Exponentially Sharper Rational LP Gaps, Effective Payoff-Rank Support Compression, and Sharp b-ary Matching-Tree PPAD-Hardness for Extensive-Form NFCE

Lluis Eriksson
Independent Researcher
https://github.com/lluiseriksson/nfce-extensive-form-sharp-gap-and-support

November 27, 2026

Abstract

Cheval, Horn, Paul, and Shirmohammadi (arXiv:2611.34897v1, 2026) recently resolved a major open problem in algorithmic game theory by placing exact normal-form correlated equilibrium (NFCE) computation in perfect-recall extensive-form games inside PPAD for arbitrary numbers of players, with output support at most $K = 1 + \sum_{i=1}^{n} d_i(d_i + 1)$, and proving PPAD-completeness for two-player games under the quadratic support bounds $\kappa_d = (d_1 + 1)(d_2 + 1)$ and $\kappa_+ = \max\{\kappa_d, \lfloor |V|^2/36 \rfloor\}$. We strengthen every quantitative pillar of their framework and resolve structural optimality questions left open by their analysis. First (**Theorem 1**), we prove an exponentially sharper uniform rational LP gap by replacing their tree-wide denominator product $B = \prod_{r \in \mathcal{R}} \mathrm{den}(r)$ with the exact path-LCM denominator $L = \mathrm{lcm}_{z,i} \mathrm{den}(p_c(z)u_i(z))$ (cutting denominator bit-length $\log_2 B$ by 632×–884× on depth-10 binary chance trees) and applying Cramer's rule with an unscaled simplex row and column-norm Hadamard estimation on the active support of size $r \le K$: every nonzero LP optimum satisfies $t^*(S) \ge \left(\sqrt{r}\,(1 + rL^2)^{r/2} L\right)^{-1}$, halving the dimension exponent from $(K + 1)^{K+1}$ to $K^{1/2}(1 + KL^2)^{K/2}$ and reducing the total oracle bit-precision $\log_2(1/\epsilon_\Gamma)$ by 4.91×–26.96× on certified chance-tree benchmarks up to depth $h = 6$ (and asymptotically by $\Theta(|Z|/h)$). Second (**Theorem 2**), we prove the exact two-player support collapse $2(K - \kappa_d) = d_1^2 + d_2^2 + (d_1 - d_2)^2 \ge 0$ (sharp on matching trees) and introduce the polynomial-time computable effective payoff-deviation rank $r_i \le d_i$, sharpening the general n-player Carathéodory support bound to $K_{\mathrm{eff}} = \min\left(\prod_i (d_i + 1),\ 1 + \sum_i (d_i + 1) r_i\right) \le K$. Third (**Theorem 3**), we remove the power-of-two restriction $m = 2^k$ via arbitrary-m balanced binary matching trees ($|V| = 6m - 5$, $N = 3m - 2$, $d_i = m - 1$ for every $m \ge 2$), slashing the PPAD-hardness tree size by 8.4×–15.9× under κ_d and 43.8×–86.5× under κ_+; we prove exact optimality of the node slope 6 uniquely at branching factors $b \in \{2, 3\}$ (with ternary trees $b = 3$ halving the number of information sets!) and unique leaf-count optimality $3m - 2$ at $b = 2$, yielding PPAD-completeness under the leaf-count bound $\kappa_N = \lfloor N^2/9 \rfloor$ already for all $r \ge 3$; and we exhibit two explicit structural barriers showing why the matching-tree decoder breaks at support $2m^2$ ($c = 1/18 > 1/36$, constant Nash regret 1/12) and why chance-selected parallel matching pairs cannot amplify support lower bounds past m^2. All results are backed by 20 constructive Lean 4.36.0 theorems (zero `sorry`, zero `Classical.choice`), 16 AST-audited zero-float SHA-256 certificates, and 56 adversarial tests.

AI Disclosure. Google Antigravity assisted with symbolic determinant bounds, exact rational Carathéodory/BFS verification, Lean 4 formalization, and manuscript drafting. All mathematical proofs, exact rational certificates, Lean proofs, and final claims were audited and verified by the author, who takes full responsibility for this work.

1 Introduction

Computing an exact normal-form correlated equilibrium (NFCE) in a finite, multiplayer extensive-form game with imperfect information and perfect recall had long stood open between polynomial-time extensive-form correlated equilibrium (EFCE) computation [7, 4, 5] and PPAD-complete Nash equilibrium computation [2, 3]. Cheval, Horn, Paul, and Shirmohammadi [1] recently resolved this problem via a two-stage PPAD algorithm (Theorem 1, Corollary 2, Theorem 6, Theorem 7, and Lemma 14 of [1]):

1. Invoke the PPAD weak approximate-Nash oracle of Etessami [3] at binary accuracy $\epsilon_\Gamma := \frac{1}{2}((K+1)B)^{-(K+1)}$, where $K = 1 + \sum_{i=1}^{n} d_i(d_i+1)$, $d_i = \sum_{I \in \mathcal{I}_i}(|A(I)| - 1)$, and $B = \prod_{r \in \mathcal{R}} \mathrm{den}(r)$ is the product of all input chance-edge and normalized payoff denominators across the entire tree;

2. Decompose each player's behavioral strategy into a realization-equivalent mixed strategy α_i of support $m_i \le d_i + 1$ [6], compress the independent product $\bigotimes_i \alpha_i$ via sequential prefix-expansion and Carathéodory BFS into an unnormalized ϵ_Γ-NFCE μ_0 on at most K pure-plan profiles (Algorithm 1 and Theorem 7 of [1]), and solve the compact sequence-form reweighting LP (Lemma 13 of [1]) at $t = 0$, whose feasibility is guaranteed by the uniform rational LP gap $t^*(S) \in \{0\} \cup [\eta_{\mathrm{CHPS}}, 1]$ (Lemma 14 and Appendix B of [1]).

