
Near-Optimal No-Regret Learning in General Games

Constantinos Daskalakis
MIT CSAIL
costis@csail.mit.edu

Maxwell Fishelson
MIT CSAIL
maxfish@mit.edu

Noah Golowich
MIT CSAIL
nzg@mit.edu

Abstract

We show that Optimistic Hedge – a common variant of multiplicative-weights-updates with recency bias – attains $\text{poly}(\log T)$ regret in multi-player general-sum games. In particular, when every player of the game uses Optimistic Hedge to iteratively update her strategy in response to the history of play so far, then after T rounds of interaction, *each player* experiences total regret that is $\text{poly}(\log T)$. Our bound improves, exponentially, the $O(T^{1/2})$ regret attainable by standard no-regret learners in games, the $O(T^{1/4})$ regret attainable by no-regret learners with recency bias [SALS15], and the $O(T^{1/6})$ bound that was recently shown for Optimistic Hedge in the special case of two-player games [CP20]. A corollary of our bound is that Optimistic Hedge converges to coarse correlated equilibrium in general games at a rate of $\tilde{O}(\frac{1}{T})$.

1 Introduction

Online learning has a long history that is intimately related to the development of game theory, convex optimization, and machine learning. One of its earliest instantiations can be traced to Brown’s proposal [Bro49] of fictitious play as a method to solve two-player zero-sum games. Indeed, as shown by [Rob51], when the players of (zero-sum) matrix game use fictitious play to iteratively update their actions in response to each other’s history of play, the resulting dynamics converge in the following sense: the product of the empirical distributions of strategies for each player converges to the set of Nash equilibria in the game, though the rate of convergence is now known to be exponentially slow [DP14]. Moreover, such convergence to Nash equilibria fails in non-zero-sum games [Sha64].

The slow convergence of fictitious play to Nash equilibria in zero-sum matrix games and non-convergence in general-sum games can be mitigated by appealing to the pioneering works [Bla54, Han57] and the ensuing literature on no-regret learning [CBL06]. It is known that if both players of a zero-sum matrix game experience regret that is at most $\varepsilon(T)$, the product of the players’ empirical distributions of strategies is an $O(\varepsilon(T)/T)$ -approximate Nash equilibrium. More generally, if each player of a general-sum, multi-player game experiences regret that is at most $\varepsilon(T)$, the empirical distribution of joint strategies converges to a coarse correlated equilibrium¹ of the game, at a rate of $O(\varepsilon(T)/T)$. Importantly, a multitude of online learning algorithms, such as the celebrated Hedge and Follow-The-Perturbed-Leader algorithms, guarantee adversarial regret $O(\sqrt{T})$ [CBL06]. Thus, when such algorithms are employed by all players in a game, their $O(\sqrt{T})$ regret implies convergence to coarse correlated equilibria (and Nash equilibria of matrix games) at a rate of $O(1/\sqrt{T})$.

While standard no-regret learners guarantee $O(\sqrt{T})$ regret for each player in a game, the players can do better by employing specialized no-regret learning procedures. Indeed, it was established by [DDK11] that there exists a somewhat complex no-regret learner based on Nesterov’s excessive gap technique [Nes05], which guarantees $O(\log T)$ regret to each player of a two-player zero-sum game.

¹In general-sum games, it is typical to focus on proving convergence rates for weaker types of equilibrium than Nash, such as coarse correlated equilibria, since finding Nash equilibria is PPAD-complete [DGP06, CDT09].

Table 1: Overview of prior work on fast rates for learning in games. m denotes the number of players, and n denotes the number of actions per player (assumed to be the same for all players). For Optimistic Hedge, the adversarial regret bounds in the right-hand column are obtained via a choice of adaptive step-sizes. The $\tilde{O}(\cdot)$ notation hides factors that are polynomial in $\log T$.

Algorithm	Setting	Regret in games	Adversarial regret
Hedge (& many other algs.)	multi-player, general-sum	$O(\sqrt{T \log n})$ [CBL06]	$O(\sqrt{T \log n})$ [CBL06]
Excessive Gap Technique	2-player, 0-sum	$O(\log n(\log T + \log^{3/2} n))$ [DDK11]	$O(\sqrt{T \log n})$ [DDK11]
DS-OptMD, OptDA	2-player, 0-sum	$\log^{O(1)}(n)$ [HAM21]	$\sqrt{T \log^{O(1)}(n)}$ [HAM21]
Optimistic Hedge	multi-player, general-sum	$O(\log n \cdot \sqrt{m} \cdot T^{1/4})$ [RS13b, SALS15]	$\tilde{O}(\sqrt{T \log n})$ [RS13b, SALS15]
Optimistic Hedge	2-player, general-sum	$O(\log^{5/6} n \cdot T^{1/6})$ [CP20]	$\tilde{O}(\sqrt{T \log n})$
Optimistic Hedge	multi-player, general-sum	$O(\log n \cdot m \cdot \log^4 T)$ (Theorem 3.1)	$\tilde{O}(\sqrt{T \log n})$ (Corollary D.1)

This represents an exponential improvement over the regret guaranteed by standard no-regret learners. More generally, [SALS15] established that if players of a multi-player, general-sum game use any algorithm from the family of Optimistic Mirror Descent (MD) or Optimistic Follow-the-Regularized-Leader (FTRL) algorithms (which are analogues of the MD and FTRL algorithms, respectively, with recency bias), each player enjoys regret that is $O(T^{1/4})$. This was recently improved by [CP20] to $O(T^{1/6})$ in the special case of two-player games in which the players use Optimistic Hedge, a particularly simple representative from both the Optimistic MD and Optimistic FTRL families.

The above results for general-sum games represent significant improvements over the $O(\sqrt{T})$ regret attainable by standard no-regret learners, but are not as dramatic as the logarithmic regret that has been shown attainable by no-regret learners, albeit more complex ones, in 2-player zero-sum games (e.g., [DDK11]). Indeed, despite extensive work on no-regret learning, understanding the optimal regret that can be guaranteed by no-regret learning algorithms in general-sum games has remained elusive. This question is especially intriguing in light of experiments suggesting that polylogarithmic regret should be attainable [SALS15, HAM21]. In this paper we settle this question by showing that no-regret learners can guarantee polylogarithmic regret to each player in general-sum multi-player games. Moreover, this regret is attainable by a particularly simple algorithm – Optimistic Hedge:

Theorem 1.1 (Abbreviated version of Theorem 3.1). *Suppose that m players play a general-sum multi-player game, with a finite set of n strategies per player, over T rounds. Suppose also that each player uses Optimistic Hedge to update her strategy in every round, as a function of the history of play so far. Then each player experiences $O(m \cdot \log n \cdot \log^4 T)$ regret.*

An immediate corollary of Theorem 1.1 is that the empirical distribution of play is a $O\left(\frac{m \log n \log^4 T}{T}\right)$ -approximate coarse correlated equilibrium (CCE) of the game. We remark that Theorem 1.1 bounds the total regret experienced by *each player* of the multi-player game, which is the most standard regret objective for no-regret learning in games, and which is essential to achieve convergence to CCE. For the looser objective of the *average* of all players' regrets, [RS13b] established a $O(\log n)$ bound for Optimistic Hedge in two-player zero-sum games, and [SALS15] generalized this bound, to $O(m \log n)$ in m -player general-sum games. Note that since some players may experience *negative regret* [HAM21], the average of the players' regrets cannot be used in general to bound the maximum regret experienced by any individual player. Finally, we remark that several results in the literature posit no-regret learning as a model of agents' rational behavior; for instance, [Rou09, ST13, RST17] show that no-regret learners in smooth games enjoy strong Price-of-Anarchy bounds. By showing that *each agent* can obtain very small regret in games by playing Optimistic Hedge, Theorem 1.1

strengthens the plausability of the common assumption made in this literature that each agent will choose to use such a no-regret algorithm.

1.1 Related work

Table 1 summarizes the prior works that aim to establish optimal regret bounds for no-regret learners in games. We remark that [CP20] shows that the regret of Hedge is $\Omega(\sqrt{T})$ even in 2-player games where each player has 2 actions, meaning that optimism is necessary to obtain fast rates. The table also includes a recent result of [HAM21] showing that when the players in a 2-player zero-sum game with n actions per player use a variant of Optimistic Hedge with adaptive step size (a special case of their algorithms DS-OptMD and OptDA), each player has $\log^{O(1)} n$ regret. The techniques of [HAM21] differ substantially from ours: the result in [HAM21] is based on showing that the joint strategies $x^{(t)}$ rapidly converge, pointwise, to a Nash equilibrium x^* . Such a result seems very unlikely to extend to our setting of general-sum games, since finding an approximate Nash equilibrium even in 2-player games is PPAD-complete [CDT09]. We also remark that the earlier work [KHSC18] shows that each player’s regret is at most $O(\log T \cdot \log n)$ when they use a certain algorithm based on Optimistic MD in 2-player zero-sum games; their technique is heavily tailored to 2-player zero-sum games, relying on the notion of duality in such a setting.

[FLL⁺16] shows that one can obtain fast rates in games for a broader class of algorithms (e.g., including Hedge) if one adopts a relaxed (approximate) notion of optimality. [WL18] uses optimism to obtain adaptive regret bounds for bandit problems. Many recent papers (e.g., [DP19, GPD20, LGNPw21, HAM21, WLZL21, AIMM21]) have studied the *last-iterate* convergence of algorithms from the Optimistic Mirror Descent family, which includes Optimistic Hedge. Finally, a long line of papers (e.g., [HMcW⁺03, DFP⁺10, KLP11, BCM12, PP16, BP18, MPP18, BP19, CP19, VGFL⁺20]) has studied the dynamics of learning algorithms in games. Essentially all of these papers do not use optimism, and many of them show *non-convergence* (e.g., divergence or recurrence) of the iterates of various learning algorithms such as FTRL and Mirror Descent when used in games.

2 Preliminaries

Notation. For a positive integer n , let $[n] := \{1, 2, \dots, n\}$. For a finite set \mathcal{S} , let $\Delta(\mathcal{S})$ denote the space of distributions on \mathcal{S} . For $\mathcal{S} = [n]$, we will write $\Delta^n := \Delta(\mathcal{S})$ and interpret elements of Δ^n as vectors in \mathbb{R}^n . For a vector $v \in \mathbb{R}^n$ and $j \in [n]$, we denote the j th coordinate of v as $v(j)$. For vectors $v, w \in \mathbb{R}^n$, write $\langle v, w \rangle = \sum_{j=1}^n v(j)w(j)$. The base-2 logarithm of $x > 0$ is denoted $\log x$.

No-regret learning in games. We consider a game G with $m \in \mathbb{N}$ players, where player $i \in [m]$ has *action space* \mathcal{A}_i with $n_i := |\mathcal{A}_i|$ actions. We may assume that $\mathcal{A}_i = [n_i]$ for each player i . The *joint action space* is $\mathcal{A} := \mathcal{A}_1 \times \dots \times \mathcal{A}_m$. The specification of the game G is completed by a collection of *loss functions* $\mathcal{L}_1, \dots, \mathcal{L}_m : \mathcal{A} \rightarrow [0, 1]$. For an action profile $a = (a_1, \dots, a_m) \in \mathcal{A}$ and $i \in [m]$, $\mathcal{L}_i(a)$ is the loss player i experiences when each player $i' \in [m]$ plays $a_{i'}$. A *mixed strategy* $x_i \in \Delta(\mathcal{A}_i)$ for player i is a distribution over \mathcal{A}_i , with the probability of playing action $j \in \mathcal{A}_i$ given by $x_i(j)$. Given a mixed strategy profile $x = (x_1, \dots, x_m)$ (or an action profile $a = (a_1, \dots, a_m)$) and a player $i \in [m]$ we let x_{-i} (or a_{-i} , respectively) denote the profile after removing the i th mixed strategy x_i (or the i th action a_i , respectively).

The m players play the game G for a total of T rounds. At the beginning of each round $t \in [T]$, each player i chooses a mixed strategy $x_i^{(t)} \in \Delta(\mathcal{A}_i)$. The *loss vector* of player i , denoted $\ell_i^{(t)} \in [0, 1]^{n_i}$, is defined as $\ell_i^{(t)}(j) = \mathbb{E}_{a_{-i} \sim x_{-i}^{(t)}}[\mathcal{L}_i(j, a_{-i})]$. As a matter of convention, set $\ell_i^{(0)} = \mathbf{0}$ to be the all-zeros vector. We consider the *full-information setting* in this paper, meaning that player i observes its full loss vector $\ell_i^{(t)}$ for each round t . Finally, player i experiences a loss of $\langle \ell_i^{(t)}, x_i^{(t)} \rangle$. The goal of each player i is to minimize its *regret*, defined as: $\text{Reg}_{i,T} := \sum_{t \in [T]} \langle x_i^{(t)}, \ell_i^{(t)} \rangle - \min_{j \in [n_i]} \sum_{t \in [T]} \ell_i^{(t)}(j)$.

Optimistic hedge. The Optimistic Hedge algorithm chooses mixed strategies for player $i \in [m]$ as follows: at time $t = 1$, it sets $x_i^{(1)} = (1/n_i, \dots, 1/n_i)$ to be the uniform distribution on \mathcal{A}_i . Then for all $t < T$, player i ’s strategy at iteration $t + 1$ is defined as follows, for $j \in [n_i]$:

$$x_i^{(t+1)}(j) := \frac{x_i^{(t)}(j) \cdot \exp(-\eta \cdot (2\ell_i^{(t)}(j) - \ell_i^{(t-1)}(j)))}{\sum_{k \in [n_i]} x_i^{(t)}(k) \cdot \exp(-\eta \cdot (2\ell_i^{(t)}(k) - \ell_i^{(t-1)}(k)))}. \quad (1)$$

Optimistic Hedge is a modification of Hedge, which performs the updates $x_i^{(t+1)}(j) := \frac{x_i^{(t)}(j) \cdot \exp(-\eta \cdot \ell_i^{(t)}(j))}{\sum_{k \in [n_i]} x_i^{(t)}(k) \cdot \exp(-\eta \cdot \ell_i^{(t)}(k))}$. The update (1) modifies the Hedge update by replacing the loss vector $\ell_i^{(t)}$ with a predictor of the *following iteration's* loss vector, $\ell_i^{(t)} + (\ell_i^{(t)} - \ell_i^{(t-1)})$. Hedge corresponds to FTRL with a negative entropy regularizer (see, e.g., [Bub15]), whereas Optimistic Hedge corresponds to *Optimistic* FTRL with a negative entropy regularizer [RS13b, RS13a].

Distributions & divergences. For distributions P, Q on a finite domain $[n]$, the *KL divergence* between P, Q is $\text{KL}(P; Q) = \sum_{j=1}^n P(j) \cdot \log \left(\frac{P(j)}{Q(j)} \right)$. The *chi-squared divergence* between P, Q is $\chi^2(P; Q) = \sum_{j=1}^n Q(j) \cdot \left(\frac{P(j)}{Q(j)} \right)^2 - 1 = \sum_{j=1}^n \frac{(P(j) - Q(j))^2}{Q(j)}$. For a distribution P on $[n]$ and a vector $v \in \mathbb{R}^n$, we write $\text{Var}_P(v) := \sum_{j=1}^n P(j) \cdot (v(j) - \sum_{k=1}^n P(k)v(k))^2$. Also define $\|v\|_P := \sqrt{\sum_{j=1}^n P(j) \cdot v(j)^2}$. If further P has full support, then define $\|v\|_P^* = \sqrt{\sum_{j=1}^n \frac{v(j)^2}{P(j)}}$. The above notations will often be used when P is the mixed strategy profile x_i for some player i and v is a loss vector ℓ_i ; in such a case the norms $\|v\|_P$ and $\|v\|_P^*$ are often called *local norms*.

3 Results

Below we state our main theorem, which shows that when all players in a game play according to Optimistic Hedge with appropriate step size, they all experience polylogarithmic individual regrets.

Theorem 3.1 (Formal version of Theorem 1.1). *There are constants $C, C' > 1$ so that the following holds. Suppose a time horizon $T \in \mathbb{N}$ and a game G with m players and n_i actions for each player $i \in [m]$ is given. Suppose all players play according to Optimistic Hedge with any positive step size $\eta \leq \frac{1}{C \cdot m \log^4 T}$. Then for any $i \in [m]$, the regret of player i satisfies*

$$\text{Reg}_{i,T} \leq \frac{\log n_i}{\eta} + C' \cdot \log T. \quad (2)$$

In particular, if the players' step size is chosen as $\eta = \frac{1}{C \cdot m \log^4 T}$, then the regret of player i satisfies

$$\text{Reg}_{i,T} \leq O(m \cdot \log n_i \cdot \log^4 T). \quad (3)$$

A common goal in the literature on learning in games is to obtain an algorithm that achieves fast rates when played by all players, and so that each player i still obtains the optimal rate of $O(\sqrt{T})$ in the adversarial setting (i.e., when i receives an arbitrary sequence of losses $\ell_i^{(1)}, \dots, \ell_i^{(T)}$). We show in Corollary D.1 (in the appendix) that by running Optimistic Hedge with an adaptive step size, this is possible. Table 1 compares our regret bounds discussed in this section to those of prior work.

4 Proof overview

In this section we overview the proof of Theorem 3.1; the full proof may be found in the appendix.

4.1 New adversarial regret bound

The first step in the proof of Theorem 3.1 is to prove a new regret bound (Lemma 4.1 below) for Optimistic Hedge that holds for an adversarial sequence of losses. We will show in later sections that when *all* players play according to Optimistic Hedge, the right-hand side of the regret bound (4) is bounded by a quantity that grows only poly-logarithmically in T .

Lemma 4.1. *There is a constant $C > 0$ so that the following holds. Suppose any player $i \in [m]$ follows the Optimistic Hedge updates (1) with step size $\eta < 1/C$, for an arbitrary sequence of losses $\ell_i^{(1)}, \dots, \ell_i^{(T)} \in [0, 1]^{n_i}$. Then*

$$\text{Reg}_{i,T} \leq \frac{\log n_i}{\eta} + \sum_{t=1}^T \left(\frac{\eta}{2} + C\eta^2 \right) \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t)} - \ell_i^{(t-1)} \right) - \sum_{t=1}^T \frac{(1 - C\eta)\eta}{2} \cdot \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t-1)} \right). \quad (4)$$

The detailed proof of Lemma 4.1 can be found in Section A, but we sketch the main steps here. The starting point is a refinement of [RS13a, Lemma 3] (stated as Lemma A.5), which gives an upper

bound for $\text{Reg}_{i,T}$ in terms of local norms corresponding to each of the iterates $x_i^{(t)}$ of Optimistic Hedge. The bound involves the difference between the Optimistic Hedge iterates $x_i^{(t)}$ and iterates $\tilde{x}_i^{(t)}$ defined by $\tilde{x}_i^{(t)} = \frac{x_i^{(t)}(j) \cdot \exp(-\eta \cdot (\ell_i^{(t)}(j) - \ell_i^{(t-1)}(j)))}{\sum_{k \in [n_i]} x_i^{(t)}(k) \cdot \exp(-\eta \cdot (\ell_i^{(t)}(k) - \ell_i^{(t-1)}(k)))}$:

$$\text{Reg}_{i,T} \leq \frac{\log n_i}{\eta} + \sum_{t=1}^T \left\| x_i^{(t)} - \tilde{x}_i^{(t)} \right\|_{x_i^{(t)}}^* \sqrt{\text{Var}_{x_i^{(t)}} \left(\ell_i^{(t)} - \ell_i^{(t-1)} \right)} - \frac{1}{\eta} \sum_{t=1}^T \text{KL}(\tilde{x}_i^{(t)}; x_i^{(t)}) - \frac{1}{\eta} \sum_{t=1}^T \text{KL}(x_i^{(t)}; \tilde{x}_i^{(t-1)}). \quad (5)$$

We next show (in Lemma A.2) that $\text{KL}(\tilde{x}_i^{(t)}; x_i^{(t)})$ and $\text{KL}(x_i^{(t)}; \tilde{x}_i^{(t-1)})$ may be lower bounded by $(1/2 - O(\eta)) \cdot \chi^2(\tilde{x}_i^{(t)}; x_i^{(t)})$ and $(1/2 - O(\eta)) \cdot \chi^2(x_i^{(t)}; \tilde{x}_i^{(t-1)})$, respectively. Note it is a standard fact that the KL divergence between two distributions is upper bounded by the chi-squared distribution between them; by contrast, Lemma A.2 can exploit that $x_i^{(t)}$, $\tilde{x}_i^{(t)}$ and $\tilde{x}_i^{(t-1)}$ are close to each other to show a reverse inequality. Finally, exploiting the exponential weights-style functional relationship between $x_i^{(t)}$ and $\tilde{x}_i^{(t-1)}$, we show (in Lemma A.3) that the χ^2 -divergence $\chi^2(x_i^{(t)}; \tilde{x}_i^{(t-1)})$ may be lower bounded by $(1 - O(\eta)) \cdot \eta^2 \cdot \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t-1)} \right)$, leading to the term $\frac{(1-C\eta)\eta}{2} \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t-1)} \right)$ being subtracted in (4). The χ^2 -divergence $\chi^2(\tilde{x}_i^{(t)}; x_i^{(t)})$, as well as the term $\left\| x_i^{(t)} - \tilde{x}_i^{(t)} \right\|_{x_i^{(t)}}^*$ in (5) are bounded in a similar manner to obtain (4).

4.2 Finite differences

Given Lemma 4.1, in order to establish Theorem 3.1, it suffices to show Lemma 4.2 below. Indeed, (6) below implies that the right-hand side of (4) is bounded above by $\frac{\log n_i}{\eta} + \eta \cdot O(\log^5 T)$, which is bounded above by $O(m \log n_i \log^4 T)$ for the choice $\eta = \Theta \left(\frac{1}{m \cdot \log^4 T} \right)$ of Theorem 3.1.²

Lemma 4.2 (Abbreviated; detailed version in Section C.3). *Suppose all players play according to Optimistic Hedge with step size η satisfying $1/T \leq \eta \leq \frac{1}{Cm \cdot \log^4 T}$ for a sufficiently large constant C . Then for any $i \in [m]$, the losses $\ell_i^{(1)}, \dots, \ell_i^{(T)} \in \mathbb{R}^{n_i}$ for player i satisfy:*

$$\sum_{t=1}^T \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t)} - \ell_i^{(t-1)} \right) \leq \frac{1}{2} \cdot \sum_{t=1}^T \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t-1)} \right) + O(\log^5 T). \quad (6)$$

The definition below allows us to streamline our notation when proving Lemma 4.2.

