

Reaction-minimal cusps in planar bimolecular mass-action networks: an exhaustive certified classification

14 September 2026

Abstract

A cusp of equilibria is the simplest generic mechanism by which a two-parameter family of vector fields passes between one and three nearby equilibria, and it organizes bistability in chemical reaction networks. We classify the two-species, five-reaction, rank-two bimolecular mass-action networks that admit a positive nondegenerate cusp transversely unfolded by two independent rate constants. Up to reaction permutation, positive rescaling of stoichiometric columns with reactants fixed, and exchange of the two species, there are exactly 52 such networks, and every one of them uses five distinct reactant complexes. Combined with the four-reaction exclusion of Banaji, Boros and Hofbauer, five reactions is the minimum reaction count for a rate-transverse cusp in planar bimolecular mass-action kinetics. The proof is exhaustive and certified. Necessary structural conditions reduce the 142,506 literal five-reaction sets to 60,036 eligible ones, and a reactant-labelled primitive-ray key collapses these to 9,999 mechanism classes. Every class is decided by an exact rational certificate: 9,063 determinant obstructions, 744 quadratic-degeneracy obstructions, 140 cubic-degeneracy obstructions and 52 positive templates. The algebraic classification, including source-faithful coverage of every eligible literal network, is an unconditional theorem verified in Lean 4 against Mathlib. The passage to ordinary analytic cusps is a Lyapunov–Schmidt reduction of the planar vector field that makes explicit a correction to the parameter-unfolding derivative missing from the formal predicate; the corrected minors of all 52 templates are certified by exact rational-polynomial Bézout identities. Every one of the 52 classes organizes nearby positive stable–saddle–stable triples, and we exhibit a fully rational template whose equilibrium equation reduces exactly to $\mu + \nu z - z^3$. The paper separates the Lean-verified core from the conventional arguments and exact computer-algebra checks, records the native-evaluation trust boundary, and lists all 52 representatives.

1 Introduction

1.1 Cusps and minimal networks

Under mass-action kinetics, a chemical reaction network with n species and m reactions defines a polynomial vector field on the positive orthant, parameterized by m positive rate constants. Which qualitative behaviours a network can display, and how few reactions are needed to display them, are classical questions of chemical reaction network theory [12]. For multiple positive equilibria the questions have precise answers in several settings: injectivity criteria exclude multistationarity in large classes of networks [10], parameter regions for multistationarity can be described algebraically [9], and the smallest networks with multiple or bistable equilibria have been enumerated in specific model classes [16, 19]. Inheritance results explain why small networks matter: a nondegenerate behaviour of a subnetwork is passed to any enlargement obtained by a list of admissible operations, so the small networks are the atoms from which the behaviour of larger networks can be read off [4, 5, 8].

The bifurcation that organizes bistability is the cusp. At a cusp of equilibria a two-parameter family of vector fields has a fold curve with a cusp point, and inside the cusp region an equilibrium with a simple zero eigenvalue has split into three nearby equilibria [13, 17]. For mass-action systems the relevant question is whether a cusp can occur at positive concentrations and be unfolded transversely by the rate constants themselves, since the rate constants are the only parameters the network provides. Banaji, Boros and Hofbauer classified local bifurcations in planar quadratic mass-action networks with few reactions and, in particular, proved that no full-rank planar network with four reactions admits a cusp unfolded by its rates [6]; earlier, Banaji and Boros determined the smallest bimolecular networks admitting Andronov–Hopf bifurcation [2], and three-reaction quadratic networks with oscillations were analysed in [7]. Five reactions is therefore the first count at which a planar bimolecular cusp can occur.

1.2 The result

We give an exhaustive certified classification of the five-reaction case. The model is the set of networks with two species, five distinct directed reactions between the six bimolecular complexes $0, X, Y, 2X, X + Y, 2Y$, no self-reactions, stoichiometric rank two, and independent positive rate constants. The equivalence is the natural one under which the family of vector fields is preserved: reaction permutation, positive rescaling of each stoichiometric column with its reactant fixed (which is an invertible reparameterization of the rates), and exchange of the two species. The main theorem (Theorem 2.3) states that a network in the model admits a positive nondegenerate cusp transversely unfolded by two of its rate constants if and only if its class is one of 52 listed classes. Every one of these classes has five distinct reactant complexes (Corollary 2.4), and five reactions is minimal (Corollary 2.5). The classification is a count of network classes, not of cusp points or parameter regions.

The proof is a decision procedure with certificates. Necessary conditions from the equilibrium and jet equations cut the $\binom{30}{5} = 142,506$ literal reaction sets down to 60,036 eligible ones; a reactant-labelled primitive-ray key that is a complete invariant of the equivalence (Lemma 2.2) reduces these to 9,999 mechanism classes. Each class is decided by one of three exact rational obstruction schemes or by an exact positive template. The obstruction schemes are sound for every positive equilibrium flux of the network, not for sampled rate values: a determinant certificate proves that the Jacobian cannot be singular at any positive equilibrium, a quadratic-degeneracy certificate proves that a singular Jacobian cannot have vanishing quadratic center coefficient, and a cubic-degeneracy certificate proves that the required center equations force the cubic coefficient to vanish. A positive template supplies an exact algebraic equilibrium, rates and jet vectors, with a root polynomial and a rational isolating interval.

1.3 Two levels of verification

The classification has two complementary verification levels, and we keep them distinct throughout.

The *formal core* is an unconditional theorem verified in Lean 4 [11] against Mathlib [18]: for every literal network Q in the model, an algebraic cusp certificate for Q exists if and only if Q is source-faithfully covered by one of the 52 positive classes. The algebraic certificate is the jet system of Section 3 together with a raw two-parameter rank condition. The formal theorem quantifies over the literal networks, not over a hand-chosen list of representatives, and its coverage step proves that each of the 60,036 eligible literal reaction sets is connected by an explicit equivalence to the actual stored record that classifies it. This is the part of the proof for which a machine check is most valuable: it is a large finite case analysis whose correctness depends on exact bookkeeping.

The *analytic interpretation* is a conventional local reduction. The formal predicate’s parameter row is the raw derivative $pJ_k(u)q$, whereas the Lyapunov–Schmidt reduction of the vector field produces the corrected row $pJ_k(u)q - pB(q, f_k(u))/\tau$. We prove that every ordinary transverse cusp in the model satisfies the formal predicate (Lemma 3.3), so the formal negative direction excludes ordinary cusps, and we verify the corrected minor for each of the 52 positive templates by an exact Bézout identity in $\mathbb{Q}[t]$ (Proposition 6.1). The analytic reduction and the 52 corrected-minor instances are exact mathematics plus reproducible computer algebra; they are not relabelled as Lean theorems. Section 9 tabulates which statements are formal and which are conventional, and records the trust boundary of the formal part, which uses Lean’s compiled native evaluation for finite computations.

1.4 Related and concurrent work

Banaji and Feliu’s bounds on positive equilibria in terms of the number of distinct sources [3] give an independent explanation of one of our structural reductions: a full-rank planar network with at most three sources cannot produce the three nearby isolated positive equilibria that a transverse cusp requires. Banaji’s preprint of 7 September 2026 [1] establishes a correspondence between fold and cusp conditions for the original and dual equations of a mass-action system, and its Section 6.3 announces a forthcoming classification of the same $(2, 5, 2)$ bimolecular class. We make no priority claim and do not claim agreement with a representative list that has not been compared. Our analytic conditions are derived directly from the planar vector field, and our verification levels are stated so that any future comparison can be made class by class using the keys of Appendix A. Large certificate-based formal classifications have precedents in the formal proofs of the four-colour theorem [14] and the Kepler conjecture [15]; the proof-engineering point of the present work is the separation of compact keys from large algebraic payloads in the coverage step (Section 7).

1.5 Organization

Section 2 fixes the model and the equivalence, and states the main theorem and its corollaries. Section 3 derives the local cusp criterion and relates it to the formal predicate. Section 4 proves the structural reductions and describes the finite enumeration. Section 5 presents the three obstruction schemes and their soundness. Section 6 presents the positive templates and a fully rational example. Section 7 describes the formal coverage theorem and proves the main theorem. Section 8 derives consequences: five sources are necessary, every class organizes bistability, ray geometry alone does not decide cusp existence, and cusps are inherited by admissible enlargements. Section 9 records the verification and reproducibility data, and Section 10 lists open problems. Appendix A contains the complete list of 52 representatives.

2 Networks, equivalence and the main theorem

2.1 The model

Let

$$\mathcal{C} = \{(0, 0), (1, 0), (0, 1), (2, 0), (1, 1), (0, 2)\} = \{0, X, Y, 2X, X + Y, 2Y\}$$

be the six complexes of molecularity at most two in two species. A *network* is a set of five distinct ordered pairs $a_r \rightarrow b_r$ with $a_r, b_r \in \mathcal{C}$ and $a_r \neq b_r$. Reactions are indexed by $r = 1, \dots, 5$ only to write formulas; the index order is immaterial. Set

$$N_r = b_r - a_r \in \mathbb{Z}^2, \quad N = (N_1 \ \cdots \ N_5) \in \mathbb{Z}^{2 \times 5}, \quad A = \begin{pmatrix} a_1^\top \\ \vdots \\ a_5^\top \end{pmatrix} \in \mathbb{Z}^{5 \times 2}.$$

We require rank $N = 2$; in particular both species occur in some reaction. With rate constants $k = (k_1, \dots, k_5)$, $k_r > 0$, the mass-action system is

$$\dot{x} = f(x, k) = \sum_{r=1}^5 N_r k_r x^{a_r}, \quad x^{a_r} = x_1^{a_{r1}} x_2^{a_{r2}}, \quad x \in \mathbb{R}_{>0}^2. \quad (2.1)$$

Since every complex is bimolecular, f is a quadratic polynomial in x and linear in k . Self-reactions and duplicate directed reactions are excluded, but distinct reactions may share a reactant, and no relation between rate constants and no conservation law is imposed.

