AIRR:ARR-2026-67BH4NQ5618SQREM v1 [Game theory] 27 Nov 2026 · AIRR.SCIENCE

Transcendental-Free Rational Cauchy–Schwarz Decoding, Inverse-Polynomial Approximate NFCE PPAD-Completeness, Exponential Support-Reweighting Barriers, and Subquadratic-Slack Support Thresholds in Extensive-Form Games

Lluis Eriksson
Independent Researcher
ORCID: 0009-0006-7494-9040

November 27, 2026

Abstract

Cheval, Horn, Paul, and Shirmohammadi (2026, arXiv:2611.34897v1) resolved a major open problem in algorithmic game theory by placing the computation of a normal-form correlated equilibrium (NFCE) in imperfect-information extensive-form games with perfect recall in PPAD for an arbitrary number of players (outputting a rational NFCE of support size at most $K = 1+\sum_{i=1}^{n} d_i(d_i+1) \le 1+N(N-1)$), and by proving PPAD-completeness of exact two-player NFCE under the quadratic support bounds $\kappa_d = (d_1+1)(d_2+1)$ and $\kappa_+ = \max\{\kappa_d, \lfloor |V|^2/36 \rfloor\}$. We establish four sharp, transcendental-free structural refinements and impossibility barriers directly extending their framework:

1. **Rational Cauchy–Schwarz Decoding & $8\times/64\times$ Gadget Sharpening:** Replacing Shannon entropy, ln, D_{KL}, and Pinsker's inequality in their Lemmas 18–19 with a purely algebraic two-bin Cauchy–Schwarz partition over $\mathbb{Q}$, we prove $\mathrm{TV}(\pi, x\otimes y)^2 \le \mathrm{TV}(\mu, \mu_1\otimes\mu_2)^2 \le \frac{1}{4}(\frac{\kappa}{m^2}-1)$ (asymptotically tight as $m\to\infty$ for every fixed ratio $\kappa/m^2 \in (1,2)$). This is $\sqrt{2}\times$ sharper than Pinsker's $\sqrt{\frac{1}{2}\ln(\kappa/m^2)}$ on $(1, 3.51286\dots)$ and reduces Theorem 3's gadget parameter m by an exact factor of $8\times$ (r^3 vs. $8r^3$; r^4 vs. $8r^4$), reducing $|V|$ by asymptotically $8\times$ and κ_d by asymptotically $64\times$.
2. **Robust Inverse-Polynomial δ-NFCE PPAD-Completeness:** Extending Theorem 3 beyond exact ($\delta = 0$) NFCE, we prove via reweighted Cauchy–Schwarz that any δ-NFCE on support $\le \kappa$ decodes to a weak $\zeta(\kappa, m, \delta)$-Nash equilibrium with $\zeta = 2(m^2+1)\delta + \frac{1}{2}\sqrt{\kappa/m^2 - (1-2m^2\delta)^2}$, proving PPAD-completeness for δ-NFCE at $\delta = (2048r^8)^{-1} = \Theta(|V|^{-8/3})$ under κ_d and $\delta = (2048r^{10})^{-1} = \Theta(|V|^{-5/2})$ under κ_+.
3. **Sharpened Hadamard LP Gap & Exact Exponential Necessity Family:** Replacing $B = \prod_{r\in R}\mathrm{den}(r)$ with the leaf-product lcm L_Γ and keeping the simplex row unscaled halves the exponent in Lemma 14 from $((K+1)B)^{-(K+1)}$ to $\eta_{\mathrm{sharp}} = (\sqrt{K}(K+1)^{K/2}L_\Gamma^K)^{-1}$. Conversely, we construct an explicit $\{0,1\}$-payoff two-player bimatrix family $(P_m, Q_m) \in \{0,1\}^{(2m+3)\times(4m+2)}$ ($B = L_\Gamma = 1$) with support $|S_m| = 6m+2$ whose exact minimum regret is $t^\star(S_m) = (2^{m+5} - 14m - 30)^{-1} = 2^{-\Omega(|S_m|)} > 0$, certified by closed-form rational primal and Farkas-dual witnesses.
4. **Universal $\lfloor |V|^2/4 \rfloor$ Totality, Infinite κ_V-Failure, $1/36$ b-ary Optimality, & Subquadratic-Slack PPAD-Completeness:** We prove universal two-player totality at $\kappa_d \le \lfloor (|V|-1)^2/4 \rfloor \le \lfloor |V|^2/4 \rfloor$; exhibit an infinite failure family $\Gamma_{\mathrm{match}}(k)$ ($|V| = 6m-5$, $m = 2^k$) where $\kappa_V = \lfloor |V|^2/36 \rfloor < m^2$ is infeasible for all $k \ge 1$; prove $\frac{b^2+b}{b-1} \ge 6$ (equality iff $b \in \{2,3\}$) makes $1/36$ uniquely optimal over b-ary prefix-matching trees; and resolve

Section 5's open problem for every subquadratic additive slack exponent $\gamma \in (0, 2]$ by proving PPAD-completeness under $\kappa_d + \lfloor C(d_1 + d_2)^{2-\gamma} \rfloor$ and $\max\{\kappa_d, \lfloor |V|^2/36 + C|V|^{2-\gamma} \rfloor\}$, while pinpointing $\gamma = 0$ (constant multiplicative slack) as an exact decoding barrier.

All core algebraic inequalities, gadget parameter bounds, primal/dual identities, totality theorems, and slack barriers are machine-checked in constructive Lean 4.36.0 (20 theorems, zero external imports, zero `sorry`, zero `Classical.choice`) and accompanied by a zero-float AST-audited exact-rational Python certificate suite (16 SHA-256 certificates; 56 adversarial tests).

AI Assistance Disclosure. This research manuscript, its constructive Lean 4.36.0 formalization, and its exact-rational Python verification suite were developed by Lluis Eriksson with interactive assistance from Google Antigravity for proof exploration, formal verification, and typesetting. All mathematical proofs, formal Lean theorems, and exact-rational certificates were mechanically verified; the author takes full responsibility for all claims.

1 Introduction

The computational complexity of finding an arbitrary normal-form correlated equilibrium (NFCE) in imperfect-information extensive-form games with perfect recall has stood for nearly two decades as a central open problem in algorithmic game theory [14, 17, 10, 8, 15, 6, 5, 1]. Unlike extensive-form correlated equilibria (EFCE), which are polynomial-time computable [10], an NFCE requires a mediator to commit before play to a full pure-plan profile $\sigma \in \Sigma = \prod_{i=1}^{n} \Sigma_i$, where $|\Sigma_i|$ is exponential in the game tree. While threshold/optimal NFCE problems were recently shown to be PSPACE-hard [5] and even EXPTIME-complete with chance [9], the complexity of computing *any* NFCE was resolved in a landmark paper by Cheval, Horn, Paul, and Shirmohammadi [4]:

- **Theorem 1 & Corollary 2 of [4]:** Computing an exact rational NFCE of a perfect-recall extensive-form game lies in PPAD for an arbitrary number of players n (part of the input), via Etessami's weak ϵ-Nash PPAD oracle [7], iterative Carathéodory compression (Theorem 7 / Algorithm 1) to support size at most $K := 1 + \sum_{i=1}^{n} d_i(d_i + 1) \leq 1 + N(N - 1)$, and exact rational LP reweighting (Theorem 6, Lemmas 13–14) at threshold $\epsilon_\Gamma := \frac{1}{2((K+1)B)^{K+1}}$.
- **Theorem 3 of [4]:** Computing an exact two-player NFCE under either quadratic support bound $\kappa_d := (d_1 + 1)(d_2 + 1)$ or $\kappa_+ := \max\{\kappa_d, \kappa_V\}$ with $\kappa_V := \lfloor |V|^2/36 \rfloor$ is PPAD-complete, proved by reducing from weak $1/r$-Nash in $r \times r$ bimatrix games [3] using a public-coin gadget $\Gamma = \frac{1}{2}G + \frac{1}{2}M$ with binary prefix-matching branch M ($m = 2^k$ labels) and decoding via conditional uniformity (Lemma 17), mutual information, and Pinsker's inequality [16] (Lemmas 18–19, 21).

1.1 Questions Left Open and Summary of Contributions

Close inspection of [4] reveals four fundamental quantitative and structural questions:

1. *Are transcendental information-theoretic tools (Shannon entropy,* ln*, D_{KL}, Pinsker) necessary to decode approximate Nash equilibria from small-support NFCEs, and is the $m = 2^{\lceil \log_2(8r^3) \rceil} / 2^{\lceil \log_2(8r^4) \rceil}$ gadget overhead optimal?*
2. *Theorem 3 of [4] requires* exact *($\delta = 0$) NFCE because Lemma 17 collapses when conditional uniformity fails. Does* PPAD*-completeness persist for inverse-polynomial δ-approximate NFCEs?*
3. *Lemma 14 of [4] bounds the minimum nonzero LP regret $t^\star(S)$ from below by $((K + 1)B)^{-(K+1)}$. Can the exponent and denominator dependence be tightened, and is exponential decay $2^{-\Omega(|S|)}$ genuinely unavoidable even for $\{0, 1\}$-payoff normal-form games ($B = 1$)?*
4. *Section 5 of [4] explicitly asks whether* PPAD*-hardness holds under more permissive support bounds, and Corollary 20 notes via matching pennies ($|V| = 7$) that $\kappa_V = \lfloor |V|^2/36 \rfloor$ fails*

Table 1: Quantitative and structural comparison between Cheval et al. [4] and the present work.

