Constructive Lean 4 Algebraic Kernel and Exact-Rational Certificate Auditor for Sub-$\sqrt{T}$ Linear Swap Regret in Polyhedral Games

Lluis Eriksson*

November 25, 2026

Abstract

Vasilis Pollatos and Alberto Marchesi (`arXiv:2611.33081`, November 24, 2026) broke the $\sqrt{T}$ cumulative regret barrier for linear swap regret in self-play over polyhedral games: their deterministic uncoupled optimistic log-barrier dynamics achieve $\tilde{O}(\mathrm{poly}(D)T^{1/3})$ cumulative linear Φ-regret on arbitrary compact H-polytopes (independent of vertex count) and $O(n^2M^4\log(T+1))$ in n-player perfect-recall extensive-form games (EFGs) with at most M terminal sequences per player, extending logarithmic trigger-regret rates to all untimed-communication (UTC) linear deviations. We present a companion zero-floating-point verification suite combining an 18-theorem zero-dependency Lean 4.35.0 kernel (`LinearSwapRegretPolyhedral.lean`; 0 `sorry`, 0 warnings, 0 `Classical.choice`) with a pure-stdlib exact-$\mathbb{Q}$ runtime certificate auditor (`SWAP-POLYHEDRAL-CERT-1.0`). The Lean 4 kernel mechanizes cleared-integer scalar inequalities for Lemma 3, Theorem 4, Eq. (13), and Appendix G, exact 2×2 adjugate Perron fixed-point and monotonicity identities, finite $r \le 6$ exponential-factor checks, inductive Cauchy–Schwarz, and an exhaustive 16-trigger proof of the Student–Professor UTC gain $3T/2$ (explicitly leaving general-m matrix Perron theory, FTRL regret dynamics, and recursive treeplex induction outside Lean 4). Complementarily, `SWAP-POLYHEDRAL-CERT-1.0` audits concrete $m\times m$ instances across four exact-$\mathbb{Q}$ modes using Gaussian determinants, cofactor adjugates, and certified 5-term Mercator Bregman brackets, replaying 16 canonical certificates in 0.192–6.384 ms and rejecting 62/62 live adversarial forgeries.

1 Introduction

In online learning in games, minimizing linear Φ-regret [4, 3] over compact polyhedral strategy sets drives play toward linear correlated equilibria, which in perfect-recall extensive-form games (EFGs) correspond exactly to untimed-communication (UTC) deviations [6]. Prior efficient algorithms scaled as $O(\sqrt{T})$, while logarithmic self-play rates [1] were restricted to trigger deviations. Pollatos and Marchesi [5] broke both barriers via three structural pillars whose notation we fix explicitly:

1. *Self-play energy cancellation (Theorem 4):* For strictly positive $m\times m$ column-stochastic lifts A, B, define the scalar log-barrier Bregman generator $g(u) := u - 1 - \log u$ $(u > 0)$ and entry-wise log-barrier Bregman divergence $D_R(B, A) := \sum_{a,b=1}^{m} g(B_{ab}/A_{ab}) \ge 0$. For reset rate $\lambda \in [0, 1)$, let x_A^λ solve the Gordon reset fixed-point equation $x_A^\lambda = (1-\lambda)Ax_A^\lambda + \lambda x_0$ with total-variation norm $\|s_B - s_A\|_P := \frac{1}{2}\|s_B - s_A\|_1 = \|s_B - s_A\|_{\mathrm{TV}}$. Coupling optimistic log-barrier FTRL with fixed-point stability $\|x_B^\lambda - x_A^\lambda\|_i^2 \le C_* D_{R_i}(B, A)$ cancels path-length prediction error whenever $\gamma^2 \ge 128n^2C_*$, yielding

$$\sum_i \mathrm{LinReg}_i(T) \le \gamma\left(\sum_i m_i^2\right)\log(T+1) + \frac{4n}{\gamma} + 2n(\lambda T + 1). \tag{1}$$