The four defining conditions (reactant and product molecularity at most two, rank two, both species used) are exactly the hypotheses of the formal predicate `IsPlanarBimolecular252` in the Lean development, and the five-reaction structure with distinct reactions and no self-reactions is the type `SmallPlanarNetwork 5` (Section 7).

2.2 Equivalence

Definition 2.1. Two networks are *simply equivalent up to species exchange* if, after permuting reactions and optionally exchanging the two species, their reactants agree reaction by reaction and their stoichiometric columns satisfy $N'_r = \lambda_r N_r$ for some $\lambda_r > 0$.

The substitution $k'_r = k_r / \lambda_r$ identifies the families of vector fields of two simply equivalent networks, and species exchange is a linear change of state coordinates. Every property considered in this paper is therefore invariant under Definition 2.1. We do not quotient by anything larger: not by arbitrary dynamical equivalence, not by diagonal row rescaling of N , and not by general changes of state coordinates. The count below is a count of classes under Definition 2.1.

For a nonzero integer vector v let $\bar{v} = v / \gcd(|v_1|, |v_2|)$ be its primitive vector; the denominator is positive, so v and $-v$ have distinct primitive vectors. Define the *key* $K(Q)$ of a network Q as the lexicographically smaller of the two sorted multisets

$$\{(a_r, \bar{N}_r) : 1 \leq r \leq 5\} \quad \text{and its species-exchanged copy.} \quad (2.2)$$

Lemma 2.2. *Two networks in the model have equal keys if and only if they are equivalent in the sense of Definition 2.1.*

Proof. Positive column scaling fixes the reactant and the primitive directed vector of each reaction, reaction permutation fixes the multiset, and species exchange sends each multiset to its exchanged copy; hence equivalent networks have equal keys. Conversely, equality of the sorted multisets, after at most one species exchange, matches the reactions one by one with equal reactants and equal primitive vectors. Each integer column is a positive integer multiple of its primitive vector, so the ratio of the two multiples is the required $\lambda_r > 0$. Multiplicities are part of the multiset, so repeated labelled rays are matched correctly. $\square$

2.3 The main theorem

Here and below an *ordinary transverse cusp* of a network is a nondegenerate cusp of equilibria of the smooth family (2.1) at a positive state x^* and positive rates k^* : the Jacobian has a simple zero eigenvalue with nonzero complementary eigenvalue, the reduced equilibrium equation has a zero of order exactly three, and two of the rate constants unfold it regularly. Section 3 gives the precise local test.

Theorem 2.3 (Classification). *For a network Q in the model, the following are equivalent.*

- (i) Q admits an ordinary transverse cusp at positive concentrations and positive rate constants.

(ii) $K(Q)$ is one of the 52 keys represented by the rows of Table 4 in Appendix A.

The 52 representatives are pairwise inequivalent under Definition 2.1.

Corollary 2.4 (Five sources are necessary). *Every network in the model that admits an ordinary transverse cusp has five distinct reactant complexes. Networks with four distinct reactants, which form 8,000 of the 9,999 mechanism classes and satisfy the general lower bound of four sources, never admit one.*

Corollary 2.5 (Reaction minimality). *Five reactions is the minimum number of reactions for which a planar bimolecular mass-action network admits a cusp transversely unfolded by its rate constants, and each of the 52 classes attains it.*

Corollary 2.6 (Bistability). *Every one of the 52 classes admits positive rate constants at which the system has three positive equilibria, two of which are asymptotically stable and one a saddle, all three in a neighbourhood of a cusp point.*

The corollaries are proved in Section 8. Theorem 2.3 is proved in Section 7 after the ingredients have been assembled: the local criterion (Section 3), the structural reductions (Section 4), the soundness of the three obstruction schemes (Section 5), the positive templates (Section 6) and the formal coverage theorem (Section 7).

3 The planar quadratic cusp criterion

All derivatives in this section are evaluated at a positive equilibrium (x^*, k^*) of (2.1). Write $J = f_x$ for the Jacobian and $B = f_{xx}$ for the symmetric bilinear second derivative without a factorial factor, so that $f(x^* + \xi, k^*) = J\xi + \frac{1}{2}B(\xi, \xi)$ exactly, because f is quadratic in x . Write $f_k(u)$ and $J_k(u)$ for the derivatives of f and J in the rate direction $u \in \mathbb{R}^5$; both are linear in u because f is linear in k .

3.1 The jet system

Suppose that

$$\begin{aligned} Jq = 0, \quad pJ = 0, \quad pq = 1, \quad \tau = \text{tr } J \neq 0, \\ pB(q, q) = 0, \quad Jh = -B(q, q), \quad ph = 0, \end{aligned} \quad (3.1)$$

where p is a row vector and q, h are column vectors. The first line says that 0 is a simple eigenvalue of J with complementary eigenvalue τ ; then $P = I - qp$ is the spectral projection onto $\text{im } J = \ker p$ and $J = \tau P$. Consequently the center-correction equation $Jh = -B(q, q)$ has a solution precisely when $pB(q, q) = 0$, and then a unique solution with $ph = 0$. The second line says that the quadratic coefficient of the reduced equation vanishes and fixes the center correction.

For a rate direction $u \in \mathbb{R}^5$ define

$$a(u) = p f_k(u), \quad b(u) = p J_k(u) q - \frac{p B(q, f_k(u))}{\tau}, \quad c = \frac{1}{2} p B(q, h). \quad (3.2)$$

Proposition 3.1 (Local analytic test). *Assume (3.1). The equilibrium (x^*, k^*) is a nondegenerate cusp of equilibria of (2.1) transversely unfolded by the rate constants k_i, k_j if and only if $c \neq 0$ and*

$$\det \begin{pmatrix} a(e_i) & a(e_j) \\ b(e_i) & b(e_j) \end{pmatrix} \neq 0, \quad (3.3)$$

where e_i, e_j are the coordinate directions in rate space. The equilibrium is a cusp transversely unfolded by some pair of rates if and only if $c \neq 0$ and the 2×5 matrix with rows a and b has rank two.

Table 1: The standard equilibrium-cusp conditions of a planar vector field and their exact specialization to the quadratic mass-action system (2.1). The labels B1, B2, B4, B5 follow [1].

Standard condition	Exact specialization
B1: simple zero eigenvalue, no other eigenvalue on the imaginary axis	$\det J = 0, \tau = \text{tr } J \neq 0$; then $J = \tau(I - qp)$
B2: vanishing quadratic coefficient of the reduced equation	$pB(q, q) = 0$
Center correction	$Jh = -B(q, q), ph = 0$
B4: nonvanishing cubic coefficient	$c = \frac{1}{2}pB(q, h) \neq 0$
B5: two-parameter transversality by rates k_i, k_j	$\det$ of the columns i, j of the rows $u \mapsto pf_k(u)$ and $u \mapsto pJ_k(u)q - pB(q, f_k(u))/\tau$ is nonzero

Proof. Resolve nearby states as $x = x^* + zq + y$ with $y \in \ker p$, so that $z = p(x - x^*)$. The transverse component $Pf(x^* + zq + y, k) = 0$ of the equilibrium equation is an equation in $\ker p$ whose derivative in y at the base point is $J|_{\ker p} = \tau \text{id}$, which is invertible. By the implicit function theorem it has a unique smooth local solution $y = y(z, k)$ with $y(0, k^*) = 0$, and equilibria of (2.1) near (x^*, k^*) correspond exactly to the zeros of the reduced scalar function

$$g(z, k) = p f(x^* + zq + y(z, k), k).$$

This is the Lyapunov–Schmidt reduction [13, Chapter I]. Implicit differentiation at the base point gives

$$y_z = 0, \quad y_{zz} = h, \quad y_k(u) = -\frac{Pf_k(u)}{\tau}.$$

Indeed $\tau y_z = -PJq = 0$; $\tau y_{zz} = -PB(q, q) = -B(q, q)$ because $pB(q, q) = 0$, and $y_{zz} \in \ker p$ identifies it with h ; and $\tau y_k(u) = -Pf_k(u)$. Differentiating g , and using that f has no third derivative in x ,

$$g = g_z = g_{zz} = 0, \quad g_{zzz} = 3pB(q, y_{zz}) = 3pB(q, h) = 6c,$$

and

$$g_k(u) = p f_k(u) = a(u),$$

$$g_{zk}(u) = p J_k(u) q + p B(q, y_k(u)) = p J_k(u) q - \frac{pB(q, f_k(u))}{\tau} = b(u),$$

where the last step expands $pB(q, y_k(u)) = -\tau^{-1}(pB(q, f_k(u)) - pB(q, q)pf_k(u))$ and uses $pB(q, q) = 0$. A scalar equation $g(z, \mu) = 0$ with $g = g_z = g_{zz} = 0$ and $g_{zzz} \neq 0$ at the origin has a versal two-parameter unfolding by (k_i, k_j) exactly when the map $(k_i, k_j) \mapsto (g, g_z)$ has invertible derivative, which is (3.3) [13, Chapter III, the cusp or hysteresis normal form z^3]. Because the complementary eigenvalue τ is nonzero, the zero set of g is the full local equilibrium set of the vector field, and these are the defining conditions of a nondegenerate cusp bifurcation of equilibria [17, Section 8.2]. The last assertion follows because the pairs of columns of the 2×5 matrix $(a; b)$ span it if and only if it has rank two. $\square$

Remark 3.2 (Conventions). The normalization $pq = 1$ suffices; no Euclidean unit convention on q is needed. Replacing q by αq with $\alpha \neq 0$ forces $p \mapsto p/\alpha$ and $h \mapsto \alpha^2 h$, so $c \mapsto \alpha^2 c$: the vanishing and the sign of c are independent of the normalization. The sign of c enters only in Corollary 2.6. In the convention of [1] the center correction solves $Jh_B = B(q, q)$, so $h_B = -h$ and the cubic condition is the same. Table 1 records the correspondence with the standard cusp conditions.

