

Thermodynamic compatibility of interacting autocatalytic cores: an exact phase diagram and a layered obstruction theory

Anonymous manuscript

September 7, 2026

Abstract

A potential autocatalytic core (PAC) is a minimal stoichiometric motif that admits a strictly productive reaction current; a consistent autocatalytic core (CAC) is a PAC whose productive current is realized by a positive concentration vector under thermodynamically consistent mass-action kinetics. Kosc, Kuperberg, Rajon and Charlat (*Proc. Natl. Acad. Sci. USA*, 2025) proved that every isolated PAC is a CAC, and exhibited two overlapping cores that admit a common productive current but no common concentration vector. We study when a family of interacting cores admits one common thermodynamic realization, with all cores sharing a single activity vector and fixed transition-state factors. First, we refute the natural strengthening of the isolated-core theorem: there is a two-core family with a common productive current and a common strict reaction-direction potential that nevertheless has no common realization. The obstruction is metric rather than directional: it already defeats the linear relaxation in which every distinct stoichiometric complex receives an independent activity. Second, for the responsible topology, two autocatalytic triangles $A \rightleftharpoons B \rightleftharpoons C \rightleftharpoons 2A$ and $A \rightleftharpoons B \rightleftharpoons D \rightleftharpoons 2A$ sharing $A \rightleftharpoons B$, we prove an exact phase diagram: with private transition-state factors in ratio k , a common realization exists if and only if $1/2 < k < 2$, and for arbitrary positive factors if and only if the reciprocal-factor sums of the two private branches are within a factor of two. In transition-state language the symmetric window is $|\Delta G^\ddagger| < RT \ln 2$. Third, we derive an interface calculus: eliminating a gain- m triangle with factors b_0, b_1, b_2 leaves exactly the open response interval $(R/m, R)$ with $R = b_0(1/b_1 + 1/b_2)$, and two eliminated cores compose precisely when their intervals overlap. Fourth, we show that the linear relaxation is not the whole story: a two-species, three-reaction family passes the flow, direction and independent-complex tests but fails the monomial lift to a common species-activity vector. Together with the direction obstruction of the original example, this yields three provably distinct failure mechanisms. All headline theorems are formalized in Lean 4 against Mathlib with warnings promoted to errors and no admitted statements.

Keywords. autocatalysis; autocatalytic cores; chemical reaction networks; mass-action kinetics; thermodynamic consistency; transition-state theory; phase diagram; toric geometry; formal verification.

MSC 2020. 92E20, 80A30, 90C05, 14M25, 68V20.

1 Introduction

Autocatalysis, the ability of a set of chemical species to catalyse its own net production, is widely regarded as a prerequisite for the emergence of Darwinian dynamics in chemistry [1–3]. Blokhuis, Lacoste and Nghe [1] identified autocatalysis with a stoichiometric property of a reaction network: a motif is autocatalytic when some reaction current strictly increases the concentration of each of its species, and they classified the minimal such motifs, the *autocatalytic cores*, into five topological types. Cores are purely combinatorial objects. They

know nothing about rate constants, concentrations or free energies, and the question of whether a stoichiometrically autocatalytic motif can actually run under thermodynamically consistent kinetics was left open by that classification.

Kosc, Kuperberg, Rajon and Charlat [4] addressed this question. They defined a *potential autocatalytic core* (PAC) as a minimal autocatalytic motif of a reaction network and a *consistent autocatalytic core* (CAC) as a PAC whose productive current is the current induced, under mass-action kinetics constrained by local detailed balance, by some positive concentration vector. They proved that PAC detection is NP-complete, and, more importantly for the present paper, that *every PAC is a CAC*: for any activation barriers and any standard free energies, an isolated PAC can always be made to run in the productive direction if concentrations are unbounded. They also observed that this realizability does not survive interaction. Their Figure 4 exhibits two PACs which share reactions, which admit a common productive current (a *multiPAC*), and which cannot be simultaneously consistent: no single concentration vector makes both cores productive, whatever the barriers.

This paper is about the structure hidden in that observation. When several cores share species and reactions, a common realization has to satisfy simultaneously the polyhedral productivity constraints of each core, the sign constraints on each reaction direction, the magnitude constraints imposed by fixed transition-state factors, and the algebraic constraints that tie the activity of each stoichiometric complex to a monomial in the species activities. It is natural to ask whether the polyhedral and sign constraints already decide the matter, as they do for one core. We call this the Strong Compatibility Hypothesis: *if a family of cores admits a common productive current and a common strict direction potential, then it admits a common thermodynamic realization*. Our first result is that this hypothesis is false, and the way it fails organizes the rest of the paper.

1.1 Main results

We work with a finite reversible reaction network Q in which reaction r has reactant complex α_r , product complex β_r and a positive *transition-state factor* b_r ; for a positive activity vector z the oriented current is $j_r(z) = b_r(z^{\alpha_r} - z^{\beta_r})$. This is the Kosc et al. model with $b_r = e^{-G_r^\ddagger/RT}$ and $z_s = e^{\mu_s/RT}$ (Section 2). A *core family* is a finite collection of PACs of Q together with a global orientation of every reaction. Four predicates on a core family are studied:

- **MultiPAC**: some current v has the prescribed sign on every reaction and makes every core productive;
- **Dir** (direction compatibility): some potential $\mu \in \mathbb{R}^S$ has strictly positive oriented affinity $\varepsilon_r \langle \alpha_r - \beta_r, \mu \rangle$ on every reaction;
- **LinCx** (linear-complex compatibility): some assignment of an independent positive activity y_c to every distinct complex c makes the currents $b_r(y_{\alpha_r} - y_{\beta_r})$ correctly signed and every core productive;
- **MultiCAC**: some positive species vector z makes the currents $j_r(z)$ correctly signed and every core productive.

MultiCAC implies each of the other three (Lemma 2.6); the three are otherwise independent filters on different projections of the problem. Our results are as follows; each is stated informally here and precisely in the body of the paper.

- (A) **Strong Compatibility fails (Theorem 3.4)**. The two autocatalytic triangles $A \rightleftharpoons B \rightleftharpoons C \rightleftharpoons 2A$ and $A \rightleftharpoons B \rightleftharpoons D \rightleftharpoons 2A$ sharing $A \rightleftharpoons B$, with factors $(1, 1, 1, 10, 10)$, form a family that is MultiPAC and Dir but not MultiCAC. The obstruction already defeats LinCx: it is a conflict between fixed factor magnitudes, not between reaction directions.

The original example of Kosc et al. fails for a different reason: it is not Dir, a signed circuit of the six reactions forbids the required orientation (Proposition 3.2).

- (B) **Exact phase diagram (Theorem 4.1).** For the same two triangles with private factors $(1, 1)$ on the left and (k, k) on the right, both cores are PACs and the family is MultiPAC and Dir for every $k > 0$, while it is MultiCAC exactly when $1/2 < k < 2$. The endpoints are excluded because a strict productivity residual degenerates there. With arbitrary factors $b_0; b_1, b_2; b_3, b_4$ the condition becomes $\frac{1}{2} < (1/b_3 + 1/b_4)/(1/b_1 + 1/b_2) < 2$ (Proposition 4.4).
- (C) **Interface calculus (Theorem 5.2, Corollary 5.3).** Eliminating a gain- m triangle $A \rightarrow B \rightarrow C \rightarrow mA$ with factors b_0, b_1, b_2 from the literal productive-current equations leaves the exact open interval $(R/m, R)$, $R = b_0(1/b_1 + 1/b_2)$, of admissible values for the ratio q of the private complex-activity drop to the shared drop. Two such intervals with gains m_L, m_R and coefficients R_L, R_R intersect exactly when $1/m_L < R_R/R_L < m_R$. The phase diagram of (B) is the case $m_L = m_R = 2$.
- (D) **Monomial-lift obstruction (Theorem 6.1).** The family $\{A \rightleftharpoons B, B \rightleftharpoons 2A\}$, $\{A \rightleftharpoons B, B \rightleftharpoons 3A\}$ with factors $(1, 1, 2)$ on two species and three reactions is MultiPAC, Dir and LinCx but not MultiCAC. At the linear-complex level the two private response intervals are independent; the monomial lift couples them through the common activity of A and empties their intersection.
- (E) **Formal verification (Section 7).** The statements in (A)–(D), including PAC minimality of every core involved, have been compiled in Lean 4 [5] against Mathlib [6] with warnings treated as errors and no unproved declarations. The few hand-proved extensions are listed explicitly in Section 7.4.

