

Certified localization of reaction deletions in autocatalytic networks: ranked support witnesses, witness selection, and exact robustness laws

Anonymous manuscript

September 7, 2026

Abstract

A reflexively autocatalytic and food-generated (RAF) set is the standard combinatorial model of a self-sustaining reaction network. We study what happens to the maximum RAF of a finite catalytic reaction system when reactions are deleted, from three directions. First, we introduce *ranked support witnesses*: for every reaction of the baseline maximum RAF, a set of producer parents whose ranks decrease along reactant edges, together with one unrestricted catalytic parent. We prove that the part of the baseline not reachable from a deletion along selected edges survives every deletion, and that the exact new maximum RAF is recovered by one residual computation inside the reachable cone. This localizes any deletion query to a source-checkable region and admits an exact charge-model theorem: on an explicit unbounded family, a packet evaluator beats a specified fresh evaluator by any prescribed constant factor. Second, we ask which witness to store. The reachable cone always contains the true loss, and the excess above that intrinsic floor is the only thing witness selection can change. Every single deletion admits a perfect witness, the true loss equals the intersection of all witness cones, singleton dependence is a preorder, and synergistic loss under combined deletions prevents any one witness from being perfect for every singleton. Weighted selection is exactly separable on a one-layer source class, becomes NP-complete after one downstream aggregator through an exact reduction from set cover, is repaired by portfolios that keep exact answers through a residual solve and admit the classical greedy coverage guarantee, and yet can require 2^k witnesses for universally exact regions on a source with $2k + 1$ reactions although two suffice for all singletons. Third, for independent random deletion we derive exact robustness laws: the first derivative of expected surviving size is the total singleton loss; the second derivative is twice the sum, over pairs, of overlapping singleton damage minus cooperative pair damage; on functional sources with one alternative catalytic site, two forward exposure sets decide survival under every availability set, yield the complete classification of inclusion-minimal external cuts, an exact objective for choosing one catalyst addition, and an exact variance formula whose overlap count is the obstruction to concentration. Disjoint catalytic modules and a common gateway separate size bias from collective criticality and concentration from persistent fluctuations. The finite statements are compiled in Lean 4 against Mathlib with warnings promoted to errors and no admitted proofs; the few hand-proved consequences are listed explicitly. The results concern structural RAF survival, not kinetic persistence.

Keywords. autocatalytic sets; RAF theory; reaction deletion; incremental computation; certificates; set cover; Russo–Margulis formula; robustness; formal verification.

MSC 2020. 92E20, 68Q17, 68Q25, 60C05, 90C27, 68V20.

1 Introduction

A catalytic reaction system consists of molecules, reactions with reactant and product sets, a food set of externally supplied molecules, and a relation recording which molecules catalyse which reactions. Kauffman [1] proposed that self-sustaining networks of mutually catalytic reactions could emerge in such systems, and Hordijk and Steel [2] gave the now standard formalization: a *reflexively autocatalytic and food-generated* (RAF) set is a nonempty set of reactions each of which has all reactants and at least one catalyst in the closure of the food under the set. RAF sets are closed under union, so every availability set of reactions contains a unique maximum RAF, and it is computable in polynomial time by alternating food closure with pruning [2, 3]. RAF theory has been applied to metabolic networks [4], to the emergence of autocatalysis in random chemistries [5], and, in its structural and algorithmic aspects, to a growing body of results on generation orderings, irreducible RAFs and residual constructions [6, 7].

This paper is about *reaction deletion*. Given a baseline maximum RAF S , deleting a set K of reactions leaves the maximum RAF of $S \setminus K$, and the reactions lost are $L(K) = S \setminus M(S \setminus K)$. Deletion analysis is the basic sensitivity tool of the field: Sousa et al. [4] measured how the *E. coli* maximum RAF shrinks when individual catalysts or reactions are removed, and essential reactions are a recurring theme in [6, 7]. Three questions organize the present work.

- (i) *Localization*. A deletion query can be answered by recomputing the maximum RAF from scratch. Can the recomputation be confined, with an exact answer, to a region determined by a checkable witness derived from the source?
- (ii) *Selection*. Such witnesses are far from unique. Which one should be stored, how hard is the choice, and what does a collection of witnesses buy?
- (iii) *Structure*. Independently of any algorithm, which features of the source determine the damage caused by deletion, in expectation and in fluctuation?

The three questions share one object, the *ranked support witness*, and one semantic contract, the ordinary RAF semantics in which a catalyst need not be present before the reaction it catalyses fires. We answer them in turn with theorems that are, with the exceptions listed in Section 6.3, compiled in Lean 4 [8] against Mathlib [9].

1.1 Main results

Throughout, Q is a finite catalytic reaction system, $M(A)$ is the maximum RAF contained in A (or $\emptyset$), and $S = M(A_0)$ is a fixed baseline with $n = |S|$.

Certified localization (Section 3). A ranked support witness $W = (p, \lambda)$ on S assigns to each reaction $r \in S$ a parent set $p(r)$ and a rank $\lambda(r) \in \mathbb{N}$ such that every nonfood reactant of r is produced by a parent of strictly smaller rank and at least one catalyst of r is food or is produced by a parent of any rank (Definition 3.1). Witnesses exist for every RAF (Lemma 3.2). Writing $C_W(K)$ for the set of baseline reactions reachable from K along selected parent edges, we prove:

- **Retention and residual identity (Theorems 3.3 and 3.5).** Any subset of S closed under selected parents is a RAF or empty and survives every deletion avoiding it; consequently, for any region $E \supseteq D$ whose complement $O = S \setminus E$ is parent-closed,

$$M(A_0 \setminus D) = O \cup M_{Q_O}((S \setminus D) \cap E),$$

where Q_O is Q with the products of O adjoined to the food. In particular $L(D) \subseteq C_W(D)$ for every valid W .

- **Closed-seed shortcut (Proposition 3.6).** If no selected edge leaves K , then $L(K) = K$ with no residual computation.
- **Charge-model amortization (Theorem 3.7).** On an explicit connected family with $2n + 1$ reactions whose stoichiometric columns lie on distinct positive rays, a sequential packet evaluator returns the literal maximum-RAF membership trace and, for any prescribed factor t with $1150t \leq 2n + 1$ and any valid edit sequence longer than an explicit setup cap, its total charge is less than $1/t$ of the charge of a specified fresh evaluator.

Witness selection (Section 4). For nonnegative singleton weights w , the cost of a witness is $J_w(W) = \sum_{d \in S} w_d |C_W(\{d\})|$.

- **Intrinsic floor and perfect queries (Proposition 4.1, Theorem 4.2).** Expected cone size decomposes as expected actual loss plus expected excess, and only the excess depends on W . Every deletion set K admits a witness W_K with $C_{W_K}(K) = L(K)$; hence $L(K) = \bigcap_W C_W(K)$ over all valid witnesses.
- **Intrinsic dependence and synergy (Theorems 4.3 and 4.4).** The relation $d \preceq r \iff r \in L(\{d\})$ is a preorder on S . If some deletion set has *synergistic* loss $L(K) \setminus \bigcup_{d \in K} L(\{d\}) \neq \emptyset$, no single witness is perfect for every singleton.
- **An exactly solvable class and a hardness wall (Theorems 4.6, 4.7, 4.8).** On a one-layer producer–consumer source class, weighted selection is separable and a minimum-weight rule is globally optimal over all valid witnesses. Adding one downstream aggregator gives an exact equivalence between witness cost thresholds and set cover; the uniform-mean threshold problem for ranked witnesses is NP-complete.
- **Portfolios and their limits (Theorems 4.9, 4.10, 4.11).** A finite portfolio of witnesses answers every deletion exactly through one residual solve on the intersection of its cones; greedy selection from a finite pool attains the $1 - (1 - 1/b)^b$ coverage guarantee on a training workload. Nevertheless there is a food-catalysed source with $2k + 1$ reactions on which exact intersection regions for every deletion set require exactly 2^k witnesses, although two suffice for all singleton deletions.

Exact robustness laws (Section 5). Retain each reaction of S independently with probability p and let $\Phi(p) = \mathbb{E}|M(A_p)|$ and $F(p) = \mathbb{P}[M(A_p) \neq \emptyset]$.

- **First and second order (Proposition 5.1, Theorem 5.2).** $\Phi'(1) = \sum_i |L(\{i\})|$, $F'(1)$ counts the singletons whose deletion extinguishes the RAF, and

$$\Phi''(1) = 2 \sum_{\{i,j\} \subseteq S} \left(|L(\{i\}) \cap L(\{j\})| - |C_{ij}| \right), \quad C_{ij} = L(\{i,j\}) \setminus (L(\{i\}) \cup L(\{j\})).$$

Curvature therefore separates two mechanisms that the mean singleton loss cannot distinguish: overlapping damage, which makes repeated deletions redundant, and cooperative failure, which makes a pair worse than the sum of its parts.

- **Exposure laws on functional sources (Lemma 5.4, Theorems 5.5–5.7).** On sources where each reaction consumes food, produces one private molecule and is catalysed by the product of one parent reaction, with one distinguished reaction allowed a second catalyst, the survival of a target under any availability set is decided by two forward orbits E_r, G_r : r survives iff $E_r \subseteq A$ or $G_r \subseteq A$. This yields the complete list of inclusion-minimal external cuts, and the exact survival gain $\Delta_r(p) = p^{|G_r|}(1 - p^{|E_r \setminus G_r|})$ of adding one catalyst, which is positive precisely when the new orbit is not contained in the old one.

