfies $\mathbb{E}[X_t] = x_t$. The update of OMD with Entropic Regularization is the same as PolyExp.*

Proof It is easy to see that $E[X_{i,t}] = \Pr(X_{i,t} = 1) = x_{i,t}$. Hence $E[X_t] = x_t$.

The update equation is $y_{t+1} = \nabla F^*(\nabla F(x_t) - \eta \tilde{l}_t)$. Evaluating ∇F and using ∇F^* from Lemma 13:

$$\begin{aligned} y_{t+1,i} &= \frac{1}{1 + \exp(-\log(x_{t,i}) + \log(1 - x_{t,i}) + \eta \tilde{l}_{t,i})} \\ &= \frac{1}{1 + \frac{1-x_{t,i}}{x_{t,i}} \exp(\eta \tilde{l}_{t,i})} \\ &= \frac{x_{t,i}}{x_{t,i} + (1 - x_{t,i}) \exp(\eta \tilde{l}_{t,i})} \end{aligned}$$

Since $0 \leq (1 + \exp(-\theta))^{-1} \leq 1$, we have that $y_{i,t+1}$ is always in $[0, 1]$. Bregman projection step is not required. So we have $x_{i,t+1} = y_{i,t+1}$ which gives the same update as PolyExp. ■

C.3 PolyExp Regret Proofs

C.3.1 FULL INFORMATION

Lemma 14 (see Bubeck and Cesa-Bianchi, 2012, Theorem. 5.5) For any $x \in \bar{\mathcal{X}}$, OMD with Legendre regularizer $F(x)$ with domain $\bar{\mathcal{X}}$ and F^* is differentiable on \mathbb{R}^n satisfies:

$$\sum_{t=1}^T x_t^\top l_t - \sum_{t=1}^T x^\top l_t \leq \frac{F(x) - F(x_1)}{\eta} + \frac{1}{\eta} \sum_{t=1}^T D_{F^*}(\nabla F(x_t) - \eta l_t \| \nabla F(x_t))$$

Lemma 15 If $|\eta l_{t,i}| \leq 1$ for all $t \in [T]$ and $i \in [n]$, OMD with entropic regularizer $F(x) = \sum_{i=1}^n x_i \log x_i + (1 - x_i) \log(1 - x_i)$ satisfies for any $x \in [0, 1]^n$,

$$\sum_{t=1}^T x_t^\top l_t - \sum_{t=1}^T x^\top l_t \leq \frac{n \log 2}{\eta} + \eta \sum_{t=1}^T x_t^\top l_t^2$$

Proof We start from Lemma 14. Using the fact that $x \log(x) + (1 - x) \log(1 - x) \geq -\log 2$, we get $F(x) - F(x_1) \leq n \log 2$. Next we bound the Bregmen term using Lemma 13

$$D_{F^*}(\nabla F(x_t) - \eta l_t \| \nabla F(x_t)) = F^*(\nabla F(x_t) - \eta l_t) - F^*(\nabla F(x_t)) + \eta l_t^\top \nabla F^*(\nabla F(x_t))$$

Using that fact that $\nabla F^* = (\nabla F)^{-1}$, the last term is $\eta x_t^\top l_t$. The first two terms can be simplified as:

$$\begin{aligned} F^*(\nabla F(x_t) - \eta l_t) - F^*(\nabla F(x_t)) &= \sum_{i=1}^n \log \frac{1 + \exp(\nabla F(x_t)_i - \eta l_{t,i})}{1 + \exp(\nabla F(x_t)_i)} \\ &= \sum_{i=1}^n \log \frac{1 + \exp(-\nabla F(x_t)_i + \eta l_{t,i})}{\exp(\eta l_{t,i})(1 + \exp(-\nabla F(x_t)_i))} \end{aligned}$$

Using the fact that $\nabla F(x_t)_i = \log x_i - \log(1 - x_i)$:

$$\begin{aligned} &= \sum_{i=1}^n \log \frac{x_{t,i} + (1 - x_{t,i}) \exp(\eta l_{t,i})}{\exp(\eta l_{t,i})} \\ &= \sum_{i=1}^n \log(1 - x_{t,i} + x_{t,i} \exp(-\eta l_{t,i})) \end{aligned}$$