Aspect	**Cheval et al. [4]**	**This Paper (Exact Refinements)**
Decoding ($\delta = 0$)	$\sqrt{\frac{1}{2}\ln(\kappa/m^2)}$ (Pinsker)	$\frac{1}{2}\sqrt{\kappa/m^2 - 1}$ (rational CS, $\sqrt{2}\times$ sharper near 1^+, Thm. 3.1)
Thm. 3 gadget size	$m = 2^{\lceil \log_2(8r^3)\rceil}, 2^{\lceil \log_2(8r^4)\rceil}$	$m = 2^{\lceil \log_2(r^3)\rceil}, 2^{\lceil \log_2(r^4)\rceil}$ ($8\times$ smaller m; asymp. $8 \times \lvert V\rvert, 64 \times \kappa_d$)
Approx. δ-NFCE	Open (exact $\delta = 0$ only)	PPAD-complete for $\delta = \Theta(\lvert V\rvert^{-8/3})$ (κ_d) and $\Theta(\lvert V\rvert^{-5/2})$ (κ_+, Thm. 4.1)
Uniform LP gap η	$((K+1)B)^{-(K+1)}$ (Lem. 14)	$(\sqrt{K}(K+1)^{K/2}L_\Gamma^K)^{-1}$ (halved exponent + lcm $L_\Gamma \mid B$, Thm. 5.1A)
LP gap necessity	No lower bound	$\{0,1\}$ bimatrix (P_m, Q_m) ($B = L_\Gamma = 1$): exact $t^\star(S_m) = (2^{m+5} - 14m - 30)^{-1}$ (Thm. 5.1B)
Node totality	$\kappa_V = \lfloor \lvert V\rvert^2/36 \rfloor$ fails at $\lvert V\rvert = 7$	Universal totality at $\lfloor \lvert V\rvert^2/4 \rfloor$; infinite κ_V-failure $\Gamma_{\text{match}}(k)$; unique $1/36$ over b-ary trees (Thm. 6.1)
Permissive slack	Open (Section 5)	PPAD-complete for all $\gamma \in (0, 2]$ ($\kappa_d + \lfloor C(d_1+d_2)^{2-\gamma} \rfloor$, $\kappa_+^{(C,\gamma)}$); barrier at $\gamma = 0$

totality. What is the universal node-based totality threshold, is $1/36$ *uniquely optimal across b-ary matching trees, and where is the exact boundary between permissive support bounds that remain* PPAD*-complete and those where matching-tree decoding fails?*

Table 1 summarizes how our four main theorems resolve each of these questions.

2 Preliminaries and Notation

We adopt the exact notation of [4]. A finite imperfect-information extensive-form game with perfect recall Γ is a finite rooted directed tree with node set V, terminal nodes $Z \subseteq V$ ($N := |Z|$), and n players. Each player $i \in [n]$ partitions its decision nodes into information sets $\mathcal{I}_i$ with action sets $A(I)$, and its realization polytope has affine dimension $d_i := \sum_{I \in \mathcal{I}_i}(|A(I)| - 1)$. Full pure plans are $\sigma_i \in \Sigma_i = \prod_{I \in \mathcal{I}_i} A(I)$, with joint space $\Sigma = \prod_{i=1}^n \Sigma_i$. By Kuhn's theorem [13] and the sequence-form sparsification of [11, 12], any rational behavioral strategy admits a polynomial-time computable realization-equivalent mixed strategy of support at most $d_i + 1$.

A correlation plan $\mu \in \Delta(\Sigma)$ has unnormalized deviation gains

$$R_{i,\sigma_i,\sigma_i'}(\mu) := \sum_{\tau \in \Sigma:\, \tau_i = \sigma_i} \mu(\tau)\big[U_i(\sigma_i', \tau_{-i}) - U_i(\tau)\big].$$

It is an *unnormalized ϵ-NFCE* if $R_{i,\sigma_i,\sigma_i'}(\mu) \le \epsilon$ for all i, σ_i, σ_i', and an *exact NFCE* if $\epsilon = 0$. In the deterministic deviation-rule formulation of [2, 4], μ is a *δ-NFCE* if for every player i and every deterministic deviation rule $\phi_i : \Sigma_i \to \Sigma_i$,

$$\mathbb{E}_{\sigma \sim \mu}\big[U_i(\phi_i(\sigma_i), \sigma_{-i}) - U_i(\sigma)\big] = \sum_{\sigma_i \in \Sigma_i} R_{i,\sigma_i,\phi_i(\sigma_i)}(\mu) \le \delta. \tag{1}$$

Notice that at $\delta = 0$, (1) coincides with exact NFCE, while for $\delta > 0$ it is the standard total-deviation (swap-regret) approximation notion (stronger than unnormalized δ-NFCE, since summing over σ_i aggregates regrets across all recommendations).

In Section 4 of [4], given an $r \times r$ rational bimatrix game $(P, Q) \in [0,1]^{r \times r}$ and $m = 2^k$ ($k \ge 1$), the public-coin gadget $\Gamma = \frac{1}{2}G + \frac{1}{2}M$ flips a fair root coin between the source bimatrix game G and the k-round binary prefix-matching game M (payoff $(1,0)$ on exact match of k-bit strings $X_1 = X_2 \in [m]$, $(0,1)$ on first mismatch). By Lemma 16 of [4], Γ has perfect recall,

$$|V| = 6m + r^2 + r - 3, \qquad N = 3m - 2 + r^2, \qquad d_1 = d_2 = m + r - 2, \qquad \kappa_d = (m + r - 1)^2,$$

and admits an exact rational NFCE of support at most κ_d. For any correlation plan $\mu \in \Delta(\Sigma)$ with marginals μ_1, μ_2, let $Y_i \in [r]$ be player i's Branch-G action, $X_i \in [m]$ be the unique label of $[m]$ consistent with player i's Branch-M plan, $\pi \in \Delta([r] \times [r])$ be the joint law of (Y_1, Y_2), and $x, y \in \Delta([r])$ be its marginals.

3 Transcendental-Free Rational Cauchy–Schwarz Decoding

In [4], Lemmas 18–19 bound $\mathrm{TV}(\pi, x \otimes y)$ via mutual information $I(\tau_1; \tau_2) \le \ln(\kappa/m^2)$ and Pinsker's inequality $\mathrm{TV}(\pi, x \otimes y) \le \sqrt{\frac{1}{2}\ln(\kappa/m^2)}$. We prove that the exact same conditional uniformity hypothesis of Lemma 17 implies a strictly sharper, purely rational quadratic bound over $\mathbb{Q}$ without any logarithms, entropy, or Pinsker's inequality.

Theorem 3.1 (Transcendental-Free Rational Cauchy–Schwarz Decoding)**.** *Let $\mu \in \Delta(S)$ ($|S| \le \kappa$) be any correlation plan of $\Gamma = \frac{1}{2}G + \frac{1}{2}M$ satisfying Lemma 17's conditional uniformity $\Pr(X_2 = j \mid \tau_1 = \sigma_1) = 1/m$ and $\Pr(X_1 = j \mid \tau_2 = \sigma_2) = 1/m$ for all positive-probability recommendations ($m \ge 1$). Then:*

1. (Support Lower Bound & Rational TV^2 Bound) *$\kappa \ge m^2$ and*

$$\mathrm{TV}(\pi, x \otimes y)^2 \;\le\; \mathrm{TV}(\mu, \mu_1 \otimes \mu_2)^2 \;\le\; \frac{1}{4}\left(\frac{\kappa}{m^2} - 1\right). \tag{2}$$