Definition 4.1 (Finite differences). Suppose $L = (L^{(1)}, \dots, L^{(T)})$ is a sequence of vectors $L^{(t)} \in \mathbb{R}^n$. For integers $h \geq 0$, the *order- h finite difference sequence* for the sequence L , denoted by $D_h L$, is the sequence $D_h L := ((D_h L)^{(1)}, \dots, (D_h L)^{(T-h)})$ defined recursively as: $(D_0 L)^{(t)} := L^{(t)}$ for all $1 \leq t \leq T$, and

$$(D_h L)^{(t)} := (D_{h-1} L)^{(t+1)} - (D_{h-1} L)^{(t)} \quad (7)$$

for all $h \geq 1, 1 \leq t \leq T - h$.³

Remark 4.3. Notice that another way of writing (7) is: $D_h L = D_1 D_{h-1} L$. We also remark for later use that $(D_h L)^{(t)} = \sum_{s=0}^h \binom{h}{s} (-1)^{h-s} L^{(t+s)}$.

Let $H = \log T$, where T denotes the fixed time horizon from Theorem 3.1 (and thus Lemma 4.2). In the proof of Lemma 4.2, we will bound the finite differences of order $h \leq H$ for certain sequences. The bound (6) of Lemma 4.2 may be rephased as upper bounding $\sum_{t=1}^T \text{Var}_{x_i^{(t)}} \left((D_1 \ell_i)^{(t-1)} \right)$, by $\frac{1}{2} \sum_{t=1}^T \text{Var}_{x_i} \left(\ell_i^{(t-1)} \right)$; to prove this, we proceed in two steps:

²Notice that the factor $\frac{1}{2}$ in (6) is not important for this argument – any constant less than 1 would suffice.

³We remark that while Definition 4.1 is stated for a 1-indexed sequence $L^{(1)}, L^{(2)}, \dots$, we will also occasionally consider 0-indexed sequences $L^{(0)}, L^{(1)}, \dots$, in which case the same recursive definition (7) holds for the finite differences $(D_h L)^{(t)}, t \geq 0$.

1. (*Upwards induction step*) First, in Lemma 4.4 below, we find an upper bound on $\left\| (D_h \ell_i)^{(t)} \right\|_\infty$ for all $t \in [T]$, $h \geq 0$, which decays exponentially in h for $h \leq H$. This is done via *upwards induction* on h , i.e., first proving the base case $h = 0$ using boundedness of the losses $\ell_i^{(t)}$ and then $h = 1, 2, \dots$ inductively. The main technical tool we develop for the inductive step is a weak form of the chain rule for finite differences, Lemma 4.5. The inductive step uses the fact that all players are following Optimistic Hedge to relate the h th order finite differences of player i 's loss sequence $\ell_i^{(t)}$ to the h th order finite differences of the strategy sequences $x_{i'}^{(t)}$ for players $i' \neq i$; then we use the exponential-weights style updates of Optimistic Hedge and Lemma 4.5 to relate the h th order finite differences of the strategies $x_{i'}^{(t)}$ to the $(h - 1)$ th order finite differences of the losses $\ell_{i'}^{(t)}$.
2. (*Downwards induction step*) We next show that for all $0 \leq h \leq H$, $\sum_{t=1}^T \text{Var}_{x_i^{(t)}} \left((D_{h+1} \ell_i)^{(t-1)} \right)$ is bounded above by $c_h \cdot \sum_{t=1}^T \text{Var}_{x_i^{(t)}} \left((D_h \ell_i)^{(t-1)} \right) + \mu_h$, for some $c_h < 1/2$ and $\mu_h < O(\log^5 T)$. This shown via *downwards induction* on h , namely first establishing the base case $h = H$ by using the result of item 1 for $h = H$ and then treating the cases $h = H - 1, H - 2, \dots, 0$. The inductive step makes use of the discrete Fourier transform (DFT) to relate the finite differences of different orders (see Lemmas 4.7 and 4.8). In particular, Parseval's equality together with a standard relationship between the DFT of the finite differences of a sequence to the DFT of that sequence allow us to first prove the inductive step in the frequency domain and then transport it back to the original (time) domain.

In the following subsections we explain in further detail how the two steps above are completed.

4.3 Upwards induction proof overview

Addressing item 1 in the previous subsection, the lemma below gives a bound on the supremum norm of the h -th order finite differences of each player's loss vector, when all players play according to Optimistic Hedge and experience losses according to their loss functions $\mathcal{L}_1, \dots, \mathcal{L}_m : \mathcal{A} \rightarrow [0, 1]$.

Lemma 4.4 (Abbreviated). *Fix a step size $\eta > 0$ satisfying $\eta \leq o\left(\frac{1}{m \log T}\right)$. If all players follow Optimistic Hedge updates with step size η , then for any player $i \in [m]$, integer h satisfying $0 \leq h \leq H$, and time step $t \in [T - h]$, it holds that $\| (D_h \ell_i)^{(t)} \|_\infty \leq O(m\eta)^h \cdot h^{O(h)}$.*

A detailed version of Lemma 4.4, together with its full proof, may be found in Section B.4. We next give a proof overview of Lemma 4.4 for the case of 2 players, i.e., $m = 2$; we show in Section B.4 how to generalize this computation to general m . Below we introduce the main technical tool in the proof, a ‘‘boundedness chain rule,’’ and then outline how it is used to prove Lemma 4.4.

Main technical tool for Lemma 4.4: boundedness chain rule. We say that a function $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$ is a *softmax-type* function if there are real numbers ξ_1, \dots, ξ_n and some $j \in [n]$ so that for all $(z_1, \dots, z_n) \in \mathbb{R}^n$, $\phi((z_1, \dots, z_n)) = \frac{\exp(z_j)}{\sum_{k=1}^n \xi_k \cdot \exp(z_k)}$. Lemma 4.5 below may be interpreted as a ‘‘boundedness chain rule’’ for finite differences. To explain the context for this lemma, recall that given an infinitely differentiable vector-valued function $L : \mathbb{R} \rightarrow \mathbb{R}^n$ and an infinitely differentiable function $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$, the higher order derivatives of the function $\phi(L(t))$ may be computed in terms of those of L and ϕ using the chain rule. Lemma 4.5 considers an analogous setting where the input variable t to L is discrete-valued, taking values in $[T]$ (and so we identify the function L with the sequence $L^{(1)}, \dots, L^{(T)}$). In this case, the *higher order finite differences* of the sequence $L^{(1)}, \dots, L^{(T)}$ (Definition 4.1) take the place of the higher order derivatives of L with respect to t . Though there is no generic chain rule for finite differences, Lemma 4.5 states that, at least when ϕ is a softmax-type function, we may *bound* the higher order finite differences of the sequence $\phi(L^{(1)}), \dots, \phi(L^{(T)})$. In the lemma's statement we let $\phi \circ L$ denote the sequence $\phi(L^{(1)}), \dots, \phi(L^{(T)})$.

Lemma 4.5 (‘‘Boundedness chain rule’’ for finite differences; abbreviated). *Suppose that $h, n \in \mathbb{N}$, $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$ is a softmax-type function, and $L = (L^{(1)}, \dots, L^{(T)})$ is a sequence of vectors in \mathbb{R}^n satisfying $\|L^{(t)}\|_\infty \leq 1$ for $t \in [T]$. Suppose for some $\alpha \in (0, 1)$, for each $0 \leq h' \leq h$ and $t \in [T - h']$, it holds that $\|D_{h'} L^{(t)}\|_\infty \leq O(\alpha^{h'}) \cdot (h')^{O(h')}$. Then for all $t \in [T - h]$,*

$$\| (D_h (\phi \circ L))^{(t)} \| \leq O(\alpha^h) \cdot h^{O(h)}.$$

A detailed version of Lemma 4.5 may be found in Section B.3. While Lemma 4.5 requires ϕ to be a softmax-type function for simplicity (and this is the only type of function ϕ we will need to consider

for the case $m = 2$) we remark that the detailed version of Lemma 4.5 allows ϕ to be from a more general family of analytic functions whose higher order derivatives are appropriately bounded. The proof of Lemma 4.4 for all $m \geq 2$ requires that more general form of Lemma 4.5.

The proof of Lemma 4.5 proceeds by considering the Taylor expansion $P_\phi(\cdot)$ of the function ϕ at the origin, which we write as follows: for $z = (z_1, \dots, z_n) \in \mathbb{R}^n$, $P_\phi(z) := \sum_{k \geq 0, \gamma \in \mathbb{Z}_{\geq 0}^n: |\gamma|=k} a_\gamma z^\gamma$, where $a_\gamma \in \mathbb{R}$, $|\gamma|$ denotes the quantity $\gamma_1 + \dots + \gamma_n$ and z^γ denotes $z_1^{\gamma_1} \dots z_n^{\gamma_n}$. The fact that ϕ is a softmax-type function ensures that the radius of convergence of its Taylor series is at least 1, i.e., $\phi(z) = P_\phi(z)$ for any z satisfying $\|z\|_\infty \leq 1$. By the assumption that $\|L^{(t)}\|_\infty \leq 1$ for each t , we may therefore decompose $(D_h(\phi \circ L))^{(t)}$ as:

$$(D_h(\phi \circ L))^{(t)} = \sum_{k \geq 0, \gamma \in \mathbb{Z}_{\geq 0}^n: |\gamma|=k} a_\gamma \cdot (D_h L^\gamma)^{(t)}, \quad (8)$$

where L^γ denotes the sequence of scalars $(L^\gamma)^{(t)} := (L^{(t)})^\gamma$ for all t . The fact that ϕ is a softmax-type function allows us to establish strong bounds on $|a_\gamma|$ for each γ in Lemma B.5. The proof of Lemma B.5 bounds the $|a_\gamma|$ by exploiting the simple form of the derivative of a softmax-type function to decompose each a_γ into a sum of $|\gamma|!$ terms. Then we establish a bijection between the terms of this decomposition and graph structures we refer to as *factorial trees*; that bijection together with the use of an appropriate generating function allow us to complete the proof of Lemma B.5.

Thus, to prove Lemma 4.5, it suffices to bound $|(D_h L^\gamma)^{(t)}|$ for all γ . We do so by using Lemma 4.6.

Lemma 4.6 (Abbreviated; detailed version in Section B.2). *Fix any $h \geq 0$, a multi-index $\gamma \in \mathbb{Z}_{\geq 0}^n$ and set $k = |\gamma|$. For each of the k^h functions $\pi: [h] \rightarrow [k]$, and for each $r \in [k]$, there are integers $h'_{\pi,r} \in \{0, 1, \dots, h\}$, $t'_{\pi,r} \geq 0$, and $j'_{\pi,r} \in [n]$, so that the following holds. For any sequence $L^{(1)}, \dots, L^{(T)} \in \mathbb{R}^n$ of vectors, it holds that, for each $t \in [T - h]$,*

$$(D_h L^\gamma)^{(t)} = \sum_{\pi: [h] \rightarrow [k]} \prod_{r=1}^k (D_{h'_{\pi,r}} L(j'_{\pi,r}))^{(t+t'_{\pi,r})}. \quad (9)$$

Lemma 4.6 expresses the h th order finite differences of the sequence L^γ as a sum of k^h terms, each of which is a product of k finite order differences of a sequence $L^{(t)}(j'_{\pi,r})$ (i.e., the $j'_{\pi,r}$ th coordinate of the vectors $L^{(t)}$). Crucially, when using Lemma 4.6 to prove Lemma 4.5, the assumption of Lemma 4.5 gives that for each $j' \in [n]$, each $h' \in [h]$, and each $t' \in [T - h']$, we have the bound $|(D_{h'} L(j'))^{(t')}| \leq O(\alpha^{h'}) \cdot (h')^{O(h')}$. These assumed bounds may be used to bound the right-hand side of (9), which together with Lemma 4.6 and (8) lets us complete the proof of Lemma 4.5.

Proving Lemma 4.4 using the boundedness chain rule. Next we discuss how Lemma 4.5 is used to prove Lemma 4.4, namely to bound $\|(D_h \ell_i)^{(t)}\|_\infty$ for each $t \in [T - h]$, $i \in [m]$, and $0 \leq h \leq H$. Lemma 4.4 is proved using induction, with the base case $h = 0$ being a straightforward consequence of the fact that $\|(D_0 \ell_i)^{(t)}\|_\infty = \|\ell_i^{(t)}\|_\infty \leq 1$ for all $i \in [m]$, $t \in [T]$. For the rest of this section we focus on the inductive case, i.e., we pick some $h \in [H]$ and assume Lemma 4.4 holds for all $h' < h$.

The first step is to reduce the claim of Lemma 4.4 to the claim that the upper bound $\|(D_h x_i)^{(t)}\|_1 \leq O(m\eta)^h \cdot h^{O(h)}$ holds for each $t \in [T - h]$, $i \in [m]$. Recalling that we are only sketching here the case $m = 2$ for simplicity, this reduction proceeds as follows: for $i \in \{1, 2\}$, define the matrix $A_i \in \mathbb{R}^{n_1 \times n_2}$ by $(A_i)_{a_1 a_2} = \mathcal{L}_i(a_1, a_2)$, for $a_1 \in [n_1]$, $a_2 \in [n_2]$. We have assumed that all players are using Optimistic Hedge and thus $\ell_i^{(t)} = \mathbb{E}_{a_{i'} \sim x_{i'}^{(t)}, \forall i' \neq i} [\mathcal{L}_i(a_1, \dots, a_n)]$; for our case here ($m = 2$), this may be rewritten as $\ell_1^{(t)} = A_1 x_2^{(t)}$, $\ell_2^{(t)} = A_2^\top x_1^{(t)}$. Thus

$$\|(D_h \ell_1)^{(t)}\|_\infty = \left\| A_1 \cdot \sum_{s=0}^h \binom{h}{s} (-1)^{h-s} x_2^{(t+s)} \right\|_\infty \leq \left\| \sum_{s=0}^h \binom{h}{s} (-1)^{h-s} x_2^{(t+s)} \right\|_1 = \|(D_h x_2)^{(t)}\|_1,$$

where the first equality is from Remark 4.3 and the inequality follows since all entries of A_1 have absolute value ≤ 1 . A similar computation allows us to show $\|(D_h \ell_2)^{(t)}\|_\infty \leq \|(D_h x_1)^{(t)}\|_1$.

To complete the inductive step it remains to upper bound the quantities $\|(\mathbb{D}_h x_i)^{(t)}\|_1$ for $i \in [m]$, $t \in [T-h]$. To do so, we note that the definition of the Optimistic Hedge updates (1) implies that for any $i \in [m]$, $t \in [T]$, $j \in [n_i]$, and $t' \geq 1$, we have

$$x_i^{(t+t')}(j) = \frac{x_i^{(t)}(j) \cdot \exp\left(\eta \cdot \left(\ell_i^{(t-1)}(j) - \sum_{s=0}^{t'-1} \ell_i^{(t+s)}(j) - \ell_i^{(t+t'-1)}(j)\right)\right)}{\sum_{k=1}^{n_i} x_i^{(t)}(k) \cdot \exp\left(\eta \cdot \left(\ell_i^{(t-1)}(k) - \sum_{s=0}^{t'-1} \ell_i^{(t+s)}(k) - \ell_i^{(t+t'-1)}(k)\right)\right)}. \quad (10)$$

For $t \in [T]$, $t' \geq 0$, set $\bar{\ell}_{i,t}^{(t')} := \eta \cdot \left(\ell_i^{(t-1)} - \sum_{s=0}^{t'-1} \ell_i^{(t+s)} - \ell_i^{(t+t'-1)}\right)$. Also, for each $i \in [m]$, $j \in [n_i]$, $t \in [T]$, and any vector $z = (z(1), \dots, z(n_i)) \in \mathbb{R}^{n_i}$ define $\phi_{t,i,j}(z) := \frac{x_i^{(t)}(j) \cdot \exp(z(j))}{\sum_{k=1}^{n_i} x_i^{(t)}(k) \cdot \exp(z(k))}$.

Thus (10) gives that for $t' \geq 1$, $x_i^{(t+t')}(j) = \phi_{t,i,j}(\bar{\ell}_{i,t}^{(t')})$. Viewing t as a fixed parameter and letting t' vary, it follows that for $h \geq 0$ and $t' \geq 1$, $(\mathbb{D}_h x_i^{(t+t')}(j))^{(t')} = (\mathbb{D}_h (\phi_{t,i,j} \circ \bar{\ell}_{i,t}^{(t')}))^{(t')}$.

Recalling that our goal is to bound $|(\mathbb{D}_h x_i(j))^{(t+1)}|$ for each t , we can do so by using Lemma 4.5 with $\phi = \phi_{t,i,j}$ and $\alpha = O(m\eta)$, if we can show that its precondition is met, i.e. that $\|(\mathbb{D}_{h'} \bar{\ell}_{i,t}^{(t')})^{(t')}\|_\infty \leq \frac{1}{B_1} \cdot \alpha^{h'} \cdot (h')^{B_0 h'}$ for all $h' \leq h$, the appropriate α and appropriate constants B_0, B_1 . Helpfully, the definition of $\bar{\ell}_{i,t}^{(t')}$ as a partial sum allows us to relate the h' -th order finite differences of the sequence $\bar{\ell}_{i,t}^{(t')}$ to the $(h'-1)$ -th order finite differences of the sequence $\ell_i^{(t)}$ as follows:

$$(\mathbb{D}_{h'} \bar{\ell}_{i,t}^{(t')})^{(t')} = \eta \cdot (\mathbb{D}_{h'-1} \ell_i)^{(t+t'-1)} - 2\eta \cdot (\mathbb{D}_{h'-1} \ell_i)^{(t+t')}. \quad (11)$$

Since $h' - 1 < h$ for $h' \leq h$, the inductive assumption of Lemma 4.4 gives a bound on the ℓ_∞ -norm of the terms on the right-hand side of (11), which are sufficient for us to apply Lemma 4.5. Note that the inductive assumption gives an upper bound on $\|(\mathbb{D}_{h'-1} \ell_i)^{(t)}\|_\infty$ that only scales with $\alpha^{h'-1}$, whereas Lemma 4.5 requires scaling of $\alpha^{h'}$. This discrepancy is corrected by the factor of η on the right-hand side of (11), which gives the desired scaling $\alpha^{h'}$ (since $\eta < \alpha$ for the choice $\alpha = O(m\eta)$).

4.4 Downwards induction proof overview

In this section we discuss in further detail item 2 in Section 4.2; in particular, we will show that there is a parameter $\mu = \tilde{\Theta}(\eta m)$ so that for all integers h satisfying $H - 1 \geq h \geq 0$,

$$\sum_{t=1}^{T-h-1} \text{Var}_{x_i^{(t)}} \left((\mathbb{D}_{h+1} \ell_i)^{(t)} \right) \leq O(1/H) \cdot \sum_{t=1}^{T-h} \text{Var}_{x_i^{(t)}} \left((\mathbb{D}_h \ell_i)^{(t)} \right) + \tilde{O}(\mu^{2h}), \quad (12)$$

where \tilde{O} hides factors polynomial in $\log T$. The validity of (12) for $h = 0$ implies Lemma 4.2. On the other hand, as long we choose the value μ in (12) to satisfy $\mu \geq m\eta H^{\Omega(1)}$, then Lemma 4.4 implies that $\sum_{t=1}^{T-H} \text{Var}_{x_i^{(t)}} \left((\mathbb{D}_H \ell_i)^{(t)} \right) \leq O(\mu^{2H})$. This gives that (12) holds for $h = H - 1$. To show that (12) holds for all $H - 1 > h \geq 0$, we use downwards induction; fix any h , and assume that (12) has been shown for all h' satisfying $h < h' \leq H - 1$. Our main tool in the inductive step is to apply Lemma 4.7 below. To state it, for $\zeta > 0$, $n \in \mathbb{N}$, we say that a sequence of distributions $P^{(1)}, \dots, P^{(T)} \in \Delta^n$ is ζ -consecutively close if for each $1 \leq t < T$, it holds that $\max \left\{ \left\| \frac{P^{(t)}}{P^{(t+1)}} \right\|_\infty, \left\| \frac{P^{(t+1)}}{P^{(t)}} \right\|_\infty \right\} \leq 1 + \zeta$.⁴ Lemma 4.7 shows that given a sequence of vectors for which the variances of its second-order finite differences are bounded by the variances of its first-order finite differences, a similar relationship holds between its first- and zeroth-order finite differences.

Lemma 4.7. *There is a sufficiently large constant $C_0 > 1$ so that the following holds. For any $M, \zeta, \alpha > 0$ and $n \in \mathbb{N}$, suppose that $P^{(1)}, \dots, P^{(T)} \in \Delta^n$ and $Z^{(1)}, \dots, Z^{(T)} \in [-M, M]^n$ satisfy the following conditions:*

1. *The sequence $P^{(1)}, \dots, P^{(T)}$ is ζ -consecutively close for some $\zeta \in [1/(2T), \alpha^4/C_0]$.*
2. *It holds that $\sum_{t=1}^{T-2} \text{Var}_{P^{(t)}} \left((\mathbb{D}_2 Z)^{(t)} \right) \leq \alpha \cdot \sum_{t=1}^{T-1} \text{Var}_{P^{(t)}} \left((\mathbb{D}_1 Z)^{(t)} \right) + \mu$.*

⁴Here, for distributions $P, Q \in \Delta^n$, $\frac{P}{Q} \in \mathbb{R}^n$ denotes the vector whose j th entry is $P(j)/Q(j)$.

Then $\sum_{t=1}^{T-1} \text{Var}_{P^{(t)}} \left((D_1 Z)^{(t)} \right) \leq \alpha \cdot (1 + \alpha) \sum_{t=1}^T \text{Var}_{P^{(t)}} (Z^{(t)}) + \frac{\mu}{\alpha} + \frac{C_0 M^2}{\alpha^3}$.

Given Lemma 4.7, the inductive step for establishing (12) is straightforward: we apply Lemma 4.7 with $P^{(t)} = x_i^{(t)}$ and $Z^{(t)} = (D_h \ell_i)^{(t)}$ for all t . The fact that $x_i^{(t)}$ are updated with Optimistic Hedge may be used to establish that precondition 1 of Lemma 4.7 holds. Since $(D_1 Z)^{(t)} = (D_{h+1} \ell_i)^{(t)}$ and $(D_2 Z)^{(t)} = (D_{h+2} \ell_i)^{(t)}$, that the inductive hypothesis (12) holds for $h+1$ implies that precondition 2 of Lemma 4.7 holds for appropriate $\alpha, \mu > 0$. Thus Lemma 4.7 implies that (12) holds for the value h , which completes the inductive step.

