CERTIFIED BEHAVIORAL MAXMIN: SEMIDEFINITE CERTIFICATION SEPARATES ABSTRACTION LOSS FROM SOLVER LOSS IN IMPERFECT-RECALL GAMES

Anonymous authors
Paper under double-blind review

ABSTRACT

A player who forgets, facing an adversary who does not, is the normal case in deployed game solving: a poker abstraction merges information sets, and a team that cannot communicate knows less than its members do. Her best guaranteed payoff, the behavioral maxmin, is a polynomial optimization problem that is $\exists\mathbb{R}$-hard even when the adversary has perfect recall, and first-order methods return a strategy without saying how far it is from that value. We add a certificate that runs once after the solver and changes nothing about it: a sum-of-squares feasibility program in which the adversary appears as a polynomial response map, so it contributes only linear variables and the semidefinite blocks carry only the player's variables. The degree of the response map defines a ladder: at degree zero the certificate bounds the minmax value, and raising the degree closes the behavioral duality gap. A posterior step makes the floating-point solution rigorous by evaluation alone, and that bound splits the deployment loss of the solver's point into an abstraction part and a solver part, which exploitability cannot distinguish: on Kuhn poker two abstractions, one with perfect recall and one without, lose the same 0.167 against the full game, and the certificate assigns the first loss to the abstraction and the second to CFR stalling. On Leduc abstractions up to 126 variables it bounds CFR's gap by 10^{-4} in minutes where the global solver SCIP leaves 10^{-3} to 10^{-2} after an hour, and on three-player Kuhn teams it certifies a team-maxmin value and an interval for the value of coordination. One certificate thus separates two failures that the standard diagnostic reports as one number.

1 INTRODUCTION

Extensive-form games with imperfect information are the standard model of sequential decision making against an adversary, and what poker and security-game solvers are built on. The full game is rarely what gets solved. Its tree is too large to enumerate, so the methods that make instances tractable are first-order: CFR and its variants converge toward equilibrium by regret minimization (Zinkevich et al., 2007; Tammelin, 2014), and their convergence guarantees assume perfect recall. Deployed players do not play the full game either. A poker solver plays an *abstraction*, a coarsening of the tree in which information sets are merged, chosen because the unabstracted tree is out of reach (Waugh et al., 2009); a team whose members cannot communicate is a single player who does not know what the others know, a structure that arises whenever information is dispersed across agents. Both produce a player who *forgets*: she acts at information sets that were merged or never joined, and no strategy recovers what she has lost. Regret minimization still returns a strategy on such a tree, but nothing in its guarantee survives the loss of recall (Lanctot et al., 2012). What such a player loses by forgetting, and how much of that loss belongs to the abstraction rather than to the solver, is what this paper sets out to measure.

The best payoff such a player can guarantee is the *behavioral maxmin*, the maximum over her behavioral strategies of the minimum over the adversary's. It is a polynomial optimization problem over a product of simplices, and deciding it is $\exists\mathbb{R}$-hard (Gimbert et al., 2020), even when the adversary has perfect recall (Tewolde et al., 2024). CFR returns a strategy regardless; it does not say how far that strategy is from the maxmin. What practitioners compute instead is exploitability, the loss of the returned strategy against a best response, and the two ways a deployed strategy can lose are

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

indistinguishable in it. Two abstractions of Kuhn poker merge the first player's cards into buckets: one at both decisions, keeping perfect recall, the other only at the second, so the player forgets her card. On each, CFR against a perfect-recall opponent returns a strategy that loses 0.167 per hand against the unabstracted game. In the first, the abstraction itself cannot do better; in the second, the best strategy in the abstraction is as good as that of the full game and CFR has stalled. The remedies are opposite, more buckets against more iterations, and exploitability reports one number.

A per-instance bound on the gap would settle the choice, and it must be a posteriori: the value of the abstraction is exactly what is hard to compute. The split matters as much as the bound, because it says which remedy applies.

This paper adds a certificate that runs once, after the solver and without changing it, and bounds the gap between the solver's point and the behavioral maxmin. It is a sum-of-squares feasibility program, but not the one obtained by writing the maxmin as a polynomial program and applying the moment hierarchy of Lasserre (2001): folding the adversary's linear program into the outer problem by strong duality would put dual variables into the moment basis. We replace them by a *polynomial response map* τ, which sends the player's strategy μ to a sequence-form plan of the adversary, under a restricted ansatz (Section 4). The adversary then contributes only linear variables and the semidefinite blocks carry only the player's variables, which is what makes Leduc-sized instances solvable. The degree r of the map is a design parameter with a game-theoretic reading: at $r = 0$ the adversary plays a constant plan and the certificate bounds the *minmax* value, and raising r descends the behavioral duality gap. The level of the cone is a second parameter of a different kind: it controls the polynomial representation of the constraints, and raising it at a fixed degree tightens the bound. A posterior step turns the solver's floating-point certificate into a rigorous bound by evaluation alone.

On 25 Kuhn instances – 13 imperfect-recall abstractions, 11 perfect-recall abstractions (7 of them lossy) and the full game – 1800 certificates are sound against the global solver SCIP. The ladder is visible: on the abstraction that forgets the card at the second decision, degree 0 certifies 0.250, the minmax gap, and degree 1 closes to the true gap 0.222. On Leduc abstractions with 84 and 126 variables against 1093 adversary sequences, the certificate returns numerical gaps below 10^{-4} in 6 and 26 minutes, which independent exact re-solves then prove rigorously, where SCIP's dual bound is still 3×10^{-3} and 10^{-2} away after an hour. On the Kuhn instances a certificate costs at most 0.2 s at the lowest degree and 23 s at the highest. For three-player Kuhn teams playing without communication, the same program gives a convex certificate of a team-maxmin value and an interval for the value of coordination. Pruning the adversary to a support S certifies the maxmin against an adversary restricted to S, no more.

The paper makes three contributions.

- We introduce a certificate for behavioral maxmin whose adversary is a polynomial response map, so that the semidefinite blocks carry only the player's variables. The degree of the map gives a hierarchy that starts at the minmax value and converges to the behavioral maxmin, and the level of the cone is a second limit that refines the same bound (Sections 4 and 6).
- We turn the solver's floating-point output into a rigorous upper bound on the behavioral maxmin by evaluation alone, charging the response map's infeasibility through row-wise weak duality, and we characterize exactly what restricting the adversary to a support certifies (Theorem 2 and Section 5).
- We separate the loss of a deployed strategy into what the abstraction or the missing coordination costs and what the solver leaves behind, and apply the split to imperfect-recall poker abstractions and to three-player teams that cannot communicate (Section 6).

2 Related Work

Imperfect recall and its maxmin. The interpretational puzzles of forgetting go back to Piccione & Rubinstein (1997) and Aumann et al. (1997); the computational side is settled as hard: Koller & Megiddo (1992) for two-player maxmin, Gimbert et al. (2020) through the correspondence with polynomial optimization, Tewolde et al. (2023; 2024) for the single-player and multi-player equilibrium landscapes, Berker et al. (2025) for the value of restoring recall. Bosansky et al. (2017) and Cermák et al. (2017) compute the two-player maxmin by bilinear programming with branch-and-bound, which yields anytime bounds; the certificate here is a single convex program whose

tightness is governed by a level and a response degree, and it takes any candidate and bounds its gap. Carminati et al. (2026) give exact algorithms for timeable two-player imperfect-recall games in mixed strategies (the correlated team semantics); our object is the behavioral maxmin, the value a deployed abstraction or an uncoordinated team obtains, which is smaller and $\exists\mathbb{R}$-hard even for perfect-recall opponents. Imperfect recall as an abstraction device is due to Waugh et al. (2009); Lanctot et al. (2012) study regret minimization on such trees and Kroer & Sandholm (2016; 2014) give a priori bounds on abstraction quality in terms of error terms of the abstraction. Ours is a posteriori, per instance, and bounds the part of the loss those analyses cannot see: the solver's.

Sums of squares in games. Laraki & Lasserre (2012) relax $\min_y \max_x p(x, y)$ over products of compact semialgebraic sets by a moment hierarchy in which both variables enter the moment basis, and prove its convergence. The certificate here is not an instance of that hierarchy: it restricts the adversary to a polynomial decision rule $x \mapsto \tau(x) \in \mathcal{Y}$ of the kind used for multistage adaptive optimization (Bertsimas et al., 2011), and keeps only the player's moments, so the adversary's sequence-form constraints enter as linear coefficients; its convergence (Proposition 1) is proved by Bernstein approximation of a continuous near-optimal response and Handelman's theorem. Parrilo (2006) and Nie & Tang (2024) apply sums of squares to polynomial games and Nash problems; Zheng et al. (2026) propose moment hierarchies for the single-player ex-ante problem and for behavioral Nash equilibria of imperfect-recall games and prove finite convergence on sum-of-squares-concave subclasses, but do not treat the two-player maxmin. Rigorous rounding of floating-point sum-of-squares certificates is classical (Peyrl & Parrilo, 2008); our posterior bound follows the evaluation route rather than the rounding route and adds the weak-duality term for the adversary.

Teams. Team-maxmin equilibria were introduced by Von Stengel & Koller (1997); Basilico et al. (2017) bound their inefficiency relative to coordinated play, Celli & Gatti (2018) and Zhang et al. (2023) compute the coordinated variant, and Zhang & An (2020); Zhang et al. (2022) compute the uncoordinated team-maxmin by piecewise-linear and McCormick relaxations solved as mixed-integer programs.

First-order methods. CFR (Zinkevich et al., 2007), CFR$^+$ (Tammelin, 2014), and predictive RM$^+$ (Farina et al., 2023) are the solvers of practice; on imperfect-recall trees they stop at points that need not be optimal (Lanctot et al., 2012; Anagnostides et al., 2026). The certificate is the missing check.