   Consequently, (x, y) is a weak ζ_{CS}-Nash equilibrium of (P, Q) with $\zeta_{\mathrm{CS}} := \frac{1}{2}\sqrt{\kappa/m^2 - 1}$.
2. (Asymptotic Tightness on $\kappa/m^2 \in (1, 2)$) *For any fixed excess ratio $\varepsilon := \kappa/m^2 - 1 \in (0, 1)$, over all discrete distributions $p \in \Delta(S)$ ($|S| \le \kappa$) satisfying $\sum_{e \in S} m^2 p_e^2 \le 1$, the supremum of $\sum_{e \in S}(p_e - m^2 p_e^2)_+$ equals $\frac{1}{2}\sqrt{\varepsilon} - O_\varepsilon(m^{-2})$ as $m \to \infty$ (with exact equality $\frac{1}{2}\sqrt{\varepsilon}$ whenever $2m^2a^2$ and $2m^2(1-a)^2$ are integers for $a = \frac{1}{2}(1 + \sqrt{\varepsilon})$).*
3. (Strict Superiority over Pinsker) *For every $u = \kappa/m^2 \in (1, 3)$, the rational Padé lower bound $\ln u \ge \frac{2(u-1)}{u+1}$ certifies $\frac{1}{4}(u-1) < \frac{u-1}{u+1} \le \frac{1}{2}\ln u$ purely over $\mathbb{Q}$ (and numerically $\zeta_{\mathrm{CS}} < \sqrt{\frac{1}{2}\ln u}$ on the entire interval $u \in (1, 3.51286\ldots)$, with $\sqrt{2}\times$ sharper leading constant as $u \to 1^+$).*
4. (8×/64× Sharper Theorem 3 Gadget Parameters) *Decoding a weak $1/r$-Nash equilibrium ($\zeta_{\mathrm{CS}} < 1/r \iff r^2(\kappa - m^2) < 4m^2$) holds already for:*
   - *$m = 2^{\lceil \log_2(r^3) \rceil}$ under κ_d ($r \ge 2$), reducing m by an exact factor of $8\times$ (and asymptotically reducing $|V|$ by $8\times$ and κ_d by $64\times$) relative to [4]'s $m = 2^{\lceil \log_2(8r^3) \rceil}$;*
   - *$m = 2^{\lceil \log_2(r^4) \rceil}$ under κ_+ ($r \ge 5$), reducing m by an exact factor of $8\times$ (and $|V|$ by asymptotically $8\times$) relative to [4]'s $m = 2^{\lceil \log_2(8r^4) \rceil}$.*

Proof. For each $e = (\sigma_1, \sigma_2) \in S$, write $p_e := \mu(e) > 0$ and $q_e := \mu_1(\sigma_1)\mu_2(\sigma_2)$. Since X_i is a deterministic function of σ_i, conditional uniformity gives

$$p_e \le \Pr(X_2 = X_2(\sigma_2), \tau_1 = \sigma_1) = \frac{1}{m}\mu_1(\sigma_1), \qquad p_e \le \Pr(X_1 = X_1(\sigma_1), \tau_2 = \sigma_2) = \frac{1}{m}\mu_2(\sigma_2).$$

Multiplying these two inequalities yields the pointwise quadratic domination

$$q_e = \mu_1(\sigma_1)\mu_2(\sigma_2) \ge m^2 p_e^2 \qquad (\forall e \in S). \tag{3}$$

Summing (3) over $e \in S$ gives $\sum_{e \in S} m^2 p_e^2 \le \sum_{e \in S} q_e \le 1$, and Cauchy–Schwarz $(\sum_{e \in S} p_e)^2 \le |S| \sum_{e \in S} p_e^2$ immediately gives $\kappa \ge |S| \ge m^2$.

To prove (2), partition $S = S_1 \sqcup S_2$ where $S_1 := \{e \in S : m^2 p_e \le 1\}$ and $S_2 := \{e \in S : m^2 p_e > 1\}$. Set $k_i := |S_i|$ ($k_1 + k_2 = |S| \le \kappa$), $a := \sum_{e \in S_1} p_e \in [0, 1]$, and $t := \sum_{e \in S_1} m^2 p_e^2 \in [0, 1]$. On S_2, $p_e - q_e \le p_e - m^2 p_e^2 < 0$, so every positive term of $p_e - q_e$ lies in S_1:

$$\mathrm{TV}(\mu, \mu_1 \otimes \mu_2) = \sum_{e \in S}(p_e - q_e)_+ \le \sum_{e \in S_1}(p_e - m^2 p_e^2) = a - t. \tag{4}$$

If $S_1 = \emptyset$ then $a = t = 0$ and $\mathrm{TV} = 0$; if $S_2 = \emptyset$ then $a = 1$ and $k_1 t \geq m^2$ gives $t \geq m^2/\kappa$, so setting $x := \kappa/m^2 - 1 \geq 0$ yields $(1-t)^2 \leq (1-\frac{1}{1+x})^2 = \frac{x^2}{(1+x)^2} \leq \frac{x}{4}$ (since $(1+x)^2 - 4x = (1-x)^2 \geq 0$). Now suppose both S_1, S_2 are nonempty (so $t > 0$ and $1 - t \geq \sum_{e \in S_2} m^2 p_e^2 > 0$). Applying Cauchy–Schwarz separately on S_1 and S_2 gives $k_1 t \geq m^2 a^2$ and $k_2(1-t) \geq m^2(1-a)^2$. Using the exact rational algebraic identity (cleared over $\mathbb{Z}$ in Lean 4, `thm01_two_bin_cs_identity`):

$$\frac{a^2}{t} + \frac{(1-a)^2}{1-t} = 1 + \frac{(a-t)^2}{t(1-t)}, \tag{5}$$

we obtain

$$\frac{\kappa}{m^2} \geq \frac{k_1 + k_2}{m^2} \geq \frac{a^2}{t} + \frac{(1-a)^2}{1-t} = 1 + \frac{(a-t)^2}{t(1-t)}.$$

Rearranging and using $4t(1-t) = 1 - (2t-1)^2 \leq 1$ yields $(a-t)^2 \leq t(1-t)(\frac{\kappa}{m^2} - 1) \leq \frac{1}{4}(\frac{\kappa}{m^2} - 1)$ (`thm02_rational_tv_quadratic_bound`). Data processing for the deterministic projection $(\tau_1, \tau_2) \mapsto (Y_1, Y_2)$ gives $\mathrm{TV}(\pi, x \otimes y) \leq \mathrm{TV}(\mu, \mu_1 \otimes \mu_2)$, proving (2). By the exact same Branch-G deviation calculation as Lemma 19 of [4], $e_a^\top P y - x^\top P y \leq \mathrm{TV}(\pi, x \otimes y)$ and $x^\top Q e_b - x^\top Q y \leq \mathrm{TV}(\pi, x \otimes y)$, so (x, y) is weak ζ_{CS}-Nash.

For Part 2 (tightness), fix $\varepsilon = \kappa/m^2 - 1 \in (0, 1)$ and take m large enough that $\varepsilon' := \varepsilon - 2/m^2 > 0$. Set $a' := \frac{1+\sqrt{\varepsilon'}}{2} \in (\frac{1}{2}, 1)$, $k_1 := \lceil 2m^2 a'^2 \rceil$, and $k_2 := \lceil 2m^2(1-a')^2 \rceil$, assigning uniform weight a'/k_1 on S_1 and $(1-a')/k_2$ on S_2. Since $2a'^2 + 2(1-a')^2 = 1 + \varepsilon'$, we have $k_1 + k_2 < m^2(1+\varepsilon') + 2 = \kappa$, so $|S| = k_1 + k_2 \leq \kappa$; furthermore $m^2 a'^2/k_1 \leq \frac{1}{2}$ and $m^2(1-a')^2/k_2 \leq \frac{1}{2}$, so $\sum_{e \in S} m^2 p_e^2 \leq 1$ and $m^2(a'/k_1) \leq \frac{1}{2a'} < 1$ (confirming $S_1 \subseteq \{m^2 p_e \leq 1\}$). Hence $\sum_{e \in S}(p_e - m^2 p_e^2)_+ \geq a' - m^2 a'^2/k_1 \geq a' - \frac{1}{2} = \frac{1}{2}\sqrt{\varepsilon - 2/m^2} = \frac{1}{2}\sqrt{\varepsilon} - O_\varepsilon(m^{-2})$. Moreover, whenever $2m^2 a^2$ and $2m^2(1-a)^2$ are integers at $a = \frac{1+\sqrt{\varepsilon}}{2}$ (e.g., $m = 256, \varepsilon = 1/16, a = 5/8$, giving $k_1 = 51200, k_2 = 18432, k_1 + k_2 = \kappa = 69632$), exact equality $(a-t)^2 = \frac{1}{64} = \frac{1}{4}\varepsilon$ is attained.

For Part 3, for $x = u - 1 \in (0, 2)$, $\frac{1}{4}x < \frac{x}{2+x} \iff x(2+x) < 4x \iff x < 2$ (`thm03_cs_beats_pinsker_small_eps`), while $\ln(1+x) \geq \frac{2x}{2+x}$ is elementary (their difference vanishes at 0 and has derivative $\frac{x^2}{(1+x)(2+x)^2} > 0$).

For Part 4, under $\kappa_d = (m + r - 1)^2$ with $r \geq 2$ and $m \geq r^3$:

$$r^2(\kappa_d - m^2) = 2mr^2(r-1) + r^2(r-1)^2 < 2m^2 + 2m^2 = 4m^2$$

(`thm04_gadget_8x_sharpening_kappa_d`). Under $\kappa_V = \lfloor |V|^2/36 \rfloor$ with $|V| = 6m + A$, $A = r^2 + r - 3$, $r \geq 5$, and $m \geq r^4$: first $|V| - 6(m + r - 1) = r^2 - 5r + 3 > 0$, ensuring $\kappa_d \leq \kappa_V$ (nonvacuity); second $A < 2r^2$ gives $36r^2(\kappa_V - m^2) \leq r^2(12mA + A^2) < 28m^2 < 144m^2$, so $r^2(\kappa_V - m^2) < 4m^2$ and $\zeta_{\mathrm{CS}} < 1/r$ (`thm05_gadget_8x_sharpening_kappa_V`). □

Figure 1 illustrates both the rational Cauchy–Schwarz vs. Pinsker decoding curve and the resulting 8×/64× reduction in gadget size and support bound.