- **Fluctuations (Theorem 5.8, Corollary 5.9, Propositions 5.10 and 5.11).** $\text{Var}|M_f(A_p)| = \sum_{r,s} (p^{|E_r \cup E_s|} - p^{|E_r|+|E_s|})$; the normalized variance is at most the density of intersecting exposure pairs, so sparse overlap forces concentration. Disjoint catalytic modules give explicit Φ , F and variance and exhibit pure size bias; a common gateway keeps $F(p) = p$ at every size and defeats concentration.

Formal verification (Section 6). The finite statements above are compiled in Lean 4.30.0 against Mathlib at the matching tagged release, in three module families: `RAFQueryCompilation` (localization, residual identity, certificate checking, charge model), `RAFSupportSelection` (selection) and `RAFReactionCriticality` (robustness laws). Section 6 gives the concordance and lists the hand-proved items: the polynomial encoding in the NP argument, the Taylor and Chebyshev consequences, the minimum-cardinality cases of the cut classification, and the closed-form module and gateway variances.

1.2 Relation to previous work

The food-closure and pruning characterization of maximum RAFs, its polynomial-time algorithm, and the pseudo-RAF and producer-counting refinements that already reduce the number of closure computations, are due to Hordijk, Steel and coauthors [2, 3]. Admissible food-generation orderings and the residual-RAF construction obtained by adjoining the products of a retained subset to the food appear in Huson, Xavier and Steel [7]; our ranks are a selected sparse instance of such an ordering and our residual identity (Theorem 3.4) is the version of their construction that we need, proved for an arbitrary fixed set rather than a RAF. We do not claim these foundations, sparse witness selection by itself, or finite fixed-point evaluation as new. What is new is the combination: a certificate whose validity is source-checkable, whose retained part is provably stable under every deletion avoiding it, and whose residual computation is exact, together with the analysis of which certificate to keep.

Incremental maintenance of recursive derivations has a substantial literature in databases and knowledge representation, including delete/rederive [10] and the backward/forward algorithm [11]. The RAF fixed point is a least closure nested inside a greatest pruning, and a retained region is certified here by a static witness rather than by rederivation; the relationship to those algorithms is discussed in Section 7. The distinction between representation succinctness and query tractability that organizes Section 4 is the one drawn by Darwiche and Marquis [12].

The set-cover reduction of Theorem 4.7 uses the classical NP-completeness of set cover [13]; the greedy portfolio guarantee is the cardinality-constrained submodular maximization bound of Nemhauser, Wolsey and Fisher [14]. We note explicitly that the additive term in our reduction prevents the transfer of the $\ln n$ approximation-hardness threshold of Feige [15]. The first- and second-order formulas of Section 5 are instances of the Russo–Margulis derivative formula [16–18], applied to the maximum-RAF size and nonemptiness observables. Elementary RAF graph structure on sources whose reactants are all food, and its cycle characterization, is the setting of Steel and Hordijk [19]; our functional sources are a special case, and the two-exposure law is the one-site extension.

1.3 Organization

Section 2 fixes the semantics and the loss operator. Section 3 develops ranked support witnesses, localization, certificate checking and the charge-model theorem, with implementation measurements. Section 4 treats witness selection: floor, perfect queries, dependence, synergy, exact and hard classes, portfolios and the exponential bound. Section 5 proves the robustness

laws. Section 6 documents the formal verification, and Section 7 discusses scope, limitations and open problems.

2 Semantics and the loss operator

Definition 2.1 (catalytic reaction system). A catalytic reaction system is $Q = (X, R, F, \rho, \pi, C)$ with finite molecule set X , finite reaction set R , food set $F \subseteq X$, reactant and product maps $\rho, \pi : R \rightarrow 2^X$, and a catalyst map $C : R \rightarrow 2^X$. For $A \subseteq R$, the *food closure* $\text{cl}_A(F)$ is the least $Y \subseteq X$ with $F \subseteq Y$ such that $\pi(r) \subseteq Y$ whenever $r \in A$ and $\rho(r) \subseteq Y$. Catalysis plays no role in this closure. A reaction r is *supported* by A if $\rho(r) \subseteq \text{cl}_A(F)$ and $C(r) \cap \text{cl}_A(F) \neq \emptyset$. A nonempty A is a *RAF* if every $r \in A$ is supported by A .

Reaction identities are retained throughout; two reactions with the same reactants, products and catalysts are distinct elements of R . The union of RAFs is a RAF, so each $A \subseteq R$ has a maximum RAF $M(A)$, with $M(A) = \emptyset$ when none exists. Let $P(A) = \{r \in A : r \text{ is supported by } A\}$. Iterating P from A gives a decreasing sequence whose fixed point is $M(A)$; every RAF contained in A survives every step by monotonicity of closure [2]. We call any A with $P(A) = A$ a *fixed set*; a fixed set is a RAF or empty. The map M is contractive ($M(A) \subseteq A$), monotone and idempotent.

This is the ordinary RAF semantics: a catalyst need not be available before the reaction it catalyses fires in the generation order. Imposing that requirement would exclude the catalytic cycles that the theory is meant to capture. None of the results below assert kinetic persistence, positive steady states, thermodynamic feasibility, or inhibition tolerance.

Definition 2.2 (baseline and loss). Fix an availability set A_0 and the baseline $S = M(A_0)$, so that $M(S) = S$ and $n = |S|$. For $K \subseteq R$ the *actual loss* is

$$L(K) = S \setminus M(S \setminus K).$$

All deletions in this paper are independent queries from S with food and catalysis unchanged. The loss includes the deleted reactions of S themselves; the *secondary loss* is $L(K) \setminus K$.

Proposition 2.3 (loss closure). *On subsets of S , the map L is extensive ($K \subseteq L(K)$), monotone and idempotent, with $L(\emptyset) = \emptyset$. If $t \in L(\{r\})$ then $L(\{t\}) \subseteq L(\{r\})$.*

Proof. Contractivity of M gives extensivity, and $L(\emptyset) = S \setminus M(S) = \emptyset$. Monotonicity of M reverses twice under complementation, so L is monotone. The identity $S \setminus L(K) = M(S \setminus K)$ and idempotence of M give $L(L(K)) = S \setminus M(M(S \setminus K)) = L(K)$. The last claim follows from monotonicity applied to $\{t\} \subseteq L(\{r\})$ and idempotence. $\square$

3 Ranked support witnesses and certified localization

3.1 Witnesses

Definition 3.1 (ranked support witness). A *ranked support witness* on a set $S \subseteq R$ is a pair $W = (p, \lambda)$ with $p : R \rightarrow 2^R$ and $\lambda : R \rightarrow \mathbb{N}$ such that for every $r \in S$:

(W1) for every reactant $x \in \rho(r) \setminus F$ there is $t \in p(r) \cap S$ with $x \in \pi(t)$ and $\lambda(t) < \lambda(r)$;

(W2) there is $x \in C(r)$ with $x \in F$, or there are $t \in p(r) \cap S$ and $x \in C(r) \cap \pi(t)$.

The *selected graph* of W on S has an edge $t \rightarrow r$ whenever $r \in S$ and $t \in p(r) \cap S$. For $K \subseteq R$ the *cone* $C_W(K)$ is the set of reactions of S reachable from $K \cap S$ by a path of length ≥ 0 in the selected graph, and $\text{Anc}_W(r) = \{d \in S : r \in C_W(\{d\})\}$.

Condition (W1) forbids reactant cycles among selected parents; condition (W2) deliberately places no rank constraint on the catalytic parent. Extra parents are permitted, and removing a parent that is not needed for (W1)–(W2) can only shrink cones. Parents outside S are irrelevant and do not appear in the selected graph. Validity is decidable by scanning the finitely many incidences; this is the Lean checker `checkRankedSupport`.

Lemma 3.2 (existence). *Every fixed set S admits a ranked support witness whose selected parent sets are contained in S and have total cardinality at most the number of nonfood reactant incidences of S plus $|S|$.*

Proof. Since every reaction of S has its reactants in $\text{cl}_S(F)$, the closure can be built by a food-generation ordering $r_1, \dots, r_{|S|}$ of S in which every nonfood reactant of r_k is produced by some r_j with $j < k$. Set $\lambda(r_k) = k$, and for each nonfood reactant select one such producer. Each $r \in S$ has a catalyst in $\text{cl}_S(F)$, which is food or produced by some reaction of S ; select that producer if needed. The count follows. $\square$

Sparsity of a witness says nothing about the size of its cones: one selected producer can support many consumers, and catalytic cycles are allowed.

3.2 Retention, residual evaluation and localization

Theorem 3.3 (retention). *Let W be a ranked support witness on S and let $O \subseteq S$ be closed under selected parents: $p(r) \cap S \subseteq O$ for every $r \in O$. Then $P(O) = O$. Consequently O is a RAF or empty, and $O \subseteq M(A)$ for every $A \supseteq O$.*

Proof. We show by strong induction on $\lambda(r)$ that every $r \in O$ has $\rho(r) \subseteq \text{cl}_O(F)$. A reactant x of r is food or, by (W1), is produced by some $t \in p(r) \cap S$ of smaller rank; closure under parents puts $t \in O$, the induction hypothesis puts $\rho(t) \subseteq \text{cl}_O(F)$, so $\pi(t) \subseteq \text{cl}_O(F)$. Once this holds for every reaction of O , the catalytic witness of (W2) is food or a product of a parent in O , whose reactants are already generated; so $C(r) \cap \text{cl}_O(F) \neq \emptyset$ regardless of the rank of that parent. Hence every reaction of O is supported by O . The final claims follow from the fixed-point characterization of M . $\square$

Theorem 3.4 (exact residual evaluation). *Let $O \subseteq A$ be a fixed set and $I \subseteq A$ with $M(A) \subseteq O \cup I$. Let Q_O denote Q with food $F \cup \bigcup_{t \in O} \pi(t)$. Then*

$$M(A) = O \cup M_{Q_O}(I).$$

Proof. For any fixed set O and any $T \subseteq R$, food generation gives $\text{cl}_T^{Q_O}(F_O) = \text{cl}_{O \cup T}^Q(F)$: generating O first from food and then executing T realizes the right side inside the left, and replacing every use of a product of O by food in a generation ordering of $O \cup T$ gives the converse. Hence a reaction is supported by T in Q_O exactly when it is supported by $O \cup T$ in Q .