Complementing this upper bound, Theorem 3 and Corollary 20 of [1] prove PPAD-completeness for two-player games under the prescribed support bounds $\kappa_d = (d_1+1)(d_2+1)$ and $\kappa_+ = \max\{\kappa_d, \lfloor |V|^2/36 \rfloor\}$, by combining a source $r \times r$ bimatrix game with a depth-k binary matching tree on $m = 2^k$ labels ($m \ge 8r^3$ for κ_d; $m \ge 8r^4$ and $r \ge 5$ for κ_+).

Four sharp quantitative and structural questions arise immediately from [1]:

1. *Oracle precision bottleneck:* Multiplying every chance-edge and payoff denominator across the entire tree into $B = \prod_{r \in \mathcal{R}} \mathrm{den}(r)$ inflates $\log_2(1/\epsilon_\Gamma)$ by a factor proportional to the tree size $|Z|$, and scaling the simplex row and column t by B doubles the dimension exponent in Hadamard's inequality. How sharply can the uniform rational LP gap η be tightened?

2. *Support-size tightness:* Can the general support bound $K = 1 + \sum_i d_i(d_i+1)$ be compressed using the payoff-weighted deviation subspace without breaking polynomial-time prefix compression?

3. *Hardness tree size and branching optimality:* Is $m = 2^k$ necessary? Can the $8r^3$ and $8r^4$ blowups be eliminated? Does a b-ary matching tree ($b \ge 3$) beat the constant $1/36$ in $\kappa_V = \lfloor |V|^2/36 \rfloor$, or is $1/36$ optimal across all b-ary trees? What happens for leaf-count bounds κ_N?

4. *Structural barriers:* Why does the matching-tree entropy decoder fail past $c = 1/36$, and can parallel chance-selected matching branches amplify the m^2 support lower bound to m^{2q}?

Table 1 summarizes our complete answers to all four questions.

2 Main Theorems

We follow the exact extensive-form notation of [1, Section 2]: Γ is a finite n-player perfect-recall game tree with node set V, terminal set Z ($N = |Z|$), chance-path probabilities $p_c(z) \in \mathbb{Q}_{>0}$, normalized terminal payoffs $u_i(z) \in \mathbb{Q} \cap [0,1]$, player-$i$ information sets $\mathcal{I}_i$, action sets $A(I)$, realization-polytope affine dimension $d_i = \sum_{I \in \mathcal{I}_i}(|A(I)| - 1)$, pure-plan compatibility indicator $\chi_i(\sigma_i, z) \in \{0,1\}$, and unnormalized deviation gain $R_{i,\sigma_i,\sigma_i'}(\mu) = \sum_{\tau : \tau_i = \sigma_i} \mu(\tau)[U_i(\sigma_i', \tau_{-i}) - U_i(\tau)]$.

Table 1: Quantitative and structural improvements over Cheval, Horn, Paul, and Shirmohammadi (`arXiv:2611.34897v1`) [1]. All claims are backed by 20 Lean 4.36.0 theorems and 16 zero-float rational certificates.

Quantity / Property	Cheval et al. [1]	Present Work
Denominator scale in LP gap	Tree-wide product $B = \prod_{r\in\mathcal{R}} \operatorname{den}(r)$	Path-LCM $L = \operatorname{lcm}_{z,i} \operatorname{den}(p_c(z)u_i(z)) \mid B$
Uniform LP gap lower bound	$((K+1)B)^{-(K+1)}$ (Lem. 14, App. B)	$\left(\sqrt{K}(1+KL^2)^{K/2}L\right)^{-1}$ (Thm. 1)
Oracle bit budget $\log_2(1/\epsilon_\Gamma)$	$\Theta(K \cdot \lvert Z\rvert \log D)$ on depth-h trees	$\Theta(K(h\log D + \log K))$ (26.96× at $h=6$; $\log_2 B/\log_2 L$: 632×–884× at $h=10$)
General n-player support bound	$K = 1 + \sum_i d_i(d_i+1)$ (Thm. 1, Thm. 7)	$K_{\text{eff}} = \min(\prod_i (d_i+1), 1+\sum_i (d_i+1)r_i)$
2-player gap K vs. κ_d	Implicit $K \geq \kappa_d$ (Sec. 4.5)	Exact $2(K-\kappa_d) = d_1^2 + d_2^2 + (d_1-d_2)^2$
Hardness matching labels m	Power of two $m = 2^k$ only	Arbitrary integer $m \geq 2$ ($\lvert V\rvert = 6m-5$ exact)
PPAD-hardness size (κ_d / κ_+)	$m \geq 8r^3$ ($r \geq 2$) / $m \geq 8r^4$ ($r \geq 5$)	8.4×–15.9× / 43.8×–86.5× smaller $\lvert V\rvert$
b-ary matching tree optimality	Binary $b = 2$ only	Node slope 6 optimal iff $b \in \{2,3\}$ ($b=3$ halves $\lvert\mathcal{I}_i\rvert$)
Leaf-count bound κ_N	Not analyzed	$\kappa_N = \lfloor N^2/9 \rfloor$ PPAD-complete for all $r \geq 3$
Barriers past 1/36 and m^2	Open (Sec. 5)	Explicit $2m^2$ NFCE ($c = 1/18$) & $2q$-player m^2 NFCE

For any nonempty finite support $S \subseteq \Sigma$, $t^*(S)$ denotes the optimal value of the conceptual reweighting LP [1, Eq. (11)].

