SHARP ACCURACY THRESHOLDS FOR ENTROPY-REGULARIZED BINARY GAMES

Anonymous authors
Paper under double-blind review

ABSTRACT

Entropy regularization yields smooth, unique best responses, yet equilibrium computation can remain hard. For binary games, we quantify this limitation through an explicit threshold for the largest unilateral improvement in regularized payoff. On every fixed bounded interval of interaction strength, tolerances above the threshold admit polynomial-time algorithms, whereas finding solutions below it is PPAD-complete, with inverse-polynomial margins. Both bounds follow from a scalar identity linking contraction algorithms and hardness reductions. Moreover, hardness persists with fixed regularization, bounded payoffs, and sparse interactions, showing that local smoothness alone does not ensure tractability.

1 INTRODUCTION

An equilibrium requires every player's mixed strategy to be a best response to the others. Adding a positive multiple of entropy to each player's payoff makes this local maximization strictly concave, with a unique, smooth best response. This supports efficient learning in structured games (Cen et al., 2021; 2023; Sokota et al., 2023), but the responses remain coupled. We ask how accurately their joint consistency can be computed when a uniform uniqueness guarantee is lost, even at fixed regularization strength.

Model and Accuracy. For two actions, write $x_i \in [-1, 1]$ for player i's probability of the first action minus that of the second. The two-action softmax best response then becomes the binary logit map

$$F(x) = \tanh(Ax + b), \qquad x \in [-1, 1]^n, \qquad \kappa = \rho(|A|). \tag{1}$$

Here $\tanh$ acts coordinatewise, $|A|$ takes entrywise absolute values, and ρ denotes spectral radius. The scalar $(Ax + b)_i$ is the *field*, or input to player i's response. In a polymatrix game, each payoff is a sum of pairwise interactions; its payoff differences give A, b, with $A_{ii} = 0$. Our algorithms also allow arbitrary rational A, b.

The gain κ bounds feedback after a suitable rescaling of coordinates (Section 3). At $\kappa \leq 1$, the fixed point is unique and iteration $x^{t+1} = F(x^t)$ converges globally. Every $\kappa > 1$ admits an instance with three fixed points (Appendix B). This is a worst-case boundary: a particular game with $\kappa > 1$ can still have a unique, easily computed equilibrium.

We seek *any* approximate equilibrium. The residual $\mathrm{Res}_F(x) = \|x - F(x)\|_\infty$ measures the largest coordinate inconsistency. To measure inconsistency between action distributions, define

$$\pi_z = \left(\frac{1+z}{2}, \frac{1-z}{2}\right), \quad d(x, y) = D_{\mathrm{KL}}(\pi_x \| \pi_y), \quad \mathcal{L}_\ell(x; F) = \max_i \ell(x_i, F_i(x)). \tag{2}$$

Here $D_{\mathrm{KL}}(p\|q) = \sum_j p_j \log(p_j/q_j)$ is Kullback–Leibler divergence, with natural logarithms and $0 \log 0 = 0$. The loss ℓ denotes absolute error or d; Section 2 also considers a broader family of losses.

Let $\lambda > 0$ be the inverse temperature, also called rationality: the entropy coefficient is $1/\lambda$. If $\mathrm{Gap}_i(x)$ is player i's largest unilateral improvement in payoff plus this entropy term, then

$$\mathrm{Gap}_i(x) = \lambda^{-1} d(x_i, F_i(x)). \tag{3}$$

This standard entropy-duality identity is proved in Appendix F (Blondel et al., 2020). Thus a KL tolerance g certifies that no player can improve its regularized objective by more than g/λ.

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.

1.1 Contributions and Main Results

The frontier is governed by two scalar functions. For $a > 1$, let $s_a > 0$ solve $s_a = \tanh(as_a)$ and define

$$\mathfrak{B}(a) = \max_{z \geq 0} \{\tanh(az) - z\}, \qquad \mathfrak{G}(a) = \max_{0 \leq z \leq s_a} d(z, \tanh(az)). \tag{4}$$

Set both to zero for $a \leq 1$. These *barriers* govern the accuracy bounds below. They measure the same scalar mismatch on $[0, s_a]$ using different losses. A larger tolerance g permits a less accurate answer. The margins δ and Δ below measure distance above and below a barrier; a fixed cap K_0 does not grow with the input size.

All input numbers are rational and encoded in binary; write B_{in} for their total bit length. Our algorithms run in $\text{poly}(B_{\text{in}}, 1/\delta)$ time, hence in input-polynomial time when $\delta \geq B_{\text{in}}^{-c}$ for a fixed $c > 0$. This is an *inverse-polynomial margin*. The reduction below measures its margin Δ against the source size N instead. An answer includes rational bounds certifying loss at most g, avoiding exact comparisons with irrational response values (Appendix G). We first give the algorithms.

Theorem 1.1 (Polynomial-Time Upper Bounds). *Given rational A, b and $g, \delta > 0$, put $\kappa = \rho(|A|)$:*

(a) if $g \geq \mathfrak{B}(\kappa) + \delta$, an algorithm finds a rational point with $\text{Res}_F(x) \leq g$;

(b) for each fixed $K_0 \geq 1$, if $\kappa \leq K_0$ and $g \geq \mathfrak{G}(\kappa) + \delta$, it finds a rational point with $\mathcal{L}_d(x; F) \leq g$, and hence regularized gap at most g/λ.

Both algorithms are deterministic and run in time polynomial in the input length and $1/\delta$.

PPAD is a class of search problems with guaranteed solutions, including approximate Brouwer fixed points (Papadimitriou, 1994). A *reduction* translates a source problem into ours and any solution back. PPAD-hardness means a polynomial-time solver would solve every PPAD problem; membership as well gives *PPAD-completeness*.

Theorem 1.2 (Hardness at a Prescribed Gain). *Fix $K_0 > 1$. For rational $1 < K \leq K_0$ and $g, \Delta > 0$, the following problems are PPAD-hard when Δ is at least inverse polynomial in N, the size of the source Boolean constraint problem:*

(a) finding x with $\mathcal{L}_d(x; F) \leq g$ if $g \leq \mathfrak{G}(K) - \Delta$;

(b) finding x with $\text{Res}_F(x) \leq g$ if $g \leq \min\{\bar{g}, \mathfrak{B}(K) - \Delta\}$, for any fixed rational $0 < \bar{g} < 2/3$.

In both cases, $A_{ii} = 0$, $\rho(|A|) = K$ exactly, and $|A|$ is irreducible: directed influence paths connect every ordered pair of coordinates. Each coordinate has at most six neighbors, counting either direction. The instances are binary-action polymatrix games with payoffs in $[-1, 1]$ and fixed rationality. Bounds on $|(Ax + b)_i|$ and rationality depend only on K_0 in (a), and on $K_0, \bar{g}$ in (b). The reductions have time and output length polynomial in source size, parameter encoding lengths, and $1/\Delta$. Both certified search problems belong to PPAD for arbitrary rational A, b and positive tolerances, giving PPAD-completeness in the stated regimes.

Table 1 separates the matching frontier from the remaining gap: whenever $\mathfrak{B}(K) \geq 2/3$, residual tolerances between $2/3$ and $\mathfrak{B}(K)$ are not classified. The residual upper bound has no gain cap. At $K \leq 1$, both barriers vanish and inverse-polynomial accuracy is tractable. Equality at a positive barrier and smaller-than-inverse-polynomial margins remain unclassified. PPAD membership at finer accuracy does not assert a polynomial-time solver there. The remaining question is whether hardness relies on weakening the entropy penalty as the problem grows. The following specialization rules this out.

Corollary 1.3 (Fixed-Temperature Games). *For rational $1 < K \leq 101/100$, under the stated margin conditions, both matching frontiers hold already for binary-action polymatrix games with maximum degree six, payoffs in $[-1, 1]$, and rationality $\lambda_0 = 4$. The frontier for the maximum entropy-regularized best-response gap is $\mathfrak{G}(K)/4$; at $K = 1 + \eta$ it is $\eta^3/18\,(1 + o(1))$.*

The entropy coefficient stays at $1/4$ as size grows. Appendix E.1 proves the field bound.

Table 1: Accuracy guarantees at gain $K > 1$. Upper bounds use $g \geq H(K) + \delta$; lower bounds use $g \leq H(K) - \Delta$. H is the barrier in the corresponding row. The gaps are inverse polynomial, K_0 and $0 < \bar{g} < 2/3$ are fixed, and $\mathfrak{B}(K_{\mathrm{aff}}) = 2/3$ with $K_{\mathrm{aff}} \approx 6.27099$.

Loss	Gain range	Polynomial-time side	PPAD-complete side
Bernoulli KL	$1 < K \leq K_0$	Above $\mathfrak{G}(K)$	Below $\mathfrak{G}(K)$
Residual	$1 < K \leq K_0 < K_{\mathrm{aff}}$	Above $\mathfrak{B}(K)$	Below $\mathfrak{B}(K)$
Residual	$1 < K \leq K_0$	Above $\mathfrak{B}(K)$	Below $\mathfrak{B}(K)$, $g \leq \bar{g}$

Regularization and Solve Accuracy. For payoff differences $2(Cx+c)$, rationality λ gives $A = \lambda C$ and $b = \lambda c$. Thus Theorem 1.1 computes a requested regularized gap ϵ whenever

$$\lambda\epsilon \geq \mathfrak{G}(\lambda\rho(|C|)) + \delta, \tag{5}$$

at fixed bounded gain and positive margin δ. The matching fixed-temperature lower bound quantifies over games, not over temperatures of one fixed game.

1.2 Related Work

Convergence guarantees do not by themselves classify attainable accuracy. We position our results by distinguishing these two questions.

Logit responses and continuation methods are classical (McKelvey and Palfrey, 1995; Goeree et al., 2016; Turocy, 2005). Regularized learning guarantees use zero-sum structure or sufficient regularization (Cen et al., 2021; 2023; Sokota et al., 2023; Sun et al., 2024); Zhao et al. (2026) assume conditions ensuring uniqueness in general-sum games. Stability criteria (Melo, 2022; Wang, 2026) address a different question. In particular, Wang's projected Euclidean norm allows opposing influences to cancel; our absolute spectral gain does not dominate every instance-specific certificate.

Černý et al. (2022, Proposition 7) prove PPAD-hardness for quantal equilibria with correlated signals, using a logit coefficient that depends on the target Nash accuracy. We obtain matching bounds at fixed entropy coefficient and prescribed gain in sparse binary games. Pure-Circuit yields tight bounds for threshold games (Deligkas et al., 2024); our smooth-response construction must also control signal amplitudes and global feedback. The smooth equilibria of Daskalakis et al. (2024) restrict deviations, rather than adding entropy.

Classical Lipschitz-constrained uniform approximation recovers our residual scalar value (del Barrio et al., 2020, Section 3); Appendix A.1 gives the specialization. Our scalar result gives one optimal response for a broad class of losses, including KL, and identifies its loss with the tolerance for transmitting a binary signal.

1.3 Proof Roadmap

The matching bounds come from two uses of the same scalar quantity. The proof proceeds in three steps.

1. Scalar Threshold (Section 2). We find the minimum cost of limiting a response's slope and prove that it equals the largest error under which the response preserves a binary value encoded by its sign.

2. Algorithm Above the Threshold (Section 3). We apply the optimal approximation to each response and rescale coordinates to obtain a contraction. Its iteration and rounding errors fit within the remaining margin, certifying accuracy for the original map.

3. Hardness Below the Threshold (Section 4). Robust signs encode a PPAD-hard Boolean constraint problem. We control the combined gain and sparsity, then realize the map as a game with fixed regularization. This reduction establishes hardness for arbitrary solvers, beyond the scalar approximation class.

2 A Scalar Threshold for Approximation and Sign Preservation

We now explain why the algorithm and hardness proof have the same barrier. Both begin with a single response before returning to the coupled system.

For the algorithm, we approximate $\tanh$ by a function whose Lipschitz constant, or maximum allowed slope, is at most $1/a$. For hardness, consider a *relay*: an input x produces an output y with $\ell(y, \tanh(ax)) \le g$. We encode the two binary values by $[P, 1]$ for 1 and $[-1, -P]$ for 0, with $P > 0$. Preserving a signal means keeping the output in the same interval despite every error allowed by g. Both questions use only the following natural properties of a loss.

Definition 2.1. A continuous loss $\ell : [-1, 1] \times (-1, 1) \to [0, \infty)$ is admissible if it vanishes exactly when its arguments agree and satisfies (i) $\ell(-x, -y) = \ell(x, y)$; (ii) $\ell(x, y) > \ell(x', y)$ for $x < x' \le y$; and (iii) $\ell(x, y) \le \ell(x, y')$ for $x \le y \le y'$.

Condition (i) makes the two signs equivalent. Condition (ii) says that approaching the target strictly lowers the loss; condition (iii) says that moving the target farther away cannot lower it. The conditions include $|x - y|^p$ for every $p > 0$ and Bernoulli KL (Appendix A.1). Define

$$\mathfrak{H}_\ell(a) = \max_{0 \le z \le s_a} \ell(z, \tanh(az)), \qquad a > 1. \tag{6}$$

For $0 < P \le s_a$, let $\Gamma_\ell(a, P)$ be the largest g such that every relay with input in $[P, 1]$ has output in $[P, 1]$.

The candidate approximation keeps $\tanh$ where its slope is small and replaces its central part by a line:

$$\chi_a(t) = \operatorname{sgn}(t) \min\{|t|/a, \tanh|t|\}, \qquad \operatorname{sgn}(0) = 0. \tag{7}$$

The next theorem says that this *clipped response* is optimal and that its loss is exactly the best relay tolerance. Here $\mathrm{Lip}(f) \le 1/a$ means $|f(t) - f(t')| \le |t - t'|/a$ for all t, t'.

Theorem 2.2 (Scalar Approximation and Relay Threshold)**.** *For every admissible loss, $a > 1$, and $0 < P \le s_a$,*

$$\inf_{\substack{f:\mathbb{R}\to[-1,1] \\ \mathrm{Lip}(f)\le 1/a}} \sup_{t \in \mathbb{R}} \ell(f(t), \tanh t) = \mathfrak{H}_\ell(a), \tag{8}$$

$$\Gamma_\ell(a, P) = \ell(P, \tanh(aP)), \tag{9}$$

$$\max_{0 < P \le s_a} \Gamma_\ell(a, P) = \mathfrak{H}_\ell(a). \tag{10}$$

The function χ_a in (7) *attains the infimum for every such loss. Sign preservation holds at equality $g = \Gamma_\ell(a, P)$.*

Thus the approximation cost and the sign-preserving tolerance are equal. Both are controlled by the gap between $\tanh(az)$ and z, as the proof shows.

Proof. Put $R_a(z) = \tanh(az) - z$. Strict concavity on $z > 0$ gives $R_a(z) \ge 0$ exactly for $0 \le z \le s_a$ and $R_a'(s_a) < 0$. Consequently $\chi_a(t) = t/a$ for $|t| \le as_a$ and equals $\tanh t$ outside this interval, where its derivative $1/\cosh^2 t \le 1/a$. It is therefore $1/a$-Lipschitz, and its worst-case loss is $\mathfrak{H}_\ell(a)$.

For the lower bound, fix $0 \le t \le as_a$. Lipschitz continuity gives $f(t) - f(-t) \le 2t/a$, so at least one of $f(t)$ and $-f(-t)$ is at most $t/a \le \tanh t$. Conditions (i)–(ii) therefore give

$$\max_{\epsilon \in \{-1,1\}} \ell(f(\epsilon t), \tanh(\epsilon t)) \ge \ell(t/a, \tanh t).$$

Maximizing over t proves the lower bound.

For the relay, let $x \ge P$ and $y < P$. The same strict decrease, followed by target monotonicity, gives

$$\ell(y, \tanh(ax)) > \ell(P, \tanh(ax)) \ge \ell(P, \tanh(aP)).$$

This proves preservation up to and including the stated threshold. Any larger tolerance admits an output just below P when the input is P, by continuity. Maximizing over P proves the last equality. □

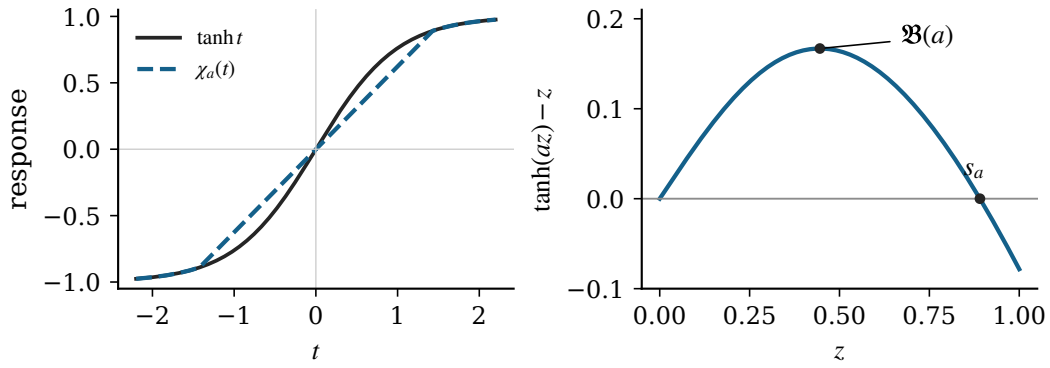


Figure 1: Response clipping and the accuracy frontiers. Left: replacing the central response by t/a limits its slope. Middle/right: classification on fixed gain intervals, with the margins in Table 1; equality at a barrier is excluded. *Residual hardness is uniform only under each fixed cap $\bar{g} < 2/3$; the grey region is unclassified. The dotted KL curve is $\mathfrak{B}(K)^2/2$; its Pinsker hardness interpretation also requires the residual cap. The plots illustrate analytic formulas.