3.2 The formal predicate and why it is necessary

The correction term in $b(u)$ is essential, and it is the point at which the formal and analytic levels of the paper differ. The Lean predicate used by the algebraic classification asks for an exact certificate consisting of a positive state, positive rates, vectors p, q, h and two rate directions u^1, u^2 such that the equilibrium equation holds, all of (3.1) holds, $pB(q, h) \neq 0$, and the raw 2×2 matrix

$$\begin{pmatrix} p f_k(u^1) & p f_k(u^2) \\ p J_k(u^1) q & p J_k(u^2) q \end{pmatrix} \quad (3.4)$$

has nonzero determinant. We denote this existential predicate by $\text{AlgCusp}(Q)$; it is the predicate $\text{AdmitsTransverseCusp}$ of Section 7. Nonvanishing of (3.4) for some pair of directions is equivalent to rank two of the 2×5 matrix of raw rows.

Lemma 3.3 (Necessary formal certificate). *Every ordinary transverse cusp of a network in the model satisfies $\text{AlgCusp}(Q)$.*

Proof. An ordinary transverse cusp supplies positive (x^*, k^*) , the jet system (3.1), and $pB(q, h) = 2c \neq 0$. It remains to produce two rate directions on which the raw matrix (3.4) is nonsingular. Define

$$u_r = k_r^* \sum_{\ell=1}^2 a_{r\ell} \frac{h_\ell}{x_\ell^*} \quad (r = 1, \dots, 5), \quad w_\ell = \frac{q_\ell h_\ell}{x_\ell^*} \quad (\ell = 1, 2).$$

Termwise differentiation of (2.1) gives the identities

$$f_k(u) = Jh, \quad J_k(u)q = B(q, h) + Jw.$$

The first holds because $\sum_r N_r u_r (x^*)^{a_r} = \sum_r N_r k_r^* (x^*)^{a_r} \sum_\ell a_{r\ell} h_\ell / x_\ell^* = Jh$; the second follows by differentiating the first identity in the state direction. Hence $p f_k(u) = p Jh = 0$ and $p J_k(u)q = pB(q, h) \neq 0$. Since $\text{rank } N = 2$, the row $u \mapsto p f_k(u) = \sum_r p N_r (x^*)^{a_r} u_r$ is nonzero: $p \neq 0$, the columns N_r span $\mathbb{R}^2$, and every $(x^*)^{a_r}$ is positive. Choose v with $p f_k(v) \neq 0$. The raw matrix on (v, u) is lower triangular with nonzero diagonal, so its determinant $p f_k(v) \cdot pB(q, h)$ is nonzero. Both directions are admissible because the base rates lie in the interior of the positive orthant. $\square$

We use only this direction of the relationship. The converse is false in general: the raw rank condition does not imply the corrected condition (3.3). For the positive direction of Theorem 2.3 the corrected minors are checked exactly on every positive template (Section 6).

4 Structural reductions and finite enumeration

4.1 Positive flux balance

At a positive equilibrium let $v_r = k_r(x^*)^{a_r} > 0$ be the reaction fluxes. Then

$$v > 0, \quad Nv = 0, \quad J = N \text{diag}(v) A \text{diag}(1/x^*). \quad (4.1)$$

Lemma 4.1. *Let $N \in \mathbb{Z}^{2 \times 5}$ have rank two. There is a strictly positive v with $Nv = 0$ if and only if for every column N_r some column has positive cross product with N_r and some column has negative cross product with N_r .*

Proof. If $Nv = 0$ with $v > 0$, take the cross product of both sides with N_r : a sum $\sum_s v_s (N_r \times N_s)$ with strictly positive weights vanishes, and rank two ensures that at least one term is nonzero, so both signs occur. Conversely, the condition says that no closed half-plane bounded by a line through the origin contains all columns; since the columns span the plane, the origin then lies in the interior of their convex hull, which is equivalent to a strictly positive balancing vector. $\square$




Figure 1: The levels of the classification. The first two arrows are quotients of the eligible literal reaction sets, by species exchange and by the complete invariant of Lemma 2.2; only the last arrow selects the cusp-admitting classes. The formal coverage records retain the exact literal source and its species orientation.

4.2 At least four distinct reactants

Lemma 4.2. *A network in the model that admits an ordinary transverse cusp has at least four distinct reactant complexes.*

Proof. Inside the cusp region of the unfolding of Proposition 3.1, the reduced equation has three simple zeros, so for suitable positive rates arbitrarily close to k^* the system has three isolated positive equilibria. By the few-source bound of Banaji and Feliu [3, Corollary 4.25], a rank-two planar mass-action network whose reactions have at most three distinct sources cannot have three isolated positive equilibria for any choice of rates. $\square$

The Lean development proves the same conclusion for `AlgCusp` by a direct algebraic argument that does not use [3]: after grouping reactions by equal reactants, the vector field has at most three monomials, and for each of four kernel charts covering the possible nonzero singular Jacobians the mixed cubic contraction $pB(q, h)$ is forced to vanish by the singular-equilibrium equations. The two arguments are independent; either suffices for the classification.

4.3 The eligible catalogue

There are 30 directed non-self reactions on $\mathcal{C}$, hence $\binom{30}{5} = 142,506$ literal five-reaction sets. Call a set *eligible* if it satisfies rank two, the surrounding condition of Lemma 4.1, and at least four distinct reactants. By (4.1), Lemma 4.1 and Lemma 4.2, every cusp-admitting network is eligible. (The Lean enumeration also records affine spanning of the reactants; in the six-point configuration $\mathcal{C}$ a line contains at most three points, so four distinct reactants already span the plane affinely, and the condition is implied.) The set $\mathcal{E}$ of eligible literal reaction sets has 60,036 elements; its quotient by species exchange has 30,051 elements, and its quotient by the key (2.2) has 9,999 elements (Figure 1).

These counts are exact finite computations. The number 60,036 is a Lean theorem (`structurallyEligibleSourceCatalogue_card`), and the check script accompanying the paper re-enumerates all 142,506 sets independently, applies the three filters, and recovers 60,036, 30,051 and 9,999; it also verifies that the keys of the networks extracted from the four Lean certificate layers are pairwise distinct and exhaust the keys of $\mathcal{E}$.

4.4 Unit-state normalization

The obstruction certificates of Section 5 are stated for the polynomial vector field

$$F_v(\xi) = \sum_r N_r v_r \xi^{a_r} \quad \text{at } \xi = \mathbf{1},$$

Table 2: The exhaustive partition of the 9,999 mechanism classes by the certificate that decides them. The source count is the number of distinct reactant complexes.

Certificate type	Four sources	Five sources	Total
Determinant obstruction	7,205	1,858	9,063
Quadratic-degeneracy obstruction	663	81	744
Cubic-degeneracy obstruction	132	8	140
Positive cusp template	0	52	52
Total	8,000	1,999	9,999

whose Jacobian and Hessian at $\mathbf{1}$ we denote $J_v = N \operatorname{diag}(v)A$ and B_v . With $D_x = \operatorname{diag}(x^*)$ these are related to the derivatives of f at (x^*, k^*) by

$$J_v = JD_x, \quad B_v(s, t) = B(D_x s, D_x t).$$

Hence if (p, q, h) satisfy (3.1) for f , then $q' = D_x^{-1}q$, $h' = D_x^{-1}h$ and $p' = p$ satisfy

$$J_v q' = 0, \quad p' J_v = 0, \quad p' B_v(q', q') = 0, \quad J_v h' = -B_v(q', q'), \quad p' B_v(q', h') = p B(q, h) \neq 0.$$

The normalization preserves the kernel, quadratic-degeneracy, center-correction and nonzero mixed-cubic conditions. It does *not* preserve $p'q' = 1$, $p'h' = 0$ or the trace, and the obstruction proofs use only the preserved, trace-free conditions. This is recorded because a change of normalization must not silently change the class of singularities being excluded: the negative certificates exclude a singular positive equilibrium with vanishing quadratic and nonvanishing cubic center coefficients, a strictly larger class than the cusps of Proposition 3.1.