3 Behavioral Maxmin as Polynomial Optimization

The game. A zero-sum two-player extensive-form game has a finite tree with chance nodes, decision nodes of player 1 and of player 2, and leaves z with payoff $u_1(z)$ to player 1. Player 1 has *imperfect recall*: her decision nodes are partitioned into information sets $j \in J$ with action sets A_j, and a path may visit the same information set more than once or reach it with different histories of her own actions. A behavioral strategy is $\mu \in \Delta := \prod_{j \in J} \Delta(A_j)$; we write μ_{ja} for the probability of action a at j and $n = \sum_j |A_j|$ for the number of variables. Player 2 has perfect recall and is described in sequence form (Koller et al., 1996): a plan $y \in \mathcal{Y} := \{y \geq 0 : Ey = e\}$ indexed by sequences s, with one row of E per information set of player 2 and e the unit vector of the empty sequence. The reach probability of a leaf factors as $\pi_c(z)\, \pi_1^\mu(z)\, y_{s(z)}$ with $\pi_1^\mu(z) = \prod \mu_{ja}$ over player 1's decisions on the path, a monomial whose degree can exceed one in a variable when the path revisits j. Hence

$$u(\mu, y) = \sum_s y_s A_s(\mu),$$
$$A_s(\mu) := \sum_{z:\, s(z)=s} \pi_c(z) u_1(z) \pi_1^\mu(z), \tag{1}$$

where A_s is a polynomial in μ of degree at most $D := \max_z \deg \pi_1^\mu(z)$, the largest number of player-1 decisions on a path.

Two values. The *behavioral maxmin* and *minmax* values are

$$v^\star = \max_{\mu \in \Delta} \min_{y \in \mathcal{Y}} u(\mu, y), \quad \bar{v} = \min_{y \in \mathcal{Y}} \max_{\mu \in \Delta} u(\mu, y). \tag{2}$$

Under perfect recall Kuhn's theorem makes Δ equivalent to a sequence-form polytope, u is bilinear, and $v^\star = \bar{v}$ is the value of a linear program. Under imperfect recall Δ is a product of simplices on which $u(\cdot, y)$ is multilinear or of higher degree, the inner minimum is a piecewise-polynomial function, and $v^\star \le \bar{v}$ with strict inequality in general (Koller & Megiddo, 1992). The maxmin is what a deployed abstraction or a team without communication obtains. Gimbert et al. (2020) show that (2) is a polynomial optimization problem over Δ and that deciding it is $\exists\mathbb{R}$-hard; hardness persists when the adversary has perfect recall (Tewolde et al., 2024).

The two applications. An imperfect-recall *abstraction* of a poker game merges information sets of player 1 while the opponent keeps the full game. A strategy of the abstraction is a legal strategy of the full game, so its deployment value is $\min_y u(\hat{\mu}, y)$ against the unabstracted opponent, and no lifting is needed. An adversarial *team* of perfect-recall members who cannot communicate is, in behavioral strategies, a single player whose information sets are the union of the members' and who forgets what the other members know (Von Stengel & Koller, 1997; Celli & Gatti, 2018); the team's value is the team-maxmin value $v^\star$, and the value with pooled information (the linear-program value of the merged player, TMECom) is at least as large. The difference is the value of coordination, which Basilico et al. (2017) bound in the worst case and Zhang & An (2020) show can be arbitrarily large.

Solver output. Let $\hat{\mu} \in \Delta$ be the output of any solver, with deployment value $\hat{v} := \min_{y \in \mathcal{Y}} u(\hat{\mu}, y)$ computed exactly by backward induction over the adversary's tree. We use two solvers: CFR$^+$ self-play with the adversary at perfect recall (Zinkevich et al., 2007; Tammelin, 2014), which on an imperfect-recall tree stops at a point that need not be optimal (Lanctot et al., 2012), and predictive regret matching$^+$ on μ with the exact best-response subgradient (Farina et al., 2023). Neither bounds $v^\star - \hat{v}$; the certificate below does.

Reduced coordinates. Eliminating the last action of each information set, $\mu_{j,a_{\text{last}}} = 1 - \sum_{a \neq a_{\text{last}}} \mu_{ja}$, maps Δ to a full-dimensional product of simplices $\mathcal{D} \subset \mathbb{R}^m$, $m = n - |J|$, cut out by the n affine functions $g_v(x) \ge 0$, one per original action v: $g_v = x_v$ for a non-last action and $g_v = 1 - \sum_{a \neq a_{\text{last}}} x_{ja}$ for the last action of j. Every polynomial in μ becomes a polynomial in x of the same degree, and $g_v(x) = \mu_v$. All certificates are built in these coordinates; the simplex equalities disappear and the moment basis shrinks from $\binom{n+d}{d}$ to $\binom{m+d}{d}$.

4 The Response-Map Certificate

Fix a candidate $\hat{\mu}$ with value $\hat{v}$, written in reduced coordinates as $\hat{x}$. Strong duality of the inner linear program turns (2) into $v^\star = \max\{\lambda_\emptyset : x \in \mathcal{D},\ E^\top \lambda \le A(x)\}$, a polynomial program in (x, λ). A Positivstellensatz certificate of $\hat{v} + \varepsilon - \lambda_\emptyset \ge 0$ on this set carries, in general, multipliers that depend on both x and λ, and λ enters the moment basis. We restrict the ansatz: the multiplier of the s-th dual-feasibility constraint is a polynomial $\tau_s(x)$ in x alone, nonnegative on $\mathcal{D}$, and the remainder $q(x)$ is free of λ,

$$\hat{v} + \varepsilon - \lambda_\emptyset \;=\; q(x) + \sum_s \tau_s(x)\big(A_s(x) - (E^\top \lambda)_s\big).$$

Matching the coefficients of λ on both sides forces $E\tau(x) = e$ identically, and what remains is $\hat{v} + \varepsilon - \sum_s \tau_s(x) A_s(x) = q(x)$. This restriction is the definition below. It is not lossless, and its soundness is proved directly (Theorem 1); what it buys is the block structure: λ never enters a moment basis.

Definition 1 (Response-map certificate)**.** Let $P_\beta := \prod_v g_v^{\beta_v}$ denote the Handelman products of the affine functions defining $\mathcal{D}$. For a level d and a response degree r with $2d \ge D + r$, the program $(C_{d,r})$ minimizes ε over

$$\hat{v} + \varepsilon - \sum_s \tau_s(x)\, A_s(x) = \sigma_0(x) + \sum_{v \in V} g_v(x)\, \sigma_v(x) + \sum_{|\alpha| \le 2d} t_\alpha P_\alpha(x), \tag{3}$$

$$\sum_s E_{is}\, \tau_s(x) = e_i \quad \text{for every row } i, \tag{4}$$

$$\tau_s(x) = \sum_{|\beta| \le r} \theta_{s\beta}\, P_\beta(x), \qquad \theta \ge 0, \tag{5}$$

where (3)–(4) are polynomial identities, σ_0 is a sum of squares of degree $\le 2d$, each σ_v a sum of squares of degree $\le 2d-2$, $t \ge 0$, and $V \subseteq \{1,\dots,n\}$ is the set of kept localizing blocks (all actions unless pruned, Section 5).

The right-hand side of (3) is nonnegative on $\mathcal{D}$ because every g_v and every P_α is, and τ is a polynomial function $\mathcal{D} \to \mathcal{Y}$: its entries are nonnegative by (5) and satisfy the flow constraints by (4). At every x it names an adversary plan that holds the player below $\hat{v} + \varepsilon$.

Theorem 1 (Soundness)**.** *Every feasible point of $(C_{d,r})$ satisfies $v^\star \le \hat{v} + \varepsilon$; hence $\hat{\mu}$ is ε-optimal for the behavioral maxmin.*

Proof. For $x \in \mathcal{D}$, $\tau(x) \in \mathcal{Y}$, so $\min_{y\in\mathcal{Y}} u(x,y) \le u(x,\tau(x)) = \sum_s \tau_s(x) A_s(x) \le \hat{v} + \varepsilon$ by (3). Maximize over x. □

The adversary enters through θ only, a nonnegative vector of $|S|\binom{n+r}{r}$ scalars tied by the linear identities (4). The semidefinite blocks are σ_0 and the σ_v, exactly those of a problem in the player's variables alone over $\mathcal{D}$. This is the structural fact that carries the method to 1093 adversary sequences: adding adversary sequences adds linear variables, not moment rows.

The response ladder. Write $\varepsilon_{d,r}$ for the optimal value of $(C_{d,r})$.

Proposition 1 (Ladder)**.** *(i) $\varepsilon_{d,r}$ is nonincreasing in d and in r. (ii) At $r = 0$ the map is a constant plan $\bar{y} \in \mathcal{Y}$ and every feasible point proves $\max_x u(x,\bar{y}) \le \hat{v} + \varepsilon$, so $\varepsilon_{d,0} \ge \bar{v} - \hat{v}$ and $\lim_{d\to\infty} \varepsilon_{d,0} = \bar{v} - \hat{v}$. (iii) $\lim_{r\to\infty} \lim_{d\to\infty} \varepsilon_{d,r} = v^\star - \hat{v}$.*

Part (ii) says that the bottom of the ladder certifies the minmax value; the gap $\bar{v} - v^\star$ is the behavioral duality gap, and the ladder descends it as r grows. The proof of (iii) (Appendix A) takes a continuous near-optimal response $x \mapsto y(x)$, approximates it by its Bernstein polynomial of degree r, which is a convex combination of plans and hence a map into $\mathcal{Y}$ of the form (5), and applies Handelman's theorem (Handelman, 1988) to the resulting strictly positive polynomial. On the instances of Section 6 the ladder is short: $r \le 2$ closes every gap that the tested levels close, and on a hand-made family with analytic values the certified value is $\bar{v}$ at $r = 0$ and $v^\star$ at $r \ge 1$ to 2×10^{-7} at the highest tested level. Abstraction can lose value without opening a behavioral duality gap: on the perfect-recall Kuhn abstractions $r = 0$ already closes.

Posterior verification. A conic solver returns (3)–(5) up to tolerances: the identities hold up to a coefficient residual, the Gram matrices may have slightly negative eigenvalues, and t, θ may dip below zero. The following bound repairs the output by evaluation alone.