On general H-polytopes, reset contraction yields $C_*(\lambda) = O(\mathrm{poly}(D)/\lambda)$ and hence $\gamma = \Theta(n\sqrt{C_*}) = \Theta(\mathrm{poly}(D)\lambda^{-1/2})$; balancing $\lambda^{-1/2}$

*Independent researcher; manuscript prepared with transparent AI assistance via Google Antigravity. Source code, Lean 4 proofs, and exact certificates: https://github.com/lluiseriksson/linear-swap-regret-polyhedral-certificates.

against λT in (1) at $\lambda = T^{-2/3}$ yields $\tilde{O}(\mathrm{poly}(D)T^{1/3})$.

2. *Adjugate Perron monotonicity & treeplex stability (Lemmas 7–8, Theorem 9):* Given two $m \times m$ column-stochastic treeplex lift matrices A_+, A_- (with $U = (1-\epsilon)\tilde{A}$ for the convex combination $\tilde{A} = \frac{\alpha A_+ + \beta A_-}{\alpha+\beta}$ and $\epsilon \in (0,1)$), define the homogeneous degree-$(m-1)$ Laplacian pencil $L(\alpha,\beta) := \alpha(I - A_+) + \beta(I - A_-)$ and trace denominator $a(\alpha,\beta) := \mathrm{tr}(\mathrm{adj}\, L(\alpha,\beta))$. The resolvent trace inequality $\mathrm{tr}((I-U)^{-1}(I-A_\pm)) \geq 0$ implies $\partial_\alpha a, \partial_\beta a \geq 0$, which via Bernstein's inequality and $(5/2)^2 \cdot 32 = 200$ yields the $\lambda = 0$ treeplex stability bound (with $C_* = 200M^4$ and $\gamma = 320nM^2$):
$$\|s_B - s_A\|_P^2 \leq 200m^4 D_R(B,A), \tag{2}$$
giving $O(n^2M^4 \log(T+1))$ cumulative linear swap regret in n-player EFGs.

3. *UTC vs. trigger separation (Section 1, p. 3):* In the 4-state Student–Professor EFG, each cycle visits 4 equiprobable recommendation pairs $(g,b) \in \{0,1\}^2$ with payoff $4 \cdot \mathbf{1}[a_g = b] + 3 \cdot \mathbf{1}[a_b = g]$. Every one of the 16 pure block-diagonal trigger deviations $(a_g, a_b) = (f_g(g), f_b(b))$ achieves cycle sum 14 (expected payoff 7/2 per cycle, $7T/2$ over T cycles, zero trigger regret), whereas the UTC linear deviation $(a_g, a_b) = (g, g)$ (Bad copies Good's recommendation) achieves cycle sum 20 (expected payoff $5T$)—a strict UTC swap gain of $3T/2$.

Contributions and Explicit Scope. We mechanize the finite scalar/algebraic invariants in Lean 4 while verifying concrete general-m matrix instances in a pure-stdlib exact-$\mathbb{Q}$ Python auditor. We explicitly scope our Lean 4 kernel *away* from FTRL/RVU regret dynamics, general $m \times m$ Perron monotonicity, continuous-interval log-barrier calculus, Lemma 5's Legendre dual, and Lemma 7's recursive treeplex-norm induction:

1. **Zero-dependency Lean 4 algebraic kernel (§2):** 18 constructive theorems in `LinearSwapRegretPolyhedral.lean` (Lean 4.35.0 [2]), 0 `sorry`, 0 warnings, and 0 `Classical.choice`, with exact formalized statements reported in Table 1.
2. **Pure-stdlib exact-$\mathbb{Q}$ verifier (§3):** `SWAP-POLYHEDRAL-CERT-1.0`, auditing finite $m \times m$ adjugate Perron traces, Theorem 9 stability via certified Mercator Bregman brackets, reset-RVU cancellation, and finite EFG UTC-vs-trigger separation with zero floats.
3. **Empirical atlas & red-team audit (§4):** 16 canonical certificates (0.192–6.384 ms), exact Perron trace profiles, Theorem 9 headroom curves, closed-form theoretical regret-rate scaling curves, and 62/62 rejected live adversarial forgeries.

