AIRR.SCIENCE · ARR-2026-5V2B5Y9RCH9PT8VY v1 [math.GT] 29 Nov 2026

Cross-Ratio Hyperbolic Contraction, Fisher–Rao Adversarial Shock Dampening, Rényi Nash-Gap Duality, and Universal Damped Best-Response Laws for Entropic Zero-Sum Matrix Games

Lluis Eriksson
Independent Researcher, Östersund, Sweden
editor@airr.science

November 29, 2026

Abstract

Consider a two-player zero-sum matrix game with arbitrary payoff matrix $A \in \mathbb{R}^{m\times n}$ $(m, n \geq 2)$, strictly positive reference priors $p \in \Delta_m^\circ$, $q \in \Delta_n^\circ$, and player regularization temperatures $\tau_1, \tau_2 > 0$, defined by the saddle-point problem $\min_{x\in\Delta_m} \max_{y\in\Delta_n} \{x^\top Ay + \tau_1 D_{\mathrm{KL}}(x\|p) - \tau_2 D_{\mathrm{KL}}(y\|q)\}$ with unique Quantal-Response Equilibrium (QRE) $(x^*, y^*) \in \Delta_m^\circ \times \Delta_n^\circ$ and logit best-response maps $T_1 : \Delta_n \to \Delta_m^\circ$, $T_2 : \Delta_m \to \Delta_n^\circ$. We establish a unified geometric, spectral, thermodynamic, and algorithmic theory governed by a single strategic-invariant 2×2 cross-ratio diameter $\Delta_{\mathrm{cr}}(A) := \max_{i,i',j,j'} |A_{ij} - A_{i'j} - A_{ij'} + A_{i'j'}|$ (and dimensionless game coupling $\kappa := \Delta_{\mathrm{cr}}(A)^2/(16\tau_1\tau_2)$) together with the Fisher–Rao whitened interaction operator $K_* := (\tau_1\tau_2)^{-1/2}\Pi_{x^*} D_{x^*}^{1/2} A D_{y^*}^{1/2}\Pi_{y^*}$.

First, via an elementary sharp bound relating Total Variation to Hilbert's projective metric d_H, we prove exact nonlinear hyperbolic-tangent contraction laws for logit best responses, $\|T_1(y) - T_1(y')\|_{\mathrm{TV}} \leq \tanh(\frac{\Delta_{\mathrm{cr}}(A)}{4\tau_1}\|y - y'\|_{\mathrm{TV}})$ and $d_H(T_1(y), T_1(y')) \leq \frac{\Delta_{\mathrm{cr}}(A)}{\tau_1}\tanh(\frac{1}{4}d_H(y, y'))$. Thus $F = T_1 \circ T_2$ is globally contractive in both metrics whenever $\kappa < 1$ $(\sqrt{\tau_1\tau_2} > \Delta_{\mathrm{cr}}(A)/4)$, and this threshold is universally sharp: on 2×2 matching pennies, F undergoes an exact supercritical period-2 bifurcation at $\kappa = 1$ with limit-cycle TV amplitude solving $r = \tanh(\sqrt{\kappa}\, r)$.

Second, in Fisher–Rao tangent coordinates we derive the exact closed-form eigenvalue spectrum of continuous-time entropic replicator–logit dynamics, $\lambda_k^\pm = -\frac{\tau_1+\tau_2}{2} \pm \sqrt{(\frac{\tau_1-\tau_2}{2})^2 - \tau_1\tau_2\sigma_k^2}$, revealing an asymmetric-temperature overdamped-to-underdamped phase transition at $\sigma_c(\tau_1, \tau_2) = \sinh(\frac{1}{2}|\ln(\tau_1/\tau_2)|)$. Using normality of the skew-perturbed operator $I + S_*$, we prove the *Adversarial Shock-Dampening Law*: in temperature-weighted Fisher–Rao norm, strategic zero-sum coupling never amplifies payoff shocks and strictly dampens every active interaction mode by the exact cross-term-free factor $(1 + \sigma_k^2)^{-1}$.

Third, we prove an *exact* zero-error Rényi ∞-to-1 divergence identity for the unregularized Nash gap of any QRE, $\mathrm{NashGap}(x^*, y^*) = \tau_1(D_\infty(x^*\|p) - D_{\mathrm{KL}}(x^*\|p)) + \tau_2(D_\infty(y^*\|q) - D_{\mathrm{KL}}(y^*\|q))$ (reducing for uniform priors to $\tau_1(H(x^*) - H_\infty(x^*)) + \tau_2(H(y^*) - H_\infty(y^*))$), together with exact Schur-complement temperature heat capacities proving global concavity of $v(\tau_1, \tau_2)$ in τ_1 and convexity in τ_2.

Fourth, we establish the exact Pythagorean duality gap $\mathrm{Gap}_\tau(x, y) = \tau_1 D_{\mathrm{KL}}(x\|T_1(y)) + \tau_2 D_{\mathrm{KL}}(y\|T_2(x))$, a universal one-step KL equilibrium certificate, exact continuous-time reverse-KL excess dissipation, and an unconditional discrete-time geometric contraction law for damped best response ($\eta^* = 1$ for $\kappa \leq \frac{1}{2}$ and $\eta^* = 1/(2\kappa)$ for $\kappa > \frac{1}{2}$), guaranteeing $\mathrm{Gap}_\tau(x_{k+1}, y_{k+1}) \leq (1 - \frac{1}{4\kappa})\,\mathrm{Gap}_\tau(x_k, y_k)$ for all $\kappa > \frac{1}{2}$ and eliminating limit cycles at every temperature.

Keywords: zero-sum matrix games, quantal-response equilibrium, entropic regularization, Hilbert projective metric, Fisher–Rao geometry, Rényi divergence, damped best-response dynamics.

1 Introduction

Entropically regularized two-player zero-sum matrix games—equivalently, Quantal-Response Equilibria (QRE) [7], smoothed fictitious play [5, 6], and KL-regularized minimax saddle-point problems in reinforcement learning and online optimization [8, 3, 9]—are defined for an arbitrary payoff matrix $A \in \mathbb{R}^{m\times n}$ $(m, n \geq 2)$, strictly positive reference priors $p \in \Delta_m^\circ$, $q \in \Delta_n^\circ$, and regularization temperatures $\tau_1, \tau_2 > 0$ by

$$v(\tau_1, \tau_2) \;:=\; \min_{x\in\Delta_m} \max_{y\in\Delta_n} \mathcal{L}(x, y), \qquad \mathcal{L}(x, y) \;:=\; x^\top A y + \tau_1 D_{\mathrm{KL}}(x\|p) - \tau_2 D_{\mathrm{KL}}(y\|q), \tag{1}$$

where $\Delta_m := \{x \in \mathbb{R}^m_{\geq 0} : \mathbf{1}_m^\top x = 1\}$, $\Delta_m^\circ := \Delta_m \cap \mathbb{R}^m_{>0}$, and $D_{\mathrm{KL}}(u\|v) := \sum_i u_i \log(u_i/v_i)$. Because $\mathcal{L}$ is strictly convex in x and strictly concave in y with steep entropic barriers at the boundary, (1) admits a unique interior saddle point $(x^*, y^*) \in \Delta_m^\circ \times \Delta_n^\circ$ characterized by the fixed-point equations $x^* = T_1(y^*)$ and $y^* = T_2(x^*)$ of the logit best-response operators

$$T_1(y)_i \;:=\; \frac{p_i \exp\big(-(Ay)_i/\tau_1\big)}{\sum_{i'=1}^m p_{i'} \exp\big(-(Ay)_{i'}/\tau_1\big)}, \qquad T_2(x)_j \;:=\; \frac{q_j \exp\big((A^\top x)_j/\tau_2\big)}{\sum_{j'=1}^n q_{j'} \exp\big((A^\top x)_{j'}/\tau_2\big)}. \tag{2}$$

Despite extensive literature on asymptotic convergence of smoothed fictitious play [6], extra-gradient / optimistic mirror descent [3, 4], and replicator dynamics [8], four sharp non-asymptotic geometric questions have remained fragmented or uncharacterized:

1. **Exact nonlinear contraction and bifurcation threshold.** What strategic-invariant payoff functional governs the exact global Total Variation and Hilbert projective contraction of T_1, T_2, and where does undamped best response $F = T_1 \circ T_2$ bifurcate into limit cycles?
2. **Fisher–Rao spectrum and adversarial payoff-shock sensitivity.** How do asymmetric player temperatures (τ_1, τ_2) govern oscillatory versus overdamped convergence in Fisher–Rao geometry, and does adversarial zero-sum coupling amplify or dampen equilibrium sensitivity to payoff perturbations δA?
3. **Exact unregularized Nash gap and temperature thermodynamics.** Can the unregularized Nash duality gap $\mathrm{NashGap}(x^*, y^*) := \max_y (x^*)^\top A y - \min_x x^\top A y^*$ be expressed *identically* in information-theoretic terms of (x^*, y^*), and what are the exact second derivatives of $v(\tau_1, \tau_2)$?
4. **Universal discrete-time best-response contraction across all temperatures.** When $\kappa \geq 1$ causes undamped best response to cycle, what explicit damping schedule η^* guarantees monotone geometric contraction of $\mathrm{Gap}_\tau(x_k, y_k)$ for every game and every temperature $\tau_1, \tau_2 > 0$?