Let $L' = M_{Q_O}(I)$. Every reaction of L' is supported by L' in Q_O , hence by $O \cup L'$ in Q , and every reaction of O is supported by $O \subseteq O \cup L'$; so $O \cup L'$ is a fixed set contained in A and $O \cup L' \subseteq M(A)$. Conversely, $M(A) \setminus O \subseteq I$, and each of its reactions is supported by $M(A) = O \cup (M(A) \setminus O)$ in Q , hence by $M(A) \setminus O$ in Q_O ; so $M(A) \setminus O$ is a fixed set of Q_O inside I and is contained in L' . $\square$

Theorem 3.5 (certified localization). *Let $S = M(A_0)$ carry a ranked support witness W , let $D \subseteq R$ be a deletion set and let $E \supseteq D \cap S$ be a region such that $O = S \setminus E$ is closed under selected parents. Then*

$$M(A_0 \setminus D) = O \cup M_{Q_O}((S \setminus D) \cap E).$$

In particular $L(D) \subseteq C_W(D)$ for every valid W .




Figure 1: Certified localization of one deletion query. The witness is computed once from the source and checked; a query traverses the selected graph from the deleted reactions, retains the complement of the cone without any recomputation, and resolves the cone by one residual maximum-RAF computation with the retained products adjoined to the food. If no selected edge leaves D , the loss is D itself.

Proof. By Theorem 3.3, O is a fixed set, and $O \subseteq S \setminus D \subseteq A_0 \setminus D$ because $D \cap S \subseteq E$. Monotonicity gives $M(A_0 \setminus D) \subseteq M(A_0) \cap (A_0 \setminus D) \subseteq S \setminus D \subseteq O \cup ((S \setminus D) \cap E)$. Apply Theorem 3.4 with $A = A_0 \setminus D$ and $I = (S \setminus D) \cap E$. For the last claim take $E = C_W(D)$: if $r \notin C_W(D)$ then no selected parent of r lies in $C_W(D)$, so $S \setminus C_W(D)$ is parent-closed and survives. $\square$

The parent-closure condition on E is source-checkable: reverse the selected edges and require E to contain $D \cap S$ and to be successor-closed. Nothing is assumed about producers outside the selected ones; alternative producers remain available to the residual evaluator. Figure 1 summarizes the mechanism.

Proposition 3.6 (closed-seed shortcut). *Let W be valid on S and $K \subseteq S$. If $t \in K$ whenever $t \in S$, $r \in K$ and $r \in p(t)$ (no selected edge leaves K), then $L(K) = K$.*

Proof. The hypothesis gives $C_W(K) \subseteq K$, so $L(K) \subseteq K$ by Theorem 3.5; extensivity gives the converse. $\square$

3.3 Checking the residual computation

The implementation stores the baseline as a mask together with, for each molecule, the number of baseline reactions producing it. For a molecule needed as a reactant or catalyst inside the region, subtracting the producers that lie in E leaves a count; a positive remainder certifies supply by O . These molecules, with food, form the residual boundary; products of O that no local reaction needs may be omitted, since induction on local closure preserves every reactant and catalyst test.

The residual maximum RAF is itself certified. A *closure certificate* is a finite reaction schedule: replay starts from food and adds the products of a scheduled reaction only when it is available and its reactants are present, so the replayed pool is always contained in the true closure; the checker then verifies that the pool is closed under every available reaction, and leastness gives the reverse inclusion. A *pruning certificate* supplies such schedules for the successive decreasing reaction sets and is accepted only on a round that is a fixed point. Induction over accepted rounds shows that the returned set is the residual maximum; on rejection the implementation falls back to a literal fresh evaluation, so a malformed proposal

cannot change an answer. Theorem 3.5 then proves exactness of every accepted localized answer. The sequential engine additionally supports insertions under stronger source-derived dependency conditions and commits no state before certificate acceptance.

3.4 A charge-model amortization theorem

We now state the one result of this paper that concerns cost rather than correctness. It is a theorem in a declared charge model, not a native runtime claim.

Consider the family with n pairs and $m = 2n + 1$ reactions: food $\{f\}$, a hub reaction $f \rightarrow h$ catalysed by f , and for each i the reactions $h \rightarrow a_i$ catalysed by b_i and $h \rightarrow b_i$ catalysed by a_i , all species distinct. Every reaction is incident with h , so the reaction–molecule incidence graph is connected, and every stoichiometric column has a unique positive product coordinate, so no two columns are positive scalar multiples of one another. Each pair is a catalytic two-cycle: neither reaction has its catalyst before the other fires, which is exactly the situation that ordinary RAF semantics admits and that condition (W2) accommodates.

An *edit packet* is five words: a pair index, two removal bits and two addition bits, removal applied before addition. The packet generates its two-reaction region, a two-round pruning certificate whose schedules are $[\text{left}, \text{right}]$ and $[\]$, and the right reaction as a membership target. The *charge semantics* prices identifier comparisons, cached reads, Boolean branches, bounded-word arithmetic, finite-set comparisons and copies, certificate replay, state changes and output construction; all tariffs are executable definitions in the Lean development. The *fresh point baseline* scans its m -entry availability representation on every query, so its charge is at least mq on q queries.

Theorem 3.7 (charge-model amortization). *Let $m = 2n + 1$ and let $P(n)$ be the explicit setup cap below. For every positive integer t with $1150t \leq m$, every initial availability A and every valid pair-edit sequence of length $q > P(n)$, the arithmetic packet evaluator returns the literal sequential maximum-RAF membership trace and*

$$t \cdot \text{charge}(\text{packet evaluator}) < \text{charge}(\text{fresh point baseline}).$$

The conclusion also records hub incidence and the distinct-positive-ray property of the family. With $p = m + 1$,

$$\begin{aligned} T(e, i, o, p) &= 1 + e(i + 1)(p + 1) + 2(p + eo + 1)^2, \\ H(e, i, c, p) &= 1 + e(i + c + 2)(p + 1) + 2(e + 1)^2, \\ R(e, i, o, c, p) &= p + pT(e, i, o, p) + H(e, i, c, p), \\ P(n) &= m(14m + 13 + 2p) + (m + mR(m, 1, 1, 1, p)) \\ &\quad + (2m + p + m + 2m(m + 1) + p(m + 1)^2). \end{aligned}$$

Proof sketch. Source compilation constructs dependency and needs vectors once, computes the initial maximum, and builds masks and producer counts; its charge is bounded by $P(n)$ by direct substitution of the executable tariffs. A pair region contains two reactions with needed-row total 4 and product-row total 2; substitution into the charged local bounds gives region checking at most 19, boundary charge 155 and total local certificate charge at most 818, hence a local cap of 831 with metadata, an accepted-transaction budget of 1011 after edit preparation and state update, a cached per-query charge of 1014 after dispatch, target lookup and output construction, and at most 135 more for five-word decoding, arithmetic reaction codes and generated cells. The total is at most $P(n) + 1149q < 1150q$ when $q > P(n)$. Multiplying by t and using $1150t \leq m \leq (\text{baseline charge per query})$ proves the inequality. Packet refinement proves that both evaluators return the same literal trace. The formal statement is `RAFQueryCompilation.ModuleFamily.family_terminal_factor`. $\square$

Profile	Queries	Witness (s)	Reference (s)	Ratio
Two-step, 512 and 1024 pairs; producer plus checker	3074	0.99 / 1.79	—	5.10 / 7.61
same runs; complete cold session	3074	25.96	18.67	0.72
Two-step, 2048 pairs (4097 reactions); pipeline	4097	5.11	58.18	11.38
Reexport, 4096 peripheral (4097 reactions); pipeline	4097	4.21	61.16	14.54
both sources; end to end with one cold launch	8194	35.37	119.35	3.37

Table 1: Single-run wall-clock observations on one Windows machine. All 3074 + 8194 loss lists matched both references exactly. The second row shows why warm-only measurements are insufficient: a cold checker launch dominates a moderate workload. Timings are single observations, not confidence intervals.

Factors are unbounded: take $n = 575t$ and any sufficiently long valid edit list. The cap $P(n)$ is conservative, so the sufficient q can be enormous; the theorem establishes a structural amortization factor, not a numerical prediction of native break-even. It is a comparison with a specified fresh algorithm, not an information-theoretic lower bound on all RAF algorithms or on a family-specialized answer table.

3.5 Implementation measurements

The measured implementation is the ranked-witness deletion-loss route of Theorem 3.5 with the certificate bridge of Section 3.3, not the packet interpreter of Theorem 3.7. Two synthetic balanced source generators were used. In the *two-step* family a hub $f \rightarrow x$ feeds reactions $U_i : f + x \rightarrow b_i + c_i$ and $V_i : f + b_i \rightarrow x + a_i$, with V_i food-catalysed and U_i catalysed by the preceding a in a fixed group; in the *reexport* family reactions $f + x \rightarrow x + a_i$ use preceding group products as catalysts. Both have distinct nonproportional columns and connected incidence graphs; neither is biological data. For each source the complete singleton-deletion profile, one loss list per baseline reaction, was computed and compared entry by entry with two independently implemented Python references, active-producer closure/pruning and pseudo-RAF pruning after [3]; these are not the authors’ native implementation.