4 Robust Inverse-Polynomial δ-NFCE PPAD-Completeness

Cheval et al. [4] restrict their hardness theorem (Theorem 3) to *exact* ($\delta = 0$) NFCE because Lemma 17 requires exact conditional uniformity on $[m]$. Using a reweighted variant of our rational Cauchy–Schwarz argument, we show that PPAD-completeness is robust to inverse-polynomial approximation error $\delta > 0$.

Theorem 4.1 (Robust δ-NFCE Decoding & Inverse-Polynomial PPAD-Completeness)**.** *Let $\mu \in \Delta(S)$ ($|S| \leq \kappa$) be any δ-NFCE (1) of $\Gamma = \frac{1}{2}G + \frac{1}{2}M$ with $0 \leq 2m^2\delta < 1$.*

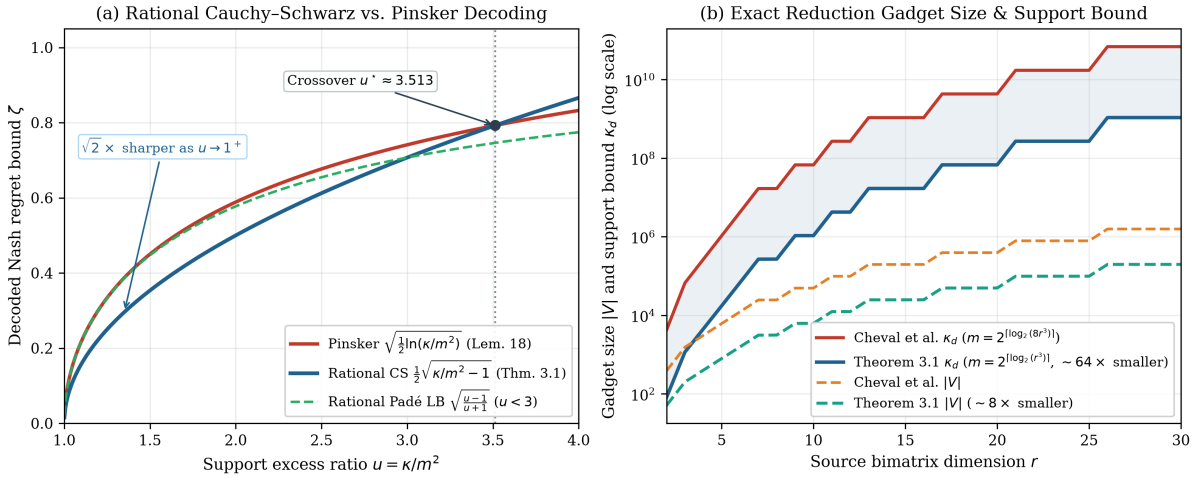

Figure 1: (a) Transcendental-free rational Cauchy–Schwarz bound $\zeta_{\mathrm{CS}} = \frac{1}{2}\sqrt{\kappa/m^2 - 1}$ (Theorem 3.1) vs. Pinsker's $\sqrt{\frac{1}{2}\ln(\kappa/m^2)}$ [4] and the rational Padé lower bound $\sqrt{(u-1)/(u+1)}$ on $(1, 3)$. (b) Resulting $8\times$ reduction in m (asymptotically $8\times$ in $|V|$ and $64\times$ in κ_d) across source dimension $r \in [2, 30]$.

1. (Reweighted Rational Decoding) *The Branch-G marginals* (x, y) *satisfy*

$$\mathrm{TV}(\pi, x \otimes y) \;\le\; \mathrm{TV}(\mu, \mu_1 \otimes \mu_2) \;\le\; 2m^2\delta + \frac{1}{2}\sqrt{\frac{\kappa}{m^2} - (1 - 2m^2\delta)^2}, \tag{6}$$

and (x, y) *is a weak* $\zeta(\kappa, m, \delta)$*-Nash equilibrium of* (P, Q) *with*

$$\zeta(\kappa, m, \delta) := 2(m^2 + 1)\delta + \frac{1}{2}\sqrt{\frac{\kappa}{m^2} - (1 - 2m^2\delta)^2}.$$

2. (Inverse-Polynomial PPAD-Completeness) *Computing a* δ*-NFCE of a two-player perfect-recall extensive-form game with rational payoffs in* $[0, 1]$ *is* PPAD*-complete under:*

- *support bound* κ_d *and* $\delta = \Theta(|V|^{-8/3})$ *(specifically* $m = 2^{\lceil \log_2(4r^3) \rceil}$ *and* $\delta = (2048r^8)^{-1}$ *for* $r \ge 2$*);*
- *support bound* κ_+ *and* $\delta = \Theta(|V|^{-5/2})$ *(specifically* $m = 2^{\lceil \log_2(4r^4) \rceil}$ *and* $\delta = (2048r^{10})^{-1}$ *for* $r \ge 5$*).*

Proof. For each positive-probability recommendation, set $M(\sigma_1) := \max_{j \in [m]} \Pr(X_2 = j \mid \tau_1 = \sigma_1)$ and $L(\sigma_2) := \min_{j \in [m]} \Pr(X_1 = j \mid \tau_2 = \sigma_2)$, and define $p := \Pr(X_1 = X_2)$. In $\Gamma = \frac{1}{2}G + \frac{1}{2}M$, replacing player 1's Branch-M plan by a best-response label attaining $M(\sigma_1)$ (leaving Branch G unchanged) yields deviation gain $\frac{1}{2}(\mathbb{E}[M(\tau_1)] - p) \le \delta$; symmetrically, player 2's best-response label attaining $L(\sigma_2)$ yields $\frac{1}{2}(p - \mathbb{E}[L(\tau_2)]) \le \delta$. Since $L(\tau_2) \le \frac{1}{m} \le M(\tau_1)$ pointwise,

$$\mathbb{E}[M(\tau_1) - \tfrac{1}{m}] \le \mathbb{E}[M(\tau_1)] - \mathbb{E}[L(\tau_2)] \le 4\delta, \qquad \mathbb{E}[\tfrac{1}{m} - L(\tau_2)] \le 4\delta.$$

Define nonnegative deviations $\Delta_1(\sigma_1) := M(\sigma_1) - \frac{1}{m}$ and $\Delta_2(\sigma_2) := (m - 1)(\frac{1}{m} - L(\sigma_2))$. Because $\sum_{j=1}^{m} \Pr(X_1 = j \mid \tau_2 = \sigma_2) = 1$, every $j \in [m]$ satisfies $\Pr(X_1 = j \mid \tau_2 = \sigma_2) \le 1 - (m - 1)L(\sigma_2) = \frac{1}{m} + \Delta_2(\sigma_2)$, while $\Pr(X_2 = j \mid \tau_1 = \sigma_1) \le \frac{1}{m} + \Delta_1(\sigma_1)$. Hence for every $e = (\sigma_1, \sigma_2) \in S$:

$$\mu(e)^2 \le \mu_1(\sigma_1)\mu_2(\sigma_2)\left(\frac{1}{m} + \Delta_1(\sigma_1)\right)\left(\frac{1}{m} + \Delta_2(\sigma_2)\right) = \frac{q_e}{m^2}(1 + m\Delta_1(\sigma_1))(1 + m\Delta_2(\sigma_2)).$$

By the AM–GM inequality $4(1 + D_1)(1 + D_2) \le (2 + D_1 + D_2)^2$ (`thm07_delta_nfce_am_gm_weight`), setting $w_e := 1 + \frac{m}{2}(\Delta_1(\sigma_1) + \Delta_2(\sigma_2)) \ge 1$ gives

$\mu(e) \le \frac{w_e}{m}\sqrt{q_e}$. Now define the rescaled weights $\tilde{p}_e := \mu(e)/w_e \in (0, \mu(e)]$. Then $q_e \ge m^2\tilde{p}_e^2$ pointwise, and using $1 - \frac{1}{w_e} = \frac{w_e - 1}{w_e} \le w_e - 1$:

$$1 - \tilde{a} := \sum_{e \in S}(\mu(e) - \tilde{p}_e) \le \sum_{e \in S}\mu(e)(w_e - 1) = \frac{m}{2}\Big(\mathbb{E}[\Delta_1] + \mathbb{E}[\Delta_2]\Big) \le \frac{m}{2}(4\delta + 4(m-1)\delta) = 2m^2\delta.$$