We now specialize the common barrier to the two losses in the main results. Their continuity with respect to gain will let us reserve small error margins in both constructions.

Corollary 2.3. *For every $p > 0$, $\mathfrak{H}_{|\cdot|^p} = \mathfrak{B}^p$; for Bernoulli KL, $\mathfrak{H}_d = \mathfrak{G}$. The functions $\mathfrak{B}, \mathfrak{G}$ are increasing and 1-Lipschitz on $[0, \infty)$ after extension by zero below one. Moreover,*

$$\mathfrak{B}(a) = \sqrt{\frac{a-1}{a}} - \frac{\operatorname{arsinh}\sqrt{a-1}}{a}, \qquad a > 1. \tag{11}$$

The residual formula follows by maximizing $R_a(z)$; arsinh denotes inverse hyperbolic sine. For KL, the maximizing problem is one-dimensional and admits certified bisection; $\mathfrak{G}$ is even $1/4$-Lipschitz (Appendix A). These facts make the scalar thresholds usable with finite precision.

Why KL Needs a Direct Reduction. Pinsker's inequality gives $\mathcal{L}_d(x; F) \geq \mathrm{Res}_F(x)^2/2$: a KL solution at tolerance g is a residual solution at tolerance $\sqrt{2g}$. Thus residual hardness transfers only below $\mathfrak{B}(a)^2/2$ in KL loss, within the residual cap and with the corresponding separation. At the residual maximizer $z_a = \operatorname{arsinh}\sqrt{a-1}/a > 0$, Appendix A gives the stronger comparison

$$\mathfrak{G}(a) \geq \frac{\mathfrak{B}(a)^2}{2(1-z_a^2)} > \frac{\mathfrak{B}(a)^2}{2}. \tag{12}$$

The direct KL reduction reaches this strictly higher barrier and also covers gains beyond the matching residual regime. Its contribution is the exact finite-gain frontier; the local cubic law below is consistent with KL being locally quadratic in the residual.

2.1 Critical Scaling

The rate at which barriers vanish determines how much accuracy is difficult near gain one. Write $a = 1 + \eta$ with $\eta \downarrow 0$. The next result converts the local order of a loss into that rate; uniformity makes the expansion valid at the moving maximizer.

Theorem 2.4 (Critical Scaling)**.** *Suppose an admissible loss satisfies $\ell(x, x + r) = c|r|^p(1 + o(1))$ as $(x, r) \to (0, 0)$, $r \neq 0$, uniformly within its domain, for constants $p, c > 0$. Then, as $\eta \downarrow 0$,*

$$\mathfrak{H}_\ell(1 + \eta) = c\mathfrak{B}(1 + \eta)^p(1 + o(1)) = c(2/3)^p\eta^{3p/2}(1 + o(1)). \tag{13}$$

The maximizing interval shrinks toward zero, so maximizing the assumed local loss comparison proves the result (Appendix A). Residual and KL therefore have critical orders $\eta^{3/2}$ and η^3, with coefficients $2/3$ and $2/9$. The scalar mechanism extends to other concave saturating responses (Proposition A.5); the complexity theorems remain specific to binary logit and the two stated losses.

3 ALGORITHMS ABOVE THE THRESHOLDS

The scalar result gives the price of reducing response slopes. We construct a convergent iteration and certify its original-map accuracy by controlling coordinate scaling and finite-precision error.

Damping and Response Approximation. Damping averages x and $F(x)$ but can leave a derivative eigenvalue above one. Input scaling $t \mapsto \tanh(t/a)$ has larger worst scalar loss than clipping under the same slope bound (Appendix B.1). We therefore use clipped responses.

3.1 A RATIONAL SPECTRAL CERTIFICATE

A row-sum bound can overestimate feedback when coordinates have different scales. We therefore use the weighted max norm $\|z\|_v = \max_i |z_i|/v_i$ for $v > 0$. If $M = |A|$ satisfies $Mv \le qv$, then F is q-Lipschitz in this norm. In fact $\rho(M) = \inf_{v>0} \max_i (Mv)_i/v_i$, which explains the gain used throughout the paper.

We need rational weights that can actually be computed. For any rational $\tau > 0$, Appendix B constructs q, v in polynomial time with

$$\rho(M) < q \le \rho(M) + \tau, \qquad v \ge \mathbf{1}, \qquad Mv \le qv - \mathbf{1}. \tag{14}$$

Here $\mathbf{1}$ is the all-ones vector and inequalities are coordinatewise. This rational matrix inequality is a checkable *spectral certificate*. Testing positivity of $(qI - M)^{-1}\mathbf{1}$ and bisecting in q constructs it (Berman and Plemmons, 1994).

Given only A, b, g with $g > H(\kappa)$, $H = \mathfrak{B}$ or $\mathfrak{G}$, successively halving the uncertainty in this certificate and scalar bounds finds a safe δ comparable to $\min\{1, g - H(\kappa)\}$ (Proposition B.6). The cost depends on this separation; no exact threshold evaluation is required.

3.2 THE CONTRACTION AND ITS ERROR BUDGET

With the coordinate scaling available, choose a response slope strictly below $1/q$. The resulting contraction must leave enough of the requested margin δ for iteration and rounding errors. Set $\delta_0 = \min\{\delta, 1\}$ and $\tau = \delta_0/16$, apply (14) to $M = |A|$, and define

$$\alpha = \begin{cases} (1+\tau)\max\{1, q\}, & \text{residual,} \\ \max\{1, q\} + \tau, & \text{KL,} \end{cases} \qquad S(x) = \chi_\alpha(Ax + b), \qquad r = q/\alpha < 1. \tag{15}$$

Then $\|S(x) - S(y)\|_v \le r\|x - y\|_v$, so exact iteration contracts distances. Keeping α close to $\max\{1, q\}$ limits the added approximation loss. For residuals, $a\mathfrak{B}'(a) \le 1/2$ and hence

$$\mathfrak{B}(\alpha) - \mathfrak{B}(\kappa) \le \tau, \qquad (1 - r)^{-1} \le 1 + \tau^{-1}. \tag{16}$$

For KL, the fixed cap $\kappa \le K_0$ gives $\mathfrak{G}(\alpha) - \mathfrak{G}(\kappa) \le \tau/2$ and $(1 - r)^{-1} \le (K_0 + 2)/\tau$.

The error for the original map splits into approximation bias and iteration error: $\mathrm{Res}_F(x) \le \mathfrak{B}(\alpha) + \|x - S(x)\|_\infty$. Thus a modified-map residual, or *defect*, $\|x - S(x)\|_\infty \le \delta_0/2$ proves the residual claim. For KL we need a comparison valid even at $x_i = \pm 1$. If $z = \chi_\alpha(h)$ lies between 0 and $\tanh h$, then for every $x \in [-1, 1]$,

$$d(x, \tanh h) \le d(z, \tanh h) + (|h| + \log 2)|x - z|. \tag{17}$$

The first term is the scalar approximation loss; the second charges for incomplete iteration. With the uniform field bound $R = \max_i(|b_i| + \sum_j |A_{ij}|)$, defect $\delta_0/[8(1 + R)]$ is enough to obtain $\mathcal{L}_d(x; F) < \mathfrak{G}(\kappa) + \delta \le g$. Only $\log(1 + R)$ enters the required precision.

Algorithm 1 implements this split with target defect e. Write $v_{\max} = \max_i v_i$; $\Pi_{[-1,1]^n}$ clips each coordinate to the cube. Here $\mathrm{Eval}_\zeta(S, x)$ returns a rational vector within ζ of $S(x)$ in max norm. Since $v \ge \mathbf{1}$, its accumulated error is at most $\zeta/(1 - r)$ in the weighted norm; the selected T, ζ give defect at most e. Substituting r, e gives the iteration bounds

$$T_{\mathrm{res}} = O\left(\frac{1}{\delta_0} \log \frac{2v_{\max}}{\delta_0}\right), \quad T_{\mathrm{KL}} = O\left(\frac{K_0 + 2}{\delta_0} \log \frac{16 v_{\max}(1 + R)}{\delta_0}\right). \tag{18}$$

These count iterations, not total bit operations or optimal rates. The certificate and evaluation precision have polynomial bit length; evaluating χ_α needs only division, a minimum, and finite-precision $\tanh$ (Ko, 1991). The final dyadic rounding (power-of-two denominators) and loss certificate fit within δ. Appendix B proves the budget for Theorem 1.1.

Algorithm 1 Clipped-Response Solver

Require: Rational A, b and $g, \delta > 0$ with the promised slack for the chosen loss.
1: Set $\delta_0 = \min\{\delta, 1\}$, $\tau = \delta_0/16$; obtain q, v from (14).
2: Form S, α, r by (15); set $R = \max_i(|b_i| + \sum_j |A_{ij}|)$.
3: Set $e = \delta_0/2$ for residual or $e = \delta_0/[8(1+R)]$ for KL.
4: Choose $\zeta \le e(1-r)/(4v_{\max})$ and integer T with $r^T \le e/(4v_{\max})$.
5: Initialize $x = 0$; repeat T times: $x \leftarrow \Pi_{[-1,1]^n}(\mathrm{Eval}_\zeta(S, x))$.
6: **return** A dyadic rounding of x and a loss certificate at most g (Appendix G).

4 Hardness Below the Thresholds

We now show that a uniformly more accurate solver would solve a PPAD-hard Boolean constraint problem. Each constraint becomes a small group of logit coordinates whose accurate responses reveal a satisfying assignment. Figure 2 summarizes the encoding, error robustness, and gain control developed below.

4.1 The Source Problem and the Meaning of a Signal

Our starting point is Pure-Circuit, a network of local constraints that may contain cycles. Its task is to assign every node a value in $\{0, 1, \perp\}$, where $0, 1$ are called *pure* and $\perp$ is unspecified. A *gate* is one of three local constraints: NOT reverses a pure input; NOR outputs 1 on two 0 inputs and 0 if either input is 1; PURIFY copies a pure input to both outputs and always makes at least one output pure. These requirements apply simultaneously to all gates.

Each node is the output of exactly one gate. The search problem remains PPAD-complete when each node is also the input of exactly one gate and has at most three neighbors (Deligkas et al., 2024, Corollary 2.3). This bounded connectivity will keep the resulting game sparse.

We represent node v by a mutually coupled pair (x_v, y_v), called a *latch*, followed by L scalar relays $r_{v,1}, \dots, r_{v,L}$ as in Section 2. Only the last, or *terminal*, coordinate is decoded: values in $[P, 1]$ mean 1, those in $[-1, -P]$ mean 0, and values in $(-P, P)$ mean $\perp$. The unspecified symbol is part of the decoding, not a third action. Gates must remain valid even when some inputs decode to $\perp$.

Reserving Error and Gain. Fix K, g, Δ from Theorem 1.2, and write $\mathfrak{H} = \mathfrak{B}$ for residuals and $\mathfrak{H} = \mathfrak{G}$ for KL. The preliminary construction, or *core*, must have gain slightly below the requested K to allow a final adjustment. Reserve a field perturbation $\tau = \Delta/16$ for KL and $\tau = \min\{\Delta/16, (2/3 - \bar{g})/4\}$ for residuals, and set

$$a = K - 4\tau, \quad K_c = K - 2\tau, \quad g_c = \begin{cases} g + \tau, & \text{residual,} \\ g + 2\tau, & \text{KL.} \end{cases} \tag{19}$$

Here a is the local relay gain, K_c bounds the core's total gain, and g_c is the error its gates must tolerate. Since $\Delta \le \mathfrak{H}(K) \le K - 1$ and $\mathfrak{H}$ is 1-Lipschitz, $1 < a < K_c < K$ and $\mathfrak{H}(a) - g_c \ge \Delta/2$. Thus both the gain budget and a scalar error margin survive. Table 2 records all parameter dependencies.

4.2 Forcing and Transmitting the Encoded Values

The encoding needs two operations: force a latch to a chosen sign, then carry that sign through a relay chain. We first quantify the field needed for forcing. A latch (x, y) has targets $\tanh(ay + H)$ and $\tanh(ax + H)$. A sufficiently positive external field H forces both coordinates above P, and a sufficiently negative field forces both below $-P$. To quantify the required field, define

$$\mathfrak{T}_\ell(a, g; P) = \sup\{h - az : -1 \le z < P,\ h \in \mathbb{R},\ \ell(z, \tanh h) \le g\}. \tag{20}$$

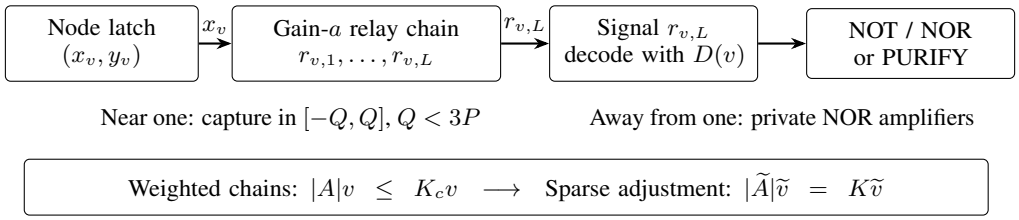


Figure 2: Reduction structure. A gate sets the fields of its output latches; only terminal signals are decoded. Near one, every terminal is captured; away from one, private amplifiers strengthen pure NOR inputs. Weighted chains enforce the spectral budget, and the final adjustment changes each old field by at most τ.

This supremum is the largest external field compatible with a coordinate below P at error g. We use it only where it is finite. Exceeding it forces both coordinates into the positive encoding interval.

Lemma 4.1 (Forcing a Latch)**.** *If $\ell(x, \tanh(ay + H)) \leq g$ and $\ell(y, \tanh(ax + H)) \leq g$, then $H > \mathfrak{T}_\ell(a, g; P)$ implies $x, y \geq P$. By symmetry, $H < -\mathfrak{T}_\ell(a, g; P)$ implies $x, y \leq -P$.*

Proof. Relabel so that $x = \min\{x, y\}$. If $x < P$, the first inequality makes $(z, h) = (x, ay + H)$ feasible in (20). Hence $H \leq ay + H - ax \leq \mathfrak{T}_\ell(a, g; P)$, a contradiction. □

This argument works for either loss; only the required field changes. It remains to choose rational parameters with a positive margin γ below the scalar barrier. Near $a = 1$, we also need every terminal signal to enter a smaller interval $[-Q, Q]$. We call this *capture*; it prevents an uncontrolled input of large magnitude from overwhelming a valid NOR input.

Lemma 4.2 (Signal Parameters Near Criticality)**.** *Let $a = 1 + \eta$, $g_c \geq 0$, and $\gamma > 0$ be rational, with $0 < \eta \leq 10^{-2}$ and $\mathfrak{H}(a) - g_c \geq \gamma$, for either loss. In time polynomial in the input length and $\log(2 + 1/\gamma)$, one can find rational P, Q, h such that*

$$\frac{24}{25}\sqrt{\eta} \leq P \leq \sqrt{\eta}, \qquad Q = 5P/2 < 1,$$
$$\ell(P, \tanh(aP)) > g_c, \qquad h > \mathfrak{T}_\ell(a, g_c; P), \qquad h = O(\eta^{3/2}).$$

At loss at most g_c, the intervals $[P, 1]$ and $[-1, -P]$ are preserved. Every relay sequence enters $[-Q, Q]$ and stays there after at most $L_0 = \lceil 12/Q^2 \rceil + 1 = O(\eta^{-1})$ steps.

Here steps count chain depth. Preservation transmits the sign; capture limits competing inputs to magnitude $Q < 3P$, as NOR requires. Lemmas C.4 and D.1 choose P near a scalar maximizer and bound the forcing field; Lemma C.3 proves the capture length.

Away from one, Lemmas D.1 and D.2 still supply a sign threshold P and bounded forcing fields. Instead of capturing all inputs, we will amplify pure NOR inputs so they dominate any unspecified input. The next subsection explains both routes.

4.3 Boolean Constraints and Decoding

The signal components now implement the source constraints. We must decode every allowed approximate profile correctly. For each circuit node v, let (x_v, y_v) be its latch, and write $r_{v,0} = x_v$ for the start of its chain. Each subsequent coordinate has target $\tanh(ar_{v,t-1})$. Decode the terminal value by

$$D(v) = \begin{cases} 1, & r_{v,L} \geq P, \\ 0, & r_{v,L} \leq -P, \\ \perp, & |r_{v,L}| < P. \end{cases} \tag{21}$$

Near $a = 1$, choose L at least the capture length, so every terminal lies in $[-Q, Q]$. Write r, r_1, r_2 for input terminal signals. For $h > \mathfrak{T}_\ell(a, g_c; P)$, the following fields enforce the three constraints:

$$\begin{aligned} H_{\mathrm{NOT}} &= -2hr/P, \qquad H^{\pm}_{\mathrm{PURIFY}} = 4hr/P \pm 2h, \\ H_{\mathrm{NOR}} &= -c_{\mathrm{N}}(r_1 + r_2) - c_{\mathrm{N}}(P + Q)/2, \quad c_{\mathrm{N}} = 8h/P. \end{aligned} \tag{22}$$

For NOR, the two required cases give fields at least $c_{\mathrm{N}}(3P-Q)/2 \geq 2h$ and at most $-c_{\mathrm{N}}(3P-Q)/2 \leq -2h$. These follow from two negative signals or one positive signal with any other signal in $[-Q, Q]$. NOT reverses a pure input's sign. For PURIFY, both fields have a pure input's sign, and at least one has magnitude at least $2h$ for any input. Forcing fixes the output latch's sign; relay preservation carries it to the terminal. Thus every decoded assignment satisfies all three constraints.