The picture that emerges is a layered one (Figure 1). A common realization can fail at the level of reaction directions (a signed circuit), at the level of fixed magnitudes (a linear-complex wall) or at the level of the monomial lift (a toric-image wall), and the three examples of this paper show that each layer can fail while the previous ones hold. For the smallest interacting topology, two triangles with one shared edge, the complete answer is a single scalar interval, and its endpoints have a direct physical reading: in the symmetric family the two private branches must have Gibbs activation energies within $RT \ln 2$, about 1.7 kJ mol^{-1} at 298 K, of each other (Section 8.1).

1.2 Relation to previous work

The framework of autocatalytic cores is due to Blokhuis, Lacoste and Nghe [1]; its diluted-regime dynamics were developed by Unterberger and Nghe [7] and Nandan, Nghe and Unterberger [8], and cores with explicit catalysis are treated by Golnik, Gatter, Stadler and Vassena [9]. The polyhedral geometry of minimal autocatalytic subnetworks is studied by Gagrani, Blanco, Smith and Baum [10]; the NP-completeness of recognizing autocatalysis goes back to Andersen, Flamm, Merkle and Stadler [11] and, for PAC detection, to Kosc et al. [4]. Vassena and Stadler [12] identify unstable cores as the source of dynamical instability; that question is orthogonal to ours, which concerns existence of a positive static witness. The relationship between cores and the RAF sets of Hordijk and Steel [2, 3] is clarified in [13]. Despons, De Decker and Lacoste [14] derive stoichiometric bounds on affinities at optimal flux; their regime, a steady state at maximal production, differs from the existence question treated here. The word “toric” is used below only for the monomial image of the species activities; it does not refer to the toric dynamical systems of Craciun, Dickenstein, Shiu and Sturmfels [15], which concern complex-balanced steady states in the sense of Feinberg [16].

To our knowledge the exact factor-of-two window, the source-level interface interval and the separation of the three obstruction layers are new. This assessment is based on a bounded primary-source search completed on 31 August 2026 and is not a proof of absence.

1.3 Organization

Section 2 fixes the model and the four compatibility predicates. Section 3 refutes the Strong Compatibility Hypothesis. Section 4 proves the phase diagram, Section 5 the interface calculus and Section 6 the monomial-lift obstruction. Section 7 documents the formal verification and Section 8 discusses the physical reading, limitations and open problems.

2 Setting

2.1 Reversible networks and barrier-weighted currents

Let $\mathcal{S}$ and $\mathcal{R}$ be finite sets of species and reactions. A *complex* is a vector $c \in \mathbb{N}^{\mathcal{S}}$; for $z \in \mathbb{R}_{>0}^{\mathcal{S}}$ write $z^c = \prod_s z_s^{c_s}$.

Definition 2.1 (Reversible network). A reversible reaction network is $Q = (\mathcal{S}, \mathcal{R}, \alpha, \beta, b)$ where $\alpha_r, \beta_r \in \mathbb{N}^{\mathcal{S}}$ are the reactant and product complexes of $r \in \mathcal{R}$ and $b_r > 0$ is its transition-state factor. Its stoichiometric matrix is $N_{sr} = \beta_r(s) - \alpha_r(s)$. For $z \in \mathbb{R}_{>0}^{\mathcal{S}}$ the *current* of r is

$$j_r(z) = b_r(z^{\alpha_r} - z^{\beta_r}). \quad (1)$$

Equation (1) is the mass-action current of Kosci et al. [4, Eq. 5] in the ideal-solution regime with $RT = 1$: if μ_s is the chemical potential of species s and $G_r^\ddagger$ the Gibbs activation energy of reaction r , local detailed balance gives $v_r = e^{-G_r^\ddagger} (e^{\langle \alpha_r, \mu \rangle} - e^{\langle \beta_r, \mu \rangle})$, which is (1) with $z_s = e^{\mu_s}$ and $b_r = e^{-G_r^\ddagger}$. The factor b_r is a kinetic prefactor: it scales both directions of r equally and does not affect the equilibrium constant. Throughout, b_r is fixed and the free variable is z . Because μ_s absorbs the standard free energy of s , the results below hold for all standard free energies; the only thermodynamic input is the set of factors b_r .

2.2 Motifs and potential autocatalytic cores

The following definitions transcribe Definitions 1 and 2 of [4] in the notation of Definition 2.1; they are the definitions used in the formalization.

Definition 2.2 (Motifs). A *motif* of Q is a pair $C = (\mathcal{S}_C, \mathcal{R}_C)$ with $\mathcal{S}_C \subseteq \mathcal{S}$ and $\mathcal{R}_C \subseteq \mathcal{R}$. It is *side-incident* if $\mathcal{S}_C$ and $\mathcal{R}_C$ are nonempty and every $r \in \mathcal{R}_C$ has a reactant and a product in $\mathcal{S}_C$. A current $v \in \mathbb{R}^{\mathcal{R}}$ is *productive* for C if

$$\sum_{r \in \mathcal{R}_C} N_{sr} v_r > 0 \quad \text{for every } s \in \mathcal{S}_C.$$

C is an *autocatalytic motif* if it is side-incident and some current is productive for it. A motif C' is a *strict submotif* of C if $\mathcal{S}_{C'} \subseteq \mathcal{S}_C$, $\mathcal{R}_{C'} \subseteq \mathcal{R}_C$ and at least one inclusion is proper. A *potential autocatalytic core* (PAC) is an autocatalytic motif with no autocatalytic strict submotif.

Productivity only involves the reactions of the motif; reactions of Q outside $\mathcal{R}_C$ may consume species of $\mathcal{S}_C$, as they do in every example below. A productive current for a PAC may have entries of either sign, so PAC minimality must be proved with both species balances of each reaction in hand; Section 6 records a case where a one-balance argument fails.

Lemma 2.3 (PAC status depends only on chemistry). *Let Q and Q' have the same species, reactions, reactant map and product map, and arbitrary positive factors. A motif is a PAC of Q if and only if it is a PAC of Q' .*

Proof. Side incidence, productivity and strict submotifs are defined from α , β and N alone, and N is determined by α and β . Hence autocatalytic motifs of Q and Q' coincide, and so do PACs. $\square$

Lemma 2.3 is trivial but it is used essentially: it lets the parameterized phase diagram of Section 4 be stated for literal core families at every value of the parameter, rather than for a fixed fixture with the factors then varied informally.

2.3 Core families and the four compatibility predicates

Definition 2.4 (Core family). A *core family* on Q consists of a finite index set $\mathcal{K}$, a PAC C_k of Q for each $k \in \mathcal{K}$, an *orientation* $\varepsilon : \mathcal{R} \rightarrow \{+1, -1\}$, and an *owner* map $o : \mathcal{R} \rightarrow \mathcal{K}$ with $r \in \mathcal{R}_{C_{o(r)}}$ for every r . In particular every reaction of Q belongs to at least one core.

The orientation records the direction in which each reaction is required to run; it is global, so a shared reaction has one direction for all cores containing it. The owner map is bookkeeping and plays no role in the theorems.

Definition 2.5 (Compatibility predicates). Let $\mathcal{F}$ be a core family on Q .