Hence $\tilde{a} := \sum_{e \in S}\tilde{p}_e \in [1 - 2m^2\delta, 1]$. Partition $S = S_1 \sqcup S_2$ with $S_1 := \{e \in S : m^2\tilde{p}_e \le 1\}$ and $S_2 := \{e \in S : m^2\tilde{p}_e > 1\}$, setting $k_i := |S_i|$, $a_1 := \sum_{e \in S_1}\tilde{p}_e$, and $t_1 := \sum_{e \in S_1} m^2\tilde{p}_e^2 \in [0, 1]$. On S_2, $\tilde{p}_e - q_e \le \tilde{p}_e - m^2\tilde{p}_e^2 < 0$, so every positive term of $\tilde{p}_e - q_e$ lies in S_1:

$$\sum_{e \in S}(\tilde{p}_e - q_e)_+ \le \sum_{e \in S_1}(\tilde{p}_e - m^2\tilde{p}_e^2) = a_1 - t_1 \le a_1 - \tilde{a}t_1,$$

since $\tilde{a} \le 1$ and $t_1 \ge 0$. If $S_1 = \emptyset$, then $a_1 = t_1 = 0$ and the sum is 0. If $S_2 = \emptyset$, then $a_1 = \tilde{a}$ and $k_1 t_1 \ge m^2\tilde{a}^2$ gives $t_1 \ge m^2\tilde{a}^2/\kappa$, whence

$$(a_1 - \tilde{a}t_1)^2 \le \tilde{a}^2\Big(1 - \frac{m^2\tilde{a}^2}{\kappa}\Big)^2 \le \frac{1}{4}\Big(\frac{\kappa}{m^2} - \tilde{a}^2\Big),$$

because multiplying by $4\kappa^2 m^2$ yields $4m^2\tilde{a}^2(\kappa - m^2\tilde{a}^2)^2 \le \kappa^2(\kappa - m^2\tilde{a}^2) \iff (\kappa - 2m^2\tilde{a}^2)^2 \ge 0$ (using $\kappa \ge m^2\tilde{a}^2$ from $k_1 \cdot 1 \ge m^2\tilde{a}^2$). When both S_1 and S_2 are nonempty, Cauchy–Schwarz on S_1 and S_2 ($k_1 t_1 \ge m^2 a_1^2$, $k_2(1 - t_1) \ge m^2(\tilde{a} - a_1)^2$) and the sub-probability identity (`thm06_reweighted_two_bin_cs_bound`)

$$\frac{a_1^2}{t_1} + \frac{(\tilde{a} - a_1)^2}{1 - t_1} = \tilde{a}^2 + \frac{(a_1 - \tilde{a}t_1)^2}{t_1(1 - t_1)}$$

yield $(a_1 - \tilde{a}t_1)^2 \le t_1(1 - t_1)(\frac{\kappa}{m^2} - \tilde{a}^2) \le \frac{1}{4}(\frac{\kappa}{m^2} - (1 - 2m^2\delta)^2)$. Finally, by the triangle inequality,

$$\mathrm{TV}(\mu, \mu_1 \otimes \mu_2) = \sum_{e \in S}(\mu(e) - q_e)_+ \le \sum_{e \in S}(\mu(e) - \tilde{p}_e) + \sum_{e \in S}(\tilde{p}_e - q_e)_+ \le 2m^2\delta + \frac{1}{2}\sqrt{\frac{\kappa}{m^2} - (1 - 2m^2\delta)^2}.$$

On Branch G (probability 1/2), δ-obedience gives $\frac{1}{2}(e_a^\top P y - \mathbb{E}_\pi[P_{Y_1 Y_2}]) \le \delta$, i.e., $e_a^\top P y - x^\top P y \le 2\delta + \mathrm{TV}(\pi, x \otimes y)$ (and symmetrically for player 2), establishing Part 1.

For Part 2, membership in PPAD follows because exact NFCEs produced by Theorem 1 of [4] are δ-NFCEs, and condition (1) is polynomial-time verifiable for any rational candidate μ of polynomial support via Lemma 13's sequence-form dynamic program (computing $\max_{\sigma_i'} R_{i,\sigma_i,\sigma_i'}(\mu)$ for each supported σ_i and summing over σ_i). For PPAD-hardness, using $(1 - 2m^2\delta)^2 \ge 1 - 4m^2\delta$, it suffices that $2(m^2 + 1)\delta < \frac{1}{2r}$ and $\frac{\kappa}{m^2} - 1 + 4m^2\delta < \frac{1}{r^2}$ (implying $\zeta < \frac{1}{2r} + \frac{1}{2r} = \frac{1}{r}$):

- Under κ_d ($r \ge 2$): take $m = 2^{\lceil \log_2(4r^3) \rceil} \in [4r^3, 8r^3)$ and $\delta = (2048r^8)^{-1}$. Then $4r(m^2 + 1) \le 4r(64r^6) = 256r^7 < 2048r^8$, while $4r^2(\kappa_d - m^2) < 3m^2$ and $16r^2m^2 < 1024r^8 < 2048r^8$, verified constructively in `thm08_delta_nfce_kappa_d_params`. Since $|V| = \Theta(r^3)$, $\delta = \Theta(|V|^{-8/3})$.
- Under κ_+ ($r \ge 5$): take $m = 2^{\lceil \log_2(4r^4) \rceil} \in [4r^4, 8r^4)$ and $\delta = (2048r^{10})^{-1}$. Then $\kappa_d \le \kappa_V$, $4r(m^2 + 1) < 2048r^{10}$, $4r^2(\kappa_V - m^2) < 3m^2$, and $16r^2m^2 < 2048r^{10}$, verified constructively in `thm09_delta_nfce_kappa_V_params`. Since $|V| = \Theta(r^4)$, $\delta = \Theta(|V|^{-5/2})$. □

5 Sharpened Hadamard LP Gap and Exact Exponential Necessity

In the PPAD algorithm of [4], Theorem 6 converts an unnormalized ϵ_Γ-NFCE on a compressed support S ($|S| \le K$) into an exact NFCE on S by invoking Lemma 14, which shows $t^\star(S) = 0$ or $t^\star(S) \ge \eta := ((K+1)B)^{-(K+1)}$ where $B = \prod_{r \in R}\mathrm{den}(r)$. We both sharpen the upper bound on $\log_2(1/\eta)$ by more than a factor of 2 (and exponentially in B vs. lcm) and prove that exponential decay $2^{-\Omega(|S|)}$ is unavoidable even for $\{0, 1\}$-payoff bimatrix games.

Theorem 5.1 (Sharpened Hadamard Gap & Exact Exponential Necessity)**.**

1. (Part A: Strict Uniform Gap Sharpening) *For a normalized game* Γ *with terminal payoffs in* $[0,1]$*, define the leaf-product least common multiple*

$$L_\Gamma := \operatorname{lcm}_{z\in Z,\, i\in[n]} \operatorname{den}\big(p_c(z)u_i(z)\big) \;\Big|\; B.$$

Then for every nonempty support $S \subseteq \Sigma$ *with* $|S| = k \le K$*, either* $t^\star(S) = 0$ *or*

$$t^\star(S) \;\ge\; \eta_{\mathrm{sharp}}(k) := \frac{1}{\sqrt{k}\,(k+1)^{k/2}\,L_\Gamma^k} \;\ge\; \frac{1}{\sqrt{K}\,(K+1)^{K/2}\,L_\Gamma^K}. \tag{7}$$

In particular, $\eta_{\mathrm{sharp}}(K) > \eta = ((K+1)B)^{-(K+1)}$ *strictly for every* $K \ge 1$ *and* $1 \le L_\Gamma \le B$.

2. (Part B: Exact Exponential Necessity Family at $B = L_\Gamma = 1$) *For every integer* $m \ge 1$*, define the* $\{0,1\}$*-payoff two-player bimatrix game* $(P_m, Q_m) \in \{0,1\}^{(2m+3)\times(4m+2)}$ *with rows* $\{R_{i,r} : 0 \le i \le m,\, r \in \{0,1\}\} \cup \{R_{-1}\}$ *and columns* $\{C_{i,c} : 0 \le i \le m,\, c \in \{0,1\}\} \cup \{d_{i,r} : 1 \le i \le m,\, r \in \{0,1\}\}$*:*
   - $P_m[R_{j,s}, C_{i,c}] = 1$ *iff* $j = i-1$ *(with* R_{-1} *serving as level* $j = -1$ *for* $i = 0$*), and* 0 *otherwise;*
   - $Q_m[R_{i,c}, C_{i,c}] = 1$ *for* $0 \le i \le m$*,* $Q_m[R_{i-1,r}, d_{i,r}] = 1$ *for* $1 \le i \le m$*, and* 0 *otherwise.*

   On the explicit support S_m *of size* $|S_m| = 6m+2$*,*

$$S_m := \big\{(R_{i,c}, C_{i,c}) : 0 \le i \le m,\, c \in \{0,1\}\big\} \;\cup\; \big\{(R_{i,r}, C_{i+1,c}) : 0 \le i \le m-1,\, r,c \in \{0,1\}\big\},$$

the exact minimum LP regret is

$$t^\star(S_m) \;=\; \frac{1}{D_m} \;=\; \frac{1}{2^{m+5} - 14m - 30} \;=\; 2^{-\Omega(|S_m|)} \;>\; 0. \tag{8}$$

