

A flow reduction for equilibrium uniqueness in rank-one bimatrix games

Alec Kriebel

Independent researcher

ORCID: [0009-0001-9320-500X](https://orcid.org/0009-0001-9320-500X)

13 October 2026 Version 1.0

Abstract

For explicitly listed rational bimatrix games of variable positive dimensions and exact rank $\text{rank}(A+B) = 1$, deciding whether there is exactly one mixed-strategy equilibrium pair is coNP-complete. Its nonuniqueness counterpart is NP-complete. Arbitrary degeneracy is allowed. A polynomial reduction from nonempty positive-integer PARTITION produces a known strict pure equilibrium and an additional equilibrium exactly when the source instance admits a partition. The construction applies the signed flow counter of Disser and Skutella and the auxiliary zero-sum representation of Adsul and coauthors. A comparison-preserving rational perturbation isolates every real-demand flow optimum, an injective simplex encoding transfers diagnostic arc flow to a scalar consistency equation, and a uniform polynomial-bit exact penalty controls every auxiliary minimizer. Rational support-polytope certificates establish membership, including positive-dimensional equilibrium sets. A bounded source audit found no matching proof among the inspected bodies; worldwide priority remains unknown.

1 Statement, scope, and lineage

A rational bimatrix game is given by explicitly listed matrices $A, B \in \mathbb{Q}^{m \times N}$, with $m, N \geq 1$. Both players maximize their own expected payoffs. Its equilibrium set $\mathcal{E}(A, B)$ consists of pairs $(x, y) \in \Delta_m \times \Delta_N$ satisfying the two best-response conditions. The word *unique* refers to the entire strategy pair. Define the ordinary languages

$$\begin{aligned}\text{UNIQUENESS}_1 &= \{(A, B) : \text{rank}(A+B) = 1, |\mathcal{E}(A, B)| = 1\}, \\ \text{NONUNIQUENESS}_1 &= \{(A, B) : \text{rank}(A+B) = 1, |\mathcal{E}(A, B)| > 1\}.\end{aligned}$$

Malformed encodings are rejected; rank-invalid games belong to neither language. Dimensions vary with the input, whose size includes all rational entries in binary.

Theorem 1. *UNIQUENESS₁ is coNP-complete and NONUNIQUENESS₁ is NP-complete under polynomial-time many-one reductions. Hardness holds for games supplied with a known strict pure equilibrium. The construction permits degeneracy and imposes no nondegeneracy promise.*

We prove this by reducing PARTITION on nonempty lists of positive integers $w_1, \dots, w_n$, $n \geq 1$. This restriction retains its usual NP-completeness. Every output has one specified strict pure equilibrium; it is the sole equilibrium on no instances, and has a distinct companion on yes instances. No claim is made about strong hardness, fixed dimensions, nondegenerate-promise hardness, or uniqueness of equilibrium payoffs.

The historical rank-one uniqueness conjecture appears in von Stengel’s contribution to Oberwolfach Report 38/2018, printed p. 2332, as joint work with Adsul, Garg, Mehta, and Sohoni [9]. The signed counting network, designated bridge, and parametric flow/LP activity hardness used here are due to Disser and Skutella [4, 5]. Their journal Lemmas 3.1–3.2 and Figures 1–3 control our source argument. The auxiliary zero-sum family and scalar hyperplane representation are established by Adsul and coauthors [1, 2]. We reprove the needed specialization to control every equilibrium, including degeneracy. Gilboa and Zemel already proved unrestricted bimatrix uniqueness coNP-completeness and rational certificate membership [7]; the known-default-plus-additional-witness pattern also occurs in Conitzer and Sandholm [3]. Neither of those displayed hardness constructions imposes sum rank one. Generic cost perturbation and all-real-demand uniqueness have earlier precedents, including Erickson, Fox, and Lkhamsuren [6]. LP duality, determinant bounds, and exact-penalty techniques are classical tools. The role of the present proof is their specific quantitative connection to the diagnostic bridge and complete rank-one equilibrium exhaustion.