Using the inequality: $e^{-x} \leq 1 - x + x^2$ when $x \geq -1$. So when $|\eta l_{t,i}| \leq 1$:

$$\leq \sum_{i=1}^n \log(1 - \eta x_{t,i} l_{t,i} + \eta^2 x_{t,i} l_{t,i}^2)$$

Using the inequality: $\log(1 - x) \leq -x$

$$\leq -\eta x_t^\top l_t + \eta^2 x_t^\top l_t^2$$

The Bregman term can be bounded by $-\eta x_t^\top l_t + \eta^2 x_t^\top l_t^2 + \eta x_t^\top l_t = \eta^2 x_t^\top l_t^2$. Hence, we have:

$$\sum_{t=1}^T x_t^\top l_t - \sum_{t=1}^T x^\top l_t \leq \frac{n \log 2}{\eta} + \eta \sum_{t=1}^T x_t^\top l_t^2$$

■

Theorem 16 *In the full information setting, if $\eta = \sqrt{\frac{\log 2}{T}}$, PolyExp attains the regret bound:*

$$E[R_T] \leq 2n\sqrt{T \log 2}$$

Proof Applying expectation with respect to the randomness of the player to definition of regret, we get:

$$\mathbb{E}[R_T] = \mathbb{E}\left[\sum_{t=1}^T X_t^\top l_t - \min_{X^* \in \{0,1\}^n} \sum_{t=1}^T X^{*\top} l_t\right] = \sum_{t=1}^T x_t^\top l_t - \min_{X^* \in \{0,1\}^n} \sum_{t=1}^T X^{*\top} l_t$$

Applying Lemma 15, we get $E[R_T] \leq \frac{n \log 2}{\eta} + \eta \sum_{t=1}^T x_t^\top l_t^2$. Using the fact that $|l_{i,t}| \leq 1$, we get $\sum_{t=1}^T x_t^\top l_t^2 \leq nT$.

$$\mathbb{E}[R_T] \leq \eta nT + \frac{n \log 2}{\eta}$$

Optimizing over the choice of η , we get that the regret is bounded by $2n\sqrt{T \log 2}$ if we choose $\eta = \sqrt{\frac{\log 2}{T}}$. ■

C.3.2 SEMI BANDIT

Lemma 17 *Let $\tilde{l}_{t,i} = l_{t,i} X_{t,i} / x_{t,i}$. If $|\eta \tilde{l}_{t,i}| \leq 1$ for all $t \in [T]$ and $i \in [n]$, OMD with entropic regularization satisfies for any $x \in [0, 1]^n$:*

$$\sum_{t=1}^T x_t^\top l_t - \sum_{t=1}^T x^\top l_t \leq \eta \mathbb{E}\left[\sum_{t=1}^T x_t^\top \tilde{l}_t^2\right] + \frac{n \log 2}{\eta}$$

Proof Since the algorithm runs OMD on \tilde{l}_t and $|\eta \tilde{l}_t| \leq 1$, we can apply Lemma 15:

$$\sum_{t=1}^T x_t^\top \tilde{l}_t - \sum_{t=1}^T x^\top \tilde{l}_t \leq \eta \sum_{t=1}^T x_t^\top \tilde{l}_t^2 + \frac{n \log 2}{\eta}$$

Applying expectation with respect to X_t and using $E[\tilde{l}_t] = l_t$:

$$\sum_{t=1}^T x_t^\top l_t - \sum_{t=1}^T x^\top l_t \leq \eta \mathbb{E}\left[\sum_{t=1}^T x_t^\top \tilde{l}_t^2\right] + \frac{n \log 2}{\eta}$$

■

Theorem 18 *In the semi-bandit setting, if $\eta = \sqrt{\frac{\log 2}{T}}$, PolyExp attains the regret bound:*

$$E[R_T] \leq 2n\sqrt{T \log 2}$$

Proof Applying expectation with respect to the randomness of the player to the definition of regret:

$$\mathbb{E}[R_T] = \mathbb{E}\left[\sum_{t=1}^T X_t^\top l_t - \min_{X^* \in \{0,1\}^n} \sum_{t=1}^T X^{*\top} l_t\right] = \sum_{t=1}^T x_t^\top l_t - \min_{X^* \in \{0,1\}^n} \sum_{t=1}^T X^{*\top} l_t$$

Applying Lemma 17, we get:

$$\mathbb{E}[R_T] \leq \eta \mathbb{E}\left[\sum_{t=1}^T x_t^\top \tilde{l}_t^2\right] + \frac{n \log 2}{\eta}$$