Summary of Main Results

We resolve all four questions cleanly through four unified theorems:

- **Theorem 2.3 (Cross-Ratio Hyperbolic Contraction and Period-2 Bifurcation):** Defining the shift-invariant 2×2 cross-ratio diameter $\Delta_{\mathrm{cr}}(A)$ and coupling $\kappa = \Delta_{\mathrm{cr}}(A)^2/(16\tau_1\tau_2)$, we prove exact nonlinear tanh contraction bounds in both Total Variation and Hilbert projective metrics, global contraction of $F = T_1 \circ T_2$ for $\kappa < 1$, and exact supercritical period-doubling bifurcation of matching pennies at $\kappa = 1$ with amplitude $r = \tanh(\sqrt{\kappa}\, r)$.
- **Theorem 3.1 (Fisher–Rao Spectrum, Oscillatory Transition, and Adversarial Shock Dampening):** Using the whitened interaction matrix K_*, we determine every eigenvalue of continuous-time entropic replicator flow in closed form, locate the sharp overdamped-to-underdamped phase boundary at $\sigma_c = \sinh(\frac{1}{2}|\ln(\tau_1/\tau_2)|)$, and prove the exact cross-term-free *Adversarial Shock-Dampening Law* showing that zero-sum coupling dampens payoff shocks by $(1 + \sigma_k^2)^{-1} \leq 1$.

- **Theorem 4.1 (Exact Rényi Nash-Gap Identity and Temperature Thermodynamics):** We prove the zero-error identity $\mathrm{NashGap}(x^*, y^*) = \tau_1(D_\infty(x^*\|p) - D_{\mathrm{KL}}(x^*\|p)) + \tau_2(D_\infty(y^*\|q) - D_{\mathrm{KL}}(y^*\|q))$ and derive exact Schur-complement temperature Hessians establishing concavity of v in τ_1 and convexity in τ_2.
- **Theorem 5.1 (Pythagorean Gap, One-Step KL Certificates, and Universal Damped Contraction):** We prove the exact Pythagorean gap identity, an unconditional one-step KL equilibrium certificate, continuous-time reverse-KL excess dissipation, and discrete-time damped contraction $\mathrm{Gap}_\tau(x_{k+1}, y_{k+1}) \le (1 - \eta + \kappa\eta^2)\,\mathrm{Gap}_\tau(x_k, y_k)$ with optimal damping $\eta^* = 1$ ($\kappa \le \frac{1}{2}$) and $\eta^* = 1/(2\kappa)$ ($\kappa > \frac{1}{2}$).

All proofs and numerical benchmarks (4/4 pytest modules) are openly archived at: https://github.com/lluiseriksson/entropic-games-qre-geometry.

2 Cross-Ratio Hyperbolic Contraction and Period-2 Bifurcation

For any $u, v \in \Delta_m$, let $\|u - v\|_{\mathrm{TV}} := \frac{1}{2}\sum_{i=1}^m |u_i - v_i|$ denote Total Variation distance, and for $u, v \in \Delta_m^\circ$, let

$$d_H(u, v) \;:=\; \mathrm{osc}\left(\log \frac{u}{v}\right) \;=\; \max_{1\le i,i'\le m} \log \frac{u_i v_{i'}}{u_{i'} v_i}$$

denote Hilbert's projective metric [1, 2], where $\mathrm{osc}(w) := \max_i w_i - \min_i w_i$.

Definition 2.1 (Strategic Cross-Ratio Diameter and Coupling Invariant)**.** For $A \in \mathbb{R}^{m\times n}$ ($m, n \ge 2$) and $\tau_1, \tau_2 > 0$, define the 2×2 cross-ratio diameter and dimensionless game coupling invariant

$$\Delta_{\mathrm{cr}}(A) \;:=\; \max_{\substack{1\le i,i'\le m\\ 1\le j,j'\le n}} |A_{ij} - A_{i'j} - A_{ij'} + A_{i'j'}|, \qquad \kappa \;:=\; \frac{\Delta_{\mathrm{cr}}(A)^2}{16\,\tau_1\tau_2}. \tag{3}$$

Observe that $\Delta_{\mathrm{cr}}(A + r\mathbf{1}_n^\top + \mathbf{1}_m c^\top) = \Delta_{\mathrm{cr}}(A)$ for every $r \in \mathbb{R}^m$, $c \in \mathbb{R}^n$: $\Delta_{\mathrm{cr}}(A)$ is invariant under all strategically equivalent row/column shifts, and $\Delta_{\mathrm{cr}}(A) = 0$ if and only if A is strategically separable ($A = r\mathbf{1}_n^\top + \mathbf{1}_m c^\top$), in which case T_1, T_2 are constant maps.

Lemma 2.2 (Sharp Total-Variation-from-Hilbert Bound)**.** *For every $m \ge 2$ and every $u, v \in \Delta_m^\circ$,*

$$\|u - v\|_{\mathrm{TV}} \;\le\; \tanh\left(\frac{d_H(u, v)}{4}\right), \tag{4}$$

with equality attained for every $m \ge 2$ and every $R = d_H(u, v) \ge 0$ by any partition $[m] = S_+ \sqcup S_-$ into two nonempty blocks with total masses $\sum_{i\in S_+} u_i = \frac{e^{R/2}}{1+e^{R/2}}$, $\sum_{i\in S_+} v_i = \frac{1}{1+e^{R/2}}$ and constant coordinatewise ratios $u_i/v_i = e^{+R/2}$ on S_+ and $u_i/v_i = e^{-R/2}$ on S_- (reducing to the 2-point symmetric pair when $m = 2$).

Proof. Set $R := d_H(u, v)$, $S_+ := \{i \in [m] : u_i \ge v_i\}$, $S_- := [m] \setminus S_+$, $a := \sum_{i\in S_+} u_i$, and $b := \sum_{i\in S_+} v_i$. If $a = b$, then $\|u - v\|_{\mathrm{TV}} = 0 \le \tanh(R/4)$. Otherwise $0 < b < a < 1$, so by averaging there exist $i_+ \in S_+$ with $u_{i_+}/v_{i_+} \ge a/b$ and $i_- \in S_-$ with $u_{i_-}/v_{i_-} \le (1 - a)/(1 - b)$. Hence

$$e^R \;\ge\; \frac{u_{i_+}/v_{i_+}}{u_{i_-}/v_{i_-}} \;\ge\; \frac{a(1-b)}{b(1-a)} \quad\Longrightarrow\quad \frac{a}{1-a} \;\le\; e^R \frac{b}{1-b}.$$

Setting $s := b/(1 - b) > 0$, we obtain $a \le \frac{e^R s}{1+e^R s}$ and $b = \frac{s}{1+s}$, so

$$\|u - v\|_{\mathrm{TV}} \;=\; a - b \;\le\; \frac{e^R s}{1 + e^R s} - \frac{s}{1 + s} \;=\; \frac{(e^R - 1)s}{(1 + e^R s)(1 + s)} \;=\; \frac{e^R - 1}{e^R s + s^{-1} + e^R + 1}.$$

By AM–GM, $e^R s + s^{-1} \ge 2e^{R/2}$ with equality at $s^* = e^{-R/2}$, yielding $\|u - v\|_{\mathrm{TV}} \le \frac{e^R-1}{(e^{R/2}+1)^2} = \frac{e^{R/2}-1}{e^{R/2}+1} = \tanh(R/4)$, attained at u, v above. □

Theorem 2.3 (Cross-Ratio Hyperbolic Contraction and Exact Period-2 Bifurcation)**.** *Let $A \in \mathbb{R}^{m\times n}$ $(m, n \geq 2)$, $p \in \Delta_m^\circ$, $q \in \Delta_n^\circ$, and $\tau_1, \tau_2 > 0$, with $\Delta_{\mathrm{cr}}(A)$ and κ as in (3).*

(a) ***Nonlinear Hyperbolic TV and Hilbert Contraction.*** *For all $y, y' \in \Delta_n$ (and $y, y' \in \Delta_n^\circ$ for d_H),*

$$d_H(T_1(y),\, T_1(y')) \;\leq\; \frac{\Delta_{\mathrm{cr}}(A)}{\tau_1}\|y - y'\|_{\mathrm{TV}} \;\leq\; \frac{\Delta_{\mathrm{cr}}(A)}{\tau_1}\tanh\!\left(\frac{d_H(y, y')}{4}\right), \tag{5}$$

$$\|T_1(y) - T_1(y')\|_{\mathrm{TV}} \;\leq\; \tanh\!\left(\frac{\Delta_{\mathrm{cr}}(A)}{4\tau_1}\|y - y'\|_{\mathrm{TV}}\right) \;\leq\; \frac{\Delta_{\mathrm{cr}}(A)}{4\tau_1}\|y - y'\|_{\mathrm{TV}}, \tag{6}$$

and symmetrically for T_2 with τ_2 replacing τ_1. In particular, the images $T_1(\Delta_n)$ and $T_2(\Delta_m)$ have TV diameters at most $\tanh(\Delta_{\mathrm{cr}}(A)/(4\tau_1))$ and $\tanh(\Delta_{\mathrm{cr}}(A)/(4\tau_2))$.