A bounded primary-source audit through 13 October 2026 located no exact matching theorem or reduction in the inspected bodies. Global priority remains unknown: the full Murty 1980 article, exact SODA 2015 proceedings body, Ruta Mehta’s 2012 thesis, and Nikolas Patris’s 2026 thesis were not recovered, and further source coverage remains incomplete. These are priority gaps, not assumptions of the proof below. We do not assert firstness or current worldwide openness. The supplement records source editions, reading scopes, and these gaps.

2 The diagnostic flow network

Write $S = \sum_i w_i > 0$, $\alpha_i = w_i/(24S)$, $\eta = 1/(24S)$, and $\delta = \eta/5$. Thus $T = \sum_i \alpha_i = 1/24 < 1/12$, and every α_i is a positive multiple of η . For $0 \leq j < 2^n$, put

$$z_j = \sum_{i=1}^n (-1)^{\text{bit}_{i-1}(j)} \alpha_i.$$

A partition exists exactly when some $z_j = 0$; otherwise $|z_j| \geq \eta = 5\delta$.

For each sign $\sigma \in \{+1, -1\}$ make distinct nodes s_i^σ, t_i^σ for $0 \leq i \leq n$. The base arc $s_0^\sigma \rightarrow t_0^\sigma$ has capacity one and cost δ . At each level $i \geq 1$ add the following four arcs, each of capacity 2^{i-1} :

Arc	Cost
$s_i^\sigma \rightarrow s_{i-1}^\sigma, t_{i-1}^\sigma \rightarrow t_i^\sigma$	$\sigma \alpha_i / 2$
$s_i^\sigma \rightarrow t_{i-1}^\sigma, s_{i-1}^\sigma \rightarrow t_i^\sigma$	$(2^i - 1 - \sigma \alpha_i) / 2$

Add a global source s , sink t , and arcs $s \rightarrow s_n^\sigma, t_n^\sigma \rightarrow t$ of capacity 2^n and cost zero. The distinguished bridge

$$e : s_0^+ \longrightarrow t_0^-$$

has capacity one and cost zero. There are

$$V = 4n + 6, \quad p = 8n + 7, \quad F = 2^{n+1}.$$

The graph is a DAG: order s first, both families of s_i in decreasing i , both families of t_i in increasing i , and t last. Every arc increases this order. The source cut bounds the flow value by F . A zero-bridge flow attaining F puts 2^{n-1} on each of the four level- n arcs in each gadget, zero on lower-level arcs, and 2^n on each exterior arc. This is a conserved flow, including when $n = 1$. Scaling it gives a zero-bridge feasible flow at every real value in $[0, F]$.

Successive shortest paths (SSP) starts from zero and augments along a shortest simple residual s - t path by its bottleneck. An original arc of cost c_i contributes a reverse residual arc of cost $-c_i$. The initial zero flow is value-zero optimal because the original graph is acyclic, even though some cheap arc costs are negative.

Lemma 2 (Disser–Skutella source event). *Every legitimate SSP execution on this unperturbed network augments in units and uses e if and only if PARTITION is a yes instance. Any forward use raises f_e from zero to one and is immediately followed by a reverse use lowering it to zero. At a forward midpoint the flow value is $q = 2j + \frac{1}{2} \leq F - \frac{3}{2}$ and $f_e = \frac{1}{2}$.*

Proof. This is the state/two-step argument in the complete journal Lemmas 3.1–3.2 of [4]. We spell out the comparisons needed for its universal SSP quantifier. In a gadget, the cheap connectors first run the lower gadget forward; once they are exhausted, the crossed connectors run it in reverse. The recursive state j has marginal path cost $j + v^{[j]} + (-1)^j \delta$ for signed weights $v = \sigma \alpha$, and the base arc is empty at even j and full at odd j . The inductive phases are strictly separated: after the base adjustment their cost envelope is $[j - T - \delta, j + T + \delta]$, and $2(T + \delta) < 1$. Thus shortest ties do not permit a departure from this recursive state sequence.