On the proof of Lemma 4.7. Finally we discuss the proof of Lemma 4.7. One technical challenge is the fact that the vectors $P^{(t)}$ are not constant functions of t , but rather change slowly (as constrained by being ζ -consecutively close). The main tool for dealing with this difficulty is Lemma C.1, which shows that for a ζ -consecutively close sequence $P^{(t)}$, for any vector $Z^{(t)}$, $\frac{\text{Var}_{P^{(t)}}(Z^{(t)})}{\text{Var}_{P^{(t+1)}}(Z^{(t)})} \in [1 - \zeta, 1 + \zeta]$. This fact, together with some algebraic manipulations, lets us to reduce to the case that all $P^{(t)}$ are equal. It is also relatively straightforward to reduce to the case that $\langle P^{(t)}, Z^{(t)} \rangle = 0$ for all t , i.e., so that $\text{Var}_{P^{(t)}}(Z^{(t)}) = \|Z^{(t)}\|_{P^{(t)}}^2$. We may further separate $\|Z^{(t)}\|_{P^{(t)}}^2 = \sum_{j=1}^n P^{(t)}(j) \cdot (Z^{(t)}(j))^2$ into its individual components $P^{(t)}(j) \cdot (Z^{(t)}(j))^2$, and treat each one separately, thus allowing us to reduce to a one-dimensional problem. Finally, we make one further reduction, which is to replace the finite differences $D_h(\cdot)$ in Lemma 4.7 with *circular finite differences*, defined below:

Definition 4.2 (Circular finite difference). Suppose $L = (L^{(0)}, \dots, L^{(S-1)})$ is a sequence of vectors $L^{(t)} \in \mathbb{R}^n$. For integers $h \geq 0$, the *level- h circular finite difference sequence* for the sequence L , denoted by $D_h^\circ L$, is the sequence defined recursively as: $(D_0^\circ L)^{(t)} = L^{(t)}$ for all $0 \leq t < S$, and

$$(D_h^\circ L)^{(t)} = \begin{cases} (D_{h-1}^\circ L)^{(t+1)} - (D_{h-1}^\circ L)^{(t)} & : 0 \leq t \leq S-2 \\ (D_{h-1}^\circ L)^{(1)} - (D_{h-1}^\circ L)^{(T)} & : t = S-1. \end{cases} \quad (13)$$

Circular finite differences for a sequence $L^{(0)}, \dots, L^{(S-1)}$ are defined similarly to finite differences (Definition 4.1) except that unlike for finite differences, where $(D_h L)^{(S-h)}, \dots, (D_h L)^{(S-1)}$ are not defined, $(D_h^\circ L)^{(S-h)}, \dots, (D_h^\circ L)^{(S-1)}$ are defined by “wrapping around” back to the beginning of the sequence. The above-described reductions, which are worked out in detail in Section C.2, allow us to reduce proving Lemma 4.7 to proving the following simpler lemma:

Lemma 4.8. Suppose $\mu \in \mathbb{R}, \alpha > 0$, and $W^{(0)}, \dots, W^{(S-1)} \in \mathbb{R}$ is a sequence of reals satisfying

$$\sum_{t=0}^{S-1} \left((D_2^\circ W)^{(t)} \right)^2 \leq \alpha \cdot \sum_{t=0}^{S-1} \left((D_1^\circ W)^{(t)} \right)^2 + \mu. \quad (14)$$

Then $\sum_{t=0}^{S-1} \left((D_1^\circ W)^{(t)} \right)^2 \leq \alpha \cdot \sum_{t=1}^{S-1} (W^{(t)})^2 + \mu/\alpha$.

To prove Lemma 4.8, we apply the discrete Fourier transform to both sides of (14) and use the Cauchy-Schwarz inequality in frequency domain. For a sequence $W^{(0)}, \dots, W^{(S-1)} \in \mathbb{R}$, its (discrete) *Fourier transform* is the sequence $\widehat{W}^{(0)}, \dots, \widehat{W}^{(S-1)}$ defined by $\widehat{W}^{(s)} = \sum_{t=0}^{S-1} W^{(t)} \cdot e^{-\frac{2\pi i s t}{S}}$. Below we prove Lemma 4.8 for the special case $\mu = 0$; we defer the general case to Section C.1.

Proof of Lemma 4.8 for special case $\mu = 0$. We have the following:

$$S \cdot \sum_{t=1}^T \left((D_1^\circ W)^{(t)} \right)^2 = \sum_{s=1}^T \left| \widehat{D_1^\circ W}^{(s)} \right|^2 = \sum_{s=1}^T \left| \widehat{W}^{(s)} (e^{2\pi i s/T} - 1) \right|^2 \leq \sqrt{\sum_{s=1}^T \left| \widehat{W}^{(s)} \right|^2} \sqrt{\sum_{s=1}^T \left| \widehat{W}^{(s)} \right|^2 |e^{2\pi i s/T} - 1|^4},$$

where the first equality uses Parseval’s equality, the second uses Fact C.3 (in the appendix) for $h = 1$, and the inequality uses Cauchy-Schwarz. By Parseval’s inequality and Fact C.3 for $h = 2$, the right-hand side of the above equals $S \cdot \sqrt{\sum_{t=1}^T (W^{(t)})^2} \cdot \sqrt{\sum_{t=1}^T \left((D_2^\circ W)^{(t)} \right)^2}$, which, by assumption, is at most $S \cdot \sqrt{\sum_{t=1}^T (W^{(t)})^2} \cdot \sqrt{\alpha \cdot \sum_{t=1}^T \left((D_1^\circ W)^{(t)} \right)^2}$. Rearranging terms completes the proof. \square

Acknowledgments and Disclosure of Funding

C.D. is Supported by NSF Awards CCF-1901292, DMS-2022448 and DMS-2134108, by a Simons Investigator Award, by the Simons Collaboration on the Theory of Algorithmic Fairness, by a DSTA grant, and by the DOE PhILMs project (No. DE-AC05-76RL01830). N.G. is supported by a Fannie & John Hertz Foundation Fellowship and an NSF Graduate Fellowship.

References

- [AIMM21] Waïss Azizian, Franck Iutzeler, Jérôme Malick, and Panayotis Mertikopoulos. The last-iterate convergence rate of optimistic mirror descent in stochastic variational inequalities. In *Conference on Learning Theory*, pages 1–32, 2021.
- [BCM12] Maria-Florina Balcan, Florin Constantin, and Ruta Mehta. The Weighted Majority Algorithm does not Converge in Nearly Zero-sum Games. In *ICML Workshop on Markets, Mechanisms, and Multi-Agent Models*, 2012.
- [Bla54] David Blackwell. Controlled Random Walks. In *Proceedings of the International Congress of Mathematicians*, volume 3, pages 336–338, 1954.
- [BP18] James P. Bailey and Georgios Piliouras. Multiplicative Weights Update in Zero-Sum Games. In *Proceedings of the 2018 ACM Conference on Economics and Computation - EC '18*, pages 321–338, Ithaca, NY, USA, 2018. ACM Press.
- [BP19] James P. Bailey and Georgios Piliouras. Fast and furious learning in zero-sum games: Vanishing regret with non-vanishing step sizes. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 12977–12987, 2019.
- [Bro49] George W Brown. Some Notes on Computation of Games Solutions. Technical report, RAND CORP SANTA MONICA CA, 1949.
- [Bub15] Sébastien Bubeck. Convex Optimization: Algorithms and Complexity. *Found. Trends Mach. Learn.*, 8(3–4):231–357, November 2015.
- [CBL06] Nicolo Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge university press, 2006.
- [CDT09] Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player nash equilibria. *Journal of the ACM (JACM)*, 56(3):1–57, 2009.
- [CP19] Yun Kuen Cheung and Georgios Piliouras. Vortices instead of equilibria in minmax optimization: Chaos and butterfly effects of online learning in zero-sum games. In *Proceedings of the Thirty-Second Conference on Learning Theory*, pages 807–834, 2019.
- [CP20] Xi Chen and Binghui Peng. Hedging in games: Faster convergence of external and swap regrets. In *Advances in Neural Information Processing Systems*, volume 33, pages 18990–18999. Curran Associates, Inc., 2020.
- [CS04] Imre Csiszár and Paul C. Shields. Information Theory and Statistics: A Tutorial. *Commun. Inf. Theory*, 1(4):417–528, December 2004.
- [DDK11] Constantinos Daskalakis, Alan Deckelbaum, and Anthony Kim. Near-Optimal No-Regret Algorithms for Zero-Sum Games. In *Proceedings of the twenty-second annual ACM-SIAM symposium on Discrete Algorithms (SODA)*, 2011.
- [DFP⁺10] Constantinos Daskalakis, Rafael Frongillo, Christos H. Papadimitriou, George Pierrakos, and Gregory Valiant. On learning algorithms for nash equilibria. In *Proceedings of the Third International Conference on Algorithmic Game Theory, SAGT'10*, page 114–125, Berlin, Heidelberg, 2010. Springer-Verlag.
- [DGP06] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a nash equilibrium. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing (STOC)*, 2006.
- [DP14] Constantinos Daskalakis and Qinxuan Pan. A counter-example to Karlin’s strong conjecture for fictitious play. In *Proceedings of the 55th Annual Symposium on Foundations of Computer Science (FOCS)*, 2014.
- [DP19] Constantinos Daskalakis and Ioannis Panageas. Last-iterate convergence: Zero-sum games and constrained min-max optimization. In *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, volume 124 of *LIPIcs*, pages 27:1–27:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

- [FLL⁺16] Dylan J Foster, Zhiyuan Li, Thodoris Lykouris, Karthik Sridharan, and Eva Tardos. Learning in games: Robustness of fast convergence. In *Advances in Neural Information Processing Systems*, volume 29. Curran Associates, Inc., 2016.
- [GKP89] Ronald L. Graham, Donald E. Knuth, and Oren Patashnik. *Concrete Mathematics: A Foundation for Computer Science*. Addison-Wesley, Reading, 1989.
- [GPD20] Noah Golowich, Sarath Pattathil, and Constantinos Daskalakis. Tight last-iterate convergence rates for no-regret learning in multi-player games. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*, 2020.
- [HAM21] Yu-Guan Hsieh, Kimon Antonakopoulos, and Panayotis Mertikopoulos. Adaptive learning in continuous games: Optimal regret bounds and convergence to nash equilibrium. In *Conference on Learning Theory*, 2021.
- [Han57] James Hannan. Approximation to Bayes risk in repeated play. *Contributions to the Theory of Games*, 3:97–139, 1957.
- [HMcW⁺03] Hart, Andreu Mas-colell, Of Jörgen W. Weibull, O Vega, Drew Fudenberg, David K. Levine, Josef Hofbauer, Karl Sigmund, Eric Maskin, Motty Perry, and Er Vasin. Uncoupled dynamics do not lead to nash equilibrium. *Amer. Econ. Rev.*, pages 1830–1836, 2003.
- [KHSC18] Ehsan Asadi Kangarshahi, Ya-Ping Hsieh, Mehmet Fatih Sahin, and Volkan Cevher. Let’s be honest: An optimal no-regret framework for zero-sum games. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2488–2496. PMLR, 10–15 Jul 2018.
- [KLP11] Robert Kleinberg, Katrina Ligett, and Georgios Piliouras. Beyond the nash equilibrium barrier. In *Innovations in Computer Science (ICS)*, pages 125–140, 2011.
- [LGNPw21] Qi Lei, Sai Ganesh Nagarajan, Ioannis Panageas, and xiao wang. Last iterate convergence in no-regret learning: constrained min-max optimization for convex-concave landscapes. In Arindam Banerjee and Kenji Fukumizu, editors, *Proceedings of The 24th International Conference on Artificial Intelligence and Statistics*, volume 130 of *Proceedings of Machine Learning Research*, pages 1441–1449. PMLR, 13–15 Apr 2021.
- [MPP18] Panayotis Mertikopoulos, Christos Papadimitriou, and Georgios Piliouras. Cycles in adversarial regularized learning. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms*, SODA ’18, page 2703–2717, USA, 2018. Society for Industrial and Applied Mathematics.
- [Nes05] Yu Nesterov. Excessive gap technique in nonsmooth convex minimization. *SIAM Journal on Optimization*, 16(1):235–249, 2005.
- [PP16] Christos Papadimitriou and Georgios Piliouras. From nash equilibria to chain recurrent sets: Solution concepts and topology. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, ITCS ’16, page 227–235, New York, NY, USA, 2016. Association for Computing Machinery.
- [Rob51] Julia Robinson. An Iterative Method of Solving a Game. *Annals of mathematics*, pages 296–301, 1951.
- [Rou09] Tim Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the Forty-First Annual ACM Symposium on Theory of Computing*, STOC ’09, page 513–522, New York, NY, USA, 2009. Association for Computing Machinery.
- [RS13a] Alexander Rakhlin and Karthik Sridharan. Online learning with predictable sequences. In *Proceedings of the 26th Annual Conference on Learning Theory*, pages 993–1019, 2013.
- [RS13b] Alexander Rakhlin and Karthik Sridharan. Optimization, Learning, and Games with Predictable Sequences. *arXiv:1311.1869 [cs]*, November 2013. arXiv: 1311.1869.
- [RST17] Tim Roughgarden, Vasilis Syrgkanis, and Éva Tardos. The price of anarchy in auctions. *J. Artif. Int. Res.*, 59(1):59–101, May 2017.
- [SALS15] Vasilis Syrgkanis, Alekh Agarwal, Haipeng Luo, and Robert E Schapire. Fast convergence of regularized learning in games. In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015.
- [Sha64] L. Shapley. Some Topics in Two-Person Games. *Advances in Game theory*, 1964.
- [ST13] Vasilis Syrgkanis and Eva Tardos. Composable and efficient mechanisms. In *Proceedings of the Forty-Fifth Annual ACM Symposium on Theory of Computing*, STOC ’13, page 211–220, New York, NY, USA, 2013. Association for Computing Machinery.

- [VGFL⁺20] Emmanouil-Vasileios Vlatakis-Gkaragkounis, Lampros Flokas, Thanasis Lianeas, Panayotis Mertikopoulos, and Georgios Piliouras. No-regret learning and mixed nash equilibria: They do not mix. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 1380–1391. Curran Associates, Inc., 2020.
- [WL18] Chen-Yu Wei and Haipeng Luo. More adaptive algorithms for adversarial bandits. In *Proceedings of the 31st Conference On Learning Theory*, pages 1263–1291, 2018.
- [WLZL21] Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, and Haipeng Luo. Linear last-iterate convergence in constrained saddle-point optimization. In *International Conference on Learning Representations*, 2021.

Checklist

1. For all authors...
 - (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
 - (b) Did you describe the limitations of your work? [Yes]
 - (c) Did you discuss any potential negative societal impacts of your work? [N/A]
 - (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
 - (a) Did you state the full set of assumptions of all theoretical results? [Yes] See Sections 2 and 3.
 - (b) Did you include complete proofs of all theoretical results? [Yes] See Section 4 and the appendix (in the supplementary material).
3. If you ran experiments...
 - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [N/A]
 - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [N/A]
 - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [N/A]
 - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [N/A]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
 - (a) If your work uses existing assets, did you cite the creators? [N/A]
 - (b) Did you mention the license of the assets? [N/A]
 - (c) Did you include any new assets either in the supplemental material or as a URL? [N/A]
 - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
 - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
 - (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
 - (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
 - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

A Proofs for Section 4.1

In this section we prove Lemma 4.1. Throughout the section we use the notation of Lemma 4.1: in particular, we assume that any player $i \in [m]$ follows the Optimistic Hedge updates (1) with step size $\eta > 0$, for an arbitrary sequence of losses $\ell_i^{(1)}, \dots, \ell_i^{(T)}$.

A.1 Preliminary lemmas

The first few lemmas in this section pertain to vectors $P, Q \in \Delta^n$, for some $n \in \mathbb{N}$; note that such vectors P, Q may be viewed as distributions on $[n]$. Let $P/Q \in \mathbb{R}^n$ denote the Radon-Nikodym derivative, i.e., the vector whose j th component is $P(j)/Q(j)$.

Lemma A.1. *If $\|P/Q\|_\infty \leq A$, then $\chi^2(P; Q) \leq A \cdot \chi^2(Q; P)$.*

Proof. The lemma is immediate from the definition of the χ^2 divergence:

$$\chi^2(P; Q) = \sum_{j=1}^n \frac{(P(j) - Q(j))^2}{Q(j)} \leq A \cdot \sum_{j=1}^n \frac{(P(j) - Q(j))^2}{P(j)} = A \cdot \chi^2(Q; P).$$

□

It is a standard fact (though one which we do not need in our proofs) that for all $P, Q \in \Delta^n$, $\text{KL}(P; Q) \leq \chi^2(P; Q)$. The below lemma shows an inequality in the opposite direction when $\|P/Q\|_\infty, \|Q/P\|_\infty$ are bounded:

Lemma A.2. *There is a constant C so that the following holds. Suppose that for $A \leq \frac{3}{2}$ we have $\|P/Q\|_\infty \leq A$ and $\|Q/P\|_\infty \leq A$. Then $(1/2 - C(A - 1)) \cdot \chi^2(P; Q) \leq \text{KL}(P; Q)$.*

Proof. There is a constant $C > 0$ so that for any $0 < \beta \leq 1/2$, for all $|x| \leq \beta$, we have

$$\log(1 + x) \geq x - (1/2 + C\beta)x^2.$$

Set $a = A - 1$, so that $|P(j)/Q(j) - 1| \leq a$ for all j by assumption. Then for $C' = C + 1/2$, we have

$$\begin{aligned} \text{KL}(P; Q) &= \sum_j P(j) \log \frac{P(j)}{Q(j)} \\ &\geq \sum_j P(j) \cdot \left(\left(\frac{P(j)}{Q(j)} - 1 \right) - (1/2 + Ca) \left(\frac{P(j)}{Q(j)} - 1 \right)^2 \right) \\ &\geq \chi^2(P; Q) - (1/2 + Ca) \sum_j P(j) \cdot \frac{(P(j) - Q(j))^2}{Q(j)^2} \\ &\geq \chi^2(P; Q) - \frac{1/2 + Ca}{A} \chi^2(P; Q) \\ &\geq \frac{1/2 - aC}{1 + a} \cdot \chi^2(P; Q) \\ &\geq (1/2 - a \cdot (C + 1/2)) \cdot \chi^2(P; Q) \\ &= (1/2 - C'a) \cdot \chi^2(P; Q). \end{aligned}$$

□

The next lemma considers two vectors $x, x' \in \Delta^n$ which are related by a multiplicative weights-style update with loss vector $w \in \mathbb{R}^n$; the lemma relates $\chi^2(x'; x)$ to $\|w\|_x^2$.

Lemma A.3. *There is a constant $C > 0$ so that the following holds. Suppose that $w \in \mathbb{R}^n$, $\alpha > 0$, $\|w\|_\infty \leq \alpha/2 \leq 1/C$, and $x, x' \in \Delta^n$ satisfy, for each $j \in [n]$,*

$$x'(j) = \frac{x(j) \cdot \exp(w(j))}{\sum_{k \in [n]} x(k) \cdot \exp(w(k))}. \quad (15)$$

Then

$$(1 - C\alpha) \cdot \text{Var}_x(w) \leq \chi^2(x'; x) \leq (1 + C\alpha) \text{Var}_x(w).$$

Proof. Let $w' = w - \langle x, w \rangle \mathbf{1}$, where $\mathbf{1}$ denotes the all-1s vector. Note that $\text{Var}_x(w) = \text{Var}_x(w')$, and that if we replace w with w' , (15) remains true. Moreover, $\|w'\|_\infty \leq 2\|w\|_\infty \leq \alpha$. Thus, by replacing w with w' , we may assume from here on that $\langle w, x \rangle = 0$ and that $\|w\| \leq \alpha$.

Note that

$$\chi^2(x'; x) = -1 + \sum_{i=1}^n x(i) \cdot (x'(i)/x(i))^2 = -1 + \mathbb{E} \left(\frac{\exp(W)}{\mathbb{E} \exp(W)} \right)^2,$$

where W is a random variable that takes values $w(j)$ with probability $x(j)$. As long as C is a sufficiently large constant, we have that, for all z satisfying $|z| \leq \alpha$,

$$1 + z + (1 - C\alpha)z^2/2 \leq \exp(z) \leq 1 + z + (1 + C\alpha)z^2/2. \quad (16)$$

Thus, for a sufficiently large constant C'_0 , we have, for all z satisfying $|z| \leq \alpha$,

$$1 + 2z + (2 - C'_0\alpha)z^2 \leq \exp(z)^2 \leq 1 + 2z + (2 + C'_0\alpha)z^2. \quad (17)$$

Moreover, since $\mathbb{E}W = 0$, we have from (16) that $1 + (1 - C\alpha)\mathbb{E}W^2/2 \leq \mathbb{E} \exp(W) \leq 1 + (1 + C\alpha)\mathbb{E}W^2/2$. For a sufficiently large constant C'_1 it follows that

$$1 + (1 - C'_1\alpha)\mathbb{E}W^2 \leq (\mathbb{E} \exp(W))^2 \leq 1 + (1 + C'_1\alpha)\mathbb{E}W^2. \quad (18)$$

Combining (17) and (18) and again using the fact that $\mathbb{E}W = 0$, we get, for some sufficiently large constant C'' , as long as $\alpha < 1/C'_1$,

$$\begin{aligned} (1 - C''\alpha)\mathbb{E}W^2 &\leq -1 + \frac{1 + (2 - C'_0\alpha)\mathbb{E}W^2}{1 + (1 + C'_1\alpha)\mathbb{E}W^2} \\ &\leq -1 + \frac{\mathbb{E}(\exp(W)^2)}{(\mathbb{E} \exp(W))^2} \\ &\leq -1 + \frac{1 + (2 + C'_0\alpha)\mathbb{E}W^2}{1 + (1 - C'_1\alpha)\mathbb{E}W^2} \\ &\leq (1 + C''\alpha)\mathbb{E}W^2. \end{aligned}$$

By the assumption that $\langle w, x \rangle = 0$, we have $\mathbb{E}W^2 = \text{Var}_x(w)$, and thus the above gives the desired result. \square

We will need the following standard lemma:

Lemma A.4 ([RS13a], Eq. (26)). *For any $n \in \mathbb{N}$, $\ell \in \mathbb{R}^n$, $y \in \Delta^n$, if it holds that $x = \arg \min_{x' \in \Delta^n} \langle x', \ell \rangle + \text{KL}(x'; y)$, then for any $z \in \Delta^n$,*

$$\langle x - z, \ell \rangle \leq \text{KL}(z; y) - \text{KL}(z; x) - \text{KL}(x; y).$$

For $t \in [T]$, we define the vector $\tilde{x}_i^{(t)} \in \Delta^{n_i}$ by

$$\tilde{x}_i^{(t)}(j) := \frac{x_i^{(t)}(j) \cdot \exp(-\eta \cdot (\ell_i^{(t)}(j) - \ell_i^{(t-1)}(j)))}{\sum_{k \in [n_i]} x_i^{(t)}(k) \cdot \exp(-\eta \cdot (\ell_i^{(t)}(k) - \ell_i^{(t-1)}(k)))}. \quad (19)$$

Additionally define $\tilde{x}_i^{(0)} := (1/n_i, \dots, 1/n_i)$ to be the uniform distribution over $[n_i]$.

