# A HORIZON-INDEPENDENT REGRET BOUND FOR OPTIMISTIC HEDGE IN GENERAL-SUM GAMES

**Anonymous authors**
Paper under double-blind review

## ABSTRACT

Can simple learning rules keep their regret bounded in self-play? Recent work achieves constant regret bounds through modified regularization and higher-order prediction (Liu et al., 2026; Abbadi et al., 2026). Yet for Optimistic Hedge, arguably the most canonical method in games, the best known individual regret bound remains logarithmic (Ha, 2026). In this work, we prove that plain Optimistic Hedge with a constant step size can attain $O_{n,d}(1)$ individual regret in general-sum games with $n$ players and $d = (d_1, \ldots, d_n)$ actions, under expected loss-vector feedback. As a corollary, its time-averaged play enjoys an $O_{n,d}(1/T)$ coarse correlated equilibrium (CCE) gap. Our analysis represents Optimistic Hedge as a real-analytic recurrence on a compact space, which yields an exact finite-order difference relation that eliminates horizon dependence. Our proof hinges on nonconstructive Noetherianity argument of Frisch (1967), so the $(n, d)$-dependence remains implicit.

## 1 INTRODUCTION

A central question in game theory (von Neumann & Morgenstern, 1944; Nash, 1950) is whether independent learning can lead to an equilibrium when opponents are adapting. Such interdependent learning dynamics underpin modern applications such as poker self-play (Brown & Sandholm, 2019), adversarial training (Goodfellow et al., 2014), and game-theoretic approaches to language-model alignment (Munos et al., 2024; Rosset et al., 2024; Wu et al., 2024; Gölz et al., 2025; Wu et al., 2026).

No-regret learning (Cesa-Bianchi & Lugosi, 2006) offers a decentralized approach: players seek low regret against the best fixed action in hindsight (Hannan, 1957). Indeed, sublinear external regret drives time-averaged play toward an equilibrium (Moulin & Vial, 1978; Hart & Mas-Colell, 2001).

While classical online learning guarantees $O(\sqrt{T})$ regret against arbitrary losses (Freund & Schapire, 1997), self-play adds an additional structure: a common update rule shared by all players. Optimistic methods exploit this structure and use recent losses to predict the next one and improve regret bounds (Rakhlin & Sridharan, 2013a; Syrgkanis et al., 2015). Optimistic Hedge, the most canonical optimistic method for games, has recently been shown to enjoy polylogarithmic (Daskalakis et al., 2021) and logarithmic (Ha, 2026) regret bounds. On the other hand, recent alternatives have shown that modified regularization and higher-order prediction methods can produce horizon-independent regret bounds (Liu et al., 2026; Abbadi et al., 2026). In this paper, we ask the following natural question:

*Can Optimistic Hedge attain horizon-independent individual regret in general-sum games?*

**Contributions.** Our main result gives an affirmative answer. Our contributions are two-fold.

(i) **An analytic finite-order difference relation.** Based on the high-order smoothness framework (Daskalakis et al., 2021) and the weighted-gap reduction (Ha, 2026), we represent Optimistic Hedge as an analytic recurrence on a compact domain. A Noetherianity argument then yields an exact finite-order difference relation controlling the first-order weighted-gap differences.

(ii) **A horizon-independent regret bound.** Based on the finite-order difference relation and a simple finite-difference interpolation argument, we obtain $O_{n,d}(1)$ individual regret and an $O_{n,d}(1/T)$ CCE gap. The dimension dependence remains nonconstructive.

Section 2 reviews prior work; Sections 3–4 introduce the setting and main result. Section 5 explains the proof strategy, with detailed proofs in the appendices.



Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

## 2 RELATED WORK

**No-regret learning and equilibrium.** The link between no-regret learning and equilibrium computation dates back to Blackwell (1956) and Hannan (1957). External regret compares cumulative loss with the best fixed action in hindsight. When all players achieve sublinear regret, averaging yields approximate Nash equilibria in two-player zero-sum games (Freund & Schapire, 1999) and CCE in general-sum games (Moulin & Vial, 1978; Hart & Mas-Colell, 2001). Controlling recommendation-dependent deviations via internal or swap regret yields correlated equilibrium (Hart & Mas-Colell, 2000; Blum & Mansour, 2007).

**Individual regret in self-play.** Improving upon standard no-regret learning algorithms, such as follow-the-regularized-leader (FTRL), optimistic algorithms exploit predictable loss sequences (Chiang et al., 2012; Rakhlin & Sridharan, 2013a). Notably, the seminal work of Syrgkanis et al. (2015) showed that optimistic follow-the-regularized-leader (OFTRL) and optimistic mirror descent (OMD) can exploit the self-play-induced predictability in games to achieve $O(\sqrt{n}\log d_{\max}T^{1/4})$ individual regret, where $d_{\max} := \max_i d_i$ is the maximum number of actions across players. For Optimistic Hedge, Chen & Peng (2020) obtained $O(\log^{5/6} d_{\max}T^{1/6})$ in two-player games, and Daskalakis et al. (2021) established $O(n \log d_i \log^4 T)$ in $n$-player games based on a high-order smoothness analysis.

A recent line of work improves regret bounds by modifying regularization, step sizes, or prediction methods. Log-regularized lifted OFTRL (LRL-OFTRL) (Farina et al., 2022) enlarges the regularized strategy space and achieves $O(nd_{\max}\log T)$ in finite general-sum games, and extends to general convex strategy sets. Cautious Optimistic MWU (COMWU) (Soleymani et al., 2025a;b) achieves $O(n\log^2 d_{\max}\log T)$ through adaptive step sizes. ECHO-OFTRL (Liu et al., 2026) and HOOD (Abbadi et al., 2026) combine higher-order prediction with lifted regularization to obtain $O(n^{21}\log^4 d_{\max})$ and $O(n^3\log^2 d_{\max})$ regret bounds, respectively. Table 1 compares these individual regret bounds.

**Other regret notions and update rules.** A parallel line studies swap regret to obtain guarantees for correlated equilibrium (CE) (Aumann, 1974). Anagnostides et al. (2022a) extend the high-order smoothness framework of Daskalakis et al. (2021) to obtain polylogarithmic internal and swap regret. Log-barrier regularization (Anagnostides et al., 2022b) and hybrid regularization (Tsuchiya, 2026) yield logarithmic and sublogarithmic swap regret, respectively. Clairvoyant Multiplicative Weights Update (CMWU) (Piliouras et al., 2022) obtains horizon-independent regret via implicit updates; however, its uncoupled implementation guarantees this bound only on a selected subsequence. For two-player general-sum games, Tao & Zheng (2026) achieve $O(\log d_{\max})$ alternating regret and $O(\log d_{\max}/T)$ CCE error using sequential updates, unlike our simultaneous external-regret setting.

**Closest prior work.** The closest works are Daskalakis et al. (2021) and Ha (2026). We build on the high-order smoothness framework of Daskalakis et al. (2021) and the recurrence-based weighted-gap analysis of Ha (2026). In particular, Ha (2026) represents Optimistic Hedge as a recurrent map acting on a state space equipped with a fixed Euclidean norm. Crucially, the high-order estimates of Ha (2026) leave a remainder that grows with the horizon at a fixed step size. In this work, we instead obtain an exact relation among difference orders and eliminate the horizon dependence in our final regret bound.

## 3 PRELIMINARIES

We study full-information repeated play in finite general-sum games within the classical no-regret learning framework (Hannan, 1957; Cesa-Bianchi & Lugosi, 2006).

**Notation.** For an integer $q \geq 1$, write $[q] := \{1, \dots, q\}$ and define a probability simplex $\Delta_q := \{x \in \mathbb{R}^q_+ : \sum_a x_a = 1\}$. Fix $n \geq 1$ players with action counts $d_i \geq 1$, action sets $A_i = [d_i]$, and mixed strategies $x_i \in \Delta_{d_i}$. Set action dimensions $d := (d_1, \dots, d_n)$ and $d_{\max} := \max_i d_i$. The notation $O_{n,d}(\cdot)$ hides dependence on $n, d$, but not on the horizon; all logarithms are natural.

**Loss functions.** We denote a loss table by $\ell : \prod_{j=1}^n A_j \to [0,1]^n$ and define player $i$'s expected loss $\ell_i : \prod_{j=1}^n \Delta_{d_j} \to [0,1]$ as $\ell_i(x) := \mathbb{E}_{a\sim\bigotimes_j x_j}[\ell(a)]_i$. For a strategy profile $x = (x_1, \dots, x_n)$,

Table 1: Individual external regret in full-information general-sum normal-form games; $d_{\max} := \max_i d_i$. The Chen–Peng bound assumes two players; LRL-OFTRL also applies to convex games. "Single loop" means closed-form updates without inner optimization or root finding, excluding feedback cost. The OFTRL/OMD check refers to its entropy-regularized Optimistic Hedge instantiation. The notation $O_{n,d}(1)$ hides dependence on the numbers of players and actions, but not on $T$.

| Reference | Algorithm | Single loop | Regret bound |
|---|---|---|---|
| Syrgkanis et al. (2015) | OFTRL/OMD | ✓ | $O(\sqrt{n}\log d_{\max} T^{1/4})$ |
| Chen & Peng (2020) | Optimistic Hedge | ✓ | $O(\log^{5/6} d_{\max} T^{1/6})$ |
| Daskalakis et al. (2021) | Optimistic Hedge | ✓ | $O(n \log d_i \log^4 T)$ |
| Farina et al. (2022) | LRL-OFTRL | ✗ | $O(n d_{\max} \log T)$ |
| Soleymani et al. (2025a;b) | COMWU | ✗ | $O(n \log^2 d_{\max} \log T)$ |
| Liu et al. (2026) | ECHO-OFTRL | ✗ | $O(n^{21} \log^4 d_{\max})$ |
| Abbadi et al. (2026) | HOOD | ✗ | $O(n^3 \log^2 d_{\max})$ |
| Ha (2026) | Optimistic Hedge | ✓ | $O(\sqrt{n}\log d_i \log T)$ |
| This work | Optimistic Hedge | ✓ | $O_{n,d}(1)$ |

write $x_{-i} := (x_j)_{j\neq i}$. We define the expected loss vector $\ell_i(x_{-i}) \in [0,1]^{d_i}$ by

$$[\ell_i(x_{-i})]_{a_i} := \sum\nolimits_{a_{-i}\in\prod_{j\neq i} A_j} [\ell(a_i, a_{-i})]_i \prod\nolimits_{j\neq i} x_{j,a_j}. \tag{1}$$

Hence, the identity $\ell_i(x) = \langle x_i, \ell_i(x_{-i})\rangle$ holds. We also write $\ell(x) := (\ell_i(x_{-i}))_{i=1}^n$ for the stacked expected-loss profile; the argument distinguishes it from the loss table. At round $t$, write $\ell^t := \ell(x^t)$; player $i$ incurs $\langle x_i^t, \ell_i^t\rangle$ and observes the full vector $\ell_i^t = \ell_i(x_{-i}^t)$.

Computing a Nash equilibrium is PPAD-complete even for two-player general-sum games (Daskalakis et al., 2009; Chen et al., 2009). Coarse correlated equilibrium (CCE) (Moulin & Vial, 1978) is a tractable relaxation: every Nash equilibrium induces a CCE, but a CCE allows correlated actions and tests only unconditional deviations. A CCE can be computed by linear programming with a time complexity polynomial in $n\prod_i d_i$ and the payoff bit length (Papadimitriou & Roughgarden, 2008).

**Definition 3.1** ($\varepsilon$-approximate coarse correlated equilibrium)**.** For a joint distribution $\mu$ on $\prod_j A_j$, define its CCE gap by

$$\mathrm{Gap}_{\mathrm{CCE}}(\mu) := \max_{i\in[n]} \max_{a_i'\in A_i} \mathbb{E}_{a\sim\mu}[[\ell(a)]_i - [\ell(a_i', a_{-i})]_i]\,. \tag{2}$$

For $\varepsilon \geq 0$, $\mu$ is an $\varepsilon$-approximate CCE if $\mathrm{Gap}_{\mathrm{CCE}}(\mu) \leq \varepsilon$.

No-regret learning approaches CCE by controlling each player's external regret (Hart & Mas-Colell, 2001; Cesa-Bianchi & Lugosi, 2006), which is defined as the difference between its cumulative loss and that of its best fixed action in hindsight.