When the two gadgets are synchronized at state j , their whole-gadget paths and the bridge path have the following costs, where $z = z_j$:

State	+ path	– path	Forward bridge	Next reverse if $z = 0$
Even j , base empty	$j + z + \delta$	$j - z + \delta$	j	$j + 2\delta$
Odd j , base full	$j + z - \delta$	$j - z - \delta$	$j - 2\delta$	j

If $z \neq 0$, one whole-gadget path is strictly cheaper than the bridge because $|z| \geq 5\delta$. Its companion augmentation restores synchronization, and the same comparisons exclude bridge use in that companion step. If $z = 0$, the forward bridge wins strictly; the next shortest path reverses it and restores synchronization at the next state. The recursive arcs and base/bridge capacities make each bottleneck one. This state induction applies to any chosen shortest path and has no special tie-breaking hypothesis. The forward use starts at value $2j$, giving the claimed midpoint, with $j \leq 2^n - 1$. $\square$

Remark 3 (Source arithmetic clarifications). The printed $i = 1$ base discussion in [4], p. 5:8, gives the second path cost as $-v_1$; its two crossed arcs give $1 - v_1$, consistent with the general state formula. Its post-base-adjustment envelope on p. 5:9 needs the additional $\pm\delta$ used above. The underlying arc formulas and strict state separation are unchanged. These are source-exposition clarifications, not a different gadget or a new source-hardness theorem. The journal body is the controlling edition; the exact SODA body was not certified in the priority audit.

3 Isolation for every real demand

Let D be the least common multiple of the reduced denominators of the original p arc costs. Enumerate the arcs deterministically, for example by sign $+$ then $-$, base then levels in increasing order and the four arcs in the table’s order, then exterior arcs, then the bridge. Set

$$\rho_i = \frac{1}{4D3^i}, \quad c'_i = c_i + \rho_i \quad (1 \leq i \leq p).$$

Lemma 4. *The perturbed SSP execution is a legitimate unperturbed execution in the same residual states. For each real $q \in [0, F]$, the perturbed fixed-value minimum-cost flow is unique. Its interpolated SSP curve is that optimum at every real value. On no instances every optimum has $f_e = 0$; on yes instances the unique optimum has $f_e = \frac{1}{2}$ at some $q \leq F - \frac{3}{2}$.*

Proof. A simple residual path has projected original-arc incidence in $\{-1, 0, 1\}^p$. A difference of two such paths has coordinates between -2 and 2 . Since $\sum_i \rho_i < 1/(8D)$, perturbing a path comparison changes it by less than $1/(4D)$. Every nonzero original comparison is in $(1/D)\mathbb{Z}$ and has magnitude at least $1/D$. Its sign persists. Consequently a perturbed shortest path is an original shortest path in that *same* residual state. Starting at zero, select that same path in the original execution; its bottleneck depends only on the common capacities and flow. Induction keeps the two states identical. Lemma 2 therefore governs the perturbed execution. Comparisons between different executions or states would not suffice.

No simple residual cycle with nonzero projected vector $\sigma \in \{-1, 0, 1\}^p$ has zero perturbed cost. A nonzero original cost has magnitude at least $1/D$ and retains its sign. If the original cost is zero, the first nonzero term of $\sum_i \sigma_i / (4D3^i)$ strictly dominates the finite tail, so the sum is nonzero. The formal two-cycle made of an original arc and its reverse has zero projection and changes no original flow.

Fix a real q and suppose $f \neq g$ were two optimal flows. The difference $g - f$ is a sign-conformal residual circulation of f : on each original arc use only the sign of $g_i - f_i$, and at most its residual capacity. Decompose it into positive multiples of simple directed cycles. At least one has nonzero projection; sign conformity excludes formal zero-projection two-cycles. Optimality of f makes every occurring cycle cost nonnegative, since otherwise a sufficiently small augmentation improves f . Their weighted sum is $c' \cdot (g - f) = 0$, so an occurring nontrivial cycle has zero cost, a contradiction. Compactness of the capacitated fixed-value feasible set supplies an optimum. This proves uniqueness for real, not just integral, q .