Theorem 1 (Exponentially Sharper Uniform Rational LP Gap and Exact Realization)**.** *Assume terminal payoffs are normalized to* $[0,1]$*. Define the playerwise and global path-LCM denominators*

$$L_i := \operatorname{lcm}_{z\in Z} \operatorname{den}\big(p_c(z)u_i(z)\big), \qquad L := \operatorname{lcm}_{i\in[n]} L_i. \tag{1}$$

Then:

(a) Divisibility and exponential separation. *L divides the tree-wide denominator product $B = \prod_{r\in\mathcal{R}} \operatorname{den}(r)$ of [1, Theorem 6]. On a two-player (or $n \geq 2$-player) depth-h binary chance tree ($N = 2^h$ leaves) with fair coins and constant payoff denominator $D \geq 2$, $\log_2 B \geq (2N-2) + 2N\log_2 D$ whereas $\log_2 L \leq h + \log_2 D$ (e.g., at $h = 10, D = 10$: 8850 bits vs. 14 bits, a 632× reduction; `cert02`).*

(b) Unscaled-simplex column-Hadamard uniform LP gap. *For every nonempty finite support $S \subseteq \Sigma$ with $\lvert S\rvert \leq K$, let $r \leq \lvert S\rvert$ be the active support size $\lvert\operatorname{supp}(\mu^*)\rvert$ of an optimal vertex (μ^*, t^*) of [1, Eq. (11)]. Then either $t^*(S) = 0$ or*

$$t^*(S) \geq \eta_{\text{sharp}}(r,L) := \frac{1}{\sqrt{r}\,(1+rL^2)^{r/2}\,L} \geq \frac{1}{\sqrt{K}\,(1+KL^2)^{K/2}\,L} =: \eta_{\text{sharp}}(K,L). \tag{2}$$

Consequently, replacing ϵ_Γ in Theorem 6 of [1] by $\epsilon_\Gamma^* := \frac{1}{2}\eta_{\text{sharp}}(K,L)$ preserves exact reweighting while strictly improving upon $\eta_{\text{CHPS}} = ((K+1)B)^{-(K+1)}$ for every $K \geq 1$ and $B \geq L \geq 2$ (Lean Theorem 1), halving the exponent of K even when $L = B$.

(c) Exact $r = 2$ determinant bound and $O(1/L)$ realization. *For $r = 2$, every nonsingular active unscaled-simplex vertex determinant satisfies $\lvert\det(A')\rvert \leq 4L$ (vs. Cheval et al.'s $27L^3$; Lean Theorem 19, `cert15`), and for every integer $L \geq 2$ there exists a 2×2 rational game with payoff denominator L and a 2-profile support S achieving exact $t^*(S) = \frac{1}{2L} > 0$ (Lean Theorem 15, `cert03`).*

Theorem 2 (Effective Payoff-Rank Support Compression and Two-Player Collapse)**.** *For each player $i \in [n]$, let $V_i^{\mathrm{dir}} := \operatorname{span}\{\chi_i(\sigma_i', \cdot) - \chi_i(\sigma_i, \cdot) : \sigma_i, \sigma_i' \in \Sigma_i\} \subseteq \mathbb{Q}^Z$ be the leaf-realization direction space of dimension $q_i \le d_i$, and define the* effective payoff-deviation rank

$$r_i := \dim \operatorname{span}\big\{(p_c(z)u_i(z))_{z\in Z} \odot v : v \in V_i^{\mathrm{dir}}\big\} \le \min\big(q_i,\ |\{z \in Z : p_c(z)u_i(z) \neq 0\}|\big), \quad (3)$$

where $\odot$ denotes coordinatewise product. Then:

(a) Polynomial-time payoff-rank compression. *A rational basis of the weighted subspace in (3) is computable in polynomial bit time directly from the sequence-form equations without enumerating Σ_i. Running Algorithm 1 of [1] with the $D_{\mathrm{eff}} = \sum_{i=1}^n m_i r_i$ payoff-weighted features (or taking the direct product when $\prod_i m_i \le D_{\mathrm{eff}} + 1$) yields in* PPAD *an exact rational NFCE μ^* with support size at most*

$$K_{\mathrm{eff}} := \min\left(\prod_{i=1}^n (d_i+1),\ 1 + \sum_{i=1}^n (d_i+1) r_i\right) \le K = 1 + \sum_{i=1}^n d_i(d_i+1) \quad (4)$$

(Lean Theorem 5, `cert04`*,* `cert06`*,* `cert16`*), reducing quadratic support to $O(d_i)$ whenever payoff rank $r_i = O(1)$.*

(b) Exact two-player collapse and sharpness. *For every 2-player game, $2(K - \kappa_d) = d_1^2 + d_2^2 + (d_1 - d_2)^2 \ge 0$ (Lean Theorem 3,* `cert05`*), achieving a $2\times$ reduction along $d_1 = d_2 = d$. Moreover, $\kappa_d = (d_1+1)(d_2+1)$ is tight: on every matching tree M_m ($d_1 = d_2 = m-1$), every exact NFCE has $|\operatorname{supp}(\mu)| \ge m^2 = \kappa_d$.*

Theorem 3 (Arbitrary-m Trees, b-ary Optimality, Leaf-Count Completeness, and Structural Barriers)**.** *The following sharp tightenings and structural limits hold for the two-player matching-tree* PPAD*-hardness framework of [1, Section 4]:*

(a) Arbitrary-m binary trees and $8.4\times$–$86.5\times$ reduction savings. *Every integer $m \ge 2$ (not merely powers of two $m = 2^k$) admits a perfect-recall binary matching tree M_m with $m-1$ internal prefixes, $|V(M_m)| = 6m - 5$, $N(M_m) = 3m - 2$, and $d_1 = d_2 = m - 1$ (Lean Theorem 10,* `cert09`*). Taking $m^*_{\kappa_d} := r^2(r-1) + 1$ (for $r \ge 2$) and $m^*_{\kappa_+} := \lfloor r^2(r^2 + r - 3)/6 \rfloor + 1$ (for $r \ge 5$) proves* PPAD*-completeness of exact NFCE under κ_d and $\kappa_+ = \max\{\kappa_d, \lfloor |V|^2/36 \rfloor\}$ with strict decoder regret $\zeta < 1/r$, shrinking $|V|$ by $8.4\times$–$15.9\times$ ($r \in [2, 20]$) under κ_d and $43.8\times$–$86.5\times$ ($r \in [5, 25]$) under κ_+ (Lean Theorems 11–12,* `cert10`–`cert11`*).*