2 Zero-Dependency Lean 4 Kernel

We formalize 18 constructive theorems in `lean/LinearSwapRegretPolyhedral.lean` (Lean 4.35.0, zero external libraries), audited via `lean/AuditAxioms.lean` with 0 warnings, 0 `sorry`, 0 `Classical.choice`, and 0 `Lean.ofReduceBool` (Table 1: 1 zero-axiom, 3 `[propext]`, and 14 `[propext, Quot.sound]`).

Proposition 1 (Constructive Linear-Swap Algebraic Kernel)**.**

1. Self-play RVU & rate-balance scalar bounds (Theorems 1–5, 17–18): *Path-length bridge* $(a+b)^2 \leq 2(a^2+b^2)$*; energy cancellation* $(128n^2C_* - \gamma^2)E \leq 0$*; Eq. (13) choice* $\gamma = 320nM^2$ *attaining* $\gamma^2 = 4(128n^2 \cdot 200M^4)$ *and* $4n/\gamma = 1/(80M^2)$*; inductive nonnegativity bound* $x \in xs \Rightarrow x \leq B$*;* $T^{1/3}$ *balance* $T = k \cdot k^2$ *with* $k^2 < T$ *for* $k \geq 2$*; and scalar reset-contraction step* $q\,d \leq (q-p)(d+\delta) \Rightarrow p\,d \leq (q-p)\delta$*.*
2. Appendix G stability constants & two-sided Mercator check (Theorems 6–10, 14): *Quotient bound* $4r^2 + r \leq 5r^2$*; segment floor* $2((D-t)A + tB) \geq DA$*; exact constant* $25 \cdot 32 = 4 \cdot 200$ *and large-divergence bound* $16 \leq 200m^4$*; two-sided Mercator arithmetic at endpoints* $u \in \{\frac{1}{2}, 2\}$*: the 2-term lower bound* $2(\frac{1}{3} + \frac{1}{81}) = \frac{56}{81} < \log 2$ *certifies* $g(\frac{1}{2}) = \log 2 - \frac{1}{2} > \frac{31}{162} > \frac{1}{16}$*, while the 3-term plus* $k = 3$ *geometric-tail upper bound* $\log 2 < \frac{23581}{34020} < \frac{11}{15}$ *certifies* $g(2) = 1 - \log 2 > \frac{4}{15} > \frac{1}{16}$*; Taylor denominator bound* $2\xi^2 \leq 8D^2$*; and* $(2r+1)^r < 2(2r)^r$ *for the six base cases* $r = 1 \ldots 6$ *plus* $(2r+1)/(2r) \leq 3/2$ *for all* $r \geq 1$*.*
3. Exact 2×2 adjugate Perron, Cauchy–Schwarz, & 16-trigger UTC gap (Theorems 11–13, 15–16): *Exact* 2×2 *column-stochastic fixed point* $A \cdot (q,p) = D \cdot (q,p)$ *and monotone denominator* $\alpha(p_1 + q_1) + \beta(p_2 + q_2)$*; Euler monotone derivative bound* $|d_\alpha - d_\beta| \leq r\rho_0$ *from* $d_\alpha + d_\beta = r\rho_0$ $(d_\alpha, d_\beta \geq 0)$*; inductive Cauchy–Schwarz* $(\sum d_j)^2 \leq |ds| \sum d_j^2$*; and exhaustive case-split proof over all 16 pure Student–Professor triggers (*14 *vs. UTC* 20*, gain* $6T$ *over* $4T$ *states* $= 3T/2$*).*

Table 1: Complete 18-theorem Lean 4.35.0 kernel in `LinearSwapRegretPolyhedral.lean` (0 `sorry`, 0 warnings, 0 `Classical.choice`), with exact formalized statements and `#print axioms` footprints.