Table 1 reports the observations. Every deployed producer permits cone regions up to 64 reactions before falling back to global closure; the threshold is a heuristic, and every proposed certificate is still checked. The gains are a scoped crossing on two synthetic families with one baseline pair; independent reproduction, heterogeneous perturbations and comparison with pinned native author implementations are not claimed.

4 Witness selection

Theorem 3.5 holds for every valid witness, and different legal producer choices can produce very different cones. This section asks which witness to store for a workload of independent deletion queries against a fixed baseline. Neither sparsity nor chemical redundancy alone identifies a good choice.

4.1 The objective and the intrinsic floor

For nonnegative singleton weights $w : S \rightarrow \mathbb{R}_{\geq 0}$, transposing the finite reachability relation gives

$$J_w(W) = \sum_{d \in S} w_d |C_W(\{d\})| = \sum_{r \in S} \sum_{d \in \text{Anc}_W(r)} w_d. \quad (1)$$

Weights summing to one make this an expectation; all weights one give the total singleton cost, whose division by n is the uniform mean. The objective is transitive reachability, not immediate out-degree or parent count. We also consider the worst singleton cone $\max_d |C_W(\{d\})|$ and, in Section 4.5, complete implementation cost; these objectives can prefer different witnesses, and reducing the mean can increase the maximum.

Proposition 4.1 (intrinsic floor). *For any finite deletion workload with probabilities $\mu(K)$ and any valid W ,*

$$\mathbb{E}_\mu |C_W(K)| = \mathbb{E}_\mu |L(K)| + \mathbb{E}_\mu |C_W(K) \setminus L(K)|.$$

Proof. Theorem 3.5 gives $L(K) \subseteq C_W(K)$; partition and take weighted sums. The Lean statement `expected_cone_decomposition` holds for any finite real weighting. $\square$

Selecting W therefore changes only the excess above a structural floor that no witness can lower.

4.2 Perfect queries, intrinsic dependence and synergy

Theorem 4.2 (perfect individual queries). *For every $K \subseteq S$ there is a valid witness W_K with $C_{W_K}(K) = L(K)$. Equivalently,*

$$L(K) = \bigcap_{W \text{ valid on } S} C_W(K).$$

Proof. Put $T = M(S \setminus K)$, a fixed set with $T \cap K = \emptyset$. By Lemma 3.2, T has a witness with parents inside T ; keep its ranks and parents on T . Choose a food-generation ordering of S and assign each $r \in S \setminus T$ a rank above the maximum rank on T , increasing along the ordering. For such r and a nonfood reactant x , the ordering supplies an earlier producer $t \in S$ of x : if $t \in T$ its rank is smaller because all ranks on T are; if $t \in S \setminus T$ it is earlier in the ordering. Select it. Select any catalytic witness of r in S . No selected edge enters T from $S \setminus T$ because parents of reactions of T were chosen in T . The cone of K starts in $S \setminus T$ and cannot enter T , so $C_{W_K}(K) \subseteq S \setminus T = L(K)$; Theorem 3.5 gives equality. The intersection formula follows since every valid cone contains $L(K)$. $\square$

The construction is an existence argument: it computes T first, so it does not accelerate that first query. It does say that the true loss is exactly what all witnesses agree on.

Theorem 4.3 (intrinsic singleton dependence). *On S define $d \preceq r$ when $r \in L(\{d\})$. This relation is reflexive and transitive; mutual dependence is an equivalence relation.*

Proof. Reflexivity is extensivity of L . For transitivity, by Theorem 4.2, $d \preceq r$ means that $r \in C_W(\{d\})$ for every valid W . If $a \preceq b$ and $b \preceq c$ then in each fixed valid selected graph there is a path from a to b and one from b to c , hence one from a to c ; applying the intersection formula again yields $a \preceq c$. (Alternatively, use the last clause of Proposition 2.3.) $\square$

The equivalence classes consist of reactions mutually indispensable for one another's survival within the fixed baseline. A class need not itself be a RAF. Dependence is intrinsic even though any one selected graph may contain additional paths.

Theorem 4.4 (synergy forces singleton overestimation). *If a valid witness W is perfect for every singleton, $C_W(\{d\}) = L(\{d\})$ for all $d \in S$, then it is perfect for every deletion set. Consequently, defining the synergy*

$$\Delta(K) = L(K) \setminus \bigcup_{d \in K} L(\{d\}),$$

if $\Delta(K) \neq \emptyset$ for some $K \subseteq S$ then no single valid witness is perfect for every singleton.

Proof. A path from a seed set starts at one of its elements, so $C_W(K) = \bigcup_{d \in K} C_W(\{d\})$. Under singleton perfection this equals $\bigcup_{d \in K} L(\{d\}) \subseteq L(K)$ by monotonicity, while $L(K) \subseteq C_W(K)$ by Theorem 3.5. Hence equality holds throughout and $\Delta(K) = \emptyset$. The second claim is the contrapositive. No converse is claimed. $\square$

Example 4.5 (two interchangeable producers). Let two food-consuming, food-catalysed producers P_0, P_1 both produce x , and let m food-catalysed consumers $e_1, \dots, e_m$ consume x . Every actual singleton loss has size one. A normalized witness assigns k consumers to P_0 and $m - k$ to P_1 ; the producer cones have sizes $1 + k$ and $1 + m - k$ and each consumer cone has size one. Since normalization removes parents without enlarging cones,

$$\min_W \max_{d \in S} |C_W(\{d\})| = 1 + \lceil m/2 \rceil.$$

The best uniform mean is $(2m + 2)/(m + 2) \rightarrow 2$ against an intrinsic floor of 1, whereas the producer-weighted mean is $1 + m/2$. Deleting both producers loses every consumer although each producer deletion alone loses nothing else, so $\Delta(\{P_0, P_1\}) \neq \emptyset$ and Theorem 4.4 applies. The example proves an unbounded worst-case or producer-weighted gap, not an unbounded uniform-mean gap.

4.3 An exactly solvable class and the downstream-convergence wall

Theorem 4.6 (one-layer weighted optimality). *Let P be a finite producer set and U a finite consumer set, with a nonempty allowed set $A_i \subseteq P$ for each $i \in U$. Construct the reactions*

$$p_j : f \rightarrow u_j + \{x_i : j \in A_i\}, \quad e_i : x_i \rightarrow y_i,$$

where f is food, all other displayed molecules are distinct, and f catalyses every reaction. The full reaction set is a RAF. For nonnegative singleton weights w , choosing for each e_i a producer of minimum weight w_{p_j} among $j \in A_i$ yields a witness that is globally optimal for J_w over all valid witnesses, of any parents and ranks.

Proof. Any valid witness includes among the parents of e_i an allowed producer of x_i . Retain one such parent per consumer and delete all other parents: this is valid with producers at rank 0 and consumers at rank 1, and deleting parents cannot increase reachability. For a normalized choice $a(i) \in A_i$ the ancestors of a producer are itself and the ancestors of e_i are $\{e_i, p_{a(i)}\}$, so by (1)

$$J_w(a) = \sum_{r \in P \cup U} w_r + \sum_{i \in U} w_{p_{a(i)}}.$$

The second sum separates over i , and a least-weight allowed producer minimizes each term. Normalization gives optimality against arbitrary witnesses, not only normalized ones. $\square$

For uniform weights the one-layer total is $|P| + 2|U|$ regardless of assignment. One additional reaction destroys this independence.

Theorem 4.7 (exact set-cover equivalence). *Augment the source of Theorem 4.6 by one reaction t requiring every y_i , producing a fresh molecule z , and catalysed by f ; let $m = |P|$ and $n = |U|$. There is a valid witness with total uniform singleton cost at most $m + 3n + 1 + q$ if and only if the incidence system $(A_i)_{i \in U}$ has a cover $B \subseteq P$ with $|B| \leq q$ and $A_i \cap B \neq \emptyset$ for every i .*

Proof. Each y_i has the unique producer e_i , so every e_i is a forced parent of t in every witness. Normalize the remaining parents as before; the ranks are 0, 1, 2. The ancestor counts are 1 for each producer, 2 for each e_i , and $1 + n + |\text{im } a|$ for t , so $J_1(a) = m + 3n + 1 + |\text{im } a|$. The image of a is a cover; conversely a cover B supplies a choice $a(i) \in A_i \cap B$ with $|\text{im } a| \leq |B|$.

Normalization never increases cost, which gives both directions. Nonempty A_i make the whole constructed source its own maximum RAF. The aggregator makes shared producer ancestry observable and couples the previously independent choices. $\square$

Theorem 4.8 (NP-completeness of mean selection). *The problem: given a finite catalytic reaction system in explicit incidence encoding, a baseline maximum RAF S , and a rational threshold θ in binary, decide whether some ranked support witness on S has uniform mean singleton cone size at most θ ; is NP-complete. Hardness already holds for the food-catalysed three-level sources of Theorem 4.7.*

Proof. Membership. A witness is encoded by a parent adjacency matrix restricted to S (parents outside S are irrelevant) and one rank per reaction; replacing ranks by their position among the distinct ranks on S preserves all strict comparisons and bounds ranks by $|S| - 1$, so the certificate has $O(|S|^2 + |S| \log |S|)$ bits. Validity is checked by scanning reactant, product, parent and catalyst incidences. Reachability from each baseline reaction is computed in $O(|S|^3)$ time; the total cone size is an integer at most $|S|^2$, and comparing its ratio to $|S|$ with θ is polynomial-bit integer arithmetic. If the baseline is supplied as $M(A_0)$ it is computed first by closure and pruning in polynomial time.

