

A certified two-point equilibrium range for the reciprocal-xi contest

Oliver D’Souza

18 October 2026

Abstract

We give an unconditional equilibrium characterization for Ewerhart’s reciprocal-xi contest on an explicit prize interval. Let f be the normalized inverse Fourier transform of $\xi(1/2)/\xi(1/2 + u)$ and let $G(y) = \int_0^y f(t) dt$. The unique equilibrium is zero effort for $0 < W \leq 1/f(0)$. Above this threshold and through $W_* = Y/G(Y)$, where $Y = 36591/80000$, each player uses an explicit two-point distribution on $\{0, y\}$, with $WG(y) = y$. We prove $W_* > 0.94817826$. Uniqueness holds among all Borel mixed strategies, including asymmetric profiles. The proof reduces uniqueness to a strict payoff inequality, establishes that inequality analytically for small y , and extends it to Y by a uniform finite certificate. The retained Lean project checks the density assumptions, certificate bounds, endpoint comparison, and connection to the original game. This gives a bounded-prize result without the Riemann hypothesis assumed in the known all-prize uniqueness theorem.

Keywords. Difference-form contest; mixed Nash equilibrium; reciprocal-xi density; interval arithmetic; formal verification.

1 Introduction and main result

Difference-form contests reward a player according to the difference between her effort and her opponent’s. Their equilibria need not be pure, and a symmetric game may in principle have asymmetric equilibria. For the reciprocal-xi contest, the question is particularly striking: Ewerhart proved uniqueness for every positive prize under the Riemann hypothesis. We study an explicit prize interval on which the equilibrium can instead be determined unconditionally. The equilibrium passes from zero effort to a two-point mixture, and the main task is to prove that no third effort is profitable throughout the stated interval. The reciprocal-xi specification serves here as an analytic test case for the relation between the shape of a noise density and equilibrium uniqueness. The bounded two-point branch makes that relation explicit without assuming the stronger global positivity property used in the RH-based theorem.

Two players independently choose costly efforts $x_1, x_2 \geq 0$ to compete for a prize of value $W > 0$. A symmetric random performance difference has density f and distribution function F . Player i wins with probability $F(x_i - x_j)$ and pays effort regardless of the outcome, so the payoff is

$$\Pi_i(x_i, x_j) = WF(x_i - x_j) - x_i, \quad j \neq i. \quad (1)$$

Mixed strategies are arbitrary Borel probability measures on $[0, \infty)$. All equilibrium and uniqueness statements concern this strategy space and allow asymmetric profiles. Since $0 < F(t) < 1$ at every finite t , any effort $x \geq W$ yields a negative payoff, whereas zero effort yields a positive payoff. Such

efforts are strictly dominated. A strategy with infinite mean effort has payoff $-\infty$ and cannot be an equilibrium strategy. Thus it suffices to analyze the compact action interval $[0, W]$; deviations in the original game remain covered by the same domination argument.

The density studied here is

$$\xi(s) = \frac{1}{2}s(s-1)\pi^{-s/2}\Gamma(s/2)\zeta(s), \quad (2)$$

$$h(u) = \frac{\xi(1/2)}{\xi(1/2+u)}, \quad f(t) = \frac{1}{2\pi} \int_{\mathbb{R}} h(u)e^{itu} du = \frac{1}{\pi} \int_0^\infty h(u) \cos(tu) du. \quad (3)$$

The expression for ξ is interpreted by analytic continuation at its removable singularities. The multiplier $h(u)$ is evaluated at the real argument $1/2+u$. The functional equation makes h even, and its real-axis positivity and rapid decay justify the cosine representation. The density results used below show that f is smooth, even, strictly positive, strictly decreasing on $(0, \infty)$, and integrable. Fourier inversion then gives the normalization and the identities for G and F in (4).

$$\int_{\mathbb{R}} f(t) dt = h(0) = 1, \quad G(t) = \int_0^t f(u) du, \quad F(t) = \frac{1}{2} + G(t). \quad (4)$$

In particular, G is odd and $F(-t) = 1 - F(t)$.

Theorem 1.1 (Certified reciprocal-xi range). *For (3), set $W_0 = 1/f(0)$ and $Y = 36591/80000$. The following conclusions hold without assuming RH and are verified in the retained Lean project.*

1. *If $0 < W \leq W_0$, the unique equilibrium of (1) is (δ_0, δ_0) .*
2. *If $W_0 < W \leq W_* := Y/G(Y)$, there is a unique $y > 0$ solving $WG(y) = y$, and the unique equilibrium is (μ_y, μ_y) , where*

$$\mu_y = q(y)\delta_y + (1 - q(y))\delta_0, \quad q(y) = \frac{1/W - f(y)}{f(0) - f(y)}. \quad (5)$$

3. *The second conclusion includes every*

$$W_0 < W \leq \frac{47408913}{50000000} = 0.94817826. \quad (6)$$

The rational number $Y = 36591/80000$ is the endpoint of the verified parameter range. It is not asserted to be a support-transition point. Because $W(y) = y/G(y)$ is strictly increasing, a certificate for every $0 < y \leq Y$ gives the exact prize endpoint $W_* = W(Y)$. The shorter decimal $0.94817826 = 47408913/50000000$ is a rigorously verified lower bound for W_* , so it can safely be included in the advertised prize range. The exact formula $W_0 = 1/f(0)$ defines the lower threshold; its numerical value is approximately 0.5264762120.

This is exactly the reciprocal-xi model of Ewerhart [1]. His Theorem 1 gives all-prize uniqueness conditional on RH; the zero-effort threshold is already stated in his Example 1. Ewerhart and Sun [2, eqs. (14)–(15), p. 53] give the same two-point candidate equations in the logistic case. The strict best-response criterion below is a self-contained two-support specialization of the established interchangeability method [3, Lemma A.1, p. 170]. Our contribution is the explicit unconditional reciprocal-xi range and its uniform best-response certificate, connected to the original game by a Lean-checked proof.

Positivity and logarithmic concavity of this density are also proved in the cited project *xi-log-concavity-lean* [6]. In its notation, $\mathcal{F}(z) = \xi(1/2 + iz/2)/4$, so $\mathcal{F}(0)/\mathcal{F}(iu) = h(u/2)$ and its density is $\Lambda(x) = 2f(2x)$. We reuse those analytic sources through a recorded compatibility port and prove

this change of normalization. The quantitative contest theorem does not require a new proof of the underlying density properties. Ewerhart’s typology [4] assumes real analyticity and proper all-order Pólya-frequency positivity for global uniqueness. Gröchenig’s RH criterion [5] also concerns all-order total positivity, a property not implied by strict logarithmic concavity.

Why RH enters this model. RH asserts that every nontrivial zero of ζ has real part $1/2$. For a nonzero integrable function g , all-order Pólya-frequency positivity means

$$\det[g(a_i - b_j)]_{i,j=1}^n \geq 0 \quad \text{for every } n \geq 1, \quad a_1 < \cdots < a_n, \quad b_1 < \cdots < b_n.$$

The property is called proper when all these determinants are strictly positive. Theorem 4 of Gröchenig’s 2020 preprint identifies RH with the all-order property for the inverse transform in (3), up to its positive normalizing constant [5]. Ewerhart proves that RH gives the proper property for this density; together with analyticity, it implies uniqueness for every prize [1, Lemmas 1–3]. In that argument, translation determinants control the nonsingularity of the finite equilibrium equations. Strict log-concavity controls the order-two determinants, but supplies none of the missing higher-order conditions. Our proof verifies best responses directly on a bounded prize interval. It does not prove all-order positivity, RH, or a converse from equilibrium uniqueness to RH.