#	Lean 4 Theorem Identifier	Exact Formalized Statement (Cleared Integer / Inductive Type)	Axioms
1	`path_length_bridge_ineq`	$\forall a, b : \mathbb{Z},\ (a+b)^2 \le 2(a^2+b^2)$	`[p, Q]`
2	`self_play_energy_cancellation`	$0 \le E \wedge 128n^2C_* \le \gamma^2 \Rightarrow (128n^2C_* - \gamma^2)E \le 0$	`[p, Q]`
3	`treeplex_eq13_parameter_choice`	$\gamma = 320nM^2, C_* = 200M^4 \Rightarrow \gamma^2 = 4(128n^2C_*) \wedge 4n \cdot 80M^2 = \gamma$	`[p]`
4	`individual_regret_le_sum`	$x \in xs \wedge$ `AllNonnegInt` $xs \wedge$ `sumListInt` $xs \le B \Rightarrow x \le B$	`[p]`
5	`t_one_third_exponent_balance`	$2 \le k \Rightarrow k^3 = k \cdot k^2 \wedge k^2 < k^3$	`[p, Q]`
6	`quotient_rule_five_halves_bound`	$1 \le r \Rightarrow 4r^2 + r \le 5r^2$	`[p, Q]`
7	`segment_entry_half_lb`	$0 \le A, 0 \le t \le D, A \le 2B \Rightarrow DA \le 2((D-t)A + tB)$	`[p, Q]`
8	`theorem9_constant_200_and_diameter_branch`	$1 \le m \Rightarrow 25 \cdot 32 = 4 \cdot 200 \wedge (m-1)^4 \le m^4 \wedge 16 \le 200m^4$	`[p, Q]`
9	`log_barrier_g_threshold_1_16`	Two-sided Mercator arithmetic: $\frac{56}{81} - \frac{1}{2} = \frac{31}{162} > \frac{1}{16}$ & $1 - \frac{23581}{34020} > \frac{4}{15} > \frac{1}{16}$	`none`
10	`log_barrier_second_deriv_floor`	$0 \le \xi \le 2D \Rightarrow 2\xi^2 \le 8D^2$	`[p, Q]`
11	`perron_2x2_adjugate_exact`	$s_{01}+s_{02} = 1 \Rightarrow L\,\mathrm{adj}(L) = 0 \wedge \mathrm{adj}(L)s_0 = (q,p) \wedge A(q,p) = D(q,p)$	`[p, Q]`
12	`perron_2x2_monotone_denominator`	$p_i, q_i \ge 0, \alpha_1 \le \alpha_2, \beta_1 \le \beta_2 \Rightarrow a(\alpha_1,\beta_1) \le a(\alpha_2,\beta_2)$ (2×2)	`[p, Q]`
13	`euler_monotone_derivative_bound`	$0 \le d_\alpha, d_\beta \wedge d_\alpha + d_\beta = r\rho_0 \Rightarrow -r\rho_0 \le d_\alpha - d_\beta \le r\rho_0$	`[p, Q]`
14	`rho_interval_factor_lt_two`	$1 \le r \Rightarrow 2(2r+1) \le 3(2r)$, plus $(2j+1)^j < 2(2j)^j$ for $j \in \{1..6\}$	`[p, Q]`
15	`cauchy_schwarz_sum_sq`	$\forall ds :$ `List` $\mathbb{Z}$, (`sumListInt` $ds)^2 \le$ `lenInt` $ds \cdot$ `sumSqListInt` ds	`[p, Q]`
16	`student_professor_utc_vs_trigger_gap`	$\forall f_g, f_b : \mathrm{Bool} \to \mathrm{Bool}, T \ge 1 :$ Trig $= 14 \wedge$ UTC $= 20 \wedge \Delta \cdot T = 6T > 0$	`[p, Q]`
17	`reset_fixed_point_contraction`	$q\,d \le (q-p)(d+\delta) \Rightarrow p\,d \le (q-p)\delta$	`[p, Q]`
18	`efg_logarithmic_regret_constants`	$(320nM^2)(nM^2) = 320n^2M^4 \wedge 320nM^2 = (4n)(80M^2)$	`[p]`