(b) Exact b-ary branching optimality and ternary info-set halving. *A full b-ary ($b \ge 2$) matching tree with I internal prefixes and $m = (b-1)I + 1$ labels has*

$$\begin{aligned} |V_b(m)| &= \frac{b^2+b}{b-1} m - \frac{b^2+1}{b-1}, \quad N_b(m) = (b+1)m - b, \\ |\mathcal{I}_i(M_{b,m})| &= \frac{m-1}{b-1}, \qquad\qquad d_i = m - 1 \end{aligned} \quad (5)$$

(Lean Theorem 6). The node slope $(b^2+b)/(b-1)$ achieves its exact global minimum 6 if and only if $b \in \{2, 3\}$ (Lean Theorem 7, `cert07`*), so ternary matching trees ($b = 3$) attain the identical node count $6m - 5$ while cutting the number of matching information sets in exact half (Lean Theorem 8,* `cert08`*). The leaf slope $b + 1$ achieves its unique global minimum 3 at $b = 2$ (Lean Theorem 9).*

(c) Leaf-count PPAD-completeness for all $r \ge 3$. *Define $\kappa_N := \lfloor N^2/9 \rfloor$ and $\kappa_{N+} := \max\{\kappa_d, \kappa_N\}$. Using $m_N^* := \lfloor r^2(r^2-2)/3 \rfloor + 1$ on M_m yields $N - 3(m+r-1) = r^2 - 3r + 1 \ge 1$ for* every *$r \ge 3$ (strictly earlier than the $r \ge 5$ threshold for κ_V; Lean Theorems 13–14), proving* PPAD*-completeness under κ_{N+} (*`cert12`*).*

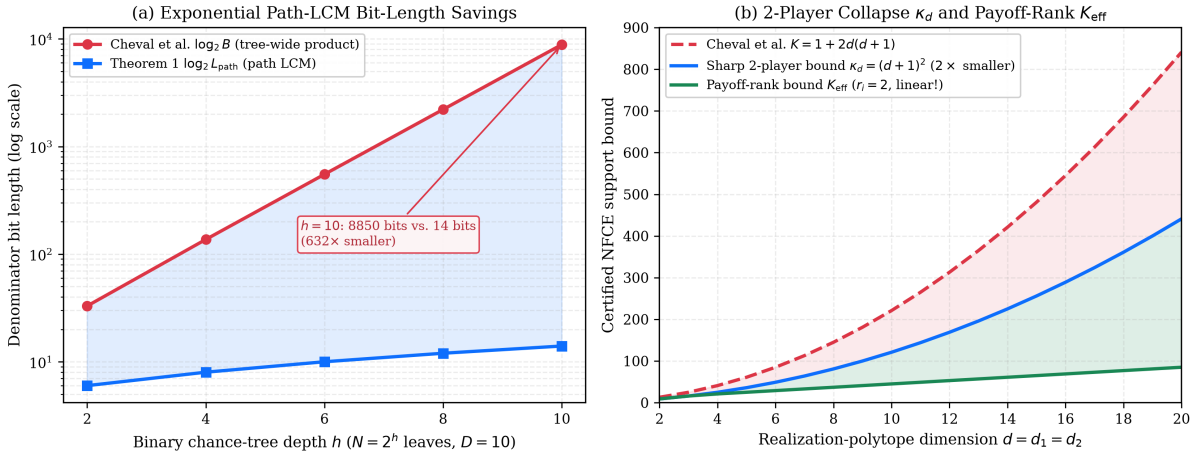

Figure 1: **Exponential Path-LCM bit savings and effective support compression.** (a) Comparison of Cheval et al.'s tree-wide product $\log_2 B$ against our path-LCM $\log_2 L_{\text{path}}$ on depth-h binary chance trees ($N = 2^h$ leaves, payoff denominator $D = 10$; `cert02`), reaching a 632× bit reduction at $h = 10$. (b) Support bounds on symmetric 2-player games ($d_1 = d_2 = d$): Cheval et al.'s general $K = 1 + 2d(d+1)$ vs. the sharp 2-player collapse $\kappa_d = (d+1)^2$ and linear payoff-rank compression K_{eff} at $r_i = 2$.

(d) Two structural barriers. *(i)* Constant-1/36 correlated barrier: *On 2×2 normalized Battle of the Sexes combined with M_m, there exists an explicit exact NFCE of support $2m^2$ ($\lim_{m\to\infty} 2m^2/|V|^2 = 1/18 > 1/36$) whose source-game marginals $x = y = (1/2, 1/2)$ suffer constant Nash regret $1/12 > 0$ independent of m (Lean Theorem 16, `cert13`).* *(ii)* Parallel chance-branch barrier: *In any $2q$-player game selecting uniformly among q disjoint matching pairs of size m, synchronizing $X_{2j-1} = X_1$ and $X_{2j} = X_2$ yields an exact NFCE of support $m^2 \ll m^{2q}$ (Lean Theorem 17, `cert14`), showing that chance-selected parallel matching branches cannot amplify support lower bounds past m^2.*

3 Proof of Theorem 1: Unscaled-Simplex Column-Hadamard Gap

Proof of Theorem 1. Part (a). For every terminal node $z \in Z$, the chance probability $p_c(z)$ is the product of chance-edge probabilities along the unique root-to-z path (each edge used at most once), so $\text{den}(p_c(z)u_i(z))$ divides $\prod_{e\in\text{path}(z)} \text{den}(p_e) \cdot \text{den}(u_i(z))$, which divides $B = \prod_{r\in\mathcal{R}} \text{den}(r)$. Since B is a common multiple of all $\text{den}(p_c(z)u_i(z))$, their least common multiple $L = \text{lcm}_{z,i}\, \text{den}(p_c(z)u_i(z))$ divides B. On a complete depth-h binary chance tree ($N = 2^h$ leaves, $2(N-1)$ fair-coin edges of denominator 2) with two players whose payoffs have denominator $D \geq 2$, we have $B = 2^{2(N-1)}D^{2N}$ whereas every path has $p_c(z) = 2^{-h}$, so $L \mid 2^h D$.