Proof strategy. Indifference between zero effort and a positive effort y forces $WG(y) = y$; stationarity at y then determines the mixing weight. These identities produce a candidate, but do not rule out profitable intermediate efforts. We prove that every other effort gives a strictly smaller payoff against the candidate. The constant-sum structure then forces both players in every equilibrium to use the same two support points, and the stationarity equation identifies their weights. For the reciprocal-xi density, a local curvature bound proves the required payoff inequality for $0 < y \leq 3/20$. A finite interval certificate covers $3/20 \leq y \leq Y$. The exact endpoint comparison completes the proof of the stated numerical range. The original FLINT/Arb calculations are retained separately as a reproducibility record; the theorem’s formal certificate is the one described in Section 5.

2 The two-point candidate and uniqueness

For this section, assume only

$$\begin{aligned} f &\text{ is a continuous, even, strictly positive probability density,} \\ f &\text{ is strictly decreasing on } (0, \infty). \end{aligned} \tag{U}$$

For (3), the density results described below establish these assumptions. Strict logarithmic concavity is stronger than this reduction requires.

To see where the candidate comes from, consider a symmetric law $q\delta_y + (1 - q)\delta_0$ with $y > 0$ and $0 < q < 1$, and write P_y for the payoff against this law. By the symmetry of F , the difference between the payoffs at its two support points is

$$P_y(y) - P_y(0) = WG(y) - y,$$

independently of q . Indifference therefore forces $W = y/G(y)$. Since y is an interior action, stationarity at that point requires

$$W\{qf(0) + (1 - q)f(y)\} = 1,$$

which gives (5). These necessary conditions determine a unique candidate once $W > W_0$. The remaining question is whether all other actions yield smaller payoffs.

Lemma 2.1 (Threshold and parameterization). *Under (U), zero effort is the unique best response to every opponent strategy if $0 < W \leq W_0 = 1/f(0)$. For every $W > W_0$ there is a unique positive solution of $WG(y) = y$. Equivalently,*

$$W(y) = \frac{y}{G(y)} \quad (7)$$

is a strictly increasing bijection from $(0, \infty)$ onto (W_0, ∞) . For this solution, $0 < y < W/2$ and $0 < q(y) < 1$.

Proof. Against a probability measure ν , the pure-action payoff is

$$P_\nu(x) = W \int F(x - z) d\nu(z) - x, \quad P'_\nu(x) = W \int f(x - z) d\nu(z) - 1.$$

The bound $f \leq f(0)$ justifies differentiation under the integral for any Borel probability measure ν . Since $f(t) \leq f(0)$, $P'_\nu \leq Wf(0) - 1$. When $W < W_0$ this is negative. At $W = W_0$, equality at a point x requires $\nu = \delta_x$, because $f(t) < f(0)$ for $t \neq 0$. Thus the derivative can vanish at at most one point; its integral over every nontrivial interval is strictly negative. The payoff is strictly decreasing also at the critical value.

For $y > 0$, strict decrease and positivity give

$$f(y) < \frac{G(y)}{y} < f(0), \quad 0 < G(y) < \frac{1}{2}. \quad (8)$$

Consequently $W'(y) = (G(y) - yf(y))/G(y)^2 > 0$. The limits at zero and infinity are W_0 and infinity, respectively, because $G(y)/y \rightarrow f(0)$ and $G(y) \rightarrow 1/2$. This proves the bijection and $y < W/2$. Substitution of (8) in (5) gives $0 < q < 1$. $\square$

Fix $y > 0$, and use $W = W(y)$ and $q = q(y)$. Write the payoff against μ_y and its displacement from the zero-effort payoff as

$$P_y(x) = W\{qF(x - y) + (1 - q)F(x)\} - x, \quad (9)$$

$$D_y(x) = P_y(x) - P_y(0) \quad (10)$$

$$= W\{q[G(x - y) + G(y)] + (1 - q)G(x)\} - x. \quad (11)$$

The parameter definitions give the exact identities

$$D_y(0) = D_y(y) = 0, \quad P'_y(y) = W\{qf(0) + (1 - q)f(y)\} - 1 = 0. \quad (12)$$

These identities hold exactly for the correlated parameters y , $W(y)$, and $q(y)$.

For every $x > y$, assumption (U) and $0 < q < 1$ imply

$$P'_y(x) = W\{qf(x - y) + (1 - q)f(x)\} - 1 < W\{qf(0) + (1 - q)f(y)\} - 1 = 0. \quad (13)$$

Hence $D_y(x) < D_y(y) = 0$ for all $x > y$. The only unresolved efforts are therefore those strictly between the two support points. In particular, (H) is equivalent to $D_y(x) < 0$ for every $0 < x < y$.

Theorem 2.2 (Strict best-response criterion). *Assume (U). For fixed $y > 0$, suppose*

$$D_y(x) < 0 \quad \text{for every } x \in [0, W(y)] \setminus \{0, y\}. \quad (H)$$

Then (μ_y, μ_y) is the unique equilibrium. For every $W > 0$, an equilibrium exists even when (H) is not known.

Proof. Working on $[0, W]$ by the preceding domination argument, adding the opponent's effort to each player's payoff does not change best responses. The transformed payoffs are

$$\widehat{\Pi}_i(x_i, x_j) = WF(x_i - x_j) - x_i + x_j, \quad \widehat{\Pi}_1 + \widehat{\Pi}_2 = W.$$

The game is therefore strategically equivalent to a continuous constant-sum game on a compact product. Minimax for its mixed extension gives existence. Its equilibrium strategy sets form a Cartesian product: every maximizing strategy of player 1 can be paired with every minimizing strategy of player 2. This is the interchangeability property. Symmetry makes the constant-sum value $W/2$.

Equations (12) and (H) say that the complete best-response set to μ_y is $\{0, y\}$; hence (μ_y, μ_y) is an equilibrium. Let (ν_1, ν_2) be any equilibrium. Interchangeability gives equilibria (ν_1, μ_y) and (μ_y, ν_2) . A measure maximizing a continuous payoff is supported on its closed maximizer set, so

$$\nu_i = p_i \delta_y + (1 - p_i) \delta_0, \quad 0 \leq p_i \leq 1.$$

Moreover, μ_y is a best response to each ν_i . It gives positive mass to the interior action y , so the first-order condition against ν_i is

$$W\{p_i f(0) + (1 - p_i) f(y)\} - 1 = 0.$$

Since $f(0) > f(y)$, this equation uniquely gives $p_i = q$. Thus $\nu_1 = \nu_2 = \mu_y$. $\square$

The argument specializes the established interchangeability method to two support points [3, Lemma A.1]; see also [2, Lemma 2] and [1, proof of Lemma 1]. The remaining task is to verify (H) uniformly for the reciprocal-xi candidate over the stated parameter range.

3 The initial two-point range

The onset of the two-point branch follows from a local concavity argument. More importantly for the quantitative proof, the same argument covers a fixed initial parameter interval once a negative upper bound for f'' is available. It therefore avoids interval division by quantities that vanish as $y \downarrow 0$.

Lemma 3.1 (Local branch). *Assume (U), $f \in C^2(\mathbb{R})$, and $f''(0) < 0$. If $r > 0$ satisfies $f''(t) < 0$ for $|t| \leq r$, then (H) holds for every $0 < y \leq r$. Consequently the unique equilibrium is (5) for $W_0 < W \leq r/G(r)$.*

Proof. Put $a(x) = P'_y(x)$. From (12), $a(y) = 0$ and $\int_0^y a(x) dx = 0$. On $[0, y]$,

$$a''(x) = W\{q f''(x - y) + (1 - q) f''(x)\} < 0.$$

Thus a is strictly concave. If $a(0) \geq 0$, strict concavity makes $a(x) > 0$ for $0 < x < y$, contradicting the integral identity. Hence $a(0) < 0$, and the same identity implies that a is positive somewhere. Strict concavity gives exactly one zero $z \in (0, y)$ in addition to y , with $a < 0$ before z and $a > 0$ between z and y . The payoff first decreases, then rises to its original value at y , proving $D_y < 0$ on $(0, y)$.