**Definition 3.2** (External regret)**.** At horizon $T \geq 1$, player $i$'s individual external regret is

$$\mathrm{Reg}_i(T) := \sum_{t=1}^T \langle x_i^t, \ell_i^t\rangle - \min_{a_i'\in A_i} \sum_{t=1}^T \ell_{i,a_i'}^t. \tag{3}$$

Let $\widehat{\mu}_T := T^{-1}\sum_{t=1}^T \bigotimes_{j=1}^n x_j^t$ be the time-averaged joint distribution, the uniform mixture of the product distributions played at each round. Under $\widehat{\mu}_T$, each fixed-action deviation gain is its average gain over the played rounds. As a result, the following standard result holds:

$$\mathrm{Gap}_{\mathrm{CCE}}(\widehat{\mu}_T) = \frac{1}{T}\max_{i\in[n]} \mathrm{Reg}_i(T). \tag{4}$$

Therefore, individual-regret bounds directly control the CCE approximation error.

**Optimistic Hedge.** Hedge is an entropy-regularized follow-the-regularized-leader (FTRL) algorithm (Freund & Schapire, 1997). Optimistic Hedge uses $\ell_i^t$ to predict $\ell_i^{t+1}$ (Rakhlin & Sridharan,

2013a). All players use a constant step size $\eta > 0$, with $\ell_i^0 = 0$ and $x_i^1 = \mathbf{1}_{d_i}/d_i$, where $\mathbf{1}_{d_i}$ is the all-ones vector. Eliminating cumulative losses from its closed-form update gives

$$x_{i,a}^{t+1} = \frac{x_{i,a}^t \exp\{-\eta(2\ell_{i,a}^t - \ell_{i,a}^{t-1})\}}{\sum_{b\in A_i} x_{i,b}^t \exp\{-\eta(2\ell_{i,b}^t - \ell_{i,b}^{t-1})\}}, \qquad t \geq 1, \quad a \in A_i. \tag{5}$$

For the update in Equation (5), Chen & Peng (2020) obtain $O(\log^{5/6} d_{\max} T^{1/6})$ individual regret in two-player games. For $n$ players, Daskalakis et al. (2021) establish $O(n \log d_i \log^4 T)$, and Ha (2026) sharpens this to $O(\sqrt{n} \log d_i \log T)$. We show that plain Optimistic Hedge can attain horizon-independent $O_{n,d}(1)$ regret for a sufficiently small constant step size.

## 4 Main Result

The following theorem shows that Optimistic Hedge admits a horizon-independent individual-regret bound at sufficiently small constant step sizes.

**Theorem 4.1** (Horizon-independent individual regret)**.** *Fix $n$ players with action counts $d_1, \ldots, d_n$. There exist $\eta_\star > 0$ and $C_\star < \infty$, depending only on $(n, d_1, \ldots, d_n)$, such that, for every game in the setup of Section 3 and every common constant step size $0 < \eta \leq \eta_\star$, the Optimistic Hedge iterates satisfy, for every $i \in [n]$ and $T \geq 1$,*

$$\mathrm{Reg}_i(T) \leq \frac{\log d_i}{\eta} + C_\star \eta. \tag{6}$$

The logarithmic guarantee for Optimistic Hedge uses a horizon-tuned constant step size (Ha, 2026). Specifically, the parameterized bound of Ha (2026) leaves a remainder proportional to $\eta T \exp\{-c/(\sqrt{n}\,\eta)\}$ (for a universal $c > 0$), which is exponentially small in the inverse step size but still grows with the horizon. In contrast, Theorem 4.1 yields a horizon-independent $O_{n,d}(1)$ individual-regret bound.

The following corollary translates this regret bound into a convergence guarantee for time-averaged play, using Equation (4).

**Corollary 4.2** (Coarse-correlated-equilibrium convergence)**.** *Under the conditions of Theorem 4.1, the time-averaged joint distribution $\widehat{\mu}_T$ satisfies, for every $T \geq 1$,*

$$\mathrm{Gap}_{\mathrm{CCE}}(\widehat{\mu}_T) \leq \frac{1}{T} \max_{i\in[n]} \left\{ \frac{\log d_i}{\eta} + C_\star \eta \right\}. \tag{7}$$

Corollary 4.2 improves the horizon dependence of the $O(\sqrt{n} \log d_{\max} \log T/T)$ CCE bound of Ha (2026) from $\log T/T$ to $1/T$. The improvement is in $T$, not in game dimensions: our $O_{n,d}(1/T)$ guarantee leaves $n, d$ unquantified. We leave explicit bounds on $\eta_\star, C_\star$ and a dimension-explicit horizon-independent individual regret bound to future work.

*Remark* 4.1 (Comparison with modified dynamics). ECHO-OFTRL achieves $O(n^{21} \log^4 d_{\max})$ individual external regret (Liu et al., 2026), and HOOD achieves $O(n^3 \log^2 d_{\max})$ (Abbadi et al., 2026). Hybrid-regularized OFTRL with the Blum–Mansour reduction gives $O(n d_{\max}^2 \sqrt{\log d_{\max} \log T})$ swap regret and a correlated-equilibrium guarantee (Tsuchiya, 2026). However, these modified dynamics involve inner optimization or root-finding subroutines. Our result instead establishes a horizon-independent regret bound for plain Optimistic Hedge, which admits a closed-form update.

## 5 Proof Overview

We follow the high-order smoothness framework of Daskalakis et al. (2021) and the weighted-gap reduction of Ha (2026). Our real-analytic state representation of Optimistic Hedge yields an exact finite-order relation through the Noetherianity theorem of Frisch (1967). Finite-difference interpolation then converts this relation into a horizon-independent regret bound.

### 5.1 Reduction to a First-Order Difference Estimate

For $x \in \Delta_q$ and $v \in \mathbb{R}^q$, write the variance of a vector as $\mathrm{Var}_x(v) := \sum_a x_a(v_a - \langle x, v\rangle)^2$. Ha (2026) gives the following small-step-size specialization of Lemma 4.1 in Daskalakis et al. (2021).

**Lemma 5.1** (Optimistic regret bound (Ha, 2026, Lemma 5.1))**.** *There is a universal $\eta_{\mathrm{rvu}} > 0$ such that, for every player, horizon, and $0 < \eta \le \eta_{\mathrm{rvu}}$,*

$$\mathrm{Reg}_i(T) \le \frac{\log d_i}{\eta} + \frac{2\eta}{3}\sum_{t=1}^{T}\mathrm{Var}_{x_i^t}(\ell_i^t - \ell_i^{t-1}) - \frac{\eta}{3}\sum_{t=1}^{T}\mathrm{Var}_{x_i^t}(\ell_i^{t-1}). \tag{8}$$

The first sum charges prediction error; the second supplies the negative variance term. Bounding the first by half the second, up to a horizon-independent remainder, controls regret. The time-varying quadratic form $\mathrm{Var}_{x_i^t}(\cdot)$ prevents a direct comparison in a fixed sequence norm; hence, we incorporate its probability weights into the vectors themselves. The pairwise identity

$$\mathrm{Var}_x(v) = \sum_{a<b} x_a x_b (v_a - v_b)^2 \tag{9}$$

motivates the weighted pairwise loss-gap vector $h_i^t \in \mathbb{R}^{\binom{d_i}{2}}$ of Ha (2026), defined for $a < b$ by

$$[h_i^t]_{ab} := \sqrt{x_{i,a}^t x_{i,b}^t}\big(\ell_{i,a}^{t-1} - \ell_{i,b}^{t-1}\big). \tag{10}$$

The following lemma from Ha (2026) controls the sum of prediction-error variances $\mathrm{Var}_{x_i^t}(\ell_i^t - \ell_i^{t-1})$ in terms of the weighted-gap differences $\Delta h_i^t := h_i^{t+1} - h_i^t$, accounting for time-varying weights.

**Lemma 5.2** (Weighted-gap reduction (Ha, 2026, Lemma 5.2))**.** *There are universal constants $C > 0$ and $\eta_{\mathrm{gap}} > 0$ such that, whenever $0 < \eta \le \eta_{\mathrm{gap}}$,*

$$\begin{aligned} \sum_{t=1}^{T}\mathrm{Var}_{x_i^t}(\ell_i^{t-1}) &= \sum_{t=1}^{T}\left\|h_i^t\right\|_2^2, \\ \sum_{t=1}^{T}\mathrm{Var}_{x_i^t}(\ell_i^t - \ell_i^{t-1}) &\le \frac{3}{2}\sum_{t=1}^{T}\left\|h_i^{t+1} - h_i^t\right\|_2^2 + C\eta^2\sum_{t=1}^{T}\left\|h_i^t\right\|_2^2 + C\eta^2. \end{aligned} \tag{11}$$

Lemmas 5.1 and 5.2 reduce the remaining challenge to bounding the sum of first-order differences of weighted gaps $\sum_{t=1}^{T}\|\Delta h_i^t\|_2^2$ by the sum of weighted gaps themselves $\sum_{t=1}^{T+1}\|h_i^t\|_2^2$, up to a horizon-independent remainder. The following key lemma supplies this bound.

**Lemma 5.3** (First-order difference estimate)**.** *Fix $n$ and the action counts $d$. There are constants $C, C' < \infty$ and $\eta_{\mathrm{temp}} > 0$, depending only on $n, d$, such that, for every game of these dimensions, every player $i \in [n]$, every horizon $T \ge 1$, and every common constant step size $0 < \eta \le \eta_{\mathrm{temp}}$,*

$$\sum_{t=1}^{T}\left\|h_i^{t+1} - h_i^t\right\|_2^2 \le C\eta^2\sum_{t=1}^{T+1}\left\|h_i^t\right\|_2^2 + C'. \tag{12}$$

Unlike the first-order difference estimate of Ha (2026), this bound is horizon-independent: $C'$ depends only on the game dimensions $(n, d)$, but not on the horizon $T$. Lemmas 5.1 to 5.3 yield Theorem 4.1, so our goal now reduces to proving Lemma 5.3.

### 5.2 Reduction to a High-Order Difference

For a sequence $y = (y^t)_{t\in\mathbb{Z}}$, define finite differences by $\Delta^0 y^t := y^t$, $\Delta y^t := y^{t+1} - y^t$, and $\Delta^{k+1} y^t := \Delta(\Delta^k y^t)$. For a sequence with a finite support, write $\|y\|_{\ell_2}^2 := \sum_t \|y^t\|^2$.

To prove Lemma 5.3, we use the following lemma to control the first-order difference $\|\Delta y\|_{\ell_2}$ in terms of the sequence norm $\|y\|_{\ell_2}$ and the order-$m$ difference $\|\Delta^m y\|_{\ell_2}$.

**Lemma 5.4** (Finite-difference interpolation)**.** *For every finitely supported sequence $y = (y^t)_{t\in\mathbb{Z}}$ in a finite-dimensional Hilbert space, the following holds for all integers $m \ge 1$ and $0 \le j \le m$:*

$$\left\|\Delta^j y\right\|_{\ell_2} \le \|y\|_{\ell_2}^{1-j/m}\left\|\Delta^m y\right\|_{\ell_2}^{j/m}. \tag{13}$$

Intuitively, Parseval's identity expresses the difference norms in a frequency domain, and Hölder's inequality interpolates between orders 0 and $m$. We defer the proof to Appendix B.

Lemma 5.4 suggests the following route. Suppose the weighted gaps satisfy the finite-order relation

$$\Delta^m h_i^{t+m} = \sum_{k=0}^{m-1} c_{k,t}\eta^{m-k}\Delta^k h_i^{t+k}, \qquad |c_{k,t}| \le A, \quad t \ge 1, \tag{14}$$

where $m, A$ depend only on $(n, d)$. At each horizon, we can construct a finitely supported cutoff $y$ of $(h_i^t)$ whose error in (14) is bounded by $O_{n,d}(\eta^{m-1/2})$ in $\ell_2$ norm (Lemma 5.9). Taking norms gives

$$\|\Delta^m y\|_{\ell_2} \le \tfrac{1}{2}\|\Delta^m y\|_{\ell_2} + O_{n,d}(\eta^m)\|y\|_{\ell_2} + O_{n,d}(\eta^{m-1/2}).$$

Here we use Lemma 5.4 and Young's inequality. Absorbing the first term on the right and applying Lemma 5.4 with $j = 1$ and Young's inequality once more yields

$$\|\Delta y\|_{\ell_2}^2 \le \|y\|_{\ell_2}^{2-2/m}\|\Delta^m y\|_{\ell_2}^{2/m} \le O_{n,d}(\eta^2)\|y\|_{\ell_2}^2 + O_{n,d}(\eta).$$

The cutoff construction $y$ in Lemma 5.9 transfers this estimate back to $h_i^t$, yielding Lemma 5.3, and hence, Theorem 4.1. Notice that the above derivation crucially hinges on the exact finite-order relation (14); we prove that precise finite-order relation in the following subsections.