- (i) $\mathcal{F}$ is **MultiPAC** if there is $v \in \mathbb{R}^{\mathcal{R}}$ with $\varepsilon_r v_r > 0$ for all r such that v is productive for every C_k .
- (ii) $\mathcal{F}$ is **MultiCAC** if there is $z \in \mathbb{R}_{>0}^{\mathcal{S}}$ with $\varepsilon_r j_r(z) > 0$ for all r such that $j(z)$ is productive for every C_k .
- (iii) $\mathcal{F}$ is *direction compatible*, **Dir**, if there is $\mu \in \mathbb{R}^{\mathcal{S}}$ with $\varepsilon_r \langle \alpha_r - \beta_r, \mu \rangle > 0$ for all r .
- (iv) $\mathcal{F}$ is *linear-complex compatible*, **LinCx**, if there is a positive function y on the set of complexes $\{\alpha_r, \beta_r : r \in \mathcal{R}\}$ such that the linear currents $j_r^{\text{lin}}(y) = b_r(y_{\alpha_r} - y_{\beta_r})$ satisfy $\varepsilon_r j_r^{\text{lin}}(y) > 0$ for all r and are productive for every C_k .

MultiPAC is the multiPAC notion of [4], with the orientation made explicit; MultiCAC is the corresponding notion of simultaneous consistency. Dir asks only that the reaction directions be jointly realizable by some chemical potential; it ignores magnitudes and productivity. LinCx keeps the fixed factors and the productivity inequalities but forgets that complex activities are monomials in species activities: each distinct complex c receives an independent unknown y_c in place of z^c . It is a strict linear feasibility problem and therefore decidable by linear programming. MultiCAC is LinCx together with the requirement that y lie in the image of the monomial map $z \mapsto (z^c)_c$; we refer to this last requirement as the *monomial lift*.

Lemma 2.6 (Necessary conditions). *MultiCAC implies MultiPAC, Dir and LinCx.*

Proof. Let z witness MultiCAC. Then $v = j(z)$ witnesses MultiPAC, and $y_c = z^c$ witnesses LinCx because $j_r^{\text{lin}}(y) = j_r(z)$. For Dir take $\mu = \log z$: since $b_r > 0$, $\varepsilon_r j_r(z) > 0$ holds if and only if $\varepsilon_r (z^{\alpha_r} - z^{\beta_r}) > 0$, if and only if $\varepsilon_r \langle \alpha_r - \beta_r, \log z \rangle > 0$. $\square$

Direction compatibility has a classical certificate of failure.

Lemma 2.7 (Signed circuits). *Let $\mathcal{F}$ be a core family. A signed circuit is a nonzero $w \in \mathbb{R}_{\geq 0}^{\mathcal{R}}$ with $\sum_r w_r \varepsilon_r (\alpha_r - \beta_r) = 0$ in $\mathbb{R}^{\mathcal{S}}$. If $\mathcal{F}$ has a signed circuit then it is not Dir; conversely, if $\mathcal{F}$ is not Dir then it has a signed circuit.*




Figure 1: The four compatibility predicates. MultiCAC implies the three others (Lemma 2.6); the examples of this paper show that a family can satisfy any prefix of the sequence MultiPAC, Dir, LinCx and fail the next requirement.

Proof. If w is a signed circuit and μ a direction potential then $0 = \sum_r w_r \varepsilon_r \langle \alpha_r - \beta_r, \mu \rangle > 0$, a contradiction. The converse is Gordan’s theorem of the alternative: the strict system $M\mu > 0$ with rows $M_r = \varepsilon_r(\alpha_r - \beta_r)$ is infeasible if and only if $w^\top M = 0$ has a nonzero nonnegative solution [17, Cor. 7.1e]. $\square$

2.4 Why isolated realizability does not compose

Theorem 2 of [4] states that for a single PAC and any current v compatible with the directions of its reactions, there exist $\lambda > 0$ and a positive concentration vector inducing λv . The proof solves for concentrations reaction by reaction along the core, which is possible because a single PAC never has to reconcile two different currents through one reaction. Definition 2.5 isolates the three ways this reconciliation can fail when cores interact: the directions may be jointly unrealizable (Dir fails), the productive current magnitudes required by the different cores through the shared reactions may be inconsistent with the fixed factors (LinCx fails), or the complex activities forced by the currents may fail to be monomials in a common species vector (the monomial lift fails). The remainder of the paper exhibits each failure and proves that, for the smallest interacting topology, the second failure is governed by an exact scalar window.


3 The Strong Compatibility Hypothesis is false

Hypothesis 3.1 (Strong Compatibility). For every core family $\mathcal{F}$ on every reversible network, $\text{MultiPAC} \wedge \text{Dir} \Rightarrow \text{MultiCAC}$.

Before refuting the hypothesis we revisit the example that motivates it.


3.1 The example of Kosc et al. is a direction obstruction

Figure 4 of [4] consists of two cores on the species $e_1, e_2, e_3, e_4, e'_1, e'_2, e_A, e_B$ and six reactions, read in the productive direction:



The two cores are

$$(\{e_1, e_2, e_3, e_4\}, \{r_1, r_2, r_3, r_4\}) \quad \text{and} \quad (\{e'_1, e'_2, e_3, e_4\}, \{r'_1, r_2, r_3, r'_4\});$$



left core $C_L = (\{A, B, C\}, \{r_0, r_1, r_2\})$
 right core $C_R = (\{A, B, D\}, \{r_0, r_3, r_4\})$
 shared reaction r_0

symmetric family Q_k :
 $(b_0, b_1, b_2, b_3, b_4) = (1, 1, 1, k, k)$

Figure 2: The two-triangle network. Each triangle is a Type-I autocatalytic core with gain two; the cores share the reaction $A \rightleftharpoons B$.

they share r_2 and r_3 , and the external species e_A, e_B are retained in the ambient network.


Proposition 3.2. *For every choice of positive factors, the two cores above are PACs, the family with all reactions oriented forward is MultiPAC, and it is neither Dir nor MultiCAC.*

Proof. The current $v = (v_1, v_2, v_3, v_4, v'_1, v'_4) = (6, 5, 4, 7, 6, 7)$ gives production residual exactly 1 on every species of each core: for instance e_4 receives $2v_3 - v_4 = 1$ from the left core and $2v_3 - v'_4 = 1$ from the right. PAC minimality is a finite case analysis on submotifs, carried out in the formalization (Section 7). For direction incompatibility, note that the six displacement vectors $\alpha_r - \beta_r$ sum to zero: the multiset of consumed species $\{e_1, e_A, e_4, e'_1, e_B, e_4, e_2, e'_2, e_3\}$ equals the multiset of produced species. Hence $w = (1, 1, 1, 1, 1, 1)$ is a signed circuit and Lemma 2.7 applies. Finally MultiCAC fails by Lemma 2.6; the formalization proves it directly by multiplying the monomial inequalities $z_{e_B} z_{e_2} < z_{e_A} z_{e_4}$, $z_{e_A} z_{e'_2} < z_{e_B} z_{e_4}$ and $z_{e_4}^2 < z_{e_2} z_{e'_2}$ that positive forward currents would force. $\square$

So the original incompatibility is purely topological: the six reactions form a closed stoichiometric cycle, and no chemical potential can drive every reaction of such a cycle forward. The barriers play no role. Hypothesis 3.1 is precisely the statement that direction obstructions are the only ones. They are not.

3.2 Two triangles with a shared edge

Let Q be the network of Figure 2 on species A, B, C, D with reactions



factors $b_0, \dots, b_4 > 0$, and the two motifs $C_L = (\{A, B, C\}, \{r_0, r_1, r_2\})$ and $C_R = (\{A, B, D\}, \{r_0, r_3, r_4\})$. All reactions are oriented forward, from left to right as written.

Lemma 3.3. *C_L and C_R are PACs of Q for every choice of factors. The family $\mathcal{F} = \{C_L, C_R\}$ is MultiPAC, witnessed by $v = (1, \frac{4}{5}, \frac{3}{5}, \frac{4}{5}, \frac{3}{5})$, and Dir, witnessed by $\mu = (\mu_A, \mu_B, \mu_C, \mu_D) = (-1, -\frac{3}{2}, -\frac{9}{5}, -\frac{9}{5})$.*

Proof. For v as stated, each of the six production residuals equals $\frac{1}{5}$: in C_L , species A receives $-v_0 + 2v_2$, B receives $v_0 - v_1$ and C receives $v_1 - v_2$, and symmetrically in C_R . The oriented affinities of μ are $\mu_A - \mu_B = \frac{1}{2}$, $\mu_B - \mu_C = \frac{3}{10}$, $\mu_C - 2\mu_A = \frac{1}{5}$, $\mu_B - \mu_D = \frac{3}{10}$ and $\mu_D - 2\mu_A = \frac{1}{5}$, all positive. Both witnesses are independent of the factors.