Equation (13) gives (H) for $x > y$. Apply Theorem 2.2 and Lemma 2.1. $\square$

Ordinary strict log-concavity alone does not imply $f''(0) < 0$: a density proportional to e^{-t^4} illustrates the distinction. For (3), differentiation under the integral instead gives

$$f''(0) = -\frac{1}{\pi} \int_0^\infty u^2 h(u) \, du < 0.$$

For the reciprocal-xi density, write

$$L_k = \frac{1}{\pi} \int_0^\infty u^k h(u) \, du.$$

Then $f''(0) = -L_2$, $f'''(0) = 0$, and $|f^{(4)}(t)| \leq L_4$. Taylor's theorem gives

$$f''(t) \leq -L_2 + \frac{L_4 t^2}{2}.$$

The formal moment certificates establish $L_2 \geq 44$ and $L_4 \leq 3400$. Consequently, for $|t| \leq 3/20$,

$$f''(t) \leq -44 + \frac{3400}{2} \left(\frac{3}{20}\right)^2 = -\frac{23}{4} < 0.$$

Lemma 3.1 therefore proves (H) for every $0 < y \leq 3/20$, including the endpoint. This explains the initial cutoff in the formal certificate. The historical computation used a separate rescaled Taylor argument near zero; its formulas and coverage are preserved in the numerical appendix.

4 A uniform certificate between the support points

Direct interval evaluation of D_y is unsuitable near $x = 0$ and $x = y$, where D_y vanishes exactly. The certificate uses these equalities as anchors. Near zero it can propagate negativity using $P'_y < 0$; near y it uses $P''_y < 0$ and $P'_y(y) = 0$; between these regions it may certify either the payoff difference itself or a derivative with the appropriate sign. This flexibility proves the payoff inequality without requiring a global count of stationary points.

For a fixed y , choose $0 < s < b_- < y < b_+$. On $[b_-, b_+]$, certify $P''_y < 0$. On adjoining closed cells covering $[0, s]$, certify either $D_y < 0$ or $P'_y < 0$. On adjoining closed cells covering $[s, b_-]$, certify either $D_y < 0$ or $P'_y > 0$. The middle curvature condition and the exact anchors imply $D_y < 0$ on $[b_-, b_+] \setminus \{y\}$. Forward propagation from zero proves negativity on $(0, s]$; backward propagation from b_- proves it on $[s, b_-]$. Equation (13) covers every $x > y$. Thus the three finite covers imply (H). For an entire parameter interval $I = [\ell, u]$, the same proof applies provided $b_- < \ell \leq u < b_+$ and each inequality holds uniformly over both the parameter interval and its effort cell. The exact identities at 0 and the actual y are used after this uniform enclosure step.

From finite data to uniform bounds. For each parameter interval, enclosing polynomial models bound f , f' , and G throughout the required effort intervals. These are function bounds, rather than pointwise samples of the payoff. Proved quadrature and Taylor-remainder estimates connect the stored coefficients to the Fourier integral. Exact rational calculations check the signs and the joins of the cells. Enclosing W and q independently may weaken a sign test, but does not invalidate it. Once a whole rectangle of parameters and efforts is covered, the exact anchors are applied to each actual triple $(y, W(y), q(y))$.

The formal regular certificate covers $[3/20, Y]$ with 261 adjoining closed parameter intervals and 4,765 checked effort cells. Bounds for 13,601 transform samples, analytic quadrature and

Taylor remainders, and exact integration anchors supply the models’ analytic inputs. Every row is checked in Lean’s kernel; the proved soundness theorem converts its finite sign checks into the three-region cover above. The right-hand effort region is covered by (13). The exact parameter joins include both $3/20$ and Y . Appendix C.3 displays one actual parameter row, all its effort cells, and outward-rounded sign bounds. It also specifies what the stored integer intervals mean.

Proposition 4.1 (Uniform best-response inequality). *For the density in (3), condition (H) holds for every $0 < y \leq Y = 36591/80000$.*

Proof. For $0 < y \leq 3/20$, apply Lemma 3.1 and the moment bounds in Section 3. For $3/20 \leq y \leq Y$, the certified closed parameter cover and its sign-cover soundness theorem give the conclusion. The ranges meet at $3/20$ and include the endpoint Y . $\square$

Completion of the main theorem. The density satisfies (U). Lemma 2.1 proves the zero-effort regime, and it maps $(0, Y]$ bijectively to $(W_0, W_*]$. Proposition 4.1 and Theorem 2.2 therefore identify the unique equilibrium throughout the latter interval. The endpoint certificate proves

$$\frac{47408913}{50000000} < \frac{Y}{G(Y)}. \quad (14)$$

This one-sided inequality gives the inclusive decimal cap in Theorem 1.1. The complete finite data, exact analytic interfaces, and theorem declarations are included in the accompanying source project; Appendix C records the verification scope. The separate constant-sum existence argument applies for every positive prize.

5 Formal verification

The formal proof is connected to the density in (3), with its real shift and angular Fourier normalization, and to the original contest on all Borel probability measures on nonnegative efforts. It establishes normalization and assumption (U), proves the compact reduction and constant-sum argument, and verifies the function bounds and certificate semantics used above. Its main conclusion has no remaining density, best-response, or numerical-enclosure hypothesis.

The density input reuses the cited analytic sources through the documented compatibility port. The original interval-arithmetic argument in Appendix B provides a separate account of the required positivity and monotonicity: compact Fourier models are joined to a finite-contour tail estimate. The completed formal route has different analytic certificates and does not treat that historical contour calculation as an oracle.

The archive also proves the stronger pointwise inequality $(\log f)'' < 0$ on $\mathbb{R}$, strict concavity of $\log f$, and the strict order-two translation determinant. These results are distinct from the quantitative best-response certificate; the equilibrium reduction uses only (U). The source reuse, normalization, coverage, and exact formal targets are described in Appendix C.

The retained complete Lean build and dependency audit finished on 13 October 2026. This release relies on that build and subsequent source-alignment checks; it includes no new full Lean build or FLINT/Arb replay. The accompanying source project supplies the pinned sources, compatibility port, certificate data, and executable build and axiom checks. The earlier execution records are retained separately. The cited public repository [6] provides the pinned upstream analytic sources, rather than the complete contest formalization. Appendix D distinguishes the retained records and the materials required to rebuild them.

6 Further questions

The theorem determines the equilibrium through an explicit certified endpoint. Determining the first change in its support is the next natural question. Uncertified numerical diagnostics suggest a further contact near $W = 0.94822$, but a sharp transition requires a separate proof. A third effort tying the existing support payoffs would invalidate the strict criterion (H) without necessarily destroying equilibrium; a strictly profitable effort would invalidate the candidate itself. Neither event alone determines the number of equilibria beyond the certified interval. All-prize uniqueness remains outside the result proved here.

The appendices preserve the historical numerical derivation separately from the completed formal proof. Numerical subdivision counts, contour bounds, and printed decimal intervals retain the scope of the route that produced them. The dated reproduction record below is not a claim that those computations were rerun for this release.

A Historical numerical best-response certificates

This appendix preserves the original FLINT/Arb implementation, including its four-region effort cover and rescaled small-parameter calculation. These differ from the three-region formal cover and local-moment argument used in the main text. The following numerical counts and printed enclosures belong to the historical calculation.

A.1 A sign-cover criterion

The payoff is smooth in the reciprocal- ξ case. For a fixed parameter y , choose

$$0 < b_- < y < b_+ < \widehat{W}, \quad W(y) \leq \widehat{W},$$

and a split point $s \in (0, b_-)$. Suppose finite adjoining closed cells cover the four intervals below, and the indicated strict inequality or disjunction is valid throughout each cell:

Interval	Permitted certificate on each cell
$[b_-, b_+]$	$P_y'' < 0$
$[0, s]$	$D_y < 0$ or $P_y' < 0$
$[s, b_-]$	$D_y < 0$ or $P_y' > 0$
$[b_+, \widehat{W}]$	$D_y < 0$ or $P_y' < 0$

Together with (12), these conditions imply (H). Indeed, $P_y'' < 0$ and $P_y'(y) = 0$ give a strict payoff maximum at y on the middle interval and $D_y(b_{\pm}) < 0$. On the first interval, start from $D_y(0) = 0$ and proceed forward: a cell with negative payoff is immediate, and a cell with negative derivative propagates strict negativity from its left endpoint. On $[s, b_-]$ apply the analogous backward induction using $P_y' > 0$ and the negative right anchor. On the last interval use forward induction from b_+ . Shared endpoints leave no gap.