For completeness, interpolation is optimal rather than an integrality heuristic. If f is optimal at value q , P is a shortest residual s - t path of cost ℓ , and $0 \leq \Delta$ is at most its bottleneck, any feasible value- $(q + \Delta)$ flow g yields a residual flow $g - f$. Its path/cycle decomposition has s - t paths of total amount Δ and nonnegative-cost cycles. Each path costs at least ℓ , so $c' \cdot (g - f) \geq \Delta\ell$. The interpolation $f + \Delta\chi_P$ attains equality and is optimal. Starting at the value-zero optimum proves this along the entire SSP curve. The source event supplies the half-bridge midpoint in the yes case; uniqueness transfers its no-case zero-bridge property from the curve to *every* optimum. $\square$

Choose a topological order r and $H = 1 + \max_i |c'_i|$. Replace costs by

$$\bar{c}_{uv} = c'_{uv} + H(r(v) - r(u)) > 0.$$

For a value- q flow this adds $H(r(t) - r(s))q$, constant at fixed value. Residual s - t paths all receive the same cost shift, and residual cycles receive zero. All conclusions of Lemma 4 remain valid. From now on, an *unrestricted optimum* always means the optimum on this *post-perturbation*, *post-potential* network, before imposing the simplex restriction below.

4 An injective simplex and a uniform exact penalty

Let $L = V - 1 = 4n + 5$, $K = 1/(2L)$, and $\varepsilon = K/2 = 1/(4L)$. A nonnegative feasible DAG flow of value q decomposes into paths of at most L arcs, hence

$$\sum_i f_i \leq Lq. \tag{1}$$

There are no original circulation components; in particular $q = 0$ forces $f = 0$.

Use simplex coordinates $x = (x_0, x_*, x_1, \dots, x_p) \in \Delta_{p+2}$, and define linear functions on the simplex by

$$h = F(1 - x_0) = F\left(x_* + \sum_i x_i\right), \quad f_i = LFx_i, \quad q = h - Kf_e. \quad (2)$$

They satisfy $f \geq 0$, $0 \leq h \leq F$, $\sum_i f_i \leq Lh$, and

$$q \geq h - K \sum_i f_i \geq h/2 \geq 0. \quad (3)$$

Moreover $q > 0$ for every simplex point other than pure row 0. Conversely, these bounds on (f, h) give the unique representation

$$x_i = \frac{f_i}{LF}, \quad x_* = \frac{h}{F} - \frac{\sum_i f_i}{LF}, \quad x_0 = 1 - \frac{h}{F}. \quad (4)$$

Thus the slack coordinate creates no multiplicity for a fixed (f, h) .

Let $\mathcal{D}$ be the network incidence matrix with +1 at an arc's tail and -1 at its head, and d have +1 at s , -1 at t , zero elsewhere. Impose $f_i \leq u_i$, $\mathcal{D}f = qd$, and $q = \lambda$. Replace each equality by its two opposite inequalities, retaining all V conservation equations, to obtain

$$Rx \leq r + s\lambda.$$

Here s denotes the parameter-coefficient vector, not the network source. There are $k = p + 2V + 2$ inequalities. Define $c \in \mathbb{Q}^{p+2}$ by $c_0 = c_* = 0$, $c_i = LF\bar{c}_i$ on arc coordinates, so $c^T x = \sum_i \bar{c}_i f_i$. Consider

$$v(\lambda) = \min\{c^T x : x \in \Delta_{p+2}, Rx \leq r + s\lambda\}. \quad (5)$$

For every real $\lambda \in [0, F]$, scale the zero-bridge maximum flow. It has $h = q = \lambda$ and satisfies (1), proving feasibility and compactness. On a no instance, the unique unrestricted optimum has $f_e = 0$, so $h = q = \lambda \leq F$ and the optimum survives the restriction. Every restricted optimum is therefore that unrestricted optimum and has $f_e = 0$. On a yes instance, the half-bridge unrestricted optimum has

$$q = \lambda \leq F - \frac{3}{2}, \quad h = \lambda + K/2 < F, \quad \lambda < F - \varepsilon.$$

It obeys (1), so its unique representation is retained and remains optimal in (5).