Theorem 2 (Posterior bound)**.** *Let a solver return ε, symmetric G_0 (Gram matrix of σ_0 in the monomial basis b_d of size $\mathcal{B}$), H_v (Gram matrices of σ_v in b_{d-1}, size $\mathcal{L}$), and vectors t, θ. Let t^+, θ^+ be their positive parts, with $|a|_- := \max(-a, 0)$, $\pi_t = \|t^-\|_1$, $\pi_\lambda = |\lambda_{\min}(G_0)|_- \mathcal{B} + \sum_v |\lambda_{\min}(H_v)|_- \mathcal{L}$, let ρ_1 be the ℓ_1 norm of the coefficient residual of* (3) *evaluated with t^+, θ^+, and let $\delta_i(x)$ be the residual polynomial of row i of* (4) *with θ^+, with coefficient norm $\|\delta_i\|_1$. Write E in the standard sequence form: one row $y_\emptyset = 1$ for the empty sequence and, for each information set I of the adversary with parent sequence $p(I)$, one row $y_{p(I)} - \sum_{a\in A_I} y_{Ia} = 0$, so e is the unit vector of the first row. Let $U_i \ge |\lambda_i(x)|$ bound the i-th entry of the backward-induction dual of $\min_{y\in\mathcal{Y}} A(x)^\top y$ uniformly in $x \in \mathcal{D}$; one may take $U_I = M \sum_{h\in I} \pi_c(h)$ for the row of information set I and $U_\emptyset = M$, where $M = \max_z |u_1(z)|$ and $\pi_c(h)$ is the chance reach of history h. Then*

$$v^\star \;\le\; \hat{v} + \varepsilon_{\mathrm{rig}}, \qquad \varepsilon_{\mathrm{rig}} := \varepsilon + \rho_1 + \pi_\lambda + \pi_t + \sum_i U_i \|\delta_i\|_1.$$

The last term comes from the response map: the repaired map $\tilde{\tau}$ is nonnegative but violates the flow constraints by $\delta(x)$, and weak duality of the inner linear program charges the violation at the price of the adversary's subtree values. The row-wise bound matters in practice: with the global bound $\max_z |u_1(z)|$ on every row the term dominated ε by two orders of magnitude on Leduc, with the subtree masses it is below 10^{-8}. The proof is in Appendix A; the quantities are evaluated in double precision and, for the 483 certificates checked in Section 6, once more in exact rational arithmetic with interval Cholesky bounds on the eigenvalues; on the Leduc and team instances the check runs on a re-solve, not on the stored certificate (Appendix B).

Scale. When $\binom{m+d}{d}$ exceeds a few hundred we restrict every block of (3) to a clique of the chordal extension of the correlative sparsity graph of the leaf monomials (Waki et al., 2006), keeping τ dense; the resulting cone is a sub-cone of $(C_{d,r})$, so Theorems 1 and 2 apply unchanged and the certified value can only increase. Keeping τ dense costs only linear variables: at $r = 1$ on the 126-variable Leduc abstraction with 1093 adversary sequences, θ has 138,811 nonnegative coefficients and (4) is 34,237 coefficient identities over the 468 adversary information sets and the root row.

5 Pruning and the Loss Decomposition

Pruning the adversary. Restricting θ to a set S of sequences containing the support of some plan removes $|S^c|\binom{n+r}{r}$ variables; every row of (4) is kept, and a row whose children all leave S forces its parent sequence to zero. Let $\mathcal{Y}_S := \mathcal{Y} \cap \{y_s = 0,\ s \notin S\}$, the plans of an adversary forbidden to leave S, and $v^\star_S := \max_x \min_{y \in \mathcal{Y}_S} u(x, y) \geq v^\star$.

Theorem 3 (Adversary-side pruning)**.** *The program $(C_{d,r})$ with $\theta_{s\cdot} = 0$ for $s \notin S$ is the response-map certificate of the game in which the adversary is restricted to $\mathcal{Y}_S$. Consequently every feasible point proves $v^\star_S \leq \hat{v} + \varepsilon$, and $\lim_{r,d\to\infty} \varepsilon^S_{d,r} = v^\star_S - \hat{v}$. The pruning loss $v^\star_S - v^\star$ is a property of S alone and does not vanish with the level.*

Proof. $\mathcal{Y}_S$ is a compact polytope of the same flow form, so Theorem 1 and Proposition 1(iii) apply verbatim to the restricted game. □

The natural choice $S = \operatorname{supp}$ of the best response at $\hat{\mu}$ is therefore only safe when no better strategy of the player makes the adversary leave that support. On Kuhn abstractions with the support of CFR's average adversary strategy the loss is zero on all but the instances where CFR itself stalled, and there the certificate reports exactly $v^\star_S - \hat{v}$ (Section 6); on Leduc the support of a single run is too small and pruning the adversary is lossy. We therefore keep all sequences by default; they are cheap.

Pruning the player. Complementary slackness also suggests dropping the localizing blocks $g_v\sigma_v$ of the actions in the support of $\hat{\mu}$, so that V consists of the off-support actions. The pruned cone is a sub-cone of $(C_{d,r})$, so Theorems 1 and 2 apply unchanged and the certified value can only increase. The increase is not always small: on our team and Leduc instances it is 0.03 to 0.90 against unpruned certificates of 5.6×10^{-5} to 0.11 (Table 7), so we keep every block by default.

The loss decomposition. Let v_{ref} be a reference value that the player would obtain with more information: the full-game value for an abstraction, the pooled-information value for a team. The deployment loss of $\hat{\mu}$ is

$$v_{\mathrm{ref}} - \hat{v} \;=\; \underbrace{(v_{\mathrm{ref}} - v^\star)}_{\text{abstraction / coordination}} \;+\; \underbrace{(v^\star - \hat{v})}_{\text{solver}}, \tag{6}$$

and the certificate bounds the solver part by $\varepsilon_{\mathrm{rig}}$ and hence the abstraction part from below by $v_{\mathrm{ref}} - \hat{v} - \varepsilon_{\mathrm{rig}}$. The first part can only be recovered by a different abstraction, or for a team by more information; the second by a better point of the same abstraction. The solver part concerns an $\exists\mathbb{R}$-hard quantity.

6 Experiments

Setup. *Kuhn abstractions*: the first player's three cards are merged into buckets ($\{JQ\}|K$, $J|\{QK\}$, $\{JK\}|Q$, $\{JQK\}$) at her first decision, her second, or both, and a second family applies a different bucket map at each decision; all have $D = 2$, 2–6 information sets and 13 adversary sequences. Bucketing at the second decision alone destroys recall (4 instances), and cross-bucketing in nine of 12; the rest refine information over time and keep perfect recall (Appendix B). *Leduc abstractions*: two or three card buckets, the board card forgotten (kept in one), and the first betting round remembered fully, only as raised-or-not, or not at all; 28 to 126 variables against 1093 adversary sequences, $D = 4$. *Teams*: three-player Kuhn with each two-player team as one imperfect-recall player (64 variables, 33 sequences, $D = 3$ or 4); the reference value is the pooled-information linear program. *Hand-made family*: four small games with analytic values (Appendix B). Candidates come from

CFR$^+$ self-play (5000 iterations, deterministic), from predictive RM$^+$ with exact best responses (three seeds), and from DCFR (Brown & Sandholm, 2019) at 5000 and 50000 iterations on the imperfect-recall abstractions. Certificates are solved with CLARABEL (Goulart & Chen, 2026) in one process with its default thread count; SCIP (Bestuzheva et al., 2023) on the bilinear reach formulation, started from the same candidate, supplies ground truth and, at fixed budgets, an anytime dual bound in place of the specialized branch-and-bound methods (Bosansky et al., 2017; Zhang & An, 2020), which we do not run. A candidate is *certified* when $\varepsilon_{\text{rig}} \le 10^{-4}\max(1, |\hat{v}|)$. Times are wall-clock for that process and include model construction; the 3600 s limit (7200 for the level-3 certificate of seats $1, 2$) applies to the conic solve.

Kuhn: soundness, ladder, and the decomposition. Over 25 instances, two solvers and three seeds (150 runs; CFR$^+$ is deterministic, so 100 distinct candidates), response degrees $r \le 3$ and three pruning variants, 1800 certificates were solved, excluding the baselines of Tables 5 and 6. None violates the SCIP bound. Without pruning the certified gap and SCIP's gap agree to 10^{-4} on all 150 runs, and 100 of them are certified optimal, which SCIP also proves optimal. A certificate takes at most 0.2, 11 and 23 s at $r = 0, 1, 2$. The ladder appears on the abstraction that forgets the card at the second decision: CFR's point has gap 0.222, degree 0 certifies $0.250 = \bar{v} - \hat{v}$. A joint (x, λ) Putinar hierarchy at level 2 has floating-point bounds consistent with SCIP on all 25 candidates, but its cone carries 55–105 monomials against 6–28, takes 17 s at the median against 0.07 s, and is never tighter: it matches the response map's $r = 1$ bound to 10^{-6} on 9 of the 25 instances and is looser on the other 16, by up to a factor of 316 (Table 5). Figure 1 plots the deployment loss against the certified solver gap. The points split into two populations. On the diagonal the whole loss is the solver's: RM$^+$ stops at gaps of 10^{-4} to 10^{-3} that the certificate reports to 10^{-4}, and on $\{JK\}|Q$-second CFR stalls 0.167 below a value the abstraction can reach, which RM$^+$ finds to within 10^{-3}. The gap survives a tenfold budget and a different solver: on the seven abstractions whose certified gap exceeds 10^{-3}, DCFR at 5000 and 50000 iterations and CFR$^+$ at 50000 certify the same gap to 1.7×10^{-5} (Table 6). On the floor the whole loss is the abstraction's; on $\{JK\}|Q$-both CFR's point is optimal to 3×10^{-6}.

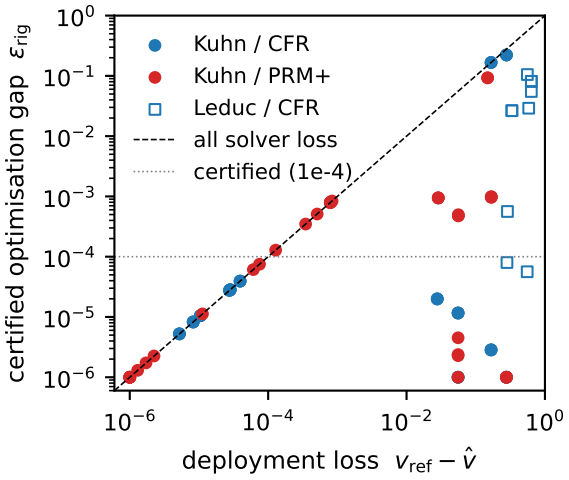

Figure 1: Deployment loss $v_{\text{ref}} - \hat{v}$ against the certified solver gap ε_{rig} for every candidate. Points on the diagonal lose only to the solver; points on the floor lose only to the abstraction. Exploitability places both populations on the same vertical line.