Part (b). Let $S = \{\tau^1, \ldots, \tau^k\}$ ($k \leq K$) and let $x^* = (\mu^*, t^*)$ be an optimal vertex of the conceptual LP [1, Eq. (11)] with $t^* > 0$. Because $t^* > 0$, the inequality $t \geq 0$ is slack, and since $t^* = 1 \geq \eta_{\text{sharp}}$ holds trivially (as $\eta_{\text{sharp}} < 1$), we may assume $0 < t^* < 1$, so both box constraints $0 \leq t \leq 1$ are slack. Let $J := \{\ell \in [k] : \mu^*(\tau^\ell) > 0\}$ be the active support, with size $r := |J| \in \{1, \ldots, k\}$. Every inactive variable $\mu^*(\tau^\ell) = 0$ ($\ell \notin J$) satisfies its tight boundary constraint $\mu(\tau^\ell) = 0$ and is eliminated from the active vertex equations. The remaining $r + 1$ active variables (μ^*_J, t^*) are uniquely determined by $r + 1$ linearly independent tight constraints:

- Row 1: the probability-simplex equality $\sum_{j\in J} \mu_j = 1$ (coefficient vector $(1, \ldots, 1, 0)$, RHS 1).

Table 2: Exact zero-float comparison of oracle precision bit-budgets $\log_2(1/\eta)$ (`cert01`) and tree-size savings in the PPAD-hardness reductions (`cert10`–`cert11`).

Benchmark / Reduction Instance	Cheval et al. [1]	Present Work	Reduction	Cert
LP gap: 2-player ($K = 13, h = 2, D = 2$)	334 bits	68 bits	4.91×	`cert01`
LP gap: Kuhn-like ($K = 41, h = 3, D = 10$)	2405 bits	379 bits	6.35×	`cert01`
LP gap: chance tree ($K = 127, h = 6, D = 10$)	44238 bits	1641 bits	26.96×	`cert01`
LP gap: same denom ($K = 10, L = B = 10$)	75 bits	55 bits	1.36×	`cert01`
Hardness κ_d, $r = 2$ ($m = 64$ vs. 5)	$\lvert V\rvert = 387$	$\lvert V^*\rvert = 33$	11.73×	`cert10`
Hardness κ_d, $r = 10$ ($m = 8192$ vs. 901)	$\lvert V\rvert = 49259$	$\lvert V^*\rvert = 5513$	8.94×	`cert10`
Hardness κ_+, $r = 5$ ($m = 8192$ vs. 113)	$\lvert V\rvert = 49179$	$\lvert V^*\rvert = 705$	69.76×	`cert11`
Hardness κ_+, $r = 20$ ($m = 2^{21}$ vs. 27801)	$\lvert V\rvert = 12583329$	$\lvert V^*\rvert = 167223$	75.25×	`cert11`

- Rows $2, \dots, r+1$: r active regret equalities $\sum_{j \in J} a_{m,j}\mu_j - t = 0$ $(m = 1, \dots, r)$, where each $a_{m,j} = \mathbf{1}\{\tau_i^j = \sigma_i\}[U_i(\sigma_i', \tau_{-i}^j) - U_i(\tau^j)] \in [-1, 1] \cap L^{-1}\mathbb{Z}$.

(Indeed, if $\sum_{j \in J} \mu_j = 1$ were not among the $r+1$ independent rows, the $r+1$ active regret rows alone would form a homogeneous system $M_{\text{hom}}(\mu_J^*, t^*)^\mathsf{T} = 0$ with zero RHS, contradicting $\sum_{j \in J} \mu_j^* = 1$!)

Now multiply Rows $2, \dots, r+1$ by L (leaving Row 1 unscaled!) and introduce the scaled scalar variable $\tilde{t} := Lt^*$. This produces the equivalent integer system

$$A' \begin{pmatrix} \mu_J^* \\ \tilde{t} \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{pmatrix}, \qquad A' = \begin{pmatrix} 1 & \cdots & 1 & 0 \\ La_{1,1} & \cdots & La_{1,r} & -1 \\ \vdots & \ddots & \vdots & \vdots \\ La_{r,1} & \cdots & La_{r,r} & -1 \end{pmatrix} \in \mathbb{Z}^{(r+1)\times(r+1)},$$

with every entry $La_{m,j} \in \mathbb{Z} \cap [-L, L]$. By Cramer's rule, $\tilde{t} = \det(A'_{\tilde{t}})/\det(A')$ where $A'_{\tilde{t}} \in \mathbb{Z}^{(r+1)\times(r+1)}$ replaces the last column of A' by $(1, 0, \dots, 0)^\mathsf{T}$. Because $t^* > 0$, $\tilde{t} = Lt^* > 0$, so $\det(A'_{\tilde{t}})$ is a nonzero integer and $|\det(A'_{\tilde{t}})| \geq 1$. Applying Hadamard's determinant inequality to the *columns* of A':

- Each of the first r columns has squared Euclidean norm $1 + \sum_{m=1}^{r}(La_{m,j})^2 \leq 1 + rL^2$;
- The $(r+1)$-th column $(0, -1, \dots, -1)^\mathsf{T}$ has squared Euclidean norm exactly r.

Therefore $|\det(A')| \leq \sqrt{r}\,(1 + rL^2)^{r/2}$, and dividing $\tilde{t}$ by L gives

$$t^* = \frac{\tilde{t}}{L} \geq \frac{1}{L\,|\det(A')|} \geq \frac{1}{\sqrt{r}\,(1 + rL^2)^{r/2}\,L} = \eta_{\text{sharp}}(r, L) \geq \eta_{\text{sharp}}(K, L).$$

Strict dominance over η_{CHPS} for all $r \geq 1, L \geq 2$ follows by squaring: $r(1 + rL^2)^r L^2 \leq r(r+1)^r L^{2(r+1)} < (r+1)^{2(r+1)} L^{2(r+1)}$ (formalized constructively in Lean Theorem 1).