Hardness. Reduce from set cover [13] with m sets, n elements and ℓ incidences. Theorem 4.7 constructs $m + n + 1$ reactions, $m + 2n + 2$ molecules and $O(m + n + \ell)$ incidences, with threshold $\theta = (m + 3n + 1 + q)/(m + n + 1)$ of polynomial bit length. Instances with an uncovered element are recognized in polynomial time and mapped to a fixed no-instance, so all A_i may be assumed nonempty and the full reaction set is the baseline. Theorem 4.7 is the yes/no equivalence. $\square$

The additive term $m + 3n + 1$ matters: a multiplicative approximation to the total singleton cost does not translate into the same multiplicative approximation to $|\text{im } a|$ after subtracting it, so no set-cover approximation threshold [15] is transferred. Nor does the reduction classify the worst-singleton objective. The finite source, normalization, baseline identity and threshold equivalence are compiled in Lean; the encoding argument is conventional.

4.4 Portfolios: exact answers, greedy selection, and an exponential wall

Let H be a finite pool of valid witnesses and $\mathcal{P} \subseteq H$ a portfolio. For a deletion $K \subseteq S$ define

$$O_{\mathcal{P}}(K) = \bigcup_{W \in \mathcal{P}} (S \setminus C_W(K)), \quad D_{\mathcal{P}}(K) = S \setminus O_{\mathcal{P}}(K).$$

For a nonempty portfolio $D_{\mathcal{P}}(K) = \bigcap_{W \in \mathcal{P}} C_W(K)$; for the empty portfolio $D_{\emptyset}(K) = S$ and the method is a fresh solve.

Theorem 4.9 (exact portfolio residual). *For every nonempty portfolio $\mathcal{P}$ and every $K \subseteq S$, with $O = O_{\mathcal{P}}(K)$,*

$$M(S \setminus K) = O \cup M_{Q_O}((S \setminus K) \setminus O).$$

Proof. Each $S \setminus C_W(K)$ is parent-closed for its own witness, hence a fixed set by Theorem 3.3, and it avoids K . A union of fixed sets is a fixed set, since the supporting closures of each member are contained in the closure of the union. Apply Theorem 3.4 with $A = S \setminus K$ and $I = (S \setminus K) \setminus O$. $\square$

The answer is exact even when $D_{\mathcal{P}}(K)$ strictly contains $L(K)$, and the intersection need not be a cone of any one witness; soundness comes from the retained union.

For a finite training set $\mathcal{T}$ of deletion queries, the *rescue events* of W are the pairs (K, r) with $K \in \mathcal{T}$ and $r \in S \setminus C_W(K)$, and $G(\mathcal{P})$ counts the events covered by at least one member of $\mathcal{P}$.

three independent producer choices, one conjunctive consumer




Figure 2: The paired source for $k = 3$. Each producer consumes food and produces its own private molecule together with the shared x_i ; food, catalyst and private-product incidences are suppressed. The consumer c requires all x_i . A witness must choose one parent per pair for c , and each choice is defeated by the deletion of exactly the chosen producers.

Theorem 4.10 (greedy portfolio selection). *For nonempty H and $b \geq 1$, repeatedly add a member of H with maximum new rescue count. After b steps the portfolio $\mathcal{P}_b$ has at most b members, yields the exact answer of Theorem 4.9 for every deletion query, and*

$$G(\mathcal{P}_b) \geq \left[1 - \left(1 - \frac{1}{b}\right)^b\right] \max_{\mathcal{B} \subseteq H, |\mathcal{B}| \leq b} G(\mathcal{B}).$$

Proof. Coverage is monotone with diminishing marginal gains. For a comparison portfolio $\mathcal{B}$ of size at most b , at step j the gap $G(\mathcal{B}) - G(\mathcal{P}_j)$ is at most the sum of the marginal gains of the members of $\mathcal{B}$, hence at most b times the largest available gain; greedy reduces the gap by at least a $1/b$ fraction, and iterating from zero gives the bound [14]. Every selected candidate belongs to H and is valid, so Theorem 4.9 applies regardless of the workload. $\square$

The guarantee concerns rescued training mass relative to the best same-budget portfolio in the supplied pool; it is not a residual-cost, runtime, all-witness, or unseen-query guarantee. Exactness of unseen answers is unconditional.

Theorem 4.11 (sharp portfolio size). *For each integer $k \geq 0$ there is a food-catalysed source with $2k + 1$ reactions whose maximum RAF is the full set and on which the minimum number of witnesses whose cone intersection equals $L(K)$ for every deletion set K is exactly 2^k . For $k \geq 1$, two complementary witnesses achieve exact intersections for every singleton deletion, and one does not.*

Proof. Construct producer pairs a_i, b_i for $i = 1, \dots, k$, each consuming food f and producing x_i together with a private product, and one reaction c consuming all x_i and producing a fresh product; every reaction is catalysed by f (Figure 2). Producers have rank 0 and c rank 1; the full set is a RAF.

Lower bound. A valid witness must contain in $p(c)$ at least one of a_i, b_i for each i , because nothing else produces x_i ; choosing one per pair extracts an assignment $\alpha \in \{0, 1\}^k$ from every witness, whatever its extra parents or ranks. For $\beta \in \{0, 1\}^k$ let K_β delete the producer of each pair *not* chosen by β . Every undeleted producer survives and supplies its pair’s input, so c survives and $L(K_\beta) = K_\beta$. If $\alpha \neq \beta$, some selected parent of c lies in K_β and $c \in C_W(K_\beta)$. A fixed witness therefore excludes c for at most one of the 2^k queries; a portfolio with fewer members misses some β , and its intersection contains c although the loss does not.

Attainability. Store every normalized assignment witness W_α . Its cone of K contains K and contains c exactly when $c \in K$ or a selected producer lies in K . If c is neither deleted nor deprived of some pair entirely, choose an undeleted producer per pair; that witness’s cone

Source mechanism	One witness	Portfolio 3	Portfolio 5	Active fresh
Overlapping inputs	1.013	2.881	3.978	7.690
Catalytic feedback	3.478	7.842	14.147	15.261
Redundant reexport	0.918	3.522	5.724	5.803
Mixed depth	5.388	7.624	11.450	7.686

Table 2: Median complete Python session times in milliseconds over four distinct 48-reaction source instances per row; witness strategies include the closed-seed test. Descriptive medians, not confidence intervals or native timings.

excludes c . Otherwise c is actually lost and lies in every cone. Producers have no parents, so their cone membership is exactly membership in K . Hence the intersection over the full pool is $L(K)$ for every K .

Singletons. For a producer deletion one of the two constant assignments avoids it, so their cone intersection is exact; a singleton deletion of c is exact in both. For $k \geq 1$ any single witness chooses a parent of c , and deleting that parent alone leaves its mate and c alive while the cone contains c . For $k = 0$ one witness suffices, consistent with $2^0 = 1$. $\square$

These are representation-size bounds for exact regions obtained by intersecting cones. By Theorem 4.9 a small portfolio still returns exact answers; the theorem shows that singleton tests can substantially understate the storage required for universal exact-region compilation. The Lean lower bound quantifies over arbitrary valid witness families indexed by any finite type of cardinality less than 2^k ; the consumer has arity k , and a bounded-arity refinement is not claimed.

4.5 Implementation evidence and adoption rule

The implementation treats proposed witnesses as untrusted. A canonical one-layer constructor implements Theorem 4.6; for general sources a bounded proposer enumerates source-allowed witness slots, separates reactant rank inequalities from catalytic choices, rejects reactant cycles and scores actual reachability, with a default assignment limit of 4096 beyond which a validated default is returned labelled feasible rather than optimal.

A held-out study used sixteen distinct 48-reaction sources, four independent seeds in each of four synthetic mechanisms, with every distinct singleton deletion, 768 queries per strategy. All loss outputs agreed with the two independent references. Table 2 reports complete session medians. Portfolios of three and five charge two and four survivor-construction solves respectively, together with memory, cone discovery, boundary work and outputs. A separate five-seed 128-reaction overlap profile gave a median single-witness session of 2.756 ms against 55.157 ms for active fresh evaluation, while a seed-first three-witness portfolio took 10.470 ms even though 127 of 128 queries per source qualified for the closed-seed shortcut: portfolio setup and repeated discovery consumed the region savings.

Applying the mean-optimizing selector to the 32 held-out sources of Section 5 plus a cooperative fixture, all 196 deletion queries per strategy returned identical exact loss sets; the sum of cone sizes fell from 499 to 480 while measured setup rose from 1.80 to 4.09 ms and query time changed from 2.68 to 2.72 ms. On an eight-reaction aggregator the native checker accepted 22 singleton queries under default and mean-optimal witnesses with zero mismatches; total singleton cones fell from 17 to 15 while the worst cone rose from 3 to 5.

For a homogeneous workload of Q queries, an extra setup cost H is repaid only when $H + Q t_{\text{reuse}} < Q t_{\text{fresh}}$, with per-query terms including discovery, residual evaluation and output. This is accounting, not a theorem. The measured decision is to prefer one checked witness with the closed-seed shortcut in the tested Python profiles, and fresh evaluation for a short

cold native workload; larger portfolios are justified only by a declared coverage objective or complete measured savings.