This criterion proves the payoff inequality. It does not count all stationary points or imply a single global decrease–increase–decrease derivative pattern. A successful direct $D_y < 0$ cell can contain additional stationary points without affecting the proof.

A.2 Uniform positive-parameter cells

For a parameter interval $I = [\ell, u]$ with $\ell > 0$, all function evaluations are enclosing ball evaluations over the full interval. Put $y_c = (\ell + u)/2$ and $r_I = (u - \ell)/2$. We use

$$W'(y) = \frac{G(y) - yf(y)}{G(y)^2}, \quad (15)$$

$$A(y) = G(y)/y - f(y), \quad B(y) = f(0) - f(y), \quad q(y) = A(y)/B(y), \quad (16)$$

$$A'(y) = \frac{yf(y) - G(y)}{y^2} - f'(y), \quad B'(y) = -f'(y), \quad (17)$$

$$q'(y) = \frac{A'(y)B(y) - A(y)B'(y)}{B(y)^2}. \quad (18)$$

Thus an enclosure for W is $W(y_c) + [-r_I, r_I]W'(I)$, and similarly for q . The implementation encloses the derivative magnitude before multiplication. Successful parameter checks prove $0 < q < 1$, $2y < W$, and $W < 9/8$ throughout the cell.

Set $\delta = \min(1/128, \ell/5)$, $b_- = \ell - \delta$, and $b_+ = u + \delta$. The middle interval contains every actual $y \in I$. A rigorously checked dyadic upper bound $\widehat{W}$ extends the effort cover past every actual endpoint $W(y)$. Each effort cell is accepted only when one of the required strict ball inequalities holds. Adaptive bisection refines cells that remain inconclusive; it does not interpret failure as a false mathematical inequality. Uniform intervals may lose correlation between y, W, q , which only weakens the tests. The exact anchors are then applied pointwise to the actual correlated parameters.

A.3 The parameter interval approaching zero

Direct division in q degenerates as $y \downarrow 0$. The small-parameter computation removes those powers symbolically. Let $a_j = f^{(j)}(0)/j!$ and let L_8 bound $|f^{(8)}|$ globally. Evenness removes the odd coefficients. The expansions through degree six satisfy

$$\begin{aligned} f(t) &= a_0 + a_2t^2 + a_4t^4 + a_6t^6 + R_f(t), & |R_f(t)| &\leq L_8|t|^8/8!, \\ f'(t) &= 2a_2t + 4a_4t^3 + 6a_6t^5 + R_{f'}(t), & |R_{f'}(t)| &\leq L_8|t|^7/7!, \\ G(t) &= a_0t + a_2t^3/3 + a_4t^5/5 + a_6t^7/7 + R_G(t), & |R_G(t)| &\leq L_8|t|^9/9!. \end{aligned}$$

Define continuous removable quotients

$$A_0(t) = \frac{f(t) - a_0}{t^2}, \quad B_0(t) = \frac{G(t) - a_0t}{t^3}, \quad C_0(t) = \frac{f'(t)}{t}.$$

At zero their values are $a_2, a_2/3, 2a_2$. Then

$$W = \frac{1}{a_0 + y^2B_0(y)}, \quad \varepsilon = \frac{Wa_0 - 1}{y^2} = -\frac{B_0(y)}{a_0 + y^2B_0(y)}, \quad q = \frac{B_0(y) - A_0(y)}{-A_0(y)}.$$

The computation encloses these quotients with denominators bounded away from zero, including over intervals whose lower endpoint is zero. Under $x = ys$, the sign tests use

$$\frac{D_y(ys)}{y^3} = s\varepsilon + W\{q[(s-1)^3B_0(y(s-1)) + B_0(y)] + (1-q)s^3B_0(ys)\}, \quad (19)$$

$$\frac{P'_y(ys)}{y^2} = \varepsilon + W\{q(s-1)^2A_0(y(s-1)) + (1-q)s^2A_0(ys)\}, \quad (20)$$

$$\frac{P''_y(ys)}{y} = W\{q(s-1)C_0(y(s-1)) + (1-q)sC_0(ys)\}. \quad (21)$$

For each positive y , the scaling preserves signs. The continuous limiting functions at $y = 0$ are, respectively,

$$\frac{a_2}{3a_0}s(s-1)^2, \quad \frac{a_2}{3a_0}(s-1)(3s-1), \quad \frac{2a_2}{3a_0}(3s-2),$$

and $q \rightarrow 2/3$. These limits regularize the computation; they do not create two distinct atoms at $y = 0$.

The sign-cover test uses a middle interval $s \in [4/5, 6/5]$, a split to its left, and a right interval ending at $s = 2$. Beyond $x = 2y$, monotonicity gives $P'_y(x) < Wf(y) - 1 < 0$, extending the payoff bound analytically to W . An outward enclosure of the image under squaring is used when a ball crosses zero. Absolute values in the odd-power remainder bounds are essential because $x - y$ can be negative.

A.4 Historical coverage and endpoint evidence

The recorded FLINT/Arb replay passed on the three small-parameter pieces

$$[0, 1/10], \quad [1/10, 1/8], \quad [1/8, 3/20],$$

and on 269 adjoining positive-width parameter cells covering $[1/8, Y]$, where $Y = 36591/80000$. Exact-rational readback verifies their joins. The small and regular covers overlap by $1/40$. Within each regular cell the effort sweeps were executed successfully. This historical numerical supplement does not include a standalone list of every accepted effort or rescaled-effort leaf. The later Lean archive contains its own complete regular effort-leaf certificates, described in Appendix C.

Since $W(y)$ is strictly increasing, this proves coverage through the actual endpoint $W_* = Y/G(Y)$. The repaired driver evaluates G at the exact rational Y and exports an enclosure whose outward-rounded display is

$$\begin{aligned} 0.94817826045417481115 &\leq W_* \\ &\leq 0.94817826045485639200. \end{aligned} \tag{22}$$

A separately written standard-library-only exact-rational integration of the fresh Taylor cache gives a compatible, slightly tighter interval, with endpoints

$$\begin{aligned} 0.948178260454217414514036505000860690305365385, \\ 0.948178260454813788639215319425199082484255178. \end{aligned}$$

These intervals are historical numerical enclosures; the final Lean theorem does not assert all their printed digits or their upper endpoints. Both historical lower bounds exceed the rational cap (6). The driver checks this direction with explicit exceptions, both in directed Arb arithmetic and by exact dyadic readback. The old export used an upper enclosure near 0.948178260825 as an inclusive cutoff; that longer value is invalid and is superseded. An upper endpoint for W_* cannot justify a theorem through that endpoint. This correction does not change the successful parameter sweeps.

B The original numerical density proof

This section records the finite and analytic ingredients of the original density computation. Its finite-contour estimates, displayed margins, and 450-cell compact reconstruction belong to the FLINT/Arb route. The Lean route in Appendix C establishes the same required density conclusions by different certificates. Positivity and logarithmic concavity are dependencies, not first-discovery claims.