Hand-made family. *Coordination* ($\bar{v} = 1/2$, $v^\star = 1/4$): player 1 chooses $a \in \{0, 1\}$; the adversary, seeing nothing, chooses b; player 1, having forgotten a, chooses c; the payoff is 1 iff $a = b = c$. With $p = \Pr[a = 1]$ and $q = \Pr[c = 1]$ the adversary's two plans give $f_1 = pq$ and $f_0 = (1-p)(1-q)$, so $v^\star = \max\min(pq, (1-p)(1-q)) = \frac{1}{4}$ while the uniform plan gives $\max_\mu \frac{1}{2}(pq+(1-p)(1-q)) = \frac{1}{2}$. The certified value is 0.500 at $r = 0$ and 0.250 at $r = 1$ for both solvers, one of which stopped at 0.162.

Three-way coordination adds a third forgetful decision ($v^\star = 1/8$, $\bar{v} = 1/2$), and here an affine response map already attains $v^\star$. With p, q, t the three probabilities of choosing 1, the adversary's plans give $P = pqt$ and $Q = (1-p)(1-q)(1-t)$. Set $s = (p+q+t)/3$ and let the map play plan 1 with probability $1 - s$ and plan 0 with probability s, which is of the form (5) with $r = 1$. By the AM–GM inequality $P \le s^3$ and $Q \le (1-s)^3$, so the payoff is at most $(1-s)s^3 + s(1-s)^3 = a(1-2a)$ with $a = s(1-s) \in [0, \frac{1}{4}]$, which is at most $\frac{1}{8}$. Degree one therefore suffices in the limit, and the excess $0.173 - 0.125$ that the level-2 cone returns at $r = 1$ is a limit of the level, not of the degree: raising the level at fixed $r = 1$ removes it. The crossed sweep separates $r = 0$ from $r \ge 1$ but not

Table 1: Three-way coordination, CFR$^+$ candidate ($\hat{v} = 0.125$), floating-point posterior bound $\hat{v} + \varepsilon_{\text{rig}}$ (not exact-verified) against SOS level d and response degree r; the minimum admissible level is $d = \lceil (3 + r)/2 \rceil$. Times in seconds, single-threaded.

r	$d = 2$	$d = 3$	$d = 4$	$d = 5$
0	0.5000 (1.2)	0.5000 (3.7)	0.5000 (17)	–
1	0.1729 (0.0)	0.1250 (3.9)	0.1250 (23)	–
2	–	0.1250 (5.0)	0.1250 (25)	0.1250 (143)

Table 2: Hand-made family: posterior bound $\hat{v} + \varepsilon_{\text{rig}}$ at the minimum admissible level for each response degree r, from the CFR and RM$^+$ points; the two solvers agree to 10^{-7}.

game	$\hat{v}$ (CFR)	$\hat{v}$ (RM$^+$)	$r = 0$	$r = 1$	$r = 2$	$v^\star$
coordination	0.250	0.162	0.500	0.250	0.250	$1/4$
three-way coordination	0.125	0.031	0.500	0.173	0.125	$1/8$
adversarial driver	1.000	1.333	1.333	1.333	1.333	$4/3$
signal and forget	0.000	0.500	0.500	0.500	0.500	$1/2$

$r = 1$ from $r = 2$: at $r = 1$ the returned bound is 0.1729 at level 2 and 0.1250 at levels 3 and 4 (to 2×10^{-7}), at $r = 2$ it is 0.1250 at levels 3, 4, 5, and at $r = 0$ it stays 0.5000 through level 4 (Table 1).

Adversarial absent-minded driver: the adversary first selects one of two payoff maps $(0, 4, 1)$ or $(1, 0, 4)$ for (exit now, exit later, never exit); the driver then faces the same information set twice. CFR stalls at 1.000 while the certificate reports the optimum 1.333 from either solver's point. *Signal and forget*: chance deals one of two cards c equiprobably; player 1 chooses x knowing c; the adversary chooses y knowing x and not c; player 1 chooses z at a single information set, knowing neither c nor y. All actions are binary and the payoff is $(3{\cdot}\mathbf{1}[x = c]{-}1){+}(1{-}3{\cdot}\mathbf{1}[y = x]){+}\mathbf{1}[y = 1]\,(2{\cdot}\mathbf{1}[z = c]{-}1)$. The plan $y = x$ holds every strategy to $\frac{1}{2}(0 + 1) = \frac{1}{2}$, and $x = c$, $z = 1$ guarantees it against every y, so $v^\star = \bar{v} = \frac{1}{2}$: the game has no behavioral duality gap, which the certificate reports. Table 2 collects the posterior bounds of all four games from both solvers' points. At a fixed candidate-independent configuration the optimal certificate value depends on the game, not on the candidate; the returned bounds may differ.

Leduc: certificates where the global solver stops. Table 3 lists the Leduc abstractions. On the two largest that remember the first round, 84 and 126 variables, CFR's point is certified optimal within 4.92×10^{-5} and 8.01×10^{-5} (exact-verified re-solves, Table 8); the stored solves took 6 and 26 minutes. SCIP's dual gap from the same points is 0.0028 and 0.011 after an hour. On the 63-variable abstraction SCIP proves CFR's point optimal and the certificate reaches 5.6×10^{-4}, and a single-threaded re-solve exact-verified to 5.633×10^{-4} (Appendix B). On the 28-variable abstraction the certificate is loose, 0.055 against a true gap of 0.046; the dense cone at level 2 (basis 153) is looser, 0.19, and level 3 dense is out of reach. The decomposition is one-sided: the abstraction part accounts for at least 0.55 and 0.28 of the deployment losses; these abstractions are bad, not under-solved.

Teams. For each of the three two-player teams of three-player Kuhn the certificate gives $\hat{v} \le v^\star \le \hat{v} + \varepsilon_{\text{rig}}$, and the pooled-information value $v_{\text{pool}} \ge v^\star$ is a linear program, so the value of coordination $v_{\text{pool}} - v^\star$ lies in the interval of Table 4. For the team of seats 1 and 2 the interval is $[0.045, 0.053]$ (level 3, 1.7 hours) and for seats 0 and 2 it is $[0.024, 0.050]$ at level 2 (5 minutes); SCIP confirms both team-maxmin values. The third team is the hard case: SCIP does not close it in an hour and the level-2 certificate leaves $\varepsilon_{\text{rig}} = 0.087$ (single-threaded re-solve exact-verified to 8.693×10^{-2}), a coordination interval of $[0, 0.042]$ after clipping. RM$^+$ with exact best responses lands within 2×10^{-4} and 5×10^{-4} of the team-maxmin on the first two teams, so there the interval is almost entirely the value of coordination.

Table 3: Leduc abstractions, CFR$^+$ candidate, no pruning, sparse cone at level 3. SCIP is started from the CFR point and its gap is its dual bound minus $\hat{v}$ after 900 s, an upper bound on the true gap when SCIP has not closed; an asterisk marks instances SCIP had not closed at that budget. Times in minutes, at the degree r reported. $\varepsilon_{\mathrm{rig}}$ is the stored floating-point bound; exact-verified values in Table 8.

abstraction	vars	SCIP gap	$\varepsilon_{\mathrm{rig}}$	r	time
b2, forget board+round	28	0.046	0.055	1	13.4
b2, forget board, raise flag	42	0.022	0.029	0	8.3
b3, forget board+round	42	0.026	0.026	1	32.4
b3, forget board, raise flag	63	0.000	5.6×10^{-4}	0	15.8
b2, forget board	84	0.015*	5.6×10^{-5}	0	5.9
b2, forget round	98	0.411*	0.105	0	62.5
b3, forget board	126	0.025*	8.0×10^{-5}	1	25.6

Table 4: Three-player Kuhn teams, CFR$^+$ candidate, no pruning. $v^\star$ from SCIP started at the CFR point, after 3600 s, reported as a dual bound with an asterisk where its gap exceeds 10^{-6}; the interval for the value of coordination is $[\, v_{\mathrm{pool}} - \hat{v} - \varepsilon_{\mathrm{rig}},\ v_{\mathrm{pool}} - \hat{v}\,]$, clipped at 0, endpoints rounded outward. $\varepsilon_{\mathrm{rig}}$ is the stored floating-point bound (exact-verified re-solves in Appendix B); the level-3 certificate of seats $1, 2$ is exact-verified to 6.74×10^{-3}.

team	D	$\hat{v}$	$\hat{v} + \varepsilon_{\mathrm{rig}}$	$v^\star$ (SCIP)	v_{pool}	interval
seats 1,2	3	0.0313	0.0380	0.0365	0.0833	$[0.045, 0.053]$
seats 0,2	3	0.0143	0.0396	0.0237	0.0641	$[0.024, 0.050]$
seats 0,1	4	-0.0417	0.0450	$\le -0.002^*$	0.0000	$[0, 0.042]$

Pruning. Adversary-side pruning to the support of CFR's average strategy is lossless on Kuhn except where CFR stalled; on those certificates at level two and above the reported gap equals the restricted maxmin $v^\star_S - \hat{v}$ computed by SCIP to 1.1×10^{-7} at level two and 5.5×10^{-7} at level three, as Theorem 3 predicts in the limit. Player-side pruning is not free and we do not use it: on cubic and quartic instances it raises the returned bound to 0.094 vs. 0.014 on the team of seats $1, 2$ at level 2 and 0.12 vs. 5.6×10^{-5} on the 84-variable Leduc abstraction (Table 7).

Exact verification. Re-solving the stored certificates and re-evaluating their bounds in exact rational arithmetic raises each bound by $\le 3 \times 10^{-6}$ and changes no label on 483 certificates (450 Kuhn of degree at most two, 24 hand-made, 9 of the 10 in Tables 3 and 4; the tenth is verified by the controlled re-solve of Appendix B), and the Kuhn re-solves reproduce the stored $\varepsilon_{\mathrm{rig}}$ to 10^{-7} absolute.

7 Conclusion

A player who forgets, facing an adversary who does not, obtains the behavioral maxmin, and first-order solvers stop at points they cannot certify. One semidefinite program certifies how far such a point is: the adversary enters as a polynomial response map whose degree climbs from the minmax to the maxmin value, the floating-point solution is repaired into a rigorous bound by evaluation, and the bound splits the point's deployment loss into what the abstraction or the lack of coordination costs and what the solver left on the table. On Kuhn the split is sharp to 10^{-4} and exposes two abstractions with the same exploitability and opposite diagnoses; on Leduc abstractions of up to 126 variables the certificate bounds CFR's gap by 10^{-4} in minutes where a global solver does not finish in an hour; on three-player teams it bounds the value of coordination. It composes with any solver: the response map it returns is the polynomial adversary, and the posterior step repairs only the bound. It needs only the candidate and the game: the reference value enters the accounting, not the bound.