(b) ***Global Contraction Threshold.*** *The two-step best-response map $F := T_1 \circ T_2 : \Delta_m \to \Delta_m^\circ$ satisfies*

$$\begin{aligned} \|F(x) - F(x')\|_{\mathrm{TV}} &\leq \kappa\,\|x - x'\|_{\mathrm{TV}} \qquad (\forall x, x' \in \Delta_m), \\ d_H(F(x), F(x')) &\leq \kappa\, d_H(x, x') \qquad (\forall x, x' \in \Delta_m^\circ), \end{aligned} \tag{7}$$

and is therefore a strict global contraction in both TV (on Δ_m) and Hilbert projective metrics (on Δ_m°) whenever $\kappa < 1$, i.e., $\sqrt{\tau_1\tau_2} > \Delta_{\mathrm{cr}}(A)/4$.

(c) ***Universal Sharpness and Exact Matching-Pennies Period-2 Bifurcation.*** *The threshold $\kappa = 1$ is universally sharp. On the 2×2 matching-pennies game $A = \left[\begin{smallmatrix} a & -a \\ -a & a \end{smallmatrix}\right]$ $(a > 0,\ \Delta_{\mathrm{cr}}(A) = 4a)$ with uniform priors $p = q = (\frac{1}{2}, \frac{1}{2})$ and $\tau_1 = \tau_2 = \tau$ $(\kappa = a^2/\tau^2)$:*

- *For $\kappa \leq 1$, $x^* = (\frac{1}{2}, \frac{1}{2})$ is globally attracting under F.*
- *For $\kappa > 1$, F undergoes a supercritical period-doubling bifurcation: every orbit initialized at $x_0 \neq x^*$ converges to a unique period-2 limit cycle $\{x^+, x^-\}$ with $x^\pm = (\frac{1\pm r_*}{2}, \frac{1\mp r_*}{2})$ whose exact TV amplitude $r_* = \|x^+ - x^-\|_{\mathrm{TV}} \in (0, 1)$ is the unique positive root of*

$$r_* \;=\; \tanh\!\big(\sqrt{\kappa}\, r_*\big). \tag{8}$$

Proof. **(a)** By (2), $\log(T_1(y)_i/T_1(y')_i) = -\tau_1^{-1}(A(y - y'))_i + \text{const}$, so $d_H(T_1(y), T_1(y')) = \tau_1^{-1}\operatorname{osc}(A(y - y'))$. For any $i, i' \in [m]$, let $w_j := A_{ij} - A_{i'j}$ and $c := \frac{1}{2}(\max_j w_j + \min_j w_j)$, so $\|w - c\mathbf{1}_n\|_\infty = \frac{1}{2}\operatorname{osc}(w) \leq \frac{1}{2}\Delta_{\mathrm{cr}}(A)$. Since $\mathbf{1}_n^\top(y - y') = 0$ and $\|y - y'\|_1 = 2\|y - y'\|_{\mathrm{TV}}$,

$$|(A(y - y'))_i - (A(y - y'))_{i'}| \;=\; |\langle w - c\mathbf{1}_n,\, y - y'\rangle| \;\leq\; \frac{\operatorname{osc}(w)}{2}\|y - y'\|_1 \;\leq\; \Delta_{\mathrm{cr}}(A)\|y - y'\|_{\mathrm{TV}}.$$

Taking the maximum over i, i' yields the first inequality in (5); applying Lemma 2.2 to $\|y - y'\|_{\mathrm{TV}}$ yields the second, and applying Lemma 2.2 to $T_1(y), T_1(y')$ yields (6) (with $\tanh u \leq u$ for $u \geq 0$). The symmetric bounds for T_2 follow identically using $\Delta_{\mathrm{cr}}(A^\top) = \Delta_{\mathrm{cr}}(A)$.

(b) Composing the linear upper bounds from (6) and (5) (using $\tanh(u/4) \leq u/4$) yields (7) with product constant $\frac{\Delta_{\mathrm{cr}}(A)}{4\tau_1} \cdot \frac{\Delta_{\mathrm{cr}}(A)}{4\tau_2} = \kappa$.

(c) Parameterize Δ_2 by $x(r) := (\frac{1+r}{2}, \frac{1-r}{2})$ for $r \in [-1, 1]$. For $A = \left[\begin{smallmatrix} a & -a \\ -a & a \end{smallmatrix}\right]$ and uniform priors, $A^\top x(r) = (ar, -ar)^\top$, hence $T_2(x(r)) = y(g(r))$ where $g(r) := \tanh(\frac{a}{\tau}r) = \tanh(\sqrt{\kappa}\, r)$. Similarly, $Ay(s) = (as, -as)^\top$, hence $T_1(y(s)) = x(-g(s))$. Thus $F(x(r)) = x(-g(g(r)))$ and $F^2(x(r)) = x(g^{\circ 4}(r))$. Because g is odd, strictly increasing, and strictly concave on $(0, \infty)$ with $g'(0) = \sqrt{\kappa}$ and $\lim_{r\to\infty} g(r) = 1$, for $\kappa \leq 1$ we have $0 < g(r) < r$ for all $r > 0$, so $g^{\circ k}(r) \to 0$ globally. For $\kappa > 1$, g has a unique positive fixed point $r_* \in (0, 1)$ solving (8), with $g(r) > r$ on $(0, r_*)$ and $g(r) < r$ on $(r_*, 1]$, so $g^{\circ 4k}(r) \to r_*$ for every $r \in (0, 1]$ (and $-r_*$ for $r \in [-1, 0)$). Since $F(x(r_*)) = x(-g(g(r_*))) = x(-r_*)$ and $\|x(r_*) - x(-r_*)\|_{\mathrm{TV}} = r_*$, $\{x^+, x^-\}$ is the unique attracting period-2 orbit. □

3 Fisher–Rao Interaction Spectrum and Adversarial Shock Dampening

Equip the open simplex Δ_m° with the Fisher–Rao Riemannian metric $\langle v, w\rangle_{\mathrm{FR},x} := \sum_{i=1}^m \frac{v_i w_i}{x_i}$ on $T_x\Delta_m = \{v \in \mathbb{R}^m : \mathbf{1}_m^\top v = 0\}$. At the unique QRE $(x^*, y^*) \in \Delta_m^\circ \times \Delta_n^\circ$, write $D_{x^*} := \mathrm{diag}(x^*)$, $D_{y^*} := \mathrm{diag}(y^*)$, let $\Pi_{x^*} := I_m - \sqrt{x^*}\sqrt{x^*}^\top$ and $\Pi_{y^*} := I_n - \sqrt{y^*}\sqrt{y^*}^\top$ be the orthogonal projections onto $(\sqrt{x^*})^\perp$ and $(\sqrt{y^*})^\perp$ (note that $v \mapsto D_{x^*}^{-1/2} v$ maps $T_{x^*}\Delta_m$ isometrically onto $(\sqrt{x^*})^\perp$), and define the *Fisher–Rao whitened interaction operator*

$$K_* \;:=\; \frac{1}{\sqrt{\tau_1\tau_2}}\,\Pi_{x^*} D_{x^*}^{1/2} A D_{y^*}^{1/2} \Pi_{y^*} \;\in\; \mathbb{R}^{m\times n}. \tag{9}$$

Let $r := \mathrm{rank}(K_*) \le \min(m-1, n-1)$ and let $\sigma_1 \ge \cdots \ge \sigma_r > 0$ be the nonzero singular values of K_*. Consider the continuous-time entropic replicator–logit (Fisher–Rao natural gradient) flow on $\Delta_m^\circ \times \Delta_n^\circ$ [8]:

$$\dot{x}_i = -x_i\Big(\xi_i - \sum_{i'=1}^m x_{i'}\xi_{i'}\Big), \qquad \dot{y}_j = -y_j\Big(\zeta_j - \sum_{j'=1}^n y_{j'}\zeta_{j'}\Big), \tag{10}$$

where $\xi := \tau_1 \log(x/p) + Ay$ and $\zeta := \tau_2 \log(y/q) - A^\top x$.