Lemma 5 (Uniform exact penalty). *One can compute a positive integer M of polynomial binary length such that, for every real $\lambda \in [0, F]$, all minimizers over the simplex of*

$$\phi_\lambda(x) = c^T x + M \max\{0, \max_j (R_j x - r_j - s_j \lambda)\} \quad (6)$$

are precisely the feasible optima of (5).

Proof. The dual of (5) is

$$\max_{\mu, \nu} \nu - (r + s\lambda)^T \mu, \quad \mu \geq 0, \quad \nu \mathbf{1} \leq c + R^T \mu.$$

Its feasible polyhedron is rational, independent of λ , and nonempty ($\mu = 0$, $\nu \leq \min_i c_i$). It is pointed: any two-sided line has zero μ direction by $\mu \geq 0$, then zero ν direction by its upper bounds. Strong LP duality supplies an attained optimal face for every real λ in the specified interval.

A nonempty pointed polyhedron has a vertex; a vertex of this optimal face is a vertex of the whole feasible polyhedron, because a feasible decomposition of an optimum has optimal endpoints.

Let a positive integer Q clear all denominators in the *dual constraints*, and let $H_0 \geq 1$ bound the absolute integer coefficients and right-hand sides after clearing, including nonnegativity. This uses only R, c , not λ . Put $d_0 = k + 1$. Every dual vertex has d_0 independent active equations. Cramer's rule bounds each coordinate's magnitude by $d_0!H_0^{d_0}$: the denominator is a nonzero integer determinant of magnitude at least one, and expansion bounds each numerator determinant. This also applies to a lower-dimensional optimal face or redundant constraints. Thus an optimal dual vertex can always be chosen with

$$\|\mu\|_1 \leq C := k d_0! H_0^{d_0}.$$

Compute $M = C + 1$ directly, without enumerating vertices.

For any simplex point let $\Delta_x = \max\{0, \max_j(R_j x - r_j - s_j \lambda)\}$. The chosen optimal dual vertex gives

$$c^T x \geq \nu - \mu^T R x \geq v(\lambda) - C \Delta_x, \quad \phi_\lambda(x) \geq v(\lambda) + (M - C) \Delta_x.$$

Every infeasible point is strictly worse than a feasible optimum. Feasible minimizers attain exactly $v(\lambda)$, proving equality of the entire minimizer sets. Irrational λ changes only the dual objective; the same rational vertex bound is uniform. The product of constraint denominators is an admissible Q , with polynomial bit length. Thus $\log H_0$ and $\log M = O(k \log k + k \log H_0)$ are polynomial. $\square$

5 The rank-one game and all its equilibria

Adjoin a baseline penalty column $j = 0$ with $R_0 = r_0 = s_0 = 0$. For $0 \leq j \leq k$ define

$$G_j = c + M R_j^T - M r_j \mathbf{1}, \quad b_j = -M s_j, \quad A_{\cdot j} = -G_j, \quad B = -A + a b^T, \quad (7)$$

where $a_0 = -\varepsilon$ and every other row entry of a is $F - \varepsilon$. Therefore

$$a^T x = h - \varepsilon.$$

The opposite inequalities $q \leq \lambda$ and $-q \leq -\lambda$ supply entries $-M, +M$ in b . Since $a_0 \neq 0$ and $b \neq 0$, $A + B = a b^T$ has *exact* rank one. Its dimensions are

$$m = p + 2 = 8n + 9, \quad N = k + 1 = p + 2V + 3 = 16n + 22.$$

For real λ consider the auxiliary zero-sum game with row matrix

$$A_\lambda = A - \mathbf{1} \lambda b^T$$

and column matrix $-A_\lambda$. Its column- j row payoff against x is

$$x^T A_{\lambda, \cdot j} = -c^T x - M(R_j x - r_j - s_j \lambda).$$

Thus its row player's minimax-optimal strategies are exactly the minimizers of (6), including all optimal faces. Finite zero-sum minimax guarantees an optimal opposing strategy and that any optimal row strategy can be paired with any optimal column strategy.