5 Three exact obstruction schemes

All statements in this section concern arbitrary positive unit-equilibrium fluxes v of a fixed network, that is, all $v > 0$ with $Nv = 0$; nothing is sampled. Let

$$d(v) = \det J_v, \quad J_v = N \operatorname{diag}(v)A,$$

a quadratic form in v with integer coefficients. The certificates have rational coefficients, their generic soundness lemmas are Lean theorems, and each instance is checked in Lean by exact rational arithmetic. Table 2 gives the numbers of mechanism classes decided by each scheme, split by the number of distinct reactants.

5.1 Determinant obstructions

A determinant certificate consists of a rational multiplier $\sigma \neq 0$ and linear forms L_1, L_2 such that

$$\sigma d(v) - \sum_{\ell=1}^2 L_\ell(v) (Nv)_\ell = \sum_{r \leq s} c_{rs} v_r v_s, \quad c_{rs} \geq 0, \quad (5.1)$$

with at least one $c_{rs} > 0$. The identity is verified coefficientwise as an identity of quadratic forms. If $Nv = 0$ and $v > 0$, the right-hand side is strictly positive, so $d(v) \neq 0$: the Jacobian is nonsingular at every positive equilibrium, contradicting the necessary condition $\det J = 0$. This scheme decides 9,063 mechanism classes.

5.2 Quadratic-degeneracy obstructions

The closed nonnegative flux cone $\{v \geq 0 : Nv = 0\}$ is a pointed rational polyhedral cone generated by finitely many rational extreme rays. Write $v = Gt$ with $t \geq 0$, where the columns of G are the complete list of generators. Completeness of the generator list is part of the certificate and is checked, so the argument is not restricted to a chosen subcone.

For the planar matrix J_v , the vector perpendicular to its first row is a right-kernel candidate $q_0(v)$ and the vector perpendicular to its first column is a left-kernel candidate $p_0(v)$; both are linear in v . Let $H(v) = p_0(v) B_v(q_0(v), q_0(v))$, a homogeneous quartic in v . A quadratic-degeneracy certificate is an identity

$$\sigma H(Gt) = R_4(t) + L_2(t) d(Gt), \quad (5.2)$$

where R_4 is a quartic with nonnegative rational coefficients and L_2 is a quadratic form, together with a support condition: for every set S of generator indices such that the generators in S jointly activate every reaction coordinate, R_4 has a strictly positive coefficient on some monomial in the variables t_s , $s \in S$. The support condition ensures $R_4(t) > 0$ whenever $Gt > 0$, even if some weights t_s vanish.

Suppose $v = Gt > 0$, $Nv = 0$ and $d(v) = 0$. Then (5.2) gives $H(v) \neq 0$. In particular the first row and first column of J_v are nonzero, so $q_0(v)$ and $p_0(v)$ are nonzero multiples of the actual right and left kernel vectors of the singular matrix J_v , and $H(v) \neq 0$ says $pB_v(q, q) \neq 0$: the quadratic center coefficient does not vanish. This contradicts the second line of (3.1). The scheme decides 744 mechanism classes. It excludes a degenerate fold at a singular equilibrium; it says nothing about ordinary nondegenerate folds, which these networks may well have.

5.3 Cubic-degeneracy obstructions

A nonzero singular 2×2 matrix has a nonzero row and a nonzero column, so the perpendiculars to a row and to a column give four right/left kernel charts, each valid under a guard stating that the chosen row and column are nonzero. A cubic-degeneracy certificate proves, for each of the four charts and under its guard, the polynomial implication

$$\left. \begin{array}{l} v > 0, \quad Nv = 0, \quad d(v) = 0, \\ p_0 B_v(q_0, q_0) = 0, \quad J_v h_0 = -B_v(q_0, q_0) \end{array} \right\} \implies p_0 B_v(q_0, h_0) = 0. \quad (5.3)$$

The stored proofs use the equilibrium identities of the specific network to reduce the mixed contraction to zero; coverage of all four charts avoids dividing by an unverified nonzero entry.

To see that (5.3) excludes a cusp, take a chart whose candidates are nonzero and write $q_0 = \alpha q'$, $p_0 = \beta p'$ for the normalized kernel vectors of Section 4.4 and nonzero scalars α, β . Then $h_0 = \alpha^2 h'$ solves the center equation, and $p_0 B_v(q_0, h_0) = \beta \alpha^3 p' B_v(q', h')$, which is nonzero for a cusp. The scheme decides 140 mechanism classes.

None of the three schemes uses a parameter-unfolding condition; they all contradict state-jet conditions that hold for every singular equilibrium with the required center degeneracies. This is why the formal negative direction transfers to ordinary cusps without any analytic correction (Lemma 3.3).

6 Exact positive templates

Each of the 52 positive records supplies an equilibrium at $x^* = \mathbf{1}$, positive rates, vectors p, q, h and a selected pair of rate indices (i, j) . Every entry is a polynomial with rational coefficients in a real root t of a rational polynomial $m(t)$, together with a rational isolating interval for the root. Sixteen records are rational: they use $m(t) = t$ and $t = 0$. Fifteen use a quadratic



Counts remaining after removal by the selected certificate layers

Figure 2: Removal of mechanism classes by successive certificate layers. The intermediate numbers count classes not yet decided by the preceding schemes; they do not assert that every remaining class realizes the preceding degeneracy conditions. The label of a class names the certificate that decides it, not the first jet condition the network can satisfy.

m and twenty-one a cubic m . The Lean proofs establish the algebraic certificate `AlgCusp` for each record: the equilibrium equations, all identities in (3.1), positivity of the state and rates, nonzero trace, nonzero cubic contraction and the raw rank condition (3.4).

For the analytic conclusion we additionally need the corrected minor (3.3). At unit state, with e_r the r -th coordinate direction in rate space,

$$\alpha_r = pN_r = pf_k(e_r), \quad \eta_r = pN_r(a_r \cdot q) = pJ_k(e_r)q,$$

and the trace-scaled corrected minor of the selected pair is

$$D_{ij} = \alpha_i(\tau\eta_j - pB(q, N_j)) - \alpha_j(\tau\eta_i - pB(q, N_i)) = \tau \cdot \det \begin{pmatrix} a(e_i) & a(e_j) \\ b(e_i) & b(e_j) \end{pmatrix}. \quad (6.1)$$

All quantities are reduced modulo m . Each record contains polynomials $U, V \in \mathbb{Q}[t]$ with

$$U(t)D_{ij}(t) + V(t)m(t) = 1 \quad (6.2)$$

identically in $\mathbb{Q}[t]$. At the chosen root of m this gives $D_{ij} \neq 0$, and since $\tau \neq 0$ the corrected determinant D_{ij}/τ is nonzero. The trace-scaled form of the minor and its relation to the corrected matrix are the content of the Lean file `ReducedUnfolding.lean`, which proves the generic identity behind (6.1) for any certificate with nonzero trace.

Proposition 6.1 (Positive existence). *Every representative in Table 4 admits an ordinary transverse cusp at positive concentrations and positive rates.*


Proof. For each of the 52 records the exact checker verifies, by rational polynomial arithmetic modulo m and exact root isolation, that t is a root of m in the stated interval, that the state and rates are positive, that the equilibrium and all jet identities (3.1) hold, that $\tau \neq 0$ and $c \neq 0$, and that (6.2) holds identically for the recomputed D_{ij} . The record payloads agree with the Lean source files and their bound hashes. Hence every record satisfies the hypotheses of Proposition 3.1 for its selected pair. No floating-point tolerance enters any existence or nonvanishing test. $\square$

The 52 Bézout identities and root data are part of the machine-readable supplement, and the check script accompanying the paper re-verifies every identity (6.2) symbolically. They are compact exact evidence for the analytic correction; we do not describe them as 52 additional Lean analytic theorems.

In all 52 records the signs are the same: $\tau < 0$, $c < 0$ and $D_{ij} > 0$, proved by exact rational interval arithmetic on the isolating interval. The sign of τ and c is used in Section 8.

6.1 A rational representative

Template 02 makes the positive argument fully explicit:



with rates $(k_1, \dots, k_5) = (1, 3, 3, 3, 1)/11$ and vector field

$$\dot{x} = k_1 - k_2x + 2k_3y - k_4x^2 - k_5xy, \quad \dot{y} = k_4x^2 - k_3y. \quad (6.4)$$

At $x^* = (1, 1)$ one may take

$$J = \frac{1}{11} \begin{pmatrix} -10 & 5 \\ 6 & -3 \end{pmatrix}, \quad q = \frac{1}{11} \begin{pmatrix} 5 \\ 10 \end{pmatrix}, \quad p = \left(\frac{33}{65}, \frac{11}{13}\right), \quad h = \frac{1}{1573} \begin{pmatrix} -250 \\ 150 \end{pmatrix},$$

and direct substitution gives

$$\tau = -\frac{13}{11}, \quad c = -\frac{75}{17303}, \quad D_{1,5} = \frac{27}{65},$$

so the corrected determinant for the pair (k_1, k_5) is $D_{1,5}/\tau = -297/845$, and $U = 65/27$, $V = 0$ certify (6.2).