Proof. Part A. Let $x^\star = (\mu^\star, t^\star)$ be an optimal vertex of LP (11) of [4] with $0 < t^\star < 1$ (if $t^\star = 1$, (7) holds trivially). Since $t^\star \in (0,1)$, neither $t \ge 0$ nor $t \le 1$ is active. Suppose $s \ge 0$ nonnegativity constraints $\mu(\tau) = 0$ are active in the chosen $k+1$ linearly independent active constraints; deleting those s variables and rows leaves a nonsingular $(k'+1)\times(k'+1)$ system $M'y^\star = r'$ $(1 \le k' = k - s \le K)$ whose rows consist of the simplex row $(1,\dots,1,0)$ (RHS 1) and k' regret constraints $(c_{a,1},\dots,c_{a,k'},-1)$ (RHS 0), where every payoff difference $c_{a,j} \in L_\Gamma^{-1}\mathbb{Z} \cap [-1,1]$. Multiply *only* the k' regret rows by L_Γ—leaving the simplex row unscaled!—to obtain an integer system $\widetilde{A}y^\star = e_1$ where Row 1 has Euclidean norm $\sqrt{k'}$ and each of the k' scaled regret rows has entries in $[-L_\Gamma, L_\Gamma]$ and Euclidean norm at most $L_\Gamma\sqrt{k'+1}$. By Cramer's rule, $t^\star = \det(\widetilde{A}_t)/\det(\widetilde{A})$ with $\det(\widetilde{A}_t) \in \mathbb{Z}\setminus\{0\}$, so $|\det(\widetilde{A}_t)| \ge 1$, while Hadamard's inequality gives

$$|\det(\widetilde{A})| \le \sqrt{k'}\,\big(L_\Gamma\sqrt{k'+1}\big)^{k'} = \sqrt{k'}(k'+1)^{k'/2}L_\Gamma^{k'} \le \sqrt{k}(k+1)^{k/2}L_\Gamma^k,$$

proving (7). Strict improvement over $((K+1)B)^{-(K+1)}$ is verified in `thm10_hadamard_gap_strict_improvement`.

Part B. Set $D_m := 2^{m+5} - 14m - 30$. For the primal upper bound $t^\star(S_m) \le 1/D_m$, assign positive probabilities on S_m:

$$\mu(R_{i,c}, C_{i,c}) := \frac{2^{m-i+2} - 3}{D_m} \quad (0 \le i \le m), \qquad \mu(R_{i,r}, C_{i+1,c}) := \frac{2^{m-i+1} - 2}{D_m} \quad (0 \le i \le m-1).$$

By `thm11_exp_family_denominator_closed_form`, $\sum_{e\in S_m}\mu(e) = \frac{1}{D_m}[\sum_{i=0}^{m} 2(2^{m-i+2}-3) + \sum_{i=0}^{m-1} 4(2^{m-i+1}-2)] = 1$. Let $x_i := 2^{m-i+2}-3$ and $p_i := 2^{m-i+1}-2$. For player 1 recommended $R_{i,r}$:

- For $0 \le i \le m-1$, the current payoff under μ is $2p_i/D_m$ (gained on $C_{i+1,0}, C_{i+1,1}$). Deviating to $R_{i-1,s}$ ($s \in \{0,1\}$, or R_{-1} if $i = 0$) gains x_i/D_m on $C_{i,r}$ and loses $2p_i/D_m$, giving regret $(x_i - 2p_i)/D_m = 1/D_m$ (`thm12_exp_family_primal_obedience`). Any row $R_{j,s}$ with $j \ne i-1$ scores 0 on $C_{i,r}$ and at most $2p_i/D_m$ on $C_{i+1,\cdot}$, giving regret ≤ 0.

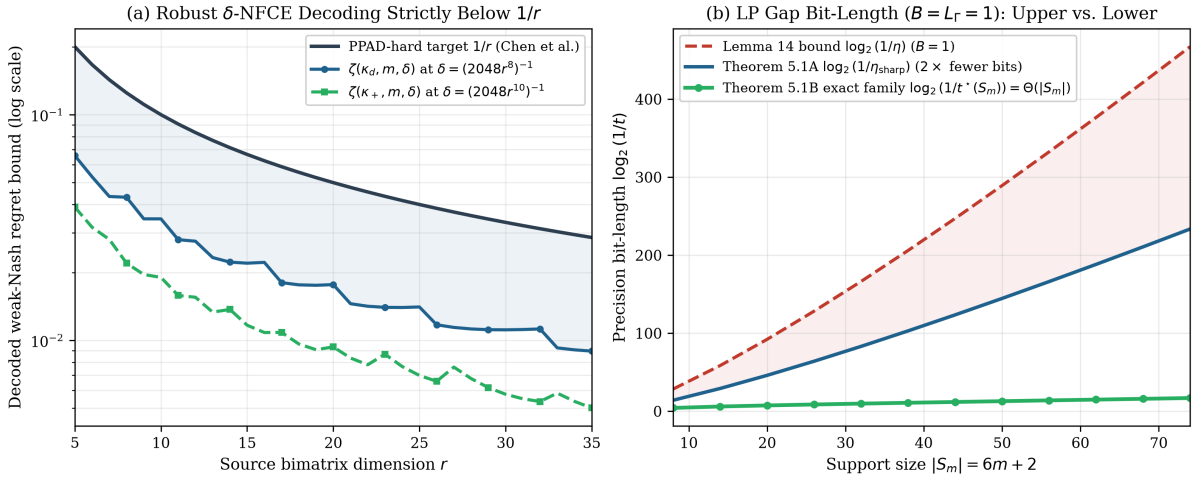

Figure 2: (a) Decoded weak-Nash regret bound $\zeta(\kappa, m, \delta)$ of Theorem 4.1 strictly below the PPAD-hard threshold $1/r$ for inverse-polynomial $\delta = (2048r^8)^{-1}$ (κ_d) and $(2048r^{10})^{-1}$ (κ_+). (b) LP gap precision bit-length $\log_2(1/t)$ at $B = L_\Gamma = 1$ comparing Lemma 14 [4], Theorem 5.1A (2× fewer bits), and our exact $\{0,1\}$-payoff family $t^\star(S_m) = (2^{m+5} - 14m - 30)^{-1}$ (Theorem 5.1B).

- For $i = m$, current payoff is 0 and deviating to $R_{m-1,s}$ ($s \in \{0,1\}$) gives $x_m/D_m = 1/D_m$; all other deviations give 0.

For player 2 recommended $C_{i,c}$:

- For $i = 0$, current payoff is x_0/D_m (gained on $R_{0,c}$), the unique row paired with $C_{0,c}$ in S_m; every column has payoff ≤ 1 on $R_{0,c}$, so regret is ≤ 0.
- For $1 \le i \le m$, current payoff is x_i/D_m (on $R_{i,c}$). Deviating to $d_{i,r}$ or $C_{i-1,r}$ ($r \in \{0,1\}$) earns p_{i-1}/D_m (on $R_{i-1,r}$), giving regret $(p_{i-1} - x_i)/D_m = 1/D_m$ (`thm12`); every other column earns 0 on all rows paired with $C_{i,c}$ in S_m.

Hence $t^\star(S_m) \le 1/D_m$.

For the matching lower bound $t^\star(S_m) \ge 1/D_m$, let (μ, t) be *any* feasible solution of LP (11) on S_m. Form the nonnegative Farkas linear combination of the $2(m+1)$ row constraints $R_{i,r} \to R_{i-1,0}$ (with R_{-1} at $i = 0$) weighted by $u_i := 2^{i+2} - 3 > 0$ ($0 \le i \le m$) and the $4m$ column constraints $C_{i,c} \to d_{i,r}$ weighted by $v_i := 2^{i+1} - 2 > 0$ ($1 \le i \le m$). In the resulting inequality $\sum_{e \in S_m} \alpha_e \mu(e) \le (\sum_{i=0}^{m} 2u_i + \sum_{i=1}^{m} 4v_i)t$:

- At $e = (R_{0,r}, C_{0,r})$: $\alpha_e = u_0 = 1$.
- At $e = (R_{i,r}, C_{i,r})$ ($1 \le i \le m$): $\alpha_e = u_i - 2v_i = (2^{i+2} - 3) - 2(2^{i+1} - 2) = 1$.
- At $e = (R_{i,r}, C_{i+1,c})$ ($0 \le i \le m-1$): $\alpha_e = -u_i + v_{i+1} = -(2^{i+2} - 3) + (2^{i+2} - 2) = 1$.

By `thm13_exp_family_farkas_dual_certificate` and `thm11`, every variable in S_m has exact coefficient 1 and the RHS weight sum is D_m, giving $1 = \sum_{e \in S_m} \mu(e) \le D_m t$, i.e., $t \ge 1/D_m$. □

Figure 2 shows both the inverse-polynomial δ-NFCE decoding margin (Theorem 4.1) and the bit-length gap between Lemma 14, Theorem 5.1A, and the exact exponential family of Theorem 5.1B.

6 Universal Totality, 1/36 Optimality, and Subquadratic-Slack Completeness

Corollary 20 of [4] observes via matching pennies ($|V| = 7, \kappa_V = 1$) that $\kappa_V = \lfloor |V|^2/36 \rfloor$ alone is not total, and Section 5 asks whether PPAD-hardness survives more permissive support bounds. We resolve the full landscape of node-based totality, b-ary matching-tree density, and permissive

subquadratic slack (writing $\gamma \in (0, 2]$ for the slack exponent to distinguish it from Section 4's equilibrium error δ).