Axiom key: `none` = zero axioms (1); `[p]` = `[propext]` (3); `[p, Q]` = `[propext, Quot.sound]` (14).

3 Exact-$\mathbb{Q}$ Verifier Architecture

`SWAP-POLYHEDRAL-CERT-1.0` (`src/swap_certifier.py`) is a pure-stdlib zero-float verifier using exact Gaussian determinants, classical cofactor adjugates, Gauss–Jordan inverses, and certified 5-term Mercator rational brackets for $\log(u)$ across four modes:

1. **Mode 1 (`perron_adjugate_monotone_cert`):** Independently computes $L(\alpha,\beta)$, $\mathrm{adj}\,L$, $a(\alpha,\beta) = \mathrm{tr}(\mathrm{adj}\,L)$, verifies $\det L = 0$, exact Gordon fixed point $\tilde{A}s = s$ for the normalized convex combination $\tilde{A} = \frac{\alpha A_+ + \beta A_-}{\alpha+\beta}$ (generalizing the 2×2 unnormalized equation $A(q,p) = D(q,p)$), monotonicity in (α,β), and Lemma 7's resolvent traces $\mathrm{tr}((I-U)^{-1}(I-A_\pm)) \ge 0$.
2. **Mode 2 (`treeplex_stability_theorem9_cert`):** Computes exact fixed points s_A, s_B, TV-norm squared $\|s_B - s_A\|_P^2$, certified rational interval $[D_R^{\mathrm{lb}}, D_R^{\mathrm{ub}}]$ containing $D_R(B,A)$, and verifies $\|s_B - s_A\|_P^2 \le 200(m-1)^4 D_R^{\mathrm{lb}} \le 200m^4 D_R^{\mathrm{lb}}$ (2).
3. **Mode 3 (`reset_stability_and_rvu_regret_cert`):** Solves $(I-(1-\lambda)M)x^\lambda = \lambda x_0$ exactly over $\mathbb{Q}$, certifies $\lambda\|x_B^\lambda - x_A^\lambda\|_{\mathrm{TV}} \le (1-\lambda)\|(B-A)x_A^\lambda\|_{\mathrm{TV}}$, checks $C_* = 200M^4$ and $\lambda T \ge 1$, and audits Theorem 4's energy margin $\gamma^2 - 128n^2C_* \ge 0$.
4. **Mode 4 (`utc_vs_trigger_efg_separation_cert`):** Audits finite treeplex flow partitions, enforces block-diagonal column-stochasticity on all supplied triggers, independently computes the exact global optimum over all $\prod_g |g|^{|g|}$ pure block-diagonal triggers to certify zero trigger regret, verifies vertex-wide invariance of the UTC map, and certifies strict positive UTC swap gain.

4 Empirical Atlas and Red-Team Audit

Running `scripts/run_all_experiments.py` produces 16 canonical certificates (4 per mode) and four exact-$\mathbb{Q}$ studies, audited end-to-end by `scripts/verify_all.py` with SHA-256 verification across all 41 tracked artifacts.