The same network admits a direct scalar reduction that displays the cusp without any local theory. Keep $k_3 = k_4 = 3/11$ and $k_5 = 1/11$ and write $a = k_1$, $b = k_2$. The second equilibrium equation gives $y = x^2$, and substituting into the first, with

$$x = 1 + z, \quad a = \frac{1 + \mu - \nu}{11}, \quad b = \frac{3 - \nu}{11},$$

yields exactly

$$11 \dot{x}|_{y=x^2} = \mu + \nu z - z^3. \quad (6.5)$$

The map $(\mu, \nu) \mapsto (a, b)$ is affine and invertible, the base rates are positive, and (6.5) is the universal unfolding of the cusp singularity. The fold curve is $(\mu, \nu) = (-2z^3, 3z^2)$, and the region $27\mu^2 < 4\nu^3$ contains three positive equilibria (Figure 3). This is a regular cusp unfolding by the pair (k_1, k_2) , a different admissible pair from the certified (k_1, k_5) .

7 Exhaustive coverage and proof of the main theorem

7.1 The formal statement

The Lean development represents a literal network as a structure with five reactant and five product complexes, injective reaction pairs and no self-reactions, and the model as the predicate `IsPlanarBimolecular252`. The algebraic certificate `AlgCusp` is the predicate `AdmitsTransverseCusp`. The following listing reproduces the definition of a valid certificate, with the network's mass-action field, Jacobian, Hessian and their rate derivatives as defined functions.

```

structure CuspCertificate {m : ℕ} (Q : SmallPlanarNetwork m) where
  state : Species → ℝ
  rates : Fin m → ℝ
  rightKernel : Species → ℝ
  leftKernel : Species → ℝ
  centerCorrection : Species → ℝ
  unfoldingDirections : Fin 2 → Fin m → ℝ

```




Figure 3: Exact equilibrium reduction of template 02, equation (6.5). Left: the fold curve $(\mu, \nu) = (-2z^3, 3z^2)$ bounds the three-root region. Right: at the marked parameter value $\mu = 0$, $\nu = 1/100$ the roots are $z = -1/10, 0, 1/10$, all corresponding to positive concentrations. The plots illustrate an exact polynomial identity, not fitted continuation data.

A certificate is valid when its state and rates are positive, the state is an equilibrium, the stored vectors satisfy the jet system (3.1) with nonzero trace and nonzero mixed cubic contraction, and the raw unfolding matrix on the two stored directions is nonsingular:

```
def CuspCertificate.Valid (C : CuspCertificate Q) : Prop :=
  PositiveVector C.state ^ PositiveVector C.rates ^
  (∀ i, Q.massAction C.rates C.state i = 0) ^
  (∀ i, Q.jacobianApply C.rates C.state C.rightKernel i = 0) ^
  (∀ j, ∑ i : Species, C.leftKernel i * Q.jacobian C.rates C.state i j = 0) ^
  dot C.leftKernel C.rightKernel = 1 ^
  Q.jacobian C.rates C.state 0 0 + Q.jacobian C.rates C.state 1 1 ≠ 0 ^
  dot C.leftKernel (Q.hessianApply C.rates C.state C.rightKernel C.rightKernel)
    = 0 ^
  (∀ i, Q.jacobianApply C.rates C.state C.centerCorrection i =
    -Q.hessianApply C.rates C.state C.rightKernel C.rightKernel i) ^
  dot C.leftKernel C.centerCorrection = 0 ^
  dot C.leftKernel
    (Q.hessianApply C.rates C.state C.rightKernel C.centerCorrection) ≠ 0 ^
  Matrix.det C.unfoldingMatrix ≠ 0

def AdmitsTransverseCusp (Q : SmallPlanarNetwork m) : Prop :=
  ∃ C : CuspCertificate Q, C.Valid
```

Here `unfoldingMatrix` is the raw matrix (3.4) evaluated on the two stored directions. The public theorem and its negative direction are unconditional:

```
theorem no_cusp_if_not_in_52 (Q : SmallPlanarNetwork 5)
  (hQ : IsPlanarBimolecular252 Q)
  (hnot : ¬ BelongsToBimolecularCuspClass Q) : ¬ AdmitsTransverseCusp Q

theorem classify_planar_bimolecular_2_5_2_cusps
  (Q : SmallPlanarNetwork 5) (hQ : IsPlanarBimolecular252 Q) :
  AdmitsTransverseCusp Q ↔ BelongsToBimolecularCuspClass Q
```

Membership `BelongsToBimolecularCuspClass Q` means that the literal network Q has an exact coded realization whose source-lookup record is valid and carries the positive tag; such a record names the positive template, the reaction matching and the species orientation

that connect Q to it. It is source-faithful: no representative is silently identified with a network, and every step is witnessed by exact reaction-set equality, species exchange or positive column scaling, each of which is proved to preserve and reflect `AdmitsTransverseCusp`.

7.2 Structure of the formal proof

The formal proof has the following layers; every layer is a separate Lean module (Table 3).

- (1) *Necessary structure*. From a valid certificate the development derives the structural conditions of Section 4: rank two, a strictly positive kernel vector, and at least four reactant complexes. These define the eligible catalogue of 60,036 literal reaction sets.
- (2) *Mechanism partition*. Four lists of mechanism representatives, of lengths 52, 9,063, 744 and 140, carry positive certificates or obstruction certificates checked instance by instance. The four outcome tags are proved exhaustive and pairwise exclusive, and $52 + 9063 + 744 + 140 = 9999$ is a theorem.
- (3) *Source records*. 30,051 records, one per species-exchange orbit of $\mathcal{E}$, each contain a literal source, the index of its mechanism representative, matching and orientation data and the outcome tag. Docking (positive simple equivalence of the source to the oriented representative, proved semantically from a division-free integer ray check) and target consistency (agreement of the tag with the layer) are checked separately and combined into record validity.
- (4) *Lookup*. A table of 60,036 oriented source keys maps every eligible set to a record index and an orientation flag. The lookup is proved complete (every eligible set has an entry) and sound (every entry is eligible), and key injectivity converts key equality into exact equality of reaction sets.
- (5) *Transport*. A valid record transports its outcome to its literal source: positive records yield `AdmitsTransverseCusp`, obstruction records yield its negation, along the proved equivalence transports.

The size of the certificate payloads made a naive coverage proof impractical, and the solution is a reusable proof-engineering idea. The coverage argument never elaborates the large algebraic records while it reasons about the lookup table. Instead, sixty-one bounded Lean batch proofs establish that mapping the actual records, in their stored order, to their pair of compact keys reproduces a packed key list; list-map and concatenation identities assemble the batches; a generic indexed-lookup lemma then recovers the actual record for any key; and key injectivity supplies the reaction-set equality. Numerical preflights guided the representation, but the key/record equalities and the coverage implication are Lean theorems.

7.3 Proof of Theorem 2.3

Proof. Suppose first that Q admits an ordinary transverse cusp. Lemma 3.3 gives `AlgCusp`(Q), so by `classify_planar_bimolecular_2_5_2_cusps` the literal network Q is source-faithfully covered by a positive record. Unwinding the record, Q is simply equivalent up to species exchange to one of the 52 representatives, and by Lemma 2.2 its key is that representative’s key. (Equivalently, and without appeal to the analytic level: were Q ’s record an obstruction record, the corresponding scheme of Section 5 would contradict a necessary state-jet condition of the cusp.)

Conversely, if $K(Q)$ equals a listed key, Lemma 2.2 makes Q simply equivalent up to species exchange to the representative, Proposition 6.1 gives that representative an ordinary transverse cusp, and reaction permutation, positive rate reparameterization and species exchange transport the cusp, together with its transverse unfolding by the corresponding pair of rates, to Q .

The 52 listed keys are pairwise distinct by exact finite comparison, so the representatives are pairwise inequivalent by Lemma 2.2. $\square$

8 Consequences

8.1 Five distinct reactants

Proof of Corollary 2.4. All 52 representatives have five distinct reactant complexes (Table 4), and the number of distinct reactants is invariant under Definition 2.1. Table 2 shows that the 8,000 mechanism classes with four reactants are all decided by obstructions. Thus the general bound of Lemma 4.2 is not sharp in this five-reaction class: four sources are necessary in general, but five are necessary here. $\square$

The Lean development exports the general lower bound (`four_reactants_necessary_for_bimolecular_cusp`) and the existence of a five-reactant witness (`five_distinct_reactants_suffice`); the sharpened statement of Corollary 2.4 is read off from the exact partition, which is verified by the check script, and is not a separate Lean export.

8.2 Reaction minimality

Proof of Corollary 2.5. Banaji, Boros and Hofbauer proved that no full-rank planar quadratic mass-action network with four reactions admits a cusp unfolded by its rate constants [6, Theorem 36]; this covers the rank-two planar bimolecular networks with four reactions. With at most three reactions there are at most three sources, and Lemma 4.2 applies. Rank-one planar quadratic systems have a one-dimensional stoichiometric class on which the equilibrium equation is a polynomial of degree at most two, so they have at most two isolated positive equilibria and no cubic degeneracy; rank zero is trivial. Template 02 with five bimolecular reactions attains the bound, so five reactions is the minimum and each of the 52 classes realizes it. $\square$

8.3 Every class organizes bistability

Recall that in all 52 records $\tau < 0$ and $c < 0$.