The next lemma, Lemma A.5 is very similar to [RS13a, Lemma 3], and is indeed essentially shown in the course of the proof of that lemma. Note that no boundedness assumption is placed on the vectors $\ell_i^{(t)}$ in Lemma A.5. For completeness we provide a full proof of the lemma.

Lemma A.5 (Refinement of Lemma 3, [RS13a]). *Suppose that any player $i \in [m]$ follows the Optimistic Hedge updates (1) with step size $\eta > 0$, for an arbitrary sequence of losses $\ell_i^{(1)}, \dots, \ell_i^{(T)} \in \mathbb{R}^{n_i}$. For any vector $x^* \in \Delta^{n_i}$, it holds that*

$$\sum_{t=1}^T \langle x_i^{(t)} - x^*, \ell_i^{(t)} \rangle \leq \frac{\log n_i}{\eta} + \sum_{t=1}^T \left\| x_i^{(t)} - \tilde{x}_i^{(t)} \right\|_{x_i^{(t)}}^* \sqrt{\text{Var}_{x_i^{(t)}} \left(\ell_i^{(t)} - \ell_i^{(t-1)} \right)} - \frac{1}{\eta} \sum_{t=1}^T \text{KL}(\tilde{x}_i^{(t)}; x_i^{(t)}) - \frac{1}{\eta} \sum_{t=1}^T \text{KL}(x_i^{(t)}; \tilde{x}_i^{(t-1)}). \quad (20)$$

Proof. For any $x^* \in \Delta^{n_i}$, it holds that

$$\langle x_i^{(t)} - x^*, \ell_i^{(t)} \rangle = \langle x_i^{(t)} - \tilde{x}_i^{(t)}, \ell_i^{(t)} - \ell_i^{(t-1)} \rangle + \langle x_i^{(t)} - \tilde{x}_i^{(t)}, \ell_i^{(t-1)} \rangle + \langle \tilde{x}_i^{(t)} - x^*, \ell_i^{(t)} \rangle. \quad (21)$$

For $t \in [T]$, set $c^{(t)} = \langle x_i^{(t)}, \ell_i^{(t)} - \ell_i^{(t-1)} \rangle$. Using the definition of the dual norm and the fact $\langle x_i^{(t)} - \tilde{x}_i^{(t)}, \mathbf{1} \rangle = 0$, we have

$$\begin{aligned} \langle x_i^{(t)} - \tilde{x}_i^{(t)}, \ell_i^{(t)} - \ell_i^{(t-1)} \rangle &= \langle x_i^{(t)} - \tilde{x}_i^{(t)}, \ell_i^{(t)} - \ell_i^{(t-1)} - c^{(t)} \mathbf{1} \rangle \\ &\leq \left\| x_i^{(t)} - \tilde{x}_i^{(t)} \right\|_{x_i^{(t)}}^* \cdot \left\| \ell_i^{(t)} - \ell_i^{(t-1)} - c^{(t)} \mathbf{1} \right\|_{x_i^{(t)}} \\ &\leq \left\| x_i^{(t)} - \tilde{x}_i^{(t)} \right\|_{x_i^{(t)}}^* \cdot \sqrt{\text{Var}_{x_i^{(t)}} \left(\ell_i^{(t)} - \ell_i^{(t-1)} \right)}. \end{aligned} \quad (22)$$

It is immediate from the definitions of $\tilde{x}_i^{(t)}$ (in (19)) and $x_i^{(t)}$ (in (1)) that for $j \in [n_i]$,

$$x_i^{(t)}(j) = \frac{\tilde{x}_i^{(t-1)}(j) \cdot \exp(-\eta \cdot \ell_i^{(t-1)}(j))}{\sum_{k \in [n_i]} \tilde{x}_i^{(t-1)}(k) \cdot \exp(-\eta \cdot \ell_i^{(t-1)}(k))} = \arg \min_{x \in \Delta^{n_i}} \langle x, \eta \cdot \ell_i^{(t-1)} \rangle + \text{KL}(x; \tilde{x}_i^{(t-1)}) \quad (23)$$

Using Lemma A.4 with $x = x_i^{(t)}$, $\ell = \eta \ell_i^{(t-1)}$, $y = \tilde{x}_i^{(t-1)}$, $z = \tilde{x}_i^{(t)}$, we obtain

$$\langle x_i^{(t)} - \tilde{x}_i^{(t)}, \ell_i^{(t-1)} \rangle \leq \frac{1}{\eta} \text{KL}(\tilde{x}_i^{(t)}; \tilde{x}_i^{(t-1)}) - \frac{1}{\eta} \text{KL}(\tilde{x}_i^{(t)}; x_i^{(t)}) - \frac{1}{\eta} \text{KL}(x_i^{(t)}; \tilde{x}_i^{(t-1)}). \quad (24)$$

Next, we note that, again by (19) and (1), for $j \in [n_i]$,

$$\tilde{x}_i^{(t)}(j) = \frac{\tilde{x}_i^{(t-1)}(j) \cdot \exp(-\eta \cdot \ell_i^{(t)}(j))}{\sum_{k \in [n_i]} \tilde{x}_i^{(t-1)}(k) \cdot \exp(-\eta \cdot \ell_i^{(t)}(k))} = \arg \min_{x \in \Delta^{n_i}} \langle x, \eta \cdot \ell_i^{(t)} \rangle + \text{KL}(x; \tilde{x}_i^{(t-1)}).$$

Using Lemma A.4 with $x = \tilde{x}_i^{(t)}$, $\ell = \eta \ell_i^{(t)}$, $y = \tilde{x}_i^{(t-1)}$, $z = x^*$, we obtain

$$\langle \tilde{x}_i^{(t)} - x^*, \ell_i^{(t)} \rangle \leq \frac{1}{\eta} \text{KL}(x^*; \tilde{x}_i^{(t-1)}) - \frac{1}{\eta} \text{KL}(x^*; \tilde{x}_i^{(t)}) - \frac{1}{\eta} \text{KL}(\tilde{x}_i^{(t)}; \tilde{x}_i^{(t-1)}). \quad (25)$$

By (21), (22), (24), and (25), we have

$$\begin{aligned} \langle x_i^{(t)} - x^*, \ell_i^{(t)} \rangle &\leq \left\| x_i^{(t)} - \tilde{x}_i^{(t)} \right\|_{x_i^{(t)}}^* \cdot \sqrt{\text{Var}_{x_i^{(t)}} \left(\ell_i^{(t)} - \ell_i^{(t-1)} \right)} \\ &\quad + \frac{1}{\eta} \text{KL}(\tilde{x}_i^{(t)}; \tilde{x}_i^{(t-1)}) - \frac{1}{\eta} \text{KL}(\tilde{x}_i^{(t)}; x_i^{(t)}) - \frac{1}{\eta} \text{KL}(x_i^{(t)}; \tilde{x}_i^{(t-1)}) \\ &\quad + \frac{1}{\eta} \text{KL}(x^*; \tilde{x}_i^{(t-1)}) - \frac{1}{\eta} \text{KL}(x^*; \tilde{x}_i^{(t)}) - \frac{1}{\eta} \text{KL}(\tilde{x}_i^{(t)}; \tilde{x}_i^{(t-1)}) \\ &= \left\| x_i^{(t)} - \tilde{x}_i^{(t)} \right\|_{x_i^{(t)}}^* \cdot \sqrt{\text{Var}_{x_i^{(t)}} \left(\ell_i^{(t)} - \ell_i^{(t-1)} \right)} + \frac{1}{\eta} \text{KL}(x^*; \tilde{x}_i^{(t-1)}) - \frac{1}{\eta} \text{KL}(x^*; \tilde{x}_i^{(t)}) \\ &\quad - \frac{1}{\eta} \text{KL}(\tilde{x}_i^{(t)}; x_i^{(t)}) - \frac{1}{\eta} \text{KL}(x_i^{(t)}; \tilde{x}_i^{(t-1)}). \end{aligned} \quad (26)$$

The statement of the lemma follows by summing (26) over $t \in [T]$ and using the fact that for any choice of x^* , $\text{KL}(x^*; x_i^{(0)}) \leq \log n_i$. \square

A.2 Proof of Lemma 4.1

Now we are ready to prove Lemma 4.1. For convenience we restate the lemma.

Lemma 4.1 (restated). *There is a constant $C > 0$ so that the following holds. Suppose any player $i \in [m]$ follows the Optimistic Hedge updates (1) with step size $0 < \eta < 1/C$, for an arbitrary sequence of losses $\ell_i^{(1)}, \dots, \ell_i^{(T)} \in [0, 1]^{n_i}$. Then for any vector $x^* \in \Delta^{n_i}$, it holds that*

$$\sum_{t=1}^T \langle x_i^{(t)} - x^*, \ell_i^{(t)} \rangle \leq \frac{\log n_i}{\eta} + \sum_{t=1}^T \left(\frac{\eta}{2} + C\eta^2 \right) \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t)} - \ell_i^{(t-1)} \right) - \sum_{t=1}^T \frac{(1 - C\eta)\eta}{2} \cdot \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t-1)} \right). \quad (27)$$

Proof. Lemma A.5 gives that, for any $x^* \in \Delta^{n_i}$,

$$\sum_{t=1}^T \langle x_i^{(t)} - x^*, \ell_i^{(t)} \rangle \leq \frac{\log n_i}{\eta} + \sum_{t=1}^T \left\| x_i^{(t)} - \tilde{x}_i^{(t)} \right\|_{x_i^{(t)}}^* \sqrt{\text{Var}_{x_i^{(t)}} \left(\ell_i^{(t)} - \ell_i^{(t-1)} \right)} - \frac{1}{\eta} \sum_{t=1}^T \text{KL}(\tilde{x}_i^{(t)}; x_i^{(t)}) - \frac{1}{\eta} \sum_{t=1}^T \text{KL}(x_i^{(t)}; \tilde{x}_i^{(t-1)}). \quad (28)$$

Note that for any vectors $x, x' \in \Delta^{n_i}$, if there is a vector $\ell \in \mathbb{R}^{n_i}$ so that for all $j \in [n_i]$, $x'(j) = \frac{x(j) \cdot \exp(\eta \cdot \ell(j))}{\sum_k x(j) \cdot \exp(\eta \cdot \ell(k))}$, we have that

$$\exp(-2\eta \|\ell\|_\infty) \leq \left\| \frac{x'}{x} \right\|_\infty \leq \exp(2\eta \|\ell\|_\infty).$$

Therefore, by (19) and (23), respectively, we obtain that, for $\eta \leq 1/4$,

$$\begin{aligned} \exp(-2\eta \|\ell_i^{(t)} - \ell_i^{(t-1)}\|_\infty) &\leq \left\| \frac{\tilde{x}_i^{(t)}}{x_i^{(t)}} \right\|_\infty \leq \exp(2\eta \|\ell_i^{(t)} - \ell_i^{(t-1)}\|_\infty) \leq \exp(4\eta) \leq 1 + 8\eta \\ \exp(-2\eta \|\ell_i^{(t-1)}\|_\infty) &\leq \left\| \frac{x_i^{(t)}}{\tilde{x}_i^{(t-1)}} \right\|_\infty \leq \exp(2\eta \|\ell_i^{(t-1)}\|_\infty) \leq \exp(2\eta) \leq 1 + 4\eta. \end{aligned} \quad (29)$$

(Above we have also used that $\|\ell_i^{(t)}\|_\infty \leq 1$ for all t .) Thus, for $\eta \leq \frac{1}{16}$, we can apply Lemma A.2 and show, for a sufficiently large constant C_0 ,

$$\text{KL}(\tilde{x}_i^{(t)}; x_i^{(t)}) \geq \chi^2(\tilde{x}_i^{(t)}; x_i^{(t)}) \cdot (1/2 - C_0\eta) \quad (30)$$

$$\text{KL}(x_i^{(t)}; \tilde{x}_i^{(t-1)}) \geq \chi^2(x_i^{(t)}; \tilde{x}_i^{(t-1)}) \cdot (1/2 - C_0\eta). \quad (31)$$

Note also that for vectors x, y we have that $\chi^2(x; y) = \left(\|x - y\|_y^* \right)^2$. By Lemma A.3 and (19), we have that, for a sufficiently large constant C_1 , as long as $\eta \leq 1/C_1$,

$$\left(\left\| x_i^{(t)} - \tilde{x}_i^{(t)} \right\|_{x_i^{(t)}}^* \right)^2 = \chi^2(\tilde{x}_i^{(t)}; x_i^{(t)}) \leq (1 + C_1\eta)\eta^2 \cdot \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t)} - \ell_i^{(t-1)} \right) \quad (32)$$

and

$$\chi^2(\tilde{x}_i^{(t)}; x_i^{(t)}) \geq (1 - C_1\eta)\eta^2 \cdot \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t)} - \ell_i^{(t-1)} \right). \quad (33)$$

Next we lower bound $\chi^2(x_i^{(t)}; \tilde{x}_i^{(t-1)})$ as follows, where C_2 denotes a sufficiently large constant: as long as $\eta \leq 1/C_2$,

$$\chi^2(x_i^{(t)}; \tilde{x}_i^{(t-1)}) \geq \chi^2(\tilde{x}_i^{(t-1)}; x_i^{(t)}) \cdot \exp(-2\eta) \quad (34)$$

$$\geq (1 - C_2\eta)\eta^2 \cdot \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t-1)} \right), \quad (35)$$

where (34) follows from Lemma A.1 and (29), and (35) follows from Lemma A.3 and (23).

Combining (28), (30), (31), (32), (33), and (35) gives that for a sufficiently large constant C , as long as $\eta < 1/C$,

$$\sum_{t=1}^T \langle x_i^{(t)} - x^*, \ell_i^{(t)} \rangle \leq \frac{\log n_i}{\eta} + \sum_{t=1}^T (\eta/2 + C\eta^2) \cdot \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t)} - \ell_i^{(t-1)} \right) - \frac{(1 - C\eta)\eta}{2} \cdot \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t-1)} \right),$$

as desired. \square

B Proofs for Section 4.3

In this section we give the full proof of Lemma 4.4. In Section B.1 we introduce some preliminaries. In Section B.2 we prove Lemma 4.5, the ‘‘boundedness chain rule’’ for finite differences. In Section B.4 we show how to use this lemma to prove Lemma 4.4.

B.1 Additional preliminaries

In this section we introduce some additional notations and basic combinatorial lemmas. Definition B.1 introduces the *shift operator* E_s , which like the finite difference operator D_h , maps one sequence to another sequence.

Definition B.1 (Shift operator). Suppose $L = (L^{(1)}, \dots, L^{(T)})$ is a sequence of vectors $L^{(t)} \in \mathbb{R}^n$. For integers $s \geq 0$, the *s-shift sequence* for the sequence L , denoted by $E_s L$, is the sequence $E_s L = ((E_s L)^{(1)}, \dots, (E_s L)^{(T-s)})$, defined by $(E_s L)^{(t)} = L^{(t+s)}$ for $1 \leq t \leq T - s$.

For sequences $L = (L^{(1)}, \dots, L^{(T)})$ and $K = (K^{(1)}, \dots, K^{(T)})$ of real numbers, we will denote the *product sequence* as $L \cdot K$ as the sequence of vectors $L \cdot K := (L^{(1)}K^{(1)}, \dots, L^{(T)}K^{(T)})$. Lemmas B.1 and B.2 below are standard analogues of the product rule for finite differences. The (straightforward) proofs are provided for completeness.

Lemma B.1 (Product rule; Eq. (2.55) of [GKP89]). Suppose $L = (L^{(1)}, \dots, L^{(T)})$ and $K = (K^{(1)}, \dots, K^{(T)})$ are sequences of real numbers. Then the product sequence $L \cdot K$ satisfies

$$D_1(L \cdot K) = L \cdot D_1 K + D_1 L \cdot E_1 K.$$

Proof. We compute

$$\begin{aligned} D_1(L \cdot K)^{(t)} &= L^{(t+1)}K^{(t+1)} - L^{(t)}K^{(t)} \\ &= L^{(t+1)}K^{(t+1)} - L^{(t)}K^{(t+1)} + L^{(t)}K^{(t+1)} - L^{(t)}K^{(t)} \\ &= (L \cdot D_1 K + D_1 L \cdot E_1 K)^{(t)}. \end{aligned}$$

□

Lemma B.2 (Multivariate product rule). Suppose that $m \in \mathbb{N}$ and for $1 \leq i \leq m$, $L_i = (L_i^{(1)}, \dots, L_i^{(T)})$ are sequences of real numbers. Then the product sequence $\prod_{i=1}^m L_i$ satisfies

$$D_1 \prod_{i=1}^m L_i = \sum_{i=1}^m \left(\prod_{i' < i} L_{i'} \right) \cdot D_1 L_i \cdot \left(\prod_{i' > i} E_1 L_{i'} \right).$$

Proof. We compute

$$\begin{aligned} \left(D_1 \prod_{i=1}^m L_i \right)^{(t)} &= \prod_{i=1}^m L_i^{(t+1)} - \prod_{i=1}^m L_i^{(t)} \\ &= \sum_{i=1}^m \left(\prod_{i' < i} L_{i'}^{(t+1)} \prod_{i' > i} L_{i'}^{(t)} - \prod_{i' < i} L_{i'}^{(t+1)} \prod_{i' \geq i} L_{i'}^{(t)} \right) \\ &= \sum_{i=1}^m \left(\prod_{i' < i} L_{i'}^{(t+1)} \cdot \prod_{i' > i} L_{i'}^{(t)} \cdot (L_i^{(t+1)} - L_i^{(t)}) \right) \\ &= \left(\sum_{i=1}^m \left(\prod_{i' < i} L_{i'} \right) \cdot D_1 L_i \cdot \left(\prod_{i' > i} E_1 L_{i'} \right) \right)^{(t)}. \end{aligned}$$

□

Lemma B.4 and Lemma B.3, which is used in the proof of the former, are used to bound certain sums with many terms in the proof of Lemma 4.5. To state Lemma B.3 we make one definition. For positive integers k, m and any $h, C > 0$, define

$$R_{h,m,k,C} = \sum_{0 \leq n_1, \dots, n_k \leq m} \left(\frac{\prod_{i=1}^k n_i^{n_i}}{h^{\sum_{i=1}^k n_i}} \right)^C,$$

where the sum is over integers n_1, \dots, n_k satisfying $0 \leq n_i \leq m$ for $i \in [k]$. In the definition of $R_{h,m,k,C}$, the quantity 0^0 (which arises when some $n_i = 0$) is interpreted as 1.

Lemma B.3. For any positive integers k, m and any $h, C > 0$ so that $m \leq h/2$, $C \geq 2$, and $h \geq 8$, then

$$R_{h,m,k,C} \leq \exp\left(\frac{2k}{h^C}\right).$$

Proof of Lemma B.3. We may rewrite $R_{h,m,k,C}$ and then upper bound it as follows:

$$\begin{aligned} R_{h,m,k,C} &= \left(\sum_{j=0}^m \left(\frac{j}{h}\right)^{Cj} \right)^k \\ &\leq \left(1 + \left(\frac{1}{h}\right)^C + (m-1) \max\left(\left(\frac{2}{h}\right)^{2C}, \left(\frac{m}{h}\right)^{mC}\right) \right)^k \\ &\leq \left(1 + \left(\frac{1}{h}\right)^C + (h/2) \max\left(\left(\frac{2}{h}\right)^{2C}, \left(\frac{1}{2}\right)^{hC/2}\right) \right)^k \end{aligned} \quad (36)$$

where (36) follows since $\left(\frac{i}{h}\right)^{Ci}$ is convex in i for $i \geq 0$, and therefore, in the interval $[2, m] \subseteq [2, h/2]$, takes on maximal values at the endpoints. We see

$$(h/2) \left(\frac{2}{h}\right)^{2C} = \left(\frac{2}{h}\right)^{2C-1} \leq \left(\frac{1}{h}\right)^C$$

for $h \geq 8$ when $C \geq 2$. Also,

$$(h/2) \left(\frac{1}{2}\right)^{hC/2} \leq \left(\frac{1}{h}\right)^C$$

for $h \geq 8$ when $C \geq 2$. (This inequality is easily seen to be equivalent to the fact that $(C+1) \log h - \frac{Ch}{2} \leq 1$, which follows from the fact that $\log h - h/2 \leq 0$ for $h \geq 8$ and $3 \log h - h \leq 1$ for $h \geq 8$.) Therefore,

$$\begin{aligned} R_{h,m,k,C} &\leq \left(1 + \left(\frac{1}{h}\right)^C + (h/2) \max\left(\left(\frac{2}{h}\right)^{2C}, \left(\frac{1}{2}\right)^{hC/2}\right) \right)^k \\ &\leq \left(1 + 2 \left(\frac{1}{h}\right)^C \right)^k \\ &\leq \exp\left(\frac{2k}{h^C}\right). \end{aligned}$$

□

Lemma B.4. Fix integers $h \geq 0, k \geq 1$. For any function $\pi : [h] \rightarrow [k]$, define, for each $i \in [k]$, $h_i(\pi) = |\{q \in [h] \mid \pi(q) = i\}|$. Then, for any $C \geq 3$,

$$\sum_{\pi: [h] \rightarrow [k]} \frac{\prod_{i=1}^k h_i(\pi)^{Ch_i(\pi)}}{h^{Ch}} \leq \max\left\{k^7, (hk+1) \cdot \exp\left(\frac{2k}{h^{C-1}}\right)\right\}. \quad (37)$$

Proof. In the case that $h \leq 7$, we simply use the fact that the number of functions $\pi : [h] \rightarrow [k]$ is $k^h \leq k^7$, and each term of the summation on the left-hand side of (37) is at most 1. In the remainder of the proof we may thus assume that $h \geq 8$.