### 5.3 State Recurrence and Finite-Order Difference Relation

To obtain the finite-order relation (14), we write Optimistic Hedge as a recurrence on a compact state space. Specifically, we encode the scaled differences $\eta^{-k}\Delta^k h_i^{t+k}$ by analytic maps. Then, Frisch's Noetherianity theorem (Frisch, 1967), which guarantees finite generation of analytic spans near compact semianalytic sets, gives the finite relation we need.

**Square-root state.** The factor $\sqrt{x_{i,a}x_{i,b}}$ in Equation (10) is not analytic in the probabilities at the boundary. Hence, we set $r_i(a) = \sqrt{x_i(a)}$, which turns the pair weights into the polynomial products $r_i(a)r_i(b)$ and places each $r_i$ on a compact nonnegative unit sphere. We store the loss profile as $\ell = (\ell_i)_{i=1}^n$ and use the state space

$$\mathcal{S} := \left\{(r, \ell) : r_i \in \mathbb{R}_+^{d_i},\ \|r_i\|_2 = 1, \ell_i \in [0,1]^{d_i} \text{ for every } i\right\}.$$

**State recurrence.** For a nonzero $r_i$, define the mixed strategy $x_i(r_i) \in \Delta_{d_i}$ by $[x_i(r_i)]_a := r_i(a)^2/\|r_i\|_2^2$. Recall that $\ell$ denotes the stacked expected-loss map from Section 3, and write $\ell(r) := \ell((x_i(r_i))_{i=1}^n)$. For $s = (r, \ell) \in \mathcal{S}$, define the strategy update $F_i : \mathcal{S} \times \mathbb{R} \to \mathbb{R}^{d_i}$ and the full state update $G : \mathcal{S} \times \mathbb{R} \to \mathcal{S}$ by

$$[F_i(r, \ell; \eta)]_a := \frac{\|r_i\|_2\ r_i(a)\exp\{-\eta(2[\ell(r)]_i(a) - [\ell_i]_a)/2\}}{\left(\sum_b r_i(b)^2\exp\{-\eta(2[\ell(r)]_i(b) - [\ell_i]_b)\}\right)^{1/2}}, \tag{15}$$

$$G(r, \ell; \eta) := \left(\left(F_i(r, \ell; \eta)\right)_{i=1}^n, \ell(r)\right). \tag{16}$$

The factor $\|r_i\|_2$ leaves the update unchanged on $\mathcal{S}$ and extends it to a neighborhood of the unit spheres. We keep $\eta$ fixed under composition: $G^0(s; \eta) = s$ and $G^{j+1}(s; \eta) = G(G^j(s; \eta); \eta)$. For Optimistic Hedge, setting $r_i^t(a) = \sqrt{x_i^t(a)}$ and $s^t = (r^t, \ell^{t-1})$ gives $s^{t+1} = G(s^t; \eta)$.

**Weighted gaps.** We map the square-root state back to the weighted gaps with $H_{i,0} : \mathcal{S} \times \mathbb{R} \to \mathbb{R}^{\binom{d_i}{2}}$:

$$[H_{i,0}(r, \ell; \eta)]_{ab} := r_i(a)r_i(b)\left([\ell_i]_a - [\ell_i]_b\right), \qquad a < b. \tag{17}$$

For $\eta \ne 0$ and $k \ge 0$, define the scaled-difference maps recursively by

$$H_{i,k+1}(s; \eta) := \frac{H_{i,k}(G^2(s; \eta); \eta) - H_{i,k}(G(s; \eta); \eta)}{\eta}. \tag{18}$$

Induction along $s^{t+1} = G(s^t; \eta)$ gives $\Delta^k h_i^{t+k} = \eta^k H_{i,k}(s^t; \eta)$ for $k \ge 0$ and $t \ge 1$.

Notice that a relation $H_{i,m} = \sum_{k<m} c_k H_{i,k}$ immediately yields the finite-order relation (14) after multiplication by $\eta^m$. A natural question is therefore when such a relation would exist among the maps $H_{i,k}$. The following consequence of the Noetherianity theorem of Frisch (1967) shows that real analyticity of these maps near a compact semianalytic set suffices; see also Lønsted (1973). Recall that a semianalytic set is locally specified by finitely many real-analytic equalities and inequalities.

**Theorem 5.5** (Finite relation among analytic maps)**.** *Let $K \subset \mathbb{R}^q$ be compact and semianalytic, and let $V$ be a finite-dimensional real vector space. For every sequence $\Phi_0, \Phi_1, \dots$ of $V$-valued maps, each real analytic on a neighborhood of $K$, there exist an integer $m \geq 1$, an open neighborhood $U$ of $K$, and scalar real-analytic functions $c_0, \dots, c_{m-1}$ on $U$ such that the following holds on $U$:*

$$\Phi_m = \sum_{k=0}^{m-1} c_k \Phi_k. \tag{19}$$

Locally, analytic maps behave like polynomials: after finitely many $\Phi_k$, new maps can be expressed using earlier ones with analytic coefficients. The number of earlier maps needed locally may vary across $K$. Compactness and semianalyticity rule out unbounded growth of this number; we defer the proof to Appendix C.4. We verify these hypotheses for $H_{i,k}$ and transfer the relation to $h_i^t$.

### 5.4 Proof of the Finite-Order Difference Relation

Recall that our goal is to derive the finite-order relation (14) for the weighted gaps $h_i^t$ in arbitrary games with dimensions $(n, d)$. To make the finite relation produced by Theorem 5.5 hold uniformly over such games, we include the entire loss table in the arguments of our maps. Let

$$\mathcal{G} := [0,1]^{n \prod_{j=1}^{n} d_j}, \qquad K_\star := \mathcal{S} \times \mathcal{G} \times \{0\},$$

where $\mathcal{G}$ is the set of loss tables. We choose the singleton step size space $\{0\}$ because we only need the relation for sufficiently small step sizes. With $g \in \mathcal{G}$ fixed, we suppress it throughout this subsection:

$$F_i(r, \ell; \eta) \equiv F_i(r, \ell; g, \eta), \qquad G(s; \eta) \equiv G(s; g, \eta), \qquad H_{i,k}(s; \eta) \equiv H_{i,k}(s; g, \eta).$$

We now verify the hypotheses of Theorem 5.5 for $\Phi_k = H_{i,k}$ and prove the finite-order relation (14).

**Step 1: analyticity of $G$ and $H_{i,k}$.** We first study the analyticity of $G$ and its behavior at $\eta = 0$.

**Lemma 5.6** (Analyticity of the state update)**.** *The state update $G(s; \eta)$ is jointly real analytic in the state, the loss-table entries, and $\eta$ on one common neighborhood of $K_\star$. The state space $\mathcal{S}$ is invariant. Moreover, $G(r, \ell; 0) = (r, \ell(r))$ and $G^2(s; 0) = G(s; 0)$ throughout this neighborhood.*

At $\eta = 0$, $G$ only replaces the stored losses by $\ell(r)$, so applying it twice changes nothing further. Thus the zero-step identity makes the numerator in Equation (18) vanish at $\eta = 0$. By Lemma 5.6, the numerator is analytic whenever $H_{i,k}$ is, so the quotient extends analytically through $\eta = 0$. Thus the recursion preserves analyticity from the polynomial base $H_{i,0}$, giving Lemma 5.7.

**Lemma 5.7** (Analyticity of scaled finite differences)**.** *For every fixed $k \geq 0$, the function $H_{i,k}$ in Equation* (18) *extends real analytically through $\eta = 0$, jointly in the state and loss-table entries, on a neighborhood of $K_\star$. The neighborhood may depend on $k$, but not on the state or game.*

**Step 2: finite generation and uniform coefficient bounds.** Now that we have established the analyticity of $H_{i,k}$, we can apply Theorem 5.5 to $H_{i,0}, H_{i,1}, \dots$. For some finite $m$, Theorem 5.5 gives analytic coefficient functions $c_0, \dots, c_{m-1}$ on a common neighborhood of $K_\star$ such that

$$H_{i,m} = \sum_{k=0}^{m-1} c_k H_{i,k}. \tag{20}$$

Evaluating (20) along the Optimistic Hedge trajectory gives the finite-order relation we have sought:

$$\Delta^m h_i^{t+m} = \eta^m H_{i,m}(s^t; \eta) = \sum_{k=0}^{m-1} c_k(s^t; \eta) \eta^{m-k} \Delta^k h_i^{t+k}.$$

Compactness of $K_\star$ puts $\mathcal{S} \times \mathcal{G} \times [0, \eta_i]$ inside the common neighborhood for some $\eta_i > 0$ depending on $(n, d)$. Continuity bounds $c_k$ and $H_{i,k}$ on this compact set; invariance of $\mathcal{S}$ makes these bounds valid at every round and produces the following lemma.

**Lemma 5.8** (Exact finite-order relation)**.** *Fix the number of players $n$, number of actions $d = (d_1, \dots, d_n)$, and a player $i$. There are an integer $m \geq 1$, constants $A, M < \infty$, and $\eta_i > 0$, uniform over all games of these dimensions, such that, for every such game, $0 < \eta \leq \eta_i$, and $t \geq 1$, there exist scalars $c_{0,t}, \dots, c_{m-1,t}$, possibly depending on the game, $\eta$, and $t$, satisfying*

$$\Delta^m h_i^{t+m} = \sum_{k=0}^{m-1} c_{k,t} \eta^{m-k} \Delta^k h_i^{t+k}, \qquad |c_{k,t}| \leq A, \tag{21}$$

$$\left\|\Delta^k h_i^{t+k}\right\|_2 \leq M\eta^k, \qquad 0 \leq k \leq m. \tag{22}$$

Lemma 5.8 gives the precise finite-order relation (14) that we needed in Section 5.2. In particular, it provides a uniform bound on the coefficients $c_{k,t}$ and does not include any error term that accumulates over the horizon $T$. The exact order $m$ and constants $A, M$ remain implicit due to the nonconstructive nature of Theorem 5.5. We defer the full proof to Appendix C.4.

### 5.5 From the Finite-Order Relation to the Regret Bound

With the exact finite-order relation from Lemma 5.8 in hand, one final step unlocks Lemma 5.4: a finitely supported cutoff of $h_i^t$ with boundary error independent of $T$. The next lemma supplies this cutoff construction to complete the argument in Section 5.2; we defer the proof to Appendix D.

**Lemma 5.9** (Finite-support cutoff)**.** *Fix an integer $m \geq 1$ and $A, M < \infty$. For $0 < \eta < 1/m$, suppose a sequence $(u^t)_{t\geq 1}$ in a finite-dimensional Euclidean space satisfies, for every $t \geq 1$,*

$$\Delta^m u^{t+m} = \sum_{k=0}^{m-1} c_{k,t} \eta^{m-k} \Delta^k u^{t+k}, \qquad |c_{k,t}| \leq A \quad (0 \leq k < m), \tag{23}$$

$$\left\|\Delta^k u^{t+k}\right\| \leq M\eta^k \qquad (0 \leq k \leq m). \tag{24}$$

*There is a constant $C < \infty$, depending only on $m, A, M$, such that every horizon $T \geq 1$ admits a finitely supported sequence $y = (y^t)_{t\in\mathbb{Z}}$ with*

$$\|\Delta y\|_{\ell_2}^2 \geq \sum\nolimits_{t=1}^{T} \left\|\Delta u^t\right\|^2 - C, \tag{25}$$

$$\|y\|_{\ell_2}^2 \leq \sum\nolimits_{t=1}^{T+1} \left\|u^t\right\|^2, \tag{26}$$

$$\|\Delta^m y\|_{\ell_2} \leq A \sum\nolimits_{k=0}^{m-1} \eta^{m-k} \left\|\Delta^k y\right\|_{\ell_2} + C\eta^{m-1/2}. \tag{27}$$

This finite-support construction is similar to that of Ha (2026), but is tailored to preserve the finite-order relation up to a horizon-independent error. The inequalities (25)–(26) let us transfer a first-order difference bound for $y$ back to $h_i^t$ only at a constant cost.

*Proof of Lemma 5.3.* We combine the finite-order relation in Lemma 5.8, the cutoff bound (27) from Lemma 5.9, and the finite-difference interpolation in Lemma 5.4 to bound $\|\Delta y\|_{\ell_2}$. We then use (25) and (26) from Lemma 5.9 to transfer the bound to $h_i^t$ with error independent of $T$.