Proof of Corollary 2.6. Fix a representative and its certified pair (k_i, k_j) . By Proposition 3.1 the reduced equation $g(z, k) = 0$ is a versal unfolding of $z \mapsto cz^3$ with $c < 0$, so for rates in the interior of the cusp region, close to k^* , $g(\cdot, k)$ has three simple zeros $z_1 < z_2 < z_3$ with $g_z(z_1, k) < 0$, $g_z(z_2, k) > 0$, $g_z(z_3, k) < 0$. Let $x(z, k) = x^* + zq + y(z, k)$ be the corresponding equilibria and $J(z, k)$ their Jacobians. Differentiating the identity $f(x(z, k), k) = qg(z, k)$ in z gives $J(z, k)(q + y_z) = g_z q$. Writing $J(z, k)$ in the basis $\{q + y_z, w\}$ with w spanning $\ker p$, its matrix has first column $(g_z, -g_z\gamma)^\top$ where $y_z = \gamma w$, and second column $(*, \beta)^\top$ with $\beta = \tau$ and $\gamma = 0$ at the cusp point; hence $\det J(z, k) = g_z(\beta + \gamma*)$, whose second factor is close to $\tau < 0$ near the cusp. So $\det J < 0$ at the middle zero and $\det J > 0$ at the outer zeros, while $\text{tr } J$ is close to $\tau < 0$ at all three. The outer equilibria are asymptotically stable, the middle one is a saddle, and all three are positive because x^* is positive. $\square$

The mechanism is an exact rational instance for template 02. With $0 < \varepsilon < 1$ take $k_1 = (1 - \varepsilon^2)/11$, $k_2 = (3 - \varepsilon^2)/11$ and k_3, k_4, k_5 as in (6.3). Then $\mu = 0$, $\nu = \varepsilon^2$ in (6.5), the three equilibria are $x = 1 - \varepsilon, 1, 1 + \varepsilon$ with $y = x^2$, and their Jacobians have

$$\det J = \frac{6\varepsilon^2}{121}, \quad -\frac{3\varepsilon^2}{121}, \quad \frac{6\varepsilon^2}{121}, \quad \text{tr } J = \frac{8\varepsilon - 13}{11}, \quad \frac{\varepsilon^2 - 13}{11}, \quad -\frac{8\varepsilon + 13}{11},$$

respectively: a stable node, a saddle and a stable node. At $\varepsilon = 1/10$ the rates are $(99/1100, 299/1100, 3/11, 3/11, 1/11)$ and the equilibria $(9/10, 81/100), (1, 1), (11/10, 121/100)$ have determinants $3/6050, -3/12100, 3/6050$ and traces $-61/55, -1299/1100, -69/55$. These identities are verified symbolically by the check script. Corollary 2.6 is an existential local statement for each class; it describes neither all cusp points nor the global parameter space of the class.

8.4 Ray geometry does not determine cusp existence

A natural conjecture is that cusp existence depends only on the unlabelled geometry of the five stoichiometric rays, or at least on their oriented matroid. The classification refutes it. The 52 positive classes realize 31 distinct unlabelled multisets of primitive directed rays and only five distinct rank-two oriented matroids, so the reactant labels carry most of the information. The lexicographically smallest collision between a positive and a negative class is

positive (template 00): $0 \rightarrow X, X \rightarrow 0, Y \rightarrow 2X, 2X \rightarrow X + Y, X + Y \rightarrow 0$;
 negative (cubic layer): $0 \rightarrow X, X \rightarrow 0, Y \rightarrow 2X, X + Y \rightarrow 0, X + Y \rightarrow 2Y$.

The reactions $2X \rightarrow X + Y$ and $X + Y \rightarrow 2Y$ both have stoichiometric vector $(-1, 1)$, so the two networks have the same five directed vectors, not merely the same matroid; they differ only in the reactant of that one reaction, and the negative network has four distinct reactants.

The reason is explicit. Write the rates of the negative network, in the displayed order, as a, b, d, e, f . Its vector field is

$$\dot{x} = a - bx + 2dy - (e + f)xy, \quad \dot{y} = ((f - e)x - d)y.$$

A positive equilibrium requires $f > e$ and $x = d/(f - e)$, and there $\det J = -d(f - 3e)y$. A singular positive equilibrium therefore requires $f = 3e$, hence $x = d/(2e)$, and then the first equation reads $a - bd/(2e) = 0$ independently of y : the singular equilibrium lies on a whole line of equilibria. The degeneracy is a continuum, never an isolated cubic cusp, and the cubic-degeneracy certificate of Section 5 is its exact witness. The same ray geometry supports different nonlinear equilibrium geometry because the reactant label changed. A complete reactant-labelled key is the invariant of Lemma 2.2, and it yields the finite classification itself; we assert no smaller structural generator theorem.

8.5 Inheritance

Once analytic rate transversality is established, the inheritance theorem of Banaji, Boros and Hofbauer [8, Theorem 3.2] applies: every network in the 52 classes passes its nondegenerate cusp, with its transverse unfolding, to any enlargement obtained by the operations listed there, including rank-preserving addition of reactions, addition of inflow and outflow reactions, and the specified new-species constructions, and to any finite composition of such operations. This gives infinitely many larger mass-action networks, bimolecular or not, with certified cusps and the bistability of Corollary 2.6; the rank hypotheses on added reversible reactions and reaction splittings are essential, and arbitrary supernetworks are not covered. This is an application of a published theorem, not a new formal result.

9 Verification and reproducibility

9.1 Formal interfaces

The reusable proof sources are under `proofs/SmallCusp/` in the supplement. The positive templates and the three obstruction layers are separate modules rather than one generated monolith; Table 3 lists the critical interfaces.




Figure 4: Stoichiometric vectors of three positive templates, each arrow labelled by the reactant complex of its reaction. Templates 00 and 02 share four reactions and differ only in the product of the reaction with reactant $X + Y$. In template 02 the reactions $X \rightarrow 0$ and $X + Y \rightarrow Y$ have the same vector $(-1, 0)$, so its five reactions span only four rays; the reactant labels are what distinguish coincident rays, and Section 8.4 shows that they also decide cusp existence. Template 51 has five distinct rays and no reaction with reactant 0.

Table 3: Critical formal interfaces, with paths relative to `proofs/SmallCusp/`.

Role	Source module
Literal model and algebraic cusp predicate	<code>Source/SourceClasses.lean</code>
Trace-free unit-state normalization	<code>Normalization/TraceFreeCertificate.lean</code>
Positive templates (52)	<code>Cusp/PositiveCuspLayer.lean</code>
Trace-scaled corrected minor identity	<code>Cusp/ReducedUnfolding.lean</code>
Determinant exclusions (9,063)	<code>Obstruction/ConicDeterminantLayer.lean</code>
Quadratic exclusions (744)	<code>Obstruction/CircuitFoldLayer.lean</code>
Cubic exclusions (140)	<code>Obstruction/CubicResidualLayer.lean</code>
Outcome exhaustiveness and partition count	<code>Classification/CoverageTypes.lean</code>
Eligible catalogue cardinality (60,036)	<code>Classification/SourceCatalogueCount.lean</code>
Unconditional literal coverage	<code>Classification/SourceCoverageLookup.lean</code>
Public iff and negative theorem	<code>Classification/BimolecularClassification.lean</code>
Exported corollaries and counts	<code>Classification/BimolecularCorollaries.lean</code>

9.2 Trust boundary

The public theorems were verified with Lean 4.30.0 and the frozen Mathlib manifest of the repository, with warnings treated as errors, no `sorry`, and exit code zero; the receipt covers 1,072 dependency modules. Each of the two public exports reports the standard axioms `Classical.choice`, `Quot.sound` and `propext` together with 3,826 generated `native_decide` axioms, and no unexpected axiom. Thus the finite computations (certificate checks, list lengths, key comparisons) trust Lean’s compiled native-evaluation bridge, in addition to the Lean kernel and the frozen environment; this is not a claim of kernel-only reduction of every finite computation. The strict receipt records a 120.711-second verification of the public module; this is provenance, not a benchmark of proof development or of a clean rebuild, which materializes the dependency modules first.

9.3 What is formal and what is conventional

- *Lean theorems*: the model and certificate definitions; the necessary structural conditions for `AlgCusp`, including the four-reactant bound; the soundness lemmas and all instances

of the three obstruction schemes; the 52 positive AlgCusp certificates; exhaustiveness and exclusivity of the outcomes and the count 9,999; the cardinalities 60,036 and 30,051; completeness and soundness of the lookup; the equivalence transports; the trace-scaled minor identity; and the unconditional public iff and negative theorem.

- *Conventional proofs in this paper:* the Lyapunov–Schmidt reduction of Proposition 3.1; Lemma 3.3; Lemma 4.1; Lemma 4.2 via [3]; the key invariant, Lemma 2.2; the sign argument of Corollary 2.6; the continuum-degeneracy computation of Section 8.4; the reduction to published results in Corollary 2.5 and the inheritance remark.
- *Exact computer algebra, reproducible from the supplement:* the 52 corrected-minor Bézout identities and sign certificates; the independent enumeration of the 142,506 literal sets and the counts 60,036, 30,051, 9,999; pairwise distinctness of the 9,999 keys and of the 52 positive keys; the split of Table 2 by source count; the ray-skeleton and oriented-matroid counts of Section 8.4; and the rational examples of Sections 6.1 and 8.