For any tuple (h_1, \dots, h_k) of non-negative integers with $\sum_{i=1}^k h_i = h$, there are $\binom{h}{h_1, h_2, \dots, h_k} \leq \frac{h^h}{\prod_i h_i^{h_i}}$ (see [CS04, Lemma 2.2] for a proof of this inequality) functions $\pi : [h] \rightarrow [k]$ such that

$h_i(\pi) = h_i$ for all $i \in [k]$. Combining these like terms,

$$\begin{aligned} \sum_{\pi: [h] \rightarrow [k]} \frac{\prod_i h_i(\pi)^{C h_i(\pi)}}{h^{Ch}} &\leq \sum_{\substack{h_1, \dots, h_k \geq 0 \\ \sum h_i = h}} \frac{h^h}{\prod_i h_i^{h_i}} \cdot \left(\frac{\prod_i h_i^{h_i}}{h^h} \right)^C \\ &\leq \sum_{\substack{h_1, \dots, h_k \geq 0 \\ \sum h_i = h}} \left(\frac{\prod_i h_i^{h_i}}{h^h} \right)^{C-1}. \end{aligned} \quad (38)$$

We evaluate this sum in 2 cases: whether or not $h_{\max} := \max_i \{h_i\}$ is greater than $h/2$. The contribution to this sum coming from terms with $h_{\max} \leq h/2$ is

$$\begin{aligned} \sum_{\substack{h_1, \dots, h_k \geq 0 \\ h_1, \dots, h_k \leq h/2 \\ \sum h_i = h}} \left(\frac{\prod_i h_i^{h_i}}{h^h} \right)^{C-1} &\leq \sum_{\substack{h_1, \dots, h_k \geq 0 \\ h_1, \dots, h_k \leq h/2}} \left(\frac{\prod_i h_i^{h_i}}{h \sum h_i} \right)^{C-1} \\ &= R_{h, \lfloor h/2 \rfloor, k, C-1} \\ &\leq \exp\left(\frac{2k}{h^{C-1}}\right), \end{aligned} \quad (39)$$

by Lemma B.3.

We next consider the case where $h_{\max} > h/2$. For a specific term (h_1, \dots, h_k) with $\max_i \{h_i\} > h/2$, we know there is a unique $M \in [k]$ such that $h_M = \max_i \{h_i\}$ since $\sum_{i=1}^k h_i = h$. So, we can represent the contribution to the sum from this case as

$$\begin{aligned} \sum_{M=1}^k \sum_{\substack{h_1, \dots, h_k \geq 0 \\ h_M > h/2 \\ \sum h_i = h}} \left(\frac{\prod_i h_i^{h_i}}{h^h} \right)^{C-1} &= k \sum_{\substack{h_1, \dots, h_k \geq 0 \\ h_k > h/2 \\ \sum h_i = h}} \left(\frac{\prod_i h_i^{h_i}}{h^h} \right)^{C-1} \quad (40) \\ &\leq k \sum_{d=0}^{\lfloor h/2 \rfloor} \left(\frac{(h-d)^{h-d}}{h^{h-d}} \right)^{C-1} \sum_{\substack{h_1, \dots, h_{k-1} \geq 0 \\ \sum h_i = d}} \left(\frac{\prod_i h_i^{h_i}}{h^d} \right)^{C-1} \quad (41) \\ &\leq k \sum_{d=0}^{\lfloor h/2 \rfloor} \sum_{\substack{h_1, \dots, h_{k-1} \geq 0 \\ h_1, \dots, h_{k-1} \leq d}} \left(\frac{\prod_i h_i^{h_i}}{h \sum h_i} \right)^{C-1} \\ &= k \sum_{d=0}^{\lfloor h/2 \rfloor} R_{h, d, k-1, C-1} \\ &\leq kh \cdot \exp\left(\frac{2k}{h^{C-1}}\right), \end{aligned} \quad (42)$$

where (40) follows by symmetry, (41) follows by factoring out the contribution of $\left(\frac{h_k}{h^{h_k}}\right)^C$ and letting $d = h - h_k$, and (42) follows by Lemma B.3.

The statement of the lemma follows from (38), (39), and (42). \square

Lemma B.5. For $n \in \mathbb{N}$, let $\xi_1, \dots, \xi_n \geq 0$ such that $\xi_1 + \dots + \xi_n = 1$. For each $j \in [n]$, define $\phi_j : \mathbb{R}^n \rightarrow \mathbb{R}$ to be the function

$$\phi_j((z_1, \dots, z_n)) = \frac{\xi_j \exp(z_j)}{\sum_{k=1}^n \xi_k \cdot \exp(z_k)}$$

and let $P_{\phi_j}(z) = \sum_{\gamma \in \mathbb{Z}_{\geq 0}^n} a_{j,\gamma} \cdot z^\gamma$ denote the Taylor series of ϕ_j . Then for any $j \in [n]$ and any integer $k \geq 1$,

$$\sum_{\gamma \in \mathbb{Z}_{\geq 0}^n: |\gamma|=k} |a_{j,\gamma}| \leq \xi_j e^{k+1}.$$

Proof. Note that, for each $j \in [n]$,

$$a_{j,\gamma} = \frac{1}{\gamma_1! \gamma_2! \cdots \gamma_n!} \cdot \frac{\partial^k \phi_j(0)}{\partial z_1^{\gamma_1} \partial z_2^{\gamma_2} \cdots \partial z_n^{\gamma_n}},$$

and so

$$\begin{aligned} \sum_{\gamma \in \mathbb{Z}_{\geq 0}^n: |\gamma|=k} |a_{j,\gamma}| &= \sum_{\gamma \in \mathbb{Z}_{\geq 0}^n: |\gamma|=k} \frac{1}{\gamma_1! \gamma_2! \cdots \gamma_n!} \cdot \left| \frac{\partial^k \phi_j(0)}{\partial z_1^{\gamma_1} \partial z_2^{\gamma_2} \cdots \partial z_n^{\gamma_n}} \right| \\ &= \frac{1}{k!} \sum_{\gamma \in \mathbb{Z}_{\geq 0}^n: |\gamma|=k} \frac{k!}{\gamma_1! \gamma_2! \cdots \gamma_n!} \cdot \left| \frac{\partial^k \phi_j(0)}{\partial z_1^{\gamma_1} \partial z_2^{\gamma_2} \cdots \partial z_n^{\gamma_n}} \right| \\ &= \frac{1}{k!} \sum_{t \in [n]^k} \left| \frac{\partial^k \phi_j(0)}{\partial z_{t_1} \partial z_{t_2} \cdots \partial z_{t_k}} \right|. \end{aligned}$$

It is straightforward to see that the following equalities hold for any $i \in [n]$, $i \neq j$:

$$\begin{aligned} \frac{\partial \phi_j}{\partial z_j} &= \phi_j(1 - \phi_j) \\ \frac{\partial \phi_j}{\partial z_i} &= -\phi_i \phi_j \\ \frac{\partial(1 - \phi_j)}{\partial z_j} &= -\phi_j(1 - \phi_j) \\ \frac{\partial(1 - \phi_j)}{\partial z_i} &= \phi_i \phi_j \end{aligned}$$

We claim that for any $(t_1, \dots, t_k) \in [n]^k$, we can express $\frac{\partial^k \phi_j}{\partial z_{t_1} \cdots \partial z_{t_k}}$ as a polynomial in $\phi_1, \dots, \phi_n, (1 - \phi_1), \dots, (1 - \phi_n)$ comprised of $k!$ monomials each of degree $k + 1$. We verify this by induction, first noting that after taking zero derivatives, the function ϕ_j is a degree-1 monomial. Assume that for some sequence $b_1, \dots, b_{(\ell-1)!} \in \{0, 1\}$, we can express

$$\frac{\partial^{\ell-1} \phi_j}{\partial z_{t_1} \cdots \partial z_{t_{\ell-1}}} = \sum_{f=1}^{(\ell-1)!} (-1)^{b_f} \prod_{d=0}^{\ell-1} m_{f,d}$$

where each $m_{f,d} \in \{\phi_1, \dots, \phi_n, (1 - \phi_1), \dots, (1 - \phi_n)\}$. We see that for each f , there is some sequence of bits $b_{f,0}, \dots, b_{f,\ell-1} \in \{0, 1\}$ so that

$$\frac{\partial}{\partial z_{t_\ell}} \prod_{d=0}^{\ell-1} m_{f,d} = \sum_{d=0}^{\ell-1} (-1)^{b_{f,d}} \cdot m_{f,0} \cdots m'_{f,d} \cdots m_{f,d,\ell} \quad (43)$$

where we define, for each $0 \leq d \leq \ell - 1$,

$$m'_{f,d} \text{ and } m_{f,d,\ell} = \begin{cases} m_{f,d} \text{ and } \phi_{t_\ell} & \text{if } m_{f,d} = \phi_i \text{ with } i \neq t_\ell \\ m_{f,d} \text{ and } (1 - \phi_{t_\ell}) & \text{if } m_{f,d} = \phi_{t_\ell} \\ (1 - m_{f,d}) \text{ and } \phi_{t_\ell} & \text{if } m_{f,d} = 1 - \phi_i \text{ with } i \neq t_\ell \\ (1 - m_{f,d}) \text{ and } (1 - \phi_{t_\ell}) & \text{if } m_{f,d} = 1 - \phi_{t_\ell}. \end{cases}$$

Thus, $\frac{\partial^\ell \phi_j}{\partial z_{t_1} \cdots \partial z_{t_\ell}}$ can be expressed as a sum of $\ell!$ monomials of degree $(\ell + 1)$, completing the inductive step.

This inductive argument also demonstrates a bijection between the $k!$ monomials of $\frac{\partial^k \phi_j}{\partial z_{t_1} \cdots \partial z_{t_k}}$ and a combinatorial structure that we call *factorial trees*. Formally, we define a factorial tree to be a directed graph on vertices $\{0, 1, \dots, k\}$ such that each vertex $i \neq 0$ has a single incoming edge from one of the vertices in $[0, i - 1]$. (For a non-negative integer i , we write $[0, i] := \{0, 1, \dots, i\}$.) For a factorial tree f , let $p_f(\ell) \in [0, \ell - 1]$ denote the parent of a vertex ℓ . A particular factorial tree f represents the monomial that was generated by choosing the $p_f(\ell)^{\text{th}}$ term in (43) for derivation when taking the derivative $\frac{\partial}{\partial z_{t_\ell}}$, for each $\ell \in [k]$. (See Figure 1 for an example.)

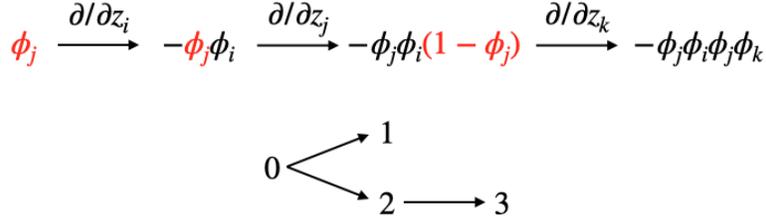


Figure 1: A monomial $-\phi_j \phi_i \phi_j \phi_k$ of $\frac{\partial^3 \phi_j}{\partial z_i \partial z_j \partial z_k}$ and its corresponding factorial tree

Each of the $k!$ monomials comprising $\frac{\partial^k \phi_j}{\partial z_{t_1} \cdots \partial z_{t_k}}$ is a product of $k + 1$ terms corresponding to indices j, t_1, \dots, t_k (i.e., the first term in the product is either ϕ_j or $1 - \phi_j$, the second term is either ϕ_{t_1} or $1 - \phi_{t_1}$, and so on). We say that a term corresponding to index $i \in [n]$ is *perturbed* if it is $(1 - \phi_i)$ (as opposed to ϕ_i). From our construction, we see that the ℓ^{th} term is perturbed if $t_\ell = t_{p_f(\ell)}$ and there is no ℓ' such that $p_f(\ell') = \ell$. That is, ℓ is a leaf in the corresponding factorial tree f and the parent of ℓ corresponds to the same index as ℓ . One can think of t_1, \dots, t_k as a coloring of all the vertices of the factorial tree with n colors, except the root of the tree (vertex 0) which has fixed color j . Then, we can say the ℓ^{th} term is perturbed if and only if ℓ is a leaf with the same color as its parent. We call such a leaf a *petal*. For $t \in [n]$, we let $P_{f,t} \subseteq [k]$ be the set of petals on tree f with color t , $L_f \subseteq [k]$ be the set of leaves of tree f , and $B_f = [k] \setminus L_f$ be the set of all non-leaves other than the fixed-color root. Therefore,

$$\begin{aligned} \sum_{\gamma \in \mathbb{Z}_{\geq 0}^n: |\gamma|=k} |a_{j,\gamma}| &= \frac{1}{k!} \sum_{t \in [n]^k} \left| \frac{\partial^k \phi_j(0)}{\partial z_{t_1} \cdots \partial z_{t_k}} \right| \\ &\leq \frac{1}{k!} \sum_{t \in [n]^k} \sum_f \prod_{\ell=0}^k (\phi_{t_\ell}(0) \cdot \mathbb{1}[\ell \notin P_{f,t}] + (1 - \phi_{t_\ell}(0)) \cdot \mathbb{1}[\ell \in P_{f,t}]) \end{aligned}$$

(where we let $t_0 = j$ for notational convenience)

$$\begin{aligned} &= \frac{1}{k!} \sum_{t \in [n]^k} \sum_f \prod_{\ell=0}^k (\xi_{t_\ell} \cdot \mathbb{1}[\ell \notin P_{f,t}] + (1 - \xi_{t_\ell}) \cdot \mathbb{1}[\ell \in P_{f,t}]) \\ &= \frac{1}{k!} \sum_f \sum_{t_{B_f} \in [n]^{B_f}} \sum_{t_{L_f} \in [n]^{L_f}} \prod_{\ell=0}^k (\xi_{t_\ell} \cdot \mathbb{1}[\ell \notin P_{f,t}] + (1 - \xi_{t_\ell}) \cdot \mathbb{1}[\ell \in P_{f,t}]), \end{aligned}$$

where in the last step we decompose, for each factorial tree f , $t \in [n]^k$ into the tuple of indices $t_{B_f} \in [n]^{B_f}$ corresponding to the non-leaves B_f , and the tuple of indices $t_{L_f} \in [n]^{L_f}$ corresponding to the leaves L_f .

We note that, fixing tree f and the colors of all non-leaves t_B ,

$$\begin{aligned}
& \sum_{t_{L_f} \in [n]^{L_f}} \prod_{\ell \in L_f} (\xi_{t_\ell} \cdot \mathbb{1}[\ell \notin P_{f,t}] + (1 - \xi_{t_\ell}) \cdot \mathbb{1}[\ell \in P_{f,t}]) \\
&= \prod_{\ell \in L_f} \left(\sum_{t_\ell \in [n]} \xi_{t_\ell} \cdot \mathbb{1}[t_\ell \neq t_{p_f(\ell)}] + (1 - \xi_{t_\ell}) \cdot \mathbb{1}[t_\ell = t_{p_f(\ell)}] \right) \\
&= \prod_{\ell \in L_f} (2 - 2\xi_{t_{p_f(\ell)}}) \\
&\leq 2^{|L_f|}
\end{aligned}$$

And so,

$$\begin{aligned}
& \frac{1}{k!} \sum_f \sum_{t_{B_f} \in [n]^{B_f}} \sum_{t_{L_f} \in [n]^{L_f}} \prod_{\ell=0}^k (\xi_{t_\ell} \cdot \mathbb{1}[\ell \notin P_{f,t}] + (1 - \xi_{t_\ell}) \cdot \mathbb{1}[\ell \in P_{f,t}]) \\
&\leq \frac{1}{k!} \sum_f 2^{|L_f|} \sum_{t_{B_f} \in [n]^{B_f}} \prod_{\ell \in B_f \cup \{0\}} (\xi_{t_\ell} \cdot \mathbb{1}[\ell \notin P_{f,t}] + (1 - \xi_{t_\ell}) \cdot \mathbb{1}[\ell \in P_{f,t}]) \\
&= \frac{1}{k!} \sum_f 2^{|L_f|} \sum_{t_{B_f} \in [n]^{B_f}} \prod_{\ell \in B_f \cup \{0\}} \xi_{t_\ell}
\end{aligned}$$

(as no non-leaf can ever be a petal)

$$\begin{aligned}
&= \frac{\xi_j}{k!} \sum_f 2^{|L_f|} \prod_{\ell \in B_f} \left(\sum_{t_\ell \in [n]} \xi_{t_\ell} \right) \\
&= \frac{\xi_j}{k!} \sum_f 2^{|L_f|} = \xi_j \mathbb{E}_{f \sim \mathcal{U}(\mathcal{F})} [2^{|L_f|}]
\end{aligned}$$

where \mathcal{F} is the set of all factorial trees and $\mathcal{U}(\mathcal{F})$ is the uniform distribution over \mathcal{F} . For a specific vertex $\ell \in [0, k]$, we note that $\ell \in L_f$ if and only if it is not the parent of any vertex $\ell + 1, \dots, k$. So,

$$\Pr_{f \sim \mathcal{U}(\mathcal{F})} [\ell \in L_f] = \prod_{i=\ell+1}^k \frac{i-1}{i} = \frac{\ell}{k} \tag{44}$$

We will show via induction that, for any vertex set $S \subseteq [0, k]$

$$\Pr_{f \sim \mathcal{U}(\mathcal{F})} [S \subseteq L_f] \leq \prod_{\ell \in S} \frac{\ell}{k} \tag{45}$$

Having established the base case for every S with $|S| = 1$, we assume (45) holds for all S with $|S| < s$. For any set of s vertices V , consider an arbitrary partition of V into two sets $S \cup T = V$ with $|S|, |T| < s$. We see

$$\begin{aligned}
\Pr_{f \sim \mathcal{U}(\mathcal{F})} [V \subseteq L_f] &= \prod_{c=1}^k \Pr [p_f(c) \notin V] \\
&= \prod_{c=1}^k \Pr [p_f(c) \notin S] \Pr [p_f(c) \notin T | p_f(c) \notin S] \\
&\leq \prod_{c=1}^k \Pr [p_f(c) \notin S] \Pr [p_f(c) \notin T] \\
&= \Pr [S \subseteq L_f] \Pr [T \subseteq L_f] \\
&\leq \prod_{\ell \in V} \frac{\ell}{k}
\end{aligned}$$

by the inductive hypothesis, as desired. Thus, $\Pr [|L_f| \geq s] \leq \sum_{S:|S|=s} \Pr [S \subseteq L_f]$ is at most the s^{th} coefficient of the polynomial

$$R(x) = \prod_{\ell=0}^k \left(1 + \frac{\ell}{k} x\right)$$

and so

$$\begin{aligned}
\mathbb{E}_{f \sim \mathcal{U}(\mathcal{F})} [2^{|L_f|}] &\leq \sum_{s=0}^k 2^s \Pr [|L_f| \geq s] \\
&\leq R(2) \\
&\leq e^{\sum_{\ell=0}^k 2\ell/k} = e^{k+1}
\end{aligned}$$

and

$$\sum_{\gamma \in \mathbb{Z}_{\geq 0}^n: |\gamma|=k} |a_{j,\gamma}| \leq \xi_j e^{k+1},$$

as desired. \square

Lemma B.6. Let ϕ_1, \dots, ϕ_m be softmax-type functions. That is, for each ϕ_i , there is some $j_i \in [n]$ and indices $\xi_{i1}, \dots, \xi_{in}$ such that

$$\phi_i((z_1, \dots, z_n)) = \frac{\exp(z_{j_i})}{\sum_{k=1}^n \xi_{ik} \cdot \exp(z_k)}$$

where $\xi_{i1} + \dots + \xi_{in} = 1$ for all i . Let $P(z) = \sum_{\gamma \in \mathbb{Z}_{\geq 0}^n} a_\gamma z^\gamma$ denote the Taylor series of $\prod_i \phi_i$. Then for any integer k ,

$$\sum_{\gamma \in \mathbb{Z}_{\geq 0}^n: |\gamma|=k} |a_\gamma| \leq (e^3 m)^k.$$

Proof. Letting $P_i(z) = \sum_{\gamma \in \mathbb{Z}_{\geq 0}^n} a_{i,\gamma} z^\gamma$ denote the Taylor series of ϕ_i for all i , we have $P(z) = \prod_i P_i(z)$ and therefore

$$\sum_{\gamma \in \mathbb{Z}_{\geq 0}^n: |\gamma|=k} |a_\gamma| \leq \sum_{\substack{k_1, \dots, k_m \in \mathbb{Z}_{\geq 0} \\ \sum k_i = k}} \prod_i \sum_{\gamma \in \mathbb{Z}_{\geq 0}^n: |\gamma|=k_i} |a_{i,\gamma}|$$

We have that $\sum_{|\gamma|=k_i} |a_{i,\gamma}| \leq e^{2k_i}$ for all k_i since, for $k_i = 0$, $a_{i,0} = \phi_i(0) = 1$, and for $k_i \geq 1$,

$$\sum_{|\gamma|=k_i} |a_{i,\gamma}| \leq \frac{\xi_{ij}}{\xi_{ij}} e^{k_i+1} \leq e^{2k_i} \quad (46)$$

from Lemma B.5. Note that the softmax-type functions discussed in Lemma B.5 have a ξ_{ij} term in the numerator, while those discussed here do not. This accounts for the extra ξ_{ij} term that appears in equation (46). Thus,

$$\begin{aligned} \sum_{\substack{k_1, \dots, k_m \in \mathbb{Z}_{\geq 0} \\ \sum k_i = k}} \prod_i \sum_{\gamma \in \mathbb{Z}_{\geq 0}^n: |\gamma| = k_i} |a_{i, \gamma}| &\leq \sum_{\substack{k_1, \dots, k_m \in \mathbb{Z}_{\geq 0} \\ \sum k_i = k}} e^{2k} \\ &= e^{2k} \binom{m+k-1}{k} \\ &\leq e^{2k} \left(\frac{e(m+k-1)}{k} \right)^k \\ &= (e^3 m)^k \end{aligned}$$

as desired. \square

Lemma B.7. *Let $\phi((z_1, \dots, z_n)) = \frac{\exp(z_j)}{\sum_{k=1}^n \xi_k \exp(z_k)}$ be any softmax-type function. Then the radius of convergence of the Taylor series of ϕ at the origin is at least 1.*

Proof. For a complex number z , write $\Re(z)$, $\Im(z)$ to denote the real and imaginary parts, respectively, of z . Note that for any $\zeta_1, \dots, \zeta_n \in \mathbb{C}$ with $|\zeta_k| \leq \pi/3$ for all $k \in [n]$, we have

$$\Re(\exp(\zeta_k)) \geq \cos(\pi/3) \cdot \exp(-\pi/3) > 1/10,$$

and thus $|\sum_{k=1}^n \xi_k \cdot \exp(\zeta_k)| \geq 1/10$. Moreover, for any such point $\zeta = (\zeta_1, \dots, \zeta_n)$, it holds that $|\exp(\zeta_j)| \leq \exp(\pi/3) < 3$. It then follows that for such ζ we have $|\phi(z)| \leq 30$. In particular, ϕ is holomorphic on the region $\{\zeta : |\zeta_k| \leq \pi/3 \ \forall k \in [n]\}$.