Fix $i, T$ and take $m, A, M, \eta_i$ from Lemma 5.8. For $0 < \eta \leq \bar{\eta}_i := \frac{1}{2}\min\{\eta_i, 1/m, 1\}$, apply Lemma 5.9 to $(h_i^t)_{t\geq 1}$ to obtain the cutoff $y$, and write $C_1 \geq 0$ for its constant $C$.

The cutoff bound (27) in Lemma 5.9 and the finite-difference interpolation (13) in Lemma 5.4 give:

$$\begin{aligned}\|\Delta^m y\|_{\ell_2} &\overset{(27)}{\leq} A \sum_{k=0}^{m-1} \eta^{m-k} \left\|\Delta^k y\right\|_{\ell_2} + C_1\eta^{m-1/2} \\ &\overset{(13)}{\leq} A\eta^m \|y\|_{\ell_2} + A \sum_{k=1}^{m-1} \left(\eta^m \|y\|_{\ell_2}\right)^{1-k/m} \|\Delta^m y\|_{\ell_2}^{k/m} + C_1\eta^{m-1/2}.\end{aligned} \tag{28}$$

For each term in the sum, Young's inequality gives

$$A\left(\eta^m \|y\|_{\ell_2}\right)^{1-k/m} \|\Delta^m y\|_{\ell_2}^{k/m} \leq \frac{1}{2m} \|\Delta^m y\|_{\ell_2} + \frac{m-k}{m}(2k)^{\frac{k}{m-k}} A^{\frac{m}{m-k}} \eta^m \|y\|_{\ell_2}. \tag{29}$$

Substituting (29) into (28) gives a coefficient $(m-1)/(2m) \le 1/2$ on $\|\Delta^m y\|_{\ell_2}$. Moving this term to the left yields the following for some constant $C_2 \ge 0$ depending only on $m, A, M$:

$$\|\Delta^m y\|_{\ell_2} \le C_2 \eta^m \|y\|_{\ell_2} + C_2 \eta^{m-1/2}. \tag{30}$$

*Case* $m = 1$. Equation (30) already bounds $\|\Delta y\|_{\ell_2}$. Squaring and using $(a+b)^2 \le 2a^2 + 2b^2$ gives

$$\|\Delta y\|_{\ell_2}^2 \le 2C_2^2 \eta^2 \|y\|_{\ell_2}^2 + 2C_2^2 \eta.$$

*Case* $m \ge 2$. Invoking Lemma 5.4 with $j = 1$ and combining it with (30) gives:

$$\begin{aligned}
\|\Delta y\|_{\ell_2}^2 &\overset{(13)}{\le} \|y\|_{\ell_2}^{2-2/m} \|\Delta^m y\|_{\ell_2}^{2/m} \overset{(30)}{\le} C_2^{2/m} \|y\|_{\ell_2}^{2-2/m} \left(\eta^m \|y\|_{\ell_2} + \eta^{m-1/2}\right)^{2/m} \\
&\le C_2^{2/m} \eta^2 \|y\|_{\ell_2}^2 + C_2^{2/m} \left(\eta^2 \|y\|_{\ell_2}^2\right)^{1-1/m} \eta^{1/m} \\
&\le C_2^{2/m} \left((2 - 1/m)\eta^2 \|y\|_{\ell_2}^2 + \eta/m\right) \\
&\le 2C_2^{2/m} \left(\eta^2 \|y\|_{\ell_2}^2 + \eta\right).
\end{aligned} \tag{31}$$

The third inequality uses $(a+b)^{2/m} \le a^{2/m} + b^{2/m}$, $\forall m \ge 2$. The fourth applies Young's inequality $(\eta^2 \|y\|_{\ell_2}^2)^{1-1/m} \eta^{1/m} \le (1 - 1/m)\eta^2 \|y\|_{\ell_2}^2 + \eta/m$.

We now return to $h_i^t$ using (25) and (26). Either case and $\eta \le 1$ give

$$\sum_{t=1}^{T} \left\|h_i^{t+1} - h_i^t\right\|_2^2 \le \|\Delta y\|_{\ell_2}^2 + C_1 \le 2C_2^{2/m} \eta^2 \sum_{t=1}^{T+1} \left\|h_i^t\right\|_2^2 + C_1 + 2C_2^{2/m}.$$

The coefficients depend only on $(n, d)$ and $i$, through $m, A, M$. Taking common upper bounds over $i$ and $\eta_{\text{temp}} := \min_i \bar{\eta}_i$ proves Lemma 5.3. □

**Completing the regret bound.** Combining Lemma 5.1, Lemma 5.2, and Lemma 5.3 directly proves Theorem 4.1. We give a detailed calculation in Appendix A.

## 6 Discussion

**Implications.** Notably, a recent line of work has obtained horizon-independent regret bounds in finite general-sum games by combining modified regularization with higher-order prediction (Liu et al., 2026; Abbadi et al., 2026). However, Theorem 4.1 shows that plain Optimistic Hedge can already attain this horizon independence with a dimension-dependent constant step size. Our result suggests that adaptive learning rates (Soleymani et al., 2025a;b), hybrid regularization (Tsuchiya, 2026), and higher-order predictions (Liu et al., 2026; Abbadi et al., 2026) are unnecessary for this guarantee.

**Comparison with Ha (2026).** Both our analysis and the analysis of Ha (2026) build on the high-order smoothness framework of Daskalakis et al. (2021). Ha's centered-logit derivative bounds control each difference order separately. Choosing $m \asymp (\sqrt{n}\eta)^{-1}$ leaves a regret contribution proportional to $\eta T \exp\{-c/(\sqrt{n}\eta)\}$, leading to a logarithmic regret bound when the step size is tuned to the horizon. In our analysis, Frisch's theorem on the compact state space yields an exact relation among difference orders (Lemma 5.8). Its order is fixed by $(n, d)$, and it has no additive error that accumulates over $T$. Finite-difference interpolation (Lemma 5.4) then yields a first-order difference bound that lets us absorb prediction error into stability, up to a horizon-independent error. Intuitively, our approach thus trades Ha's explicit dimension dependence for horizon independence.

## 7 Conclusion

In this work, we proved that plain Optimistic Hedge with a constant step size can achieve horizon-independent individual regret in finite general-sum games. An exact finite-order relation for weighted gaps removes the horizon-accumulating remainder in the logarithmic analysis of Ha (2026), and yields an $O_{n,d}(1/T)$ CCE guarantee. Quantifying the explicit dimension dependence of $\eta_\star$ and $C_\star$ in our regret bound remains open, and we leave it as an intriguing future work.

### AI USE STATEMENT

The author formulated the initial horizon-independent regret hypothesis and conducted a numerical experiment to test the hypothesis empirically. The hypothesis was motivated by the close connection between optimistic and implicit update rules (Mokhtari et al., 2020). Specifically, the authors hypothesized that a similar connection might exist between Optimistic Hedge and Clairvoyant Multiplicative Weights Update (CMWU) (Piliouras et al., 2022), and therefore, Optimistic Hedge might also enjoy horizon-independent regret bounds. Under the authors' direction, generative AI tools were used to develop and write proof arguments based on the high-order smoothness framework of Daskalakis et al. (2021) and the weighted-gap reduction of Ha (2026). AI tools also assisted with literature search, exposition, and LaTeX editing. The author takes responsibility for the final content, including all mathematical claims, proofs, and references.

### REPRODUCIBILITY STATEMENT

The assumptions and algorithm are specified in Sections 3 and 4. Complete proofs appear in Section 5 and Appendices A–D.

## REFERENCES

Omar Abbadi, Rida Laraki, and Panayotis Mertikopoulos. Constant regret in general games via higher-order optimism. *arXiv preprint arXiv:2609.04113*, 2026. URL https://arxiv.org/abs/2609.04113.

Ioannis Anagnostides, Constantinos Daskalakis, Gabriele Farina, Maxwell Fishelson, Noah Golowich, and Tuomas Sandholm. Near-optimal no-regret learning for correlated equilibria in multi-player general-sum games. In *Proceedings of the ACM Symposium on Theory of Computing*, 2022a.

Ioannis Anagnostides, Gabriele Farina, Christian Kroer, Chung-Wei Lee, Haipeng Luo, and Tuomas Sandholm. Uncoupled learning dynamics with $O(\log T)$ swap regret in multiplayer games. In *Advances in Neural Information Processing Systems*, 2022b.

Robert J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974. doi: 10.1016/0304-4068(74)90037-8.

Constantin Bănică and Octavian Stănăşilă. *Algebraic Methods in the Global Theory of Complex Spaces*. John Wiley & Sons, 1976. ISBN 0-471-01809-0.

David Blackwell. An analog of the minimax theorem for vector payoffs. *Pacific Journal of Mathematics*, 6(1):1–8, 1956. doi: 10.2140/pjm.1956.6.1.

Avrim Blum and Yishay Mansour. From external to internal regret. *Journal of Machine Learning Research*, 8:1307–1324, 2007.

Noam Brown and Tuomas Sandholm. Superhuman AI for multiplayer poker. *Science*, 365(6456): 885–890, 2019. doi: 10.1126/science.aay2400.

Nicolò Cesa-Bianchi and Gábor Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, 2006.

Xi Chen and Binghui Peng. Hedging in games: Faster convergence of external and swap regrets. In *Advances in Neural Information Processing Systems*, 2020.

Xi Chen, Xiaotie Deng, and Shang-Hua Teng. Settling the complexity of computing two-player Nash equilibria. *Journal of the ACM*, 56(3):14:1–14:57, 2009. doi: 10.1145/1516512.1516516. URL https://arxiv.org/abs/0704.1678.

Chao-Kai Chiang, Tianbao Yang, Chia-Jung Lee, Mehrdad Mahdavi, Chi-Jen Lu, Rong Jin, and Shenghuo Zhu. Online optimization with gradual variations. In *Proceedings of the 25th Annual Conference on Learning Theory*, volume 23 of *Proceedings of Machine Learning Research*, pp. 6.1–6.20, 2012. URL https://proceedings.mlr.press/v23/chiang12.html.

Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009. doi: 10.1137/070699652.

Constantinos Daskalakis, Maxwell Fishelson, and Noah Golowich. Near-optimal no-regret learning in general games. In *Advances in Neural Information Processing Systems*, 2021.

Gabriele Farina, Ioannis Anagnostides, Haipeng Luo, Chung-Wei Lee, Christian Kroer, and Tuomas Sandholm. Near-optimal no-regret learning dynamics for general convex games. In *Advances in Neural Information Processing Systems*, 2022.

Yoav Freund and Robert E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *Journal of Computer and System Sciences*, 55(1):119–139, 1997. doi: 10.1006/jcss.1997.1504.

Yoav Freund and Robert E. Schapire. Adaptive game playing using multiplicative weights. *Games and Economic Behavior*, 29(1–2):79–103, 1999. doi: 10.1006/game.1999.0738.

Jacques Frisch. Points de platitude d'un morphisme d'espaces analytiques complexes. *Inventiones Mathematicae*, 4:118–138, 1967. doi: 10.1007/BF01425245.

Paul Gölz, Nika Haghtalab, and Kunhe Yang. Distortion of AI alignment: Does preference optimization optimize for preferences? In *Advances in Neural Information Processing Systems*, volume 38, 2025. doi: 10.52202/085713-0876. URL https://arxiv.org/abs/2505.23749.

Ian J. Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, volume 27, pp. 2672–2680, 2014.

Gert-Martin Greuel, Christoph Lossen, and Eugenii Shustin. *Introduction to Singularities and Deformations*. Springer Monographs in Mathematics. Springer, Berlin, Heidelberg, first edition, 2007. doi: 10.1007/3-540-28419-2.

Junsoo Ha. A Logarithmic Regret Bound for Optimistic Hedge in General-Sum Games. arXiv preprint arXiv:2609.19677, 2026. URL https://arxiv.org/abs/2609.19677.

James Hannan. Approximation to Bayes risk in repeated play. In Melvin Dresher, Albert W. Tucker, and Philip Wolfe (eds.), *Contributions to the Theory of Games, Volume III*, volume 39 of *Annals of Mathematics Studies*, pp. 97–139. Princeton University Press, Princeton, NJ, 1957.

Sergiu Hart and Andreu Mas-Colell. A simple adaptive procedure leading to correlated equilibrium. *Econometrica*, 68(5):1127–1150, 2000. doi: 10.1111/1468-0262.00153.

Sergiu Hart and Andreu Mas-Colell. A general class of adaptive strategies. *Journal of Economic Theory*, 98(1):26–54, 2001. doi: 10.1006/jeth.2000.2746.

Lars Hörmander. *An Introduction to Complex Analysis in Several Variables*, volume 7 of *North-Holland Mathematical Library*. North-Holland, third revised edition, 1990. ISBN 0-444-88446-7.