Away from one, a *private amplifier* is an extra latch used only by one NOR gate. Its field $2hr/P$ turns a pure input into an output of the same sign with magnitude at least $P_{\mathrm{amp}} > 1/3$. Choose h large enough to force both thresholds P and P_{amp}. For amplifier outputs u_1, u_2, a field

$$H = -c_{\mathrm{amp}}(u_1 + u_2) - \theta_{\mathrm{amp}}, \quad c_{\mathrm{amp}} = \frac{2h}{3P_{\mathrm{amp}} - 1}, \quad \theta_{\mathrm{amp}} = \frac{c_{\mathrm{amp}}(1 + P_{\mathrm{amp}})}{2} \tag{23}$$

is at least h on two negative pure inputs and at most $-h$ if one input is positive pure, even when the other is anywhere in $[-1, 1]$. For KL, $P_{\mathrm{amp}} = 1/2$ works. For residuals, finite forcing additionally requires $g_c < 1 - P_{\mathrm{amp}}$. A choice with $1/3 < P_{\mathrm{amp}} < 1 - g_c$ therefore exists exactly when $g_c < 2/3$, which the reserved budget guarantees. This explains the residual cap in Theorem 1.2; it is a limitation of these gates, not a general impossibility theorem (Lemmas D.3–D.4).

4.4 Spectral Control and the Final Error Bound

Local gate correctness does not yet bound their combined gain. Relay chains preserve signs while weakening long feedback loops enough to meet the bound K_c.

Let M_0 contain the absolute couplings among latch coordinates, excluding relays, and let $\sigma = K_c - a > 0$. Appendix D constructs rational weights $v \geq \mathbf{1}$ satisfying $K_c v - M_0 v = \sigma \mathbf{1}$. Each latch row therefore has unused allowance σ. Set $\beta = a/K_c < 1$ and give the depth-t relay after coordinate j weight $v_j \beta^t$. Its weighted input is then exactly K_c times its own weight. If D bounds the sum of absolute coefficients from terminal signals into any latch, choose the chain length L so that

$$D v_{\max} \beta^L \leq \sigma. \tag{24}$$

The terminal inputs fit within the unused allowance, giving $|A|v \leq K_c v$. Chains are also at least as long as required for capture. Their lengths and all rational weight descriptions are polynomial in the input size and $1/\Delta$ (Appendix D).

This bounds the gain from above; the theorem requires equality with K. Lemma E.1 adds a sparse cycle of auxiliary coordinates so that the extended positive weights satisfy $|\widetilde{A}|\widetilde{v} = K\widetilde{v}$. For a nonnegative matrix, a positive eigenvector certifies that its eigenvalue is the spectral radius (Perron–Frobenius theory). The construction also makes the influence graph strongly connected, keeps degree at most six, and changes each old field by at most τ.

Why Decoding Survives the Adjustment. The preliminary map uses tolerance g_c from (19). Changing an old field by at most τ changes residual loss by at most τ and KL by at most 2τ, since $|\partial_h d(x, \tanh h)| \leq 2$. Every final solution at tolerance g is therefore a preliminary solution at tolerance g_c, so the decoder satisfies all Pure-Circuit constraints.

Size, Game Realization, and Completeness. For N source nodes, the reduction uses $O(N(1 + \Delta^{-1}) \log(2 + \Delta^{-1}))$ coordinates of polynomial encoding length. Theorem 1.2 combines Theorem D.5 (core), Lemma E.1 (exact gain), Proposition F.1 (bounded payoffs at fixed rationality), and Appendix G (certified answers and PPAD membership).

5 Conclusion

Together, our bounds show that smooth local responses leave a worst-case barrier to equilibrium accuracy. The shared scalar identity explains why contraction algorithms and hardness reductions match under the stated promises.

These guarantees concern response consistency; proximity to selected equilibria requires additional conditioning (Etessami and Yannakakis, 2010). Moreover, residual tolerances $2/3 \leq g < \mathfrak{B}(K)$ at large gain, equality at a barrier, and smaller-than-inverse-polynomial separation remain open. Future work can address these cases and extend the reductions to other responses and multi-action games.

Reproducibility Statement

Section 2 proves the common scalar theorem. The appendices give complete proofs of the finite-precision algorithm, reduction, exact spectral adjustment, game realization, and PPAD membership. All theoretical claims are established analytically; the figures illustrate formulas and proof structure.

AI Use Statement

Generative AI tools assisted with literature discovery, analysis of proofs and parameter bounds, development and checking of proof arguments, verification code, and drafting and editing. The authors are responsible for the correctness and originality of every mathematical claim, the accuracy of citations, the artifacts, and the final content.

References

S. Bai, J. Z. Kolter, and V. Koltun. Deep equilibrium models. In *Advances in Neural Information Processing Systems*, volume 32, 2019. https://arxiv.org/abs/1909.01377.

A. Baker. *Transcendental Number Theory*. Cambridge University Press, 1975. https://doi.org/10.1017/CBO9780511565977.

A. Berman and R. J. Plemmons. *Nonnegative Matrices in the Mathematical Sciences*. SIAM, 1994. https://doi.org/10.1137/1.9781611971262.

M. Blondel, A. F. T. Martins, and V. Niculae. Learning with Fenchel–Young losses. *Journal of Machine Learning Research*, 21(35):1–69, 2020. https://jmlr.org/papers/v21/19-021.html.

S. Cen, Y. Wei, and Y. Chi. Fast policy extragradient methods for competitive games with entropy regularization. In *Advances in Neural Information Processing Systems*, volume 34, pages 27952–27964, 2021. https://proceedings.neurips.cc/paper/2021/hash/eb1848290d5a7de9c9ccabc67fefa211-Abstract.html.

S. Cen, Y. Chi, S. S. Du, and L. Xiao. Faster last-iterate convergence of policy optimization in zero-sum Markov games. In *International Conference on Learning Representations*, 2023. https://openreview.net/forum?id=bRwBpKrNzF7.

X. Chen, Y. Li, and M. Yannakakis. Computing a fixed point of contraction maps in polynomial queries. In *STOC*, pages 1364–1373, 2024.

J. Černý, B. An, and A. N. Zhang. Quantal correlated equilibrium in normal form games. In *Proceedings of the 23rd ACM Conference on Economics and Computation (EC)*, pages 210–239, 2022. https://doi.org/10.1145/3490486.3538350.

C. Daskalakis, N. Golowich, N. Haghtalab, and A. Shetty. Smooth Nash equilibria: Algorithms and complexity. In *15th Innovations in Theoretical Computer Science Conference (ITCS)*, volume 287 of *LIPIcs*, pages 37:1–37:22, 2024. https://doi.org/10.4230/LIPIcs.ITCS.2024.37.

C. Daskalakis, G. Farina, and B. H. Zhang. Polynomial-time computation of ℓ_p-contraction fixed points for even p. *arXiv:2607.23863*, 2026. https://arxiv.org/abs/2607.23863.

E. del Barrio, H. Inouzhe, and C. Matrán. Box-constrained monotone approximations to Lipschitz regularizations, with applications to robust testing. *Journal of Optimization Theory and Applications*, 187(1):65–87, 2020. https://doi.org/10.1007/s10957-020-01743-5.

A. Deligkas, J. Fearnley, A. Hollender, and T. Melissourgos. Pure-Circuit: Tight inapproximability for PPAD. *Journal of the ACM*, 71(5), Article 31, 2024. https://arxiv.org/abs/2209.15149.

K. Etessami and M. Yannakakis. On the complexity of Nash equilibria and other fixed points. *SIAM Journal on Computing*, 39(6):2531–2597, 2010. https://doi.org/10.1137/080720826.

J. K. Goeree, C. A. Holt, and R. T. Palfrey. *Quantal Response Equilibrium: A Stochastic Theory of Games*. Princeton University Press, 2016.

S. Haslebacher, J. Lill, P. Schnider, and S. Weber. Query-efficient fixpoints of ℓ_p-contractions. In *FOCS*, 2025. https://arxiv.org/abs/2503.16089.

K.-I. Ko. *Complexity Theory of Real Functions*. Birkhäuser, 1991. https://doi.org/10.1007/978-1-4684-6802-1.

N. M. Mayer. Echo state condition at the critical point. *Entropy*, 19(1):3, 2017. https://doi.org/10.3390/e19010003.

R. D. McKelvey and T. R. Palfrey. Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1):6–38, 1995.

E. Melo. On the uniqueness of quantal response equilibria and its application to network games. *Economic Theory*, 74:681–725, 2022. https://doi.org/10.1007/s00199-021-01385-2.

M. Nowicka and A. Witkowski. Optimal bounds for the sine and hyperbolic tangent means. *Punjab University Journal of Mathematics*, 52(5):77–88, 2020. https://pu.edu.pk/images/journal/maths/PDF/Paper_6_52_5_2020.pdf.

C. H. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *Journal of Computer and System Sciences*, 48(3):498–532, 1994.

S. Sokota, R. D'Orazio, J. Z. Kolter, N. Loizou, M. Lanctot, I. Mitliagkas, N. Brown, and C. Kroer. A unified approach to reinforcement learning, quantal response equilibria, and two-player zero-sum games. In *International Conference on Learning Representations*, 2023. https://openreview.net/forum?id=ndZ42T8iUmd.

Y. Sun, T. Liu, P. R. Kumar, and S. Shahrampour. Linear convergence of independent natural policy gradient in games with entropy regularization. *arXiv:2405.02769*, 2024. https://arxiv.org/abs/2405.02769.

T. L. Turocy. A dynamic homotopy interpretation of the logistic quantal response equilibrium correspondence. *Games and Economic Behavior*, 51(2):243–263, 2005.

T. Wang. Sharp spectral thresholds for logit fixed points. *arXiv:2605.15651*, 2026. https://arxiv.org/abs/2605.15651.

E. Winston and J. Z. Kolter. Monotone operator equilibrium networks. In *Advances in Neural Information Processing Systems*, volume 33, 2020. https://arxiv.org/abs/2006.08591.

Z. Zhao, A. Shukla, R. Jain, and V. G. Subramanian. Efficient decentralized learning of generalized quantal response equilibrium. In *UAI*, volume 337 of *Proceedings of Machine Learning Research*, pages 8182–8208, 2026. https://proceedings.mlr.press/v337/zhao26a.html.

A Scalar Thresholds and Losses

Section 2 identifies the common scalar mechanism. Here we verify the loss assumptions and derive the quantitative barrier bounds used by the algorithm and reduction.

A.1 Loss Assumptions

Theorem 2.2 uses only the order conditions in Definition 2.1. Convexity is one sufficient condition for (ii): if $x' = (1-\theta)x + \theta y$ with $0 < \theta \le 1$, then

$$\ell(x', y) \le (1-\theta)\ell(x, y) < \ell(x, y).$$

Here the strict inequality follows from $x < y$ and the unique zero at $x = y$. The assumptions also allow $|x-y|^p$ for all $p > 0$, including the nonconvex case $p < 1$. For Bernoulli KL, symmetry and positivity hold, convexity in x gives (ii), and

$$\partial_y d(x, y) = \frac{y-x}{1-y^2} \ge 0 \quad (x \le y)$$

gives (iii). No inverse loss function is needed.

Relation to Classical Uniform Approximation. For absolute loss, the envelope formula of del Barrio et al. (2020, Lemmas 3.1–3.2) recovers the optimal value. Specializing it to $\tanh$ and Lipschitz constant $1/a$ gives

$$\inf_{\mathrm{Lip}(f) \le 1/a} \|f - \tanh\|_\infty = \frac{1}{2} \sup_{u \ge v} \left(\tanh u - \tanh v - \frac{u-v}{a} \right) = \sup_{t \ge 0} \left(\tanh t - \frac{t}{a} \right) = \mathfrak{B}(a).$$

Indeed, at fixed interval length $u - v = 2t$, the integral of $\tanh' = \mathrm{sech}^2$ is largest on $[-t, t]$, since this derivative is even and decreases away from zero. The formula on the real line follows by applying the compact-interval result on $[-T, T]$ and letting T grow; χ_a supplies a global matching upper bound. Restricting the range to $[-1, 1]$ does not change the infimum, since projection onto this interval cannot increase absolute error. Thus we do not claim the residual approximation value itself as new. Theorem 2.2 provides a common optimizer for the full admissible-loss family and identifies its value with the sign-preserving relay tolerance; the computational lower bound needs the separate reduction.

A.2 The Residual Threshold

We first maximize the scalar response excess $R_a(z) = \tanh(az) - z$. Besides giving the barrier explicitly, the calculation controls how much it changes when the algorithm increases the gain.

Lemma A.1. *For $0 \le a \le 1$, $\mathfrak{B}(a) = 0$. For $a > 1$, put*

$$t_a := \operatorname{arsinh} \sqrt{a-1}, \qquad z_a := t_a / a.$$

Then z_a is the unique maximizer of R_a and

$$\mathfrak{B}(a) = \sqrt{\frac{a-1}{a}} - \frac{\operatorname{arsinh} \sqrt{a-1}}{a}. \tag{25}$$

Moreover $\mathfrak{B}$ is strictly increasing on $(1, \infty)$, is 1-Lipschitz on $[0, \infty)$, and tends to 1 as $a \to \infty$. For $1 \le u \le v$,

$$0 \le \mathfrak{B}(v) - \mathfrak{B}(u) \le \tfrac{1}{2} \log(v/u). \tag{26}$$

Also

$$\mathfrak{B}(1+\eta) = \frac{2}{3}\eta^{3/2} - \frac{13}{15}\eta^{5/2} + O(\eta^{7/2}), \qquad \frac{2\eta^{3/2}}{3(1+\eta)^{3/2}} \le \mathfrak{B}(1+\eta) \le \frac{2\eta^{3/2}}{3(1+\eta)} \qquad (\eta > 0).$$

Proof. For $a \le 1$, $\tanh(az) \le az \le z$. If $a > 1$, then

$$R_a'(z) = a \operatorname{sech}^2(az) - 1, \qquad R_a''(z) = -2a^2 \operatorname{sech}^2(az) \tanh(az) < 0$$

for $z > 0$. The unique critical point satisfies $\cosh^2(az) = a$, hence $az = t_a$ and $\tanh(t_a) = \sqrt{(a-1)/a}$, proving (25).

Differentiation gives

$$0 < a\mathfrak{B}'(a) = \frac{t_a}{a} \le \frac{\sqrt{a-1}}{a} \le \frac{1}{2}.$$

The last inequality is $2\sqrt{a-1} \le a$. Integration proves (26); continuity at one and extension by zero give monotonicity and the global 1-Lipschitz bound. The closed form gives $\mathfrak{B}(a) \to 1$ as $a \to \infty$. Thus $\mathfrak{B}(K_{\mathrm{aff}}) = 2/3$ has a unique solution, and certified interval evaluation gives $6.27098 < K_{\mathrm{aff}} < 6.27099$. Differentiating gives $\frac{d}{da}[a\mathfrak{B}(a)] = \sqrt{(a-1)/a}$, hence

$$\frac{2\eta^{3/2}}{3(1+\eta)^{3/2}} \le \mathfrak{B}(1+\eta) = \frac{1}{1+\eta}\int_0^\eta \sqrt{\frac{t}{1+t}}\,dt \le \frac{2\eta^{3/2}}{3(1+\eta)}. \tag{27}$$

The inequalities follow by bounding $(1+t)^{-1/2}$ inside the integral. For the expansion, integrate $(1+t)^{-1/2} = 1 - t/2 + O(t^2)$ to obtain

$$\mathfrak{B}(1+\eta) = \frac{\frac{2}{3}\eta^{3/2} - \frac{1}{5}\eta^{5/2} + O(\eta^{7/2})}{1+\eta}.$$

□

For absolute loss, Theorem 2.2 gives

$$\mathfrak{H}_{\ell_1}(a) = \max_{0\le z\le s_a} |z - \tanh(az)| = \max_{z\ge 0} R_a(z) = \mathfrak{B}(a),$$

since $R_a \ge 0$ on $[0, s_a]$ and is negative after s_a.

A.3 The KL Threshold

For KL there is no comparable closed form for the maximizing point. Writing the target as $\tanh h$ gives a formula in the field h that supports both stable evaluation and derivative bounds. Let Φ be the negative entropy of a binary distribution:

$$\Phi(x) := \frac{1+x}{2}\log\frac{1+x}{2} + \frac{1-x}{2}\log\frac{1-x}{2}.$$

For $x \in [-1, 1]$ and $h \in \mathbb{R}$, using $0 \log 0 = 0$,

$$d(x, \tanh h) = \Phi(x) - hx + \log(2\cosh h), \tag{28}$$

so, for $|x| < 1$,

$$\partial_x d(x, \tanh h) = \operatorname{artanh}(x) - h, \qquad \partial_h d(x, \tanh h) = \tanh h - x.$$

The derivative with respect to the first argument can diverge at ± 1. The next bound instead uses convexity to control error all the way to those endpoints.

Lemma A.2. *Let $h \in \mathbb{R}$, $x \in [-1, 1]$, and let z lie between zero and $\tanh h$. Then*

$$d(x, \tanh h) \le d(z, \tanh h) + (|h| + \log 2)|x - z|. \tag{29}$$

For each h, the coefficient $|h| + \log 2$ is best possible under these hypotheses. In particular, for any affine logit map and $R = \max_i(|b_i| + \sum_j |A_{ij}|)$,

$$\mathcal{L}_d(x; F) \le (R + \log 2)\,\mathrm{Res}_F(x).$$

Proof. By symmetry assume $h \ge 0$, so $0 \le z \le \tanh h < 1$. Put $f(x) = d(x, \tanh h)$ and $H(z) = -\Phi(z) \le \log 2$. The slopes from z to the two endpoints satisfy

$$\frac{f(-1) - f(z)}{1+z} = h + \frac{H(z)}{1+z} \le h + \log 2,$$

and, since $h \ge \operatorname{artanh} z$,

$$\frac{f(1) - f(z)}{1-z} = \frac{H(z)}{1-z} - h \le \frac{\log(2/(1+z))}{1-z} \le \log 2.$$

The last inequality is $\log(1+z) \ge z \log 2$, by concavity of the logarithm. Convexity of f bounds it by the chord on each of $[-1, z]$ and $[z, 1]$, proving (29). Equality holds for $z = 0, x = -1$. For the final claim, take $z = F_i(x)$ and $h = (Ax + b)_i$. □

The preceding bound converts a modified-map residual into KL loss. We next study the best scalar KL loss itself, locating its maximizer and comparing its value with the residual barrier. Convexity below concerns the gain a; concavity concerns the maximization variable z.