B.1 Real-axis integrals and moment bounds

Real-axis positivity of $\xi(s)$ follows from its defining formula for $s > 1$, the functional equation for $s < 0$, and the sign of $\zeta(s)$ on $(0, 1)$; the endpoint values are $\xi(0) = \xi(1) = 1/2$. Thus $h(u) > 0$ for real u . For $U \geq 8$, put $s_U = 1/2 + U$ and

$$B(U) = \frac{\xi(1/2)\pi^{s_U/2}}{(s_U - 1)\Gamma(1 + s_U/2)}, \quad \lambda_U = \frac{1}{2} \left(\log(1 + s_U/2) - \frac{1}{1 + s_U/2} - \log \pi \right).$$

The inequalities $\zeta(s) \geq 1$ for $s > 1$ and $\psi(a) \geq \log a - 1/a$ show that, when $\lambda_U - k/U > 0$,

$$\int_U^\infty u^k h(u) du \leq \frac{B(U)U^k}{\lambda_U - k/U}. \quad (23)$$

The bound follows by exponential domination of $B(u)$ and $u^k \leq U^k e^{k(u-U)/U}$. Validated quadrature on $[0, U]$, followed by (23), encloses

$$f^{(k)}(t) = \frac{1}{\pi} \Re \left(i^k \int_0^\infty u^k h(u) e^{itu} du \right), \quad L_k = \frac{1}{\pi} \int_0^\infty u^k h(u) du.$$

In particular, $|f^{(k)}(t)| \leq L_k$. The relevant calculations take $U = 110$.

B.2 Compact Taylor models

Nine centers $c_k = (2k + 1)/16$, $0 \leq k \leq 8$, with radius $r = 1/16$ cover $[0, 9/8]$. Derivatives through order 12 are enclosed by the preceding quadrature. For $m = 0, 1, 2$, write

$$f^{(m)}(c + s) = P_m(s) + R_m(s), \quad P_m(s) = \sum_{j=0}^{12-m} \frac{f^{(j+m)}(c)}{j!} s^j,$$

where $|s| \leq r$ and $|R_m| \leq L_{13} r^{13-m}/(13-m)!$. If S_m bounds $|P_m|$ on the coarse interval, then

$$\left| f f'' - (f')^2 - (P_0 P_2 - P_1^2) \right| \leq S_0 R_2 + S_2 R_0 + R_0 R_2 + 2S_1 R_1 + R_1^2, \quad (24)$$

with R_m now denoting the nonnegative remainder bounds. Exact rational subcells of width $1/400$, 50 per coarse interval, certify $f > 0$ and $Q := f f'' - (f')^2 < 0$ throughout the compact region. A separately implemented Fraction-arithmetic reconstruction passed all 450 cells and yields the conservative bounds

$$f > 0.0000421026, \quad Q < -0.0000000008646 \quad (0 \leq t \leq 9/8).$$

These are compact bounds, not uniform strict bounds on the whole real line. Forty short printed Q balls in the primary output lose their negative sign through low-precision serialization. They cannot serve as a standalone sign certificate. The full derivative and moment enclosures permit exact reconstruction of every sign.

The contest cache uses the same derivative order and centers, with G obtained by exact polynomial integration plus the integrated remainder $L_{13} r^{14}/14!$ on each half interval. Adjacent pieces are chained from $G(0) = 0$. Even/odd symmetry extends the models to negative arguments. Actual contest evaluations lie strictly inside $[-9/8, 9/8]$.

B.3 Finite contour and infinite tail

The finite-contour computation takes height $H = 100$. Zero isolation and counting verify that precisely $K = 29$ nontrivial zeros satisfy $0 < \Im \rho < H$, all are on the critical line, and each is simple. Let

$$\rho_k = \frac{1}{2} + i\gamma_k, \quad c_k = \frac{i\xi(1/2)}{\xi'(\rho_k)}, \quad S(t) = \sum_{k=1}^{29} c_k e^{-\gamma_k t}.$$

The c_k are real. The computation verifies $c_1 > 0$, $c_2 < 0$, the required pair signs, and $\gamma_{29} < 100 < \gamma_{30}$. Zero existence and exhaustiveness use the library's isolation and count APIs; merely obtaining a value ball that contains zero would not prove existence.

Shifting the Fourier contour from the real axis to $\mathbb{R} + iH$ gives, for $t > 0$,

$$f(t) = S(t) + E_H(t), \quad |E_H^{(j)}(t)| \leq \frac{M_j}{2\pi} e^{-Ht}, \quad j = 0, 1, 2, \quad (25)$$

where

$$M_j = \int_{\mathbb{R}} |u + iH|^j |h(u + iH)| du.$$

Stirling estimates uniformly over the finite strip justify the vanishing vertical sides. The verified finite zero count accounts for every pole crossed; no assumption about zeros above height 100 is needed.

The contour integral is bounded by interval suprema on 13,920 exact dyadic cells: $[0, 20]$ at width $1/512$, $[20, 60]$ at width $1/64$, and $[60, 200]$ at width $1/8$, with symmetry covering the negative half. An analytic bound controls the rest. In detail, for $x \geq X = 200$, set $\sigma = x + 1/2$, $a = 1 + \sigma/2$, and

$$B_H(x) = \frac{\xi(1/2)\pi^{\sigma/2} \exp\{(H^2/8)(a^{-2} + a^{-1})\}}{\Gamma(a)(\sigma - 1)(2 - \zeta(\sigma))}.$$

The gamma product and $|\zeta(\sigma + iH)| \geq 2 - \zeta(\sigma) > 0$ imply $|h(x + iH)| \leq B_H(x)$. With $a_X = 1 + (X + 1/2)/2$ and $\lambda_X = (\log a_X - a_X^{-1} - \log \pi)/2$, one obtains

$$\int_X^\infty |x + iH|^j |h(x + iH)| dx \leq \frac{B_H(X)(X + H)^j}{\lambda_X - j/(X + H)}.$$

All denominators are checked positive. Conservative readable bounds from the recorded 11 October output are $M_0 < 1.651 \cdot 10^{30}$, $M_1 < 1.651 \cdot 10^{32}$, and $M_2 < 1.651 \cdot 10^{34}$; the calculation and readback use the tighter recorded bounds.

The leading exponential controls positivity. Its logarithmic curvature is zero, so the next exponential controls the leading nonzero curvature term. Here is the tail sign argument. Set $e_j(t) = M_j e^{-Ht}/(2\pi)$, $s_j(t) = \sum_k |c_k| \gamma_k^j e^{-\gamma_k t}$, and

$$L = |c_1 c_2| (\gamma_2 - \gamma_1)^2.$$

Positivity follows if

$$\mathcal{P}(t) = \sum_{k=2}^{29} \frac{|c_k|}{c_1} e^{-(\gamma_k - \gamma_1)t} + \frac{M_0}{2\pi c_1} e^{-(H - \gamma_1)t} < 1.$$

The finite part of logarithmic curvature is

$$SS'' - (S')^2 = \sum_{j < k} c_j c_k (\gamma_k - \gamma_j)^2 e^{-(\gamma_j + \gamma_k)t}.$$

Keep the leading negative pair (1, 2) and bound only the positive pairs above, dropping all other negative pairs. The error from (25) is at most

$$C(t) = s_0 e_2 + s_2 e_0 + e_0 e_2 + 2s_1 e_1 + e_1^2.$$

Thus $Q < 0$ follows if

$$\mathcal{R}(t) = \sum_{\substack{j < k \\ c_j c_k > 0}} \frac{c_j c_k (\gamma_k - \gamma_j)^2}{L} e^{-(\gamma_j + \gamma_k - \gamma_1 - \gamma_2)t} + \frac{C(t)}{L e^{-(\gamma_1 + \gamma_2)t}} < 1.$$

Every exponential rate in these normalized bounds is verified positive. It is therefore enough to check them at $t_1 = 21/20$. A separately implemented same-FLINT zero/tail reconstruction gives

$$1 - \mathcal{P}(t_1) > 0.9283585697, \quad 1 - \mathcal{R}(t_1) > 0.3212988937.$$

This proves positivity and $Q < 0$ for $t \geq 21/20$, overlapping the compact proof by $3/40$. Evenness covers negative t . Exponential decay from (25) gives $f \in L^1$ and justifies (4). Finally, $(\log f)'' = Q/f^2 < 0$ and evenness imply $(\log f)'(0) = 0$ and $f'(t) < 0$ for $t > 0$. This supplies (U) and completes the analytical dependency of Theorem 1.1.