For minimality of C_L (the case of C_R is symmetric), let C' be an autocatalytic strict submotif with productive current v . Every reaction of C' lies in $\{r_0, r_1, r_2\}$; the ambient

reactions r_3, r_4 are absent. If $r_0 \in \mathcal{R}_{C'}$, side incidence puts $A, B \in \mathcal{S}_{C'}$, and productivity of A and B within C' reads $-v_0 + 2[r_2 \in C']v_2 > 0$ and $v_0 - [r_1 \in C']v_1 > 0$; similarly $r_1 \in C'$ forces $B, C \in \mathcal{S}_{C'}$ with $[r_0 \in C']v_0 - v_1 > 0$ and $v_1 - [r_2 \in C']v_2 > 0$, and $r_2 \in C'$ forces $C, A \in \mathcal{S}_{C'}$ with $[r_1 \in C']v_1 - v_2 > 0$ and $-[r_0 \in C']v_0 + 2v_2 > 0$. Checking the seven nonempty subsets of $\{r_0, r_1, r_2\}$ shows that each proper subset yields contradictory sign conditions (for example $\{r_0\}$ alone requires $-v_0 > 0$ and $v_0 > 0$), so $\mathcal{R}_{C'} = \{r_0, r_1, r_2\}$, and side incidence then forces $\mathcal{S}_{C'} = \{A, B, C\}$. Thus C' is not strict. $\square$

Theorem 3.4 (Refutation of Strong Compatibility). *With factors $(b_0, \dots, b_4) = (1, 1, 1, 10, 10)$, the family $\mathcal{F}$ is MultiPAC and Dir but not LinCx, hence not MultiCAC. In particular Hypothesis 3.1 is false.*

Proof. By Lemma 3.3 it remains to show that LinCx fails. Suppose y is a positive function on the complexes $A, B, C, D, 2A$ whose linear currents $j_0 = y_A - y_B$, $j_1 = y_B - y_C$, $j_2 = y_C - y_{2A}$, $j_3 = 10(y_B - y_D)$, $j_4 = 10(y_D - y_{2A})$ are positive and productive for both cores. Productivity of C_L (species A, B, C in turn) gives $2j_2 > j_0$, $j_0 > j_1$ and $j_1 > j_2$, so

$$j_1 + j_2 > 2j_2 > j_0.$$

Productivity of C_R gives $2j_4 > j_0$, $j_0 > j_3$ and $j_3 > j_4$, and the last two inequalities give

$$j_3 + j_4 < 2j_0.$$

Now the private complex drops telescope: $j_1 + j_2 = y_B - y_{2A}$ and $j_3 + j_4 = 10(y_B - y_{2A})$. Writing $x = y_B - y_{2A}$ we obtain $x > j_0$ and $10x < 2j_0$, which is impossible for $j_0 > 0$. $\square$

The proof never uses that y comes from species activities, so the obstruction lives in the linear layer: the two branches leaving B and returning to $2A$ must carry the same complex-activity drop $y_B - y_{2A}$, the left branch needs that drop to exceed the shared drop j_0 , and with factor 10 the right branch needs it to be below $j_0/5$. This is a conflict of magnitudes, and it suggests that the ratio of the private factors is the controlling parameter.

4 The exact phase diagram

For $k > 0$ let Q_k be the network of Figure 2 with factors $(1, 1, 1, k, k)$, and let $\mathcal{F}_k$ be the family $\{C_L, C_R\}$ on Q_k with all reactions oriented forward. By Lemma 2.3 and Lemma 3.3, $\mathcal{F}_k$ is a core family, MultiPAC and Dir for every $k > 0$.

Theorem 4.1 (Two-triangle phase diagram). *For every $k > 0$,*

$$\mathcal{F}_k \text{ is MultiCAC} \iff \frac{1}{2} < k < 2.$$

Proof. Necessity. Let $z = (A, B, C, D)$ be a positive witness. The currents are

$$j_0 = A - B, \quad j_1 = B - C, \quad j_2 = C - A^2, \quad j_3 = k(B - D), \quad j_4 = k(D - A^2),$$

all positive. Put $j = j_0 > 0$ and $x = B - A^2$. As in the proof of Theorem 3.4, productivity of C_L gives $j < j_1 + j_2 < 2j$, and $j_1 + j_2 = x$; productivity of C_R gives $j < j_3 + j_4 < 2j$, and $j_3 + j_4 = kx$. Hence $x \in (j, 2j)$ and $kx \in (j, 2j)$. If $k \leq \frac{1}{2}$ then $kx \leq x/2 < j$, and if $k \geq 2$ then $kx \geq 2x > 2j$; either way a contradiction.

Sufficiency. Assume $\frac{1}{2} < k < 2$ and put

$$q = \frac{3}{1+k}, \quad p = kq = \frac{3k}{1+k}.$$

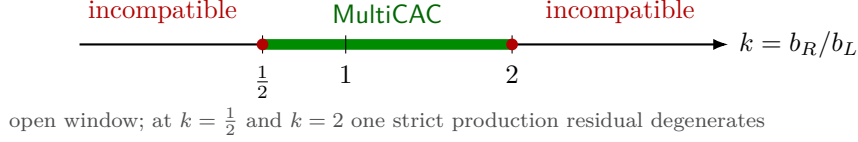


Figure 3: Theorem 4.1. For every $k > 0$ the family is MultiPAC and Dir; a common realization exists exactly in the open window $\frac{1}{2} < k < 2$. In the symmetric transition-state family this is $|G_R^\ddagger - G_L^\ddagger| < RT \ln 2$ (Section 8.1).

Then $1 < q < 2$ and $1 < p < 2$: indeed $q > 1$ iff $k < 2$, $q < 2$ iff $k > \frac{1}{2}$, $p > 1$ iff $k > \frac{1}{2}$ and $p < 2$ iff $k < 2$. For $1 < t < 2$ define

$$\delta(t) = \frac{(t-1)(2-t)}{4}, \quad u(t) = \frac{t}{2} + \delta(t), \quad v(t) = \frac{t}{2} - \delta(t),$$

so that $u(t) + v(t) = t$ and $\frac{1}{2} < v(t) < u(t) < 1$; the last chain holds because $0 < \delta(t) < \min\{\frac{t-1}{2}, \frac{2-t}{2}\}$. Let $(u_L, v_L) = (u(q), v(q))$ and $(u_R, v_R) = (u(p), v(p))$, and set

$$j = \frac{1}{4(q+1)}, \quad A = \frac{1}{2}, \quad B = A - j, \quad C = B - u_L j, \quad D = B - \frac{u_R j}{k}.$$

Then $j < \frac{1}{8}$, $B > 0$, $C > \frac{1}{2} - 2j > 0$, and $B - A^2 = \frac{1}{4} - j = qj$ by the choice of j . The currents are $j_0 = j$, $j_1 = u_L j$, $j_2 = (B - A^2) - u_L j = v_L j$, $j_3 = u_R j$ and $j_4 = k(B - A^2) - u_R j = (kq - u_R)j = v_R j$; all are positive. Since $u_R < p = kq$ we have $D > B - qj = A^2 > 0$. Productivity of C_L reads $(2v_L - 1)j > 0$, $(1 - u_L)j > 0$ and $(u_L - v_L)j > 0$, and productivity of C_R reads the same with (u_R, v_R) ; all hold. Hence $z = (A, B, C, D)$ witnesses MultiCAC. $\square$

Remark 4.2 (Endpoints). At $k = 2$ the necessity argument gives $2x \in (j, 2j)$ and $x \in (j, 2j)$, whose closures meet only at $x = j$, where the residual $j_1 + j_2 - j_0$ of species A in C_L vanishes. The same happens at $k = \frac{1}{2}$ with the roles of the cores exchanged. The window is open because productivity is strict; a weak version of productivity would close it.