Theorem 3.1 (Fisher–Rao Spectrum, Oscillatory Phase Transition, and Adversarial Shock Dampening)**.** *Let $(x^*, y^*) \in \Delta_m^\circ \times \Delta_n^\circ$ be the unique QRE of (1), with whitened interaction operator K_* in (9) having nonzero singular values $\sigma_1 \ge \cdots \ge \sigma_r > 0$.*

(a) ***Exact Closed-Form Spectrum and Oscillatory Phase Transition.*** *Linearizing (10) at (x^*, y^*) in Fisher–Rao tangent coordinates $(u, w) := (D_{x^*}^{-1/2}\delta x,\, D_{y^*}^{-1/2}\delta y) \in (\sqrt{x^*})^\perp \oplus (\sqrt{y^*})^\perp$ yields*

$$\frac{d}{dt}\begin{pmatrix} u \\ w \end{pmatrix} \;=\; J_{\mathrm{FR}} \begin{pmatrix} u \\ w \end{pmatrix}, \qquad J_{\mathrm{FR}} \;:=\; \begin{pmatrix} -\tau_1 I_m & -\sqrt{\tau_1\tau_2}\,K_* \\ \sqrt{\tau_1\tau_2}\,K_*^\top & -\tau_2 I_n \end{pmatrix}\Bigg|_{(\sqrt{x^*})^\perp \oplus (\sqrt{y^*})^\perp}. \tag{11}$$

The $(m+n-2)$ eigenvalues of J_{FR} are given in exact closed form by:

$$\lambda_k^\pm \;=\; -\frac{\tau_1+\tau_2}{2} \;\pm\; \sqrt{\Big(\frac{\tau_1-\tau_2}{2}\Big)^2 - \tau_1\tau_2\sigma_k^2} \qquad (k = 1, \ldots, r), \tag{12}$$

together with $-\tau_1$ (multiplicity $m-1-r$) and $-\tau_2$ (multiplicity $n-1-r$). Consequently, every mode undergoes a sharp overdamped-to-underdamped phase transition at the critical singular value

$$\sigma_c(\tau_1, \tau_2) \;:=\; \frac{|\tau_1 - \tau_2|}{2\sqrt{\tau_1\tau_2}} \;=\; \sinh\Big(\frac{1}{2}\Big|\ln\frac{\tau_1}{\tau_2}\Big|\Big)\,: \tag{13}$$

modes with $\sigma_k \le \sigma_c$ are purely real (overdamped), while modes with $\sigma_k > \sigma_c$ form complex-conjugate spiral pairs with universal decay rate $(\tau_1+\tau_2)/2$.

(b) ***Adversarial Shock-Dampening Law.*** *The map $A \mapsto (x^*(A), y^*(A))$ is C^∞ on $\mathbb{R}^{m\times n}$. Under any infinitesimal payoff perturbation $\delta A \in \mathbb{R}^{m\times n}$, define the uncoupled single-agent shock vectors $f_x := -\tau_1^{-1/2}\Pi_{x^*} D_{x^*}^{1/2}(\delta A) y^*$ and $f_y := \tau_2^{-1/2}\Pi_{y^*} D_{y^*}^{1/2}(\delta A)^\top x^*$. Then the temperature-weighted Fisher–Rao equilibrium response energy satisfies the exact cross-term-free identity*

$$\tau_1\|\delta x^*\|_{\mathrm{FR},x^*}^2 + \tau_2\|\delta y^*\|_{\mathrm{FR},y^*}^2 \;=\; f_x^\top\big(I_m + K_*K_*^\top\big)^{-1} f_x \;+\; f_y^\top\big(I_n + K_*^\top K_*\big)^{-1} f_y \tag{14}$$

$$\le\; \frac{1}{\tau_1}\,\mathrm{Var}_{x^*}\big((\delta A) y^*\big) + \frac{1}{\tau_2}\,\mathrm{Var}_{y^*}\big((\delta A)^\top x^*\big), \tag{15}$$

with spectral floor $\geq (1+\sigma_1^2)^{-1}(\|f_x\|_2^2+\|f_y\|_2^2)$ *(where* $\sigma_1 := \|K_*\|_2$*, with* $\sigma_1 = 0$ *when* $r = 0$*). Thus strategic zero-sum coupling never amplifies payoff shocks and strictly dampens every component aligned with an active interaction mode* $\sigma_k > 0$ *by* $(1+\sigma_k^2)^{-1} < 1$*.*

Proof. **(a)** At (x^*, y^*), both $\xi^* = c_1\mathbf{1}_m$ and $\zeta^* = c_2\mathbf{1}_n$ are constant vectors, so the brackets in (10) vanish identically. Differentiating (10) at (x^*, y^*) gives

$$\delta\dot{x} = -D_{x^*}(I_m - \mathbf{1}_m(x^*)^\top)(\tau_1 D_{x^*}^{-1}\delta x + A\,\delta y), \qquad \delta\dot{y} = -D_{y^*}(I_n - \mathbf{1}_n(y^*)^\top)(\tau_2 D_{y^*}^{-1}\delta y - A^\top\delta x).$$

Using $D_{x^*}^{1/2}(I_m - \mathbf{1}_m(x^*)^\top) = \Pi_{x^*}D_{x^*}^{1/2}$ and substituting $\delta x = D_{x^*}^{1/2}u$ ($u \in (\sqrt{x^*})^\perp$, so $\Pi_{x^*}u = u$) and $\delta y = D_{y^*}^{1/2}w$ ($w \in (\sqrt{y^*})^\perp$) yields (11). Let $K_* = \sum_{k=1}^r \sigma_k a_k b_k^\top$ be the compact SVD of K_* ($a_k \in (\sqrt{x^*})^\perp$, $b_k \in (\sqrt{y^*})^\perp$). Extending $\{a_k\}_{k=1}^r$ and $\{b_k\}_{k=1}^r$ to orthonormal bases of $(\sqrt{x^*})^\perp$ and $(\sqrt{y^*})^\perp$, J_{FR} block-diagonalizes into r invariant 2×2 blocks $\left[\begin{smallmatrix} -\tau_1 & -\sqrt{\tau_1\tau_2}\sigma_k \\ \sqrt{\tau_1\tau_2}\sigma_k & -\tau_2 \end{smallmatrix}\right]$—whose characteristic polynomial $\lambda^2 + (\tau_1+\tau_2)\lambda + \tau_1\tau_2(1+\sigma_k^2) = 0$ has exact roots (12)—plus diagonal entries $-\tau_1$ ($m-1-r$ times) and $-\tau_2$ ($n-1-r$ times). Setting the discriminant $(\frac{\tau_1-\tau_2}{2})^2 - \tau_1\tau_2\sigma_k^2 = 0$ yields (13).

(b) Differentiating the QRE stationarity equations $\tau_1\log(x^*/p) + Ay^* = c_1\mathbf{1}_m$ and $\tau_2\log(y^*/q) - A^\top x^* = c_2\mathbf{1}_n$ with respect to A in direction δA, multiplying by $\tau_1^{-1/2}\Pi_{x^*}D_{x^*}^{1/2}$ and $\tau_2^{-1/2}\Pi_{y^*}D_{y^*}^{1/2}$, and setting $\tilde{u} := \sqrt{\tau_1}\,D_{x^*}^{-1/2}\delta x^* \in (\sqrt{x^*})^\perp$ and $\tilde{w} := \sqrt{\tau_2}\,D_{y^*}^{-1/2}\delta y^* \in (\sqrt{y^*})^\perp$ gives the linear system

$$(I_{m+n} + S_*)\begin{pmatrix}\tilde{u}\\ \tilde{w}\end{pmatrix} = \begin{pmatrix} f_x \\ f_y\end{pmatrix}, \qquad S_* := \begin{pmatrix} 0 & K_* \\ -K_*^\top & 0\end{pmatrix}.$$

Because $S_*^\top = -S_*$ is skew-symmetric, $I + S_*$ is normal and satisfies

$$(I+S_*)^\top(I+S_*) = (I-S_*)(I+S_*) = I - S_*^2 = \begin{pmatrix} I_m + K_*K_*^\top & 0 \\ 0 & I_n + K_*^\top K_*\end{pmatrix} \succeq I_{m+n},$$

so $I + S_*$ is invertible on $\mathbb{R}^{m+n}$ with $\det(I+S_*) = \prod_{k=1}^r(1+\sigma_k^2) \geq 1$. Since $K_*\sqrt{y^*} = 0$ and $K_*^\top\sqrt{x^*} = 0$, $I + S_*$ leaves the $(m+n-2)$-dimensional constrained tangent subspace $(\sqrt{x^*})^\perp \oplus (\sqrt{y^*})^\perp$ invariant, so its restriction to this subspace is likewise invertible (justifying C^∞ smoothness by the implicit function theorem), and $\mathcal{R}_* := (I+S_*)^{-1}$ satisfies $\mathcal{R}_*^\top\mathcal{R}_* = (I - S_*^2)^{-1}$. Consequently,

$$\tau_1\|\delta x^*\|_{\mathrm{FR},x^*}^2 + \tau_2\|\delta y^*\|_{\mathrm{FR},y^*}^2 = \|\tilde{u}\|_2^2 + \|\tilde{w}\|_2^2 = \begin{pmatrix} f_x \\ f_y\end{pmatrix}^\top (I - S_*^2)^{-1}\begin{pmatrix} f_x \\ f_y\end{pmatrix},$$

which expands without cross-terms to (14). Since $(I + K_*K_*^\top)^{-1} \preceq I_m$, $(I + K_*^\top K_*)^{-1} \preceq I_n$, and $\|f_x\|_2^2 = \tau_1^{-1}\operatorname{Var}_{x^*}((\delta A)y^*)$, $\|f_y\|_2^2 = \tau_2^{-1}\operatorname{Var}_{y^*}((\delta A)^\top x^*)$, inequality (15) and the spectral floor $(1+\sigma_1^2)^{-1}$ follow immediately. □

4 Exact Rényi Nash-Gap Duality and Temperature Thermodynamics

For $u, v \in \Delta_m^\circ$, recall the Rényi ∞-divergence $D_\infty(u\|v) := \log\max_i(u_i/v_i) \geq D_{\mathrm{KL}}(u\|v)$, Shannon entropy $H(u) := -\sum_i u_i\log u_i$, and Rényi min-entropy $H_\infty(u) := -\log\max_i u_i \leq H(u)$. Let $v_0 := \min_{x\in\Delta_m}\max_{y\in\Delta_n} x^\top Ay$ denote the unregularized Nash minimax value and $\mathrm{NashGap}(x,y) := \max_{j\in[n]}(A^\top x)_j - \min_{i\in[m]}(Ay)_i$ the unregularized Nash duality gap.