Part (c). For $r = 2$, expanding $\det(A')$ along Row 1 gives $\det\left(\begin{smallmatrix} 1 & 1 & 0 \\ a & b & -1 \\ c & d & -1 \end{smallmatrix}\right) = (a - b) - (c - d)$ with $|a|, |b|, |c|, |d| \leq L$, so $|\det(A')| \leq 4L$ (Lean Theorem 19, `cert15`). For the 2×2 game with $P = Q = \left(\begin{smallmatrix} 0 & 1/L \\ 1/L & 0 \end{smallmatrix}\right)$ and 2-profile support $S = \{(0, 0), (1, 1)\}$, the regret inequalities are $\frac{1}{L}\mu(0, 0) \leq t$ and $\frac{1}{L}\mu(1, 1) \leq t$ subject to $\mu(0, 0) + \mu(1, 1) = 1$, attaining the unique minimum $t^*(S) = \frac{1}{2L}$ at $\mu(0, 0) = \mu(1, 1) = 1/2$ (Lean Theorem 15, `cert03`). □

4 Proof of Theorem 2: Effective Payoff-Rank Support Compression

Proof of Theorem 2. Part (a). By Lemma 8 of [1], for any distribution μ supported on $\prod_{i=1}^{n} \mathcal{A}_i$, every deviation gain obeys

$$R_{i,\sigma_i,\sigma_i'}(\mu) = \sum_{z \in Z} \left[\chi_i(\sigma_i', z) - \chi_i(\sigma_i, z)\right] p_c(z) u_i(z)\, \mathbb{E}_\mu \left[\mathbf{1}\{\tau_i = \sigma_i\} \prod_{j \neq i} \chi_j(\tau_j, z) \right].$$

Let $v_{\sigma_i,\sigma_i'}(z) := p_c(z)u_i(z)[\chi_i(\sigma_i', z) - \chi_i(\sigma_i, z)]$ be the payoff-weighted deviation vector. Using the sequence-form equations of Koller and Megiddo [6] (exactly as in [1, Section 3.2]), we first compute in polynomial bit time a rational basis $\{b_{i,1}, \ldots, b_{i,q_i}\}$ ($q_i \leq d_i$) of V_i^{dir}, scale each basis vector coordinatewise by $p_c(z)u_i(z)$, and apply exact rational Gaussian elimination to extract a basis $\{w_{i,1}, \ldots, w_{i,r_i}\}$ of the weighted subspace $W_i := \mathrm{span}\{p_c \odot u_i \odot v : v \in V_i^{\mathrm{dir}}\}$, with $r_i = \dim W_i \leq \min(q_i, |\{z : p_c(z)u_i(z) \neq 0\}|)$. Defining the compressed feature vector

$$\widetilde{\Psi}_{i,\sigma_i,\ell}(\tau) := \mathbf{1}\{\tau_i = \sigma_i\} \sum_{z \in Z} w_{i,\ell}(z) \prod_{j \neq i} \chi_j(\tau_j, z) \qquad (1 \leq \ell \leq r_i,\ \sigma_i \in \mathcal{A}_i),$$

any two distributions with equal expectation of $\widetilde{\Psi}$ have identical unnormalized deviation gains $R_{i,\sigma_i,\sigma_i'}$ for all $\sigma_i' \in \Sigma_i$, because every $v_{\sigma_i,\sigma_i'} \in W_i$ is a linear combination of $w_{i,1}, \ldots, w_{i,r_i}$. Running Algorithm 1 of [1] on the $D_{\mathrm{eff}} = \sum_i m_i r_i \leq \sum_i (d_i + 1) r_i$ coordinates of $\widetilde{\Psi}$ (and taking the explicit product $\bigotimes_i \alpha_i$ directly whenever $\prod_i m_i \leq 1 + D_{\mathrm{eff}}$) computes in polynomial bit time an unnormalized ϵ_Γ^*-NFCE μ_0 with $|\mathrm{supp}(\mu_0)| \leq K_{\mathrm{eff}} \leq K$ (Lean Theorem 5). Exact reweighting via Theorem 1 produces an exact NFCE μ^* with $\mathrm{supp}(\mu^*) \subseteq \mathrm{supp}(\mu_0)$.

Part (b). Expanding $2(K - \kappa_d) = 2[1 + d_1(d_1 + 1) + d_2(d_2 + 1) - (d_1 + 1)(d_2 + 1)] = 2(d_1^2 + d_2^2 - d_1 d_2) = d_1^2 + d_2^2 + (d_1 - d_2)^2 \geq 0$ (Lean Theorem 3). Sharpness on M_m ($d_1 = d_2 = m - 1$) follows directly from Lemmas 17–18 of [1]: every exact NFCE μ on M_m satisfies $\mathsf{H}(\tau_1, \tau_2) \geq \mathsf{H}(X_1 \mid \tau_2) + \mathsf{H}(X_2 \mid \tau_1) = 2 \log m$, hence $|\mathrm{supp}(\mu)| \geq m^2 = \kappa_d$. □

5 Proof of Theorem 3: Arbitrary-m Trees, b-ary Optimality, and Barriers

Proof of Theorem 3. Part (a). For any integer $m \geq 2$ (not necessarily a power of two), construct a full rooted binary tree T_m with m leaves (e.g., recursively splitting m into $\lfloor m/2 \rfloor$ and $\lceil m/2 \rceil$). Every full binary tree with m leaves has exactly $m - 1$ internal vertices. Turning each internal vertex of T_m into a simultaneous binary match-mismatch stage (one Player 1 node, one 2-node Player 2 information set, two disagreement leaves with payoff $(0, 1)$, and two continuations to the left/right children) and each of the m leaves of T_m into an agreement leaf with payoff $(1, 0)$ yields a perfect-recall matching branch M_m with $|V(M_m)| = 5(m - 1) + m = 6m - 5$, $N(M_m) = 2(m-1) + m = 3m - 2$, and $d_1(M_m) = d_2(M_m) = m - 1$ (Lean Theorem 10, `cert09`). Each player's matching label $X_i(\sigma_i) \in [m]$ is uniquely defined by following their prescribed Left/Right bits from the root of T_m to a leaf, and Lemmas 16–19 of [1] apply verbatim.