Reproducibility Statement

Appendix A contains the proofs, and Appendix B the construction of every instance, the certificate parameters behind each table, and the solver and hardware settings. The anonymized supplementary package ships the experiment drivers, the stored results behind the reported numbers, and a compiled certificate kernel.

AI Use Disclosure

AI tools were used for language polishing and code debugging.

References

Ioannis Anagnostides, Emanuel Tewolde, Brian Zhang, Ioannis Panageas, Vincent Conitzer, and Tuomas Sandholm. Convergence of regret matching in potential games and constrained optimization. In C. Vondrick, B. Hariharan, C. Raffel, L. Pinto, D. Yang, and A. Faust (eds.), *International Conference on Learning Representations*, volume 2026, pp. 108681–108707, 2026. URL https://proceedings.iclr.cc/paper_files/paper/2026/file/b0b932e3b588a8e25c395433dbff769f-Paper-Conference.pdf.

Robert J Aumann, Sergiu Hart, and Motty Perry. The absent-minded driver. *Games and Economic Behavior*, 20(1):102–116, 1997.

Nicola Basilico, Andrea Celli, Giuseppe De Nittis, and Nicola Gatti. Team-maxmin equilibrium: Efficiency bounds and algorithms. In *AAAI*, pp. 356–362. AAAI Press, 2017.

Ratip Emin Berker, Emanuel Tewolde, Ioannis Anagnostides, Tuomas Sandholm, and Vincent Conitzer. The value of recall in extensive-form games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 13631–13640, 2025.

Dimitris Bertsimas, Dan Andrei Iancu, and Pablo A Parrilo. A hierarchy of near-optimal policies for multistage adaptive optimization. *IEEE Transactions on Automatic Control*, 56(12):2809–2824, 2011.

Ksenia Bestuzheva, Mathieu Besançon, Wei-Kun Chen, Antonia Chmiela, Tim Donkiewicz, Jasper van Doornmalen, Leon Eifler, Oliver Gaul, Gerald Gamrath, Ambros Gleixner, et al. Enabling research through the SCIP optimization suite 8.0. *ACM Transactions on Mathematical Software*, 49(2):1–21, 2023.

Branislav Bosansky, Jiri Cermak, Karel Horak, and Michal Pechoucek. Computing maxmin strategies in extensive-form zero-sum games with imperfect recall. In *International Conference on Agents and Artificial Intelligence*, volume 2, pp. 63–74. SCITEPRESS, 2017.

Noam Brown and Tuomas Sandholm. Solving imperfect-information games via discounted regret minimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pp. 1829–1836, 2019.

Luca Carminati, Brian Hu Zhang, Federico Cacciamani, Junkang Li, Gabriele Farina, Nicola Gatti, and Tuomas Sandholm. Efficient representations for team and imperfect-recall equilibrium computation. *Artificial Intelligence*, 356:104529, 2026.

Andrea Celli and Nicola Gatti. Computational results for extensive-form adversarial team games. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

Jirí Cermák, Branislav Bosanský, and Michal Pechoucek. Combining incremental strategy generation and branch and bound search for computing maxmin strategies in imperfect recall games. In *AAMAS*, pp. 902–910. ACM, 2017.

Gabriele Farina, Julien Grand-Clément, Christian Kroer, Chung-Wei Lee, and Haipeng Luo. Regret Matching+: (In)Stability and Fast Convergence in Games. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL https://openreview.net/forum?id=nYgs0qZJ97.

Hugo Gimbert, Soumyajit Paul, and B. Srivathsan. A bridge between polynomial optimization and games with imperfect recall. In *AAMAS*, pp. 456–464. International Foundation for Autonomous Agents and Multiagent Systems, 2020.

Paul J Goulart and Yuwen Chen. Clarabel: An interior-point solver for conic programs with quadratic objectives. *Mathematical Programming Computation*, pp. 1–83, 2026.

David Handelman. Representing polynomials by positive linear functions on compact convex polyhedra. *Pacific Journal of Mathematics*, 132(1):35–62, 1988.

Daphne Koller and Nimrod Megiddo. The complexity of two-person zero-sum games in extensive form. *Games and economic behavior*, 4(4):528–552, 1992.

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.

Christian Kroer and Tuomas Sandholm. Extensive-form game abstraction with bounds. In *Proceedings of the fifteenth ACM conference on Economics and computation*, pp. 621–638, 2014.

Christian Kroer and Tuomas Sandholm. Imperfect-recall abstractions with bounds in games. In *Proceedings of the 2016 ACM Conference on Economics and Computation*, pp. 459–476, 2016.

Marc Lanctot, Richard G. Gibson, Neil Burch, Martin Zinkevich, and Michael Bowling. No-regret learning in extensive-form games with imperfect recall. In *Proceedings of the 29th International Conference on Machine Learning (ICML)*, pp. 65–72. icml.cc / Omnipress, 2012.

Rida Laraki and Jean B Lasserre. Semidefinite programming for min–max problems and games. *Mathematical programming*, 131(1):305–332, 2012.

Jean B Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on optimization*, 11(3):796–817, 2001.

Jiawang Nie and Xindong Tang. Nash equilibrium problems of polynomials. *Mathematics of Operations Research*, 49(2):1065–1090, 2024.

Pablo A Parrilo. Polynomial games and sum of squares optimization. In *Proceedings of the 45th IEEE Conference on Decision and Control*, pp. 2855–2860. IEEE, 2006.

Helfried Peyrl and Pablo A Parrilo. Computing sum of squares decompositions with rational coefficients. *Theoretical Computer Science*, 409(2):269–281, 2008.

Michele Piccione and Ariel Rubinstein. On the interpretation of decision problems with imperfect recall. *Games and Economic Behavior*, 20(1):3–24, 1997.

Oskari Tammelin. Solving large imperfect information games using cfr+. *arXiv preprint arXiv:1407.5042*, 2014.

Emanuel Tewolde, Caspar Oesterheld, Vincent Conitzer, and Paul W Goldberg. The computational complexity of single-player imperfect-recall games. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pp. 2878–2887, 2023.

Emanuel Tewolde, Brian Hu Zhang, Caspar Oesterheld, Manolis Zampetakis, Tuomas Sandholm, Paul Goldberg, and Vincent Conitzer. Imperfect-recall games: equilibrium concepts and their complexity. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, pp. 2994–3004, 2024.

Bernhard Von Stengel and Daphne Koller. Team-maxmin equilibria. *Games and Economic Behavior*, 21(1-2):309–321, 1997.

Hayato Waki, Sunyoung Kim, Masakazu Kojima, and Masakazu Muramatsu. Sums of squares and semidefinite program relaxations for polynomial optimization problems with structured sparsity. *SIAM Journal on Optimization*, 17(1):218–242, 2006.

Kevin Waugh, Martin Zinkevich, Michael Johanson, Morgan Kan, David Schnizlein, and Michael H Bowling. A practical use of imperfect recall. In *SARA*, 2009.

Brian Hu Zhang, Gabriele Farina, and Tuomas Sandholm. Team belief dag: generalizing the sequence form to team games for fast computation of correlated team max-min equilibria via regret minimization. In *International Conference on Machine Learning*, pp. 40996–41018. PMLR, 2023.

Youzhi Zhang and Bo An. Computing team-maxmin equilibria in zero-sum multiplayer extensive-form games. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pp. 2318–2325, 2020.

Youzhi Zhang, Bo An, and VS Subrahmanian. Correlation-based algorithm for team-maxmin equilibrium in multiplayer extensive-form games. In *IJCAI*, pp. 606–612, 2022.

Rui Zheng, Ryann Sim, and Antonios Varvitsiotis. Solving imperfect-recall games via sum-of-squares optimization. In *Forty-third International Conference on Machine Learning*, 2026.

Martin Zinkevich, Michael Johanson, Michael Bowling, and Carmelo Piccione. Regret minimization in games with incomplete information. *Advances in neural information processing systems*, 20, 2007.

A PROOFS

Throughout, $\mathcal{D} \subset [0,1]^m$ is the reduced domain, $g_1, \dots, g_n$ the affine functions with $\mathcal{D} = \{x : g_v(x) \geq 0 \ \forall v\}$, and $P_\beta = \prod_v g_v^{\beta_v}$. Two facts are used repeatedly: every g_v and every P_β takes values in $[0,1]$ on $\mathcal{D}$, and every monomial x^α does too, so a polynomial p satisfies $|p(x)| \leq \|p\|_1$ on $\mathcal{D}$, where $\|p\|_1$ is the ℓ_1 norm of its coefficient vector in the monomial basis.

A.1 PROOF OF THEOREM 2

Let $\tilde{\tau}_s(x) = \sum_\beta \theta^+_{s\beta} P_\beta(x)$ be the map built from the clipped coefficients; $\tilde{\tau}(x) \geq 0$ on $\mathcal{D}$. Let $r(x)$ be the residual polynomial of (3) with t^+ and θ^+, i.e.

$$\hat{v} + \varepsilon - \sum_s \tilde{\tau}_s A_s = b_d^\top G_0 b_d + \sum_{v \in V} g_v \, b_{d-1}^\top H_v b_{d-1} + \sum_\alpha t^+_\alpha P_\alpha + r(x),$$

so that $\rho_1 = \|r\|_1$. For $x \in \mathcal{D}$ the Gram terms are bounded below by $-|\lambda_{\min}(G_0)|_- \|b_d(x)\|^2 \geq -|\lambda_{\min}(G_0)|_- \mathcal{B}$ and, using $g_v \in [0,1]$, by $-|\lambda_{\min}(H_v)|_- \mathcal{L}$; the t^+ terms are nonnegative; and $|r(x)| \leq \rho_1$. Hence

$$\sum_s \tilde{\tau}_s(x) A_s(x) \ \leq \ \hat{v} + \varepsilon + \rho_1 + \pi_\lambda \qquad (x \in \mathcal{D}). \tag{7}$$

(Clipping t moves its negative part into r, so π_t is already counted in ρ_1; it is kept in ε_{rig} only to match the quantity the code reports, and it is zero on every certificate below.)