Lemma A.3. *After extension by zero below one, $\mathfrak{G}$ is increasing, convex, and $1/4$-Lipschitz on $[0,\infty)$. For $a>1$, put*

$$t_a = \operatorname{arsinh}\sqrt{a-1}, \qquad z_a = t_a/a, \qquad c_a = \tanh t_a = \sqrt{(a-1)/a}.$$

The function $J_a(z) = d(z,\tanh(az))$ is strictly concave on $[z_a, c_a] \subset (0, s_a)$, which contains its unique maximum on $[0, s_a]$. Moreover,

$$\frac{\mathfrak{B}(a)^2}{2(1-z_a^2)} \le \mathfrak{G}(a) \le \frac{a\mathfrak{B}(a)^2}{2}. \tag{30}$$

Consequently,

$$\mathfrak{G}(1+\eta) = \frac{2}{9}\eta^3 - \frac{16}{45}\eta^4 + O(\eta^5).$$

Proof. In natural coordinates,

$$d(x,\tanh h) = \Phi(x) - hx + \log(2\cosh h), \qquad \partial_a d(z,\tanh(az)) = zR_a(z). \tag{31}$$

Before $a = \operatorname{artanh}(z)/z$, the point z lies outside the maximizing interval; at that value its loss is zero. Thus, with $r_+ = \max\{r,0\}$,

$$\mathfrak{G}(a) = \max_{0\le z\le 1} \int_0^a z\big(\tanh(tz) - z\big)_+ \, dt. \tag{32}$$

Every integrand is nondecreasing in t and lies in $[0, z(1-z)] \subseteq [0,1/4]$. Each integral is therefore convex in a; taking the maximum proves that $\mathfrak{G}$ is increasing, convex, and $1/4$-Lipschitz on $[0,\infty)$.

To locate the maximum, write $J_a(z) = d(z,\tanh(az))$. For $0 < z < s_a$,

$$J_a'(z) = \int_{\operatorname{artanh} z}^{az} (a\operatorname{sech}^2 t - 1)\, dt. \tag{33}$$

The integrand changes sign at t_a. Hence $J_a' > 0$ for $z \le z_a$ and $J_a' < 0$ for $z \ge c_a$ within this interval. Also $z_a < c_a$ because $R_a(z_a) > 0$, and $\tanh(ac_a) > \tanh(az_a) = c_a$, so $c_a < s_a$. On $[z_a, c_a]$,

$$J_a''(z) = \frac{1}{1-z^2} + a^2\operatorname{sech}^2(az) - 2a < 0 :$$

each of the first two terms is at most a, and equality in both would require $z = c_a = z_a$. This proves strict concavity and uniqueness.

The lower bound retains the location of the residual maximum:

$$\mathfrak{G}(a) \ge d(z_a, c_a) = \int_{z_a}^{c_a} \frac{u - z_a}{1-u^2}\, du \ge \frac{(c_a - z_a)^2}{2(1-z_a^2)}.$$

For the upper bound, fix $z \in [0, s_a]$ and put $v = az - \operatorname{artanh} z \ge 0$. Then

$$\begin{aligned} d(z,\tanh(az)) &= \int_{az-v}^{az} [R_a(t/a) + t/a - z]\, dt \\ &\le \mathfrak{B}(a)v - \frac{v^2}{2a} \le \frac{a\mathfrak{B}(a)^2}{2}. \end{aligned}$$

This proves (30). Combining its weaker lower bound $\mathfrak{G} \ge \mathfrak{B}^2/2$ with (27) gives

$$\frac{2\eta^3}{9(1+\eta)^3} \le \mathfrak{G}(1+\eta) \le \frac{2\eta^3}{9(1+\eta)}, \qquad \eta > 0. \tag{34}$$

For the second-order expansion, put $a = 1+\eta$. Since $c_a - z_a = \mathfrak{B}(a) = O(\eta^{3/2})$ and $c_a + z_a = O(\sqrt{\eta})$, (30) gives

$$0 \le \frac{a\mathfrak{B}(a)^2}{2} - \mathfrak{G}(a) \le \frac{a\mathfrak{B}(a)^3(c_a + z_a)}{2(1-z_a^2)} = O(\eta^5). \tag{35}$$

Thus $\mathfrak{G}(1+\eta) = (1+\eta)\mathfrak{B}(1+\eta)^2/2 + O(\eta^5)$. Substituting the expansion from Lemma A.1 gives the stated coefficients. □

Proof of Corollary 2.3. The identities with $\mathfrak{B}$ and $\mathfrak{G}$ follow from Theorem 2.2 and the definitions above. The closed form and all stated properties are Lemmas A.1 and A.3. □

Proof of Theorem 2.4. Lemma C.2 gives $s_{1+\eta} = \Theta(\sqrt{\eta})$, so both z and $R_{1+\eta}(z)$ tend uniformly to zero on the maximizing interval. For every $e \in (0, c)$, the loss is therefore between $(c-e)R_{1+\eta}(z)^p$ and $(c+e)R_{1+\eta}(z)^p$ when η is small. Maximizing and then letting $e \downarrow 0$ proves the first equivalence. The closed form (11) gives $\mathfrak{B}(1+\eta) = \frac{2}{3}\eta^{3/2}(1+O(\eta))$, proving the second. □

A separate arithmetic observation clarifies what equality at the residual barrier means. An algebraic number is a root of a nonzero polynomial with rational coefficients; a transcendental number is not.

Theorem A.4 (Algebraic Inputs Do Not Attain the Residual Threshold)**.** *For every algebraic $a > 1$, $\mathfrak{B}(a)$ is transcendental. Consequently, if A has algebraic entries and $\rho(|A|) > 1$, no algebraic tolerance equals $\mathfrak{B}(\rho(|A|))$.*

Proof. Put $c = \sqrt{1 - 1/a} \in (0, 1)$. The closed form is $\mathfrak{B}(a) = c - \operatorname{artanh}(c)/a$. If both a and $\mathfrak{B}(a)$ were algebraic, then $\alpha = 2a(c - \mathfrak{B}(a)) = 2\operatorname{artanh} c$ would be nonzero and algebraic, yet $e^\alpha = (1+c)/(1-c)$ would also be algebraic. This contradicts Hermite–Lindemann (Baker, 1975, Chapter 1). If A is algebraic, so are $|A|$ and every root of its characteristic polynomial, including $\rho(|A|)$. □

Thus rationally encoded instances cannot hit this barrier exactly. They can still approach it arbitrarily closely, so this observation does not remove the margin requirement in our complexity results.

A.4 Scalar Duality Beyond the Logit Response

The following extension separates the general scalar mechanism from the logit-specific complexity reductions. No algorithmic or PPAD-hardness claim for a new response family is part of this proposition.

Proposition A.5. *Let $\phi : \mathbb{R} \to (-1, 1)$ be odd, continuously differentiable, increasing, and strictly concave on $(0, \infty)$, with $\phi'(0) = 1$ and $\lim_{t\to\infty} \phi(t) = 1$. For $a > 1$, let $s_{a,\phi} > 0$ be the unique root of $s = \phi(as)$, and put*

$$\mathfrak{H}_{\ell,\phi}(a) = \max_{0 \le z \le s_{a,\phi}} \ell(z, \phi(az)), \qquad \chi_{a,\phi}(t) = \operatorname{sgn}(t) \min\{|t|/a, \phi(|t|)\}.$$

For every admissible loss from Definition 2.1,

$$\inf_{\substack{f:\mathbb{R}\to[-1,1] \\ \operatorname{Lip}(f)\le 1/a}} \sup_t \ell(f(t), \phi(t)) = \mathfrak{H}_{\ell,\phi}(a) = \max_{0 < P \le s_{a,\phi}} \Gamma_{\ell,\phi}(a, P), \tag{36}$$

where $\Gamma_{\ell,\phi}(a, P)$ is the largest tolerance preserving $[P, 1]$ under $\ell(y, \phi(ax)) \le g$. The clipped response attains the infimum, and $\Gamma_{\ell,\phi}(a, P) = \ell(P, \phi(aP))$, including equality.

If also $\phi(t) = t - \beta t^3 + o(t^3)$ as $t \to 0$, with $\beta > 0$, and $\ell(x, x + r) = c|r|^p(1 + o(1))$ uniformly as in Theorem 2.4, then

$$\mathfrak{H}_{\ell,\phi}(1+\eta) = c\left(\frac{2}{3\sqrt{3\beta}}\right)^p \eta^{3p/2}(1 + o(1)). \tag{37}$$

Proof. Strict concavity implies that $\phi(t)/t$ decreases strictly from 1 to 0 on $(0, \infty)$. Hence it equals $1/a$ at exactly one point $t_a = as_{a,\phi}$, with $s_{a,\phi} = \phi(t_a) \in (0, 1)$. The clipped response is t/a for $|t| \le t_a$ and $\phi(t)$ outside. Concavity gives $0 \le \phi'(t) \le \phi(t_a)/t_a = 1/a$ for $t \ge t_a$; oddness gives the same bound on the negative side. Thus $\operatorname{Lip}(\chi_{a,\phi}) \le 1/a$, and its loss is precisely the middle term of (36).

Conversely, for $0 \le t \le t_a$, every $1/a$-Lipschitz f satisfies $f(t) - f(-t) \le 2t/a$. At least one of $f(t), -f(-t)$ is at most $t/a \le \phi(t)$. Symmetry and monotonicity of the loss imply that one of the two losses is at least $\ell(t/a, \phi(t))$. Maximizing gives the minimax lower bound. For a relay with $x \ge P$ and $y < P$,

$$\ell(y, \phi(ax)) > \ell(P, \phi(ax)) \ge \ell(P, \phi(aP)).$$

This proves preservation at the stated tolerance; continuity admits an output just below P at any larger tolerance. The last equality follows.

For the critical law, write $s = s_{1+\eta,\phi}$. First $s \to 0$ as $\eta \downarrow 0$: otherwise a positive limit point would satisfy $s = \phi(s)$, contradicting strict concavity and $\phi'(0) = 1$. The fixed-point equation and the expansion of ϕ then give $s^2 = \eta/\beta\,(1 + o(1))$. Uniformly for bounded nonnegative y,

$$\phi((1+\eta)\sqrt{\eta}\,y) - \sqrt{\eta}\,y = \eta^{3/2}\big(y - \beta y^3 + o(1)\big).$$

The rescaled maximizing interval tends to $[0, 1/\sqrt{\beta}]$. The maximum of $y - \beta y^3$ on that interval is $2/(3\sqrt{3\beta})$, attained at $y = 1/\sqrt{3\beta}$. Finally the uniform local expansion of ℓ compares its maximum with c times the pth power of this residual maximum, exactly as in Theorem 2.4. This proves (37). □

For example, $\phi(t) = t/\sqrt{1+t^2}$ has $\beta = 1/2$ and residual coefficient $2\sqrt{2}/(3\sqrt{3})$; the scaled arctangent $\phi(t) = (2/\pi)\arctan(\pi t/2)$ has $\beta = \pi^2/12$ and coefficient $4/(3\pi)$. Their leading KL coefficients are $4/27$ and $8/(9\pi^2)$, respectively. For $\tanh$, $\beta = 1/3$ recovers the coefficients $2/3$ and $2/9$. These examples describe scalar barriers; no matching computational lower bound for them is asserted here.

B Uniqueness, Spectral Certificates, and Finite-Precision Algorithms

The scalar approximation becomes an algorithm only after feedback, iteration error, and arithmetic precision are controlled. We first justify the uniqueness boundary used in the introduction, then construct the rational weights and evaluation bounds for Section 3. Algorithms for contractions or selected nonexpansive maps (Chen et al., 2024; Haslebacher et al., 2025; Daskalakis et al., 2026) do not directly apply beyond gain one. Implicit networks (Bai et al., 2019; Winston and Kolter, 2020) motivate related fixed-point questions; the present guarantees concern affine $\tanh$ maps only.

B.1 The Exact Worst-Case Uniqueness Threshold

Critical convergence under a singular-value condition is studied by Mayer (2017). The following comparison uses $\rho(|A|)$ and also covers reducible matrices.

Theorem B.1 (Uniqueness and Convergence at and Below One)**.** *Let $F(x) = \tanh(Ax + b)$ and $M = |A|$.*

(i) If $\rho(M) < 1$, then F is a contraction in a weighted max norm, has a unique fixed point, and simultaneous iteration converges geometrically.

(ii) If $\rho(M) = 1$, the fixed point is still unique, and simultaneous iteration converges to it from every point of $[-1, 1]^n$.

(iii) For every $\kappa > 1$, a two-coordinate affine logit system with $\rho(|A|) = \kappa$ has three fixed points.

For any M, define $u^0 = \mathbf{1}$ and $u^{t+1} = \tanh(Mu^t)$. Then

$$|F^t(x) - F^t(y)| \le 2u^t \qquad (x, y \in [-1, 1]^n). \tag{38}$$

This bound is attained by $A = M$, $b = 0$, $x = \mathbf{1}$, $y = -\mathbf{1}$. If $\rho(M) \le 1$, then $u^t \downarrow 0$.

Proof. For real r, s, the identity $2\cosh r \cosh s = \cosh(r+s) + \cosh(r-s)$ gives

$$|\tanh r - \tanh s| \le 2\tanh(|r-s|/2).$$

Hence $|F(x) - F(y)| \le 2\tanh(M|x-y|/2)$, and induction proves (38); the stated example attains equality. Monotonicity gives $\mathbf{1} = u^0 \ge u^1 \ge \cdots \to u^* \ge 0$, with $u^* = \tanh(Mu^*)$. If $S = \{i : u^*_i > 0\}$ were nonempty, then

$$0 < u^*_S < M_{SS}u^*_S,$$

The minimum coordinate ratio $(M_{SS}u_S^*)_i/u_i^*$ is therefore greater than one. The Collatz–Wielandt bound gives $\rho(M) \geq \rho(M_{SS}) > 1$. Thus $u^* = 0$ when $\rho(M) \leq 1$. Brouwer gives a fixed point, and (38) gives uniqueness and convergence to it.

For the geometric rate in (i), choose $\rho(M) < q < 1$ and $v = (I - M/q)^{-1}\mathbf{1} > 0$. Then $Mv = q(v - \mathbf{1}) \leq qv$, so F is q-Lipschitz in $\|z\|_v = \max_i |z_i|/v_i$. For (iii), take $A = \left(\begin{smallmatrix} 0 & \kappa \\ \kappa & 0 \end{smallmatrix}\right)$, $b = 0$. The symmetric fixed points are 0 and $\pm s_\kappa \mathbf{1}$. □

The theorem gives convergence at gain one without a geometric rate. Under a positive weight certificate, the next result quantifies the slower decay and shows its order is unavoidable.

Corollary B.2. *Suppose $Mv \leq v$ for some $v > 0$, normalized by $\min_i v_i = 1$. For $x^{t+1} = F(x^t)$, let $e_t = \|x^t - x^*\|_v$, where x^* is the unique fixed point. Then*

$$e_t \leq \frac{e_0}{\sqrt{1 + te_0^2/6}}. \tag{39}$$

The weight condition holds when M is irreducible and $\rho(M) \leq 1$. The order $t^{-1/2}$ cannot be improved in general.

Proof. The preceding scalar comparison and concavity of $\tanh$ on $[0, \infty)$ give $e_{t+1} \leq 2\tanh(e_t/2)$, since $v_i \geq 1$. For $0 \leq y < 1$, concavity of $w^{2/3}$ gives

$$\frac{d}{dy}\frac{y}{\sqrt{1 - 2y^2/3}} = (1 - 2y^2/3)^{-3/2} \leq (1 - y^2)^{-1}.$$

Integration and inversion yield

$$\tanh r \leq \frac{r}{\sqrt{1 + 2r^2/3}}, \qquad r \geq 0. \tag{40}$$

On $(0, 1)$ this is also the $\alpha = 2/3$ case of Nowicka and Witkowski (2020, Theorem 4.2). Thus $e_{t+1}^{-2} \geq e_t^{-2} + 1/6$ while $e_t > 0$, proving (39); zero error stays zero. For sharpness of the order, use the two-coordinate example above with $\kappa = 1$ and $x^0 = \mathbf{1}$. Both coordinates obey $z_{t+1} = \tanh z_t$. Since $z_{t+1}^{-2} - z_t^{-2} \to 2/3$, we have $z_t \sim \sqrt{3/(2t)}$. □

Why Change the Response? The uniqueness analysis also clarifies why the upper-bound algorithm modifies the response instead of relying on damping alone. The following comparisons concern contraction and worst scalar error, rather than practical performance on particular instances.

For $A = \left(\begin{smallmatrix} 0 & a \\ a & 0 \end{smallmatrix}\right)$, $b = 0$, and $a > 1$, the damped map $(1 - \omega)x + \omega F(x)$ has derivative eigenvalue $1 + \omega(a - 1) > 1$ at zero for every $0 < \omega \leq 1$. Thus damping alone cannot ensure contraction. Temperature scaling $t \mapsto \tanh(t/a)$ meets the same slope bound but has strictly larger worst scalar loss than χ_a. Indeed, a maximizer z of $\mathfrak{H}_\ell(a)$ lies in $(0, s_a)$, and $\tanh z < z$ implies $\ell(\tanh z, \tanh(az)) > \mathfrak{H}_\ell(a)$ by admissibility. Our algorithm controls approximation bias and iteration error separately.