Fix any $\gamma \in \mathbb{Z}_{\geq 0}^n$, and let $k = |\gamma|$. By the multivariate version of Cauchy's integral formula,

$$\begin{aligned} \left| \frac{d^\gamma}{dz^\gamma} \phi(z) \right| &= \left| \frac{\gamma!}{(2\pi i)^n} \int_{|\zeta_1 - z_1| = \pi/3} \cdots \int_{|\zeta_n - z_n| = \pi/3} \frac{\phi(\zeta_1, \dots, \zeta_n)}{(\zeta_1 - z_1)^{\gamma_1+1} \cdots (\zeta_n - z_n)^{\gamma_n+1}} d\zeta_1 \cdots d\zeta_n \right| \\ &\leq \frac{30\gamma!}{(\pi/3)^{k+n}} \leq \frac{30\gamma!}{(\pi/3)^k}. \end{aligned}$$

The power series of ϕ at $\mathbf{0}$ is defined as $P_\phi(z) = \sum_{\gamma \in \mathbb{Z}_{\geq 0}^n} a_\gamma \cdot z^\gamma$, where $a_\gamma = \frac{1}{\gamma!} \frac{d^\gamma}{dz^\gamma} \phi(\mathbf{0})$. For any $\gamma \in \mathbb{Z}_{\geq 0}^n$ with $k = |\gamma|$, we have $|a_\gamma|^{1/k} \leq (30/(\pi/3)^k)^{1/k} = (30)^{1/k} \cdot 3/\pi$, which tends to $3/\pi < 1$ as $k \rightarrow \infty$. Thus, by the (multivariate version of the) Cauchy-Hadamard theorem, the radius of convergence of the power series of ϕ at $\mathbf{0}$ is at least $\pi/3 \geq 1$. \square

B.2 Proof of Lemma 4.6

In this section prove Lemma 4.6, which, as explained in Section 4.3, is an important ingredient in the proof of Lemma 4.5. The detailed version of Lemma 4.6 is presented below; it includes several claims which are omitted for simplicity in the abbreviated version in Section 4.3.

Lemma 4.6 (Detailed). *Fix any integer $h \geq 0$, a multi-index $\gamma \in \mathbb{Z}_{\geq 0}^n$ and set $k = |\gamma|$. For each of the k^h functions $\pi : [h] \rightarrow [k]$, and for each $r \in [k]$, there are integers $h'_{\pi, r} \in \{0, 1, \dots, h\}$, $t'_{\pi, r} \geq 0$, and $j'_{\pi, r} \in [n]$, so that the following holds. For any sequence $L^{(1)}, \dots, L^{(T)} \in \mathbb{R}^n$ of vectors, it holds that*

$$D_h L^\gamma = \sum_{\pi: [h] \rightarrow [k]} \prod_{r=1}^k E_{t'_{\pi, r}} D_{h'_{\pi, r}} (L(j'_{\pi, r})). \quad (47)$$

Moreover, the following properties hold:

1. For each π and $r \in [k]$, $h'_{\pi, r} = |\{q \in [h] : \pi(q) = r\}|$. In particular, $\sum_{r=1}^k h'_{\pi, r} = h$.
2. For each π and $r \in [k]$, it holds that $0 \leq t'_{\pi, r} + h'_{\pi, r} \leq h$.

3. For each $\pi, r \in [k]$, and $j \in [n]$, $\gamma_j = |\{r \in [k] : j'_{\pi,r} = j\}|$.

Proof of Lemma 4.6. We use induction on h . First note that in the case $h = 0$ and for any $k \geq 0$, we have that $(D_h L^\gamma)^{(t)} = (L^{(t)})^\gamma$, and so for the unique function $\pi : \emptyset \rightarrow [k]$, for all $r \in [k]$, we may take $t'_{\pi,r} = 0$, $h'_{\pi,r} = 0$, and ensure that for each $j \in [n]$ there are γ_j values of r so that $j'_{\pi,r} = j$.

Now fix any integer $h > 0$, and suppose the statement of the claim holds for all $h' < h$. We have that

$$\begin{aligned} & D_h L^\gamma \\ &= D_1 D_{h-1} L^\gamma \\ &= D_1 \sum_{\pi: [h-1] \rightarrow [k]} \prod_{r=1}^k E_{t'_{\pi,r}} D_{h'_{\pi,r}} L(j'_{\pi,r}) \\ &= \sum_{\pi: [h-1] \rightarrow [k]} \sum_{r=1}^k D_1 E_{t'_{\pi,r}} D_{h'_{\pi,r}} L(j'_{\pi,r}) \cdot \prod_{r'=1}^{r-1} E_{t'_{\pi,r'}} D_{h'_{\pi,r'}} L(j'_{\pi,r'}) \cdot \prod_{r'=r+1}^k E_{t'_{\pi,r'}+1} D_{h'_{\pi,r'}} L(j'_{\pi,r'}) \end{aligned} \quad (48)$$

$$= \sum_{\pi: [h-1] \rightarrow [k]} \sum_{r=1}^k E_{t'_{\pi,r}} D_{h'_{\pi,r}+1} L(j'_{\pi,r}) \cdot \prod_{r'=1}^{r-1} E_{t'_{\pi,r'}} D_{h'_{\pi,r'}} L(j'_{\pi,r'}) \cdot \prod_{r'=r+1}^k E_{t'_{\pi,r'}+1} D_{h'_{\pi,r'}} L(j'_{\pi,r'}). \quad (49)$$

where (48) uses Lemma B.2 and (49) uses the commutativity of $E_{t'}$ and D_1 . For each $\pi : [h-1] \rightarrow [k]$, we construct k functions $\pi_1, \dots, \pi_k : [h] \rightarrow [k]$, defined by $\pi_r(q) = \pi(q)$ for $q < h$, and $\pi_r(h) = r$ for $r \in [k]$. Next, for $r, r' \in [k]$, we define the quantities $h'_{\pi_r, r'}, t'_{\pi_r, r'}, j'_{\pi_r, r'}$ as follows:

- Set $h'_{\pi_r, r'} = h'_{\pi, r'}$ if $r \neq r'$, and $h'_{\pi_r, r} = h'_{\pi, r} + 1$.
- Set $t'_{\pi_r, r'} = t'_{\pi, r'}$ if $r' \leq r$, and $t'_{\pi_r, r'} = t'_{\pi, r'} + 1$ if $r' > r$.
- Set $j'_{\pi_r, r'} = j'_{\pi, r'}$.

By (49) and the above definitions, we have

$$D_h L^\gamma = \sum_{\pi: [h] \rightarrow [k]} \prod_{r'=1}^k E_{t'_{\pi, r'}} D_{h'_{\pi, r'}} L(j'_{\pi, r'}),$$

thus verifying (9) for the value h .

Finally we verify that items 1 through 3 in the lemma statement hold. The definition of $h'_{\pi_r, r'}$ above together with the inductive hypothesis ensures that for all $r, r' \in [k]$, $h'_{\pi_r, r'} = |\{q \in [h] : \pi_r(q) = r'\}|$, thus verifying item 1 of the lemma statement. Since $h'_{\pi_r, r'} + t'_{\pi_r, r'} \leq h'_{\pi, r} + t'_{\pi, r} + 1$ for all r, r' , it follows from the inductive hypothesis that $0 \leq h'_{\pi_r, r'} + t'_{\pi_r, r'} \leq h$; this verifies item 2. Finally, note that for any $j \in [n]$ and $r \in [k]$, $\{r' \in [k] : j'_{\pi_r, r'} = j\} = \{r' \in [k] : j'_{\pi, r'} = j\}$, and thus item 3 follows from the inductive hypothesis. \square

B.3 Proof of Lemma 4.5

In this section we prove Lemma 4.5. To introduce the detailed version of the lemma we need the following definition. Suppose $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$ is a real-valued function that is real-analytic in a neighborhood of the origin. For real numbers $Q, R > 0$, we say that ϕ is (Q, R) -bounded if the Taylor series of ϕ at $\mathbf{0}$, denoted $P_\phi(z_1, \dots, z_n) = \sum_{\gamma \in \mathbb{Z}_{\geq 0}^n} a_\gamma z^\gamma$, satisfies, for each integer $k \geq 0$, $\sum_{\gamma \in \mathbb{Z}_{\geq 0}^n : |\gamma| = k} |a_\gamma| \leq Q \cdot R^k$. In the statement of Lemma 4.5 below, the quantity 0^0 is interpreted as 1 (in particular, $(h')^{B_0 h'} = 1$ for $h' = 0$).

Lemma 4.5 (“Boundedness chain rule” for finite differences; detailed). *Suppose that $h, n \in \mathbb{N}$, $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$ is a (Q, R) -bounded function so that the radius of convergence of its power series at*

$\mathbf{0}$ is at least 1, and $L = (L^{(1)}, \dots, L^{(T)}) \in \mathbb{R}^n$ is a sequence of vectors satisfying $\|L^{(t)}\|_\infty \leq 1$ for $t \in [T]$. Suppose for some $\alpha \in (0, 1)$, for each $0 \leq h' \leq h$ and $t \in [T - h']$, it holds that $\|D_{h'} L^{(t)}\|_\infty \leq \frac{1}{B_1} \cdot \alpha^{h'} \cdot (h')^{B_0 h'}$ for some $B_1 \geq 2e^2 R$, $B_0 \geq 3$. Then for all $t \in [T - h]$,

$$|(D_h(\phi \circ L))^{(t)}| \leq \frac{12RQe^2}{B_1} \cdot \alpha^h \cdot h^{B_0 h + 1}.$$

Proof of Lemma 4.5. Note that the h th order finite differences of a constant sequence are identically 0 for $h \geq 1$, so by subtracting $\phi(\mathbf{0})$ from ϕ , we may assume without loss of generality that $\phi(\mathbf{0}) = 0$. (Here $\mathbf{0}$ denotes the all-zeros vector.)

By assumption, the radius of convergence of the power series of ϕ at the origin is at least 1, and so for each $\gamma \in \mathbb{Z}_{\geq 0}^n$, there is a real number a_γ so that for $z = (z_1, \dots, z_n)$ with $|z_j| \leq 1$ for each j ,

$$\phi(z) = \sum_{k \in \mathbb{N}, \gamma \in \mathbb{Z}_{\geq 0}^n: |\gamma|=k} a_\gamma z^\gamma. \quad (50)$$

Let $A_k := \sum_{\gamma \in \mathbb{Z}_{\geq 0}^n: |\gamma|=k} |a_\gamma|$; by the assumption that ϕ is (Q, R) -bounded, we have that $A_k \leq Q \cdot R^k$ for all $k \in \mathbb{N}$.

For $\gamma \in \mathbb{Z}_{\geq 0}^n$, recall that L^γ denotes the sequence $((L^\gamma)^{(1)}, \dots, (L^\gamma)^{(T)})$, defined by $(L^\gamma)^{(t)} = (L^{(t)}(1))^{\gamma_1} \dots (L^{(t)}(n))^{\gamma_n}$. Then since $\|L^{(t)}\|_\infty \leq 1$ for all $t \in [T]$, we have that, for $t \in [T - h]$, $(D_h(\phi \circ L))^{(t)} = \sum_{\gamma \in \mathbb{Z}_{\geq 0}^n} a_\gamma \cdot (D_h L^\gamma)^{(t)}$.

We next upper bound the quantities $|(D_h L^\gamma)^{(t)}|$. To do so, fix some $\gamma \in \mathbb{Z}_{\geq 0}^n$, and set $k = |\gamma|$. For each function $\pi : [h] \rightarrow [k]$ and $r \in [k]$, recall the integers $h'_{\pi, r} \in \{0, 1, \dots, h\}$, $t'_{\pi, r} \geq 0$, $j'_{\pi, r} \in [n]$ defined in Lemma 4.6. By assumption it holds that for each $t \in [T - h]$, each $h' \leq h$, each $0 \leq t' \leq h$, $|(D_{h'} L(j))^{(t+t')}| \leq \frac{1}{B_1} \cdot \alpha^{h'} \cdot (h')^{B_0 h'}$. It follows that for each $t \in [T - h]$ and function $\pi : [h] \rightarrow [k]$,

$$\left| \prod_{r=1}^k \left(E_{t'_{\pi, r}} D_{h'_{\pi, r}} L(j'_{\pi, r}) \right)^{(t)} \right| \leq \prod_{r=1}^k \frac{1}{B_1} \cdot \alpha^{h'_{\pi, r}} \cdot (h'_{\pi, r})^{B_0 h'_{\pi, r}} = \frac{\alpha^h}{B_1^k} \cdot \prod_{r=1}^k (h'_{\pi, r})^{B_0 h'_{\pi, r}},$$

where the last equality uses that $\sum_{r=1}^k h'_{\pi, r} = h$ (item 1 of Lemma 4.6). Then by Lemma 4.6, we have:

$$\begin{aligned} |(D_h L^\gamma)^{(t)}| &\leq \sum_{\pi: [h] \rightarrow [k]} \left| \prod_{r=1}^k \left(E_{t'_{\pi, r}} D_{h'_{\pi, r}} L(j'_{\pi, r}) \right)^{(t)} \right| \\ &\leq \frac{\alpha^h}{B_1^k} \sum_{\pi: [h] \rightarrow [k]} \prod_{r=1}^k (h'_{\pi, r})^{B_0 h'_{\pi, r}} \\ &\leq \frac{\alpha^h}{B_1^k} \cdot h^{B_0 h} \max \left\{ k^7, (hk + 1) \cdot \exp \left(\frac{2k}{h^{B_0 - 1}} \right) \right\}, \end{aligned} \quad (51)$$

where (51) follows from Lemma B.4, the fact that $B_0 \geq 3$, and that $h'_{\pi, r} = |\{q \in [h] : \pi(q) = r\}|$ (item 1 of Lemma 4.6).

We may now bound the order- h finite differences of the sequence $\phi \circ L$ as follows: for $t \in [T - h]$,

$$\begin{aligned} |(\mathbb{D}_h(\phi \circ L))^{(t)}| &\leq \sum_{\gamma \in \mathbb{Z}_{\geq 0}^n} |a_\gamma| \cdot \left| (\mathbb{D}_h L^\gamma)^{(t)} \right| \\ &\leq \alpha^h \cdot h^{B_0 h + 1} \sum_{\gamma \in \mathbb{Z}_{\geq 0}^n} |a_\gamma| \cdot B_1^{-|\gamma|} \cdot \max \left\{ |\gamma|^{|\gamma|}, (|\gamma| + 1) \cdot \exp \left(\frac{2|\gamma|}{h^{B_0 - 1}} \right) \right\} \end{aligned} \quad (52)$$

$$\begin{aligned} &\leq \alpha^h \cdot h^{B_0 h + 1} \cdot \sum_{k \in \mathbb{N}} A_k \cdot B_1^{-k} \cdot (k^7 + 2k \cdot \exp(2k/h^{B_0 - 1})) \\ &\leq \alpha^h \cdot h^{B_0 h + 1} \cdot Q \cdot \left(\sum_{k \in \mathbb{N}} k^7 \cdot (R/B_1)^k + \sum_{k \in \mathbb{N}} 2k \cdot (R/B_1)^k \cdot e^{2k} \right) \end{aligned} \quad (53)$$

$$\begin{aligned} &\leq \frac{2RQe^2}{B_1} \cdot \alpha^h \cdot h^{B_0 h + 1} \cdot \left(\sum_{k \in \mathbb{N}} k^7 \cdot (2e^2)^{-k} + 2 \sum_{k \in \mathbb{N}} k \cdot 2^{-k} \right) \\ &= \frac{12RQe^2}{B_1} \cdot \alpha^h \cdot h^{B_0 h + 1}. \end{aligned} \quad (54)$$

where (52) uses (51), (53) uses the bound $A_k \leq QR^k$, and (54) uses the assumption $B_1 \geq 2e^2 R$. This gives the desired conclusion of the lemma. \square

B.4 Proof of Lemma 4.4

In this section we prove Lemma 4.4. The detailed version of Lemma 4.4 is stated below.

Lemma 4.4 (Detailed). *Fix a parameter $\alpha \in \left(0, \frac{1}{H+3}\right)$. If all players follow Optimistic Hedge updates with step size $\eta \leq \frac{\alpha}{36e^5 m}$, then for any player $i \in [m]$, integer h satisfying $0 \leq h \leq H$, time step $t \in [T - h]$, it holds that*

$$\|(\mathbb{D}_h \ell_i)^{(t)}\|_\infty \leq \alpha^h \cdot h^{3h+1}.$$

Proof. We have that for each agent $i \in [m]$, each $t \in [T]$, and each $a_i \in [n_i]$, $\ell_i^{(t)}(a_i) = \mathbb{E}_{a_{i'} \sim x_{i'}^{(t)}: i' \neq i} [\mathcal{L}_i(a_1, \dots, a_m)]$. Thus, for $1 \leq t \leq T$,

$$\begin{aligned} |(\mathbb{D}_h \ell_i)^{(t)}(a_i)| &= \left| \sum_{s=0}^h \binom{h}{s} (-1)^{h-s} \ell_i^{(t+s)}(a_i) \right| \\ &= \left| \sum_{a_{i'} \in [n_{i'}], \forall i' \neq i} \mathcal{L}_i(a_1, \dots, a_m) \sum_{s=0}^h \binom{h}{s} (-1)^{h-s} \cdot \prod_{i' \neq i} x_{i'}^{(t+s)}(a_{i'}) \right| \\ &\leq \sum_{a_{i'} \in [n_{i'}], \forall i' \neq i} \left| \sum_{s=0}^h \binom{h}{s} (-1)^{h-s} \cdot \prod_{i' \neq i} x_{i'}^{(t+s)}(a_{i'}) \right| \\ &= \sum_{a_{i'} \in [n_{i'}], \forall i' \neq i} \left| \left(\mathbb{D}_h \left(\prod_{i' \neq i} x_{i'}(a_{i'}) \right) \right)^{(t)} \right|, \end{aligned} \quad (55)$$

where (55) and (56) use Remark 4.3 and in (56), $\prod_{i' \neq i} x_{i'}(a_{i'})$ refers to the sequence $\prod_{i' \neq i} x_{i'}^{(1)}(a_{i'})$, $\prod_{i' \neq i} x_{i'}^{(2)}(a_{i'})$, \dots , $\prod_{i' \neq i} x_{i'}^{(T)}(a_{i'})$.

In the remainder of this lemma we will prepend to the loss sequence $\ell_i^{(1)}, \dots, \ell_i^{(T)}$ the vectors $\ell_i^{(0)} = \ell_i^{(-1)} := \mathbf{0} \in \mathbb{R}^{n_i}$. We will also prepend $x_i^{(0)} := x_i^{(1)} = (1/n_i, \dots, 1/n_i) \in \Delta^{n_i}$ to the strategy sequence $x_i^{(1)}, \dots, x_i^{(T)}$. Next notice that for any agent $i \in [m]$, any $t_0 \in \{0, 1, \dots, T\}$, and

any $t \geq 0$, by the definition (1) of the Optimistic Hedge updates, it holds that, for each $j \in [n_i]$,

$$x_i^{(t_0+t+1)}(j) = \frac{x_i^{(t_0)}(j) \cdot \exp\left(\eta \cdot \left(\ell_i^{(t_0-1)}(j) - \sum_{s=0}^t \ell_i^{(t_0+s)}(j) - \ell_i^{(t_0+t)}(j)\right)\right)}{\sum_{k=1}^{n_i} x_i^{(t_0)}(k) \cdot \exp\left(\eta \cdot \left(\ell_i^{(t_0-1)}(k) - \sum_{s=0}^t \ell_i^{(t_0+s)}(k) - \ell_i^{(t_0+t)}(k)\right)\right)}.$$

Note in particular that our definitions of $\ell_i^{(0)}$, $\ell_i^{(-1)}$, $x_i^{(0)}$ ensure that the above equation holds even for $t_0 \in \{0, 1\}$. Now an integer t_0 satisfying $0 \leq t_0 \leq T$; for $t \geq 0$, let us write

$$\bar{\ell}_{i,t_0}^{(t)} := \ell_i^{(t_0-1)} - \sum_{s=0}^{t-1} \ell_i^{(t_0+s)} - \ell_i^{(t_0+t-1)}.$$

Also, for a vector $z = (z(1), \dots, z(n_i)) \in \mathbb{R}^{n_i}$ and an index $j \in [n_i]$, define

$$\phi_{t_0,j}(z) := \frac{\exp(z(j))}{\sum_{k=1}^{n_i} x_i^{(t_0)}(k) \cdot \exp(z(k))}, \quad (57)$$

so that $x_i^{(t_0+t)}(j) = x_i^{(t_0)}(j) \cdot \phi_{t_0,j}(\eta \cdot \bar{\ell}_{i,t_0}^{(t)})$ for $t \geq 1$. In particular, for any $i \in [m]$, and any choices of $a_{i'} \in [n_{i'}]$ for all $i' \neq i$,

$$\prod_{i' \neq i} x_{i'}^{(t_0+t)}(a_{i'}) = \prod_{i' \neq i} x_{i'}^{(t_0)}(a_{i'}) \cdot \phi_{t_0,a_{i'}}(\eta \cdot \bar{\ell}_{i',t_0}^{(t)}). \quad (58)$$

Next, note that

$$(\mathbf{D}_1 \bar{\ell}_{i,t_0}^{(t)})^{(t)} = \ell_i^{(t_0+t-1)} - 2\ell_i^{(t_0+t)} = \ell_i^{(t_0+t-1)} - 2(\mathbf{E}_1 \ell_i)^{(t_0+t-1)},$$

meaning that for any $h' \geq 1$,

$$(\mathbf{D}_{h'} \bar{\ell}_{i,t_0}^{(t)})^{(t)} = (\mathbf{D}_{h'-1} \ell_i)^{(t_0+t-1)} - 2(\mathbf{E}_1 \mathbf{D}_{h'-1} \ell_i)^{(t_0+t-1)}. \quad (59)$$

We next establish the following claims which will allow us to prove Lemma 4.4 by induction.

Claim B.8. For any $t_0 \in \{0, 1, \dots, T\}$, $t \geq 0$, and $i \in [m]$, it holds that $\|\bar{\ell}_{i,t_0}^{(t)}\|_\infty \leq t + 2$.