The backward-induction dual. Fix x and write $a = A(x)$. The inner problem $\min\{a^\top y : Ey = e, \ y \geq 0\}$ has the dual $\max\{\lambda_\emptyset : E^\top \lambda \leq a\}$; with the row convention of Theorem 2 the column of a sequence Ia has entry -1 in row I and $+1$ in every row J with $p(J) = Ia$, so the dual constraints read

$$-\lambda_I + \sum_{J:\, p(J) = Ia} \lambda_J \ \leq \ a_{Ia} \quad \text{for every sequence } Ia, \qquad \lambda_\emptyset + \sum_{J:\, p(J) = \emptyset} \lambda_J \ \leq \ a_\emptyset.$$

Since the adversary has perfect recall, her information sets form a forest in which every J has a unique parent sequence, and the continuation values

$$W_I(x) \ := \ \min_{a \in A_I} \Big(a_{Ia} + \sum_{J:\, p(J) = Ia} W_J(x) \Big)$$

are well defined from the leaves up. Set $\lambda_I := -W_I(x)$ and $\lambda_\emptyset := a_\emptyset + \sum_{J:\, p(J)=\emptyset} W_J(x)$. Every constraint holds: $-\lambda_I + \sum_J \lambda_J = W_I - \sum_J W_J \leq a_{Ia}$ by the definition of W_I as a minimum, and the root constraint holds with equality. The plan y that selects a minimizing action at every I attains $a^\top y = \lambda_\emptyset$, so by weak duality $\lambda(x)$ is an optimal dual and $\lambda_\emptyset = \min_{y \in \mathcal{Y}} a^\top y$. (Only optimality is used below; the counterfactual definition of W_I matters because λ_I must be feasible also at information sets that an optimal plan never reaches.)

The bounds U_i. $\lambda_\emptyset$ is an expected payoff under (x, y), so $|\lambda_\emptyset| \leq M$. For I, unrolling the recursion shows that $W_I(x)$ is the minimum over pure continuations y' below I of $\sum_{h \in I} \sum_{z \succ h, \ y' \text{ reaches } z} \pi_c(z) u_1(z) \pi_1^x(z)$. For one $h \in I$ and one y', the leaves below h that y' reaches are distinguished only by chance and by player-1 moves below h, so their conditional reach probabilities sum to at most one: $\sum_{z \succ h, \ y' \text{ reaches } z} \pi_c(h \to z) \, \pi_1^x(h \to z) \leq 1$, and the contribution of h is at most $M \pi_c(h) \pi_1^x(h) \leq M \pi_c(h)$ in absolute value. Summing over $h \in I$ and using $|\min_{y'}(\cdot)| \leq \max_{y'} |\cdot|$ gives $|\lambda_I| = |W_I(x)| \leq M \sum_{h \in I} \pi_c(h) = U_I$.

Weak duality against the repaired map. Since $\tilde{\tau}(x) \geq 0$ and $E\tilde{\tau}(x) = e + \delta(x)$,

$$a^\top \tilde{\tau}(x) \ \geq \ \lambda(x)^\top E \tilde{\tau}(x) \ = \ \lambda_\emptyset + \lambda(x)^\top \delta(x) \ \geq \ \min_{y \in \mathcal{Y}} a^\top y - \sum_i U_i \, \|\delta_i\|_1,$$

using $E^\top \lambda \leq a$, $\tilde{\tau} \geq 0$, $|\delta_i(x)| \leq \|\delta_i\|_1$ on $\mathcal{D}$ and $|\lambda_i(x)| \leq U_i$. Together with (7), $\min_{y \in \mathcal{Y}} u(x, y) \leq \hat{v} + \varepsilon + \rho_1 + \pi_\lambda + \sum_i U_i \|\delta_i\|_1$ for every $x \in \mathcal{D}$; maximize over x. $\square$

A.2 Proof of Proposition 1

(i) The feasible set of $(C_{d,r})$ grows with d (larger Gram matrices and more products) and with r (more products in (5), the identities (4) being imposed on the larger space), so the optimal value is nonincreasing.

(ii) At $r = 0$, $\tau_s = \theta_{s0}$ is constant and (4) reads $E\theta = e$, $\theta \geq 0$, i.e. $\theta = \bar{y} \in \mathcal{Y}$. A feasible point then certifies $u(x, \bar{y}) \leq \hat{v} + \varepsilon$ on $\mathcal{D}$, so $\hat{v} + \varepsilon \geq \max_x u(x, \bar{y}) \geq \bar{v}$. For the limit, let $\bar{y}^\star$ attain the minmax and note that $f_d := \bar{v} + \eta - u(\cdot, \bar{y}^\star)$ is strictly positive on $\mathcal{D}$ for every $\eta > 0$; Handelman's theorem (Handelman, 1988) gives $f_d = \sum_\alpha t_\alpha P_\alpha$ with $t \geq 0$ for some finite degree, which is a feasible point with $\varepsilon = \bar{v} + \eta - \hat{v}$ once $2d$ exceeds that degree. Hence $\limsup_d \varepsilon_{d,0} \leq \bar{v} - \hat{v} + \eta$ for every η.

(iii) Throughout, r denotes the total degree of the Handelman products in (5); since $\varepsilon_{d,r}$ is nonincreasing in r and in d by (i), it suffices to exhibit feasible points along a subsequence. Fix $\eta > 0$ and let $M = \max_z |u_1(z)|$.

A continuous η-optimal response. For $\gamma > 0$ let $y_\gamma(x) := \arg\min_{y \in \mathcal{Y}} \big(A(x)^\top y + \gamma \|y\|_2^2\big)$, unique by strict convexity. For two objectives a, a' the optimality conditions $(a + 2\gamma y)^\top (y' - y) \geq 0$ and $(a' + 2\gamma y')^\top (y - y') \geq 0$ add up to $2\gamma \|y - y'\|_2^2 \leq (a - a')^\top (y' - y)$, so $\|y_\gamma(x) - y_\gamma(x')\|_2 \leq \|A(x) - A(x')\|_2/(2\gamma)$ and y_γ is continuous on $\mathcal{D}$ because A is polynomial. Since $\mathcal{Y} \subset [0,1]^S$, $\|y\|_2^2 \leq |S|$ on $\mathcal{Y}$, and for a minimizer $y^\star$ of $A(x)^\top y$, $A(x)^\top y_\gamma(x) \leq A(x)^\top y^\star + \gamma\big(\|y^\star\|_2^2 - \|y_\gamma(x)\|_2^2\big) \leq \min_{y \in \mathcal{Y}} A(x)^\top y + \gamma |S|$. With $\gamma = \eta/|S|$ the map $y_\eta := y_\gamma$ is continuous and η-optimal at every x.

Tensor Bernstein approximation. Write $x = (x^{(1)}, \ldots, x^{(K)})$, $K = |J|$, where $x^{(j)}$ ranges over the reduced simplex of information set j with affine functions g_v, $v \in A_j$, $\sum_{v \in A_j} g_v = 1$. For a multidegree $k = (k_1, \ldots, k_K)$ the Bernstein basis of degree k_j on factor j is $b_\beta^{(j)}(x) = \binom{k_j}{\beta} \prod_{v \in A_j} g_v(x)^{\beta_v}$ over $|\beta| = k_j$: nonnegative multiples of Handelman products of degree k_j that sum to one by the multinomial theorem. The product basis $b_\beta = \prod_j b_{\beta^{(j)}}^{(j)}$ therefore consists of nonnegative multiples of Handelman products of total degree $r := \sum_j k_j$ and sums to one. The tensor Bernstein operator $B_k[f](x) = \sum_\beta b_\beta(x)\, f(x_\beta)$, with x_β the grid point $(\beta^{(j)}/k_j)_j$, is a positive linear operator that reproduces 1 and every coordinate and satisfies $B_k[x_v^2] = x_v^2 + x_v(1 - x_v)/k_j$ for $v \in A_j$; by Korovkin's theorem $B_k[f] \to f$ uniformly on $\mathcal{D}$ for every continuous f as $\min_j k_j \to \infty$. Apply B_k to each entry of y_η and set $\tau := B_k[y_\eta]$. Each τ_s is a nonnegative combination of Handelman products of degree $\leq r$, so τ is of the form (5) with $\theta \geq 0$, and $E\tau = \sum_\beta b_\beta\, E\, y_\eta(x_\beta) = e \sum_\beta b_\beta = e$ holds as a polynomial identity, which is (4).

Feasibility. Let $C := \sum_z \pi_c(z)|u_1(z)|$, the sum over all leaves. Writing Z_s for the leaves whose adversary sequence is s, $A_s(x) = \sum_{z \in Z_s} \pi_c(z) u_1(z) \pi_1^x(z)$ with $0 \leq \pi_1^x(z) \leq 1$ on $\mathcal{D}$, so $\|A(x)\|_1 = \sum_s |A_s(x)| \leq \sum_z \pi_c(z)|u_1(z)| = C$ on $\mathcal{D}$; the constant is C and not M because the sum over s visits every leaf once, without the adversary's own reach probabilities (two terminal adversary actions paying 1 each give $M = 1$ and $\|A\|_1 = 2$). Hence $|u(x, \tau(x)) - u(x, y_\eta(x))| = |A(x)^\top(\tau(x) - y_\eta(x))| \leq C\, \|\tau(x) - y_\eta(x)\|_\infty \leq \eta$ once $\min_j k_j$ is large. Then $f(x) := v^\star + 3\eta - u(x, \tau(x)) \geq v^\star + 2\eta - u(x, y_\eta(x)) \geq v^\star + \eta - \min_y u(x, y) \geq \eta > 0$ on $\mathcal{D}$. By Handelman's theorem (Handelman, 1988), $\mathcal{D}$ being a full-dimensional polytope cut out by the g_v, $f = \sum_\alpha t_\alpha P_\alpha$ with $t \geq 0$ and some finite degree $2d_0$. Hence for every $d \geq \max(d_0, \lceil (D + r)/2 \rceil)$ the point $(\varepsilon, \sigma_0, \sigma_v, t, \theta) = (v^\star + 3\eta - \hat{v}, 0, 0, t, \theta)$ is feasible for $(C_{d,r})$, and with Theorem 1, $v^\star - \hat{v} \leq \varepsilon_{d,r} \leq v^\star - \hat{v} + 3\eta$. As η was arbitrary and $\varepsilon_{d,r}$ is monotone in both indices, $\lim_{r \to \infty} \lim_{d \to \infty} \varepsilon_{d,r} = v^\star - \hat{v}$. □

A.3 Proof of Theorem 3

Setting $\theta_{s\cdot} = 0$ for $s \notin S$ removes those sequences from (5) and from the sums in (3)–(4). The polytope $\mathcal{Y}_S = \{y \geq 0 : Ey = e, y_s = 0\ (s \notin S)\}$ is the sequence-form polytope of the adversary whose tree is pruned to S (rows whose information set has no child in S become $y_{\text{parent}} = 0$, which is consistent because S contains the support of a plan), and the restricted program is exactly $(C_{d,r})$ for the game with adversary set $\mathcal{Y}_S$ and the same A_s, $s \in S$. Theorem 1 gives $v_S^\star \leq \hat{v} + \varepsilon$. For the

limit, the proof of Proposition 1(iii) used only that $\mathcal{Y}$ is a compact polytope of plans on which the inner problem is a linear program and that convex combinations stay inside; both hold for $\mathcal{Y}_S$. □

B Instances, Protocol, and Additional Results

Hand-made family. The definitions of the four games, the AM–GM argument that degree one attains $v^\star$ in the limit, the crossed sweep separating $r = 0$ from $r \geq 1$ but not $r = 1$ from $r = 2$, and the posterior bounds themselves are in Section 6 (Tables 1 and 2).