Theorem 4.1 (Exact Rényi ∞-to-1 Nash-Gap Identity and Temperature Thermodynamics)**.** *Let* $(x^*, y^*) \in \Delta_m^\circ \times \Delta_n^\circ$ *be the unique QRE of* (1) *at temperatures* $\tau_1, \tau_2 > 0$*, with regularized value* $v(\tau_1,\tau_2) = \mathcal{L}(x^*, y^*)$ *and whitened interaction operator* K_* *in* (9)*.*

(a) ***Exact Zero-Error Rényi Nash-Gap Identity.*** *The unregularized Nash gap of* (x^*, y^*) *equals identically the temperature-weighted Rényi* ∞*-to-*1 *divergence defect:*

$$\mathrm{NashGap}(x^*, y^*) \;=\; \tau_1\Big(D_\infty(x^*\|p) - D_{\mathrm{KL}}(x^*\|p)\Big) \;+\; \tau_2\Big(D_\infty(y^*\|q) - D_{\mathrm{KL}}(y^*\|q)\Big). \tag{16}$$

In particular, $\mathrm{NashGap}(x^*, y^*) = 0$ *if and only if* $(x^*, y^*) = (p, q)$*. For uniform priors* $p = \frac{1}{m}\mathbf{1}_m$*,* $q = \frac{1}{n}\mathbf{1}_n$*,* (16) *simplifies identically to the Shannon-minus-min-entropy gap*

$$\mathrm{NashGap}(x^*, y^*) \;=\; \tau_1\Big(H(x^*) - H_\infty(x^*)\Big) \;+\; \tau_2\Big(H(y^*) - H_\infty(y^*)\Big) \;\le\; \tau_1 \log m + \tau_2 \log n. \tag{17}$$

(b) ***Exact Temperature Thermodynamics and Schur-Complement Heat Capacities.*** *The value function* $v(\tau_1, \tau_2)$ *is* C^∞ *on* $\mathbb{R}^2_{>0}$*, satisfies the universal two-sided value envelope*

$$-\tau_2 \log \frac{1}{q_{\min}} \;\le\; v(\tau_1, \tau_2) - v_0 \;\le\; \tau_1 \log \frac{1}{p_{\min}} \qquad (p_{\min} := \min_i p_i,\; q_{\min} := \min_j q_j), \tag{18}$$

and has exact first derivatives $\partial_{\tau_1} v = D_{\mathrm{KL}}(x^*\|p) \ge 0$*,* $\partial_{\tau_2} v = -D_{\mathrm{KL}}(y^*\|q) \le 0$*. Setting* $g_x := \Pi_{x^*} D_{x^*}^{1/2} \log(x^*/p)$ *and* $g_y := \Pi_{y^*} D_{y^*}^{1/2} \log(y^*/q)$*, the exact second derivatives are*

$$\frac{\partial^2 v}{\partial \tau_1^2} \;=\; -\frac{1}{\tau_1}\, g_x^\top \big(I_m + K_* K_*^\top\big)^{-1} g_x \;\le\; 0, \qquad \frac{\partial^2 v}{\partial \tau_2^2} \;=\; +\frac{1}{\tau_2}\, g_y^\top \big(I_n + K_*^\top K_*\big)^{-1} g_y \;\ge\; 0, \tag{19}$$

$$\frac{\partial^2 v}{\partial \tau_1 \partial \tau_2} \;=\; +\frac{1}{\sqrt{\tau_1 \tau_2}}\, g_y^\top K_*^\top \big(I_m + K_* K_*^\top\big)^{-1} g_x, \tag{20}$$

establishing that $v(\tau_1, \tau_2)$ *is globally concave in* τ_1 *and globally convex in* τ_2*.*

Proof. **(a)** From $x^* = T_1(y^*)$, $\log(x_i^*/p_i) = -\tau_1^{-1}(Ay^*)_i - \log Z_1$ for every $i \in [m]$ (where $Z_1 = \sum_{i'} p_{i'} e^{-(Ay^*)_{i'}/\tau_1}$). Taking the maximum over i yields $D_\infty(x^*\|p) = -\tau_1^{-1} \min_i (Ay^*)_i - \log Z_1$, while taking the expectation under x^* yields $D_{\mathrm{KL}}(x^*\|p) = -\tau_1^{-1}(x^*)^\top Ay^* - \log Z_1$. Subtracting gives

$$\tau_1 (D_\infty(x^*\|p) - D_{\mathrm{KL}}(x^*\|p)) \;=\; (x^*)^\top Ay^* - \min_{1 \le i \le m} (Ay^*)_i.$$

Symmetrically, $y^* = T_2(x^*)$ gives $\tau_2(D_\infty(y^*\|q) - D_{\mathrm{KL}}(y^*\|q)) = \max_{1 \le j \le n} (A^\top x^*)_j - (x^*)^\top Ay^*$. Summing the two identities eliminates $(x^*)^\top Ay^*$ and proves (16). Since $x^* \in \Delta_m^\circ$ has full support, $D_\infty(x^*\|p) = D_{\mathrm{KL}}(x^*\|p)$ iff x_i^*/p_i is constant in i, i.e., $x^* = p$ (and similarly $y^* = q$). For uniform priors, $D_\infty(x^*\|p) - D_{\mathrm{KL}}(x^*\|p) = (\log m - H_\infty(x^*)) - (\log m - H(x^*)) = H(x^*) - H_\infty(x^*)$, giving (17).

(b) Let (x_0, y_0) be any unregularized Nash equilibrium ($v_0 = x_0^\top A y_0$). Because $0 \le D_{\mathrm{KL}}(x\|p) \le \log(1/p_{\min})$ on Δ_m and $0 \le D_{\mathrm{KL}}(y\|q) \le \log(1/q_{\min})$ on Δ_n, saddle optimality gives

$$v(\tau_1, \tau_2) \le \max_{y \in \Delta_n} \mathcal{L}(x_0, y) \le \max_{y \in \Delta_n} x_0^\top A y + \tau_1 D_{\mathrm{KL}}(x_0\|p) \le v_0 + \tau_1 \log \frac{1}{p_{\min}},$$

and symmetrically $v(\tau_1, \tau_2) \ge \min_{x \in \Delta_m} \mathcal{L}(x, y_0) \ge v_0 - \tau_2 \log(1/q_{\min})$, proving (18). Smoothness of (x^*, y^*) in (τ_1, τ_2) follows from invertibility of $I + S_*$ (Theorem 3.1(b)). By the chain rule and stationarity ($\nabla_x \mathcal{L}(x^*, y^*) = (c_1 + \tau_1)\mathbf{1}_m$, $\mathbf{1}_m^\top \partial_{\tau_1} x^* = 0$, and symmetrically for y^*),

$$\frac{\partial v}{\partial \tau_1} = \frac{\partial \mathcal{L}}{\partial \tau_1}(x^*, y^*) = D_{\mathrm{KL}}(x^*\|p), \qquad \frac{\partial v}{\partial \tau_2} = \frac{\partial \mathcal{L}}{\partial \tau_2}(x^*, y^*) = -D_{\mathrm{KL}}(y^*\|q).$$

Differentiating the QRE stationarity equations with respect to τ_1 gives $\tau_1 D_{x^*}^{-1} \partial_{\tau_1} x^* + A\, \partial_{\tau_1} y^* = -\log(x^*/p) + \dot{c}_1 \mathbf{1}_m$ and $\tau_2 D_{y^*}^{-1} \partial_{\tau_1} y^* - A^\top \partial_{\tau_1} x^* = \dot{c}_2 \mathbf{1}_n$. Setting $u_1 := D_{x^*}^{-1/2} \partial_{\tau_1} x^* \in (\sqrt{x^*})^\perp$ and $w_1 := D_{y^*}^{-1/2} \partial_{\tau_1} y^* \in (\sqrt{y^*})^\perp$ yields

$$\tau_1 u_1 + \sqrt{\tau_1 \tau_2}\, K_* w_1 = -g_x, \qquad \tau_2 w_1 - \sqrt{\tau_1 \tau_2}\, K_*^\top u_1 = 0.$$