Proof of Claim B.8. The claim is immediate from the triangle inequality and the fact that $\|\ell_i^{(t)}\|_\infty \leq 1$ for all $t \in [T]$. \square

Claim B.9. Fix h so that $1 \leq h \leq H$. Suppose that for some $B_0 \geq 3$ and for all $0 \leq h' < h$, all $i \in [m]$, and all $t \leq T - h'$, it holds that $\|(\mathbf{D}_{h'} \ell_i)^{(t)}\|_\infty \leq \alpha^{h'} \cdot (h' + 1)^{B_0(h'+1)}$. Suppose that the step size η satisfies $\eta \leq \min\left\{\frac{\alpha}{36e^5 m}, \frac{1}{12e^5(H+3)m}\right\}$. Then for all $i \in [m]$ and $1 \leq t \leq T - h$,

$$\|(\mathbf{D}_h \ell_i)^{(t)}\|_\infty \leq \alpha^h \cdot h^{B_0 h + 1}. \quad (60)$$

Proof of Claim B.9. Set $B_1 := 12e^5 m$, so that the assumption of the claim gives $\eta \leq \min\left\{\frac{\alpha}{3B_1}, \frac{1}{B_1(H+3)}\right\}$.

We first use Lemma 4.5 to bound, for each $0 \leq t_0 \leq T - h$, $i \in [m]$, and $a_{i'} \in [n_{i'}]$ for all $i' \neq i$, the quantity $\left|(\mathbf{D}_h \left(\prod_{i' \neq i} x_{i'}(a_{i'})\right))^{(t_0+1)}\right|$. In particular, we will apply Lemma 4.5 with $n = \sum_{i' \neq i} n_{i'}$, the value of h in the statement of Claim B.9, $T = h + 1$, and the sequence $L^{(t)}$, for $1 \leq t \leq h + 1$, defined as

$$L^{(t)} = \left(\eta \cdot \bar{\ell}_{1,t_0}^{(t)}, \dots, \eta \cdot \bar{\ell}_{i-1,t_0}^{(t)}, \eta \cdot \bar{\ell}_{i+1,t_0}^{(t)}, \dots, \eta \cdot \bar{\ell}_{m,t_0}^{(t)}\right),$$

namely the concatenation of the vectors $\eta \cdot \bar{\ell}_{1,t_0}^{(t)}, \dots, \eta \cdot \bar{\ell}_{i-1,t_0}^{(t)}, \eta \cdot \bar{\ell}_{i+1,t_0}^{(t)}, \dots, \eta \cdot \bar{\ell}_{m,t_0}^{(t)}$. The function ϕ in Lemma 4.5 is set to the function that takes as input the concatenation of $z_{i'} \in \mathbb{R}^{n_{i'}}$ for all $i' \neq i$ and outputs:

$$\phi_{t_0, a_{-i}}(z_1, \dots, z_{i-1}, z_{i+1}, \dots, z_m) := \prod_{i' \neq i} \phi_{t_0, a_{i'}}(z_{i'}), \quad (61)$$

where the function $\phi_{t_0, a_{i'}}$ are as defined in (57). We first verify the preconditions of Lemma 4.5. By Lemma B.6, $\phi_{t_0, a_{-i}}$ is a $(1, e^3 m)$ -bounded function for some constant $C \geq 1$. By Lemma B.7, the radius of convergence of each function $\phi_{t_0, a_{i'}}$ at $\mathbf{0}$ is at least 1; thus the radius of convergence of $\phi_{t_0, a_{-i}}$ at $\mathbf{0}$ is at least 1. Claim B.8 gives that $\|\bar{\ell}_{i,t_0}^{(t)}\|_\infty \leq t + 2 \leq h + 3$ for all $t \leq h + 1$. Thus, since $\eta \leq \frac{1}{B_1(H+3)}$,

$$\left\| (D_0(\eta \cdot \bar{\ell}_{i,t_0}))^{(t)} \right\|_\infty = \|\eta \cdot \bar{\ell}_{i,t_0}^{(t)}\|_\infty \leq \eta \cdot (H + 3) \leq \frac{1}{B_1}$$

for $1 \leq t \leq h_0 + 1$. Next, for $1 \leq h' \leq h$ and $1 \leq t \leq h + 1 - h'$, we have

$$\left\| (D_{h'}(\eta \cdot \bar{\ell}_{i,t_0}))^{(t)} \right\|_\infty \leq \eta \cdot \left\| (D_{h'-1} \ell_i)^{(t_0+t-1)} \right\|_\infty + 2\eta \cdot \left\| (D_{h'-1} \ell_i)^{(t_0+t)} \right\|_\infty \quad (62)$$

$$\leq 3\eta \cdot \alpha^{h'-1} \cdot (h')^{B_0(h')} \quad (63)$$

$$\leq \frac{1}{B_1} \cdot \alpha^{h'} \cdot (h')^{B_0(h')}, \quad (64)$$

where (62) follows from (59), (63) follows from the assumption in the statement of Claim B.9 and $t_0 + t + h' - 1 \leq t_0 + h \leq T$, and (64) follows from the fact that $3\eta \leq \frac{\alpha}{B_1}$. It then follows from Lemma 4.5 and (58) that

$$\begin{aligned} & \frac{1}{\prod_{i' \neq i} x_{i'}^{(t_0)}(a_{i'})} \cdot \left| \left(D_h \left(\prod_{i' \neq i} x_{i'}(a_{i'}) \right) \right)^{(t_0+1)} \right| \\ &= \left| \left(D_h \left(\prod_{i' \neq i} (\phi_{t_0, a_{i'}} \circ \eta \bar{\ell}_{i', t_0}) \right) \right)^{(t_0+1)} \right| \end{aligned} \quad (65)$$

$$= \left| (D_h(\phi_{t_0, a_{-i}} \circ (\eta \bar{\ell}_{1,t_0}, \dots, \eta \bar{\ell}_{i-1,t_0}, \eta \bar{\ell}_{i+1,t_0}, \dots, \eta \bar{\ell}_{m,t_0})))^{(1)} \right| \quad (66)$$

$$\leq \frac{12e^5 m}{B_1} \cdot \alpha^h \cdot h^{B_0 h + 1} = \alpha^h \cdot (h)^{B_0 h + 1}. \quad (67)$$

(In particular, (65) uses (58), (66) uses the definition of $\phi_{t_0, a_{-i}}$ in (61), and (67) uses Lemma 4.5.)

Next we use (56), which gives that for each $i \in [m]$ and $t \geq 1$,

$$\begin{aligned} \left\| (D_h \ell_i)^{(t)} \right\|_\infty &\leq \sum_{a_{i'} \in [n_{i'}], \forall i' \neq i} \left| \left(D_h \left(\prod_{i' \neq i} x_{i'}(a_{i'}) \right) \right)^{(t)} \right| \\ &\leq \sum_{a_{i'} \in [n_{i'}], \forall i' \neq i} \prod_{i' \neq i} x_{i'}^{(t_0)}(a_{i'}) \cdot \alpha^h \cdot (h)^{B_0 h + 1} \\ &= \alpha^h \cdot (h)^{B_0 h + 1}, \end{aligned} \quad (68)$$

where (68) follows from (67) with $t = t_0 + 1$ (here we use that t_0 may be 0). This completes the proof of Claim B.9. \square

It is immediate that for all $i \in [m], t \in [T]$, we have that $\|(D_0 \ell_i)^{(t)}\|_\infty \leq 1 = \alpha^0 \cdot 1^{B_0 \cdot 1}$. We now apply Claim B.9 inductively with $B_0 = 3$, for which it suffices to have $\eta \leq \frac{\alpha}{36e^5 m}$ as long as $\alpha \leq 1/(H + 3)$. This gives that for $0 \leq h \leq H, i \in [m]$, and $t \in [T - h]$, $\|(D_h \ell_i)^{(t)}\|_\infty \leq \alpha^h \cdot h^{3h+1}$, completing the proof of Lemma 4.4. \square

C Proofs for Section 4.4

The main goal of this section is to prove Lemma 4.7. First, in Section C.1 we prove some preliminary lemmas and then we prove Lemma 4.7 in Section C.2

C.1 Preliminary lemmas

Lemma C.1 shows that $\text{Var}_P(W)$ and $\text{Var}_{P'}(W)$ are close when the entries of P, P' are close; it will be applied with P, P' equal to the strategies $x_i^{(t)} \in \Delta^{n_i}$ played in the course of Optimistic Hedge.

Lemma C.1. *Suppose $n \in \mathbb{N}$ and $M > 0$ are given, and $W \in \mathbb{R}^n$ is a vector. Suppose $P, P' \in \Delta^n$ are distributions with $\max\left\{\left\|\frac{P}{P'}\right\|_\infty, \left\|\frac{P'}{P}\right\|_\infty\right\} \leq 1 + \alpha$ for some $\alpha > 0$. Then*

$$(1 - \alpha) \text{Var}_P(W) \leq \text{Var}_{P'}(W) \leq (1 + \alpha) \text{Var}_P(W). \quad (69)$$

Proof. We first prove that $\text{Var}_{P'}(W) \leq (1 + \alpha) \text{Var}_P(W)$. To do so, note that since adding a constant to every entry of W does not change $\text{Var}_P(W)$ or $\text{Var}_{P'}(W)$, by replacing W with $W - \langle P, W \rangle \cdot \mathbf{1}$, we may assume without loss of generality that $\langle P, W \rangle = 0$. Thus $\text{Var}_P(W) = \sum_{j=1}^n P(j)W(j)^2$. Now we may compute:

$$\begin{aligned} \text{Var}_{P'}(W) &\leq \sum_j P'(j) \cdot W(j)^2 \\ &= \sum_j P(j) \cdot W(j)^2 + \sum_j (P'(j) - P(j)) \cdot W(j)^2 \\ &= (1 + \alpha) \text{Var}_P(W), \end{aligned} \quad (70)$$

where (70) uses the fact that $\left\|\frac{P'}{P}\right\|_\infty \leq 1 + \alpha$.

By interchanging the roles of P, P' , we obtain that

$$\text{Var}_{P'}(W) \geq \frac{1}{1 + \alpha} \text{Var}_P(W) \geq (1 - \alpha) \text{Var}_P(W).$$

This completes the proof of the lemma. \square

Next we prove Lemma 4.8 (recall that only the special case $\mu = 0$ was proved in Section 4.4). For convenience the lemma is repeated below.

Lemma 4.8 (Restated). *Suppose $\mu \in \mathbb{R}$, $\alpha > 0$, and $W^{(0)}, \dots, W^{(S-1)} \in \mathbb{R}$ is a sequence of reals satisfying*

$$\sum_{t=0}^{S-1} \left((D_2^\circ W)^{(t)} \right)^2 \leq \alpha \cdot \sum_{t=0}^{S-1} \left((D_1^\circ W)^{(t)} \right)^2 + \mu. \quad (71)$$

Then

$$\sum_{t=0}^{S-1} \left((D_1^\circ W)^{(t)} \right)^2 \leq \alpha \cdot \sum_{t=1}^{S-1} (W^{(t)})^2 + \mu/\alpha.$$

To prove Lemma 4.8 we need the following basic facts about the Fourier transform:

Fact C.2 (Parseval's equality). *It holds that $\sum_{t=0}^{S-1} |W^{(t)}|^2 = \frac{1}{S} \sum_{s=0}^{S-1} |\widehat{W}^{(s)}|^2$.*

The second fact gives a formula for the Fourier transform of the circular finite differences; its simple form is the reason we work with *circular* finite differences in this section:

Fact C.3. *For $h \in \mathbb{Z}_{\geq 0}$, $\widehat{D_h^\circ W}^{(s)} = \widehat{W}^{(s)} \cdot (e^{2\pi i s t / S} - 1)^h$.*

Proof of Lemma 4.8. Note that the discrete Fourier transform of $D_1^\circ W$ satisfies $\widehat{D_1^\circ W}^{(s)} = \widehat{W}^{(s)} \cdot (e^{2\pi is/T} - 1)$, and similarly $\widehat{D_2^\circ W}^{(s)} = \widehat{W}^{(s)} \cdot (e^{2\pi is/T} - 1)^2$, for $0 \leq s \leq S-1$. By the Cauchy-Schwarz inequality, Parseval's equality (Fact C.2), Fact C.3, and the assumption that (71) holds, we have

$$\begin{aligned}
\sum_{t=0}^{S-1} \left((D_1^\circ W)^{(t)} \right)^2 &= \frac{1}{S} \sum_{s=0}^{S-1} \left| \widehat{D_1^\circ W}^{(s)} \right|^2 \\
&= \frac{1}{S} \sum_{s=0}^{S-1} \left| \widehat{W}^{(s)} \cdot (e^{2\pi is/T} - 1) \right|^2 \\
&= \frac{1}{S} \sum_{s=0}^{S-1} \left| \widehat{W}^{(s)} \right| \cdot \left| \widehat{W}^{(s)} \right| \cdot \left| e^{2\pi is/T} - 1 \right|^2 \\
&\leq \sqrt{\frac{1}{S} \sum_{s=0}^{S-1} \left| \widehat{W}^{(s)} \right|^2} \cdot \sqrt{\frac{1}{S} \sum_{s=0}^{S-1} \left| \widehat{W}^{(s)} \right|^2} \cdot \left| e^{2\pi is/T} - 1 \right|^4 \\
&= \sqrt{\sum_{t=0}^{S-1} (W^{(t)})^2} \cdot \sqrt{\frac{1}{S} \sum_{s=0}^{S-1} \left| \widehat{D_2^\circ W}^{(s)} \right|^2} \\
&= \sqrt{\sum_{t=0}^{S-1} (W^{(t)})^2} \cdot \sqrt{\sum_{t=0}^{S-1} \left((D_2^\circ W)^{(t)} \right)^2} \\
&\leq \sqrt{\sum_{t=0}^{S-1} (W^{(t)})^2} \cdot \sqrt{\alpha \cdot \sum_{t=0}^{S-1} \left((D_1^\circ W)^{(t)} \right)^2 + \mu}. \tag{72}
\end{aligned}$$

Note that for real numbers $A > 0$ and ϵ with $A + \epsilon > 0$, it holds that

$$\frac{A^2}{A + \epsilon} = \frac{A}{1 + \epsilon/A} \geq A \cdot (1 - \epsilon/A) = A - \epsilon.$$

Taking $A = \sum_{t=0}^{S-1} \left((D_1^\circ W)^{(t)} \right)^2$ and $\epsilon = \mu/\alpha$ (for which $A + \epsilon > 0$ is immediate) and using (72) then gives

$$\sum_{t=0}^{S-1} \left((D_1^\circ W)^{(t)} \right)^2 - \mu/\alpha \leq \frac{\left(\sum_{t=0}^{S-1} \left((D_1^\circ W)^{(t)} \right)^2 \right)^2}{\sum_{t=0}^{S-1} \left((D_1^\circ W)^{(t)} \right)^2 + \mu/\alpha} \leq \alpha \cdot \sum_{t=0}^{S-1} (W^{(t)})^2,$$

as desired. \square

C.2 Proof of Lemma 4.7

Now we prove Lemma 4.7. For convenience we restate the lemma below with the exact value of the constant C_0 referred to in the version in Section 4.4.

Lemma 4.7 (Restated). *For any $M, \zeta, \alpha > 0$ and $n \in \mathbb{N}$, suppose that $P^{(1)}, \dots, P^{(T)} \in \Delta^n$ and $Z^{(1)}, \dots, Z^{(T)} \in [-M, M]^n$ satisfy the following conditions:*

1. *The sequence $P^{(1)}, \dots, P^{(T)}$ is ζ -consecutively close for some $\zeta \in [1/(2T), \alpha^4/8256]$.*
2. *It holds that $\sum_{t=1}^{T-2} \text{Var}_{P^{(t)}} \left((D_2 Z)^{(t)} \right) \leq \alpha \cdot \sum_{t=1}^{T-1} \text{Var}_{P^{(t)}} \left((D_1 Z)^{(t)} \right) + \mu$.*

Then

$$\sum_{t=1}^{T-1} \text{Var}_{P^{(t)}} \left((D_1 Z)^{(t)} \right) \leq \alpha \cdot (1 + \alpha) \sum_{t=1}^T \text{Var}_{P^{(t)}} \left(Z^{(t)} \right) + \frac{\mu}{\alpha} + \frac{1290M^2}{\alpha^3}. \tag{73}$$

Proof. Fix a positive integer $S < 1/(2\zeta) < T$, to be specified exactly below. For $1 \leq t_0 \leq T - S + 1$, define $\mu_{t_0} \in \mathbb{R}$ by

$$\mu_{t_0} = \sum_{s=0}^{S-3} \text{Var}_{P^{(t_0+s)}} \left((D_2 Z)^{(t_0+s)} \right) - \alpha \cdot \sum_{s=0}^{S-3} \text{Var}_{P^{(t_0+s)}} \left((D_1 Z)^{(t_0+s)} \right). \quad (74)$$

Then

$$\begin{aligned} & \sum_{t_0=1}^{T-S+1} \mu_{t_0} \\ &= \sum_{t=1}^{T-2} \text{Var}_{P^{(t)}} \left((D_2 Z)^{(t)} \right) \cdot \min\{S-2, t, T-t-1\} \\ & \quad - \alpha \cdot \sum_{t=1}^{T-1} \text{Var}_{P^{(t)}} \left((D_1 Z)^{(t)} \right) \cdot \min\{S-2, t, T-t-1\} \\ & \leq (S-2) \cdot \sum_{t=1}^{T-2} \text{Var}_{P^{(t)}} \left((D_2 Z)^{(t)} \right) - (S-2)\alpha \cdot \sum_{t=1}^{T-1} \text{Var}_{P^{(t)}} \left((D_1 Z)^{(t)} \right) + 8\alpha(S-2)^2 M^2 \\ & \leq (S-2)\mu + 2\alpha(S-2)^2 M^2, \end{aligned} \quad (75)$$

$$(76)$$

where (75) uses the fact that $\|Z^{(t)}\|_\infty \leq M$ and so $\|(D_1 Z)^{(t)}\|_\infty \leq 2M$ for all $t \in [T]$, and the final inequality (76) follows from assumption 2 of the lemma statement.

By (74) and Lemma C.1 with $P = P^{(t_0)}$, we have, for some constant $C > 0$,

$$\begin{aligned} \sum_{s=0}^{S-3} \text{Var}_{P^{(t_0)}} \left((D_2 Z)^{(t_0+s)} \right) & \leq (1 + 2\zeta S) \cdot \sum_{s=0}^{S-3} \text{Var}_{P^{(t_0+s)}} \left((D_2 Z)^{(t_0+s)} \right) \\ & = (1 + 2\zeta S)\alpha \cdot \sum_{s=0}^{S-3} \text{Var}_{P^{(t_0+s)}} \left((D_1 Z)^{(t_0+s)} \right) + (1 + 2\zeta S)\mu_{t_0} \\ & \leq (1 + 2\zeta S)^2 \alpha \cdot \sum_{s=0}^{S-3} \text{Var}_{P^{(t_0)}} \left((D_1 Z)^{(t_0+s)} \right) + (1 + 2\zeta S)\mu_{t_0}. \end{aligned} \quad (77)$$

Here we have used that for $0 \leq s \leq S$, it holds that $\max \left\{ \left\| \frac{P^{(t_0+s)}}{P^{(t_0)}} \right\|_\infty, \left\| \frac{P^{(t_0)}}{P^{(t_0+s)}} \right\|_\infty \right\} \leq (1 + \zeta)^S \leq 1 + 2\zeta S$ since $\zeta S \leq 1/2$.