Kuhn abstractions. A bucket map is applied to the first player's private card at her first decision, her second decision (after a check–bet), or both; a second family uses a different map at each decision. The player has perfect recall exactly when the second-decision partition refines the first, since then her own earlier action tells her everything she has forgotten: bucketing at the second decision alone merges cards she could tell apart (4 instances), and cross-bucketing does so whenever the first-decision map is not already the coarsest (9 of the 12). The remaining 11 abstractions and the full game keep perfect recall while losing information, so the study separates the two effects. All have $D = 2$; each certificate uses the smallest admissible level $d = \lceil (2 + r)/2 \rceil$, i.e. $d = 1, 2, 2, 3$ for $r = 0, 1, 2, 3$, and the identity (3) is imposed on every monomial of degree $\leq 2d$.

Joint (x, λ) hierarchy. The baseline of Table 5 keeps the adversary's sequence-form dual $\lambda \in \mathbb{R}^L$, one coordinate per row of E (the $L - 1$ information sets and the root row), as a variable next to x. Let $U \in \mathbb{R}^L$ be the row bounds of Theorem 2, $U_\emptyset = M$ and $U_I = M \sum_{h \in I} \pi_c(h)$, and let

$$K := \big\{(x, \lambda) : g_v(x) \geq 0 \ (v \in V),\ U_i - \lambda_i \geq 0,\\ U_i + \lambda_i \geq 0 \ (i \leq L),\ A_s(x) - (E^\top \lambda)_s \geq 0 \ (s \in S)\big\}.$$

For every $x \in \mathcal{D}$ the inner linear program $\min_{y \in \mathcal{Y}} A(x)^\top y$ has a dual optimum $\lambda^\star(x)$ with $E^\top \lambda^\star(x) \leq A(x)$, $\lambda^\star_\emptyset(x) = \min_{y \in \mathcal{Y}} A(x)^\top y$ and $|\lambda^\star_i(x)| \leq U_i$ (the counterfactual continuation values of the proof of Theorem 2), so $(x, \lambda^\star(x)) \in K$. The level-$d$ joint program is

$$\hat{v} + \varepsilon - \lambda_\emptyset = \sigma_0 + \sum_v g_v\, \sigma_v + \sum_i \big[(U_i - \lambda_i)\, \sigma_i^+ + (U_i + \lambda_i)\, \sigma_i^-\big] + \sum_s \big(A_s(x) - (E^\top \lambda)_s\big)\, \sigma_s, \quad (8)$$

an identity in $\mathbb{R}[x, \lambda]_{\leq 2d}$, where every σ is a sum of squares in (x, λ): σ_0 of degree $2d$, σ_v, $\sigma_i^\pm$ of degree $2(d - 1)$, and σ_s of degree $2(d - \lceil D/2 \rceil)$. No Handelman products and no sparsity are used; every multiplier is a dense Gram matrix in the monomial basis of (x, λ), which is what the column "joint $\mathcal{B}$" counts. The program is sound at every level: an exact identity (8) gives $\hat{v} + \varepsilon \geq \lambda_\emptyset$ on K, and at $(x, \lambda^\star(x))$ this reads $\hat{v} + \varepsilon \geq \min_y u(x, y)$ for every $x \in \mathcal{D}$, hence $\varepsilon \geq v^\star - \hat{v}$. We make no convergence claim for it (that would need an Archimedean quadratic module, e.g. the redundant ball constraint) and use it only as the sound baseline that the response map replaces.

The floating-point solution is repaired exactly as in Theorem 2, with the box on λ replacing the unit box on x. For a monomial $m_\alpha(x, \lambda) = x^{\alpha_x} \lambda^{\alpha_\lambda}$ let $W_\alpha := \prod_i U_i^{(\alpha_\lambda)_i} \geq \sup_K |m_\alpha|$ (since $x \in [0, 1]^m$ on $\mathcal{D}$ and $|\lambda_i| \leq U_i$ on K), and for each block b of (8) with multiplier polynomial q_b let $Q_b \geq \sup_K |q_b|$: $Q_0 = 1$, $Q_v = 1$, $Q_{i\pm} = 2U_i$, and $Q_s = \sum_{z \in Z_s} \pi_c(z)|u_1(z)| + \sum_i |E_{is}|\, U_i$.

Proposition 2 (Posterior bound for the joint program)**.** *Let V be the full action set (as in every joint run of Table 5, so that $K \subseteq \mathcal{D} \times \prod_i [-U_i, U_i]$), let a solver return ε and Gram matrices G_b for the blocks of* (8)*, let r_α be the coefficient residual of* (8) *at monomial m_α, and set $\rho_1 := \sum_\alpha |r_\alpha|\, W_\alpha$ and $\pi_\lambda := \sum_b \big(-\lambda_{\min}(G_b)\big)^+ Q_b \sum_{\alpha \in \mathrm{basis}(b)} W_\alpha^2$. Then $v^\star - \hat{v} \leq \varepsilon + \rho_1 + \pi_\lambda =: \varepsilon_{\mathrm{rig}}$.*

Proof. On K, where the bounds W_α and Q_b hold, replacing G_b by $G_b + (-\lambda_{\min}(G_b))^+ I$ makes every block positive semidefinite and adds $(-\lambda_{\min}(G_b))^+\, q_b(x, \lambda)\, \|b_b(x, \lambda)\|_2^2 \leq (-\lambda_{\min}(G_b))^+ Q_b \sum_{\alpha \in \mathrm{basis}(b)} W_\alpha^2$ to the right-hand side at every point of K (each $q_b \geq 0$ on K). The residual polynomial $\sum_\alpha r_\alpha m_\alpha$ is bounded by ρ_1 on K. Hence $\hat{v} + \varepsilon - \lambda_\emptyset \geq -\rho_1 - \pi_\lambda$ on K, and evaluating at $(x, \lambda^\star(x))$ gives $\min_y u(x, y) \leq \hat{v} + \varepsilon + \rho_1 + \pi_\lambda$ for every $x \in \mathcal{D}$. □

The response-map certificate needs no W_α: its monomials and Handelman products are bounded by one on $\mathcal{D}$, which is why Theorem 2 charges π_λ at $\mathcal{B}$ and $\mathcal{L}$ per unit of eigenvalue. On K the

Table 5: Joint (x, λ) Putinar hierarchy at level 2 against the response map at $r = 1$ on the thirteen imperfect-recall Kuhn abstractions, CFR$^+$ candidate, no pruning. An instance name gives the bucket map at the player's first decision and, after $\to$, at her second; "-second" is the map applied at the second decision only. $\mathcal{B}$ is the number of basis monomials of the cone; the joint column is the floating-point evaluation of Proposition 2, not exact-verified. The joint bound is consistent with SCIP on all of them and is never tighter; on the full 25-instance set it matches the response map to 10^{-6} on 9 and is looser by up to a factor 316 on the other 16.

instance	$\varepsilon_{\mathrm{rig}}$ (joint)	$\varepsilon_{\mathrm{rig}}$ $(r = 1)$	joint $\mathcal{B}$	map $\mathcal{B}$
$\{J\}\|\{QK\} \to \{JQK\}$	0.222	0.222	66	10
$\{JQK\}$-second	0.222	0.222	78	15
$\{JQ\}\|\{K\} \to \{JQK\}$	0.200	0.114	66	10
$\{JK\}\|\{Q\}$-second	0.167	0.167	91	21
$\{JQ\}\|\{K\} \to \{JK\}\|\{Q\}$	0.148	0.148	78	15
$\{J\}\|\{QK\} \to \{JK\}\|\{Q\}$	0.111	0.111	78	15
$\{JK\}\|\{Q\} \to \{JQK\}$	0.056	0.056	66	10
$\{JQ\}\|\{K\} \to \{J\}\|\{QK\}$	7.34×10^{-3}	2.33×10^{-5}	78	15
$\{J\}\|\{QK\} \to \{JQ\}\|\{K\}$	3.24×10^{-3}	2.46×10^{-5}	78	15
$\{J\}\|\{QK\}$-second	1.32×10^{-3}	2.77×10^{-5}	91	21
$\{JK\}\|\{Q\} \to \{JQ\}\|\{K\}$	2.52×10^{-5}	2.38×10^{-5}	78	15
$\{JK\}\|\{Q\} \to \{J\}\|\{QK\}$	2.69×10^{-5}	1.48×10^{-5}	78	15
$\{JQ\}\|\{K\}$-second	2.20×10^{-5}	2.00×10^{-5}	91	21

Table 6: DCFR against CFR$^+$ on the thirteen imperfect-recall Kuhn abstractions, response degree $r = 2$, no pruning. DCFR discounts the stored regrets by $(t-1)^\alpha/((t-1)^\alpha+1)$ and $(t-1)^\beta/((t-1)^\beta+1)$ for positive and negative entries and accumulates the average strategy with weight t^γ, at $\alpha = 1.5$, $\beta = 0$, $\gamma = 2$. No certificate violates the SCIP bound; on the seven rows whose gap exceeds 10^{-3} the three columns agree with each other to 1.7×10^{-5}, and DCFR at 50000 iterations is the tighter candidate on all but one of the rest.