B.2 A Rational Perron Certificate with Polynomial Encoding Length

Beyond gain one, the algorithm needs a coordinate scaling for the modified response. We now turn the spectral-radius condition into rational matrix inequalities that can be computed and checked efficiently.

Lemma B.3. *Let $M \geq 0$ and $\tau > 0$ be rational. In time polynomial in their binary encoding lengths and $\log(2 + 1/\tau)$, one can find rational q and $v \in \mathbb{Q}^n$ of polynomial encoding length such that*

$$\rho(M) < q \leq \rho(M) + \tau, \qquad v \geq \mathbf{1}, \qquad Mv \leq qv - \mathbf{1}. \tag{41}$$

Proof. For $q > 0$, solve $(qI - M)u = \mathbf{1}$ exactly and reject singular systems. When $q > \rho(M)$, the Neumann series for $(qI - M)^{-1}$ gives $u > 0$. Conversely, take a nonzero $w \geq 0$ with $w^\mathsf{T} M = \rho(M)w^\mathsf{T}$, a left Perron vector. If $u > 0$, it gives

$$(q - \rho(M))w^\mathsf{T} u = w^\mathsf{T}\mathbf{1} > 0.$$

Thus positivity is equivalent to $q > \rho(M)$. Bisect between zero and $1 + \max_i \sum_j M_{ij}$, retaining the accepted upper endpoint, until the interval has width at most τ. Set $v = u \max\{1, 1/\min_i u_i\}$. Exact elimination and Cramer's rule give polynomial time and encoding length. No linear optimization is needed. □

B.3 Inexact Contraction Iteration

The weight certificate controls the exact contraction. We next track how bounded evaluation errors accumulate and how an iterate's distance to the modified fixed point bounds its computable residual.

Lemma B.4. *Let $C = \prod_i [-T_i, T_i]$, let $v \geq \mathbf{1}$, and suppose $S : C \to C$ obeys*

$$\|S(x) - S(y)\|_v \leq r\|x - y\|_v, \qquad 0 \leq r < 1.$$

Assume an evaluator $\widetilde{S}$ satisfies $\|\widetilde{S}(x) - S(x)\|_\infty \leq \zeta$, and define $\widehat{x}^{t+1} = \Pi_C(\widetilde{S}(\widehat{x}^t))$ by coordinatewise projection. If x^* is the fixed point of S, then

$$\|\widehat{x}^t - x^*\|_v \leq r^t\|\widehat{x}^0 - x^*\|_v + \frac{\zeta}{1-r}, \tag{42}$$

$$\|\widehat{x}^t - S(\widehat{x}^t)\|_\infty \leq v_{\max}(1+r)\left(r^t\|\widehat{x}^0 - x^*\|_v + \frac{\zeta}{1-r}\right). \tag{43}$$

Proof. Coordinatewise projection onto C is nonexpansive in every weighted max norm and fixes $S(x)$. Since $v_i \geq 1$, an ℓ_∞ evaluation error of size ζ has weighted norm at most ζ. Thus

$$e_{t+1} := \|\widehat{x}^{t+1} - x^*\|_v \leq re_t + \zeta.$$

Unrolling the recurrence gives (42). Finally,

$$\|\widehat{x}^t - S(\widehat{x}^t)\|_v \leq \|\widehat{x}^t - x^*\|_v + \|S(x^*) - S(\widehat{x}^t)\|_v \leq (1+r)e_t,$$

and $\|z\|_\infty \leq v_{\max}\|z\|_v$ gives (43). □

The error recurrence assumes access to an approximate evaluator. The following lemma supplies one for the clipped response without having to locate its switching point exactly.

Lemma B.5. *Given rational $a > 1$, rational t, and an integer precision m, one can compute a rational number within 2^{-m} of $\chi_a(t)$ in time polynomial in the input length and m.*

Proof. Use (7). Division, absolute value, and minimum are polynomial-time rational operations, and minimum does not amplify an additive evaluation error. The function $\tanh$ is polynomial-time computable to m bits: for $|t|$ above a threshold of order $m + 1$, use $1 - |\tanh t| \leq 2e^{-2|t|}$; on the remaining interval, evaluate the exponential with polynomially many bits. The sign of rational t is exactly known. This computes $\chi_a(t)$ without approximating s_a. □

B.4 Finite-Precision Parameters for Theorem 1.1

The previous two lemmas provide the iteration estimate and its evaluator. We now choose their precision so that the original-map loss, not merely the modified-map residual, meets the requested tolerance. For residuals, (16) gives $(1-r)^{-1} \leq 17/\delta_0$, independently of κ. For KL the bound is $16(K_0 + 2)/\delta_0$. The residual algorithm may immediately return zero when $g \geq 1$. To make the stopping rule explicit, let e denote the desired unweighted defect and start at zero. In Lemma B.4, it suffices to take

$$\zeta \leq \frac{e(1-r)}{4v_{\max}}, \qquad r^t \leq \frac{e}{4v_{\max}}.$$

Then $\|x^0 - x^*\|_v \leq 1$ and $1 + r \leq 2$ give defect at most e. For residuals use $e = \delta_0/2$. For KL, put $R = \max_i(|b_i| + \sum_j |A_{ij}|)$ and use $e = \delta_0/[8(1+R)]$ as in Section 3. Each coordinate of $S(x)$ lies between zero and $F_i(x)$, so Lemma A.2 bounds the extra KL loss by

$$(R + \log 2)e \leq \delta_0/8.$$

Projection onto $[-1,1]^n$ works for both losses. In the KL case, $\log(1/e) = O(\log(1/\delta_0) + \log(1 + R))$, which is polynomial in the input length. No separation from the boundary of the cube is needed.

KL certificates can also be evaluated with polynomially many bits. Use (28) and the stable identity

$$\log(2\cosh h) = |h| + \log(1 + e^{-2|h|}).$$

Truncate the last exponential once its contribution is below the requested precision; otherwise its argument has polynomial size. The entropy terms at rational coordinates are evaluated directly, with $0 \log 0 = 0$. Thus large rational fields require only their encoding length and the requested precision.

B.5 Choosing a Margin from the Requested Accuracy

The preceding algorithm accepts a known safe margin. We now show how to find one from the requested tolerance, under the promise that it is strictly above the barrier. This avoids any exact comparison of irrational values.

Proposition B.6 (Certified Margin Search)**.** *Let $H = \mathfrak{B}$, or let $H = \mathfrak{G}$ with a fixed cap $\kappa \le K_0$. Given rational A, b, g with $d_* = g - H(\rho(|A|)) > 0$, one can find a rational δ satisfying*

$$\frac{4}{9}\min\{1, d_*\} \le \delta \le \min\{1, d_*\}, \qquad H(\rho(|A|)) + \delta \le g.$$

The preprocessing time is polynomial in the input bit length and $\log(2 + 1/d_*)$. Combining this search with Theorem 1.1 gives total time polynomial in the input bit length and $1/\min\{1, d_*\}$.

Proof. For $j = 0, 1, \ldots$, put $\epsilon_j = 2^{-j}$. Lemma B.3 computes a rational q_j with $\kappa < q_j \le \kappa + \epsilon_j/16$. Compute a rational upper bound

$$H(q_j) \le U_j \le H(q_j) + \epsilon_j/16.$$

For $\mathfrak{B}$, use its closed form with interval arithmetic. For $\mathfrak{G}$, Lemma A.3 locates the maximum in the concavity interval $[z_{q_j}, c_{q_j}]$. On the fixed range $q_j \le K_0 + 1$ this interval stays uniformly below 1, and J_{q_j} and its derivative have uniformly bounded Lipschitz constants there. Certified concave one-dimensional maximization therefore gives U_j in time polynomial in the input length and j. For example, bisect on the derivative; if its sign is uncertain at precision ϵ, concavity bounds the remaining improvement by ϵ times the interval length. The interval endpoints are computed to the same accuracy. For $q_j \le 1$, simply use $U_j = 0$.

Stop at the first j with $U_j + \epsilon_j \le g$, and return $\delta = \epsilon_j$. Since both barriers are increasing and 1-Lipschitz,

$$H(\kappa) \le U_j \le H(\kappa) + \epsilon_j/8.$$

Thus the test certifies the margin and succeeds whenever $9\epsilon_j/8 \le d_*$, after $O(1+\log(1/\min\{1, d_*\}))$ stages. If $j = 0$, then $\delta = 1 \le d_*$. Otherwise the preceding failure gives $d_* < 9\epsilon_{j-1}/8 = 9\delta/4$. This proves the asserted bounds and complexity. □

The promise is essential: this procedure need not terminate when $g \le H(\kappa)$. It neither decides which side of the frontier an arbitrary input lies on nor removes the solver's inverse-margin cost.

C The Reduction Near the Critical Point

We turn from constructing an approximate solution to encoding a hard constraint problem in such solutions. This section gives the explicit near-critical residual construction; Appendix D extends it to KL and larger gains.

Parameter Budget and Proof Dependencies

Table 2 collects the parameters of the full reduction, including its extension away from one in Appendix D. The construction first reserves the perturbation budget, then chooses the scalar thresholds, builds the gates and sufficiently long chains, and finally adjusts the spectral radius. In this order, no parameter depends on a solution of the target fixed-point problem.

Table 2: Parameters and certificates for Theorem 1.2. Here $\mathfrak{H} = \mathfrak{B}$ for residual loss and $\mathfrak{H} = \mathfrak{G}$ for KL; N is the source size. Caps K_0 and, for residuals, $\bar{g} < 2/3$ are fixed.

Parameter	Choice or requirement	Role and proof location
τ	$\Delta/16$ for KL; $\min\{\Delta/16, (2/3-\bar{g})/4\}$ for residuals	Reserves final field perturbation; Appendix E.
a, K_c	$a = K - 4\tau$, $K_c = K - 2\tau$	Local feedback and core spectral cap, with $1 < a < K_c < K$.
g_c, γ	$g_c = g + 2\tau$ for KL, $g_c = g + \tau$ for residuals; $\gamma = \Delta/2 \le \mathfrak{H}(a) - g_c$	Core tolerance and scalar separation; equation (19).
P, Q	$\ell(P, \tanh(aP)) > g_c$; near one, $Q = 5P/2$	Sign preservation and capture; Lemmas C.4, D.1.
$h, \alpha = P_{\mathrm{amp}}$	$h > \mathfrak{T}_\ell(a, g_c; P)$; away from one also $h > \mathfrak{T}_\ell(a, g_c; \alpha)$, $\alpha > 1/3$	Forces latches; $\alpha = 1/2$ for KL, $\alpha + g_c < 1$ for residuals; Lemma D.4.
σ, v	$\sigma = K_c - a$; $v = \mathbf{1} + (M_0 - aI)\mathbf{1}/\sigma$ on latches	$K_c v - M_0 v = \sigma \mathbf{1}$ before relay inputs; Theorem D.5.
L, β	$\beta = a/K_c$; $L \ge L_{\mathrm{cap}}$ and $D v_{\max} \beta^L \le \sigma$	Capture and global spectral control; equation (71).
$\widetilde{v}$	$\lvert\widetilde{A}\rvert\widetilde{v} = K\widetilde{v}$, $\widetilde{v} > 0$	Exact gain and irreducibility, field change $\le \tau$, degree six; Lemma E.1.

All scalar parameters are computed by certified elementary-function evaluation with strict margins. The terminal values of any final solution decode a source solution: the perturbation bound first gives core loss at most g_c, the gate fields then force the required latch signs, and relay preservation carries them to the terminals. The two length constraints on L hold simultaneously, and extending a chain preserves both its signs and the degree bound. For fixed caps,

$$L = O((1 + \Delta^{-1}) \log(2 + \Delta^{-1})), \qquad n_{\mathrm{final}} = O(N(1 + \Delta^{-1}) \log(2 + \Delta^{-1})).$$

The rational weights $v_j \beta^t$ have bit length at most the bit length of v_j plus t times that of β, up to a constant encoding overhead. Thus the chain length, rather than an exponentially long real-number description, accounts for the dependence on $1/\Delta$. The final spectral adjustment uses rational sums, products, quotients, and a maximum, preserving polynomial encoding length. Appendix E proves the full size bound.

C.1 The Source Problem

The scalar relays will carry values of a Boolean constraint problem. We use the following established hardness result; its connectivity restrictions are what allow our final game to have bounded degree.

Theorem C.1 (Bounded-Degree Pure-Circuit)**.** *In the definition of Pure-Circuit, every node is the output of exactly one gate. With the gate set {PURIFY, NOT, NOR}, the problem is PPAD-complete even under the additional restrictions that every node is the input of exactly one gate and that the interaction graph has total degree at most three (Deligkas et al., 2024, Corollary 2.3).*

We encode Boolean value 1 by a positive signal, Boolean value 0 by a negative signal, and all intermediate signals by $\perp$. The robust NOR semantics are: two 0 inputs force output 1, while any input equal to 1 forces output 0 regardless of the other input. PURIFY copies a pure input to both outputs and, on an unspecified input, forces at least one output to be pure.

C.2 Uniform Near-Critical Relay Geometry

The source constraints require both preservation of pure signals and control of unspecified ones. We establish a region that every noisy relay enters and cannot leave; this is the invariant capture interval. Let

$$R_a(z) := \tanh(az) - z.$$

For $a > 1$, R_a is strictly concave on $(0, \infty)$, first increasing and then decreasing. Throughout the finite-parameter hard-side construction, fix

$$\eta_\star := 10^{-2}.$$

First we locate the positive scalar fixed point and an outer threshold Q beyond it. The negative value of $R_a(Q)$ will make the response point inward even after a bounded perturbation.

Lemma C.2. *Let $a = 1 + \eta$ with $\eta > 0$. Then*

$$\frac{3\eta}{1+3\eta} \le s_a^2 \le 3\eta. \tag{44}$$

If $0 < \eta \le 10^{-2}$, $(24/25)\sqrt{\eta} \le P \le \sqrt{\eta}$, and $Q = 5P/2$, then

$$\frac{3}{2}\sqrt{\eta} < s_a < 2\sqrt{\eta} < Q < 1, \qquad R_a(Q) < -2\eta^{3/2}. \tag{45}$$

Proof. The fixed-point equation gives

$$\eta = \frac{1}{s_a}\int_0^{s_a} \frac{t^2}{1-t^2}\,dt, \qquad \frac{s_a^2}{3} \le \eta \le \frac{s_a^2}{3(1-s_a^2)},$$

which proves (44) and the stated root bounds. Write $Q = u\sqrt{\eta}$, where $12/5 \le u \le 5/2$. For $v = Q - 2\eta^{3/2} \in (0,1)$, the same integral gives $\operatorname{artanh} v \ge v + v^3/3$. Hence

$$\eta^{-3/2}\big(\operatorname{artanh} v - aQ\big) \ge \frac{(u-2\eta)^3}{3} - u - 2 \ge \frac{(119/50)^3}{3} - \frac{22}{5} > 0.$$

The middle expression increases with u and decreases with η on the stated range. Thus $\tanh(aQ) < v$, proving the residual bound. □

An inward response at the boundary gives invariance but not a bound on the time needed to reach it. The next lemma supplies that bound uniformly over all permitted relay errors.

Lemma C.3. *Let $a = 1 + \eta > 1$, $(12/5)\sqrt{\eta} \le Q \le 1$, and $0 \le e \le (11/10)\eta^{3/2}$. Every sequence in $[-1,1]$ satisfying $|x_{t+1} - \tanh(ax_t)| \le e$ enters the invariant interval $[-Q, Q]$ within*

$$L_0 = \left\lceil \frac{12}{Q^2} \right\rceil + 1 = O(\eta^{-1}) \tag{46}$$

steps. When $Q \le (5/2)\sqrt{\eta}$, this order is necessary as $\eta \downarrow 0$, even without noise.

Proof. For $Q \le x \le 1$, we have $\eta \le 25x^2/144$ and $e \le (275/3456)x^3 < x^3/12$. Since $a^{-2} \ge 1 - 2\eta$, (40) gives

$$\tanh(ax) \le \frac{x}{\sqrt{a^{-2} + 2x^2/3}} \le \frac{x}{\sqrt{1 + 23x^2/72}} \le x - x^3/8.$$

For the last step, the convex function $(1 + 23s/72)^{-1/2}$ lies below its chord on $[0,1]$, whose endpoint at one is $\sqrt{72/95} < 7/8$. Thus every next value satisfies $y \le \tanh(ax) + e \le x - x^3/24$. At $x = Q$ this proves invariance by monotonicity and oddness. Also $e < Q$, so a step from above Q cannot fall below $-Q$. While $y > Q$,

$$y^{-2} \ge x^{-2}(1 - x^2/24)^{-2} \ge x^{-2} + 1/12.$$

Summing proves the capture bound; symmetry handles negative values. For the lower bound, take $\eta \le 1/9$ and start without noise at $x_0 = 3\sqrt{\eta} \le 1$. Before capture, $Q < x_t \le 3\sqrt{\eta}$ and

$$x_t - \tanh(ax_t) \le x_t - \tanh x_t \le x_t^3/3 \le 9\eta^{3/2}.$$

Traversing $x_0 - Q \ge \sqrt{\eta}/2$ takes at least $1/(18\eta)$ steps. □

We can now choose rational thresholds with enough room for approximation error in their computation. The next result combines preservation and capture into the signal guarantees consumed by the gate construction.