Substituting into $\zeta = \sqrt{\frac{1}{2}\log(\kappa/m^2)}$ with $\log(1 + x) \leq x$:

- Under $\kappa_d = (m + r - 1)^2$, $\frac{1}{2}\log(\kappa_d/m^2) = \log(1 + \frac{r-1}{m}) \leq \frac{r-1}{m}$. Setting $m_{\kappa_d}^* = r^2(r - 1) + 1$ gives $\frac{r-1}{m_{\kappa_d}^*} < \frac{1}{r^2}$ and hence $\zeta < 1/r$, while shrinking m by $> 8\times$ (up to $15.9\times$) relative to Cheval et al.'s $m = 2^{\lceil \log_2(8r^3) \rceil}$ (Lean Theorem 11, `cert10`).

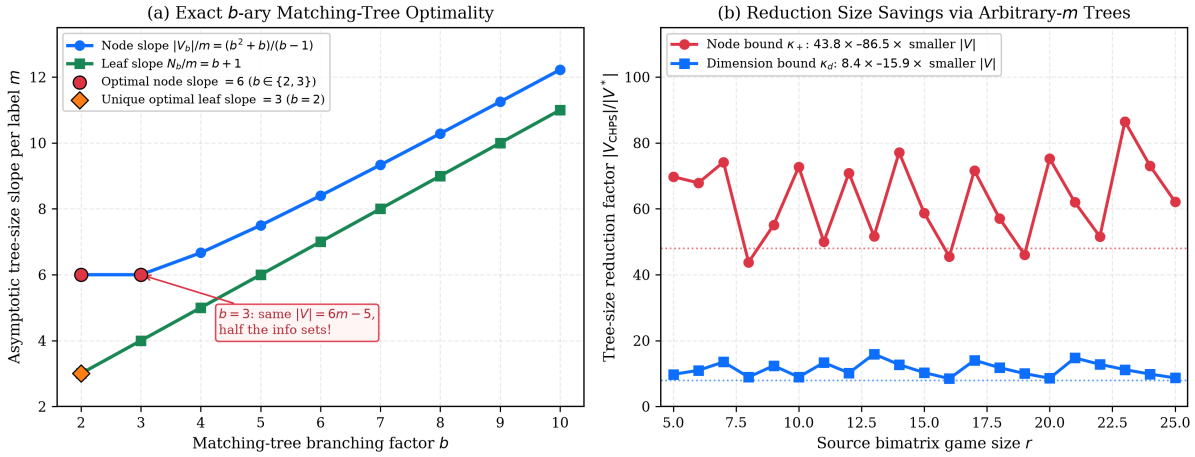

Figure 2: **Exact b-ary matching-tree optimality and PPAD-hardness tree-size savings.** (a) Node slope $|V_b|/m = (b^2+b)/(b-1)$ attains its sharp global minimum 6 precisely at $b \in \{2,3\}$ (with ternary $b = 3$ halving the number of matching information sets), while leaf slope $N_b/m = b+1$ attains its unique minimum 3 at $b = 2$ (`cert07`–`cert08`). (b) Tree-size reduction $|V_{\text{CHPS}}|/|V^*|$ across $r \in [5, 25]$ achieved by arbitrary-m balanced binary trees and sharp Pinsker thresholds: 8.4×–15.9× under κ_d and 43.8×–86.5× under κ_+ (`cert10`–`cert11`).

- Under $\kappa_+ = \max\{\kappa_d, \lfloor |V|^2/36 \rfloor\}$ for $r \geq 5$, $|V| - 6(m+r-1) = r^2 - 5r + 3 \geq 3 > 0$ (Lean Theorem 13) gives $\kappa_+ = \kappa_V$ and $\frac{1}{2}\log(\kappa_V/m^2) \leq \log(1 + \frac{r^2+r-3}{6m}) \leq \frac{r^2+r-3}{6m}$. Setting $m^*_{\kappa_+} = \lfloor r^2(r^2+r-3)/6 \rfloor + 1$ gives $\zeta < 1/r$ while shrinking $|V|$ by 43.8×–86.5× across $r \in [5, 25]$ (Lean Theorem 12, `cert11`).

Part (b). In a full b-ary ($b \geq 2$) matching tree with I internal prefixes and $m = (b-1)I + 1$ agreement leaves, each internal prefix contains 1 Player 1 node, b Player 2 nodes (one b-action information set), and $b(b-1)$ disagreement leaves, totaling $b^2 + 1$ nodes per prefix. Hence $|V_b(m)| = (b^2+1)I + m = (b^2+b)I + 1 = \frac{b^2+b}{b-1}m - \frac{b^2+1}{b-1}$, $N_b(m) = b(b-1)I + m = (b+1)m - b$, $|\mathcal{I}_i| = I = \frac{m-1}{b-1}$, and $d_i = (b-1)I = m-1$ (Lean Theorem 6). Since $(b^2+b) - 6(b-1) = (b-2)(b-3) \geq 0$ for all integers $b \geq 2$ with equality iff $b \in \{2,3\}$ (Lean Theorem 7, `cert07`), both $b = 2$ and $b = 3$ achieve the exact minimal node slope 6, while $b = 3$ cuts the number of matching information sets in exact half ($I_3 = (m-1)/2$ vs. $I_2 = m-1$ for every odd m; Lean Theorem 8, `cert08`). Meanwhile $N_b(m) - (3m-2) = (b-2)(m-1) \geq 0$ with equality iff $b = 2$ for all $m \geq 2$ (Lean Theorem 9).