instance	$\hat{v}$ (DCFR 50k)	$\varepsilon_{\mathrm{rig}}$ (5k)	$\varepsilon_{\mathrm{rig}}$ (50k)	$\varepsilon_{\mathrm{rig}}$ (CFR$^+$)
$\{JQK\}$-second	-3.33×10^{-1}	0.222	0.222	0.222
$\{J\}\|\{QK\} \to \{JQK\}$	-3.33×10^{-1}	0.222	0.222	0.222
$\{JK\}\|\{Q\}$-second	-2.22×10^{-1}	0.167	0.167	0.167
$\{JQ\}\|\{K\} \to \{JK\}\|\{Q\}$	-2.22×10^{-1}	0.148	0.148	0.148
$\{J\}\|\{QK\} \to \{JK\}\|\{Q\}$	-2.22×10^{-1}	0.111	0.111	0.111
$\{JQ\}\|\{K\} \to \{JQK\}$	-3.33×10^{-1}	0.111	0.111	0.111
$\{JK\}\|\{Q\} \to \{JQK\}$	-3.33×10^{-1}	0.056	0.056	0.056
$\{J\}\|\{QK\}$-second	-5.56×10^{-2}	2.70×10^{-5}	6.19×10^{-6}	8.60×10^{-6}
$\{JQ\}\|\{K\} \to \{J\}\|\{QK\}$	-7.41×10^{-2}	2.32×10^{-5}	4.28×10^{-6}	4.82×10^{-6}
$\{JQ\}\|\{K\}$-second	-8.33×10^{-2}	1.91×10^{-5}	2.31×10^{-6}	2.97×10^{-6}
$\{JK\}\|\{Q\} \to \{J\}\|\{QK\}$	-1.11×10^{-1}	8.16×10^{-6}	1.66×10^{-6}	2.24×10^{-6}
$\{JK\}\|\{Q\} \to \{JQ\}\|\{K\}$	-8.33×10^{-2}	1.12×10^{-5}	1.31×10^{-6}	1.86×10^{-6}
$\{J\}\|\{QK\} \to \{JQ\}\|\{K\}$	-1.11×10^{-1}	1.33×10^{-6}	1.72×10^{-6}	1.40×10^{-6}

same accounting must carry U_i up to M per power of λ_i, and Table 5 reports $\varepsilon_{\mathrm{rig}}$ computed with Proposition 2. The comparison is between the two formulations as run: the response-map cone carries Handelman products and the joint cone does not, so it measures the two hierarchies, not the elimination of λ in isolation.

Leduc abstractions. The six cards form three ranks; "b2" merges the two low ranks, "b3" keeps all three. "Forget board" removes the public card from the second-round information sets; "forget round" removes the first-round betting history, "raise only" keeps whether it contained a raise. Abstractions that keep the board have 98, 147, 434 and 651 variables (two or three buckets, first round forgotten or kept); only the 98-variable one, b2 with the round forgotten, is within the range of this study. $D = 4$; the correlative sparsity graph of the leaf monomials, chordally completed, has cliques of at most

Table 7: Player-side pruning on the Leduc and team instances, CFR$^+$ candidate, $\varepsilon_{\mathrm{rig}}$ without and with pruning at the level shown. All entries are stored floating-point posterior bounds; the two rows marked † are $r = 1$ solves that reached the 3600 s limit (`user_limit` on Leduc, `optimal_inaccurate` after 3600 s on the team) in both variants, not converged and not exact-verified.

instance	r	level	unpruned	pruned
b2, forget board+round	0	3	0.081	0.202
b2, forget board	0	3	5.6×10^{-5}	0.121
b2, forget round	0	3	0.105	0.554
b2, forget round†	1	3	4.82	2.24
b3, forget board+round	0	3	0.027	0.074
b3, forget board	0	3	1.5×10^{-4}	0.027
seats 1,2	0	2	0.014	0.094
seats 1,2	1	2	0.013	0.094
seats 0,2	0	2	0.025	0.069
seats 0,2	1	2	0.025	0.069
seats 0,1	0	2	0.087	0.989
seats 0,1†	1	3	3.61	2.10

seven variables, and all certificates use level 3 on the cliques with Handelman products of degree two. Response degree $r = 1$ is dense over all n affine functions.

Player-side pruning beyond level one. Table 7 lists every Leduc and team certificate solved both with and without player-side pruning from CFR's point; the increases of 0.03 to 0.90 in the returned bound are the ten rows where the pruned bound is larger, and the team example quoted there is seats 1 and 2 at level 2, $r = 0$. On Kuhn the six level-two exceptions are the $r = 1$ certificates of the cross-bucketed abstraction $\{JQ\}|K$ then $\{JQK\}$, both solvers and all three seeds, each larger by 1.8×10^{-4}; at level three 129 of 150 are unchanged.

Teams. In three-player Kuhn the team of seats i, j is the imperfect-recall player with 32 information sets of two actions; cliques have 11 variables. Level 2 takes 1–6 minutes per certificate; level 3 took 1.7 hours on the team of seats 1 and 2 and did not finish in 2 hours on the others (exploratory cluster jobs at a 7200 s solve limit, in no table), so those are reported at level 2.

Protocol. Every certificate is solved once with CLARABEL (interior point, default tolerances, time limit 3600 s, except the level-3 certificate of seats $1, 2$ at 7200 s); any solver status, including `optimal_inaccurate` and `user_limit`, is accepted, because the posterior bound of Theorem 2 does not depend on it. A posterior bound printed without an exact-verified companion is the double-precision evaluation of Theorem 2, not an independently checked inequality. SCIP runs the bilinear reach formulation with a relative optimality gap tolerance of 10^{-6} and the stated time limits; its dual bound is used as ground truth whenever it closes and as an anytime bound otherwise. Exact verification recomputes the coefficient residuals in rational arithmetic with game coefficients as the rationals OpenSpiel specifies, and bounds each Gram matrix's smallest eigenvalue by an interval Cholesky factorization of the matrix shifted by a candidate bound (40 decimal digits).

Verification manifest. Three numbers are attached to every re-verified certificate: the *stored* bound $\varepsilon_{\mathrm{rig}}$ printed in the tables, the *re-solved* bound of a fresh CLARABEL run from the same candidate in the same configuration (no pruning, the same level and degree), and the *exact* bound obtained by re-evaluating the posterior terms of that fresh solve in rational arithmetic with interval Cholesky bounds on the Gram eigenvalues. The candidate is rounded to a rational with denominators at most 10^9 and rebalanced so that every information set sums to one; $\hat{v}$ is recomputed exactly for that rational candidate, which moves it by at most 4×10^{-16} across the original cohort of 483 certificates, and the exact bound refers to it. The 3×10^{-6} of Section 6 is the largest difference between the exact and the re-solved bound; it is 9.8×10^{-11} on the 450 Kuhn certificates, 8.0×10^{-11} on the 24 hand-made ones, and is attained on the b2 raise-flag Leduc row. The re-solved and stored bounds agree to

Table 8: Verification manifest for the ten certificates of Tables 3 and 4: stored bound, re-solved bound, exact re-evaluation of the re-solve, and whether the certified/not-certified label ($\varepsilon_{\mathrm{rig}} \le 10^{-4}\max(1,|\hat{v}|)$) agrees between the stored and the exact bound. The seats $1,2$ row is re-solved under the pinning of the controlled re-solves below; the other nine are the plain re-solve.

instance	r	level	stored	re-solved	exact	label
b2, forget board+round	1	3	5.51×10^{-2}	6.27×10^{-2}	6.27×10^{-2}	same
b2, forget board, raise flag	0	3	2.91×10^{-2}	2.91×10^{-2}	2.91×10^{-2}	same
b3, forget board+round	1	3	2.64×10^{-2}	2.64×10^{-2}	2.65×10^{-2}	same
b3, forget board, raise flag	0	3	5.63×10^{-4}	1.54×10^{-1}	1.54×10^{-1}	same
b2, forget board	0	3	5.63×10^{-5}	4.91×10^{-5}	4.92×10^{-5}	same
b2, forget round	0	3	1.05×10^{-1}	1.06×10^{-1}	1.06×10^{-1}	same
b3, forget board	1	3	8.00×10^{-5}	8.00×10^{-5}	8.01×10^{-5}	same
seats 1,2	0	3	6.73×10^{-3}	6.73×10^{-3}	6.74×10^{-3}	same
seats 0,2	0	2	2.53×10^{-2}	2.53×10^{-2}	2.53×10^{-2}	same
seats 0,1	0	2	8.67×10^{-2}	7.26×10^{1}	7.26×10^{1}	same

8.2×10^{-8} on Kuhn (median 9×10^{-16}), where CLARABEL returns `optimal`. On the Leduc and team certificates it returns `optimal_inaccurate`, and a fresh run can stop at a different feasible point: seven of the nine re-solves land within 14% of the stored bound and two land far above it (Table 8; all nine return to the stored bound under the controlled re-solves below). Both are posterior bounds of Theorem 2 in floating point; the tables print the stored one (the re-solve is tighter on three rows, by at most 13%). The remaining certificate, seats 1 and 2 at level 3, is re-verified by the controlled re-solve below, outside the 483: its exact expansion exceeds the 64 GB of the dev machine and needs a node. The joint-hierarchy bounds of Table 5 are evaluated in floating point through Proposition 2 and are not exact-verified. The 25 joint certificates and the 9 cells of Table 1 are outside the 483.

Controlled re-solves. The rows of Table 8 were re-solved with CLARABEL pinned to one thread, at a four-hour budget (eleven hours for the level-3 seats $1,2$ certificate, whose expansion into exact polynomials exceeds the 64 GB of the dev machine); all return to within a factor of 1.3 of the stored bound. At three times the stored solve time the level-2 seats $1,2$ bound reproduces to 2×10^{-9}, while the 63-variable Leduc certificate returns exact-evaluated bounds of 1.28×10^{-3} (default tolerances and tighter tolerances), 2.42×10^{-3} (equilibration and static regularization), and 5.633×10^{-4} (a four-hour budget) against a stored 5.625×10^{-4}; seats $0,1$ returns exact-evaluated bounds of 2.06 (default tolerances) and 8.693×10^{-2} (a four-hour budget) against a stored 8.668×10^{-2}; the level-3 seats $1,2$ certificate returns an exact-evaluated bound of 6.736×10^{-3} (an eleven-hour budget) against a stored 6.730×10^{-3}. These re-solves, not the stored certificates, are the exact-verified objects: they prove $v^\star-\hat{v}\le5.633\times10^{-4}$, $\le8.693\times10^{-2}$ and $\le6.736\times10^{-3}$; the tables print the stored floating-point bounds.

Scale. The sparse cone at level 3 takes 6–63 minutes on Leduc abstractions of 28–126 variables at $r=0$ and 11–67 at $r=1$, the upper end being the 98-variable one, whose cliques are larger. Adversary sequences enter only through θ: at $r=0$ they add $|S|$ scalars, at $r=1$ $(n+1)|S|$, with no effect on the semidefinite blocks.