Lemma 6 (Full scalar correspondence). *A pair (x, y) is an equilibrium of (A, B) if and only if it is an equilibrium of an auxiliary zero-sum game at a parameter satisfying $\lambda = a^T x$. Every candidate parameter lies in $[-\varepsilon, F - \varepsilon]$.*

Proof. For any y , replacing A by A_λ shifts every row payoff by the same scalar $-\lambda b^T y$. At $a^T x = \lambda$, original and auxiliary column expected payoffs agree because $x^T B = -x^T A + \lambda b^T = -x^T A_\lambda$. Both complete best-response sets therefore agree, proving both directions without support restrictions. The range follows from the convex combination of the entries of a . In particular this correspondence includes opposing-strategy degeneracy, zero probabilities, and continua. $\square$

This is the rank-one specialization of [1, Theorem 1]. In arXiv:1812.04611v4, dated 22 December 2019, the corresponding result is Theorem 7. Its Theorems 19, 20, 21 correspond to journal Theorems 4, 5, 6; journal Theorem 7 is a different structure result. The journal appeared online 4 June 2020 and in volume 69(2) in 2021.

Proposition 7. *There is exactly one crossing at negative parameters. It occurs at $\lambda = -\varepsilon$ and gives a strict pure equilibrium pair: pure row 0 and the column for $q - \lambda$.*

Proof. For $\lambda < 0$, the diagnostic penalty $q - \lambda$ is positive throughout the simplex. Since $c^T x \geq 0$ and (3) holds,

$$\phi_\lambda(x) \geq c^T x + M(q - \lambda) \geq -M\lambda.$$

Pure row 0 attains equality: $f = h = q = 0$, capacity residuals are negative, conservation residuals vanish, the baseline is zero, and the opposite parameter residual is $\lambda < 0$. Every other x has $q > 0$, so is strictly worse. Pure row 0 is therefore the unique optimal row strategy at every negative parameter. Its scalar equation forces $\lambda = -\varepsilon$.

At that point $q - \lambda$ is the unique maximizing penalty column against row 0, so the auxiliary opposing strategy is its pure column. Directly in the original game, row 0 receives payoff zero in that column, while every other row receives $-c_i - Mq_i < 0$. The star row has $q_* = F$; each arc row has $q_i = F$, except the bridge row has $q_e = F/2$. Against row 0 the diagnostic column payoff is $M\varepsilon > 0$, while baseline and conservation columns give zero, the opposite parameter column gives $-M\varepsilon$, and capacity columns give $-Mu_i < 0$. Both best responses are strict, proving uniqueness of the pair at this crossing. $\square$

For every candidate $\lambda \geq 0$, Lemma 5 applies because $\lambda \leq F - \varepsilon < F$. It forces $q = \lambda$. The scalar equation is then

$$\lambda = h - \varepsilon = q + Kf_e - \varepsilon \iff f_e = \frac{\varepsilon}{K} = \frac{1}{2}. \quad (8)$$

At $\lambda = 0$, the DAG has only the zero flow, so no such crossing exists. In a no PARTITION instance every restricted optimum has $f_e = 0$, excluding every nonnegative crossing and all opposing-strategy choices. Proposition 7 leaves exactly the strict default pair. In a yes instance the retained half-augmentation optimum satisfies (8) at $\lambda \leq F - 3/2 < F - \varepsilon$; pair its x with any auxiliary optimal y . Lemma 6 gives a second equilibrium, distinct because its bridge flow is $1/2$. This is an actual positive witness existence argument. The yes equilibrium set may be degenerate or positive-dimensional; the no-case argument has exhausted all pairs rather than one selected path.

6 Encoding and membership

The preceding reduction is a polynomial construction. Each capacity has $O(n)$ bits and F has $n + 2$ bits. Original cost denominators divide $240S$, so D has $O(\log S)$ bits; a generic denominator-product bound would also suffice. The ternary denominators add $O(p)$ bits per arc. Costs of order 2^n , the topological potential, L, K, ε , and all R, r, s, c coefficients have polynomial encoding length.