For any integer $1 \leq t_0 \leq T - S + 1$, we define the sequence $Z_{t_0}^{(s)} := Z^{(t_0+s)} - \langle P^{(t_0)}, Z^{(t_0+s)} \rangle \mathbf{1}$, for $0 \leq s \leq S - 1$. Thus $\langle Z_{t_0}^{(s)}, P^{(t_0)} \rangle = 0$ for $0 \leq s \leq S - 1$, which implies that for all $h \geq 0$, $0 \leq s \leq S - 1$, $\langle (D_h^\circ Z_{t_0})^{(s)}, P^{(t_0)} \rangle = 0$, and thus

$$\text{Var}_{P^{(t_0)}} \left((D_h^\circ Z_{t_0})^{(s)} \right) = \sum_{j=1}^n P^{(t_0)}(j) \cdot (D_h^\circ Z_{t_0})^{(s)}(j)^2. \quad (78)$$

By the definition of the sequence Z_{t_0} , for $0 \leq s \leq S - h - 1$, we have

$$\text{Var}_{P^{(t_0)}} \left((D_h Z)^{(t_0+s)} \right) = \text{Var}_{P^{(t_0)}} \left((D_h Z_{t_0})^{(s)} \right) = \text{Var}_{P^{(t_0)}} \left((D_h^\circ Z_{t_0})^{(s)} \right). \quad (79)$$

For $1 \leq t_0 \leq T - S + 1$, let us now define

$$\nu_{t_0, j} := \sum_{s=0}^{S-1} (D_2^\circ Z_{t_0})^{(s)}(j)^2 - (1 + 2\zeta S)^2 \alpha \cdot \sum_{s=0}^{S-1} (D_1^\circ Z_{t_0})^{(s)}(j)^2, \quad (80)$$

so that, by (77), (78), and (79),

$$\begin{aligned}
& \sum_{j=1}^n P^{(t_0)}(j) \cdot \nu_{t_0,j} \\
&= \sum_{s=0}^{S-1} \text{Var}_{P^{(t_0)}} \left((D_2^\circ Z_{t_0})^{(s)} \right) - (1 + 2\zeta S)^2 \alpha \cdot \sum_{s=0}^{S-1} \text{Var}_{P^{(t_0)}} \left((D_1^\circ Z_{t_0})^{(s)} \right) \\
&\leq \left(\sum_{s=0}^{S-3} \text{Var}_{P^{(t_0)}} \left((D_2 Z)^{(t_0+s)} \right) \right) + \text{Var}_{P^{(t_0+1)}} \left((D_2^\circ Z)^{(t_0+S-2)} \right) + \text{Var}_{P^{(t_0+1)}} \left((D_2^\circ Z)^{(t_0+S-1)} \right) \\
&\quad - (1 + 2\zeta S)^2 \alpha \cdot \sum_{s=0}^{S-3} \text{Var}_{P^{(t_0)}} \left((D_1 Z)^{(t_0+s)} \right) \\
&\leq (1 + 2\zeta S) \mu_{t_0} + \text{Var}_{P^{(t_0)}} \left((D_2^\circ Z)^{(t_0+S-2)} \right) + \text{Var}_{P^{(t_0)}} \left((D_2^\circ Z)^{(t_0+S-1)} \right). \tag{81}
\end{aligned}$$

By (80) and Lemma 4.8 applied to the sequence $Z_{t_0}^{(0)}, \dots, Z_{t_0}^{(S-1)}$, it holds that, for each $j \in [n]$,

$$\sum_{s=0}^{S-1} (D_1^\circ Z_{t_0})^{(s)}(j)^2 \leq (1 + \zeta CS)^2 \alpha \cdot \sum_{s=0}^{S-1} Z_{t_0}^{(s)}(j)^2 + \frac{\nu_{t_0,j}}{(1 + \zeta CS)^2 \alpha}. \tag{82}$$

Then we have:

$$\begin{aligned}
& \sum_{s=0}^{S-2} \text{Var}_{P^{(t_0)}} \left((D_1 Z)^{(t_0+s)} \right) \\
&= \sum_{s=0}^{S-2} \text{Var}_{P^{(t_0)}} \left((D_1^\circ Z_{t_0})^{(s)} \right) \tag{83}
\end{aligned}$$

$$\leq (1 + 2\zeta S)^2 \alpha \cdot \sum_{s=0}^{S-1} \text{Var}_{P^{(t_0)}} \left(Z_{t_0}^{(s)} \right) + \sum_{j=1}^n P^{(t_0)}(j) \cdot \frac{\nu_{t_0,j}}{(1 + 2\zeta S)^2 \alpha} \tag{84}$$

$$\leq (1 + 2\zeta S)^2 \alpha \sum_{s=0}^{S-1} \text{Var}_{P^{(t_0)}} \left(Z_{t_0}^{(s)} \right) + \frac{\mu_{t_0}}{(1 + 2\zeta S) \alpha} + \frac{\text{Var}_{P^{(t_0)}} \left((D_2^\circ Z)^{(t_0+S-2)} \right) + \text{Var}_{P^{(t_0)}} \left((D_2^\circ Z)^{(t_0+S-1)} \right)}{(1 + 2\zeta S)^2 \alpha}, \tag{85}$$

where (83) follows from (79), (84) follows from (82) and (78), and (85) follows from (81). Summing the above for $1 \leq t_0 \leq T - S + 1$, we obtain, for some constant $C > 0$,

$$\begin{aligned} & (S-1) \cdot \sum_{t=1}^{T-1} \text{Var}_{P^{(t)}} \left((D_1 Z)^{(t)} \right) \\ & \leq \sum_{t_0=1}^{T-S+1} \sum_{s=0}^{S-2} \text{Var}_{P^{(t_0+s)}} \left((D_1 Z)^{(t_0+s)} \right) + 8(S-1)^2 M^2 \end{aligned} \quad (86)$$

$$\leq \sum_{t_0=1}^{T-S+1} (1 + 2\zeta S) \sum_{s=0}^{S-2} \text{Var}_{P^{(t_0)}} \left((D_1 Z)^{(t_0+s)} \right) + 8(S-1)^2 M^2 \quad (87)$$

$$\begin{aligned} & \leq (1 + 2\zeta S)^3 \alpha \sum_{t_0=1}^{T-S+1} \sum_{s=0}^{S-1} \text{Var}_{P^{(t_0)}} \left(Z_{t_0}^{(s)} \right) + \sum_{t_0=1}^{T-S+1} \frac{\mu_{t_0}}{\alpha} + 8(S-1)^2 M^2 \\ & \quad + \sum_{t_0=1}^{T-S+1} \frac{\text{Var}_{P^{(t_0)}} \left((D_2^\circ Z)^{(t_0+S-2)} \right) + \text{Var}_{P^{(t_0)}} \left((D_2^\circ Z)^{(t_0+S-1)} \right)}{(1 + 2\zeta S)\alpha} \end{aligned} \quad (88)$$

$$\begin{aligned} & \leq (1 + 2\zeta S)^4 \alpha \sum_{t_0=1}^{T-S+1} \sum_{s=0}^{S-1} \text{Var}_{P^{(t_0+s)}} \left(Z^{(t_0+s)} \right) + \sum_{t_0=1}^{T-S+1} \frac{\mu_{t_0}}{\alpha} + 8(S-1)^2 M^2 \\ & \quad + \frac{4}{(1 + 2\zeta S)\alpha} \sum_{t_0=1}^{T-S+1} \left[\text{Var}_{P^{(t_0)}} \left(Z^{(t_0+S-2)} \right) + 3 \text{Var}_{P^{(t_0)}} \left(Z^{(t_0+S-1)} \right) \right. \\ & \quad \left. + 3 \text{Var}_{P^{(t_0)}} \left(Z^{(t_0)} \right) + \text{Var}_{P^{(t_0)}} \left(Z^{(t_0+1)} \right) \right] \end{aligned} \quad (89)$$

$$\begin{aligned} & \leq (1 + 2\zeta S)^4 \alpha S \sum_{t=1}^T \text{Var}_{P^{(t)}} \left(Z^{(t)} \right) + \sum_{t_0=1}^{T-S+1} \frac{\mu_{t_0}}{\alpha} + 8(S-1)^2 M^2 \\ & \quad + \frac{32}{(1 + 2\zeta S)\alpha} \sum_{t=1}^T \text{Var}_{P^{(t)}} \left(Z^{(t)} \right) \end{aligned} \quad (90)$$

$$\leq (1 + 2\zeta S)^4 \alpha S \cdot \left(1 + \frac{32}{\alpha^2 S} \right) \sum_{t=1}^T \text{Var}_{P^{(t)}} \left(Z^{(t)} \right) + \sum_{t_0=1}^{T-S+1} \frac{\mu_{t_0}}{\alpha} + 8(S-1)^2 M^2 \quad (91)$$

$$\leq (1 + 2\zeta S)^4 \alpha S \cdot \left(1 + \frac{32}{\alpha^2 S} \right) \sum_{t=1}^T \text{Var}_{P^{(t)}} \left(Z^{(t)} \right) + \frac{(S-2)\mu}{\alpha} + 10(S-1)^2 M^2, \quad (92)$$

where:

- (86) follows since $\|Z^{(t)}\|_\infty \leq M$ and thus $\|(D_1 Z)^{(t)}\|_\infty \leq 2M$ for all t ;
- (87) follows from Lemma C.1 and the fact that for $0 \leq s \leq S-2$, $\max \left\{ \left\| \frac{P^{(t_0+s)}}{P^{(t_0)}} \right\|_\infty, \left\| \frac{P^{(t_0)}}{P^{(t_0+s)}} \right\|_\infty \right\} \leq 1 + 2\zeta S$ as established above as a consequence of the fact that the distributions $P^{(t)}$ are ζ -consecutively close.
- (88) follows from (85);
- The first term in (89) is bounded using Lemma C.1 and the fact that the distributions $P^{(t)}$ are ζ -consecutively close, and the final term in (89) is bounded using the fact that for any vectors $Z_1, \dots, Z_k \in \mathbb{R}^n$ and any $P \in \Delta^n$, we have $\text{Var}_P(Z_1 + \dots + Z_k) \leq k \cdot (\text{Var}_P(Z_1) + \dots + \text{Var}_P(Z_k))$;
- (90) and (91) by rearranging terms;
- (92) follows from (76).

Now choose $S = \lceil \frac{128}{\alpha^3} \rceil$, so that $\frac{32}{\alpha^2 S} \leq \frac{\alpha}{4}$. Therefore, as long as $2\zeta S \leq \frac{\alpha}{32}$, we have, since $\alpha \leq 1/2$, that

$$(1 + 2\zeta S)^4 \alpha \cdot \frac{S}{S-1} \cdot \left(1 + \frac{32}{\alpha^2 S}\right) \leq \alpha \cdot (1 + \alpha/4)^3 \leq \alpha \cdot (1 + \alpha).$$

Then it follows from (92) that

$$\sum_{t=1}^{T-1} \text{Var}_{P^{(t)}} \left((D_1 Z)^{(t)} \right) \leq \alpha(1 + \alpha) \cdot \sum_{t=1}^T \text{Var}_{P^{(t)}} \left(Z^{(t)} \right) + \frac{\mu}{\alpha} + 10SM^2, \quad (93)$$

Using that $S \leq \frac{129}{\alpha^3}$, the inequality $2\zeta S \leq \alpha/32$ can be satisfied by ensuring that $\zeta \leq \frac{\alpha^4}{8256} = \frac{\alpha^4}{129 \cdot 42} \leq \frac{\alpha}{64S}$. Note that our choice of S ensures that $\zeta S \leq 1/2$, as was assumed earlier. Moreover, we have $10SM^2 \leq \frac{1290M^2}{\alpha^3}$. Thus, (93) gives the desired result. \square

C.3 Completing the proof of Theorem 3.1

Using the lemma developed in the previous sections we now can complete the proof of Theorem 3.1.

We begin by proving Lemma 4.2. The lemma is restated formally below.

Lemma 4.2 (Detailed). *There are constants $C, C' > 1$ so that the following holds. Suppose a time horizon $T \geq 4$ is given, we set $H := \lceil \log T \rceil$, and all players play according to Optimistic Hedge with step size η satisfying $1/T \leq \eta \leq \frac{1}{C \cdot m H^4}$. Then for any $i \in [m]$, the losses $\ell_i^{(1)}, \dots, \ell_i^{(T)} \in [0, 1]^{n_i}$ for player i satisfy:*

$$\sum_{t=1}^T \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t)} - \ell_i^{(t-1)} \right) \leq \frac{1}{2} \cdot \sum_{t=1}^T \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t-1)} \right) + C' H^5. \quad (94)$$

Proof. Since $T \geq 4$, we have that $H \geq 2$. Set $C_0 = 1290$ (note that C_0 is the constant appearing in the inequality (73) of the statement of Lemma 4.7 in Section C.2), and $C_1 = 8256$ (note that C_1 is the constant appearing in item 1 of Lemma 4.7 in Section C.2). Set $\alpha = 1/(4H)$, and $\alpha_0 = \frac{\sqrt{\alpha/8}}{H^3} < \frac{1}{H+3}$. Finally note that our assumption on η implies that, as long as the constant C satisfies $C \geq 4^4 \cdot 7 \cdot C_1 = 14794752$,

$$\eta \leq \min \left\{ \frac{\alpha^4}{7C_1}, \frac{\alpha_0}{36e^5 m} \right\}. \quad (95)$$

By Lemma 4.4 with the parameter α in the lemma set to α_0 , for each $i \in [m]$, $0 \leq h \leq H$ and $1 \leq t \leq T - h$, it holds that $\left\| (D_h \ell_i)^{(t)} \right\|_\infty \leq H \cdot (\alpha_0 H^3)^h$. Also set $\mu = H \cdot (\alpha_0 H^3)^H$. Therefore, we have, for each $i \in [m]$,

$$\sum_{t=1}^{T-H} \text{Var}_{x_i^{(t)}} \left((D_H \ell_i)^{(t)} \right) \leq \alpha \cdot \sum_{t=1}^{T-H+1} \text{Var}_{x_i^{(t)}} \left((D_{H-1} \ell_i)^{(t)} \right) + \mu^2 T. \quad (96)$$

We will now prove, via reverse induction on h , that for all $i \in [m]$ and h satisfying $H - 1 \geq h \geq 0$,

$$\sum_{t=1}^{T-h-1} \text{Var}_{x_i^{(t)}} \left((D_{h+1} \ell_i)^{(t)} \right) \leq \alpha \cdot (1+2\alpha)^{H-h-1} \cdot \sum_{t=1}^{T-h} \text{Var}_{x_i^{(t)}} \left((D_h \ell_i)^{(t)} \right) + \frac{2C_0 \cdot H^2 \cdot (2\alpha_0 H^3)^{2h}}{\alpha^3}. \quad (97)$$

The base case $h = H - 1$ is verified by (96) and the fact that $2^{2(H-1)} \geq 2^H \geq T$. Now suppose that (97) holds for some h satisfying $H - 1 \geq h \geq 1$. We will now apply Lemma 4.7, with $P^{(t)} = x_i^{(t)}$ and $Z^{(t)} = (D_{h-1} \ell_i)^{(t)}$ for $1 \leq t \leq T - h + 1$, as well as $M = H \cdot (2\alpha_0 H^3)^{h-1}$, $\zeta = 7\eta$, $\mu = \frac{2C_0 \cdot H^2 \cdot (2\alpha_0 H^3)^{2h}}{\alpha^3}$, and the parameter α of Lemma 4.7 set to $\alpha \cdot (1 + 2\alpha)^{H-h-1}$. We first verify that precondition 1 holds. By the definition of the Optimistic Hedge updates, the sequence $x_i^{(1)}, \dots, x_i^{(T)}$ is $\exp(6\eta)$ -consecutively close, and thus $(1 + 7\eta)$ -consecutively close (since $\exp(6\eta) \leq 1 + 7\eta$ for η satisfying (95)). Our choice of η, α ensures that $1/(2(T - h + 1)) \leq 7\eta \leq \alpha^4/C_1$. This verifies

that condition 1 of Lemma 4.7 holds. Note that condition 2 of the lemma holds by (97) and our choice of μ . Therefore, by Lemma 4.7 and the fact that $1 + \alpha \cdot (1 + 2\alpha)^H \leq 1 + 2\alpha$ for our choice of $\alpha = 1/(4H)$, it follows that

$$\begin{aligned}
\sum_{t=1}^{T-h} \text{Var}_{x_i^{(t)}} \left((D_h \ell_i)^{(t)} \right) &\leq \alpha \cdot (1 + 2\alpha)^{H-h} \cdot \sum_{t=1}^{T-h+1} \text{Var}_{x_i^{(t)}} \left((D_{h-1} \ell_i)^{(t)} \right) + \frac{2C_0 \cdot H^2 \cdot (2\alpha_0 H^3)^{2h}}{\alpha^4} \\
&\quad + \frac{C_0 \cdot H^2 \cdot (2\alpha_0 H^3)^{2(h-1)}}{\alpha^3} \\
&\leq \alpha \cdot (1 + 2\alpha)^{H-h} \cdot \sum_{t=1}^{T-h+1} \text{Var}_{x_i^{(t)}} \left((D_{h-1} \ell_i)^{(t)} \right) \\
&\quad + \frac{C_0 \cdot H^2 \cdot (2\alpha_0 H^3)^{2(h-1)}}{\alpha^3} \cdot \left(1 + \frac{2(2\alpha_0 H^3)^2}{\alpha} \right) \\
&\leq \alpha \cdot (1 + 2\alpha)^{H-h} \cdot \sum_{t=1}^{T-h+1} \text{Var}_{x_i^{(t)}} \left((D_{h-1} \ell_i)^{(t)} \right) + \frac{2C_0 \cdot H^2 \cdot (2\alpha_0 H^3)^{2(h-1)}}{\alpha^3},
\end{aligned}$$

where the final equality follows since α_0 is chosen so that $2(2\alpha_0 H^3)^2 = \alpha$. This completes the proof of the inductive step. Thus (97) holds for $h = 0$. Using again that the sequence $x_i^{(t)}$ is $(1 + 7\eta)$ -exponentially close, we see that

$$\begin{aligned}
&\sum_{t=1}^T \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t)} - \ell_i^{(t-1)} \right) \\
&\leq 1 + \sum_{t=2}^T \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t)} - \ell_i^{(t-1)} \right) \\
&\leq 1 + (1 + 7\eta) \sum_{t=1}^{T-1} \text{Var}_{x_i^{(t)}} \left((D_1 \ell_i)^{(t)} \right) \tag{98}
\end{aligned}$$

$$\leq 1 + (1 + 7\eta) \cdot \left(\alpha(1 + 2\alpha)^{H-1} \sum_{t=1}^T \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t)} \right) + \frac{2C_0 H^2}{\alpha^3} \right) \tag{99}$$

$$\leq 2 + (1 + 7\eta) \cdot \left(\alpha(1 + 2\alpha)^{H-1} (1 + 7\eta) \sum_{t=2}^T \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t-1)} \right) + \frac{2C_0 H^2}{\alpha^3} \right) \tag{100}$$

$$\leq \alpha(1 + 2\alpha)^H \sum_{t=1}^T \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t-1)} \right) + \frac{2(1 + 7\eta)C_0 H^2}{\alpha^3} + 2$$

$$\leq 2\alpha \sum_{t=1}^T \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t-1)} \right) + \frac{2(1 + 7\eta)C_0 H^2}{\alpha^3} + 2, \tag{101}$$

where (98) and (100) follow from Lemma C.1, and (99) uses (97) for $h = 0$. Now, (101) verifies the statement of the lemma. In summary, it suffices to take $C = 14794752$ and $C' = 2 + 2(1 + 7/8256)C_0 \cdot 4^3 = 165262$. \square

We are finally ready to prove Theorem 3.1. For convenience the theorem is restated below.

Theorem 3.1 (Restated). *There are constants $C, C' > 1$ so that the following holds. Suppose a time horizon $T \in \mathbb{N}$ is given. Suppose all players play according to Optimistic Hedge with any positive step size $\eta \leq \frac{1}{C \cdot m \log^4 T}$. Then for any $i \in [m]$, the regret of player i satisfies*

$$\text{Reg}_{i,T} \leq \frac{\log n_i}{\eta} + C' \cdot \log T. \tag{102}$$

In particular, if the players' step size is chosen as $\eta = \frac{1}{C \cdot m \log^4 T}$, then the regret of player i satisfies

$$\text{Reg}_{i,T} \leq O(m \cdot \log n_i \cdot \log^4 T). \tag{103}$$

Proof. The conclusion of the theorem is immediate if $T < 4$, so we may assume from here on that $T \geq 4$. Moreover, the conclusion of (102) is immediate if $\eta \leq 1/T$ (as $\text{Reg}_{i,T} \leq T$ necessarily), so we may also assume that $\eta \geq 1/T$. Let C'' be the constant C of Lemma 4.1, let B be the constant called C in Lemma 4.2 and B' be the constant called C' in Lemma 4.2. As long as the constant C of Theorem 3.1 is chosen so that $C \geq B$ and $\eta \leq \frac{1}{C \cdot m \log^4 T}$ implies that $C''\eta \leq 1/6$, we have the following:

$$\text{Reg}_{i,T} \leq \frac{\log n_i}{\eta} + \sum_{t=1}^T \left(\frac{\eta}{2} + C\eta^2 \right) \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t)} - \ell_i^{(t-1)} \right) - \sum_{t=1}^T \frac{(1 - C\eta)\eta}{2} \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t-1)} \right) \quad (104)$$

$$\begin{aligned} &\leq \frac{\log n_i}{\eta} + \frac{2\eta}{3} \sum_{t=1}^T \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t)} - \ell_i^{(t-1)} \right) - \frac{\eta}{3} \sum_{t=1}^T \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t-1)} \right) \\ &\leq \frac{\log n_i}{\eta} + \frac{2\eta}{3} \cdot (B' \cdot (2 \log T)^5) \end{aligned} \quad (105)$$

$$\leq \frac{\log n_i}{\eta} + 32B' \cdot \log T, \quad (106)$$

where (104) follows from Lemma 4.1, (105) follows from Lemma 4.2, and (106) follows from the upper bound $\eta \leq \frac{1}{C \cdot m \log^4 T}$. We have thus established (102). The upper bound (103) follows immediately. \square

D Adversarial regret bounds

In this section we discuss how Optimistic Hedge can be modified to achieve an algorithm that obtains the fast rates of Theorem 3.1 when played by all players, and which still obtains the optimal rate of $O(\sqrt{T})$ in the adversarial setting. Such guarantees are common in the literature [DDK11, RS13b, SALS15, KHSC18, HAM21]. The guarantees of this modification of Optimistic Hedge are stated in the following corollary (of Lemmas 4.1 and 4.2):

Corollary D.1. *There is an algorithm \mathcal{A} which, if played by all m players in a game, achieves the regret bound of $\text{Reg}_{i,T} \leq O(m \cdot \log n_i \cdot \log^4 T)$ for each player i ; moreover, when player i is faced with an adversarial sequence of losses, the algorithm \mathcal{A} 's regret bound is $\text{Reg}_{i,T} \leq O(m \log n_i \cdot \log^4 T + \sqrt{T} \log n_i)$.*

Proof. Let C be the constant called C in Theorem 3.1 and C' be the constant called C' in Lemma 4.2. The algorithm \mathcal{A} of Corollary D.1 is obtained as follows:

1. Initially run Optimistic Hedge, with the step-size $\eta = \frac{1}{C \cdot m \log^4 T}$.
2. If, for some $T_0 \geq 4$, (94) first fails to hold at time T_0 , i.e.,

$$\sum_{t=1}^{T_0} \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t)} - \ell_i^{(t-1)} \right) > \frac{1}{2} \cdot \sum_{t=1}^{T_0} \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t-1)} \right) + C' \lceil \log T \rceil^5, \quad (107)$$

then set $\eta' = \sqrt{\frac{\log n_i}{T}}$ and continue on running Optimistic Hedge with step size η' .

If there is no $T_0 \geq 4$ so that (107) holds (and by Lemma 4.2, this will be the case when \mathcal{A} is played by all m players in a game), then the proof of Theorem 3.1 shows that the regret of each player i is bounded as $\text{Reg}_{i,T} \leq O(m \log n_i \cdot \log^4 T)$. Otherwise, since T_0 is defined as the smallest integer at least 4 so that (107) holds, we have

$$\sum_{t=1}^{T_0} \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t)} - \ell_i^{(t-1)} \right) \leq \frac{1}{2} \cdot \sum_{t=1}^{T_0} \text{Var}_{x_i^{(t)}} \left(\ell_i^{(t-1)} \right) + C' \lceil \log T \rceil^5 + 4,$$

and thus, by Lemma 4.1, for any $x^* \in \Delta^{n_i}$,

$$\sum_{t=1}^{T_0} \langle \ell_i^{(t)}, x_i^{(t)} - x^* \rangle \leq \text{Reg}_{i, T_0} \leq O(m \log n_i \cdot \log^4 T_0). \quad (108)$$

Further, by the choice of step size $\eta' = \sqrt{\frac{\log n_i}{T}}$ for time steps $t > T_0$, we have, for any $x^* \in \Delta^{n_i}$,

$$\sum_{t=T_0+1}^T \langle \ell_i^{(t)}, x_i^{(t)} - x^* \rangle \leq \frac{\log n_i}{\eta'} + \eta' \sum_{t=T_0+1}^T \|\ell_i^{(t)} - \ell_i^{(t-1)}\|_\infty^2 \quad (109)$$

$$\leq \frac{\log n_i}{\eta'} + \eta' T \leq O(\sqrt{T \log n_i}), \quad (110)$$

where (109) uses [SALS15, Proposition 7]. Adding (108) and (110) completes the proof of the corollary. \square