We have that:

$$\mathbb{E}[x_t^\top \tilde{l}_t^2] = \mathbb{E}\left[\sum_{i=1}^n x_{t,i} \frac{l_{t,i}^2 X_{t,i}}{x_{t,i}^2}\right] = \mathbb{E}\left[\sum_{i=1}^n \frac{l_{t,i}^2 X_{t,i}}{x_{t,i}}\right] = \sum_{i=1}^n \frac{l_{t,i}^2}{x_{t,i}} \mathbb{E}[X_{t,i}] = \sum_{i=1}^n l_{t,i}^2 \leq n$$

Hence, we get:

$$\mathbb{E}[R_T] \leq \eta n T + \frac{n \log 2}{\eta}$$

Rest of the proof is similar to the proof of Theorem 16. ■

C.3.3 BANDIT

Lemma 19 *Let $\tilde{l}_t = P_t^{-1} X_t X_t^\top l_t$. If $|\eta \tilde{l}_{t,i}| \leq 1$ for all $t \in [T]$ and $i \in [n]$, OMD with entropic regularization and uniform exploration satisfies for any $x \in [0, 1]^n$:*

$$\sum_{t=1}^T x_t^\top l_t - \sum_{t=1}^T x^\top l_t \leq \eta \mathbb{E}\left[\sum_{t=1}^T x_t^\top \tilde{l}_t^2\right] + \frac{n \log 2}{\eta} + 2\gamma n T$$

Proof We have that:

$$\sum_{t=1}^T x_t^\top \tilde{l}_t - \sum_{t=1}^T x^\top \tilde{l}_t = (1 - \gamma) \left(\sum_{t=1}^T x_{p_t}^\top \tilde{l}_t - \sum_{t=1}^T x^\top \tilde{l}_t \right) + \gamma \left(\sum_{t=1}^T x_\mu^\top \tilde{l}_t - \sum_{t=1}^T x^\top \tilde{l}_t \right)$$

Since the algorithm runs OMD on \tilde{l}_t and $|\eta \tilde{l}_t| \leq 1$, we can apply Lemma 15:

$$\sum_{t=1}^T x_t^\top \tilde{l}_t - \sum_{t=1}^T x^\top \tilde{l}_t \leq (1 - \gamma) \left(\eta \sum_{t=1}^T x_{p_t}^\top \tilde{l}_t^2 + \frac{n \log 2}{\eta} \right) + \gamma \left(\sum_{t=1}^T x_\mu^\top \tilde{l}_t - \sum_{t=1}^T x^\top \tilde{l}_t \right)$$

Apply expectation with respect to X_t . Using the fact that $\mathbb{E}[\tilde{l}_t] = l_t$ and $x_\mu^\top l_t - x^\top l_t \leq 2n$:

$$\begin{aligned} \sum_{t=1}^T x_t^\top l_t - \sum_{t=1}^T x^\top l_t &\leq (1-\gamma)\left(\eta\mathbb{E}\left[\sum_{t=1}^T x_{p_t}^\top \tilde{l}_t^2\right] + \frac{n \log 2}{\eta}\right) + 2\gamma nT \\ &\leq \eta\mathbb{E}\left[\sum_{t=1}^T x_t^\top \tilde{l}_t^2\right] + \frac{n \log 2}{\eta} + 2\gamma nT \end{aligned}$$

■

Theorem 20 *In the bandit setting, if $\eta = \sqrt{\frac{3 \log 2}{8nT}}$ and $\gamma = 4n\eta$, PolyExp with uniform exploration on $\{0, 1\}^n$ attains the regret bound:*

$$\mathbb{E}[R_T] \leq 4n^{3/2} \sqrt{6T \log 2}$$

Proof Applying expectation with respect to the randomness of the player to the definition of regret, we get:

$$\mathbb{E}[R_T] = \mathbb{E}\left[\sum_{t=1}^T X_t^\top l_t - \min_{X^* \in \{0,1\}^n} \sum_{t=1}^T X^{*\top} l_t\right] = \sum_{t=1}^T x_t^\top l_t - \min_{X^* \in \{0,1\}^n} \sum_{t=1}^T X^{*\top} l_t$$

Assuming $|\eta \tilde{l}_{t,i}| \leq 1$, we apply Lemma 19

$$\mathbb{E}[R_T] \leq \eta\mathbb{E}\left[\sum_{t=1}^T x_t^\top \tilde{l}_t^2\right] + \frac{n \log 2}{\eta} + 2\gamma nT$$