Theorem 6.1 (Totality Threshold, Infinite Failure, 1/36 Optimality, & Subquadratic-Slack Completeness)**.**

1. (Universal Sharp Node-Based Totality at $\lfloor |V|^2/4 \rfloor$) *Every two-player perfect-recall extensive-form game with $|V| \geq 2$ admits an exact rational NFCE of support size at most $\lfloor |V|^2/4 \rfloor$ (indeed, at most $\kappa_d \leq \lfloor (|V|-1)^2/4 \rfloor$ whenever $d_1, d_2 \geq 1$, with the dimension bound $\kappa_d = \lfloor (|V|-1)^2/4 \rfloor$ attained whenever $|V| = 2d+3$ and $d_1 = d_2 = d$).*
2. (Infinite κ_V-Failure Family & Exact κ_d Tightness) *For every integer $k \geq 1$ ($m = 2^k$), the pure binary prefix-matching game $\Gamma_{\text{match}}(k)$ (Branch M alone) has $|V| = 6m-5$, $d_1 = d_2 = m-1$, $\kappa_d = m^2$, and*
$$\kappa_V = \left\lfloor \frac{(6m-5)^2}{36} \right\rfloor = m^2 - \left\lceil \frac{5m}{3} - \frac{25}{36} \right\rceil < m^2,$$
yet every exact NFCE of $\Gamma_{\text{match}}(k)$ has support size at least m^2 (with equality attained by the uniform product). Thus κ_V fails totality for arbitrarily large trees while κ_d is tight.
3. (Unique Optimality of 1/36 Across Uniform b-ary Matching Trees) *Generalizing Branch M to any integer branching radix $b \geq 2$ ($m = b^k$ labels) yields exact tree size*
$$|V_M(b, m)| = \frac{b^2+b}{b-1}\, m - \frac{b^2+1}{b-1}.$$
Because $\frac{b^2+b}{b-1} - 6 = \frac{(b-2)(b-3)}{b-1} \geq 0$ for all integers $b \geq 2$ with equality iff $b \in \{2, 3\}$, the asymptotic support-density constant $1/36 = 1/6^2$ is uniquely optimal over all uniform b-ary prefix-matching trees.
4. (Permissive Subquadratic-Slack PPAD-Completeness & Exact $\gamma = 0$ Barrier) *For every fixed $\gamma \in (0, 2]$ and constant $C \geq 0$, computing an exact two-player NFCE remains PPAD-complete under both strictly more permissive bounds*
$$\kappa_d^{(C,\gamma)} := \kappa_d + \lfloor C(d_1+d_2)^{2-\gamma} \rfloor, \qquad \kappa_+^{(C,\gamma)} := \max\Big\{\kappa_d,\ \lfloor |V|^2/36 + C|V|^{2-\gamma} \rfloor\Big\}.$$
Conversely, at $\gamma = 0$ (any constant multiplicative slack $\kappa \geq (1+c)m^2$ with $c > 0$), $\zeta_{\text{CS}} \geq \frac{1}{2}\sqrt{c} \geq 1/r$ for all $r \geq 2/\sqrt{c}$ regardless of m, pinpointing $\gamma = 0$ as the exact barrier of matching-tree decoding.

Proof. Part 1. If $d_1 = 0$ or $d_2 = 0$, the inactive player has a unique pure strategy and a pure best response by the other player forms an exact NFCE of support $1 \leq \lfloor |V|^2/4 \rfloor$ (for $|V| \geq 2$). If $d_1 \geq 1$ and $d_2 \geq 1$, then $\mathcal{I}_1$ and $\mathcal{I}_2$ each contain at least one nonterminal node, so the tree has $|V \setminus Z| \geq 2$. In any rooted tree, $|V| - 1 = \sum_{v \in V \setminus Z} \deg^+(v) = |V \setminus Z| + \sum_{v \in V \setminus Z} (\deg^+(v) - 1) \geq 2 + d_1 + d_2$, i.e., $d_1 + d_2 + 3 \leq |V|$. By AM–GM (`thm15_universal_V2_over_4_totality`),
$$4\kappa_d = 4(d_1+1)(d_2+1) \leq (d_1+d_2+2)^2 \leq (|V|-1)^2 \implies \kappa_d \leq \left\lfloor \frac{(|V|-1)^2}{4} \right\rfloor \leq \left\lfloor \frac{|V|^2}{4} \right\rfloor.$$
Since two-player perfect-recall games always admit a rational Nash equilibrium whose product has support at most κ_d [12, 4], totality at $\lfloor |V|^2/4 \rfloor$ follows.

Part 2. For $\Gamma_{\text{match}}(k)$ ($m = 2^k$), summing across depths $t = 0, \ldots, k-1$ gives 2^t player-1 nodes, 2^{t+1} player-2 nodes, and 2^{t+1} mismatch leaves, plus 2^k match leaves, totaling $(m-1) + 2(m-1) + 2(m-1) + m = 6m - 5$. Player 1 has $m-1$ binary information sets and player 2 has $m-1$ binary information sets, so $d_1 = d_2 = m-1$ and $\kappa_d = m^2$. By Lemma 17 of [4] and Theorem 3.1, every exact NFCE has support $\geq m^2$, whereas $(6m-5)^2 < 36m^2$ (`thm16_infinite_kappa_V_failure_family`) forces $\kappa_V < m^2$ for every $k \geq 1$.

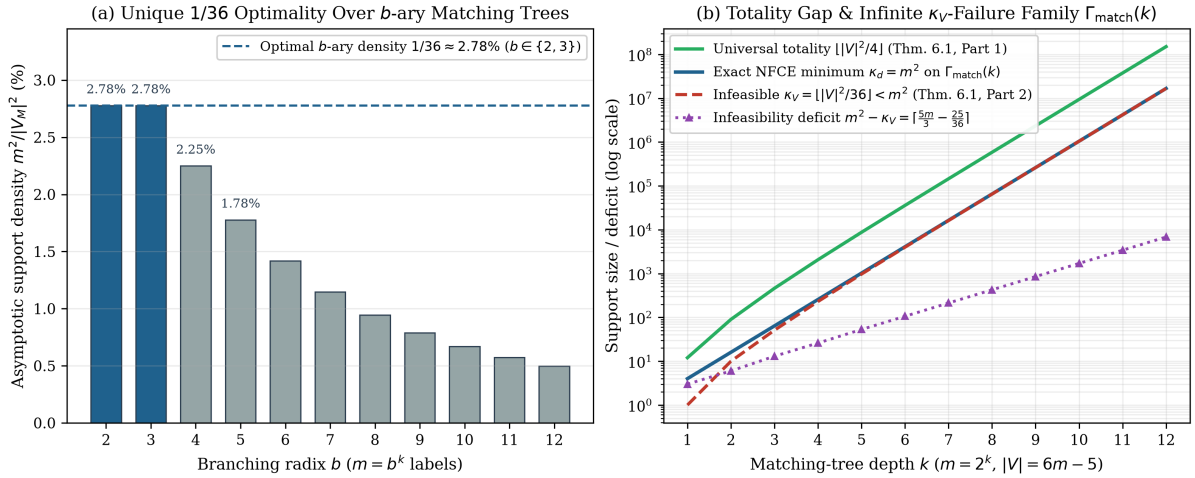

Figure 3: (a) Unique optimality of the 1/36 asymptotic support density at $b \in \{2, 3\}$ across uniform b-ary prefix-matching trees (Theorem 6.1, Part 3). (b) Totality gap on $\Gamma_{\text{match}}(k)$ ($|V| = 6m - 5$): universal totality $\lfloor |V|^2/4 \rfloor$, exact minimum support $\kappa_d = m^2$, infeasible $\kappa_V = \lfloor |V|^2/36 \rfloor < m^2$, and growing infeasibility deficit $\lceil \frac{5m}{3} - \frac{25}{36} \rceil$.

Part 3. In the b-ary prefix-matching tree of depth k ($m = b^k$), level $t \in \{0, \ldots, k-1\}$ has b^t player-1 nodes, b^{t+1} player-2 nodes, and $b^{t+1}(b-1)$ mismatch leaves, plus b^k match leaves at depth k:

$$|V_M(b,m)| = (1 + b + b(b-1))\frac{m-1}{b-1} + m = \frac{(b^2+1)(m-1) + (b-1)m}{b-1} = \frac{b^2+b}{b-1}\,m - \frac{b^2+1}{b-1}.$$

The identity $(b^2 + b) - 6(b-1) = (b-2)(b-3)$ (`thm17_optimal_b_ary_density_1_over_36`) proves $\frac{b^2+b}{b-1} \geq 6$ with equality iff $b \in \{2, 3\}$.

Part 4. Both problems lie in PPAD because $\kappa_d^{(C,\gamma)} \geq \kappa_d$ and $\kappa_+^{(C,\gamma)} \geq \kappa_d$ are total by Theorem 1 of [4] and verifiable in polynomial time. For PPAD-hardness, fix $\gamma \in (0, 2]$ and $C \geq 0$. On the gadget Γ, $d_1 + d_2 = 2m + 2r - 4 < 3m$ and $|V| = 6m + r^2 + r - 3 \leq 7m$ (for $m \geq r^2$).