Remark 4.3 (Scale invariance). Only the ratio of private to shared factors enters. Replacing all five factors by λb_r multiplies every current by λ and changes nothing. Replacing the shared factor b_0 by λb_0 scales j_0 alone and is absorbed by the shared drop $A - B$; the general-factor statement below makes this explicit.

Proposition 4.4 (Arbitrary factors). *Let $b_0, \dots, b_4 > 0$ be arbitrary and put $R_L = 1/b_1 + 1/b_2$ and $R_R = 1/b_3 + 1/b_4$. The family $\{C_L, C_R\}$ on the network of Figure 2 is MultiCAC if and only if*

$$\frac{1}{2} < \frac{R_R}{R_L} < 2.$$

Proof. For a positive witness z , write $j_0 = b_0(A - B)$, $j_1 = b_1(B - C)$, $j_2 = b_2(C - A^2)$, $j_3 = b_3(B - D)$, $j_4 = b_4(D - A^2)$, and let $q = (B - A^2)/(A - B)$ be the ratio of the private complex drop to the shared drop. Then

$$q = \frac{j_1/b_1 + j_2/b_2}{j_0/b_0} = b_0 \left(\frac{u_L}{b_1} + \frac{v_L}{b_2} \right), \quad u_L = \frac{j_1}{j_0}, \quad v_L = \frac{j_2}{j_0},$$

and productivity of C_L is exactly $\frac{1}{2} < v_L < u_L < 1$. By Theorem 5.2 below, the set of values of q compatible with C_L is the open interval $(b_0 R_L/2, b_0 R_L)$; likewise for C_R it is $(b_0 R_R/2, b_0 R_R)$. Because both branches leave B and return to $2A$, the ratio q is literally the same quantity for both cores, so a common witness forces the two intervals to intersect, which by Corollary 5.3 with




Figure 4: Theorem 5.2. Literal productive currents of a gain- m triangle, the normalized productive cone, and the exact interval of the interface response q .

$m_L = m_R = 2$ is the displayed condition. Conversely, given q in both intervals, Theorem 5.2 provides (u_L, v_L) and (u_R, v_R) with $\frac{1}{2} < v < u < 1$ realizing q on each side. Choose $A = \frac{1}{2}$ and $B = (A^2 + qA)/(1 + q) \in (A^2, A)$, so that $B - A^2 = q(A - B)$; put $j_0 = b_0(A - B)$, $C = B - u_L j_0/b_1$ and $D = B - u_R j_0/b_3$. Then $C - A^2 = (B - A^2) - u_L j_0/b_1 = v_L j_0/b_2 > 0$ by the defining identity of q , so $C > A^2 > 0$, and the currents through C_L are $j_0, u_L j_0, v_L j_0$; the same holds on the right. Productivity follows as in Theorem 4.1. $\square$

The symmetric family has $R_L = 2$ and $R_R = 2/k$, so Proposition 4.4 recovers Theorem 4.1. The proposition is proved here by hand on top of the formally verified interval theorem; the formalization covers the literal symmetric family (Section 7.4).

5 The interface calculus

The proofs of Section 4 used one scalar per core: the ratio of the complex-activity drop across the private branch to the drop across the shared reaction. This section shows that the scalar carries all the information of a triangular core, for arbitrary gain and factors, and that two eliminated cores compose by intersecting intervals.

Definition 5.1 (Literal triangle currents). Let $m > 1$ and $b_0, b_1, b_2 > 0$. Real numbers (j_0, j_1, j_2) are *literal productive currents* of the triangle $A \rightarrow B \rightarrow C \rightarrow mA$ with these factors if

$$j_0 > 0, \quad m j_2 - j_0 > 0, \quad j_1 - j_2 > 0, \quad j_0 - j_1 > 0, \quad (2)$$

that is, if j_0 is forward and the net production of A , C and B is strictly positive. The *interface response* of such currents is

$$q = b_0 \left(\frac{j_1}{b_1 j_0} + \frac{j_2}{b_2 j_0} \right). \quad (3)$$

When the currents come from species activities, $j_1/b_1 + j_2/b_2 = (B - C) + (C - A^m) = B - A^m$ and $j_0/b_0 = A - B$, so q is the ratio of the private drop $B - A^m$ to the shared drop $A - B$ in complex-activity space. Definition 5.1 does not assume this; it works with currents alone. The gain m is allowed to be any real number greater than one, although only integers arise from stoichiometry.

Theorem 5.2 (Literal interface projection). Let $m > 1$ and $b_0, b_1, b_2 > 0$, and put $R = b_0(1/b_1 + 1/b_2)$. A real number q is the interface response of some literal productive currents if and only if

$$\frac{R}{m} < q < R.$$

Proof. Normalize by $j_0 > 0$: with $u = j_1/j_0$ and $v = j_2/j_0$, the conditions (2) become $\frac{1}{m} < v < u < 1$, and (3) becomes $q = xu + yv$ with $x = b_0/b_1$ and $y = b_0/b_2$; conversely any (u, v) in this cone gives literal currents $(1, u, v)$. So it suffices to show that the image of the open cone $\{\frac{1}{m} < v < u < 1\}$ under $(u, v) \mapsto xu + yv$ is exactly the interval $((x + y)/m, x + y)$.

If $\frac{1}{m} < v < u < 1$ then $xu + yv > x/m + y/m$ and $xu + yv < x + y$ because $x, y > 0$. Conversely let $(x + y)/m < q < x + y$ and put $S = x + y$, $t = q/S$, $\ell = 1/m$ and

$$e = \frac{(t - \ell)(1 - t)}{2S}, \quad u = t + ey, \quad v = t - ex.$$

Then $\ell < t < 1$, $e > 0$, and $xu + yv = St = q$. Moreover $ex < t - \ell$ because $(1 - t)x < x < 2S$, and $ey < 1 - t$ because $(t - \ell)y < y < 2S$; hence $\ell < v$ and $u < 1$, while $u - v = e(x + y) > 0$. Thus (u, v) lies in the cone and realizes q . $\square$

Corollary 5.3 (Composition of two interfaces). *Let $m_L, m_R > 1$ and $R_L, R_R > 0$. The intervals $(R_L/m_L, R_L)$ and $(R_R/m_R, R_R)$ intersect if and only if*

$$\frac{1}{m_L} < \frac{R_R}{R_L} < m_R.$$

Proof. Two nonempty open intervals (a, b) and (c, d) intersect if and only if $a < d$ and $c < b$. Here $a < d$ reads $R_L/m_L < R_R$ and $c < b$ reads $R_R/m_R < R_L$. $\square$

Corollary 5.3 is symmetric under exchanging the cores, as it must be: $1/m_L < R_R/R_L < m_R$ is equivalent to $1/m_R < R_L/R_R < m_L$. For the two triangles of Section 4, $m_L = m_R = 2$, $R_L = 1 + 1 = 2$ and $R_R = 2/k$, and the corollary gives $\frac{1}{2} < 1/k < 2$, that is $\frac{1}{2} < k < 2$.

Remark 5.4 (Reading the interval). The interval $(R/m, R)$ has a simple meaning. The two private reactions carry the entire private drop $B - A^m$; the fastest way to spend it is to put all of it on the cheaper reaction, and the productive cone forbids either reaction from carrying less than a $1/m$ share of the shared current or more than the full share. The coefficient R is the sum of reciprocal private factors, a “resistance” of the branch, weighted by the shared factor b_0 . A branch with large private factors (low activation energies) has small R and can only carry a small drop relative to the shared one; a branch with small factors needs a large drop. Two branches sharing the same drop must therefore have resistances within the factor m of each other, and for gain two within a factor of two.


Remark 5.5 (When the scalar is not enough). The composition rule intersects two intervals of one variable q . This is the exact source-level condition whenever the two private branches have the same product complex, because then the private drops coincide. If the branches have different gains, the private drops $B - A^{m_L}$ and $B - A^{m_R}$ are different quantities related through the common activity of A , and at the species level $q_L \neq q_R$ in general. At the linear-complex level, however, y_{m_LA} and y_{m_RA} are independent unknowns and the two responses decouple. The next section exploits exactly this gap.