5 Exact robustness laws under random deletion

We now leave algorithms and ask which features of the source determine damage. Let $A_p \subseteq S$ retain each reaction of the baseline independently with probability p , and define

$$\Phi(p) = \mathbb{E}|M(A_p)|, \quad F(p) = \mathbb{P}[M(A_p) \neq \emptyset].$$

Both are polynomials in p , defined for all real p though probabilities refer to $0 \leq p \leq 1$. All probability statements use independent availability in one fixed source; catalytic relations are never resampled, and no finite robustness curve establishes a thermodynamic singularity.

5.1 First and second order at the all-retained endpoint

For a real observable H on subsets of S and a coordinate i , write A_p^{-i} for independent availability on $S \setminus \{i\}$. Grouping the product-measure expectation by one coordinate, differentiating its affine dependence on that coordinate, and then moving all coordinates together gives the Russo–Margulis formula [16–18]

$$\frac{d}{dp} \mathbb{E}H(A_p) = \sum_{i \in S} \mathbb{E}[H(A_p^{-i} \cup \{i\}) - H(A_p^{-i})]. \quad (2)$$

Proposition 5.1 (first-order sensitivity).

$$\Phi'(1) = \sum_{i \in S} |L(\{i\})|, \quad F'(1) = |\{i \in S : M(S \setminus \{i\}) = \emptyset\}|.$$

Proof. At $p = 1$ the summand of (2) is $H(S) - H(S \setminus \{i\})$. For size this is $n - |M(S \setminus \{i\})| = |L(\{i\})|$; for the nonempty indicator, monotonicity makes it the indicator of the pivotal event $M(S \setminus \{i\}) = \emptyset$. No division by $p(1 - p)$ is involved. $\square$

Theorem 5.2 (curvature balances overlap and cooperation). *For distinct $i, j \in S$ put $C_{ij} = L(\{i, j\}) \setminus (L(\{i\}) \cup L(\{j\}))$. Then, summing once over each unordered pair,*

$$\Phi''(1) = 2 \sum_{\{i, j\} \subseteq S} (|L(\{i\}) \cap L(\{j\})| - |C_{ij}|).$$

Proof. Differentiate (2) once more. Each coordinate difference is independent of its own coordinate, so the second sum runs over ordered pairs of distinct coordinates, and at $p = 1$ the pair (i, j) contributes

$$|S| - |M(S \setminus \{i\})| - |M(S \setminus \{j\})| + |M(S \setminus \{i, j\})| = |L(\{i\})| + |L(\{j\})| - |L(\{i, j\})|.$$

Monotonicity gives $L(\{i\}) \cup L(\{j\}) \subseteq L(\{i, j\})$; split $L(\{i, j\})$ into that union and its complement C_{ij} , and apply inclusion–exclusion. The summand is symmetric, and ordered pairs count it twice. $\square$

Corollary 5.3 (small-deletion expansion). *Delete each reaction independently with probability $q = 1 - p$ and let D_q be the deleted set. Then*

$$\mathbb{E}|L(D_q)| = q \sum_i |L(\{i\})| + q^2 \sum_{\{i, j\}} (|C_{ij}| - |L(\{i\}) \cap L(\{j\})|) + O(q^3).$$

Proof. The left side is $n - \Phi(1 - q)$ with $\Phi(1) = n$; expand the polynomial Φ at 1 and use Proposition 5.1 and Theorem 5.2. The remainder is for each fixed finite source and is not uniform in the source size. $\square$

Overlap makes repeated deletions partly redundant; cooperative failure makes a pair's joint effect exceed the singleton prediction. There is no universal sign. A catalytic cycle of length $\ell \geq 2$ has $\Phi(p) = \ell p^\ell$ and $\Phi''(1) = \ell^2(\ell - 1) > 0$. Two food-catalysed producers of a common molecule and one food-catalysed consumer requiring it give $\Phi(p) = 2p + 2p^2 - p^3$ and $\Phi''(1) = -2$: the only cooperative pair is the two producers, each of which spares the consumer alone.

5.2 One-site catalyst choice and forward exposures

Lemma 5.4 (one-site decomposition). *Let C_0, C_1 be two catalyst maps on the same reaction system that agree at every reaction other than v , and let $C_\vee = C_0 \cup C_1$ pointwise. Then for every availability set A ,*

$$M_\vee(A) = M_0(A) \cup M_1(A).$$

Proof. Every RAF for C_0 or C_1 is a RAF for $C_\vee$, giving one inclusion. Let T be a RAF under $C_\vee$. If $v \notin T$ the three maps agree on T . If $v \in T$, choose a catalyst of v in $\text{cl}_T(F)$; it belongs to $C_0(v)$ or $C_1(v)$, and in that branch every other reaction of T keeps its catalytic witness because the maps agree there, while food generation is unchanged. So T is a RAF in one branch, and the union over RAF witnesses gives the other inclusion. The argument assumes nothing about reactants or products. $\square$

We now specialize the source. A *functional source* on a reaction set R has every reaction r consuming common food and producing its own private nonfood molecule x_r , no food catalysis, and r catalysed by $x_{f(r)}$ for a map $f : R \rightarrow R$. At one distinguished reaction v we allow instead either $x_{f(v)}$ or $x_{g(v)}$, where $g = f$ off v . The *forward exposure* of r under f is the finite orbit

$$O_f(r) = \{r, f(r), f^2(r), \dots\}.$$

Theorem 5.5 (all-deletion exposure law). *For every $A \subseteq R$, with $E_r = O_f(r)$ and $G_r = O_g(r)$,*

$$r \in M_\vee(A) \iff E_r \subseteq A \text{ or } G_r \subseteq A, \quad \mathbb{P}[r \in M_\vee(A_p)] = p^{|E_r|} + p^{|G_r|} - p^{|E_r \cup G_r|}.$$

Proof. In a functional source every RAF containing r contains its unique catalyst producer $f(r)$, hence the whole orbit; conversely an available orbit is food-generated and supplies a catalyst for each of its members, so it is a RAF. Thus $r \in M_f(A) \iff O_f(r) \subseteq A$. Apply Lemma 5.4. The probability formula holds because retaining an orbit means retaining all its reactions, and retaining both means retaining their union; it preserves their shared randomness and is not a model of independent branch survival. $\square$

The orbits are extracted from each *resolved* source: an added catalyst can create a new self-sustaining cycle, and a witness computed before the addition can miss the rescue (Figure 3).

5.3 Complete external cuts and one catalyst addition

Fix a target r , write $E = E_r$, $G = G_r$, and exclude direct target deletion. Theorem 5.5 gives

$$r \notin M_\vee(R \setminus D) \iff D \cap E \neq \emptyset \text{ and } D \cap G \neq \emptyset.$$

Theorem 5.6 (all inclusion-minimal external cuts). *The inclusion-minimal sets D with $r \notin D$ that lose r are exactly*

$$\{\{x\} : x \in (E \cap G) \setminus \{r\}\} \cup \{\{x, y\} : x \in E \setminus G, y \in G \setminus E\}.$$




Figure 3: A new witness must be extracted after a catalyst addition. Arrows point from a reaction to its required catalytic parent. Initially $f(0) = f(1) = 0$. Adding the product of 1 as a catalyst for 1 supplies the second resolution $g(1) = 1$; after deleting 0, reaction 1 survives through its new self-loop, which the old witness cannot see. All reactions consume food.




Figure 4: Shared and exclusive exposures have different meanings. A shared external element z alone loses the target; a pair with one element from each exclusive region is inclusion-minimal even when a singleton cut exists. The target lies in the intersection but is excluded from external cuts.

Proof. Every cut meets both exposures. If it contains a shared element, that element alone is a cut and minimality forces the singleton. Otherwise choose one element in each exposure; neither is shared, their pair is a cut contained in the original, and minimality forces the pair. Conversely each displayed singleton meets both exposures and has no nonempty proper subset; each displayed pair meets both while each of its singletons misses one. The target itself lies in both exposures and is excluded. $\square$

The minimum cardinality of an external cut is one if a shared external element exists, two if none exists and both exclusive parts are nonempty, and no external cut exists if either orbit is $\{r\}$. Inclusion-minimal pairs may coexist with singleton cuts (Figure 4).

Theorem 5.7 (when a catalyst addition protects a target). *Compare the functional source f with the one-site union source. The exact target-survival gain is*

$$\Delta_r(p) = \mathbb{P}[r \in M_V(A_p)] - \mathbb{P}[r \in M_f(A_p)] = p^{|G_r|}(1 - p^{|E_r \setminus G_r|}),$$

which for $0 < p < 1$ is positive exactly when $E_r \setminus G_r \neq \emptyset$. The number of singleton deletions fatal to r falls from $|E_r|$ to $|E_r \cap G_r|$, a reduction of $|E_r \setminus G_r|$.

Proof. Subtract $p^{|E_r|}$ from the formula of Theorem 5.5 and use $|E_r \cup G_r| = |G_r| + |E_r \setminus G_r|$. The singleton count is Theorem 5.6 with direct deletion included in both sources. $\square$

Counting added catalysts alone does not identify protection: an added orbit contained in the old one gains nothing. For a supplied list of c candidate single additions (v, w) , resolve each by $g(v) = w$ and $g = f$ elsewhere, sum $\Delta_r(p)$ over targets, and choose a maximizer; by linearity of expectation this maximizes expected surviving size over the list. Forward traversal constructs all orbits in $O(n^2)$ parent visits per candidate, so the selector uses $O(cn^2)$ expected elementary operations with hashed sets; precomputed powers reduce scoring to linear arithmetic once exposure sizes are known. No optimum for simultaneous multiple additions is claimed.