Mingyang Liu, Gabriele Farina, and Asuman Ozdaglar. Constant individual regret in general games. *arXiv preprint arXiv:2608.31166*, 2026.

Knud Lønsted. Vector bundles over finite CW-complexes are algebraic. *Proceedings of the American Mathematical Society*, 38(1):27–31, 1973.

Aryan Mokhtari, Asuman Ozdaglar, and Sarath Pattathil. A unified analysis of extra-gradient and optimistic gradient methods for saddle point problems: Proximal point approach. In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pp. 1497–1507. PMLR, 2020. URL https://proceedings.mlr.press/v108/mokhtari20a.html.

Hervé Moulin and Jean-Philippe Vial. Strategically zero-sum games: The class of games whose completely mixed equilibria cannot be improved upon. *International Journal of Game Theory*, 7 (3–4):201–221, 1978. doi: 10.1007/BF01769190.

Remi Munos, Michal Valko, Daniele Calandriello, Mohammad Gheshlaghi Azar, Mark Rowland, Zhaohan Daniel Guo, Yunhao Tang, Matthieu Geist, Thomas Mesnard, Côme Fiegel, Andrea Michi, Marco Selvi, Sertan Girgin, Nikola Momchev, Olivier Bachem, Daniel J. Mankowitz, Doina Precup, and Bilal Piot. Nash learning from human feedback. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 36743–36768. PMLR, 2024.

John F. Nash. Equilibrium points in $n$-person games. *Proceedings of the National Academy of Sciences*, 36(1):48–49, 1950. doi: 10.1073/pnas.36.1.48.

Christos H. Papadimitriou and Tim Roughgarden. Computing correlated equilibria in multi-player games. *Journal of the ACM*, 55(3):14:1–14:29, 2008. doi: 10.1145/1379759.1379762.

Georgios Piliouras, Ryann Sim, and Stratis Skoulakis. Beyond time-average convergence: Near-optimal uncoupled online learning via clairvoyant multiplicative weights update. In *Advances in Neural Information Processing Systems*, 2022.

Alexander Rakhlin and Karthik Sridharan. Optimization, learning, and games with predictable sequences. In *Advances in Neural Information Processing Systems*, 2013a.

Alexander Rakhlin and Karthik Sridharan. Online learning with predictable sequences. In *Proceedings of the 26th Annual Conference on Learning Theory*, volume 30 of *Proceedings of Machine Learning Research*, pp. 993–1019, 2013b. URL https://proceedings.mlr.press/v30/Rakhlin13.html.

Corby Rosset, Ching-An Cheng, Arindam Mitra, Michael Santacroce, Ahmed Awadallah, and Tengyang Xie. Direct Nash optimization: Teaching language models to self-improve with general preferences. *arXiv preprint arXiv:2404.03715*, 2024. URL https://arxiv.org/abs/2404.03715.

Ashkan Soleymani, Georgios Piliouras, and Gabriele Farina. Cautious optimism: A meta-algorithm for near-constant regret in general games. In *Proceedings of the 26th ACM Conference on Economics and Computation*, pp. 870, 2025a. doi: 10.1145/3736252.3742639. URL https://arxiv.org/abs/2506.05005.

Ashkan Soleymani, Georgios Piliouras, and Gabriele Farina. Faster rates for no-regret learning in general games via cautious optimism. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, pp. 518–529, 2025b. doi: 10.1145/3717823.3718242. URL https://arxiv.org/abs/2503.24340.

Vasilis Syrgkanis, Alekh Agarwal, Haipeng Luo, and Robert E. Schapire. Fast convergence of regularized learning in games. In *Advances in Neural Information Processing Systems*, 2015.

Yixin Tao and Weiqiang Zheng. Optimal alternating regret for online learning and games. *arXiv preprint arXiv:2608.24731*, 2026. URL https://arxiv.org/abs/2608.24731.

Taira Tsuchiya. Sublogarithmic swap regret in multiplayer general-sum games via hybrid regularization. *arXiv preprint arXiv:2608.04149*, 2026.

John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, 1944.

Fang Wu, Xu Huang, Weihao Xuan, Zhiwei Zhang, Yijia Xiao, Frank Wan, Xiaomin Li, Bing Hu, Peng Xia, Jure Leskovec, and Yejin Choi. Multiplayer Nash preference optimization. In *International Conference on Learning Representations*, 2026. URL https://iclr.cc/virtual/2026/oral/10006670.

Yue Wu, Zhiqing Sun, Huizhuo Yuan, Kaixuan Ji, Yiming Yang, and Quanquan Gu. Self-play preference optimization for language model alignment. *arXiv preprint arXiv:2405.00675*, 2024. URL https://arxiv.org/abs/2405.00675.

## APPENDIX CONTENTS

**A Proof of the main result** **13**

**B Preliminary lemmas** **14**

B.1 Optimistic regret bound . . . . . 14

B.2 Weighted pairwise loss gaps . . . . . 15

B.3 Interpolation of finite differences . . . . . 16

**C Exact finite-order difference relation** **17**

C.1 Analytic state update . . . . . 17

C.2 Analytic scaled finite differences . . . . . 17

C.3 Frisch's Noetherianity Theorem . . . . . 18

C.4 Exact finite-order relation . . . . . 20

**D Finite-support cutoff construction** **21**

## A PROOF OF THE MAIN RESULT

We derive the horizon-independent regret guarantee from the three estimates in Section 5.1.

**Theorem 4.1** (Horizon-independent individual regret)**.** *Fix $n$ players with action counts $d_1, \dots, d_n$. There exist $\eta_\star > 0$ and $C_\star < \infty$, depending only on $(n, d_1, \dots, d_n)$, such that, for every game in the setup of Section 3 and every common constant step size $0 < \eta \le \eta_\star$, the Optimistic Hedge iterates satisfy, for every $i \in [n]$ and $T \ge 1$,*

$$\mathrm{Reg}_i(T) \le \frac{\log d_i}{\eta} + C_\star \eta. \tag{6}$$

*Proof of Theorem 4.1.* Let $\eta_{\mathrm{rvu}}, \eta_{\mathrm{gap}}, \eta_{\mathrm{temp}}$ be the thresholds in Lemmas 5.1 to 5.3, with $\eta_{\mathrm{temp}} \le 1$. Since $\left\|h_i^{T+1}\right\|_2^2 \le 1/4$, the last two lemmas give

$$\sum_{t=1}^{T} \mathrm{Var}_{x_i^t}(\ell_i^t - \ell_i^{t-1}) \le C_0 \eta^2 \sum_{t=1}^{T} \mathrm{Var}_{x_i^t}(\ell_i^{t-1}) + C_1, \tag{32}$$

with constants $C_0, C_1 \ge 0$ depending only on $n, d$. For $0 < \eta \le \min\{\eta_{\mathrm{rvu}}, \eta_{\mathrm{gap}}, \eta_{\mathrm{temp}}\}$, substitution into Equation (8) yields

$$\mathrm{Reg}_i(T) \le \frac{\log d_i}{\eta} + \frac{2C_1\eta}{3} - \frac{\eta}{3}(1 - 2C_0\eta^2) \sum_{t=1}^{T} \mathrm{Var}_{x_i^t}(\ell_i^{t-1}).$$

Set

$$\eta_\star := \min\left\{\eta_{\mathrm{rvu}}, \eta_{\mathrm{gap}}, \eta_{\mathrm{temp}}, \frac{1}{\sqrt{2\max\{C_0, 1\}}}\right\}, \qquad C_\star := \frac{2C_1}{3}.$$

The variance coefficient is nonpositive for every $0 < \eta \le \eta_\star$, proving Equation (6). Both constants depend only on $n, d$. □

**Corollary 4.2** (Coarse-correlated-equilibrium convergence)**.** *Under the conditions of Theorem 4.1, the time-averaged joint distribution $\widehat{\mu}_T$ satisfies, for every $T \ge 1$,*

$$\mathrm{Gap}_{\mathrm{CCE}}(\widehat{\mu}_T) \le \frac{1}{T} \max_{i \in [n]} \left\{\frac{\log d_i}{\eta} + C_\star \eta\right\}. \tag{7}$$

*Proof.* Apply Equation (4) to Theorem 4.1: for every $T \geq 1$,

$$\mathrm{Gap}_{\mathrm{CCE}}(\widehat{\mu}_T) = \frac{1}{T} \max_{i \in [n]} \mathrm{Reg}_i(T) \leq \frac{1}{T} \max_{i \in [n]} \left\{ \frac{\log d_i}{\eta} + C_\star \eta \right\}.$$

This proves Equation (7). □

# B Preliminary lemmas

We prove the optimistic regret bound, the weighted-gap reduction, and finite-difference interpolation. For the first two, we only use the update rule of Optimistic Hedge and boundedness of loss vectors; we use the analytic structure of self-play in Appendix C.

## B.1 Optimistic regret bound

Fix a player $i$, write $d := d_i$, and fix $\eta > 0$. Define the soft-min potential $\Phi : \mathbb{R}^d \to \mathbb{R}$ by

$$\Phi(z) := -\frac{1}{\eta} \log \sum_{a=1}^{d} e^{-\eta z_a}$$

Its gradient gives exponential weights, while its value approximates the smallest coordinate. Since $-\Phi$ is convex, we can define its corresponding Bregman divergence for $y, z \in \mathbb{R}^d$ by

$$D_{-\Phi}(y, z) := \Phi(z) + \langle \nabla \Phi(z), y - z \rangle - \Phi(y).$$

For $v \in \mathbb{R}^d$, write $\mathrm{osc}(v) := \max_a v_a - \min_a v_a$ for its coordinate span. The following auxiliary bound from Ha (2026) compares each Bregman divergence with the variance under its reference distribution.

**Lemma B.1** (Bregman divergence and local variance (Ha, 2026, Lemma B.1))**.** *Let $z, v \in \mathbb{R}^d$ and $R \geq 0$, set $p := \nabla \Phi(z)$, and suppose that $\mathrm{osc}(v) \leq R$. Then*

$$\frac{\eta}{2} e^{-\eta R} \mathrm{Var}_p(v) \leq D_{-\Phi}(z + v, z) \leq \frac{\eta}{2} e^{\eta R} \mathrm{Var}_p(v). \tag{33}$$

*Proof.* We compare the potential's curvature along the segment from $z$ to $z+v$. Set $p_s := \nabla \Phi(z+sv)$ for $s \in [0, 1]$; these are the exponential weights along that segment. Since

$$-\nabla^2 \Phi(z + sv) = \eta \big( \mathrm{Diag}(p_s) - p_s p_s^\top \big),$$

the integral remainder formula gives

$$D_{-\Phi}(z + v, z) = \eta \int_0^1 (1 - s) \, \mathrm{Var}_{p_s}(v) \, \mathrm{d}s. \tag{34}$$

Moreover,

$$\frac{p_s(a)}{p(a)} = \frac{e^{-\eta s v_a}}{\sum_b p(b) e^{-\eta s v_b}},$$

so $e^{-\eta s R} p(a) \leq p_s(a) \leq e^{\eta s R} p(a)$ coordinatewise. Using

$$\mathrm{Var}_p(v) = \min_{c \in \mathbb{R}} \sum_a p(a)(v_a - c)^2,$$

we obtain

$$e^{-\eta R} \mathrm{Var}_p(v) \leq \mathrm{Var}_{p_s}(v) \leq e^{\eta R} \mathrm{Var}_p(v).$$

Substitute into Equation (34) and use $\int_0^1 (1 - s) \mathrm{d}s = 1/2$. □

The optimistic regret bound of Ha (2026) specializes Lemma 4.1 of Daskalakis et al. (2021), whose Lemma A.5 refines Lemma 3 of Rakhlin & Sridharan (2013b). We reproduce its proof here for completeness.