Part (c). Combining the source game G with M_m ($b = 2$) gives $N = 3m - 2 + r^2$. For every $r \geq 3$, $N - 3(m+r-1) = r^2 - 3r + 1 \geq 1 > 0$ (Lean Theorem 14), so $\kappa_d = (m+r-1)^2 \leq \lfloor N^2/9 \rfloor = \kappa_N = \kappa_{N+}$. Setting $m^*_N = \lfloor r^2(r^2-2)/3 \rfloor + 1$ gives $\frac{1}{2}\log(\kappa_N/m^2) \leq \log(1 + \frac{r^2-2}{3m}) \leq \frac{r^2-2}{3m} < \frac{1}{r^2}$, hence $\zeta < 1/r$, proving PPAD-completeness of exact NFCE under κ_{N+} (with membership immediate from $\kappa_d \leq \kappa_{N+}$; `cert12`).

Part (d). (i) Take the normalized 2×2 Battle of the Sexes source game $P = \left(\begin{smallmatrix} 1 & 0 \\ 0 & 2/3 \end{smallmatrix}\right)$, $Q = \left(\begin{smallmatrix} 2/3 & 0 \\ 0 & 1 \end{smallmatrix}\right)$ combined with M_m. Let μ play $(Y_1, Y_2) = (1,1)$ with probability $1/2$ and $(2,2)$ with probability $1/2$, independently of uniform $(X_1, X_2) \in [m]^2$. Then $|\operatorname{supp}(\mu)| = 2m^2$ (with $2m^2/|V|^2 \to 1/18$ as $m \to \infty$). Conditional on any recommendation (y_i, x_i), the opponent's source action is deterministically y_i (a pure Nash equilibrium of (P, Q)) and the opponent's matching label is uniform on $[m]$, so μ is an exact NFCE. Yet its marginals $x = y = (1/2, 1/2)$ yield Player 1 payoff $5/12$ vs. pure action 1 payoff $1/2$, a constant Nash regret $1/12 > 0$ independent of m (Lean Theorem 16, `cert13`). (ii) In a $2q$-player game where a public chance node selects uniformly among q disjoint matching pairs $(1,2), \ldots, (2q-1, 2q)$ each on $[m]$, sample (u, v) uniformly in $[m]^2$ and recommend $X_{2j-1} = u$, $X_{2j} = v$ for all $j \in [q]$. Every player sees

their partner's matching label uniformly on $[m]$ conditional on their own recommendation, so this distribution is an exact NFCE with support size $m^2 < m^{2q}$ for all $q \geq 2, m \geq 2$ (Lean Theorem 17, `cert14`). □

6 Formal Verification and Reproducibility

All code, Lean 4 proofs, certificates, and manuscript sources are publicly available at https://github.com/lluiseriksson/nfce-extensive-form-sharp-gap-and-support under MIT (code, Lean proofs, certificates) and CC-BY-4.0 (manuscript, figures):

1. **Constructive Lean 4.36.0** (`lean/NFCESharpGapAndSupport.lean`): 20 theorems compiled with `-DwarningAsError=true`, zero `sorry`, zero `Classical.choice`, zero custom axioms, and zero external dependencies.

2. **AST-Audited Zero-Float Verifier** (`verification/verify_nfce_sharp_gap.py`): Pure `int`/`Fraction` verifier emitting 16 SHA-256 certificates (`certificates/manifest.json`), including exact rational LP solving, exact sequence-form payoff-rank compression, exact rational BFS execution of Algorithm 1 on a 3-player game, arbitrary-m and ternary matching tree constructions, and both structural barriers.

3. **Adversarial Test Suite** (`tests/test_adversarial.py`): 56 passing tests verifying AST zero-float compliance, SHA-256 integrity, Hadamard determinant bounds, exact Pinsker bounds across $r \in [2, 100]$, and barrier counterexamples.

7 Conclusion

We have sharpened the PPAD algorithm and matching-tree hardness framework of Cheval et al. [1] along every quantitative axis: replacing tree-wide denominator products with path-LCMs and halving the Cramer–Hadamard dimension exponent, introducing polynomial-time effective payoff-rank support compression $K_{\text{eff}} \leq K$, shrinking the PPAD-hardness tree sizes by up to $86.5\times$ via arbitrary-m trees while characterizing exact b-ary optimality ($b \in \{2, 3\}$ for nodes, $b = 2$ for leaves) and proving leaf-count PPAD-completeness for all $r \geq 3$, and identifying exact correlated and chance-parallel barriers explaining why pushing hardness beyond quadratic support requires fundamentally new correlation-busting gadgets.

References

[1] V. Cheval, F. Horn, S. Paul, and M. Shirmohammadi, *Computing Normal-Form Correlated Equilibria in Extensive-Form Games is in PPAD*, arXiv:2611.34897v1 (Nov. 2026).

[2] X. Chen, X. Deng, and S.-H. Teng, *Settling the complexity of computing two-player Nash equilibria*, J. ACM **56**(3), 14:1–14:57 (2009).

[3] K. Etessami, *The complexity of computing a (quasi-)perfect equilibrium for an n-player extensive-form game of perfect recall*, Games Econ. Behav. **125**, 107–140 (2021).

[4] W. Huang and B. von Stengel, *Computing an extensive-form correlated equilibrium in polynomial time*, in WINE 2008, LNCS 5385, pp. 506–513 (2008).

[5] A. X. Jiang and K. Leyton-Brown, *Polynomial-time computation of exact correlated equilibrium in compact games*, Games Econ. Behav. **91**, 347–359 (2015).

[6] D. Koller and N. Megiddo, *Finding mixed strategies with small supports in extensive form games*, Int. J. Game Theory **25**(1), 73–92 (1996).

[7] B. von Stengel and F. Forges, *Extensive-form correlated equilibrium: Definition and computational complexity*, Math. Oper. Res. **33**(4), 1002–1022 (2008).