C Formal sources and verification scope

C.1 Mathematical targets and strategy space

The formalization uses Lean 4.33.0-rc2 with mathlib pinned to revision 51e6992efd06126df61a496bebf8f49482a4e129. Its definition of f is connected to (3), including the real shift in ξ , the angular Fourier factor, and the identity $\int_{\mathbb{R}} f = 1$. The strategy space consists of all Borel probability measures on nonnegative efforts. Expected effort is extended-valued; the proof derives finite cost for equilibria and handles infinite-cost deviations explicitly. Compact reduction, constant-sum exchangeability, support concentration, and identification of the two atom masses are proved in this game.

The declaration `ContestLean.reciprocalXi_contest_main` establishes all three parts of Theorem 1.1. It includes $W = W_0$ in the zero-effort regime and $W = W_*$ in the two-point regime, uniquely identifies the positive solution of $WG(y) = y$, proves $0 < q < 1$, and quantifies over both players' arbitrary Borel strategies. The identity of the resulting measure with (5) is supplied by `reciprocalXi_branch_measure`. The separate declaration `reciprocalXi_all_positive_prizes_have_equilibrium` proves existence for every $W > 0$. Existence beyond W_* does not imply uniqueness there.

C.2 Actual analytic inputs and finite certificates

The completed theorem has no remaining assumptions of the form (U), (H), or an unproved numerical enclosure. Actual density positivity and strict decrease follow through a rebuilt heat-law and Fourier argument using three certified low zeros and a height-sixty finite theorem in the upstream normalization. The formal project does not use the historical 29-zero, height-100 contour calculation as an oracle.

Finite zeros and the positive heat trace. In the auxiliary normalization $\mathcal{F}(z) = \xi(1/2 + iz/2)/4$, let a range over zeros with positive real part, counted with multiplicity, and define

$$H(t) = \Re \sum_a e^{-a^2 t}, \quad t > 0.$$

Six certified endpoint signs and the intermediate-value theorem produce three genuine real zeros. An inverse-power-moment certificate then proves that these are the only zero occurrences with real part at most 60; in particular, they are simple. This is an exhaustion statement, not merely three intervals whose function enclosures contain zero. The value 60 is in the auxiliary coordinate: it corresponds to zeta ordinate 30, not ordinate 60. No claim is made about the critical-line location of all higher zeros.

The remaining zeros can contribute negatively to H , but a proved analytic estimate controls their total contribution. For a selected positive real low zero $0 < c \leq 51$, the actual tail estimate is

$$\Re \sum_{\Re a \geq 60} e^{-a^2 t} \geq -\frac{35}{36} e^{-c^2 t}, \quad t > 0.$$

It follows from the zero-strip bound $|\Im a| \leq 1$, the inverse-square sum bound $\sum_a |a|^{-2} \leq 70$, and the individual estimate

$$\Re e^{-a^2 t} \geq -e^{-c^2 t} e^{-(499/2400)\Re a}, \quad e^{-(499/2400)\Re a} \leq \frac{1}{72|a|^2} \quad (\Re a \geq 60).$$

The real low zeros have nonnegative heat terms and the selected one contributes $e^{-c^2 t}$. Therefore $H(t) \geq e^{-c^2 t}/36 > 0$. This is the finite-height input needed for (U); it is different from verifying RH at every height.

The rebuilt heat-law theorem constructs a probability law ν_0 on $[0, \infty)$ with Laplace transform

$$\int e^{-vs} d\nu_0(s) = \exp \left(\int_0^\infty (e^{-vt} - 1) \frac{H(t)}{t} dt \right), \quad v \geq 0.$$

Its product/transform identity gives $h(u) = \int e^{-4su^2} d\nu_0(s)$. Writing ν for the image under $s \mapsto 4s$, the Fourier bridge yields

$$f(x) = \int_{(0, \infty)} \frac{e^{-x^2/(4s)}}{\sqrt{4\pi s}} d\nu(s).$$

The proved integrability of h rules out an atom at zero and supplies the domination needed for inversion and integration of this mixture. Every component is an even, strictly positive probability density and is strictly decreasing for positive x . The mixture establishes all of (U), including normalization, without invoking global logarithmic concavity. The corresponding contest-side targets are `xiHeatNonnegative_of_height_sixty`, `xi_heatMixture_representation`, and `reciprocalXiDensity_U_of_heat_nonnegative`.

The proof encloses all 13,601 actual transform samples used by its numerical models. Kernel-checked integer streams, analytic quadrature and remainder estimates, and exact primitive anchors connect the finite data to f and G . The local-curvature argument with $L_2 \geq 44$ and $L_4 \leq 3400$ covers $0 < y \leq 3/20$, as explained in Section 3. The regular argument consists of 261 closed adjoining parameter rows covering $[3/20, Y]$, with 4,765 effort leaves and proved fixed-scale enclosure semantics. Its soundness theorem applies the three-region sign-cover criterion to the actual correlated parameters; the effort region $x > y$ is handled analytically. The two ranges meet at $3/20$ and reach the exact endpoint Y with no gaps.

From transform samples to function bounds. The finite analytic input consists of rational values v_k and p satisfying

$$|h(k/80) - v_k| \leq 2 \cdot 10^{-120} \quad (0 \leq k \leq 13600), \quad |p - \pi| \leq 10^{-150}.$$

Define the sampled cosine sum in the auxiliary density's coordinates by

$$\mathcal{Q}(z) = \frac{1}{40p} \left(\frac{v_0}{2} + \sum_{k=1}^{13600} v_k \cos(zk/40) \right).$$

Here $\Lambda(z)$ denotes the proved entire extension of the auxiliary density. A proved quadrature estimate gives $|\mathcal{Q}(z) - \Lambda(z)| < 10^{-85}$ on $|\Re z| \leq 4$, $|\Im z| \leq 1/100$. Cauchy estimates and $\Lambda(x) = 2f(2x)$ consequently give, for $|x| \leq 6$ and $n \in \mathbb{N}$,

$$\left| f^{(n)}(x) - \frac{\Re \mathcal{Q}^{(n)}(x/2)}{2^{n+1}} \right| \leq \frac{n! 100^n}{2^{n+1}} 10^{-85}.$$

For the derivative orders used below, kernel-checked integer sums enclose these sampled derivatives. At zero, $f^{(2n)}(0) = (-1)^n L_{2n}$ and the odd derivatives vanish. The moment checks give the local inputs $44 \leq L_2 \leq 45$, $L_4 \leq 3400$ and also

$$L_{12} \leq 7 \cdot 10^{12}, \quad L_{14} \leq 2.8 \cdot 10^{15}, \quad L_{13} \leq 10L_{12} + L_{14}/40 \leq 1.4 \cdot 10^{14}.$$

The certified derivatives through order 12 at the nine centers $c = (2k+1)/16$, $0 \leq k \leq 8$, give Taylor enclosures for f and f' . At offset s , their respective remainders are bounded by $L_{13}|s|^{13}/13!$ and $L_{13}|s|^{12}/12!$. Integrating the polynomial gives an enclosure for $G(c+s) - G(c)$ with remainder $L_{13}|s|^{14}/14!$. Exact primitive anchors propagate from $G(0) = 0$ through shared panel endpoints, and reflection covers negative arguments. Thus the local branch uses the second and fourth moments; the regular cover uses the centered function models and thirteenth-moment bound; and the endpoint comparison uses the primitive enclosure. The separate global log-concavity theorem is not an input to this chain.