Lemma C.4. *Let $1 < a \le 1 + \eta_\star$, $\varepsilon \ge 0$, and suppose*

$$\mathfrak{B}(a) - \varepsilon \ge \gamma > 0. \tag{47}$$

In time polynomial in the bit lengths of a, ε, γ and $\log(1/\gamma)$, one can find rational numbers P, Q such that

$$0 < P < Q < 1, \qquad Q \le \frac{5}{2}P, \tag{48}$$

$$R_a(P) \ge \varepsilon + \gamma/2, \tag{49}$$

$$R_a(Q) \le -\varepsilon - \gamma/2. \tag{50}$$

They have the following consequences for every noisy relay transition

$$|y - \tanh(ax)| \le \varepsilon. \tag{51}$$

(i) $[P, Q]$ and $[-Q, -P]$ are invariant.

(ii) $[-Q, Q]$ is invariant.

(iii) From every $x_0 \in [-1, 1]$, every sequence satisfying (51) enters $[-Q, Q]$ after at most $L_0 := \lceil 12/Q^2 \rceil + 1 = O((a-1)^{-1})$ steps.

(iv) If $x_0 \ge P$ (respectively, $x_0 \le -P$), the sequence enters $[P, Q]$ (respectively, $[-Q, -P]$) within L_0 steps and never loses its sign encoding.

Moreover, if $\eta = a - 1$, then

$$P \ge \frac{9}{10}\sqrt{\eta}, \qquad \gamma \le \eta^{3/2}. \tag{52}$$

Proof. Write $\eta = a - 1$ and let $z_a = \operatorname{arsinh} \sqrt{\eta}/a$ be the explicit maximizer of R_a. The integral bounds for arsinh give

$$\frac{49}{50}\sqrt{\eta} \le \frac{\sqrt{\eta}}{(1+\eta)^{3/2}} \le z_a \le \sqrt{\eta}.$$

Choose a rational lower approximation P satisfying

$$\frac{99}{100} z_a \le P \le z_a, \qquad z_a - P \le \frac{\gamma}{4(a+1)}, \qquad Q = \frac{5}{2}P.$$

The elementary-function evaluator constructs P with polynomially many bits. The $(a+1)$-Lipschitz bound on R_a gives $R_a(P) \ge \mathfrak{B}(a) - \gamma/4 \ge \varepsilon + 3\gamma/4$. Also $P \ge (24/25)\sqrt{\eta}$ and $Q \le (5/2)\sqrt{\eta} < 1$. Lemma C.2 gives $R_a(Q) < -2\eta^{3/2}$.

Since $\varepsilon + \gamma \le \mathfrak{B}(a) \le (2/3)\eta^{3/2}$, these estimates prove (49)–(50) and (52). Theorem 2.2 preserves $[P, 1]$ and $[-1, -P]$. Lemma C.3 with $e = \varepsilon$ gives invariance and capture for $[-Q, Q]$. Intersecting the intervals proves all four claims. □

C.3 A Direct Forcing Estimate

The relays preserve a sign once it has been established. We now bound the external field needed to establish that sign in a latch, using the forcing threshold from Lemma 4.1.

Lemma C.5. *Let $a > 1$, $P > 0$, and $e \ge 0$. Suppose $P + e \le s_a$ and $\ell(z, \tanh h) \le g$ implies $|z - \tanh h| \le e$. Then*

$$\mathfrak{T}_\ell(a, g; P) \le a\big(\mathfrak{B}(a) + e\big). \tag{53}$$

In particular, for residual tolerance $\varepsilon \ge 0$ and $R_a(P) > \varepsilon$, any latch satisfying

$$|x - \tanh(ay + H)| \le \varepsilon, \qquad |y - \tanh(ax + H)| \le \varepsilon \tag{54}$$

has $x, y \ge P$ when $H > a(\mathfrak{B}(a) + \varepsilon)$, and $x, y \le -P$ when $H < -a(\mathfrak{B}(a) + \varepsilon)$.

Proof. The definition of $\mathfrak{B}$ gives the useful bound

$$\operatorname{artanh} u - au \le a\mathfrak{B}(a), \qquad -1 < u \le s_a. \tag{55}$$

For $u \ge 0$ the left side is nonpositive. For $u < 0$, write $u = -\tanh t$; the left side is $a \tanh t - t = aR_a(t/a) \le a\mathfrak{B}(a)$. For any pair in (20), put $u = \tanh h < P + e \le s_a$. Then

$$h - az = (\operatorname{artanh} u - au) + a(u - z) \le a\mathfrak{B}(a) + ae.$$

Taking the supremum proves (53). For residuals, $R_a(P) > \varepsilon$ implies $P < s_a$ and $P + \varepsilon < \tanh(aP) \le s_a$. The latch claims follow from Lemma 4.1 and symmetry. □

C.4 Robust Gates

The previous estimate turns a sufficiently large field into a definite output sign. We choose affine fields that meet that estimate for every input configuration required by NOT, NOR, and PURIFY. Choose a rational upper approximation

$$\mathfrak{B}(a) \le \overline{B} \le \mathfrak{B}(a) + \gamma/16$$

and put

$$h := 3a\overline{B}.$$

Since $\varepsilon < \mathfrak{B}(a)$, $h > a(\mathfrak{B}(a) + \varepsilon)$. Given the rational P, Q from Lemma C.4, define

$$d_{\mathrm{N}} := \frac{2h}{P}, \tag{56}$$

$$d_{\mathrm{P}} := \frac{4h}{P}, \qquad \theta_{\mathrm{P}} := 2h, \tag{57}$$

$$d_{\mathrm{G}} := \frac{8h}{P}, \qquad \theta_{\mathrm{G}} := \frac{d_{\mathrm{G}}(P+Q)}{2}. \tag{58}$$

Here d_{G} is the NOR coefficient denoted c_{N} in Section 4.3. The following checks show that unspecified inputs cannot invalidate the required outputs.

Lemma C.6. *Suppose every gate input has been passed through a relay chain long enough that it belongs to $[-Q, Q]$, and call it pure positive if it lies in $[P, Q]$, pure negative if it lies in $[-Q, -P]$, and unspecified otherwise. Use a two-coordinate latch satisfying* (54) *for each gate output.*

(i) The field $H_{\mathrm{NOT}} = -d_{\mathrm{N}} r$ negates every pure input.

(ii) The field

$$H_{\mathrm{NOR}} = -d_{\mathrm{G}}(r_1 + r_2) - \theta_{\mathrm{G}}$$

forces a positive output when both inputs are pure negative, and forces a negative output when either input is pure positive, regardless of the other input.

(iii) The two fields

$$H^{\pm}_{\mathrm{PURIFY}} = d_{\mathrm{P}} r \pm \theta_{\mathrm{P}}$$

copy a pure input to both outputs; for every unspecified input, at least one field has magnitude at least $2h$ and hence at least one output is pure.

Proof. For NOT, a pure positive input gives $H_{\mathrm{NOT}} \le -2h$, and a pure negative input gives $H_{\mathrm{NOT}} \ge 2h$.

For NOR, if $r_1, r_2 \le -P$, then

$$H_{\mathrm{NOR}} \ge 2d_{\mathrm{G}}P - \theta_{\mathrm{G}} = \frac{d_{\mathrm{G}}}{2}(3P - Q) \ge 2h,$$

where $Q \le 5P/2$ was used. If, say, $r_1 \ge P$ and $r_2 \ge -Q$, then

$$H_{\mathrm{NOR}} \le d_{\mathrm{G}}(Q - P) - \theta_{\mathrm{G}} = -\frac{d_{\mathrm{G}}}{2}(3P - Q) \le -2h.$$

The other input position is symmetric.

For PURIFY, a pure positive input satisfies $d_{\mathrm{P}} r - \theta_{\mathrm{P}} \ge 2h$, and a pure negative input satisfies $d_{\mathrm{P}} r + \theta_{\mathrm{P}} \le -2h$. Finally,

$$\max\{|d_{\mathrm{P}} r + \theta_{\mathrm{P}}|, |d_{\mathrm{P}} r - \theta_{\mathrm{P}}|\} \ge \theta_{\mathrm{P}} = 2h$$

for every real r. Lemma C.5 converts all field inequalities into the stated output signs. □

C.5 Construction and Correctness

All local components are now available. The remaining task is to connect them so that every approximate profile decodes correctly while the full matrix, not just each component, has controlled gain. In this preliminary construction, K is an upper bound on gain; it plays the role of K_c when the final gain is adjusted in Appendix E.

Theorem C.7 (Reduction with a Spectral Bound)**.** *Let $a, K, \varepsilon, \gamma$ be rational parameters satisfying*

$$1 < a < K \le 1 + \eta_\star, \qquad \mathfrak{B}(a) - \varepsilon \ge \gamma > 0. \tag{59}$$

Assume γ and $K - a$ are inverse-polynomial in the size N of a given Pure-Circuit instance. In polynomial time one can construct a rational affine logit map

$$F(x) = \tanh(Ax + b)$$

with the following properties.

- *(i) From every x with $\mathrm{Res}_F(x) \le \varepsilon$, one can decode a solution of the Pure-Circuit instance in polynomial time.*
- *(ii) The interaction graph has maximum degree at most five.*
- *(iii) With $M = |A|$,*
$$a \le \rho(M) \le K. \tag{60}$$
- *(iv) The number of coordinates and the encoding length are polynomial in N, $1/\gamma$, and $1/(K - a)$.*
- *(v) Every row satisfies $|b_i| + \sum_j |A_{ij}| < 2$.*

Proof. Use the bounded-degree Pure-Circuit instance from Theorem C.1. Represent every circuit node v by a two-coordinate latch (x_v, y_v) with internal equations

$$x_v \simeq_\varepsilon \tanh(a y_v + H_v), \qquad y_v \simeq_\varepsilon \tanh(a x_v + H_v), \tag{61}$$

where $u \simeq_\varepsilon w$ means $|u - w| \le \varepsilon$. Attach to x_v a relay chain

$$r_{v,1} \simeq_\varepsilon \tanh(a x_v), \qquad r_{v,t} \simeq_\varepsilon \tanh(a r_{v,t-1}) \quad (2 \le t \le L). \tag{62}$$

The terminal signal $r_{v,L}$ is used by the unique gate for which v is an input. The field H_v is one of the NOT, NOR, or PURIFY fields from Lemma C.6, according to the unique gate that outputs v.

Let

$$D := \max\{d_{\mathrm{N}}, d_{\mathrm{P}}, 2d_{\mathrm{G}}\}.$$

Choose

$$L \ge \max\left\{L_0, \left\lceil \frac{\max\{0, \log(D/(K-a))\}}{\log(K/a)} \right\rceil \right\}, \tag{63}$$

where L_0 is from Lemma C.4.

Now consider any ε-residual fixed point. By Lemma C.4, every terminal signal lies in $[-Q, Q]$. Decode node v as 1 if $r_{v,L} \ge P$, as 0 if $r_{v,L} \le -P$, and as $\perp$ otherwise. If a gate field forces its output latch positive or negative, Lemma C.5 makes both latch coordinates have that sign and magnitude at least P; Lemma C.4(iv) then preserves the sign through the outgoing chain. Lemma C.6 therefore proves exactly the robust NOT, NOR, and PURIFY semantics. The decoded assignment is a Pure-Circuit solution, proving (i).

For the spectral bound, assign the positive weights

$$v_{x_v} = v_{y_v} = 1, \qquad v_{r_{v,t}} = \left(\frac{a}{K}\right)^t.$$

Every relay row satisfies

$$(Mv)_{r_{v,t}} = a v_{r_{v,t-1}} = K v_{r_{v,t}}.$$

Each latch row has internal weighted load a and incoming gate load at most $D(a/K)^L$. By (63),

$$(Mv)_i \le a + D(a/K)^L \le K = Kv_i.$$

Thus $Mv \le Kv$, and Collatz–Wielandt gives $\rho(M) \le K$. Every latch contains the principal submatrix $\left(\begin{smallmatrix} 0 & a \\ a & 0 \end{smallmatrix}\right)$, so $\rho(M) \ge a$. This proves (iii).

The degree bound follows directly from the construction. A latch coordinate has its partner, at most two gate-input neighbours, and possibly the first relay neighbour. A relay interior node has degree two. Because every Pure-Circuit node is the input of exactly one gate, a terminal relay has one predecessor and at most four output-latch neighbours (the PURIFY case). Hence the maximum degree is at most five.

It remains to verify size and coefficient bounds. Put $\eta = a - 1$. From Lemma C.4, $P \ge (9/10)\sqrt{\eta}$ and $\gamma \le \eta^{3/2}$. Since $\overline{B} \le \mathfrak{B}(a) + \gamma/16 \le (17/16)\mathfrak{B}(a)$, Lemma A.1 gives

$$h \le 4a\mathfrak{B}(a) \le 4a\eta^{3/2}, \qquad \frac{h}{P} \le \frac{40}{9}a\eta < 5a\eta.$$

Consequently $d_{\mathrm{G}} \le 40a\eta$ and $D \le 80a\eta$. Moreover, $\theta_{\mathrm{G}} \le (7/4)d_{\mathrm{G}}P = 14h$ and $\theta_{\mathrm{P}} = 2h$. The worst row is a NOR-output latch row, and the displayed bounds give

$$|b_i| + \sum_j |A_{ij}| \le a + 2d_{\mathrm{G}} + \theta_{\mathrm{G}} < 1.01 + 0.808 + 0.057 < 2.$$

All other row types have a smaller bound. Thus every row sum is strictly less than 2. Also $L_0 = O(\eta^{-1}) = O(\gamma^{-2/3})$, while

$$\log(K/a) \ge \frac{K-a}{K},$$

so (63) is polynomial in $1/\gamma$ and $1/(K-a)$.

The constants $\mathfrak{B}(a)$, P, and Q are used only to choose rational parameters. The closed form (25) and interval evaluation of elementary functions find all required rationals to $O(\log(1/\gamma))$ bits beyond the input precision. Since $P \ge \gamma^{1/3}/4$, divisions by P have polynomial encoding length. No transcendental number is written into the output instance. This proves (iv) and (v). □

D The Reduction on Bounded Gain Intervals

This section upgrades the explicit near-critical reduction to direct KL loss and then treats every compact gain interval away from one. The two regimes need different components: near criticality the signal and forcing scales shrink together; away from one, both stay bounded and an extra amplifier handles arbitrary unspecified inputs.

D.1 Scalar Parameters

The near-critical construction needs a sign threshold and a bounded forcing field; away from one it also needs an amplification threshold. We first supply these parameters directly for KL, so that the reduction reaches $\mathfrak{G}(a)$ rather than only the Pinsker transfer.

Lemma D.1. *Fix $K_0 > 1$ and set $\eta_0 = 10^{-2}$. There are constants $c_0, C_0 > 0$, depending only on K_0, with the following property. Given rational $1 < a \le K_0$, $g \ge 0$, and $\gamma > 0$ with $\mathfrak{G}(a) - g \ge \gamma$, one can compute rational parameters in time polynomial in the input length and $\log(2 + 1/\gamma)$ such that:*

(i) a threshold $P \in (0, s_a)$ satisfies

$$J_a(P) := d(P, \tanh(aP)) \ge g + 7\gamma/8, \qquad P \ge c_0 \min\{\sqrt{a-1}, 1\}; \tag{64}$$

(ii) a rational $h_P > \mathfrak{T}_d(a, g; P)$ is computed explicitly;

(iii) if $\eta = a - 1 \le \eta_0$, take $Q = 5P/2 < 1$. Then $J_a(Q) > \mathfrak{G}(a)$. Under every transition $d(y, \tanh(ax)) \le g$, the intervals $[-Q, Q]$, $[P, Q]$, and $[-Q, -P]$ are invariant. Every sequence enters $[-Q, Q]$ within $L_0 = \lceil 12/Q^2 \rceil + 1 = O(\eta^{-1})$ steps. Moreover,

$$h_P \le a\mathfrak{B}(a) + a\sqrt{2g} + \gamma/64 \le C_0\eta^{3/2}; \tag{65}$$

(iv) if $a - 1 \ge \eta_0/2$, then $P \ge c_0$, and rational bounds $h_P > \mathfrak{T}_d(a, g; P)$ and $h_{1/2} > \mathfrak{T}_d(a, g; 1/2)$ are bounded by C_0.

In (iii), $L_0 = O(\gamma^{-1/3})$. All output encoding lengths are polynomial in the input length and $\log(2 + 1/\gamma)$.

Proof. Use z_a, c_a from Lemma A.3, fix rational $\widehat{K}_0 > K_0$, and set $\xi = \gamma/128$, $e_0 = \xi/\widehat{K}_0$. Compute rational endpoints inside the concavity interval:

$$z_a \le L \le z_a + e_0, \qquad c_a - e_0 \le U \le c_a.$$

They satisfy $L < U$ because $c_a - z_a = \mathfrak{B}(a)$ and $\gamma \le \mathfrak{G}(a) \le a\mathfrak{B}(a)^2/2 < a\mathfrak{B}(a)$. On $[0, s_a]$, $J_a' = aR_a - (az - \operatorname{artanh} z)$ is the difference of two numbers in $[0, a]$. Thus J_a is a-Lipschitz, and restricting its maximum to the initial interval $[L, U]$ loses at most ξ.