6 A monomial-lift obstruction

Theorem 3.4 showed that LinCx can fail while MultiPAC and Dir hold. This section shows that LinCx itself is not sufficient: the requirement that complex activities be monomials of one species vector can be the only thing that fails.

6.1 Two cores on two species

Let Q^{tor} be the network on species A, B with the three distinct reactions



all oriented forward, and let $C_1 = (\{A, B\}, \{r_0, r_1\})$ and $C_2 = (\{A, B\}, \{r_0, r_2\})$. The two cores share r_0 and differ in the gain of their return reaction. No parallel or duplicate reaction channels are used; r_1 and r_2 have different product complexes.




Figure 5: The monomial-lift obstruction of Theorem 6.1. Each core admits a nonempty interval of interface responses, and independent complex activities can place q_1 and q_2 in their intervals simultaneously. A common species vector forces $A < 1$, hence $B - A^3 > B - A^2$ and $q_2 > q_1$, which is incompatible with $q_1 > \frac{1}{2} > q_2$.

Theorem 6.1 (Monomial-lift obstruction). C_1 and C_2 are PACs of Q^{tor} . The family $\{C_1, C_2\}$ is MultiPAC, Dir and LinCx, but not MultiCAC.

Proof. PAC minimality. Consider C_1 ; the argument for C_2 is identical with r_2 in place of r_1 . Let C' be an autocatalytic strict submotif with productive current v , so $\mathcal{R}_{C'} \subseteq \{r_0, r_1\}$ is nonempty. Side incidence gives $\mathcal{S}_{C'} = \{A, B\}$ whenever $\mathcal{R}_{C'} \neq \emptyset$, because each of r_0, r_1 has A on one side and B on the other. If $\mathcal{R}_{C'} = \{r_0\}$, productivity of A and B reads $-v_0 > 0$ and $v_0 > 0$; if $\mathcal{R}_{C'} = \{r_1\}$ it reads $2v_1 > 0$ and $-v_1 > 0$. Both are contradictory, so $\mathcal{R}_{C'} = \{r_0, r_1\}$ and C' is not strict. Note that both balances are needed: the A -balance alone does not exclude $\{r_1\}$, since a productive current for a PAC may be negative on some reactions.

MultiPAC. The current $v = (1, \frac{3}{4}, \frac{1}{2})$ is positive; in C_1 species A receives $-1 + 2 \cdot \frac{3}{4} = \frac{1}{2}$ and B receives $1 - \frac{3}{4} = \frac{1}{4}$; in C_2 , A receives $-1 + 3 \cdot \frac{1}{2} = \frac{1}{2}$ and B receives $1 - \frac{1}{2} = \frac{1}{2}$.

Dir. With $\mu = (\mu_A, \mu_B) = (-1, -\frac{3}{2})$ the oriented affinities are $\mu_A - \mu_B = \frac{1}{2}$, $\mu_B - 2\mu_A = \frac{1}{2}$ and $\mu_B - 3\mu_A = \frac{3}{2}$.

LinCx. Assign $y_A = 4$, $y_B = 3$, $y_{2A} = \frac{9}{4}$, $y_{3A} = \frac{11}{4}$. The linear currents are $j_0 = 4 - 3 = 1$, $j_1 = 3 - \frac{9}{4} = \frac{3}{4}$ and $j_2 = 2(3 - \frac{11}{4}) = \frac{1}{2}$, which is the MultiPAC witness above; so they are positive and productive for both cores.

Not MultiCAC. Let $z = (A, B)$ be positive with positive productive currents $j_0 = A - B$, $j_1 = B - A^2$, $j_2 = 2(B - A^3)$. Productivity of A in C_1 gives $2(B - A^2) > A - B$; productivity of B in C_2 gives $A - B > 2(B - A^3)$. Positivity of j_0 and j_1 gives $A^2 < B < A$, hence $0 < A < 1$ and therefore $A^3 < A^2$, so $B - A^3 > B - A^2$. Combining,

$$A - B > 2(B - A^3) > 2(B - A^2) > A - B,$$

a contradiction. $\square$

In the language of Section 5, each core here is a two-reaction core with a single private reaction, and its interface response $q = b_0 j_{\text{priv}} / (b_{\text{priv}} j_0)$ ranges over $(b_0 / (m b_{\text{priv}}), b_0 / b_{\text{priv}})$; so $q_1 \in (\frac{1}{2}, 1)$ and $q_2 \in (\frac{1}{6}, \frac{1}{2})$. The linear witness realizes $q_1 = \frac{3}{4}$ and $q_2 = \frac{1}{4}$. Under the monomial lift, $q_1 = (B - A^2) / (A - B)$ and $q_2 = (B - A^3) / (A - B)$ are coupled by $q_2 > q_1$ whenever $A < 1$, and $A < 1$ is forced by the directions of r_0 and r_1 . The obstruction is that the linear witness has $y_{3A} > y_{2A}$, which no species vector with $A < 1$ can reproduce (Figure 5).

6.2 How the example was found

The example is the smallest member of a family isolated by a theory-conditioned search rather than by a blind one. For two cores $\{A \rightleftharpoons B, B \rightleftharpoons mA\}$ and $\{A \rightleftharpoons B, B \rightleftharpoons nA\}$ with $2 \leq m < n$, shared factor 1 and private factors b_m, b_n , the argument in the proof of Theorem 6.1

shows that a common species witness has $0 < A < 1$, hence $q_n > q_m$, while the productive cones require $q_m > 1/(mb_m)$ and $q_n < 1/b_n$. Therefore

$$\frac{1}{mb_m} \geq \frac{1}{b_n} \quad (4)$$

is a sufficient condition for the family to fail MultiCAC, while both cores individually, and the linear relaxation, are unconstrained by (4). Enumerating $2 \leq m < n \leq 8$ and integer $b_m, b_n \in \{1, \dots, 12\}$ with exact rational arithmetic produced 441 pairs satisfying (4); ordered by $b_m + b_n$, then b_n , then (m, n) , the first is $(m, n, b_m, b_n) = (2, 3, 1, 2)$, which is the network above. The enumeration is a discovery device: only the instance $(2, 3, 1, 2)$ is formally verified, and we make no claim of global minimality beyond source-minimality of the two PACs and minimality within the enumerated dominance family.

7 Formal verification

All theorems marked as formal in this section have been compiled in Lean 4 [5], toolchain `leanprover/lean4:v4.30.0`, against the Mathlib library [6] at the matching tagged release, with a frozen `lake-manifest`. Verification treats every warning as an error, rejects `sorry`, extracts the requested declaration signatures, reports the axioms used, and records SHA-256 hashes of the source file, the compiled module and the declaration interface. Every verified declaration depends only on the standard axioms `propext`, `Classical.choice` and `Quot.sound`.

7.1 Architecture

The development lives in the namespace `ThermoCoreCompatibility`, about 2,600 lines in nine files.

- `Reaction.lean` defines `ReversibleCRN` (Definition 2.1), complexes, complex activities and currents.
- `Source.lean` defines motifs, side incidence, productivity, strict submotifs and `IsPAC` (Definition 2.2), and proves `Motif.rebase_isPAC` (Lemma 2.3).
- `CoreFamily.lean` defines `CoreFamily`, `MultiPAC` and `MultiCAC` (Definitions 2.4 and 2.5).
- `Direction.lean` defines `DirectionCompatible`, `SignedCircuit` and proves the forward half of Lemma 2.7.
- `Toric.lean` defines `LinearComplexCompatible`, the toric realizability predicate on complex activities, and proves `multiCAC_iff_positivePolyhedralToricPoint` and `multiCAC_implies_linearComplexCompatible`.
- `InterfaceProfile.lean` contains the scalar calculus of Section 5, from the literal current predicate `LiteralTriangleCurrents` to `commonResponse_iff_ratio`.
- The four files `PaperIncompatiblePair.lean`, `MagnitudeObstructionPair.lean`, `TrianglePhaseDiagram.lean` and `ToricObstructionPair.lean` in the `Examples` directory contain the four literal networks with their PAC proofs and witnesses.
- `Main.lean` re-exports the publication-facing statements.