Finally, `reciprocalXi_endpoint_enclosure` bounds the actual value $Y/G(Y)$ from below. Combining that analytic enclosure with exact rational arithmetic gives

$$\frac{47408913}{50000000} < \frac{Y}{G(Y)}.$$

This is the one-sided statement needed for the advertised inclusive cap. It does not promote every historical decimal interval in (22) to a formal conclusion.

C.3 A complete representative row

The first regular row, `xiFixedRow0`, covers

$$I = [3/20, 121/800], \quad s = 91/1280, \quad b_- = 91/640, \quad b_+ = 509/3200.$$

Thus $0 < s < b_- < 3/20 \leq y \leq 121/800 < b_+$ throughout the row. Its stored parameter enclosures are the exact terminating-decimal intervals

$$\begin{aligned} W(y) &\in [0.573318063091029355171998, 0.574153088543271512510576], \\ q(y) &\in [0.647159600603048760023597, 0.648501248977553801233101], \\ G(y) &\in [0.261616819054599617528883, 0.263448851032112783573750]. \end{aligned}$$

They hold simultaneously for every actual $y \in I$, although subsequent interval arithmetic need not retain their correlations.

A stored integer pair (L, U) denotes the closed real interval $[L/10^{24}, U/10^{24}]$. Each fixed Taylor model certifies intervals containing $f^{(j+d)}(c)/j!$ for $d = 0, 1$ and $j < 13 - d$, the integrated coefficients $f^{(j)}(c)/(j+1)!$ for $j < 13$, and an anchor enclosing $G(c)$. Horner evaluation, with outward rounding after multiplication, bounds the whole polynomial on an effort interval. For $|t - c| \leq 7/100$, the proved bound $L_{13} \leq 140000000000000$ supplies remainders at most 10^{-10} for f , 10^{-8} for f' , and 10^{-12} for G . These are enclosures of functions on intervals, not error bars on a finite set of payoff observations.

The table lists every effort cell in this row. Each enclosure holds for all $(y, x) \in I \times J$, where J is the displayed cell. Decimal endpoints in the table are exact rationals rounded outward from the retained fixed-scale output; the last column is their strictly positive distance from zero. The four first-region cells use three negative slope tests followed by a negative payoff test. The six left-region cells use positive slope tests; the two middle cells use negative curvature tests.

Effort cell J	Quantity	Enclosure	Sign margin
$[0, \frac{91}{5120}]$	P'_y	$[-0.085184, -0.031932]$	0.031932
$[\frac{91}{5120}, \frac{273}{10240}]$	P'_y	$[-0.051924, -0.021719]$	0.021719
$[\frac{273}{10240}, \frac{91}{2560}]$	P'_y	$[-0.037489, -0.010527]$	0.010527
$[\frac{91}{2560}, \frac{91}{1280}]$	D_y	$[-0.003311, -0.000135]$	0.000135
$[\frac{91}{1280}, \frac{91}{1024}]$	P'_y	$[0.001818, 0.039375]$	0.001818
$[\frac{91}{1024}, \frac{273}{2560}]$	P'_y	$[0.007541, 0.045038]$	0.007541
$[\frac{273}{2560}, \frac{637}{5120}]$	P'_y	$[0.004146, 0.044050]$	0.004146
$[\frac{637}{5120}, \frac{273}{2048}]$	P'_y	$[0.005676, 0.030215]$	0.005676
$[\frac{273}{2048}, \frac{2821}{20480}]$	P'_y	$[0.005915, 0.020868]$	0.005915
$[\frac{2821}{20480}, \frac{91}{640}]$	P'_y	$[0.002481, 0.017425]$	0.002481
$[\frac{91}{640}, \frac{241}{1600}]$	P''_y	$[-1.053462, -0.845427]$	0.845427
$[\frac{241}{1600}, \frac{509}{3200}]$	P''_y	$[-1.215786, -1.011910]$	1.011910

For the payoff cell, the checker evaluates D_y at the cell midpoint and inflates that interval by a bound for $|P'_y|$ times the half-width. It therefore keeps the effort dependence needed near the payoff maximum. All other table entries use the corresponding enclosed derivative directly. The row declaration is in `ContestLean/FixedRows/R0.lean`; `xiFixedRow0_checked` checks it, and `XiRegularRow.fixedCheck_sound` converts a checked row with valid analytic models into (H) for every parameter in its interval. The proof is precisely the three-region propagation in Section 4, followed by (13).

The displayed bounds were read back by exact integer and rational arithmetic from the retained row and model data for this revision. This finite readback does not revalidate the analytic inputs, rerun Lean, or replace access to the complete proof archive. It makes the format and strict margins of one genuine proof obligation explicit.

C.4 Global logarithmic curvature

The formal project separately proves $(\log f)''(t) < 0$ for every real t and strict concavity of $\log f$ on $\mathbb{R}$. These assertions use all 318 compact curvature panels and a proved two-exponential tail. The

compact region is $0 \leq t \leq 159/25$ in the paper's coordinates; the tail covers $|t| \geq 159/25 = 6.36$, including its boundary. Evenness covers the negative compact region.

Here is the tail estimate used in that separate result. Let $0 < a < b$ be the first two positive real zeros of $\mathcal{F}$, and let μ be the residual probability law after their two factors are removed. Positivity of its heat trace uses the third low zero in the preceding argument. Set

$$C = \frac{ab}{2(b^2 - a^2)}, \quad A = Cb \int e^{ay} d\mu(y), \quad B = Ca \int e^{by} d\mu(y),$$

$$T(x) = Ae^{-ax} - Be^{-bx}, \quad D_j = 2C \cdot 268341 (ba^j + ab^j), \quad j = 0, 1, 2.$$

The certified residual exponential-moment bound gives

$$|\Lambda^{(j)}(x) - T^{(j)}(x)| \leq D_j e^{-48x} \quad (x \geq 0, j = 0, 1, 2).$$

The leading expression has curvature $(T')^2 - TT'' = AB(b-a)^2 e^{-(a+b)x} > 0$. Writing $K_g = (g')^2 - gg''$, the error is at most

$$\begin{aligned} |K_\Lambda(x) - K_T(x)| &\leq A(D_2 + 2aD_1 + a^2D_0)e^{-(a+48)x} \\ &\quad + B(D_2 + 2bD_1 + b^2D_0)e^{-(b+48)x} \\ &\quad + (D_1^2 + D_0D_2)e^{-96x}. \end{aligned}$$

A checked logarithmic threshold makes this bound no more than $K_T(x)/2$ for $x \geq 3.18$. Hence $K_\Lambda > 0$ on that tail. The compact certificates cover the remainder, including the common boundary. This tail estimate belongs to the completed formal route; it is not the contour-error estimate in Appendix B.

For clarity, if the upstream density is $\Lambda(x) = 2f(2x)$ and

$$K_g(t) = (g'(t))^2 - g(t)g''(t),$$

then the checked dilation identity is

$$K_f(t) = \frac{1}{16} K_\Lambda(t/2).$$

Thus the upstream compact interval $[0, 3.18]$ corresponds to $[0, 6.36]$ here. These are certificate coverage endpoints, not a claimed optimal curvature threshold. The declarations `reciprocalXi_log_second_deriv_neg` and `reciprocalXi_strict_log_concavity` supply the pointwise and concavity conclusions. The additional declaration `reciprocalXi_pf2` establishes the strict order-two translation determinant; it does not assert all-order total positivity.

C.5 Audit and trust boundary

The retained complete build and dependency audit finished on 13 October 2026. All 3,422 explicit task declarations were audited using Lean's `#print axioms`; this count includes definitions, certificate data, and conditional helper lemmas, and is not a count of distinct paper conclusions. The only reported axioms are `propext`, `Classical.choice`, and `Quot.sound`. The finite checks use `decide +kernel`; candidate generation in Python is outside the trusted proof. The completed targets have no admitted proof, newly introduced project axiom, or trusted FLINT/Arb oracle.