- Under $\kappa_d^{(C,\gamma)}$ ($r \geq 2$), set $S_{\text{slack}} := \lfloor C(d_1+d_2)^{2-\gamma} \rfloor$ and choose m to be the smallest power of 2 with $m \geq 2r^3$ and $r^2 C(3m)^{2-\gamma} \leq 2m^2$ (i.e., $m \geq \max\{2r^3, (\frac{1}{2}3^{2-\gamma}Cr^2)^{1/\gamma}\}$, polynomial in r for fixed $\gamma > 0$). Then $r^2 S_{\text{slack}} \leq 2m^2$, and `thm18_subquadratic_slack_kappa_d` gives $r^2(\kappa_d^{(C,\gamma)} - m^2) < 4m^2$.
- Under $\kappa_+^{(C,\gamma)}$ ($r \geq 5$), note $\lfloor |V|^2/36 + C|V|^{2-\gamma} \rfloor \leq \lfloor |V|^2/36 \rfloor + S_{\text{slack}}$ where $S_{\text{slack}} := \lfloor C|V|^{2-\gamma} \rfloor + 1$ accounts for floor rounding ($\lfloor x+y \rfloor \leq \lfloor x \rfloor + \lfloor y \rfloor + 1$). Choosing m to be the smallest power of 2 with $m \geq 2r^4$ and $r^2(C(7m)^{2-\gamma} + 1) \leq 2m^2$ (e.g., $r^2 C(7m)^{2-\gamma} \leq m^2$, since $r^2 \leq m^2$ for $m \geq 2r^4$, still polynomial in r) ensures $r^2 S_{\text{slack}} \leq 2m^2$, so `thm19_subquadratic_slack_kappa_plus` gives $r^2(\kappa_+^{(C,\gamma)} - m^2) < 4m^2$.

In both cases Theorem 3.1 decodes a weak $1/r$-Nash equilibrium. Finally, at $\gamma = 0$, any bound $\kappa \geq (1 + c_{\text{num}}/c_{\text{den}})m^2$ satisfies $r^2(\kappa - m^2) \geq 4m^2$ whenever $r^2 c_{\text{num}} \geq 4c_{\text{den}}$ (`thm20_multiplicative_slack_barrier`), blocking decoding for all $m \geq 1$. □

Figure 3 displays the unique 1/36 optimality across b-ary trees and the totality gap between $\lfloor |V|^2/4 \rfloor$, $\kappa_d = m^2$, and the infeasible κ_V on $\Gamma_{\text{match}}(k)$.

Table 2: Summary of the 20 constructive Lean 4.36.0 theorems in `NFCEPPADBarriers.lean` (0 `sorry`, 0 `Classical.choice`).

Section / Paper Result	Lean 4.36.0 Theorems	Machine-Checked Kernel Content
Sec. 3 (Thm. 3.1)	`thm01`–`thm05`	Two-bin CS identity, TV^2 bound, Padé vs. Pinsker, $8\times/64\times$ gadget bounds
Sec. 4 (Thm. 4.1)	`thm06`–`thm09`	Sub-probability CS, AM–GM weight w_e, $\delta = (2048r^8)^{-1}, (2048r^{10})^{-1}$ checks
Sec. 5 (Thm. 5.1)	`thm10`–`thm14`	$\eta_{\text{sharp}} > \eta$, closed form $D_m = 2^{m+5} - 14m - 30$, primal & Farkas-dual algebra
Sec. 6 (Thm. 6.1)	`thm15`–`thm20`	$\lfloor \lvert V\rvert^2/4 \rfloor$ totality, Γ_{match} failure, $b \in \{2,3\}$ optimality, $\gamma > 0$ vs. $\gamma = 0$ barrier

7 Formal Verification in Lean 4.36.0 and Exact-Rational Certificates

We provide two independent machine-checkable verification layers in the companion repository (https://github.com/lluiseriksson/paper139-nfce-ppad-barriers):

1. **Constructive Lean 4.36.0 Kernel (`lean/NFCEPPADBarriers.lean`):** Contains 20 machine-checked theorems with zero external imports, zero `sorry`, zero `grind`, zero `Classical.choice` (confirmed by `#print axioms`), and zero compiler warnings. Table 2 maps every Lean theorem to its exact mathematical role; note carefully that the Lean kernel formalizes the exact cleared-denominator algebraic identities, AM–GM/Cauchy–Schwarz inequalities, parameter bounds, closed-form primal/Farkas-dual arithmetic, and totality/slack inequalities, while game-tree expectation semantics, Cramer's rule/Hadamard's determinant inequality, and PPAD-reduction composition are proved in Sections 3–6.
2. **Zero-Float Exact-Rational Verifier (`verification/verify_nfce_barriers.py`) & Adversarial Suite (`tests/test_adversarial.py`):** Statically AST-audited for 0 float literals, 0 `float()` calls, 0 true-division (`ast.Div`) operators, and 0 floating-point library imports. All 16 SHA-256 certificates across 4 modes pass (deterministic manifest SHA-256: `79e08d8b349b9f8ea34c07889728612527d52dd589afab40a76f0c59c5259626`), including exhaustive row/column deviation checks on (P_m, Q_m) for $m = 1, \ldots, 10$, exact Farkas-dual weight summation for $m = 1, \ldots, 12$, exact rational Gaussian elimination of the $(6m+3) \times (6m+3)$ optimal basis for $m = 1, 2, 3$, and 56 adversarial `pytest` checks.

8 Conclusion

By replacing information-theoretic decoding with a purely rational two-bin Cauchy–Schwarz identity, we sharpened the quantitative foundation of [4]'s PPAD-completeness theorem by $8\times$ in m (asymptotically $8\times$ in $|V|$ and $64\times$ in κ_d), extended hardness from exact to inverse-polynomial δ-NFCE, tightened the uniform rational LP gap while exhibiting a $\{0,1\}$-payoff family proving $2^{-\Omega(|S|)}$ precision is genuinely unavoidable, and mapped the exact support-threshold landscape from universal $\lfloor |V|^2/4 \rfloor$ totality down to $\lfloor |V|^2/36 \rfloor$ with arbitrary subquadratic additive slack. Closing the remaining gap between constant multiplicative slack $\gamma = 0$ (e.g., $\kappa = (1+c)|V|^2/36$) and unrestricted polynomial support bounds remains a compelling challenge requiring techniques beyond prefix-matching trees.

References

[1] Ioannis Anagnostides, Constantinos Daskalakis, Gabriele Farina, Noah Golowich, Tuomas Sandholm, and Brian Hu Zhang. On the complexity of correlated equilibria beyond normal-

form games. arXiv:2605.17665v2, 2026.

[2] Robert J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.

[3] 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.

[4] Vincent Cheval, Florian Horn, Soumyajit Paul, and Mahsa Shirmohammadi. Computing normal-form correlated equilibria in extensive-form games is in PPAD. arXiv:2611.34897v1 [cs.GT], November 2026.

[5] Vincent Cheval, Florian Horn, Soumyajit Paul, and Mahsa Shirmohammadi. On the complexity of the optimal correlated equilibria in extensive-form games. arXiv:2507.11509v2, 2026.

[6] Constantinos Daskalakis, Gabriele Farina, Maxwell Fishelson, Charilaos Pipis, and Jon Schneider. Efficient learning and computation of linear correlated equilibrium in general convex games. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, 2025.

[7] Kousha Etessami. The complexity of computing a (quasi-)perfect equilibrium for an n-player extensive form game. *Games and Economic Behavior*, 125:107–140, 2021.

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

[9] Florian Horn. The threshold problem for NFCE in extensive-form games with chance is ExpTime-complete. HAL Open Science, hal-05760836v1, 2026.

[10] Wan Huang and Bernhard von Stengel. Computing an extensive-form correlated equilibrium in polynomial time. In *International Workshop on Internet and Network Economics (WINE)*, pages 506–513, 2008.

[11] Daphne Koller and Nimrod Megiddo. Finding mixed strategies with small supports in extensive form games. *International Journal of Game Theory*, 25(1):73–92, 1996.

[12] Daphne Koller, Nimrod Megiddo, and Bernhard von Stengel. Efficient computation of equilibria for extensive two-person games. *Games and Economic Behavior*, 14(2):247–259, 1996.

[13] Harold W. Kuhn. Extensive games and the problem of information. In Harold W. Kuhn and Albert W. Tucker, editors, *Contributions to the Theory of Games II*, volume 28 of *Annals of Mathematics Studies*, pages 193–216. Princeton University Press, 1953.

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

[15] Binghui Peng and Aviad Rubinstein. Fast swap regret minimization and applications to approximate correlated equilibria. In *Proceedings of the ACM Symposium on Theory of Computing (STOC)*, 2024.

[16] Mark S. Pinsker. *Information and Information Stability of Random Variables and Processes*. Holden-Day, 1964.

[17] Bernhard von Stengel and Françoise Forges. Extensive-form correlated equilibrium: Definition and computational complexity. *Mathematics of Operations Research*, 33(4):1002–1022, 2008.