**Lemma 5.1** (Optimistic regret bound (Ha, 2026, Lemma 5.1)). *There is a universal $\eta_{\mathrm{rvu}} > 0$ such that, for every player, horizon, and $0 < \eta \le \eta_{\mathrm{rvu}}$,*

$$\mathrm{Reg}_i(T) \le \frac{\log d_i}{\eta} + \frac{2\eta}{3}\sum_{t=1}^{T}\mathrm{Var}_{x_i^t}(\ell_i^t - \ell_i^{t-1}) - \frac{\eta}{3}\sum_{t=1}^{T}\mathrm{Var}_{x_i^t}(\ell_i^{t-1}). \tag{8}$$

*Proof.* Suppress the player index from $x_i^t$, $\ell_i^t$, and $\mathrm{Reg}_i$. We first express the strategy in terms of past losses. Set $L^0 := 0$, and for $t \ge 1$ let

$$L^t := \sum_{\tau=1}^{t}\ell^\tau, \qquad q^t := L^{t-1} + \ell^{t-1}.$$

Here $L^t$ is the cumulative loss, and $q^t$ augments the observed cumulative loss with the latest loss vector as a prediction for round $t$. Unrolling Optimistic Hedge gives

$$x_a^t = \frac{e^{-\eta q_a^t}}{\sum_{b\in[d]} e^{-\eta q_b^t}}, \qquad a \in [d].$$

This discrepancy is the curvature correction in a potential increment. Using $\nabla\Phi(q^t) = x^t$ and $L^t - L^{t-1} = \ell^t$, we obtain the one-round identity and its cumulative regret consequence:

$$D_{-\Phi}(L^t, q^t) - D_{-\Phi}(L^{t-1}, q^t) = \langle x^t, \ell^t\rangle - \Phi(L^t) + \Phi(L^{t-1}),$$
$$\mathrm{Reg}(T) \le \frac{\log d}{\eta} + \sum_{t=1}^{T}\left[D_{-\Phi}(L^t, q^t) - D_{-\Phi}(L^{t-1}, q^t)\right]. \tag{35}$$

The potential terms telescope, and $\Phi(L^T) \le \min_a L_a^T$ compares the final potential with the best fixed action. Uniform initialization gives $\Phi(0) = -(\log d)/\eta$, hence the initial cost in the second line.

Apply Lemma B.1 to the positive and negative divergences in Equation (35). Since $L^t - q^t = \ell^t - \ell^{t-1}$ has coordinate span at most 2, while $L^{t-1} - q^t = -\ell^{t-1}$ has span at most 1, Equations (33) and (35) give the exact inequality

$$\mathrm{Reg}(T) \le \frac{\log d}{\eta} + \frac{\eta e^{2\eta}}{2}\sum_{t=1}^{T}\mathrm{Var}_{x^t}(\ell^t - \ell^{t-1}) - \frac{\eta e^{-\eta}}{2}\sum_{t=1}^{T}\mathrm{Var}_{x^t}(\ell^{t-1}). \tag{36}$$

Choose a universal $\eta_{\mathrm{rvu}} > 0$ such that $e^{2\eta}/2 \le 2/3$ and $e^{-\eta}/2 \ge 1/3$ for $0 < \eta \le \eta_{\mathrm{rvu}}$. This proves Lemma 5.1. □

### B.2 Weighted pairwise loss gaps

For completeness, we reproduce the weighted-gap reduction from Lemma 5.2 of Ha (2026). The pairwise identity transfers the variances to a fixed Euclidean norm; the correction comes from the first-order differences of the probability weights.

**Lemma 5.2** (Weighted-gap reduction (Ha, 2026, Lemma 5.2)). *There are universal constants $C > 0$ and $\eta_{\mathrm{gap}} > 0$ such that, whenever $0 < \eta \le \eta_{\mathrm{gap}}$,*

$$\begin{aligned}\sum_{t=1}^{T}\mathrm{Var}_{x_i^t}(\ell_i^{t-1}) &= \sum_{t=1}^{T}\left\|h_i^t\right\|_2^2,\\ \sum_{t=1}^{T}\mathrm{Var}_{x_i^t}(\ell_i^t - \ell_i^{t-1}) &\le \frac{3}{2}\sum_{t=1}^{T}\left\|h_i^{t+1} - h_i^t\right\|_2^2 + C\eta^2\sum_{t=1}^{T}\left\|h_i^t\right\|_2^2 + C\eta^2.\end{aligned} \tag{11}$$

*Proof.* The pairwise variance identity (9) follows by expanding $\frac{1}{2}\sum_{a,b} x_a x_b (v_a - v_b)^2$ and noting that the diagonal terms vanish. Applying it to Equation (10) gives the identity in Equation (11).

For prediction error, the remaining obstacle is that the strategy weights change between rounds. Fix player $i$ and pair $a < b$, and write

$$w_{ab}^t := \sqrt{x_{i,a}^t x_{i,b}^t}.$$

Uniform initialization and the multiplicative update keep $w_{ab}^t > 0$ at every finite round, so these ratios are well defined. Their bounds remain uniform as probabilities approach zero. The optimistic loss vector $m_i^t := 2\ell_i^t - \ell_i^{t-1}$ has $\mathrm{osc}(m_i^t) \le 3$, and Equation (5) gives

$$e^{-3\eta} \le \frac{w_{ab}^{t+1}}{w_{ab}^t} \le e^{3\eta}, \qquad \left| \frac{w_{ab}^t}{w_{ab}^{t+1}} - 1 \right| \le C\eta \tag{37}$$

for a universal $C$ and all sufficiently small $\eta$. Thus the relative change of each pair weight is $O(\eta)$.

Since $[h_i^t]_{ab} = w_{ab}^t(\ell_{i,a}^{t-1} - \ell_{i,b}^{t-1})$,

$$w_{ab}^t\big[(\ell_{i,a}^t - \ell_{i,b}^t) - (\ell_{i,a}^{t-1} - \ell_{i,b}^{t-1})\big] = [h_i^{t+1} - h_i^t]_{ab} + \left(\frac{w_{ab}^t}{w_{ab}^{t+1}} - 1\right)[h_i^{t+1}]_{ab}. \tag{38}$$

The first term is the weighted-gap difference $[\Delta h_i^t]_{ab}$; the second accounts for the weight ratio $w_{ab}^t/w_{ab}^{t+1}$. Its coefficient is $O(\eta)$, so squaring yields the $O(\eta^2)$ correction in the lemma. Apply $(a+b)^2 \le \frac{3}{2}a^2 + 3b^2$ to Equation (38), then sum over $t$ and $a < b$. Using Equations (9) and (37),

$$\begin{aligned}
\sum_{t=1}^T \mathrm{Var}_{x_i^t}(\ell_i^t - \ell_i^{t-1}) &= \sum_{t=1}^T \sum_{a<b} \Big(w_{ab}^t\big[(\ell_{i,a}^t - \ell_{i,b}^t) - (\ell_{i,a}^{t-1} - \ell_{i,b}^{t-1})\big]\Big)^2 \\
&\le \frac{3}{2}\sum_{t=1}^T \left\|h_i^{t+1} - h_i^t\right\|_2^2 + C\eta^2 \sum_{t=1}^T \left\|h_i^{t+1}\right\|_2^2 \\
&\le \frac{3}{2}\sum_{t=1}^T \left\|h_i^{t+1} - h_i^t\right\|_2^2 + C\eta^2 \sum_{t=1}^T \left\|h_i^t\right\|_2^2 + C\eta^2.
\end{aligned}$$

The last inequality uses $h_i^1 = 0$ and

$$\sum_{t=1}^T \left\|h_i^{t+1}\right\|_2^2 = \sum_{t=1}^T \left\|h_i^t\right\|_2^2 + \left\|h_i^{T+1}\right\|_2^2 \le \sum_{t=1}^T \left\|h_i^t\right\|_2^2 + \frac{1}{4}.$$

The terminal bound follows from $\left\|h_i^{T+1}\right\|_2^2 = \mathrm{Var}_{x_i^{T+1}}(\ell_i^T) \le 1/4$. Thus moving the summation window contributes only the final $C\eta^2$ term, and Lemma 5.2 follows. □

### B.3 Interpolation of finite differences

This interpolation step is independent of Optimistic Hedge. It applies because the reduction to weighted gaps has put all difference orders in one Hilbert space with a fixed norm.

**Lemma 5.4** (Finite-difference interpolation)**.** *For every finitely supported sequence $y = (y^t)_{t\in\mathbb{Z}}$ in a finite-dimensional Hilbert space, the following holds for all integers $m \ge 1$ and $0 \le j \le m$:*

$$\left\|\Delta^j y\right\|_{\ell_2} \le \|y\|_{\ell_2}^{1-j/m} \left\|\Delta^m y\right\|_{\ell_2}^{j/m}. \tag{13}$$

*Proof.* Fourier transformation turns finite differences into multiplication, allowing all orders to be compared in one representation. Let $\widehat{y}(\omega) := \sum_{t\in\mathbb{Z}} y^t e^{-it\omega}$ be the Fourier transform and $\lambda(\omega) := |e^{i\omega} - 1|^2$ the squared magnitude of the first-order difference multiplier. Parseval's identity gives

$$\left\|\Delta^j y\right\|_{\ell_2}^2 = \frac{1}{2\pi}\int_{-\pi}^{\pi} \lambda(\omega)^j \left\|\widehat{y}(\omega)\right\|^2 \,\mathrm{d}\omega.$$

Thus each order weights the same Fourier components by a different power of $\lambda$. The cases $j = 0, m$ are immediate. For $0 < j < m$, apply Hölder's inequality with exponents $m/(m-j)$ and $m/j$ to

$$\lambda^j \left\|\widehat{y}\right\|^2 = \left(\left\|\widehat{y}\right\|^2\right)^{1-j/m}\left(\lambda^m \left\|\widehat{y}\right\|^2\right)^{j/m}.$$

This proves Equation (13). □

# C EXACT FINITE-ORDER DIFFERENCE RELATION

We first extend the state update and scaled finite-difference maps analytically through zero step size. Appendix C.3 introduces the analytic prerequisites and proves Frisch's Noetherianity theorem from them. Appendix C.4 then proves Theorem 5.5 and applies it uniformly over states and games to obtain Lemma 5.8. We defer the finite-support construction to Appendix D.

## C.1 ANALYTIC STATE UPDATE

To preserve the quantifiers in Theorem 4.1, we include the loss table in the compact parameter set from the outset. Define

$$\mathcal{R}_i := \{r_i \in \mathbb{R}_+^{d_i} : \|r_i\|_2 = 1\}, \qquad \mathcal{R} := \prod_{i=1}^n \mathcal{R}_i, \qquad \mathcal{V} := \prod_{i=1}^n [0,1]^{d_i}.$$

Thus $\mathcal{S} = \mathcal{R} \times \mathcal{V}$, where $\mathcal{V}$ contains the stored loss profiles. Let $\mathcal{G} := [0,1]^{n\prod_j d_j}$ be the set of loss tables, and write $g = (g_1, \dots, g_n) \in \mathcal{G}$. For $r_i \neq 0$, set $[x_i(r_i)]_a := r_i(a)^2 / \|r_i\|_2^2$. For $r \in \mathcal{R}$, define the current expected-loss profile by

$$[\ell_g(r)]_i(a_i) := \sum_{a_{-i} \in \prod_{j\neq i} A_j} g_i(a_i, a_{-i}) \prod_{j \neq i} [x_j(r_j)]_{a_j}.$$

Set $K_0 := \mathcal{R} \times \mathcal{V} \times \mathcal{G}$ and $K_\star := K_0 \times \{0\}$. We restore the loss-table argument in the maps as $F_i(r, \ell; g, \eta)$, $G(s; g, \eta)$, and $H_{i,k}(s; g, \eta)$; under $G^j(s; g, \eta)$, we keep both $g$ and $\eta$ fixed. The update is

$$[F_i(r, \ell; g, \eta)]_a := \frac{\|r_i\|_2 \; r_i(a) e^{-\eta(2[\ell_g(r)]_i(a) - [\ell_i]_a)/2}}{\left(\sum_b r_i(b)^2 e^{-\eta(2[\ell_g(r)]_i(b) - [\ell_i]_b)}\right)^{1/2}},$$
$$G(r, \ell; g, \eta) := \left(\left(F_i(r, \ell; g, \eta)\right)_{i=1}^n, \ell_g(r)\right).$$

**Lemma 5.6** (Analyticity of the state update)**.** *The state update $G(s; \eta)$ is jointly real analytic in the state, the loss-table entries, and $\eta$ on one common neighborhood of $K_\star$. The state space $\mathcal{S}$ is invariant. Moreover, $G(r, \ell; 0) = (r, \ell(r))$ and $G^2(s; 0) = G(s; 0)$ throughout this neighborhood.*

*Proof.* We verify joint analyticity, invariance, and the zero-step identity in turn. On a sufficiently small neighborhood of $\mathcal{R}$, every $\|r_i\|_2$ stays bounded away from zero. Hence $x_i(r_i)$ and $\ell_g(r)$ are rational, and therefore jointly real analytic, in $(r, g)$. At $\eta = 0$, the expression under the denominator square root in Equation (15) equals $\|r_i\|_2^2$, which is uniformly positive on $K_0$. Continuity and compactness give one product neighborhood and some $\bar{\eta} > 0$ on which every denominator remains positive for $|\eta| < \bar{\eta}$. The positive square-root function is analytic on $(0, \infty)$, so the full state update is jointly real analytic there.