The parameterized phase diagram is stated for the literal family `familyAt k hk`, whose cores are PACs by `rebase_isPAC`:

```

theorem familyAt_complete_phaseDiagram (k : ℝ) (hk : 0 < k) :
  (familyAt k hk).MultiPAC ∧
  (familyAt k hk).DirectionCompatible ∧
  ((familyAt k hk).MultiCAC ↔ 1 / 2 < k ∧ k < 2)

```

Sufficiency is organized around a certificate structure `PairData` holding the nine scalars $A, B, C, D, j, u_L, v_L, u_R, v_R$ and the fourteen inequalities and identities of the proof of Theorem 4.1; a separate theorem replays the certificate against the literal network. This separation of scalar construction from network replay was decisive: a monolithic proof repeatedly exhausted the laborator’s resources, and factoring the certificate made both halves routine.

The interface theorem is stated at the level of literal currents and then projected:

```

theorem literalTriangle_interfaceProjection {m b0 b1 b2 q : ℝ}
  (hm : 1 < m) (hb0 : 0 < b0) (hb1 : 0 < b1) (hb2 : 0 < b2) :
  InterfaceProfile.LiteralBarrierTriangleResponse m b0 b1 b2 q ↔
  b0 * (1 / b1 + 1 / b2) / m < q ∧ q < b0 * (1 / b1 + 1 / b2)

```

and the monomial-lift obstruction is an existential over core families on the literal two-species network:

```

theorem exists_sourceValid_genuineToricIncompatible_pair :
  ∃ F : CoreFamily (Core := ToricPair.Core) ToricPair.network,
    F.MultiPAC ∧ F.DirectionCompatible ∧
    F.LinearComplexCompatible ∧ ¬ F.MultiCAC

```

The refutation of Hypothesis 3.1 is stated as the negation of the universal implication on the fixed five-reaction network, `strongCompatibility_refuted`.

7.2 Concordance

Table 1 lists the principal declaration behind each statement of the paper.

7.3 Receipts

Two publication-facing declarations were replayed from a clean temporary build and matched, field by field, against their recorded receipts: toolchain, source hash, compiled-artifact hash, declaration-interface hash, declaration name and empty warning streams. The recorded SHA-256 values are listed in Table 2. A tooling limit is worth recording: a probe that requested the combined biconditional of the phase diagram exhausted memory after the module had compiled successfully; the two directions were therefore authenticated separately and then combined in a further verified declaration, so that no tooling limit sits inside the mathematical trust chain.

7.4 What is proved by hand

The following statements are proved in this paper but are not part of the formal development, so that the reader can see precisely what has been machine-checked.

1. The MultiPAC and Dir halves of Lemma 2.6. The formal development proves the LinCx half and the equivalence of MultiCAC with a positive polyhedral point on the monomial image; the Dir half is the one-line logarithm argument.
2. The converse half of Lemma 2.7, which is Gordan’s theorem.
3. The signed circuit $w = (1, \dots, 1)$ of Proposition 3.2. The formal development proves $\neg \text{MultiCAC}$ for the paper pair directly, for all factors, but does not instantiate `SignedCircuit` on it.

Statement	Lean declaration (namespace ThermoCoreCompatibility)
Definitions 2.1–2.5	ReversibleCRN, Motif.IsPAC, CoreFamily, CoreFamily.MultiPAC, CoreFamily.MultiCAC, CoreFamily.DirectionCompatible, CoreFamily.LinearComplexCompatible
Lemma 2.3	Motif.rebase_isPAC
Lemma 2.6 (LinCx part)	CoreFamily.multiCAC_implies_linearComplexCompatible
Lemma 2.7 (forward)	CoreFamily.signedCircuit_not_directionCompatible
Prop. 3.2 (PAC, MultiPAC, $\neg$ MultiCAC, all factors)	PaperPair.left_isPAC, PaperPair.right_isPAC, PaperPair.paperPair_multiPAC_not_multiCAC
Lemma 3.3	MagnitudePair.left_isPAC, MagnitudePair.right_isPAC, MagnitudePair.family_multiPAC, MagnitudePair.family_directionCompatible, TrianglePhaseDiagram.left_isPACat, TrianglePhaseDiagram.right_isPACat
Theorem 3.4	MagnitudePair.magnitudeObstruction_is_nonToric, sourceValid_nonToric_magnitudeObstruction, strongCompatibility_refuted
Theorem 4.1	TrianglePhaseDiagram.familyAt_complete_phaseDiagram, triangleFamily_complete_phaseDiagram; the two directions are pairMultiCAC_onlyIf_window and window_implies_pairMultiCAC
Theorem 5.2	InterfaceProfile.weightedTriangleResponse_iff, InterfaceProfile.literalBarrierTriangleResponse_iff_interval, literalTriangle_interfaceProjection
Corollary 5.3	InterfaceProfile.commonResponse_iff_ratio
Theorem 6.1	ToricPair.left_isPAC, ToricPair.right_isPAC, ToricPair.genuineToricObstruction, exists_sourceValid_genuineToricIncompatible_pair

Table 1: Correspondence between the statements of this paper and the formal declarations.

4. Proposition 4.4, the arbitrary-factor form of the phase diagram. Its two ingredients, the interval theorem and the composition rule, are formal; the reconstruction of species activities from a common q follows the formal `PairData` construction with $b_0, \dots, b_4$ in place of $(1, 1, 1, k, k)$ but has not itself been replayed against the literal network.
5. The general dominance inequality (4) and the count of 441 enumerated pairs. Only the instance $(2, 3, 1, 2)$ is formal; the enumeration is a bounded exact-rational script.
6. The transition-state translation of Section 8.1, which is a change of variables.

Exact rational preflight scripts were used to check the witnesses of Lemma 3.3, the closed-form construction of Theorem 4.1 at sample values of k , and the interval formulas on sample factors before they were encoded; they are not part of any proof.

8 Discussion

8.1 Physical reading

In transition-state theory [18] the factor of reaction r is $b_r = \kappa_r (k_B T / h) e^{-G_r^\ddagger / RT}$, with a transmission coefficient κ_r and the Gibbs activation energy $G_r^\ddagger$; Kosc et al. absorb the nonexponential prefactor and write $b_r = e^{-G_r^\ddagger}$ in units of RT . Consider the *symmetric scaling family* in which both private reactions of each branch share one factor, $b_1 = b_2 = b_L$ and $b_3 = b_4 = b_R$, and both branches have equal nonexponential prefactors. Then $k = b_R / b_L =$

TrianglePhaseDiagram.familyAt_complete_phaseDiagram	
source	b869aee0ae4a581740858f1373d15f9106dc14d6d6c03061488f80a05a657157
artifact	bebbc2516686f3e848541b1e4d9f5845a72d83ae62243023fa010deccf177a7d
interface	b1b251239910c19142c78c07e6362bb7efe6a44a918822be9cf6adbe6d397ef3
InterfaceProfile.literalBarrierTriangleResponse_iff_interval	
source	5a09b1b55927180f32d7d71c69a566f795e119e44dc35b17f476a6dfec8f01d1
artifact	3a4f42b260a2caf157d97f860970f6e15f44f69841b96b4c1aaf3cc36f6a2e06
interface	4ef11e06c3b96cbf603513e9f3fb9e3141c283602018c5436b9a4c205fdd0336
ToricPair.genuineToricObstruction	
source	1a8c26228b2b42cce64a6d3d01fcada6640b2ac3eb013a4bc08bbf79b70e2fb8
artifact	c1b3288060bc5a3f0d3eed696b5a09f58ecf4fd738d9f9ad78c5030b10838337
interface	468fbda95f70e6c224c5c88bec56252136ce12e7a6717cd1204969bc960bde18
Main.lean (aggregate, publication-facing aliases)	
source	3a07050e80be02e8d63b3c5884c762a4b2740b630d9a9a8f8544cc545b0f1397
interface	624fda48df9581c603b0bce7df545fb85a9646314e1e68078c2fccddb5305a7d

Table 2: SHA-256 receipts of the strict verification runs.