Adjugate Perron Profile and Stability Headroom (Fig. 1). On the 4×4 coupled lift ($\rho(0) = 554/125$, degree $r = 3$), $\rho(t)/\rho(0)$ varies monotonically from 1.554621 at $t = -9/10$ to 0.579892 at $t = 9/10$, staying well below the Appendix G ceiling 2 across the entire Bernstein window $|t| \le 1/6$ (and even on $[-9/10, 9/10]$). On the asymmetric 3×3 perturbation family $B_\tau = (1-\tau)A + \tau B$ ($\tau = 2^{-1}\ldots 2^{-6}$), exact certified Mercator brackets yield empirical stability ratios $\|s_B - s_A\|_P^2 / D_R^{\mathrm{lb}}(B,A)$ rising smoothly from 0.028791 to 0.053589—over four orders of magnitude below Theorem 9's worst-case constant $200(m-1)^4 = 3200$ ($200m^4 = 16200$).

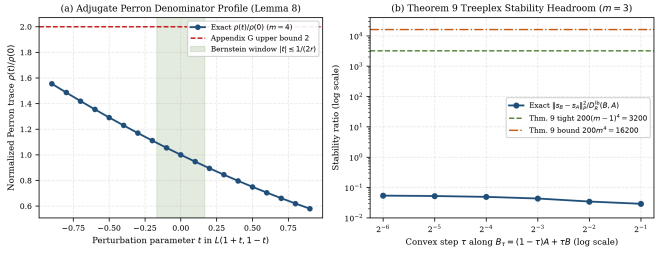

Figure 1: (a) Exact adjugate Perron trace $\rho(t)/\rho(0)$ for a 4×4 treeplex lift ($\rho(0) = 554/125$) staying strictly below 2 throughout the Bernstein window $|t| \le 1/(2r) = 1/6$. (b) Empirical Theorem 9 stability ratio $\|s_B - s_A\|_P^2 / D_R^{\mathrm{lb}}(B, A)$ on an asymmetric 3×3 Markov family across $\tau = 2^{-1} \dots 2^{-6}$ vs. theoretical ceilings $200(m-1)^4 = 3200$ and $200m^4 = 16200$.

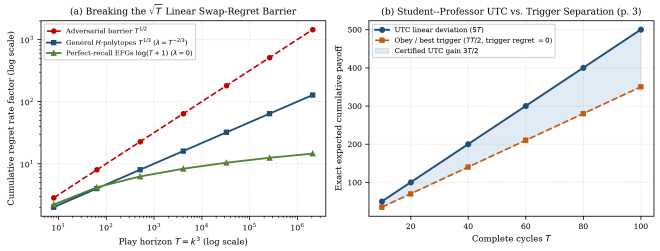

Figure 2: (a) Closed-form theoretical regret-rate curves across horizons $T = k^3$ ($k = 2 \dots 128$): adversarial $T^{1/2}$ vs. general H-polytope $T^{1/3}$ ($\lambda = 1/k^2$) vs. perfect-recall EFG $\log(T+1)$ ($\lambda = 0$). (b) Exact Student–Professor cumulative payoff ($5T$ UTC vs. $7T/2$ obey/trigger, gain $3T/2$).

Sub-$\sqrt{T}$ Scaling and Student–Professor UTC Gap (Fig. 2). Mode 3 (0.192–0.637 ms) audits strictly nonzero reset contractions and RVU cancellation across horizons $T \in \{64, 125, 1000, 8000\}$, including the sharp zero-margin boundary $\gamma^2 = 1280^2 = 128 \cdot 2^2 \cdot 3200$. Mode 4 (2.893–6.384 ms) independently certifies zero global trigger regret over all $\prod_g |g|^{|g|}$ pure block-diagonal triggers on the exact Page 3 Student–Professor separation (35 and 350 for $T = 10, 100$ vs. UTC 50 and 500, gain 15 and 150), plus bidirectional and 3-signal (6-sequence) cross-type UTC separations.