For $r_i \in \mathcal{R}_i$,

$$\sum_a [F_i(r, \ell; g, \eta)]_a^2 = \frac{\sum_a r_i(a)^2 e^{-\eta(2[\ell_g(r)]_i(a) - [\ell_i]_a)}}{\sum_b r_i(b)^2 e^{-\eta(2[\ell_g(r)]_i(b) - [\ell_i]_b)}} = 1.$$

The update also preserves nonnegativity, and $\ell_g(r) \in \mathcal{V}$; hence $G$ leaves $\mathcal{S}$ invariant. At $\eta = 0$, the denominator equals $\|r_i\|_2$, so $F_i(r, \ell; g, 0) = r_i$ throughout the ambient neighborhood. Therefore $G(r, \ell; g, 0) = (r, \ell_g(r))$. A second zero-step update leaves both components unchanged, proving the zero-step identity. □

## C.2 ANALYTIC SCALED FINITE DIFFERENCES

For $s = (r, \ell)$, define the weighted-gap map by

$$[H_{i,0}(s; g, \eta)]_{ab} := r_i(a) r_i(b) \left([\ell_i]_a - [\ell_i]_b\right), \qquad a < b,$$

and, for $\eta \neq 0$, recursively define

$$H_{i,k+1}(s; g, \eta) := \frac{H_{i,k}(G^2(s; g, \eta); g, \eta) - H_{i,k}(G(s; g, \eta); g, \eta)}{\eta}.$$

The trajectory identity follows directly by induction. The case $k = 0$ follows from $H_{i,0}(s^t; g, \eta) = h_i^t$. If it holds at order $k$, then

$$H_{i,k+1}(s^t; g, \eta) = \frac{H_{i,k}(s^{t+2}; g, \eta) - H_{i,k}(s^{t+1}; g, \eta)}{\eta}$$
$$= \eta^{-(k+1)} \left(\Delta^k h_i^{t+k+2} - \Delta^k h_i^{t+k+1}\right) = \eta^{-(k+1)} \Delta^{k+1} h_i^{t+k+1}.$$

Thus $H_{i,k}(s^t; g, \eta) = \eta^{-k} \Delta^k h_i^{t+k}$ for every $k \geq 0$ and $t \geq 1$.

The zero-step identity makes the numerator vanish at $\eta = 0$. The next proof turns that cancellation into an analytic extension at every fixed order.

**Lemma 5.7** (Analyticity of scaled finite differences)**.** *For every fixed $k \geq 0$, the function $H_{i,k}$ in Equation* (18) *extends real analytically through $\eta = 0$, jointly in the state and loss-table entries, on a neighborhood of $K_\star$. The neighborhood may depend on $k$, but not on the state or game.*

*Proof.* We argue by induction and shrink the common neighborhood of $K_\star$ only when a composition requires it. The base map $H_{i,0}$ is polynomial. Suppose $H_{i,k}$ is analytic on such a neighborhood, and define

$$N_{i,k}(s; g, \eta) := H_{i,k}(G^2(s; g, \eta); g, \eta) - H_{i,k}(G(s; g, \eta); g, \eta).$$

By Lemma 5.6, $N_{i,k}(s; g, 0) = 0$ on a neighborhood in $(s, g)$. After restricting to a product neighborhood on which $N_{i,k}$ is analytic, the fundamental theorem of calculus gives, coordinatewise,

$$N_{i,k}(s; g, \eta) = \eta \int_0^1 \partial_\eta N_{i,k}(s; g, \tau\eta)\, \mathrm{d}\tau.$$

The integral is jointly analytic on a possibly smaller neighborhood and equals $N_{i,k}/\eta = H_{i,k+1}$ for $\eta \neq 0$. It therefore extends $H_{i,k+1}$ analytically through $\eta = 0$. □

### C.3 Frisch's Noetherianity Theorem

Frisch's theorem lets us select finitely many analytic generators near a compact semianalytic set. We introduce the definitions and three classical analytic lemmas needed for its proof. We state these lemmas in terms of scalar functions and analytic coefficients; their analytic foundations are supplied by the cited results. We reproduce the proof of Frisch's theorem using these lemmas.

**Definition C.1** (Real-analytic functions)**.** *Let $U \subset \mathbb{R}^q$ be open. A function $f : U \to \mathbb{R}$ is real-analytic if, for every $x \in U$, $f(x')$ equals a convergent power series in $x' - x$ for $x'$ in a neighborhood of $x$.*

**Definition C.2** (Analytic germs)**.** *Two real-analytic functions, each defined on an open neighborhood of a compact set $K \subset \mathbb{R}^q$, represent the same germ along $K$ if they agree on an open neighborhood of $K$ contained in both domains. A germ along $K$ is an equivalence class under this relation. The set of these germs is denoted by $\mathcal{O}(K)$.*

Intuitively, $\mathcal{O}(K)$ describes real-analytic functions near $K$, identifying two functions whenever they agree on an open neighborhood of $K$. In particular, $\mathcal{O}(K)$ forms a commutative ring: the sum and product of germs represented by $f$ and $g$ are represented by $f + g$ and $fg$ on the intersection of their domains. When $K = \{x\}$, write $\mathcal{O}_x := \mathcal{O}(\{x\})$ and denote the germ of $f$ at $x$ by $f_x$. Note that agreement only at the points of $K$, rather than on a neighborhood of $K$, does not identify germs.

**Definition C.3** (Ideals)**.** *A subset $I$ of a commutative ring $R$ is an ideal if it contains zero and is closed under addition and multiplication by arbitrary elements of $R$.*

**Definition C.4** (Finitely generated ideals)**.** *An ideal $I \subseteq R$ is finitely generated if there exist $f_1, \ldots, f_r \in I$ such that*

$$I = (f_1, \ldots, f_r) := \left\{ \sum_{j=1}^{r} a_j f_j : a_1, \ldots, a_r \in R \right\}.$$

We write $f_x \in (f_{1,x}, \ldots, f_{r,x})$ in $\mathcal{O}_x$ to denote that $f = \sum_j a_j f_j$ on some neighborhood of $x$, with analytic coefficients $a_j$ there.

**Definition C.5** (Noetherian rings)**.** *A commutative ring $R$ is Noetherian if every ideal of $R$ is finitely generated.*

**Definition C.6** (Semianalytic sets)**.** *A set $S \subset \mathbb{R}^q$ is semianalytic if, for every $x \in \mathbb{R}^q$, there exists an open neighborhood $U \subset \mathbb{R}^q$ of $x$ such that $S \cap U$ is a finite union of sets, each of the form*

$$\{x' \in U : f_1(x') = \cdots = f_r(x') = 0, \quad g_1(x') > 0, \ldots, g_s(x') > 0\},$$

*where $r, s \geq 0$ are integers and the functions $f_i, g_j : U \to \mathbb{R}$ are real analytic; either list of conditions may be empty. The neighborhood $U$ and the finite families of defining functions may depend on $x$.*

In particular, a set defined by finitely many polynomial equalities and weak inequalities is semianalytic, since a weak inequality is the union of a strict inequality and an equality.

The following three lemmas, Lemma C.7, Lemma C.8, and Lemma C.9, establish when $f$ can be written as $\sum_j a_j f_j$ with analytic coefficients $a_j$. They give finite generators at a point, a common neighborhood for local representations, and analytic coefficients valid near all of $K$.

We first seek finitely many generators at one point, which would allow us to write every germ in an ideal of $\mathcal{O}_x$ as their linear combination.

**Lemma C.7** (Finite generation at a point (Greuel et al., 2007, Theorem 1.15))**.** *For every $x \in \mathbb{R}^q$, the ring $\mathcal{O}_x$ of real-analytic germs at $x$ is Noetherian.*

The lemma implies that every germ in an ideal of $\mathcal{O}_x$ can be expressed using the same finite number of generators. Hence, in $f = \sum_{j=1}^{r} a_j f_j$, the coefficients $a_j$ may vary with $f$, but the generators $f_1, \ldots, f_r$ are shared across all real-analytic functions $f$ whose germs belong to that ideal.

The following lemma allows us to extend Lemma C.7's finite generation at a single point to its neighborhood within $K$.

**Lemma C.8** (Finite generation on a neighborhood (Bănică & Stănăşilă, 1976, Chapter V, §3, pp. 186–187))**.** *Let $K \subset \mathbb{R}^q$ be compact and semianalytic, let $x \in K$, and fix $f_1, \ldots, f_r \in \mathcal{O}(K)$. There is a relative open neighborhood $W$ of $x$ in $K$ such that, for every $f \in \mathcal{O}(K)$,*

$$f_x \in (f_{1,x}, \ldots, f_{r,x}) \quad \Longrightarrow \quad f_{x'} \in (f_{1,x'}, \ldots, f_{r,x'}) \quad \text{for every } x' \in W.$$

*The neighborhood $W$ may depend on $K, x, f_1, \ldots, f_r$, but not on $f$.*

Intuitively, Lemma C.7 selects finitely many generators from an ideal at $x$, and Lemma C.8 gives a single neighborhood where those generators work locally for every element of the ideal. Compactness then allows us to collect these generators into a single finite family that works at every point of $K$. The following lemma turns these local representations into a representation with analytic coefficients valid near all of $K$.

**Lemma C.9** (Analytic coefficients near a compact set (Hörmander, 1990, Theorem 7.2.1(ii)))**.** *Let $K \subset \mathbb{R}^q$ be compact and let $f, f_1, \ldots, f_r \in \mathcal{O}(K)$. If $f_x \in (f_{1,x}, \ldots, f_{r,x})$ for every $x \in K$, then there are $a_1, \ldots, a_r \in \mathcal{O}(K)$ such that*

$$f = \sum_{j=1}^{r} a_j f_j \quad \text{in } \mathcal{O}(K).$$

Together with compactness, these three lemmas yield Frisch's finite-generation theorem.

**Theorem C.10** (Frisch's Noetherianity theorem, real-analytic form (Frisch, 1967, Theorem I.9))**.** *Let $K \subset \mathbb{R}^q$ be compact and semianalytic. Then $\mathcal{O}(K)$ is Noetherian: every ideal $I \subseteq \mathcal{O}(K)$ has finitely many generators $f_1, \ldots, f_r \in I$, so that*

$$I = \left\{ \sum_{j=1}^{r} a_j f_j : a_1, \ldots, a_r \in \mathcal{O}(K) \right\}.$$

*Proof.* Fix an ideal $I \subseteq \mathcal{O}(K)$; the case $K = \varnothing$ is immediate. We first choose generators that work locally throughout $K$, and then obtain analytic coefficients near the whole set.

*Step 1: finite generators near each point.* For $x \in K$, let $I_x \subseteq \mathcal{O}_x$ be the ideal generated by $\{f_x : f \in I\}$. By Lemma C.7, $I_x$ has finitely many generators. Each generator is a finite combination of germs coming from $I$. Collecting the elements of $I$ in these combinations gives $f_{x,1}, \dots, f_{x,r_x} \in I$ whose germs generate $I_x$. If $I_x = \{0\}$, take the single generator $0$.

Apply Lemma C.8 to this fixed finite family. It gives a relative open neighborhood $W_x$ of $x$ in $K$ such that

$$f_{x'} \in (f_{x,1,x'}, \dots, f_{x,r_x,x'}) \qquad \text{for every } f \in I \text{ and every } x' \in W_x.$$

The same $W_x$ works for every $f \in I$. Thus the selected elements generate $I_{x'}$ at every $x' \in W_x$.

*Step 2: one finite family for the compact set.* The sets $W_x$ cover $K$. Choose a finite subcover and collect the corresponding generators into $f_1, \dots, f_r \in I$. For every $x \in K$, Step 1 gives

$$I_x = (f_{1,x}, \dots, f_{r,x}) \subseteq \mathcal{O}_x.$$

The reverse inclusion needed for this equality holds because every selected $f_j$ belongs to $I$.

*Step 3: coefficients on a neighborhood of $K$.* Fix any $f \in I$. By Step 2, $f_x \in (f_{1,x}, \dots, f_{r,x})$ for every $x \in K$. Lemma C.9 gives $a_1, \dots, a_r \in \mathcal{O}(K)$ with $f = \sum_j a_j f_j$. Hence $I \subseteq (f_1, \dots, f_r)$, and the reverse inclusion follows from the ideal property. This proves that $I$ is finitely generated.