Lemma 5 computes an extremely large numerical M , but of polynomial binary length, by integer arithmetic. Explicitly listing the $m \times N$ matrices is polynomial. The exponential SSP run is used only to prove correctness; it is never executed by the reduction algorithm. This proves NP-hardness of NONUNIQUENESS₁ and, by complementing the source decision, coNP-hardness of UNIQUENESS₁.

We include the certificate proof because exact pair uniqueness with arbitrary degeneracy requires rational witnesses even when initially observed equilibria are irrational. This membership conclusion is established prior art [7, §3.13]; the explicit bounds below also clarify our encoding convention.

For nonempty sets $I \subseteq \{1, \dots, m\}$ and $J \subseteq \{1, \dots, N\}$ choose $i_0 \in I, j_0 \in J$, and define compact rational polytopes

$$\begin{aligned} X(I, J) &= \{x \in \Delta_m : x_i = 0 \ (i \notin I), \\ &\quad x^T(B_{\cdot j} - B_{\cdot j_0}) \leq 0 \text{ for all } j, \\ &\quad x^T(B_{\cdot j} - B_{\cdot j_0}) = 0 \text{ for } j \in J\}, \\ Y(I, J) &= \{y \in \Delta_N : y_j = 0 \ (j \notin J), \\ &\quad (A_{i\cdot} - A_{i_0\cdot})y \leq 0 \text{ for all } i, \\ &\quad (A_{i\cdot} - A_{i_0\cdot})y = 0 \text{ for } i \in I\}. \end{aligned}$$

Every member of $X(I, J) \times Y(I, J)$ is an equilibrium: all used actions are best responses. Every equilibrium belongs to its actual-support cell. Closed cells permit zero declared probabilities and add no false equilibria. Hence $\mathcal{E}(A, B)$ is exactly the finite union of these compact product polytopes, with no claim that the union is convex or has polynomially many cells.

Let ℓ be the total binary input length. Multiply payoffs by the product of their denominators, producing integer entries of magnitude at most 2^ℓ , and comparison coefficients of magnitude at most $H_1 = 2^{\ell+1}$. Remaining coefficients and right sides are $0, \pm 1$. At a vertex in ambient dimension d , active normals span $\mathbb{R}^d$, even for a lower-dimensional cell: otherwise a small movement in both signs along an annihilating direction preserves active constraints and the finite strict inactive inequalities, contradicting extremality. Select d independent active equations. Hadamard's bound and Cramer's rule bound each coordinate's numerator and denominator by

$$d^{d/2} H_1^d.$$

Each needs at most $\lceil d(\ell + 1) + (d/2) \log_2 d \rceil + 2$ bits, including sign. Use $d = m$ and $d = N$ for the two factors.

The vertices of a product are products of factor vertices. If the union of all nonempty cell vertices had only one pair, every nonempty cell would be that singleton, as would its union. Thus whenever the equilibrium set is nonsingleton, it contains two distinct rational vertex pairs of polynomial size, even if the initial witnesses were irrational or in a continuum. A verifier checks nonnegative unit vectors, distinctness, and the exact inequalities

$$(Ay)_i \leq x^T Ay \quad \text{for every } i, \quad (B^T x)_j \leq x^T By \quad \text{for every } j$$

for each pair. These are sufficient as well as necessary because expected payoffs are averages of the pure payoffs.

Finite-game equilibrium existence [8] makes the complement of nonuniqueness on valid rank-one inputs exactly uniqueness. Rank is polynomially checkable by exact arithmetic: for $C = A + B \neq 0$ choose a nonzero pivot $C_{i_0 j_0}$, and rank is one exactly when $C_{ij} C_{i_0 j_0} = C_{ij_0} C_{i_0 j}$ for all i, j . The ordinary complement of UNIQUENESS₁ has a polynomial-time rank-invalid (or malformed) branch and the two-pair NP certificate branch. Therefore UNIQUENESS₁ $\in$ coNP and NONUNIQUENESS₁ $\in$ NP, completing Theorem 1. The supplied-default version uses the same reduction and certificates.