The retained verification archive preserves the pinned sources, compatibility changes, certificates, build commands, and actual audit output. Its exact-source alignment was subsequently checked

against the frozen manuscript. The present revision documents those completed results without changing the formal sources. No new full Lean build or expensive numerical replay was performed for this editorial revision. The representative row in Appendix C.3 was separately read back by exact finite arithmetic; its source data were not changed. Formal verification establishes the stated mathematical propositions within Lean’s logic and checked definitions; it does not certify every historical implementation detail, printed numerical diagnostic, literature claim, or authorial decision.

D Historical evidence and reproduction

D.1 Historical numerical replay and finite checks

The recorded replay ran on 11 October 2026 UTC using Python 3.12.14, python-flint 0.9.0, and FLINT 3.6.0 on Linux x86-64. The contest calculation used 40 decimal digits (136 bits) and one FLINT thread. The foundation calculation used 60 digits for zeros and tail algebra, 30 for the contour suprema, and 40 for compact quadrature. A separate zero/tail implementation used 80 digits.

The unchanged foundation driver recomputed its zero data, contour bounds, quadratures, compact signs, and tail bounds from scratch and exited successfully. A fresh contest working directory initially contained no Taylor cache; the builder recomputed the cache, and both the original sweep and the repaired endpoint-export driver completed successfully. Those recorded executions do not establish that a new replay has occurred on another machine.

The independent finite checks have distinct scopes:

- A standard-library Fraction checker reconstructs all 450 compact sign cells from complete derivative and moment enclosures. It does not independently validate their quadratures.
- Another Fraction checker integrates the fresh Taylor cache at the exact endpoint and verifies the rational prize cap. It imports neither supplied numerical source nor a third-party arithmetic library.
- Exact coverage checks verify the contour grid and parameter-cell joins.
- A separately written same-FLINT supplement re-isolates 30 zeros, verifies $N(100) = 29$, checks nonzero derivative bounds, and reconstructs the signed tail estimates from fresh contour bounds. It does not independently reintegrate the contour suprema or validate FLINT itself.
- Full execution rechecks every accepted best-response predicate. The historical numerical archive does not export an independently checkable finite trace of every accepted effort leaf; the later formal archive supplies a separately generated trace.

These historical checks do not independently validate FLINT/Arb and are not themselves formal proofs. They remain a distinct evidence record alongside the completed formalization.

D.2 Accompanying source and result files

The accompanying source project contains the contest Lean proof, its complete analytic compatibility dependency, the numerical sources and mathematical result data, and the manuscript source and PDF. The root README.md gives the commands; PROOF_MAP.md connects the paper’s statements to named Lean declarations. The upstream revision cited in [6] is the analytic foundation, not by itself the complete contest proof. Source repository: <https://github.com/Oliverds321/reciprocal-x-i-contest-lean>.

The numerical files in `numerical/results/` contain the foundation and contest certificate enclosures, the Taylor coefficient table, and the separately implemented zero/tail result. The certificate files retain the mathematical fields of the recorded calculations; operational metadata is omitted. The Taylor table is unchanged. These supplied results are not a new numerical replay and do not independently validate their quadratures or FLINT/Arb. The numerical input-binding file records the exact Taylor-table and foundation-certificate inputs used by the contest driver; matching hashes establish file identity, not mathematical validity.

D.3 Typesetting and mathematical rebuilding

From the source-project root, `bash paper/build.sh` typesets this manuscript only. It does not rerun a numerical certificate or rebuild the formal proof.

For the formal proof, run `python3 scripts/check_sources.py` for the offline source and dependency checks, then `bash check.sh` for the full Lean build and actual axiom checks. The compatibility port is already applied. The bootstrap restores the exact toolchain and dependency revisions in the source project; do not replace them with newer revisions using `lake update`. A clean build needs network access, sufficient resources, and substantial computation. The source packaging has not completed a fresh full Lean build; the previously completed formal verification and the scope of its conclusions are described above.

For the numerical route, use the instructions in `numerical/README.md` and a working copy because rebuilding replaces result files. The dependency is pinned to python-flint 0.9.0; a hash-locked Linux x86-64 wheel requirement and a version requirement for other platforms are included. From `numerical/`, run the foundation driver `src/certify_f.py` and Taylor-table generator `src/tm_models.py`, deliberately refresh the input bindings with `python3 check_inputs.py -write`, and run the corrected contest driver `src/certify_contest.py`. The driver verifies the selected mathematical inputs before starting its numerical calculations.

Then run `bash run_checks.sh` for the independent compact, grid, endpoint, and parameter-coverage checks. Run `checks/check_zeros_tail.py` separately for the same-FLINT zero/tail check. Compare mathematical payloads rather than expecting result-file hashes to survive changes in runtime information or formatting. These checks retain the distinct scopes and limitations stated above; in particular, the numerical certificate does not export a complete finite trace of every accepted effort leaf.

The theorem uses only the sign-cover implications stated in this paper. A model-clipping tolerance near $9/8$ should be hardened before broader reuse; the actual contest argument ranges are safely below that boundary.

Author and AI assistance disclosure

Oliver D’Souza is the named author of this manuscript. Claude contributed substantially to the supplied mathematical development and computational materials. OpenAI GPT-6 Astra, accessed through dot, contributed substantially to analytical review, replay and finite-check work, literature comparison, formalization and source alignment, and preparation of this manuscript and source packet. The exact Claude model is not specified in the supplied record. These contributions include mathematical and computational work beyond copyediting. This manuscript does not represent completed line-by-line validation by the author or approval by an independent qualified human referee. Final author approval, attribution, and submission or licensing decisions remain with the author.

References

- [1] C. Ewerhart, *A game-theoretic implication of the Riemann hypothesis*, Mathematical Social Sciences **128** (2024), 52–59. doi:10.1016/j.mathsocsci.2024.01.007.
- [2] C. Ewerhart and G.-Z. Sun, *Equilibrium in the symmetric two-player Hirshleifer contest: Uniqueness and characterization*, Economics Letters **169** (2018), 51–54. doi:10.1016/j.econlet.2018.05.010.
- [3] C. Ewerhart, *Revenue ranking of optimally biased contests: The case of two players*, Economics Letters **157** (2017), 167–170. doi:10.1016/j.econlet.2017.05.012.
- [4] C. Ewerhart, *A typology of military conflict based on the Hirshleifer contest*, University of Zurich Department of Economics Working Paper 400, December 2021. <https://ewerhart.net/files/WP%202021%20Ewerhart%20A%20Typology%20of%20Military%20Conflict.pdf>.
- [5] K. Gröchenig, *Schoenberg’s theory of totally positive functions and the Riemann zeta function*, in S. D. Casey, M. M. Dodson, P. J. S. G. Ferreira and A. Zayed (eds.), *Sampling, Approximation, and Signal Analysis*, Applied and Numerical Harmonic Analysis, Birkhäuser, Cham (2023), pp. 193–210. doi:10.1007/978-3-031-41130-4_9. Earlier preprint: arXiv:2007.12889 (2020), <https://arxiv.org/abs/2007.12889>.
- [6] Paradigma, *xi-log-concavity-lean*, public source repository, inspected revision f807d86f642ec7a9404ef5d393dcd2c7b4ea17e0; final declaration ReciprocalXi.density_full_conclusions in ReciprocalXiDensityFull.lean. <https://github.com/paradigma-inc/xi-log-concavity-lean/tree/f807d86f642ec7a9404ef5d393dcd2c7b4ea17e0>. The required analytic closure was later rebuilt through the recorded compatibility port in the retained formal project; this reuses the cited sources and does not constitute an independent proof implementation.
- [7] F. Johansson, *Arb: Efficient arbitrary-precision midpoint-radius interval arithmetic*, IEEE Transactions on Computers **66** (2017), 1281–1292. doi:10.1109/TC.2017.2690633.
- [8] python-flint developers, *python-flint 0.9.0 documentation*, in particular the `acb.zeta_zeros` and `arb.zeta_nzeros` APIs. <https://python-flint.readthedocs.io/en/latest/>.