At each midpoint m, evaluate $J_a'(m)$ to error ξ. If the estimate exceeds 2ξ, replace L by m; if it is below -2ξ, replace U by m. Strict concavity preserves the maximum over the initial interval in the retained half. If neither case holds, $|J_a'(m)| \le 3\xi$; concavity bounds the remaining loss by 3ξ, since the interval has length at most one. Return $P = m$. Alternatively, once $U - L \le e_0$, return its midpoint; the Lipschitz bound makes the remaining loss at most ξ. In both cases,

$$J_a(P) \ge \mathfrak{G}(a) - 4\xi = \mathfrak{G}(a) - \gamma/32 \ge g + 31\gamma/32. \tag{66}$$

Throughout, $P \in [z_a, c_a] \subset (0, s_a)$. There are $O_{K_0}(\log(2 + 1/\gamma))$ derivative evaluations. All arguments lie in a fixed compact subinterval of $(-1, 1)$, so elementary-function evaluation uses polynomially many bits. The explicit bound $z_a \ge \sqrt{a - 1}/a^{3/2}$ also gives the required lower bound on P, including away from one.

For $\eta \le 10^{-2}$, the bounds on z_a used in Lemma C.4 give

$$\frac{24}{25}\sqrt{\eta} \le P \le \sqrt{\eta}, \quad \frac{3}{2}\sqrt{\eta} < s_a < 2\sqrt{\eta} < Q < 1. \tag{67}$$

Lemma C.2 gives $R_a(Q) < -2\eta^{3/2}$, whereas (30) gives, with $e = \sqrt{2g}$,

$$e \le \sqrt{2\mathfrak{G}(a)} \le \frac{2}{3}\eta^{3/2} < \frac{11}{10}\eta^{3/2}. \tag{68}$$

Thus $P + e < s_a$. Pinsker gives $J_a(Q) \ge R_a(Q)^2/2 > \mathfrak{G}(a)$ and $|y - \tanh(ax)| \le e$. Lemma C.3 now proves invariance and capture for $[-Q, Q]$ in $O(\eta^{-1})$ steps. Theorem 2.2 separately preserves $[P, 1]$ and $[-1, -P]$, so their intersections with $[-Q, Q]$ are invariant as well. Finally, (34) gives $\gamma \le \mathfrak{G}(a) = O(\eta^3)$, hence $L_0 = O(\gamma^{-1/3})$.

For near-critical forcing, apply Lemma C.5 using $P + e < s_a$. Approximate $a\mathfrak{B}(a) + ae$ from above within $\gamma/128$ and add $\gamma/128$. This yields (65). Away from one, the natural-coordinate identity gives, for $h \ge 0$, $d(x, \tanh h) \ge (1 - x)h - \log 2$. Therefore

$$\mathfrak{T}_d(a, g; R) \le a + \frac{g + \log 2}{1 - R} < h_R := a + \frac{g + 1}{1 - R}, \qquad 0 < R < 1. \tag{69}$$

For negative h, use $h - ax \le a$. Taking $R = P$ and $R = 1/2$ gives uniform rational bounds because $P < s_a \le s_{K_0} < 1$ and $g \le \mathfrak{G}(K_0)$. All formulas have polynomial encoding length. □

The KL parameters are now available on the full bounded gain interval. The analogous residual construction away from one needs an additional fixed tolerance cap. In this section α denotes the amplification threshold P_{amp} of Section 4.3, not the algorithm's clipping parameter.

Lemma D.2. *Fix $K_0 > 1$ and $\eta_0 > 0$. Given $1 + \eta_0 \le a \le K_0$, $\varepsilon \ge 0$, and $\mathfrak{B}(a) - \varepsilon \ge \gamma > 0$, one can compute a rational $P \in (0, s_a)$ and a rational $h_P > \mathfrak{T}_{|\cdot|}(a, \varepsilon; P)$ such that*

$$\tanh(aP) - P \ge \varepsilon + \frac{3\gamma}{4}.$$

Moreover P is bounded below and h_P is bounded above by constants depending only on K_0 and η_0.

If also $\varepsilon \le \bar{\varepsilon}$ for a fixed $0 < \bar{\varepsilon} < 2/3$ (rational), fix any rational

$$\frac{1}{3} < \alpha < 1 - \bar{\varepsilon}. \tag{70}$$

Then a rational $h_\alpha > \mathfrak{T}_{|\cdot|}(a, \varepsilon; \alpha)$ is also bounded by a constant depending only on $K_0, \eta_0, \bar{\varepsilon}$, and α.

Proof. The explicit maximizer $z_a = \operatorname{arsinh} \sqrt{a-1}/a$ is bounded away from zero on the compact gain interval. Choose a rational P close enough to z_a that $R_a(P) \ge \mathfrak{B}(a) - \gamma/4$. Lemma C.5 then gives the uniform bound $\mathfrak{T}_{|\cdot|}(a, \varepsilon; P) \le a(\mathfrak{B}(a) + \varepsilon)$. A rational upper approximation plus a positive constant supplies h_P.

For the amplifier, $\alpha + \varepsilon \le \alpha + \bar{\varepsilon} < 1$. Every feasible pair with $x < \alpha$ satisfies $\tanh h \le x + \varepsilon < \alpha + \varepsilon$, hence

$$\mathfrak{T}_{|\cdot|}(a, \varepsilon; \alpha) \le a + \operatorname{artanh}(\alpha + \varepsilon) < a + \frac{1}{1 - \alpha - \varepsilon}.$$

The last expression is a rational choice of h_α, uniformly bounded on the required gain interval. □

D.2 Gates Away from One

At larger gain we replace capture by amplification of pure NOR inputs. We first explain the required amplitude and residual cap, then show that the resulting fields implement the constraints.

Lemma D.3. *Let $0 < \alpha < 1$. Suppose an affine field $H = -d(u_1 + u_2) - \theta$, $d > 0$, must be at least $h > 0$ when $u_1, u_2 \le -\alpha$, and at most $-h$ whenever one input is at least α and the other may be any point of $[-1, 1]$. Then $\alpha > 1/3$. Under residual tolerance ε, no finite field can force a latch coordinate to be at least α unless $\varepsilon < 1 - \alpha$. Consequently this entire one-coordinate affine amplifier construction requires $\varepsilon < 2/3$.*

Proof. The input pairs $(-\alpha, -\alpha)$ and $(\alpha, -1)$ require

$$2d\alpha - \theta \ge h, \qquad d(1 - \alpha) - \theta \le -h.$$

Subtracting gives $d(3\alpha - 1) \ge 2h$, so $\alpha > 1/3$. If $\varepsilon \ge 1 - \alpha$, the continuous map $T(z) = \max\{-1, \tanh(az + H) - \varepsilon\}$ sends $[-1, \alpha]$ into $[-1, \alpha)$. It has a fixed point $z < \alpha$, and (z, z) is an admissible latch for any finite H. Thus no finite forcing field works. □

The preceding obstruction identifies the range in which this affine NOR construction can work. Within that range, an extra latch amplifies each pure input before the NOR field combines them.

Lemma D.4. *Let source signals encode pure signs by $[P, 1]$ and $[-1, -P]$. Let $\alpha > 1/3$, and choose*

$$h > \max\{\mathfrak{T}_\ell(a, g; P), \mathfrak{T}_\ell(a, g; \alpha)\}.$$

A private amplifier latch with field $2hr/P$ maps a pure source r to the same sign with magnitude at least α. NOT and PURIFY are realized by

$$H_{\mathrm{NOT}} = -\frac{2h}{P} r, \qquad H^{\pm}_{\mathrm{PURIFY}} = \frac{4h}{P} r \pm 2h.$$

For NOR, amplify its two inputs to u_1, u_2 and put

$$d_\alpha := \frac{2h}{3\alpha - 1}, \qquad \theta_\alpha := \frac{d_\alpha(1 + \alpha)}{2}, \qquad H_{\mathrm{NOR}} := -d_\alpha(u_1 + u_2) - \theta_\alpha.$$

These fields satisfy the robust NOT, NOR, and PURIFY semantics under ℓ-loss at most g.

Proof. A pure source gives the amplifier a field of magnitude at least $2h$, so Lemma 4.1 forces the claimed sign and magnitude. The NOT and PURIFY checks are identical to Lemma C.6; for every r one of the two PURIFY fields has magnitude at least $2h$. If both amplified NOR inputs are negative, then

$$H_{\mathrm{NOR}} \geq 2d_\alpha\alpha - \theta_\alpha = h.$$

If, say, $u_1 \geq \alpha$ while $u_2 \geq -1$, then

$$H_{\mathrm{NOR}} \leq d_\alpha(1-\alpha) - \theta_\alpha = -h.$$

Lemma 4.1 completes the gate semantics. □

D.3 Assembling the Reduction

We now combine the near-critical and amplified constructions. The same weighted-matrix argument controls their global gain, which lets us state a uniform result on every fixed bounded gain interval.

Theorem D.5 (Reduction on Bounded Gain Intervals)**.** *Fix $K_0 > 1$ and let $1 < a < K_c \leq K_0$. Suppose the margins $K_c - a$ and γ below are at least inverse polynomial in the source size.*

(a) If $g \geq 0$ and $\mathfrak{G}(a) - g \geq \gamma > 0$, one can construct from a bounded-degree Pure-Circuit instance a rational affine logit map such that every $\mathcal{L}_d(x; F) \leq g$ solution decodes the circuit and, for a positive rational vector v,

$$|A|v \leq K_c v.$$

The maximum degree is five, and the row sum is bounded by a constant depending only on K_0.

(b) Fix a rational $0 < \bar{\varepsilon} < 2/3$. If $0 \leq \varepsilon \leq \bar{\varepsilon}$ and $\mathfrak{B}(a) - \varepsilon \geq \gamma > 0$, the same conclusion holds for residual loss. The row bound may also depend on $\bar{\varepsilon}$.

The size and encoding length are polynomial in the source size, $1/\gamma$, and $1/(K_c - a)$.

Proof. For KL with $a - 1 \leq \eta_0$, use Lemma D.1(iii). Represent every circuit node by a gain-a two-coordinate latch and attach a relay chain. The intervals $[P, Q]$ and $[-Q, -P]$, together with $Q \leq 5P/2$, support the fields in Lemma C.6, with its residual forcing constant replaced by a rational $h = 2h_P$. Lemma 4.1 gives the gate semantics, and $h/P = O(a-1)$ keeps all coefficients uniformly bounded. Lemma C.3 gives a polynomial capture length.

For KL with $a - 1 \geq \eta_0/2$, take $\alpha = 1/2$ and apply Lemma D.4. No capture interval is needed: a pure relay remains pure by Theorem 2.2, while an arbitrary other input is absorbed by the amplifier-NOR guarantee. All coefficients are bounded on the fixed compact gain interval.

For residual loss, when $a \leq 1 + \eta_\star$, use the relay intervals and gates of Lemmas C.4–C.6; their coefficients and capture length have the bounds proved in Theorem C.7. When $a \geq 1 + \eta_\star/2$, use Lemma D.2, choose once and for all a rational α satisfying (70), and invoke Lemma D.4. The fixed tolerance cap supplies $\alpha + \varepsilon < 1$, with no further restriction on K_0.

The degree is at most five in both constructions. Each source has one outgoing chain and is used by only one gate. A terminal relay meets its predecessor and at most four output-latch coordinates (PURIFY). A latch coordinate has its partner, at most two inputs, and at most one outgoing relay; a private amplifier has its partner, one source, and two output neighbors. Interior relays have degree two.

For the spectral bound, let M_0 be the absolute matrix after deleting relay coordinates and put $w = (M_0 - aI)\mathbf{1} \geq 0$. Every pair has its internal gain-a coupling. The only remaining inputs between pairs run from private amplifiers to output latches. Thus amplifier rows have sum a and $w = 0$ there; the two output rows receive the same field, so w agrees within each pair. Multiplication by M_0 swaps paired entries with factor a, while all extra inputs come from zero entries of w. Hence $M_0 w = aw$. For $\sigma = K_c - a$, the explicit choice

$$v = \mathbf{1} + w/\sigma, \qquad M_0 v = av + w = K_c v - \sigma\mathbf{1}$$

gives the same certificate in every regime; $w = 0$ for direct gates.

Let D bound the row sums of absolute coefficients from relay endpoints into latch coordinates, and put $\beta = a/K_c$. Use a common chain length

$$L = \max\left\{1, L_{\mathrm{cap}}, \left\lceil \frac{\log \max\{1, Dv_{\max}/\sigma\}}{\log(K_c/a)} \right\rceil \right\}, \tag{71}$$

where $L_{\mathrm{cap}} = L_0$ in the near-critical construction and 1 otherwise. Each relay at depth t from coordinate j has weight $v_j\beta^t$ and weighted load exactly K_c times that weight. Each latch row has load at most $(M_0 v)_i + Dv_{\max}\beta^L \le K_c v_i$. Thus the extended vector certifies $|A|v \le K_c v$.

All local coefficients are uniformly bounded, giving $v_{\max} = 1 + O(\sigma^{-1})$. Together with $\log(K_c/a) \ge \sigma/K_c$, this yields

$$L = O\big(L_{\mathrm{cap}} + \sigma^{-1}\log(2+\sigma^{-1})\big).$$

Here and below these constants depend on K_0, and for residual loss also on the fixed cap $\bar{\varepsilon}$. Repeated rational multiplication finds a length satisfying $Dv_{\max}\beta^L \le \sigma$ within this bound. The formula for v and each $v_j\beta^t$ has polynomial encoding length. The capture length is polynomial in $1/\gamma$, while extending chains preserves the degree, row-sum bounds, and decoded signs. A latch principal block gives $\rho(|A|) \ge a$. □

E An Exact Spectral Radius and Irreducibility

The preceding construction only bounds the spectral radius from above. We now make it equal to the requested value and make the influence graph strongly connected, while perturbing each old field within the reserved error budget.

Lemma E.1 (Exact Spectral Adjustment). *Let $M = |A|$ be rational with $n \ge 1$ rows and $Mv \le (K - 2\tau)v$ for rational $v > 0$ and $K, \tau > 0$. For any integer $k \ge 1$, one can add $m = \max\{2, \lceil n/k \rceil\}$ coordinates in polynomial time such that:*

(i) $|\widetilde{A}|\widetilde{v} = K\widetilde{v}$ with $\widetilde{v} > 0$ rational, and $|\widetilde{A}|$ is irreducible;

(ii) every original field changes by at most τ;

(iii) row sums are at most $\max\{R + \tau, K + \tau\}$, where $R = \max_i \sum_j M_{ij}$;

(iv) zero diagonal is preserved; each old coordinate gains one neighbor, and each new coordinate has at most $k + 2$ neighbors.

All coefficients have polynomial encoding length. In particular, $k = 1$ retains new-coordinate degree three, whereas $k = 4$ uses $n + \max\{2, \lceil n/4 \rceil\}$ coordinates in total.

Proof. Set $r = Kv - Mv$, $W = \max_i r_i/\tau$, and partition the old indices into m groups I_j of size at most k, allowing empty groups. Add one coordinate f_j of weight W per group. For $i \in I_j$, give $f_j \to i$ coefficient $d_i = r_i/W$ and $i \to f_j$ coefficient τ/k. Join the f_j in a directed cycle with incoming coefficient

$$c_j = K - \frac{\tau}{kW}\sum_{i \in I_j} v_i.$$

The promise gives $K \ge 2\tau$ and $W \ge 2\max_i v_i$, so $0 < d_i \le \tau$ and $0 < K - \tau/2 \le c_j \le K$. The positive eigenvector is verified by

$$(Mv)_i + d_i W = Kv_i, \qquad \frac{\tau}{k}\sum_{i \in I_j} v_i + c_j W = KW.$$

The cycle and two-way group connections give irreducibility, including when some groups are empty. Keep old biases and assign new biases zero. An old field changes by at most $d_i \le \tau$; a new row sums to at most $k(\tau/k) + c_j \le K + \tau$. Since $m \ge 2$, no cycle edge is a self-loop. The degree and encoding bounds follow directly from the construction. □

The adjustment lemma preserves local behavior up to a known field error. We now account for that error in the two losses and combine the construction, decoding, size bound, and game normalization to prove the main lower bound.

Proof of Theorem 1.2. For KL use $\tau = \Delta/16$. For residual loss write $\varepsilon = g$, set $\tau = \min\{\Delta/16, (2/3 - \bar{g})/4\}$, and define the fixed core cap $\bar{\varepsilon} = \bar{g} + (2/3 - \bar{g})/4 < 2/3$. In either case set

$$a = K - 4\tau, \qquad K_c = K - 2\tau.$$

Because $\Delta \le \mathfrak{H}(K) \le K - 1$, we have $a > 1$. Since $\mathfrak{H} \in \{\mathfrak{B}, \mathfrak{G}\}$ is increasing and 1-Lipschitz,

$$\mathfrak{H}(a) \ge \mathfrak{H}(K) - 4\tau.$$

For KL, build the preliminary core of Theorem D.5 at tolerance $g + 2\tau$; for residual use $\varepsilon + \tau \le \bar{\varepsilon}$. The remaining scalar gap is at least $\Delta - 6\tau \ge \Delta/2$, and $K_c - a = 2\tau$. Thus Theorem D.5 applies, including its fixed residual-cap requirement. Apply Lemma E.1. If an augmented point has KL loss at most g, then on every original coordinate

$$d(x_i, \tanh h_i) \le d(x_i, \tanh(h_i + e_i)) + 2|e_i| \le g + 2\tau,$$

because $|\partial_h d(x, \tanh h)| = |\tanh h - x| \le 2$. For residual loss, 1-Lipschitzness of $\tanh$ gives original residual at most $\varepsilon + \tau$. Hence the preliminary decoder remains valid.