Exporting the conventional steps as Lean declarations is a separate formalization project and is not claimed here.

9.4 Supplement and check script

The supplement contains the exact positive root and witness data, all corrected-minor Bézout certificates, the four representative partitions, the critical Lean dependency sources with SHA-256 hashes, the strict verification receipt, deterministic generators for the tables and figures, and reproduction instructions. The check script distributed with this paper re-derives, with SymPy, every number printed in the text: the enumeration counts, the partition and its source split, the degree distribution of the root polynomials, the 52 Bézout identities, the jet data and corrected minor of template 02, the scalar reduction (6.5), the three-equilibria family and its rational instance, the collision computation of Section 8.4, the receipt figures quoted above, and the rows of Table 4 against the Lean records. The full positive catalogue can therefore be read without software, and its exact witnesses checked mechanically. The pinned Lean and Mathlib environment is a prerequisite for replaying the formal proof; it is not replaced by current dependencies.

10 Discussion and open problems

The classification settles the smallest planar bimolecular case of a cusp unfolded by rate constants, and it does so with the identity of every literal network preserved through the proof. Several questions are left open by design.

Comparison with the dual-equation approach. Banaji’s correspondence between original and dual fold and cusp conditions [1] is a different route to the same class, and Section 6.3 of that preprint announces a classification. When that list is available, a class-by-class comparison via the keys of Appendix A would be a valuable cross-check of two independent methods; we have not performed one, and our correctness does not depend on the dual-equation correspondence.

Larger models. The same pipeline applies to five-reaction networks with trimolecular products, to more reactions, and to three species, with the enumeration and certificate counts growing accordingly. The (2, 5, 2) trimolecular-product class is defined in the Lean development but not classified here. The separate problem of four-reaction planar networks with multiple nonnegative stable equilibria on the boundary of the orthant is also outside this theorem.

Coarser equivalences. Our count is under Definition 2.1. Under dynamical equivalence, or under diagonal equivalence that also rescales rows of N , several of the 52 classes may merge; determining the coarser count is a distinct problem.

The negative categories are not strata. The obstruction that decides a class is the first sufficient certificate found in a fixed order; it does not assert that the class realizes the preceding degeneracies. A classification of the negative classes by the strongest jet condition they can realize (nonsingular everywhere, folds only, degenerate folds only) is a natural refinement.

Formalization of the analytic bridge. The implicit-function reduction of Proposition 3.1 and the corrected-minor instances are within reach of Mathlib’s analysis library, and formalizing them would move Theorem 2.3 entirely inside the formal core.

Global dynamics. Corollary 2.6 is local. Which of the 52 classes admit globally attracting bistability, and which combine the cusp with Hopf or Bogdanov–Takens points, are questions for the methods of [2, 6].

References

- [1] Murad Banaji. On the computation of fold and cusp bifurcations in mass action networks via dual equations. arXiv:2609.07265v1, 7 September 2026, 2026. URL <https://arxiv.org/abs/2609.07265>.
- [2] Murad Banaji and Balázs Boros. The smallest bimolecular mass action reaction networks admitting Andronov–Hopf bifurcation. *Nonlinearity*, 36(2):1398–1433, 2023. doi: 10.1088/1361-6544/acb0a8.
- [3] Murad Banaji and Elisenda Feliu. Positive equilibria in mass action networks: geometry and bounds. *Journal of Mathematical Biology*, 93:37, 2026. doi: 10.1007/s00285-026-02429-8.
- [4] Murad Banaji and Casian Pantea. The inheritance of nondegenerate multistationarity in chemical reaction networks. *SIAM Journal on Applied Mathematics*, 78(2):1105–1130, 2018. doi: 10.1137/16M1103506.
- [5] Murad Banaji, Balázs Boros, and Josef Hofbauer. Adding species to chemical reaction networks: preserving rank preserves nondegenerate behaviours. *Applied Mathematics and Computation*, 426:127109, 2022. doi: 10.1016/j.amc.2022.127109.
- [6] Murad Banaji, Balázs Boros, and Josef Hofbauer. Bifurcations in planar, quadratic mass-action networks with few reactions and low molecularity. *Nonlinear Dynamics*, 112: 21425–21448, 2024. doi: 10.1007/s11071-024-10068-1.
- [7] Murad Banaji, Balázs Boros, and Josef Hofbauer. Oscillations in three-reaction quadratic mass-action systems. *Studies in Applied Mathematics*, 152(1):249–278, 2024. doi: 10.1111/sapm.12639.
- [8] Murad Banaji, Balázs Boros, and Josef Hofbauer. The inheritance of local bifurcations in mass action networks. *Journal of Nonlinear Science*, 35:71, 2025. doi: 10.1007/s00332-025-10165-4.
- [9] Carsten Conradi, Elisenda Feliu, Maya Mincheva, and Carsten Wiuf. Identifying parameter regions for multistationarity. *PLoS Computational Biology*, 13(10):e1005751, 2017. doi: 10.1371/journal.pcbi.1005751.
- [10] Gheorghe Craciun and Martin Feinberg. Multiple equilibria in complex chemical reaction networks: I. The injectivity property. *SIAM Journal on Applied Mathematics*, 65(5): 1526–1546, 2005. doi: 10.1137/S0036139904440278.
- [11] Leonardo de Moura and Sebastian Ullrich. The Lean 4 theorem prover and programming language. In *Automated Deduction – CADE 28*, volume 12699 of *Lecture Notes in Computer Science*, pages 625–635. Springer, 2021. doi: 10.1007/978-3-030-79876-5_37.

- [12] Martin Feinberg. *Foundations of Chemical Reaction Network Theory*, volume 202 of *Applied Mathematical Sciences*. Springer, 2019. doi: 10.1007/978-3-030-03858-8.
- [13] Martin Golubitsky and David G. Schaeffer. *Singularities and Groups in Bifurcation Theory, Volume I*, volume 51 of *Applied Mathematical Sciences*. Springer, 1985. doi: 10.1007/978-1-4612-5034-0.
- [14] Georges Gonthier. Formal proof—the four-color theorem. *Notices of the American Mathematical Society*, 55(11):1382–1393, 2008.
- [15] Thomas Hales, Mark Adams, Gertrud Bauer, Tat Dat Dang, John Harrison, Le Truong Hoang, Cezary Kaliszyk, Victor Magron, Sean McLaughlin, Tat Thang Nguyen, Quang Truong Nguyen, Tobias Nipkow, Steven Obua, Joseph Pleso, Jason Rute, Alexey Solovyev, Thi Hoai An Ta, Nam Trung Tran, Thi Diep Trieu, Josef Urban, Ky Vu, and Roland Zumkeller. A formal proof of the Kepler conjecture. *Forum of Mathematics, Pi*, 5:e2, 2017. doi: 10.1017/fmp.2017.1.
- [16] Badal Joshi and Anne Shiu. Atoms of multistationarity in chemical reaction networks. *Journal of Mathematical Chemistry*, 51:153–178, 2013. doi: 10.1007/s10910-012-0072-0.
- [17] Yuri A. Kuznetsov. *Elements of Applied Bifurcation Theory*, volume 112 of *Applied Mathematical Sciences*. Springer, third edition, 2004. doi: 10.1007/978-1-4757-3978-7.
- [18] The mathlib Community. The Lean mathematical library. In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs (CPP 2020)*, pages 367–381. ACM, 2020. doi: 10.1145/3372885.3373824.
- [19] Thomas Wilhelm. The smallest chemical reaction system with bistability. *BMC Systems Biology*, 3:90, 2009. doi: 10.1186/1752-0509-3-90.

A Complete list of positive representatives

Table 4 lists the 52 pairwise inequivalent cusp-admitting classes. The five reactions of each row are ordered as in its exact witness; row identifiers are zero-based and match the ordinals of the supplement. Column d is the degree of the witness root polynomial m ($d = 1$ means a rational witness with $m(t) = t$; the degree is not an assertion of irreducibility), and (i, j) is the one-based pair of rate indices whose corrected minor is certified by (6.2). The symbol 0 denotes the empty complex. All rows have five distinct reactant complexes. The key of a row is recovered directly by (2.2), and equivalence is exactly Definition 2.1. The rational example (6.3) is row 02.

Table 4: The 52 pairwise inequivalent cusp-admitting classes.