5.4 Fluctuations and sufficient conditions for concentration

Theorem 5.8 (functional-source variance). *In a functional source with $E_r = O_f(r)$ and $X = |M_f(A_p)|$,*

$$\text{Var}(X) = \sum_{r,s \in R} \left(p^{|E_r \cup E_s|} - p^{|E_r|+|E_s|} \right).$$

Proof. By Theorem 5.5 in the one-parent case, $X = \sum_r I_r$ with I_r the indicator of $E_r \subseteq A_p$. For each ordered pair $\mathbb{E}I_r I_s = p^{|E_r \cup E_s|}$ and $\mathbb{E}I_r \mathbb{E}I_s = p^{|E_r|+|E_s|}$; expand the variance of the finite sum. The identity includes the diagonal and holds at the endpoints. $\square$

Corollary 5.9 (sparse overlap suffices for concentration). *Let K_n count the ordered pairs of reactions whose exposures intersect in an n -reaction source. For $0 \leq p \leq 1$,*

$$0 \leq \text{Var}(X/n) \leq K_n/n^2.$$

If $K_n = o(n^2)$ along a source sequence then $\text{Var}(X/n) \rightarrow 0$, so X/n concentrates around its own expectation: for every $\varepsilon > 0$, $\mathbb{P}[|X/n - \mathbb{E}X/n| > \varepsilon] \rightarrow 0$. No limit for the expectations is assumed.

Proof. A disjoint pair has zero covariance because union cardinality is then the sum. Every other covariance is nonnegative, since the union has at most the sum of the cardinalities and $p \in [0, 1]$, and is at most one. Sum and divide by n^2 ; Chebyshev's inequality gives the probability statement. The bounds also hold for varying retention probabilities provided the overlap condition holds. $\square$

Orbit-size marginals do not determine the variance. The parent arrays $(0, 2, 0, 4, 0)$ and $(0, 2, 1, 1, 1)$ on five reactions both have sorted orbit sizes $(1, 2, 2, 3, 3)$ and sorted nonzero parent fan-outs $(1, 1, 3)$; at $p = 1/2$ both means are $5/4$, but the variances are $9/4$ and $33/16$. Theorem 5.8 explains the difference exactly through the pairwise unions.

5.5 Modules and gateways: two sharp controls

Proposition 5.10 (disjoint catalytic modules). *Let the source be a nonempty finite family of disjoint private-product catalytic cycles of lengths $\ell_1, \dots, \ell_m$, possibly sharing food. Then, with $X = |M(A_p)|$,*

$$\Phi(p) = \sum_j \ell_j p^{\ell_j}, \quad F(p) = 1 - \prod_j (1 - p^{\ell_j}), \quad \text{Var}(X) = \sum_j \ell_j^2 p^{\ell_j} (1 - p^{\ell_j}), \quad \Phi'(1) = \sum_j \ell_j^2,$$

Every singleton deletion destroys its entire module. If N_k modules have length k and $n = \sum_k k N_k$, a uniformly random singleton deletion has loss size k with probability $k N_k / n$, and

$$\text{Var}(X/n) \leq \frac{\sum_j \ell_j^2}{n^2} \leq \frac{\ell_{\max}}{n}.$$

Proof. A module survives exactly when all its reactions are present, and the module indicators are independent because their reaction sets are disjoint. The formulas follow; the size-bias statement counts reactions per module, and $\ell_j^2 \leq \ell_{\max} \ell_j$ gives the last bound. $\square$

The size bias is a statement about counting, not about critical branching: module frequencies proportional to $k^{-5/2}$ on a finite range produce reaction-weighted loss frequencies proportional to $k^{-3/2}$ with no cascade between modules. A largest module of size $o(n)$ suffices for concentration without any assumption on the module-size distribution.




Figure 5: Exact robustness curves, not simulation traces. Left: expected surviving fraction for the two-producer cooperative source ($\Phi = 2p + 2p^2 - p^3$ on three reactions), a catalytic three-cycle ($\Phi = 3p^3$), and a 12-reaction gateway ($\Phi = p + 11p^2$). Right: probability of any RAF for the same sources ($2p - p^2$, p^3 and p). The gateway retains $F(p) = p$ at every size; size and nonempty survival are distinct observables.

Proposition 5.11 (common gateway). *Let every reaction, including reaction 0 itself, be catalysed by the private product of reaction 0, on n reactions. Then $F(p) = p$ for every n , and conditional on retaining reaction 0 the surviving size is $1 + Y$ with $Y \sim \text{Bin}(n - 1, p)$ independent of that retention. Hence*

$$\Phi(p) = p + (n - 1)p^2, \quad \text{Var}(X) = p(1 - p)[1 + (n - 1)p]^2 + (n - 1)p^2(1 - p),$$

and for fixed $0 < p < 1$ the normalized variance $\text{Var}(X/n) \rightarrow p^3(1 - p) > 0$.

Proof. Nonempty survival occurs exactly when reaction 0 is retained, since every orbit passes through it and $\{0\}$ is a RAF. The conditional law and the law of total variance give the formulas. $\square$

A persistent shared cause defeats concentration even though the availability coordinates are independent. This does not contradict Corollary 5.9: every pair of gateway exposures intersects, so $K_n = n^2$. Figure 5 plots the three exact curves.

5.6 Discriminating computations

An exact preflight exhausts all 27 functional maps on three reactions and their one-site resolutions: 729 external-cut classifications and 2187 target-gain identities at $p \in \{1/4, 1/2, 3/4\}$, 81 functional variance identities, curvature on all 343 elementary three-reaction graphs with nonempty catalytic rows, and curvature on 40 additional four-reaction sources with varying food dependencies evaluated on their actual maxima. The cooperative fixture gives curvature -2 and the three-cycle 18; the nine candidate objectives in the selector fixture were checked against direct literal evaluation. All checks pass in rational arithmetic in 0.167s. These computations discriminate the stated formulas; they do not replace the proofs. Three scope boundaries were established the same way: an added catalyst can create a new cycle (Figure 3); with several ambiguous reactions, independent local choices need not be represented by two global parent maps; and nonfood reactants invalidate the orbit adapter although Lemma 5.4 still applies.

6 Formal verification

All theorems marked formal below are compiled in Lean 4 [8], toolchain `leanprover/lean4:v4.30.0`, against Mathlib [9] at the matching tagged release with a frozen manifest. Verification treats every warning as an error, rejects `sorry` and admitted proofs, extracts the requested declaration signatures, records the axioms used, and stores SHA-256 hashes of the source, the compiled module and the declaration interface in a receipt. Standard logical axioms (`propext`, `Classical.choice`, `Quot.sound`) are permitted; their use does not make the Python implementations verified programs.

6.1 Architecture

The development is organized in three namespaces over a common finite reaction-system library (RAF) providing closure, pruning, `evaluate` (the maximum RAF) and their monotonicity lemmas.

- **RAFQueryCompilation**: `RankedWitness` (Definition 3.1, the decidable checker, Theorem 3.3, Theorem 3.5), `RankedExistence` (Lemma 3.2), `ResidualRAF` (Theorem 3.4), the closure and pruning certificate chain of Section 3.3, `HybridQuery` (exact fallback), the charged and budgeted evaluators, `FamilyArithmeticRun` (packet refinement) and `TerminalFamily` (Theorem 3.7).
- **RAFSupportSelection**: `SelectedCone` and `Objective`, `PerfectQuery`, `IntrinsicDependence`, `OneLayerSource`, `MeanSetCover` with `Aggregator`, `PortfolioBoundary`, `PortfolioCoverage` and `SourcePortfolio`, `ExponentialPortfolio`, `SharedWitnessObstruction`, `ClosedSeed`.
- **RAFReactionCriticality**: `LossClosure`, `GeneralPivotality` and `SecondOrder` (finite product-measure derivatives), `RAFCurvature`, `CatalystChoice`, `FunctionalSource` and `Main`, `TargetCuts`, `CatalystIntervention`, `ExposureVariance`, `FunctionalVariance`, `ExposureConcentration`, `CatalyticModules`, `ModularSurvival`, `ModuleSensitivity`, `Gateway`.

6.2 Concordance

Table 3 maps each result to its declarations. Some remarks on the formal statements. Perfect-query and intrinsic-dependence statements assume the baseline is a pruning fixed point; deletion sets are arbitrary finite sets, with reactions outside the baseline ignored by intersection with S . Source-specific statements instantiate literal reactant, product, food and catalyst data. The one-layer selector assumes nonempty allowed lists and natural weights; rational distributions reduce to this by scaling. The greedy theorem assumes a nonempty finite valid pool and $b \geq 1$, and proves the cardinality bound, exactness for every K , and the coverage inequality against every $\mathcal{B} \subseteq H$ of size at most b . The derivative statements are `HasDerivAt` facts about finite product-measure polynomials; the second-order statement counts ordered pairs, which is the factor two.