Solving the second equation gives $w_1 = \sqrt{\tau_1/\tau_2}\, K_*^\top u_1$, and substituting into the first yields

$$u_1 = -\frac{1}{\tau_1}(I_m + K_* K_*^\top)^{-1} g_x, \qquad w_1 = -\frac{1}{\sqrt{\tau_1\tau_2}}\, K_*^\top (I_m + K_* K_*^\top)^{-1} g_x.$$

Finally, since $\partial_{\tau_1} D_{\mathrm{KL}}(x^*\|p) = \langle \log(x^*/p), \partial_{\tau_1} x^* \rangle = \langle g_x, u_1 \rangle$ and $\partial_{\tau_1}(-D_{\mathrm{KL}}(y^*\|q)) = -\langle g_y, w_1 \rangle$, we obtain the first formula in (19) and (20); interchanging players yields the second formula in (19). □

5 Pythagorean Gap, KL Certificates, and Universal Damped Contraction

For any $(x, y) \in \Delta_m \times \Delta_n$, define the regularized saddle-point duality gap

$$\mathrm{Gap}_\tau(x, y) \;:=\; \max_{y' \in \Delta_n} \mathcal{L}(x, y') \;-\; \min_{x' \in \Delta_m} \mathcal{L}(x', y) \;\geq\; 0. \tag{21}$$

Theorem 5.1 (Pythagorean Gap, One-Step KL Certificate, Reverse-KL Dissipation, and Damped Contraction)**.** *Let $A \in \mathbb{R}^{m\times n}$, $p \in \Delta_m^\circ$, $q \in \Delta_n^\circ$, and $\tau_1, \tau_2 > 0$, with unique QRE (x^*, y^*) and coupling $\kappa = \Delta_{\mathrm{cr}}(A)^2/(16\tau_1\tau_2)$.*

(a) ***Exact Pythagorean Duality Gap and One-Step KL Equilibrium Certificate.*** *For every $(x, y) \in \Delta_m \times \Delta_n$,*

$$\tau_1 D_{\mathrm{KL}}(x\|x^*) + \tau_2 D_{\mathrm{KL}}(y\|y^*) \;\leq\; \mathrm{Gap}_\tau(x, y) \;=\; \tau_1 D_{\mathrm{KL}}(x\|T_1(y)) + \tau_2 D_{\mathrm{KL}}(y\|T_2(x)). \tag{22}$$

(b) ***Continuous-Time Reverse-KL Excess Dissipation.*** *Along the continuous-time smoothed best-response dynamics $\dot{x} = T_1(y) - x$, $\dot{y} = T_2(x) - y$ on $\Delta_m^\circ \times \Delta_n^\circ$,*

$$\frac{d}{dt}\mathrm{Gap}_\tau(x(t), y(t)) \;=\; -\,\mathrm{Gap}_\tau(x(t), y(t)) \;-\; \tau_1 D_{\mathrm{KL}}(T_1(y(t))\|x(t)) \;-\; \tau_2 D_{\mathrm{KL}}(T_2(x(t))\|y(t)), \tag{23}$$

implying $\mathrm{Gap}_\tau(x(t), y(t)) \leq e^{-t}\,\mathrm{Gap}_\tau(x(0), y(0))$ for every $t \geq 0$ and every $\tau_1, \tau_2 > 0$.

(c) ***Universal Discrete-Time Damped Best-Response Contraction.*** *Consider the discrete-time damped logit best-response iteration with damping factor $\eta \in (0, 1]$:*

$$x_{k+1} = (1-\eta)x_k + \eta\, T_1(y_k), \qquad y_{k+1} = (1-\eta)y_k + \eta\, T_2(x_k). \tag{24}$$

For every $k \geq 0$, every $(x_k, y_k) \in \Delta_m \times \Delta_n$, and every $\eta \in (0, 1]$,

$$\mathrm{Gap}_\tau(x_{k+1}, y_{k+1}) \;\leq\; (1 - \eta + \kappa\,\eta^2)\;\mathrm{Gap}_\tau(x_k, y_k). \tag{25}$$

In particular, choosing the optimal damping $\eta^* := 1$ for $\kappa \leq \frac{1}{2}$ (including the strategically separable case $\kappa = 0$, where $\mathrm{Gap}_\tau(x_1, y_1) = 0$) and $\eta^* := 1/(2\kappa)$ for $\kappa > \frac{1}{2}$ yields unconditional global geometric contraction at every temperature $\tau_1, \tau_2 > 0$:

$$\mathrm{Gap}_\tau(x_{k+1}, y_{k+1}) \;\leq\; \rho^*(\kappa)\;\mathrm{Gap}_\tau(x_k, y_k), \qquad \rho^*(\kappa) := \begin{cases} \kappa, & \kappa \leq \frac{1}{2}, \\ 1 - \dfrac{1}{4\kappa} = 1 - \dfrac{4\tau_1\tau_2}{\Delta_{\mathrm{cr}}(A)^2}, & \kappa > \frac{1}{2}. \end{cases} \tag{26}$$

Proof. **(a)** Because $T_2(x)$ maximizes $y' \mapsto x^\top A y' - \tau_2 D_{\mathrm{KL}}(y'\|q)$ with $\log(T_2(x)_j/q_j) = \tau_2^{-1}(A^\top x)_j - \log Z_2(x)$, we have for every $y \in \Delta_n$:

$$\max_{y' \in \Delta_n} \mathcal{L}(x, y') - \mathcal{L}(x, y) \;=\; \tau_2 \log Z_2(x) - x^\top A y + \tau_2 D_{\mathrm{KL}}(y\|q) \;=\; \tau_2 D_{\mathrm{KL}}(y\|T_2(x)).$$

Symmetrically, $\mathcal{L}(x,y) - \min_{x' \in \Delta_m} \mathcal{L}(x',y) = \tau_1 D_{\mathrm{KL}}(x \| T_1(y))$. Adding the two identities yields the equality in (22). Furthermore, applying the same two identities at $y = y^*$ (where $T_1(y^*) = x^*$) and $x = x^*$ (where $T_2(x^*) = y^*$) gives $\mathcal{L}(x,y^*) - \mathcal{L}(x^*,y^*) = \tau_1 D_{\mathrm{KL}}(x\|x^*)$ and $\mathcal{L}(x^*,y^*) - \mathcal{L}(x^*,y) = \tau_2 D_{\mathrm{KL}}(y\|y^*)$. Summing and bounding $\mathcal{L}(x,y^*) \le \max_{y'} \mathcal{L}(x,y')$ and $-\mathcal{L}(x^*,y) \le -\min_{x'} \mathcal{L}(x',y)$ establishes the inequality in (22).

(b) First note that $\Delta_m^\circ \times \Delta_n^\circ$ is positively invariant under $\dot{x} = T_1(y) - x$, $\dot{y} = T_2(x) - y$ for all $t \ge 0$: because $T_1(y)_i \ge p_{\min} e^{-2\|A\|_\infty/\tau_1} > 0$ uniformly on Δ_n (and similarly for T_2), any coordinate approaching 0 has strictly positive inward drift $\dot{x}_i > 0$, ensuring $(x(t), y(t)) \in \Delta_m^\circ \times \Delta_n^\circ$ for all $t \ge 0$. Now observe the exact additive splitting $\mathrm{Gap}_\tau(x,y) = U(x) + W(y)$, where

$$U(x) := \tau_1 D_{\mathrm{KL}}(x\|p) + \psi_2(x), \quad \psi_2(x) := \tau_2 \log \sum_{j=1}^n q_j e^{(A^\top x)_j/\tau_2},$$

and $W(y) := \tau_2 D_{\mathrm{KL}}(y\|q) + \psi_1(y)$, $\psi_1(y) := \tau_1 \log \sum_{i=1}^m p_i e^{-(Ay)_i/\tau_1}$. Since $\nabla\psi_2(x) = AT_2(x)$ and $\nabla\psi_1(y) = -A^\top T_1(y)$, differentiating $U(x(t)) + W(y(t))$ along $\dot{x} = T_1(y) - x$, $\dot{y} = T_2(x) - y$ yields

$$\begin{aligned}\frac{d}{dt}\mathrm{Gap}_\tau &= \langle \tau_1 \log(x/p) + AT_2(x),\, T_1(y) - x\rangle + \langle \tau_2 \log(y/q) - A^\top T_1(y),\, T_2(x) - y\rangle \\ &= \tau_1 \langle \log \tfrac{x}{T_1(y)},\, T_1(y) - x\rangle + \tau_2 \langle \log \tfrac{y}{T_2(x)},\, T_2(x) - y\rangle,\end{aligned}$$

because substituting $\tau_1 \log(T_1(y)/p) = -Ay + c_1 \mathbf{1}_m$ and $\tau_2 \log(T_2(x)/q) = A^\top x + c_2 \mathbf{1}_n$ cancels every bilinear A-term identically: $-(T_1(y) - x)^\top Ay + T_1(y)^\top AT_2(x) - x^\top AT_2(x) + (T_2(x) - y)^\top A^\top x - T_1(y)^\top AT_2(x) + T_1(y)^\top Ay = 0$. Finally, $\langle \log(x/T_1(y)), T_1(y) - x\rangle = -D_{\mathrm{KL}}(x\|T_1(y)) - D_{\mathrm{KL}}(T_1(y)\|x)$ (and similarly for y), yielding (23) and Grönwall decay e^{-t}.