ID	Reaction 1	Reaction 2	Reaction 3	Reaction 4	Reaction 5	d	(i, j)
00	$0 \rightarrow X$	$X \rightarrow 0$	$Y \rightarrow 2X$	$2X \rightarrow X+Y$	$X+Y \rightarrow 0$	3	1,5
01	$0 \rightarrow X$	$X \rightarrow 0$	$Y \rightarrow 2X$	$2X \rightarrow X+Y$	$X+Y \rightarrow X$	3	1,5
02	$0 \rightarrow X$	$X \rightarrow 0$	$Y \rightarrow 2X$	$2X \rightarrow X+Y$	$X+Y \rightarrow Y$	1	1,5
03	$0 \rightarrow X$	$X \rightarrow 0$	$Y \rightarrow 2X$	$2X \rightarrow X+Y$	$2Y \rightarrow X$	3	1,5
04	$0 \rightarrow X$	$X \rightarrow 0$	$Y \rightarrow 2X$	$2X \rightarrow X+Y$	$2Y \rightarrow 0$	3	1,5
05	$0 \rightarrow X$	$X \rightarrow Y$	$Y \rightarrow 2Y$	$X+Y \rightarrow 0$	$2Y \rightarrow X$	3	2,5
06	$0 \rightarrow X$	$X \rightarrow Y$	$Y \rightarrow 2Y$	$X+Y \rightarrow 0$	$2Y \rightarrow 0$	3	2,5
07	$0 \rightarrow X$	$X \rightarrow Y$	$Y \rightarrow 2Y$	$X+Y \rightarrow 0$	$2Y \rightarrow 2X$	2	2,5
08	$0 \rightarrow X$	$X \rightarrow X+Y$	$Y \rightarrow 2Y$	$X+Y \rightarrow 0$	$2Y \rightarrow X$	1	2,5
09	$0 \rightarrow X$	$X \rightarrow X+Y$	$Y \rightarrow 2Y$	$X+Y \rightarrow 0$	$2Y \rightarrow 0$	1	2,5
10	$0 \rightarrow X$	$X \rightarrow X+Y$	$Y \rightarrow 2Y$	$X+Y \rightarrow 0$	$2Y \rightarrow 2X$	1	2,5
11	$0 \rightarrow X$	$X \rightarrow 2Y$	$Y \rightarrow 0$	$2X \rightarrow 0$	$X+Y \rightarrow 2X$	1	1,4
12	$0 \rightarrow X$	$X \rightarrow 2Y$	$Y \rightarrow 0$	$2X \rightarrow 0$	$2Y \rightarrow 2X$	3	1,4
13	$0 \rightarrow X$	$X \rightarrow 2Y$	$2X \rightarrow 0$	$X+Y \rightarrow 2X$	$2Y \rightarrow 0$	3	1,3
14	$0 \rightarrow X$	$X \rightarrow 2Y$	$Y \rightarrow 0$	$2X \rightarrow Y$	$X+Y \rightarrow 2X$	3	1,4
15	$0 \rightarrow X$	$X \rightarrow 2Y$	$Y \rightarrow 0$	$2X \rightarrow Y$	$2Y \rightarrow 2X$	2	1,4
16	$0 \rightarrow X$	$X \rightarrow 2Y$	$2X \rightarrow Y$	$X+Y \rightarrow 2X$	$2Y \rightarrow 0$	2	1,3
17	$0 \rightarrow X$	$X \rightarrow 2Y$	$2X \rightarrow X+Y$	$X+Y \rightarrow 2X$	$2Y \rightarrow 0$	3	1,3
18	$0 \rightarrow X$	$X \rightarrow 2Y$	$Y \rightarrow 0$	$X+Y \rightarrow 0$	$2Y \rightarrow 2X$	2	1,4
19	$0 \rightarrow X$	$X \rightarrow 2Y$	$Y \rightarrow 0$	$X+Y \rightarrow X$	$2Y \rightarrow 2X$	1	1,4
20	$0 \rightarrow X$	$X \rightarrow 2Y$	$Y \rightarrow 0$	$X+Y \rightarrow Y$	$2Y \rightarrow 2X$	2	1,4
21	$0 \rightarrow X$	$X \rightarrow 2Y$	$Y \rightarrow 2Y$	$X+Y \rightarrow 0$	$2Y \rightarrow X$	2	2,5
22	$0 \rightarrow X$	$X \rightarrow 2Y$	$Y \rightarrow 2Y$	$X+Y \rightarrow 0$	$2Y \rightarrow 0$	3	2,5
23	$0 \rightarrow X$	$X \rightarrow 2Y$	$Y \rightarrow 2Y$	$X+Y \rightarrow 0$	$2Y \rightarrow 2X$	2	2,5
24	$0 \rightarrow X$	$Y \rightarrow 2Y$	$2X \rightarrow Y$	$X+Y \rightarrow 0$	$2Y \rightarrow X$	1	3,5
25	$0 \rightarrow X$	$Y \rightarrow 2Y$	$2X \rightarrow Y$	$X+Y \rightarrow 0$	$2Y \rightarrow 0$	3	3,5
26	$0 \rightarrow X$	$Y \rightarrow 2Y$	$2X \rightarrow Y$	$X+Y \rightarrow 0$	$2Y \rightarrow 2X$	2	3,5
27	$0 \rightarrow X$	$Y \rightarrow 2Y$	$2X \rightarrow X+Y$	$X+Y \rightarrow 0$	$2Y \rightarrow X$	2	3,5
28	$0 \rightarrow X$	$Y \rightarrow 2Y$	$2X \rightarrow X+Y$	$X+Y \rightarrow 0$	$2Y \rightarrow 0$	3	3,5
29	$0 \rightarrow X$	$Y \rightarrow 2Y$	$2X \rightarrow X+Y$	$X+Y \rightarrow 0$	$2Y \rightarrow 2X$	3	3,5
30	$0 \rightarrow X+Y$	$X \rightarrow 0$	$Y \rightarrow 2X$	$2X \rightarrow X+Y$	$X+Y \rightarrow 0$	2	1,5
31	$0 \rightarrow X+Y$	$X \rightarrow 0$	$Y \rightarrow 2X$	$2X \rightarrow X+Y$	$X+Y \rightarrow X$	3	1,5
32	$0 \rightarrow X+Y$	$X \rightarrow 0$	$Y \rightarrow 2X$	$2X \rightarrow X+Y$	$X+Y \rightarrow Y$	1	1,5
33	$0 \rightarrow X+Y$	$X \rightarrow 0$	$Y \rightarrow 2X$	$2X \rightarrow X+Y$	$2Y \rightarrow X$	1	1,5
34	$0 \rightarrow X+Y$	$X \rightarrow 0$	$Y \rightarrow 2X$	$2X \rightarrow X+Y$	$2Y \rightarrow 0$	2	1,5
35	$0 \rightarrow X+Y$	$X \rightarrow 0$	$Y \rightarrow 2X$	$X+Y \rightarrow 2Y$	$2Y \rightarrow X$	1	1,5

Continued on the next page

Table 4 continued

ID	Reaction 1	Reaction 2	Reaction 3	Reaction 4	Reaction 5	d	(i, j)
36	$0 \rightarrow X+Y$	$X \rightarrow 0$	$Y \rightarrow 2X$	$X+Y \rightarrow 2Y$	$2Y \rightarrow 0$	3	1,5
37	$0 \rightarrow X+Y$	$X \rightarrow 2Y$	$2X \rightarrow 0$	$X+Y \rightarrow 2X$	$2Y \rightarrow 0$	2	1,3
38	$0 \rightarrow X+Y$	$X \rightarrow 2Y$	$2X \rightarrow Y$	$X+Y \rightarrow 2X$	$2Y \rightarrow 0$	1	1,3
39	$0 \rightarrow X+Y$	$X \rightarrow 2Y$	$2X \rightarrow X+Y$	$X+Y \rightarrow 2X$	$2Y \rightarrow 0$	3	1,3
40	$X \rightarrow 2X$	$Y \rightarrow 2Y$	$2X \rightarrow Y$	$X+Y \rightarrow 0$	$2Y \rightarrow X$	1	3,5
41	$X \rightarrow 2X$	$Y \rightarrow 2Y$	$2X \rightarrow Y$	$X+Y \rightarrow 0$	$2Y \rightarrow 2X$	3	3,5
42	$X \rightarrow 2X$	$Y \rightarrow 2Y$	$2X \rightarrow X+Y$	$X+Y \rightarrow 0$	$2Y \rightarrow 2X$	1	3,5
43	$X \rightarrow Y$	$Y \rightarrow 2X$	$2X \rightarrow 0$	$X+Y \rightarrow 2Y$	$2Y \rightarrow X$	1	1,5
44	$X \rightarrow Y$	$Y \rightarrow 2X$	$2X \rightarrow 0$	$X+Y \rightarrow 2Y$	$2Y \rightarrow 0$	1	1,5
45	$X \rightarrow Y$	$Y \rightarrow 2X$	$2X \rightarrow 0$	$X+Y \rightarrow 2Y$	$2Y \rightarrow 2X$	1	1,5
46	$X \rightarrow X+Y$	$Y \rightarrow 2X$	$2X \rightarrow 0$	$X+Y \rightarrow 2Y$	$2Y \rightarrow X$	3	1,5
47	$X \rightarrow X+Y$	$Y \rightarrow 2X$	$2X \rightarrow 0$	$X+Y \rightarrow 2Y$	$2Y \rightarrow 0$	2	1,5
48	$X \rightarrow X+Y$	$Y \rightarrow 2X$	$2X \rightarrow 0$	$X+Y \rightarrow 2Y$	$2Y \rightarrow 2X$	3	1,5
49	$X \rightarrow 2Y$	$Y \rightarrow 2X$	$2X \rightarrow 0$	$X+Y \rightarrow 2X$	$2Y \rightarrow 0$	2	2,3
50	$X \rightarrow 2Y$	$Y \rightarrow 2X$	$2X \rightarrow 0$	$X+Y \rightarrow 2Y$	$2Y \rightarrow X$	2	1,5
51	$X \rightarrow 2Y$	$Y \rightarrow 2X$	$2X \rightarrow 0$	$X+Y \rightarrow 2Y$	$2Y \rightarrow 2X$	3	1,5