6.3 What is proved by hand

The following consequences are proved in this paper and not separately claimed as Lean declarations: the polynomial encoding and NP membership in Theorem 4.8; the Taylor remainder in Corollary 5.3; the Chebyshev step in Corollary 5.9; the minimum-cardinality cases following Theorem 5.6; the strict-positivity criterion in Theorem 5.7; the closed-form variances and the size-bias reading in Propositions 5.10 and 5.11; the numerical bookkeeping in the proof sketch of Theorem 3.7, which is however contained in the formal proof through

Result	Module and principal declarations
<i>Namespace</i> RAFQueryCompilation	
Witness checker; Thm. 3.3; Thm. 3.5	RankedWitness: checkRankedSupport_sound; ranked_retained_fixed; ranked_deletion_localization
Lemma 3.2	RankedExistence: supported_ranked_exists; raf_ranked_exists
Thm. 3.4	ResidualRAF: residual_supported; evaluate_residual
Thm. 3.7	TerminalFamily: ModuleFamily.family_terminal_factor; setup cap in FamilySetup.familySetupCap
<i>Namespace</i> RAFSupportSelection	
Prop. 4.1	IntrinsicDependence: expected_cone_decomposition
Thm. 4.2	PerfectQuery: perfect_query_certificate; loss_iff_every_certificate
Thm. 4.3; Thm. 4.4	IntrinsicDependence: indispensable_transitive; mutual_indispensability_equivalence; synergy_obstructs_perfect_singletons
Example 4.5	SharedWitnessObstruction: OneLayer.two_producer_gap; OneLayer.balanced_two_producer_certificate
Thm. 4.6	OneLayerSource: OneLayer.automatic_selector_globally_optimal
Thm. 4.7	MeanSetCover: Aggregator.mean_selection_iff_set_cover; Aggregator.source_baseline
Thm. 4.9	PortfolioBoundary: portfolioQuery_correct
Thm. 4.10	SourcePortfolio: source_portfolio_selection; PortfolioCoverage: greedyPortfolio_guarantee
Thm. 4.11	ExponentialPortfolio: Paired.exponential_lower_bound; Paired.full_portfolio_exact; Paired.two_certificates_singletons
Prop. 3.6	ClosedSeed: closed_seed_exact_loss
<i>Namespace</i> RAFReactionCriticality	
Prop. 2.3; Prop. 5.1	LossClosure: loss_idempotent; singleton_loss_nested; GeneralPivotality: size_endpoint; survival_endpoint
Thm. 5.2	SecondOrder: second_endpoint; RAFCurvature: raf_second_endpoint (ordered-pair form)
Lemma 5.4	CatalystChoice: evaluate_one_site_choice
Thm. 5.5	FunctionalSource: evaluate_iff_orbit; Main: redundant_orbit_law; redundant_target_probability
Thm. 5.6	TargetCuts: classification; loss_iff_cut
Thm. 5.7	CatalystIntervention: intervention_gain; total_intervention_gain; singleton_vulnerability_reduction
Thm. 5.8; Cor. 5.9	ExposureVariance: exposure_variance; FunctionalVariance: functional_variance; normalized_concentration; ExposureConcentration: exposure_bounds
Prop. 5.10	CatalyticModules: singleton_loss; mean_size; size_bias_probability; ModularSurvival: survival_probability; ModuleSensitivity: mean_derivative_one
Prop. 5.11	Gateway: gateway_nonempty_iff; gateway_survival_probability

Table 3: Concordance between the results of this paper and the compiled declarations. Declaration names are given without their namespace prefix.

the executable tariffs. The general Python witness proposer is checked by small independent oracles, not by an implementation-refinement theorem; the one-layer and finite-pool selectors have their own compiled contracts.

7 Discussion

7.1 What the three parts say together

The results separate three things that are easily conflated: the structural loss $L(K)$, the witness used to localize it, and the computation that certifies the answer. The loss is intrinsic: it is what every witness agrees on (Theorem 4.2), it induces a preorder on the baseline (Theorem 4.3), and its first and second moments under random deletion are exact functions of singleton and pair losses (Proposition 5.1, Theorem 5.2). The witness is a choice: it can be optimized exactly on separable classes, it is NP-hard to optimize in general because a downstream aggregator makes shared ancestry observable, and no single witness can be perfect when losses are synergistic. The computation is exact whatever the witness, with a cost that is the region size plus boundary and certificate work, and with an unbounded structural amortization factor in a declared charge model but a measured break-even that depends on cold costs and workload length.

On the functional one-site class, the two exposures give a complete answer: shared exposure predicts singleton vulnerability, exclusive exposure identifies genuinely protective redundancy, and pairwise unions control joint fluctuations. Outside that class the curvature identity remains valid and separates overlap from cooperation, a distinction invisible to the mean singleton loss.

7.2 Relation to incremental maintenance

Delete/rederive [10] and backward/forward maintenance [11] recover a materialization after deletion by over-deleting the derivations that used a deleted fact and rederiving what survives through alternative derivations. The ranked support witness plays the role of a stored derivation, and the cone is the over-deletion region; the residual solve is the rederivation. Two differences are specific to RAFs. The fixed point is a least closure nested inside a greatest pruning, so rederivation must handle the pruning of reactions whose catalysts disappear; and the catalytic parent is exempt from the rank order, so the stored derivation is not a proof tree. Theorem 3.3 is the statement that the retained part is nonetheless stable, and Theorem 4.11 is a representation-size result of the kind studied in knowledge compilation [12]: singleton queries are cheap to compile and arbitrary deletion queries are not.

7.3 Limitations and open problems

- *Source classes.* The exposure laws assume elementary food generation, private products and one alternative catalytic site. Several ambiguous sites, or nonfood reactants, need a new source-derived witness structure. Lemma 5.4 is the only part that transfers unchanged.
- *Selection.* The NP-completeness concerns the uniform-mean objective. The worst-singleton objective, the complexity of the two-producer discriminator’s minimax value on general sources, and approximation algorithms with guarantees relative to the intrinsic floor are open. Whether a bounded-arity aggregator recovers the exponential wall of Theorem 4.11 is open.
- *Amortization.* Theorem 3.7 compares with a specified fresh evaluator in a charge model. A theorem against every algorithm for a natural workload, or an amortized bound on a family with nontrivial regions, would be substantially stronger.

- *Biology*. No biological network was analysed. The deletion analysis of [4] removes molecules as well as reactions and requires an explicit adapter; the present results are the structural and computational groundwork for such an analysis, not a substitute for it.
- *Kinetics*. None of the results assert kinetic persistence, positive steady states or thermodynamic feasibility of the surviving set.

7.4 Reproducibility

The Lean sources are in `proofs/RAFQueryCompilation`, `proofs/RAFSupportSelection` and `proofs/RAFReactionCriticality`; each module has a strict verification receipt recording the toolchain, the manifest hash, the requested declarations, their interfaces and the axioms used. The preflight scripts and raw measurement outputs accompany the corresponding workspaces: the functional-source preflight and the selector comparison for Section 5; the held-out profiles, closed-seed discriminator, native selector check and stage-cost audit for Section 4; the profile-frontier and cold-profile reports, with saved packets and exact mismatch assertions, for Section 3. Every measurement is a single observation or a median over distinct source instances; no query is repeated to manufacture amortization. Source hashes distinguish instances.

Acknowledgements

To be added.

References

- [1] Stuart A. Kauffman. Autocatalytic sets of proteins. *Journal of Theoretical Biology*, 119(1):1–24, 1986. doi: 10.1016/S0022-5193(86)80047-9.
- [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] Wim Hordijk, Joshua I. Smith, and Mike Steel. Algorithms for detecting and analysing autocatalytic sets. *Algorithms for Molecular Biology*, 10:15, 2015. doi: 10.1186/s13015-015-0042-8.
- [4] Filipa L. Sousa, Wim Hordijk, Mike Steel, and William F. Martin. Autocatalytic sets in *E. coli* metabolism. *Journal of Systems Chemistry*, 6:4, 2015. doi: 10.1186/s13322-015-0009-7.
- [5] Wim Hordijk and Mike Steel. Chasing the tail: The emergence of autocatalytic networks. *Biosystems*, 152:1–10, 2017. doi: 10.1016/j.biosystems.2016.12.002.
- [6] 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.
- [7] Daniel H. Huson, Joana C. Xavier, and Mike Steel. Self-generating autocatalytic networks: structural results, algorithms and their relevance to early biochemistry. *Journal of the Royal Society Interface*, 21(214):20230732, 2024. doi: 10.1098/rsif.2023.0732.
- [8] 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.

- [9] 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.
- [10] Ashish Gupta, Inderpal Singh Mumick, and V. S. Subrahmanian. Maintaining views incrementally. In *Proceedings of the 1993 ACM SIGMOD International Conference on Management of Data*, ACM SIGMOD Record 22(2), pages 157–166. ACM, 1993. doi: 10.1145/170036.170066.
- [11] Boris Motik, Yavor Nenov, Robert Piro, and Ian Horrocks. Incremental update of Datalog materialisation: the backward/forward algorithm. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pages 1560–1568. AAAI Press, 2015. doi: 10.1609/aaai.v29i1.9409.
- [12] Adnan Darwiche and Pierre Marquis. A knowledge compilation map. *Journal of Artificial Intelligence Research*, 17:229–264, 2002. doi: 10.1613/jair.989.
- [13] Richard M. Karp. Reducibility among combinatorial problems. In Raymond E. Miller and James W. Thatcher, editors, *Complexity of Computer Computations*, pages 85–103. Plenum Press, New York, 1972. doi: 10.1007/978-1-4684-2001-2_9.
- [14] George L. Nemhauser, Laurence A. Wolsey, and Marshall L. Fisher. An analysis of approximations for maximizing submodular set functions—I. *Mathematical Programming*, 14:265–294, 1978. doi: 10.1007/BF01588971.
- [15] Uriel Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM*, 45(4): 634–652, 1998. doi: 10.1145/285055.285059.
- [16] Grigori A. Margulis. Probabilistic characteristics of graphs with large connectivity. *Problemy Peredachi Informatsii*, 10(2):101–108, 1974. In Russian; English translation in Problems of Information Transmission 10 (1974), 174–179.
- [17] Lucio Russo. On the critical percolation probabilities. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete*, 56:229–237, 1981. doi: 10.1007/BF00535742.
- [18] Ryan O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014. doi: 10.1017/CBO9781139814782.
- [19] Mike Steel and Wim Hordijk. Tractable models of self-sustaining autocatalytic networks, 2018. arXiv:1801.03953v2.