(c) Write $d_x := T_1(y_k) - x_k$ and $d_y := T_2(x_k) - y_k$. By convexity of $D_{\mathrm{KL}}(\cdot\|p)$,

$$\tau_1 (D_{\mathrm{KL}}(x_{k+1}\|p) - D_{\mathrm{KL}}(x_k\|p)) \;\le\; \eta\, \tau_1 (D_{\mathrm{KL}}(T_1(y_k)\|p) - D_{\mathrm{KL}}(x_k\|p)).$$

For the smooth log-sum-exp potential ψ_2, Taylor's theorem gives $\psi_2(x_k + \eta d_x) - \psi_2(x_k) = \eta\langle AT_2(x_k), d_x\rangle + \frac{\eta^2}{2} d_x^\top \nabla^2\psi_2(\xi) d_x$ for some ξ on the segment $[x_k, x_{k+1}]$. Since $\nabla^2\psi_2(\xi) = \tau_2^{-1} A(\mathrm{diag}(T_2(\xi)) - T_2(\xi)T_2(\xi)^\top)A^\top$, Popoviciu's variance bound $\mathrm{Var}(Z) \le \frac{1}{4}\,\mathrm{osc}(Z)^2$, the cross-ratio bound $\mathrm{osc}(A^\top d_x) \le \Delta_{\mathrm{cr}}(A)\|d_x\|_{\mathrm{TV}}$ (from the proof of Theorem 2.3(a)), and Pinsker's inequality $2\|d_x\|_{\mathrm{TV}}^2 \le D_{\mathrm{KL}}(x_k\|T_1(y_k))$ yield

$$\begin{aligned}\frac{\eta^2}{2} d_x^\top \nabla^2\psi_2(\xi) d_x \;&=\; \frac{\eta^2}{2\tau_2}\,\mathrm{Var}_{T_2(\xi)}(A^\top d_x) \;\le\; \frac{\eta^2 \Delta_{\mathrm{cr}}(A)^2}{8\tau_2}\|d_x\|_{\mathrm{TV}}^2 \\ &\le\; \frac{\eta^2 \Delta_{\mathrm{cr}}(A)^2}{16\tau_1\tau_2} \cdot \tau_1 D_{\mathrm{KL}}(x_k\|T_1(y_k)).\end{aligned}$$

Applying the symmetric estimate to $W(y_{k+1}) - W(y_k)$ and summing, the first-order terms simplify identically via $\tau_1 D_{\mathrm{KL}}(T_1(y_k)\|p) + T_1(y_k)^\top Ay_k = -\psi_1(y_k)$ and $\tau_2 D_{\mathrm{KL}}(T_2(x_k)\|q) - x_k^\top AT_2(x_k) = -\psi_2(x_k)$ (with $T_1(y_k)^\top AT_2(x_k)$ cancelling) to $-\eta(U(x_k) + W(y_k)) = -\eta\,\mathrm{Gap}_\tau(x_k, y_k)$, while the quadratic remainders sum to $\kappa\eta^2\,\mathrm{Gap}_\tau(x_k,y_k)$ by (22), proving (25). Minimizing $1 - \eta + \kappa\eta^2$ over $\eta \in (0,1]$ yields $\eta^* = 1$ when $\kappa \le \frac{1}{2}$ (including $\kappa = 0$) and $\eta^* = 1/(2\kappa)$ when $\kappa > \frac{1}{2}$, proving (26). □

6 Computational Verification and Benchmark Tables

All four theorems were verified programmatically in `scripts/verify_all.py` and the automated test suite `tests/test_entropic_games.py` (4/4 pytest modules passing). Key numerical findings are summarized below, tabulated in Table 1, and illustrated in Figures 1 and 2.

- **Theorem 2.3 (120 random games across $(m, n) \in \{(2,2),(3,4),(5,5),(6,4)\}$ + matching pennies):** Strategic shift invariance of $\Delta_{\mathrm{cr}}(A)$ matched to 1.78×10^{-15}; zero violations of nonlinear TV contraction (6) (0.0) and Hilbert contraction (5) (0.0); matching-pennies period-2 limit-cycle TV amplitude matched (8) across $\tau \in \{1.5, 1.1, 0.9, 0.7, 0.5, 0.3\}$ to max error 1.06×10^{-15}.
- **Theorem 3.1 (80 random games with asymmetric temperatures $\tau_1, \tau_2 \in [0.3, 2.5]$, spanning 89 overdamped and 71 underdamped modes):** Closed-form Fisher–Rao eigenvalues (12) matched numerical eigenvalues of J_{FR} to 5.33×10^{-15}; exact shock-dampening identity (14) matched central-difference derivatives ($\varepsilon = 10^{-7}$) to relative error 3.74×10^{-9} with zero violations of the dampening ceiling (15) (0.0).
- **Theorem 4.1 (60 random games with non-uniform priors):** Exact Rényi ∞-to-1 Nash-gap identity (16) matched direct $\mathrm{NashGap}(x^*, y^*)$ to 1.78×10^{-15}; thermodynamic first derivatives $\partial_{\tau_1} v, \partial_{\tau_2} v$ and exact Schur-complement Hessians (19)–(20) matched central differences ($h = 10^{-6}$) to $\leq 3.92 \times 10^{-10}$.
- **Theorem 5.1 (200 damped best-response steps across the 5 canonical benchmark games in Table 1, spanning $\kappa \in [0.36, 45.70]$):** Pythagorean gap identity (22) matched to 1.55×10^{-15}; zero violations of the one-step KL certificate (0.0); continuous-time reverse-KL dissipation (23) matched to 9.46×10^{-10}; zero violations of discrete-time geometric contraction (26) (0.0), eliminating limit cycles even at $\kappa = 45.70$.

Explicit Benchmark and Ensemble Specifications (Table 1 and Figures 1–2). All five Table 1 benchmarks use uniform priors $p = \frac{1}{m}\mathbf{1}_m$, $q = \frac{1}{n}\mathbf{1}_n$: (1) **Matching Pennies (low-T):** $A = \left[\begin{smallmatrix} 1 & -1 \\ -1 & 1 \end{smallmatrix}\right]$, $\tau_1 = \tau_2 = 0.25$. (2) **Rock–Paper–Scissors (asymmetric T):** $A = \left[\begin{smallmatrix} 0 & -1 & 1 \\ 1 & 0 & -1 \\ -1 & 1 & 0 \end{smallmatrix}\right]$, $\tau_1 = 0.2$, $\tau_2 = 0.6$. (3) **Colonel Blotto (4×4):** $A = \left[\begin{smallmatrix} 0 & 1 & -1 & 2 \\ -1 & 0 & 2 & -1 \\ 1 & -2 & 0 & 1 \\ -2 & 1 & -1 & 0 \end{smallmatrix}\right]$, $\tau_1 = \tau_2 = 0.4$; Figure 2(a) initializes at $x_0 = (0.55, 0.25, 0.15, 0.05)^\top$, $y_0 = (0.10, 0.50, 0.25, 0.15)^\top$ ($\mathrm{Gap}_\tau(x_0, y_0) = 1.0152$). (4–5) **Random 5×4 (high-T and low-T):** two consecutive 5×4 i.i.d. Gaussian draws from `numpy.random.default_rng(202604)` (first draw $\sigma = 1.0$ at $\tau_1 = \tau_2 = 1.5$; second draw $\sigma = 2.0$ at $\tau_1 = 0.3, \tau_2 = 0.4$). In the random-game suites above, matrix dimensions are $(m, n) \in \{(2,2),(3,4),(5,5),(6,4)\}$ (30 each) for Theorem 2.3, $\{(2,2),(3,3),(4,3),(5,4)\}$ (20 each) for Theorem 3.1, and $\{(2,2),(3,4),(5,3),(4,4)\}$ (15 each) for Theorem 4.1; payoffs have i.i.d. $\mathcal{N}(0, \sigma^2)$ entries ($\sigma = 1.5, 1.2, 1.5$, seeded by $202601, 202602, 202603$), priors and test strategies are normalized $\mathrm{Exp}(1) + 0.05$ vectors, and temperatures are drawn uniformly on $[0.2, 3.0]$, $[0.3, 2.5]$, and $[0.3, 2.0]$, respectively.

Table 1: Benchmark invariants, optimal damping η^*, contraction factor $\rho^*(\kappa)$, leading Fisher–Rao singular value $\sigma_1(K_*)$, shock-dampening floor $(1 + \sigma_1^2)^{-1}$, and exact Rényi Nash gap across five games.