7 Verification and disclosure

The quantified proof above is mathematical evidence. New, separately authored exact checks accompany the supplement: the flow family checks 37 weight vectors, 660 augmentations with half-flow certificates, 864 residual cycles in small instances, six original-choice searches, and four deliberate negative controls; its normal and optimized-mode receipts are retained. The game family constructs six complete games and checks 204 auxiliary pairs, strict defaults, dual certificates, and deliberate penalty/sign/certificate failures. These finite checks supplement the proof; they do not enumerate every equilibrium or prove the all-instance claims. They are separate from the omitted historical six-game/39-network script references. The original eleven-file theoretical edition contains no executable, fixtures, or check logs, and those historical computations were not replayed.

This unrefereed preprint was developed with extensive AI assistance in mathematical solving, source inspection, verification, and writing, including fresh independent and adversarial AI review families. These reviews are not conventional human journal peer review or proof-assistant verification. Alec Kriebel is the named author and independent researcher. The accompanying support package preserves the dated AI review history and the exact input scope of each review; its current kit state is recorded in `STATUS.json`. No journal acceptance or publication service action is implied. The authored work is licensed under Creative Commons Attribution 4.0 International; cited third-party works retain their own rights and their raw bodies are excluded from the supplement.

References

- [1] B. Adsul, J. Garg, R. Mehta, M. Sohoni, and B. von Stengel, *Fast algorithms for rank-1 bimatrix games*, Operations Research 69(2) (2021), 613–631. [doi:10.1287/opre.2020.1981](https://doi.org/10.1287/opre.2020.1981). Journal Theorems 1, 4, 5, 6 correspond to Theorems 7, 19, 20, 21 in [arXiv:1812.04611v4](https://arxiv.org/abs/1812.04611v4) (22 December 2019).
- [2] B. Adsul, J. Garg, R. Mehta, and M. Sohoni, *Rank-1 bimatrix games: a homeomorphism and a polynomial time algorithm*, Proceedings of STOC (2011), 195–204. [doi:10.1145/1993636.1993664](https://doi.org/10.1145/1993636.1993664).
- [3] V. Conitzer and T. Sandholm, *New complexity results about Nash equilibria*, Games and Economic Behavior 63(2) (2008), 621–641. [Author manuscript](#); Definition 4, Theorem 1, Corollary 2.
- [4] Y. Disser and M. Skutella, *The simplex algorithm is NP-mighty*, ACM Transactions on Algorithms 15(1), Article 5 (journal footer November 2018), 1–19. [doi:10.1145/3280847](https://doi.org/10.1145/3280847); [author-hosted journal body](#). Lemmas 3.1–3.2, Figures 1–3, Corollary 1.7. Some issue bibliographies give 2019.
- [5] Y. Disser and M. Skutella, *The simplex algorithm is NP-mighty*, Proceedings of SODA (2015), 858–872. [doi:10.1137/1.9781611973730.59](https://doi.org/10.1137/1.9781611973730.59). Historical predecessor; exact proceedings body not certified in this audit.
- [6] J. Erickson, K. Fox, and L. Lkhamsuren, *Holiest minimum-cost paths and flows in surface graphs*, Proceedings of STOC (2018). [Author manuscript](#), Section 3, Lemma 3.3 and Theorem 3.4 (real capacities and demands).
- [7] I. Gilboa and E. Zemel, *Nash and correlated equilibria: some complexity considerations*, Games and Economic Behavior 1(1) (1989), 80–93. [Author-hosted journal scan](#), Sections 3.1–3.2 and 3.13.
- [8] J. F. Nash, *Non-cooperative games*, Annals of Mathematics 54(2) (1951), 286–295.
- [9] B. von Stengel, *Algorithms for rank-1 bimatrix games*, in *New Directions in Stochastic Optimisation*, Oberwolfach Reports 38/2018, 15(3), 2303–2384, contribution p. 2332. Joint work with B. Adsul, J. Garg, R. Mehta, and M. Sohoni. [doi:10.4171/OWR/2018/38](https://doi.org/10.4171/OWR/2018/38); [primary report](#). Posted 27 August 2019.