The construction also gives an explicit size bound. If the source circuit has N nodes, then $a - 1 \ge 3\Delta/4$ and $\sigma = K_c - a = 2\tau$. For KL, $\sigma = \Delta/8$. For residuals, $\sigma^{-1} = O_{\bar{g}}(1 + \Delta^{-1})$. Equation (71) therefore gives $L = O((1 + \Delta^{-1}) \log(2 + \Delta^{-1}))$, with constants depending only on the fixed caps. There are $O(N)$ latches and N chains; spectral adjustment adds only a constant factor. The final number of coordinates is thus

$$n_{\text{final}} = O\big(N(1 + \Delta^{-1}) \log(2 + \Delta^{-1})\big). \tag{72}$$

The equality $\widetilde{M}\widetilde{v} = K\widetilde{v}$ with $\widetilde{v} > 0$ implies $\rho(\widetilde{M}) = K$ by Collatz–Wielandt, and the cycle makes the matrix irreducible. The core has maximum degree five; Lemma E.1 with $k = 4$ gives maximum degree six overall. Let R_0 bound the preliminary fields. The bounded row sum gives an exact binary-action polymatrix realization: choose once and for all $\lambda_0 \ge R_0 + 2K_0 + 2$, and use

$$u_i(s_i, s_{-i}) = \frac{s_i}{\lambda_0}\left(b_i + \sum_{j \ne i} A_{ij} s_j\right).$$

All payoffs then lie in $[-1, 1]$, and the centered logit response is the constructed map. All added edges join distinct coordinates, so $A_{ii} = 0$ throughout. Pure-Circuit PPAD-hardness completes the proof for the positive tolerances in the theorem. □

E.1 Uniform Normalization Near the Critical Point

The general reduction uses a rationality depending on the fixed gain cap. Near criticality its fields are small enough to use the single value $\lambda_0 = 4$, as the following bounds show.

Proof of Corollary 1.3. Take $1 < K \le 101/100$. For residuals choose, for example, the fixed cap $\bar{g} = 1/2$; every $g < \mathfrak{B}(K)$ lies below this cap. In both losses, $\Delta \le \mathfrak{H}(K) < 1/100$, so the budget uses $\tau = \Delta/16$. Put $a = 1 + \eta = K - 4\tau$; then $0 < \eta < 1/100$. The core therefore uses the near-critical gates without amplifiers.

For residuals, the explicit choice in Appendix C gives $h \le 4a\eta^{3/2}$ and $P \ge (9/10)\sqrt{\eta}$. For KL, choose $h = 2h_P$ in Theorem D.5. Lemma D.1, with any remaining scalar gap γ, gives

$$h \le 2a\mathfrak{B}(a) + 2a\sqrt{2g_c} + \gamma/32 \le \tfrac{8}{3}a\eta^{3/2} + \gamma/32 < 4a\eta^{3/2},$$

where $\sqrt{2g_c} \le (2/3)\eta^{3/2}$ and $\gamma \le \mathfrak{G}(a) \le (2/9)\eta^3$. Also $P \ge (24/25)\sqrt{\eta}$. Hence $h/P < 5a\eta$ for either loss. Writing $d = c_{\mathrm{N}} = 8h/P$ as in (22), every core row obeys

$$|b_i| + \sum_j |A_{ij}| \le a + 2d + 14h \le a + 80a\eta + 56a\eta^{3/2} \le 1.01 + 0.808 + 0.05656 < 2.$$

This bounds the NOR row; the other rows have smaller bounds. Spectral adjustment increases an old row bound by at most τ and gives new rows at most $K + \tau$, so every final row is below 3. Proposition F.1 now allows the single rationality $\lambda_0 = 4$ for all these instances, keeping every pure payoff in $[-1, 1]$. The graph has degree at most six and the exact prescribed radius. Finally, $\mathrm{Gap} = \mathcal{L}_d/4$ by Proposition F.3; dividing both tolerance and slack by four preserves the separated promises. The expansion $\mathfrak{G}(1 + \eta) = \frac{2}{9}\eta^3(1 + o(1))$ gives $\mathfrak{G}(1 + \eta)/4 = \eta^3/18\,(1 + o(1))$. □

F GAME REALIZATION AND REGULARIZED LOSS

The reduction has so far produced a system of response equations. This section realizes those equations as a game and identifies the optimization quantity certified by its KL loss.

F.1 EXACT POLYMATRIX REALIZATION

A zero diagonal means that a player's response depends only on other players' strategies. We assign a payoff term to each pairwise coupling and scale all payoffs by one rationality parameter to meet the bounded-payoff requirement.

Proposition F.1. *Let $F(x) = \tanh(Ax + b)$ be a rational affine binary logit map with $A_{ii} = 0$ for every i and*

$$R := \max_i \left(|b_i| + \sum_j |A_{ij}| \right).$$

For every positive rational $\lambda_0 \geq R$, it is the centered simultaneous logit response map of a binary-action polymatrix game with rationality λ_0 and pure payoffs in $[-1, 1]$. The interaction graph is unchanged when every biased coordinate is incident to an edge; otherwise one may add a dummy zero-influence neighbor for each isolated biased coordinate.

Proof. For $s_i \in \{-1, +1\}$ define

$$u_i(s_i, s_{-i}) = \frac{s_i}{\lambda_0} \left(b_i + \sum_{j \neq i} A_{ij} s_j \right).$$

Put each interaction term on edge $\{i, j\}$; the two endpoint payoff tables on a polymatrix edge are independent. A bias may be placed on any incident edge, or on a dummy zero-influence neighbor if the coordinate is isolated. The row bound gives $|u_i| \leq R/\lambda_0 \leq 1$. Against centered mixed profile x, the pure-action payoff difference is

$$\Delta_i(x) = \frac{2}{\lambda_0} \left(b_i + \sum_j A_{ij} x_j \right),$$

so the centered logit response is

$$2 \Pr_i(+1) - 1 = \tanh\left(\frac{\lambda_0 \Delta_i(x)}{2} \right) = \tanh\left(b_i + \sum_j A_{ij} x_j \right).$$

□

The same game may be described using coefficients or changes in payoff differences. The next remark relates these conventions, so that the gain threshold is not confused with a normalization choice.

Remark F.2. If D_{ij} is the maximum change in player i's payoff difference when player j switches actions, then the realization above has $D_{ij} = 4|A_{ij}|/\lambda_0$. Thus for the derivative-normalized convention $C = D/4$,

$$\lambda_0 \rho(C) = \rho(|A|) = \kappa;$$

for the raw convention D, the same parameter is $\lambda_0 \rho(D)/4$. This accounts for the common factor-of-four normalization discrepancy.

F.2 EXACT KL IDENTITY

The realization identifies the response map, but we also need to interpret its error in optimization terms. We compare a player's current regularized payoff with the best payoff achievable by changing

only that player's strategy. Suppose the affine response is induced by a binary game with rationality parameter $\lambda > 0$. Let

$$q_i = \pi_{x_i}, \qquad y_i = \pi_{F_i(x)},$$

and let $u_i(q_{-i}) = (u_i^+, u_i^-)$ be the two pure expected payoffs. With $H(p) = -\sum_a p_a \log p_a$, define

$$\mathrm{Gap}_i(q) := \max_{p_i \in \Delta_2} \left\{ \langle p_i, u_i \rangle + \frac{1}{\lambda} H(p_i) \right\} - \left\{ \langle q_i, u_i \rangle + \frac{1}{\lambda} H(q_i) \right\}.$$

Write $\mathrm{Gap}(q) = \max_i \mathrm{Gap}_i(q)$.

Here Δ_2 is the probability simplex on two actions. The next identity turns response consistency into an exact bound on each player's available regularized improvement; the additional inequalities relate this quantity to the residual stopping rule.

Proposition F.3. *For every player,*

$$\mathrm{Gap}_i(q) = \frac{1}{\lambda} D_{\mathrm{KL}}(q_i \| y_i).$$

If

$$R := \max_i \left(|b_i| + \sum_j |A_{ij}| \right), \qquad r := \mathrm{Res}_F(x),$$

then

$$\frac{r^2}{2\lambda} \le \mathrm{Gap}(q) \le \frac{1}{\lambda} \min \left\{ (R + \log 2) r,\ \cosh^2 R\, r^2 \right\}. \tag{73}$$

Proof. Let $Z_i = \sum_{a \in \{+,-\}} e^{\lambda u_i^a}$, so $y_i(a) = e^{\lambda u_i^a} / Z_i$. The Gibbs variational identity gives

$$\max_{p_i} \left\{ \langle p_i, u_i \rangle + \lambda^{-1} H(p_i) \right\} = \lambda^{-1} \log Z_i,$$

and direct substitution yields $\mathrm{Gap}_i(q) = \lambda^{-1} D_{\mathrm{KL}}(q_i \| y_i)$. Put $r_i = |x_i - F_i(x)| = 2|q_i(+) - y_i(+)|$. Pinsker gives $D_{\mathrm{KL}}(q_i \| y_i) \ge r_i^2/2$. Conversely, $D_{\mathrm{KL}}(p\|y) \le \chi^2(p\|y)$, where $\chi^2(p\|y) = \sum_a (p_a - y_a)^2 / y_a$. For Bernoulli distributions,

$$\chi^2(q_i \| y_i) = \frac{(q_i(+) - y_i(+))^2}{y_i(+)(1 - y_i(+))}.$$

If $h_i = b_i + \sum_j A_{ij} x_j$, then $y_i(+)(1 - y_i(+)) = 1/(4\cosh^2 h_i) \ge 1/(4\cosh^2 R)$. Taking the maximum gives the quadratic bound; the linear bound is Lemma A.2. □

Thus KL accuracy is exactly a regularized unilateral-improvement guarantee. The linear comparison also turns a sufficiently small residual into a KL certificate, which will establish membership for the second loss below.

G Explicit PPAD Membership and Encoding

The lower bound does not by itself establish PPAD-completeness. We must also place the search problem in PPAD despite its irrational response values, which requires a finite encoding and a verifiable notion of an approximate answer.

The Certified Search Relation. A search relation specifies which finite output strings are accepted for each input. Our accepted outputs encode a point and a certified loss bound, avoiding exact comparisons with transcendental numbers. Fix a deterministic rational interval evaluator for $L(x) = \mathrm{Res}_F(x)$ or $\mathcal{L}_d(x; F)$. On input (A, b, x, m) it returns

$$L(x) \in [\underline{L}_m, \overline{L}_m], \qquad \overline{L}_m - \underline{L}_m \le 2^{-m},$$

in time polynomial in the rational input lengths and m. Lemma B.5 and the stable KL formula in Appendix B give such evaluators, including at $x_i = \pm 1$.

Here is an explicit polynomial bound on witnesses, meaning accepted output strings. Put $R = \max_i(|b_i| + \sum_j |A_{ij}|)$ and $s = \min\{1, g\}$; when an upper-bound slack δ is supplied, use $s = \min\{1, g, \delta\}$. Fix

$$k = \left\lceil \log_2 \frac{64(1+R)^2}{s} \right\rceil, \qquad m_* = \left\lceil \log_2 \frac{16}{s} \right\rceil, \qquad h = 2^{-k}. \tag{74}$$

These integers can be found by rational comparisons with powers of two. A witness consists of n signed integers z_i, each in $k+2$ bits with $|z_i| \le 2^k$, representing $x_i = 2^{-k} z_i$, and a unary precision 1^m (a string of m ones), with $1 \le m \le m_*$. The verifier recomputes the loss interval and accepts exactly when $\overline{L}_m \le g$. The witness uses $n(k+2) + m + O(n)$ bits, and

$$k, m_* = O\big(1 + \log(1+R) + \log(1/s)\big).$$

Thus the relation is polynomially balanced: accepted outputs have length bounded by a polynomial in the binary input length. A point at equality need not have a certificate; the constructions below leave strict slack. Every accepted witness satisfies the actual loss bound, which suffices for hardness. All PPAD membership statements use this certified convention.

Rounding a cube point to this dyadic grid changes it by at most h in max norm and remains in the cube. Both F and every modified map $S = \chi_\alpha(Ax + b)$ with $\alpha > 1$ have unweighted Lipschitz constant at most R. Consequently either residual changes by at most

$$(1+R)h \le \frac{s}{64(1+R)}.$$

This bound will put all outputs on the same grid, independently of the internal arithmetic used by a solver.

It remains to show that an accepting output always exists and can be found through a PPAD search procedure. The key step is to replace the original map by a continuous, rationally evaluable approximation with controlled error.

Proposition G.1. *For every positive rational ε, finding a certified rational ε-residual fixed point of a rational affine logit map belongs to PPAD. The tolerance may be inverse exponential when encoded in binary.*

Proof. The map F is L-Lipschitz in max norm, where $L = \max_i \sum_j |A_{ij}|$ has polynomial encoding length. A direct appeal to a rational-evaluation fixed-point theorem would be insufficient: F need not take rational inputs to rational outputs. We first construct a *continuous*, exactly rational-evaluable approximation.

Put $\varepsilon_0 = \min\{\varepsilon, 1\}$ and $\xi = \varepsilon_0/16$. Choose the largest mesh width $h = 2^{-k}$, $k \in \mathbb{N}$, with $h \le \varepsilon_0/[16(1+L)]$, and triangulate the cube into the standard ordered-coordinate simplices of this grid. At each grid vertex v, assign the same deterministically computed rational vector $q(v) \in [-1, 1]^n$ with $\|q(v) - F(v)\|_\infty \le \xi$. Extend these values affinely on each simplex to obtain $\widetilde{F}$.

This map is continuous because adjacent simplices use the same values on shared vertices. It maps the cube into itself by convexity. Its description is succinct: at a rational point, integer arithmetic locates the grid cell; sorting the fractional coordinates identifies a simplex and its $n+1$ vertices. Their barycentric weights and the $n+1$ rational values $q(v)$ are computed on demand. Ties lie on common faces and give the same interpolated value. No exponential vertex table is stored. Thus $\widetilde{F}$ is exactly rational-evaluable in polynomial time.

For every point x of a simplex, all its vertices are within h of x, so convex interpolation gives

$$\|\widetilde{F}(x) - F(x)\|_\infty \le Lh + \xi \le \varepsilon_0/8. \tag{75}$$

Consecutive vertices of an ordered-coordinate simplex differ in one coordinate by h. Each entry of its affine derivative therefore has magnitude at most $L + 2\xi/h$. Hence $\widetilde{F}$ is globally $n(L + 2\xi/h)$-Lipschitz in max norm, by splitting a line segment across simplices. This bound has polynomial encoding length; in fact the maximal dyadic choice gives $\xi/h < 2(1+L)$.

We apply Etessami and Yannakakis (2010, Proposition 2.2(2)): for a polynomially computable, polynomially continuous family, finding a point with small fixed-point residual (*weak approximation*)

belongs to PPAD. Polynomial continuity means that the required number of input-precision bits is polynomial in the instance length and the requested number of output-precision bits. Exact rational evaluation above proves the first hypothesis; the Lipschitz bound of polynomial encoding length proves the second. The cube is a rational polytope with a polynomial-size description, as required by that framework. On $[0,1]^n$ use $u \mapsto (\widetilde{F}(2u-\mathbf{1})+\mathbf{1})/2$ with tolerance $\varepsilon_0/16$. Undoing the affine change gives a rational point with

$$\|x-\widetilde{F}(x)\|_\infty < \varepsilon_0/8, \qquad \mathrm{Res}_F(x) < \varepsilon_0/4.$$

The grid description, evaluations, and returned point have length polynomial in the input and $\log(2+1/\varepsilon)$. Round this point to the output grid (74), using $g=\varepsilon$. The residual is then less than $\varepsilon_0/4+s/64$, and an interval at precision m_* has width at most $s/16$. Since $s \le \varepsilon_0$, its upper endpoint is below ε, so the verifier accepts. This also proves totality if a smaller positive supplied slack makes the output grid finer. Positive tolerance is essential: exact fixed points need not be rational. □

This proves membership for residual accuracy. The boundary-valid linear KL bound now transfers membership to KL with only polynomially many additional precision bits.

Corollary G.2. *For arbitrary rational A, b and positive rational g, finding a certified rational profile with $\mathcal{L}_d(x;F) \le g$ belongs to PPAD, including inverse-exponential g encoded in binary. Consequently the KL lower bounds in Theorem 1.2 are PPAD-completeness results.*

Proof. Put $g_0 = \min\{g,1\}$ and $R = \max_i(|b_i| + \sum_j |A_{ij}|)$. Invoke Proposition G.1 with the rational residual tolerance

$$e = \frac{g_0}{8(1+R)}.$$

The tolerance has polynomial encoding length, even when R is exponentially large. Round the returned point to the output grid for KL tolerance g. Its residual is at most $e + s/[64(1+R)]$, so Lemma A.2 gives

$$\mathcal{L}_d(x;F) \le (R+\log 2)\left(e + \frac{s}{64(1+R)}\right) < \frac{g_0}{8} + \frac{s}{64}.$$

The loss interval at precision m_* adds at most $s/16$. Since $s \le g_0$, its upper endpoint is below g, giving an accepting witness. □

Both search problems now have a finite encoding and a PPAD membership proof. We finally check that the constructive algorithm of Section 3 returns an answer accepted by this same convention.

Certificates for the Upper-Bound Algorithm. Before its final output rounding, Algorithm 1 has residual at most $\mathfrak{B}(\kappa) + 9\delta_0/16$ and KL loss at most $\mathfrak{G}(\kappa) + 5\delta_0/32$, where $\delta_0 = \min\{\delta, 1\}$. Use (74) with $s = \min\{1, g, \delta\}$. Rounding the returned iterate changes its defect for S by at most $s/[64(1+R)]$. The triangle inequality for residual loss and Lemma A.2, applied with the reference $S(x)$ at the rounded point, show that the corresponding upper bounds increase by at most $s/64$. Finally, precision m_* adds at most $s/16$. Because $s \le \delta_0 \le \delta$, the certified bounds are strictly below $\mathfrak{B}(\kappa)+\delta$ and $\mathfrak{G}(\kappa)+\delta$, respectively. This proves that the same explicit search relation accepts both upper-bound algorithms. Internal iterates may also be rounded to a dyadic grid, charging that error to ζ as in Lemma B.4; their precision need not equal the final output precision k.