Game ($m \times n$, τ_1, τ_2)	$\Delta_{\mathrm{cr}}(A)$	κ	η^*	$\rho^*(\kappa)$	$\sigma_1(K_*)$	$(1+\sigma_1^2)^{-1}$	NashGap
Matching Pennies (2×2, 0.25, 0.25)	4.00	16.00	0.031	0.9844	4.000	0.0588	0.0000
RPS asymmetric (3×3, 0.20, 0.60)	3.00	4.69	0.107	0.9467	1.667	0.2647	0.0000
Colonel Blotto (4×4, 0.40, 0.40)	6.00	14.06	0.036	0.9822	1.858	0.2247	0.2868
Random high-T (5×4, 1.50, 1.50)	3.58	0.36	1.000	0.3552	0.263	0.9352	0.8669
Random low-T (5×4, 0.30, 0.40)	9.37	45.70	0.011	0.9945	1.637	0.2717	0.2207

7 Discussion and Open Directions

We have established a cohesive geometric, spectral, and thermodynamic framework for entropic zero-sum matrix games anchored by the cross-ratio coupling $\kappa = \Delta_{\mathrm{cr}}(A)^2/(16\tau_1\tau_2)$ and the

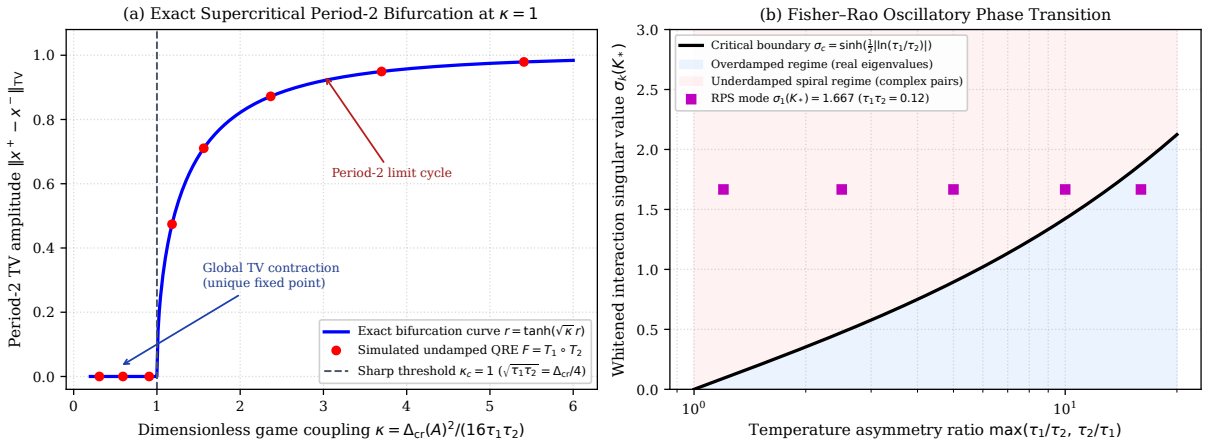


Figure 1: **(a)** Supercritical period-2 bifurcation of undamped QRE best response $F = T_1 \circ T_2$ on matching pennies at $\kappa_c = 1$: simulated TV amplitudes (red circles) match $r = \tanh(\sqrt{\kappa}\, r)$ (blue curve) to 10^{-15}. **(b)** Fisher–Rao overdamped-to-underdamped phase diagram in $(\max(\tau_1/\tau_2, \tau_2/\tau_1), \sigma_k(K_*))$, with the exact transition boundary $\sigma_c = \sinh(\frac{1}{2}|\ln(\tau_1/\tau_2)|)$ and RPS modes ($\tau_1\tau_2 = 0.12$).

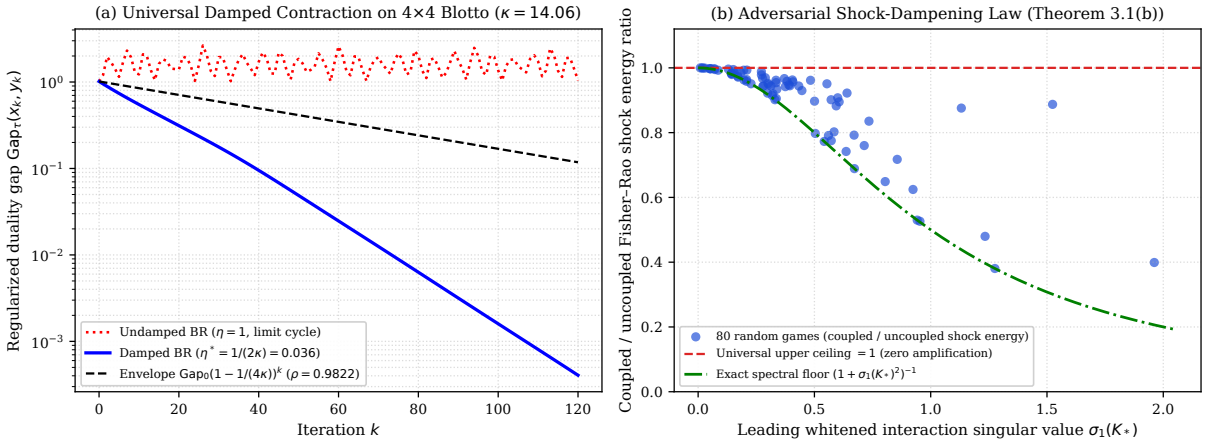


Figure 2: **(a)** Duality-gap decay on 4×4 Colonel Blotto ($\kappa = 14.06$, $\mathrm{Gap}_\tau(x_0, y_0) = 1.0152$): undamped best response ($\eta = 1$, red dotted) locks into a limit cycle, while optimal damped best response ($\eta^* = 1/(2\kappa)$, blue solid) contracts strictly below the Theorem 5.1(c) envelope $\mathrm{Gap}_\tau(x_0, y_0)(1 - 1/(4\kappa))^k$ (black dashed). **(b)** Adversarial Shock-Dampening Law (Theorem 3.1(b)) across 80 random games: every coupled/uncoupled Fisher–Rao shock energy ratio $(\tau_1\|\delta x^*\|^2_{\mathrm{FR},x^*} + \tau_2\|\delta y^*\|^2_{\mathrm{FR},y^*})/(\|f_x\|_2^2 + \|f_y\|_2^2)$ lies in the closed interval $[(1 + \sigma_1^2)^{-1}, 1]$.

Fisher–Rao interaction operator K_*. Two natural open directions follow:

1. **Multi-player polymatrix zero-sum games.** Extending the skew-symmetric normality proof of the Adversarial Shock-Dampening Law (Theorem 3.1(b)) to network zero-sum polymatrix games with graph-structured skew coupling S_*.
2. **Tightness of $\rho^*(\kappa)$ vs. alternating Gauss–Seidel damping.** Determining whether alternating damped updates (y_{k+1} using x_{k+1}) strictly improve the global contraction factor $1 - 1/(4\kappa)$ to $1 - c/\sqrt{\kappa}$ across all payoff matrices without extragradient lookahead.

Reproducibility and Open-Source Code

All numerical experiments, tables, figures, and automated tests (4/4 pytest modules) are openly archived under the MIT license at:
https://github.com/lluiseriksson/entropic-games-qre-geometry.

References

[1] G. Birkhoff, Extensions of Jentzsch's theorem, *Trans. Amer. Math. Soc.* **85** (1957), 219–227.

[2] P. J. Bushell, Hilbert's metric and positive contraction mappings in a Banach space, *Arch. Rational Mech. Anal.* **52** (1973), 330–338.

[3] S. Cen, Y. Wei, and Y. Chi, Fast policy extragradient methods for competitive games with entropy regularization, *Adv. Neural Inf. Process. Syst.* **34** (2021), 27952–27964.

[4] C. Daskalakis and I. Panageas, Last-iterate convergence: Zero-sum games and constrained min-max optimization, *ITCS* (2019), 27:1–27:18.

[5] D. Fudenberg and D. M. Kreps, Learning mixed equilibria, *Games Econom. Behav.* **5** (1993), 320–367.

[6] J. Hofbauer and E. Hopkins, Learning in perturbed asymmetric games, *Games Econom. Behav.* **52** (2005), 133–152.

[7] R. D. McKelvey and T. R. Palfrey, Quantal response equilibria for normal form games, *Games Econom. Behav.* **10** (1995), 6–38.

[8] P. Mertikopoulos and W. H. Sandholm, Learning in games via reinforcement and regularization, *Math. Oper. Res.* **41** (2016), 1297–1324.

[9] J. Perolat, R. Munos, J.-B. Lespiau, S. Omidshafiei, M. Rowland, P. Ortega, N. Burch, T. Anthony, D. Balduzzi, B. De Vylder, G. Piliouras, M. Lanctot, and K. Tuyls, From Poincaré recurrence to convergence in imperfect information games: Finding equilibrium via regularization, *ICML* (2021), 8525–8535.