$\exp[-(G_R^\ddagger - G_L^\ddagger)/RT]$ and Theorem 4.1 reads

$$\mathcal{F} \text{ is MultiCAC} \iff |G_R^\ddagger - G_L^\ddagger| < RT \ln 2,$$

about 1.7 kJ mol^{-1} at 298 K. Two autocatalytic triangles competing for the same feeder reaction can coexist, in the static sense of this paper, only if their private activation energies are matched to within this window. Outside it, the faster branch drains the shared complex-activity drop below what the slower branch needs, regardless of the standard free energies of the species. We stress that this translation is asserted only for the symmetric family; with unequal transmission coefficients or unequal factors within a branch the exact statement is Proposition 4.4 in terms of the reciprocal-factor sums, and Remark 5.4 gives its meaning.

The window is not an artefact of gain two. By Corollary 5.3, two triangles with gains m_L and m_R and equal product complexes would have the asymmetric window $1/m_L < R_R/R_L < m_R$; larger gains tolerate larger mismatches, in the direction one would expect.

8.2 The three layers

The examples of this paper occupy three different positions in Figure 1. The example of Kosc et al. fails at the direction layer, and the failure is a signed circuit, an object that depends only on stoichiometry and orientation. The two-triangle family fails at the magnitude layer: directions are fine, but the fixed factors demand inconsistent complex-activity drops, and this is detectable by linear programming over independent complex activities. The two-species family fails only at the monomial lift: the linear program is feasible, and what fails is that the feasible complex activities must lie on the monomial image $\{(z^c)_c : z > 0\}$ of the species orthant. In the terminology of algebraic geometry this image is the positive part of a toric variety, which is why we call the last failure a toric-image obstruction; we repeat that no toric dynamical system in the sense of [15] is involved.

The layers are ordered by the cost of their certificates. A signed circuit is a nonnegative kernel vector; a linear-complex wall is a Farkas certificate of an infeasible strict linear system; a monomial-lift wall is, in general, a Positivstellensatz-type certificate for the emptiness of a semialgebraic set. In the examples the last certificate is a two-line inequality, but nothing in the theory guarantees that in general.

8.3 Limitations and open problems

The results concern existence of a positive static witness: one activity vector for which all cores run productively at the same time. They say nothing about whether that state is a

steady state of the full dynamics, whether it is stable or persistent [12], whether it is reachable, or whether it survives bounded concentrations or nonideal activities. Kosci et al. note that bounding concentrations can already destroy isolated realizability; the interaction of bounds with the interface calculus is untouched here.

The interface theorem covers one shared scalar and triangular private sectors. Several questions are natural.

1. *Mixed gains at the species level.* For two triangles with different gains sharing $A \rightleftharpoons B$, the responses $q_L = (B - A^{m_L})/(A - B)$ and $q_R = (B - A^{m_R})/(A - B)$ are coupled through $A \in (0, 1)$. Is the MultiCAC region in factor space still described by finitely many rational inequalities, and what replaces Corollary 5.3?
2. *Larger overlap graphs.* When cores overlap along a tree of shared reactions, do the intervals compose by a sequence of pairwise intersections, and when the overlap graph contains a cycle, is there a holonomy-type consistency condition around the cycle?
3. *Higher-dimensional interfaces.* When two cores share more than one reaction, the interface response is a point in a polyhedron rather than an interval. The image of the productive cone under the barrier-weighted response map is then a convex set, and the monomial lift adds semialgebraic constraints. A general theory would describe exactly which such sets arise.
4. *Complexity.* Dir and LinCx are linear feasibility problems, so for a fixed family they are polynomial in the size of the network. MultiCAC asks whether a strict polynomial system has a positive solution, which is decidable but not obviously tractable. Is MultiCAC for a given family of PACs NP-hard, as PAC detection itself is [4, 11]?

8.4 Reproducibility

The formal development, the exact preflight scripts, the verification driver, the replay certificate and the LaTeX source of this paper are archived together; the receipts of Table 2 identify the verified sources bit for bit. Replaying the two publication-facing declarations requires the pinned Lean toolchain and the frozen Mathlib manifest and takes a few minutes on a workstation.

Acknowledgements

Acknowledgements to be inserted before submission.

References

- [1] Alex Blokhuis, David Lacoste, and Philippe Nghe. Universal motifs and the diversity of autocatalytic systems. *Proceedings of the National Academy of Sciences*, 117(41): 25230–25236, 2020. doi: 10.1073/pnas.2013527117.
- [2] Wim Hordijk and Mike Steel. Detecting autocatalytic, self-sustaining sets in chemical reaction systems. *Journal of Theoretical Biology*, 227(4):451–461, 2004. doi: 10.1016/j.jtbi.2003.11.020.
- [3] Mike Steel, Joana C. Xavier, and Daniel H. Huson. The structure of autocatalytic networks, with application to early biochemistry. *Journal of the Royal Society Interface*, 17(171): 20200488, 2020. doi: 10.1098/rsif.2020.0488.

- [4] Thomas Kosc, Denis Kuperberg, Etienne Rajon, and Sylvain Charlat. Thermodynamic consistency of autocatalytic cycles. *Proceedings of the National Academy of Sciences*, 122(18):e2421274122, 2025. doi: 10.1073/pnas.2421274122.
- [5] 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.
- [6] 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.
- [7] Jérémie Unterberger and Philippe Nghe. Stoichiometric and dynamical autocatalysis for diluted chemical reaction networks. *Journal of Mathematical Biology*, 85:26, 2022. doi: 10.1007/s00285-022-01791-x.
- [8] Praneet Nandan, Philippe Nghe, and Jérémie Unterberger. Autocatalytic cores in the diluted regime: classification and properties. arXiv preprint arXiv:2507.15546 [q-bio.MN], 2025.
- [9] Richard Golnik, Thomas Gatter, Peter F. Stadler, and Nicola Vassena. Autocatalytic cores in reaction networks with explicit catalysis. arXiv preprint arXiv:2603.02770 [math.CO], 2026.
- [10] Praful Gagrani, Victor Blanco, Eric Smith, and David Baum. Polyhedral geometry and combinatorics of an autocatalytic ecosystem. *Journal of Mathematical Chemistry*, 62(5): 1012–1078, 2024. doi: 10.1007/s10910-024-01576-x.
- [11] Jakob L. Andersen, Christoph Flamm, Daniel Merkle, and Peter F. Stadler. Maximizing output and recognizing autocatalysis in chemical reaction networks is NP-complete. *Journal of Systems Chemistry*, 3:1, 2012. doi: 10.1186/1759-2208-3-1.
- [12] Nicola Vassena and Peter F. Stadler. Unstable cores are the source of instability in chemical reaction networks. *Proceedings of the Royal Society A*, 480(2285):20230694, 2024. doi: 10.1098/rspa.2023.0694.
- [13] Richard Golnik, Thomas Gatter, Wim Hordijk, Peter F. Stadler, and Nicola Vassena. Bridging two theoretical frameworks of autocatalysis: RAF sets and stoichiometric autocatalysis. arXiv preprint arXiv:2605.25523, 2026.
- [14] Armand Despons, Yannick De Decker, and David Lacoste. Structural constraints limit the regime of optimal flux in autocatalytic reaction networks. *Communications Physics*, 7: 224, 2024. doi: 10.1038/s42005-024-01704-8.
- [15] Gheorghe Craciun, Alicia Dickenstein, Anne Shiu, and Bernd Sturmfels. Toric dynamical systems. *Journal of Symbolic Computation*, 44(11):1551–1565, 2009. doi: 10.1016/j.jsc.2008.08.006.
- [16] Martin Feinberg. *Foundations of Chemical Reaction Network Theory*, volume 202 of *Applied Mathematical Sciences*. Springer, 2019. doi: 10.1007/978-3-030-03858-8.
- [17] Alexander Schrijver. *Theory of Linear and Integer Programming*. John Wiley & Sons, Chichester, 1986.
- [18] Henry Eyring. The activated complex in chemical reactions. *The Journal of Chemical Physics*, 3(2):107–115, 1935. doi: 10.1063/1.1749604.