We have that:

$$\eta x_t^\top \tilde{l}_t^2 = \frac{1}{\eta} (\eta \tilde{l}_t)^\top \text{diag}(x_t) (\eta \tilde{l}_t) \leq \frac{\|\eta \tilde{l}_t\|_2^2}{\eta} \leq \frac{n}{\eta} \leq \frac{2n \log 2}{\eta}$$

This gives us:

$$\mathbb{E}[R_T] \leq \frac{3n \log 2}{\eta} + 2\gamma nT$$

To satisfy $|\eta \tilde{l}_{t,i}| \leq 1$, we need the following condition:

$$|\eta \tilde{l}_{t,i}| = \eta |\tilde{l}_t^\top e_i| = \eta |(P_t^{-1} X_t X_t^\top l_t)^\top e_i| \leq n\eta |X_t^\top P_t^{-1} e_i| \leq n\eta \|X_t^\top e_i\| \|P_t^{-1}\|$$

Since $P_t \succeq \frac{\gamma}{4} I_n$ and $|X_t^\top e_i| \leq 1$, we should have $\frac{4n\eta}{\gamma} \leq 1$. Taking $\gamma = 4n\eta$, we get:

$$\mathbb{E}[R_T] \leq \frac{3n \log 2}{\eta} + 8\eta n^2 T$$

Optimizing over η , we get $\mathbb{E}[R_T] \leq 2n^{3/2} \sqrt{24T \log 2}$ if $\eta = \sqrt{\frac{3 \log 2}{8nT}}$.

■

C.4 $\{-1, +1\}^n$ Hypercube Case

Lemma 21 *Exp2 on $\{-1, +1\}^n$ with losses l_t is equivalent to Exp2 on $\{0, 1\}^n$ with losses $2l_t$ while using the map $2X_t - \mathbf{1}$ to play on $\{-1, +1\}^n$.*

Proof Consider the update equation for Exp2 on $\{-1, +1\}^n$

$$p_{t+1}(Z) = \frac{\exp(-\eta \sum_{\tau=1}^t Z^\top l_\tau)}{\sum_{W \in \{-1, +1\}^n} \exp(-\eta \sum_{\tau=1}^t W^\top l_\tau)}$$

Using the fact that every $Z \in \{-1, +1\}^n$ can be mapped to a $X \in \{0, 1\}^n$ using the bijective map $X = (Z + \mathbf{1})/2$. So:

$$\begin{aligned} p_{t+1}(Z) &= \frac{\exp(-\eta \sum_{\tau=1}^t (2X - \mathbf{1})^\top l_\tau)}{\sum_{Y \in \{0, 1\}^n} \exp(-\eta \sum_{\tau=1}^t (2Y - \mathbf{1})^\top l_\tau)} \\ &= \frac{\exp(-\eta \sum_{\tau=1}^t X^\top (2l_\tau))}{\sum_{Y \in \{0, 1\}^n} \exp(-\eta \sum_{\tau=1}^t Y^\top (2l_\tau))} \end{aligned}$$

This is equivalent to updating the Exp2 on $\{0, 1\}^n$ with the loss vector $2l_t$. ■

Theorem 5 *Exp2 on $\{-1, +1\}^n$ using the sequence of losses l_t is equivalent to PolyExp on $\{0, 1\}^n$ using the sequence of losses $2\tilde{l}_t$. Moreover, the regret of Exp2 on $\{-1, +1\}^n$ will equal the regret of PolyExp using the losses $2\tilde{l}_t$.*

Proof After sampling X_t , we play $Z_t = 2X_t - \mathbf{1}$. So $\Pr(X_t = X) = \Pr(Z_t = 2X - \mathbf{1})$. In full information, $2\tilde{l}_t = 2l_t$ and in the bandit case $\mathbb{E}[2\tilde{l}_t] = 2l_t$. Since $2\tilde{l}_t$ is used in the bandit case to update the algorithm, by Lemma 21 we have that $\Pr(X_{t+1} = X) = \Pr(Z_{t+1} = 2X - \mathbf{1})$. By equivalence of Exp2 to PolyExp, the first statement follows immediately. Let $Z^* = \min_{Z \in \{-1, +1\}^n} \sum_{t=1}^T Z^\top l_t$ and $2X^* = Z^* + \mathbf{1}$. The regret of Exp2 on $\{-1, +1\}^n$ is:

$$\sum_{t=1}^T l_t^\top (Z_t - Z^*) = \sum_{t=1}^T l_t^\top (2X_t - \mathbf{1} - 2X^* + \mathbf{1}) = \sum_{t=1}^T (2l_t)^\top (X_t - X^*)$$

■