Replay Atlas and Red-Team Audit (Fig. 3, Table 2). All 16 canonical certificates replay in 0.192–6.384 ms with zero floats (and strictly positive fixed-point shifts on every Mode 2/3 instance). Our adversarial suite (`tests/test_suite.py`) tests

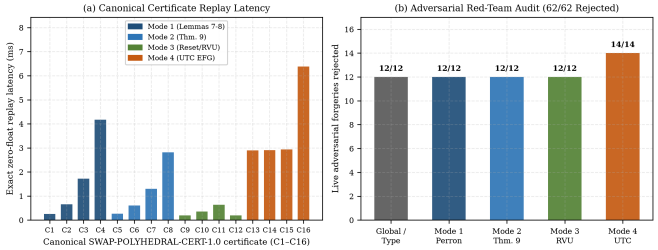

Figure 3: (a) Exact zero-float replay latency across all 16 canonical `SWAP-POLYHEDRAL-CERT-1.0` certificates (0.192–6.384 ms). (b) Red-team adversarial audit rejecting 62/62 live forgery attacks across five attack classes.

Table 2: Summary of all 16 canonical `SWAP-POLYHEDRAL-CERT-1.0` certificates and 62/62 live adversarial red-team audit.

Verification Mode	Certs	Latency (ms)	Attacks Rejected
Mode 1: Lemmas 7–8 Adjugate Perron	4	0.255–4.176	12/12
Mode 2: Thm. 9 Treeplex Stability	4	0.260–2.812	12/12
Mode 3: Reset Contraction & RVU	4	0.192–0.637	12/12
Mode 4: UTC vs. Trigger EFG Gap	4	2.893–6.384	14/14
Global Schema / Type / Rational	–	–	12/12
Total	**16**	**0.192–6.384**	**62/62**

62 live forgery attacks—including float/`None`/bool injection, unreduced fractions, oversized rational-string and integer bit-length DoS, non-stochastic or zero-sum lift matrices, broken energy cancellation ($\gamma^2 < 128n^2C_*$), cross-group or non-stochastic triggers, hidden omitted-trigger regret bypasses, infeasible UTC vertex maps, and nonzero trigger-regret instances—achieving **62/62 rejections** with zero bypasses.

5 Conclusion

We delivered an 18-theorem zero-dependency Lean 4 kernel (0 `sorry`, 0 warnings, 0 `Classical.choice`) and the exact-$\mathbb{Q}$ verifier `SWAP-POLYHEDRAL-CERT-1.0` for the adjugate-Perron, treeplex-stability, RVU energy-cancellation, and UTC-vs-trigger separation core of Pollatos and Marchesi [5]. Mechanizing the full recursive treeplex-norm induction of Lemma 7 and the exponential-family Legendre comparison of Lemma 5 in Lean 4 remain compelling targets for future formalized algorithmic game theory.

References

[1] Ioannis Anagnostides, Gabriele Farina, Christian Kroer, Chung-Wei Lee, Haipeng Luo, and Tuomas Sandholm. Near-optimal no-regret learning dynamics for general convex games. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, 2023.

[2] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, pages 625–635, 2021.

[3] Gabriele Farina and Charilaos Pipis. Polynomial-time linear-swap regret minimization in imperfect-information sequential games. In *Advances in Neural Information Processing Systems (NeurIPS)*, volume 36, 2023.

[4] Geoffrey J. Gordon, Amy Greenwald, and Casey Marks. No-regret learning in convex games. In *Proceedings of the 25th International Conference on Machine Learning (ICML)*, pages 360–367, 2008.

[5] Vasilis Pollatos and Alberto Marchesi. Breaking the square-root barrier for linear swap regret in polyhedral games. arXiv:2611.33081 [cs.GT], November 2026.

[6] Brian Hu Zhang, Gabriele Farina, Ioannis Anagnostides, Federico Cacciamani, Stephen McAleer, Andreas Haupt, Andrea Celli, Nicola Gatti, Vincent Conitzer, and Tuomas Sandholm. Mediator interpretation and faster learning algorithms for linear correlated equilibria in general sequential games. In *International Conference on Learning Representations (ICLR)*, 2024.