□

### C.4 Exact finite-order relation

Frisch's theorem gives finite generation for scalar analytic functions. We first extend Frisch's theorem to vector-valued maps and prove Theorem 5.5. Then, we apply Theorem 5.5 to the scaled finite differences and obtain Lemma 5.8.

**Theorem 5.5** (Finite relation among analytic maps)**.** *Let $K \subset \mathbb{R}^q$ be compact and semianalytic, and let $V$ be a finite-dimensional real vector space. For every sequence $\Phi_0, \Phi_1, \dots$ of $V$-valued maps, each real analytic on a neighborhood of $K$, there exist an integer $m \geq 1$, an open neighborhood $U$ of $K$, and scalar real-analytic functions $c_0, \dots, c_{m-1}$ on $U$ such that the following holds on $U$:*

$$\Phi_m = \sum_{k=0}^{m-1} c_k \Phi_k. \tag{19}$$

*Proof.* Put $R := \mathcal{O}(K)$ and $d_V := \dim V$. By Theorem C.10, every ideal of $R$ is finitely generated. We first verify the corresponding fact for a submodule $L \subseteq R^{d_V}$, a set of vectors containing zero and closed under addition and multiplication by elements of $R$. Induct on $d_V$. The case $d_V = 0$ is immediate, and the case $d_V = 1$ is the ideal property. For $d_V \geq 2$, projection onto the last coordinate sends $L$ to an ideal of $R$. Choose finitely many generators of that ideal and lift them to vectors $v_1, \dots, v_r \in L$. The kernel is a submodule of $R^{d_V - 1}$, so it has finitely many generators by induction. For any $v \in L$, subtracting a suitable $R$-linear combination of $v_1, \dots, v_r$ makes its last coordinate zero. The remainder is generated by the kernel generators. These two families therefore generate $L$.

Choose a basis of $V$, and write $[\Phi_k] \in R^{d_V}$ for the germ of $\Phi_k$ along $K$. Consider the submodule

$$M := \sum_{k \geq 0} R[\Phi_k] \subseteq R^{d_V},$$

where every sum defining an element of $M$ is finite. By the preceding argument, $M$ has finitely many generators. Each is a finite combination of the $[\Phi_k]$, so there is an $N \geq 0$ with $M = \sum_{k=0}^{N} R[\Phi_k]$. If $M = \{0\}$, take $N = 0$. Setting $m := N + 1$ gives

$$[\Phi_m] = \sum_{k=0}^{m-1} [c_k][\Phi_k] \qquad \text{for some } [c_0], \dots, [c_{m-1}] \in R.$$

Only finitely many maps and coefficient germs enter this equality. Choose their representatives on the intersection of their open neighborhoods of $K$. Equality as germs means that, after a further restriction to an open neighborhood $U \supseteq K$, the equality holds pointwise throughout $U$. This proves Equation (19). The argument does not require one common neighborhood for the entire infinite sequence. □

The loss table is part of the compact set $K_\star$. This lets us choose the finite relation once for the entire dimension class, before fixing a game or trajectory.

**Lemma 5.8** (Exact finite-order relation)**.** *Fix the number of players $n$, number of actions $d = (d_1, \dots, d_n)$, and a player $i$. There are an integer $m \geq 1$, constants $A, M < \infty$, and $\eta_i > 0$, uniform over all games of these dimensions, such that, for every such game, $0 < \eta \leq \eta_i$, and $t \geq 1$, there exist scalars $c_{0,t}, \dots, c_{m-1,t}$, possibly depending on the game, $\eta$, and $t$, satisfying*

$$\Delta^m h_i^{t+m} = \sum_{k=0}^{m-1} c_{k,t} \eta^{m-k} \Delta^k h_i^{t+k}, \qquad |c_{k,t}| \leq A, \tag{21}$$

$$\left\|\Delta^k h_i^{t+k}\right\|_2 \leq M\eta^k, \qquad 0 \leq k \leq m. \tag{22}$$

*Proof.* Fix the player $i$. Finitely many polynomial equalities and weak inequalities define $K_\star = K_0 \times \{0\}$, where $K_0 = \mathcal{R} \times \mathcal{V} \times \mathcal{G}$. Thus $K_\star$ is compact and semianalytic. By Lemma 5.7, each $H_{i,k}$ is analytic on a neighborhood of $K_\star$. Apply Theorem 5.5 to the sequence $(H_{i,k})_{k\geq 0}$ to obtain an integer $m \geq 1$ and analytic scalar coefficients $c_0, \dots, c_{m-1}$ such that

$$H_{i,m} = \sum_{k=0}^{m-1} c_k H_{i,k}$$

on a common open neighborhood $U$ of $K_\star$. Intersect $U$ with the domains of $G, H_{i,0}, \dots, H_{i,m}$ if needed; there are only finitely many of these domains.

Because $K_0$ is compact and $U$ is open, some $\eta_i > 0$ satisfies

$$K_0 \times [-\eta_i, \eta_i] \subseteq U. \tag{39}$$

Continuity on this compact set gives finite constants $A, M$ such that, for every $0 \leq k < m$, $|c_k| \leq A$, and, for every $0 \leq k \leq m$, $\|H_{i,k}\|_2 \leq M$. The compact set and analytic maps were fixed by $(n, d)$ and $i$ before any particular loss table was chosen. Hence $m, \eta_i, A, M$ are uniform over all games of these dimensions and do not depend on the horizon.

Fix any such game, $0 < \eta \leq \eta_i$, and $t \geq 1$. Invariance of $\mathcal{S}$ gives $(s^t, g, \eta) \in K_0 \times [-\eta_i, \eta_i]$. Evaluating the relation there and using the trajectory identity from Equation (18) yields

$$\eta^{-m}\Delta^m h_i^{t+m} = H_{i,m}(s^t; g, \eta) = \sum_{k=0}^{m-1} c_k(s^t; g, \eta)\eta^{-k}\Delta^k h_i^{t+k}.$$

Set $c_{k,t} := c_k(s^t; g, \eta)$ and multiply by $\eta^m$ to obtain Equation (21), with $|c_{k,t}| \leq A$. For $0 \leq k \leq m$, the same trajectory identity and the bound on $H_{i,k}$ give

$$\left\|\Delta^k h_i^{t+k}\right\|_2 = \eta^k \left\|H_{i,k}(s^t; g, \eta)\right\|_2 \leq M\eta^k,$$

which is Equation (22). □

## D FINITE-SUPPORT CUTOFF CONSTRUCTION

To localize the trajectory, we use a gradual cutoff rather than zero-padding. This preserves the $\eta^m$ scale of the recurrence and controls the cost of returning to the original segment.

**Lemma 5.9** (Finite-support cutoff)**.** *Fix an integer $m \geq 1$ and $A, M < \infty$. For $0 < \eta < 1/m$, suppose a sequence $(u^t)_{t\geq 1}$ in a finite-dimensional Euclidean space satisfies, for every $t \geq 1$,*

$$\Delta^m u^{t+m} = \sum_{k=0}^{m-1} c_{k,t} \eta^{m-k} \Delta^k u^{t+k}, \qquad |c_{k,t}| \leq A \quad (0 \leq k < m), \tag{23}$$

$$\left\|\Delta^k u^{t+k}\right\| \leq M\eta^k \qquad (0 \leq k \leq m). \tag{24}$$

*There is a constant $C < \infty$, depending only on $m, A, M$, such that every horizon $T \geq 1$ admits a finitely supported sequence $y = (y^t)_{t\in\mathbb{Z}}$ with*

$$\|\Delta y\|_{\ell_2}^2 \geq \sum\nolimits_{t=1}^{T} \left\|\Delta u^t\right\|^2 - C, \tag{25}$$

$$\|y\|_{\ell_2}^2 \le \sum\nolimits_{t=1}^{T+1} \left\|u^t\right\|^2, \tag{26}$$

$$\|\Delta^m y\|_{\ell_2} \le A \sum\nolimits_{k=0}^{m-1} \eta^{m-k} \left\|\Delta^k y\right\|_{\ell_2} + C\eta^{m-1/2}. \tag{27}$$

*Proof.* For short horizons, we pay for the whole segment directly. Otherwise, we keep a central interval unchanged and taper over $O(1/\eta)$ indices at both ends; a buffer before the first taper avoids the exceptional initial first-order difference. Choose a constant $C_0$ depending only on $m, A, M$. If $T \le C_0/\eta$, take $y^t = 0$ for every $t$. The $k = 0$ case of Equation (24) gives $\|u^t\| \le M$. Hence $\left\|\Delta u^1\right\| \le 2M$, while the $k = 1$ case gives $\|\Delta u^t\| \le M\eta$ for every $t \ge 2$. Therefore

$$\sum_{t=1}^{T} \left\|\Delta u^t\right\|^2 \le 4M^2 + C_0 M^2 \eta \le C.$$

The two norm comparisons follow, and Equation (27) is immediate because $y = 0$.

Suppose now that $T > C_0/\eta$. Put $L := \lceil 1/\eta \rceil$. Choose a fixed smooth function $\psi : \mathbb{R} \to [0, 1]$ that is zero on $(-\infty, 0]$, one on $[1, \infty)$, and has bounded derivatives. Define

$$\chi^t := \psi\left(\frac{t - 1 - 3m}{L}\right) \psi\left(\frac{T + 1 - 3m - t}{L}\right), \tag{40}$$

set $y^t := \chi^t u^t$ for $1 \le t \le T + 1$, and set $y^t := 0$ otherwise. The construction keeps $y^t = u^t$ on the central interval

$$1 + 3m + L \le t \le T + 1 - 3m - L$$

and is zero before and after the two intervals on which $\chi^t$ changes. Let $B_T$ be the set of starting indices $t$ for which $\{t, \dots, t + 2m\}$ intersects either of these two intervals. Increasing $C_0$ if necessary makes the central interval nonempty and gives

$$|B_T| \le C(L + m) \le C/\eta.$$

Repeated use of

$$\Delta\psi(t/L) = \int_{t/L}^{(t+1)/L} \psi'(s)\, \mathrm{d}s$$

and the finite-difference product rule gives $\sup_t |\Delta^\ell \chi^t| \le C_{m,\ell} \eta^\ell$ for $0 \le \ell \le m$. Moreover,

$$\Delta^j (\chi u)^t = \sum_{\ell=0}^{j} \binom{j}{\ell} (\Delta^\ell \chi)^t (\Delta^{j-\ell} u)^{t+\ell}. \tag{41}$$

The buffer of length $3m$ before the first transition ensures that every finite difference of $u$ of positive order in Equation (41) begins at an index covered by Equation (24). Thus, whenever $t \in B_T$,

$$\left\|\Delta^j y^{t+r}\right\| \le C\eta^j, \qquad 0 \le j \le m, \quad 0 \le r \le 2m - j. \tag{42}$$

We retain the original recurrence wherever the cutoff leaves the trajectory unchanged and collect its failure into a boundary residual. For each $t$, define

$$\widetilde{c}_{k,t} := \begin{cases} c_{k,t}, & t \ge 1 \text{ and } y^{t+r} = u^{t+r} \text{ for every } 0 \le r \le 2m, \\ 0, & \text{otherwise,} \end{cases}$$

and define $e = (e^t)_{t \in \mathbb{Z}}$ by

$$\Delta^m y^{t+m} = \sum_{k=0}^{m-1} \widetilde{c}_{k,t} \eta^{m-k} \Delta^k y^{t+k} + e^t. \tag{43}$$

If $t \notin B_T$, the values $y^t, \dots, y^{t+2m}$ either agree with the corresponding values of $u$ or are all zero. In the first case Equation (23) gives $e^t = 0$; in the second case this is immediate. Hence $e^t = 0$ for every $t \notin B_T$. For $t \in B_T$, Equation (42) bounds every term in Equation (43) by $C\eta^m$. Therefore

$$\|e\|_{\ell_2} \le C\eta^m \sqrt{|B_T|} \le C\eta^{m-1/2}. \tag{44}$$

Taking $\ell_2$ norms in Equation (43), using $|\widetilde{c}_{k,t}| \leq A$, shift invariance, and Equation (44), yields Equation (27).

Finally, $0 \leq \chi^t \leq 1$ gives Equation (26). For $1 + 3m + L \leq t \leq T - 3m - L$, both $t$ and $t+1$ lie in the central interval, so $\Delta y^t = \Delta u^t$. Among $2 \leq t \leq T$, at most $C(L+m) \leq C/\eta$ indices lie outside this range, including the last index of the central interval. For these indices, Equation (24) gives $\|\Delta u^t\| \leq M\eta$, so their squared contribution is at most $C\eta$. The initial difference satisfies $\left\|\Delta u^1\right\| \leq 2M$. These bounds prove Equation (25). □