Positive Cut-Word Models for the Kanade–Russell I_1 Bijection Problem

A general cancellation theorem, reversible local constructions, global obstructions, and supplementary Lean verification

Christopher Yamba

Research with a GPT-6.1 Terra agent fleet

Coordinated mathematical synthesis and independent criticism

27 October 2026

Abstract

We investigate a structural, weight-preserving bijection for the full modulo-nine Kanade–Russell identity KR1. The requested residue-to-gap bijection is not obtained. The main retained result is instead a general constructive positivity theorem for normalized divided differences. An ordered word of cuts and colored particles converts every signed factor into a matching staircase increment. A least-eligible sign-reversing involution works for arbitrary numbers of cuts and markers. Applied to Xia's auxiliary product sequences, it gives explicit positive models for all A_n, B_n, C_n, not merely bounded values of n. We also prove a quantitative maximum-part obstruction for any injection on a bounded-multiplicity source family, supply a reversible local bidegree surgery, audit the existing gap-to-cluster endpoint, and retain exact failures of several carry and recurrence recipes. Lean 4 verifies selected universal components; the source interpretation and full word encoding have separate written review. Novelty relative to the literature has not been established.

Scope of the result. The all-index positive auxiliary models are proved by a written construction. They do not yet supply the branch bijections needed to reach all KR1 gap partitions. No weight-specific matching, ranking, or apex lookup is used to define a claimed global map. The complete KR1 bijection remains unfinished.

1 The exact target and the repository baseline

Let $\mathcal{B}_N$ be the partitions of N with every part congruent to $1, 3, 6, 8 \pmod 9$. Let $\mathcal{A}_N$ be the ascending positive partitions $\lambda_1 \le \cdots \le \lambda_\ell$ satisfying

$$\lambda_{i+2} - \lambda_i \ge 3, \tag{1}$$

$$\lambda_{i+1} - \lambda_i \le 1 \implies \lambda_i + \lambda_{i+1} \equiv 0 \pmod 3. \tag{2}$$

The target consists of explicit algorithms $F : \mathcal{B}_N \to \mathcal{A}_N$ and $G : \mathcal{A}_N \to \mathcal{B}_N$, for every $N \ge 0$, with validity, weight, termination, and both inverse identities proved.

The familiar series is

$$\sum_{a,b\ge 0} \frac{q^{a^2+3ab+3b^2}}{(q;q)_a (q^3;q^3)_b} = \frac{1}{(q, q^3, q^6, q^8; q^9)_\infty}. \tag{3}$$

Here $(z;q)_r = \prod_{j=0}^{r-1}(1-zq^j)$. Kurşungöz's cluster construction identifies the sum with the gap class [4]. Xia and Mizuno give recent analytic treatments [1, 2]. The Palomar record for Mizuno's formalization reports Lean and independent nanoda checking of the three symmetric identities at a pinned source revision [3]. We inspected the registry metadata, pinned challenge statement, and solution wrapper, but did not locally rebuild that external development. The challenge template's placeholders are separate from the submitted solution. An analytic or formal power-series identity is different from the structural algorithm sought here.

The investigation starts from the corrected KR1 branch at `4ccc4fe`, which follows PR #15 at `2ccf9c5`. PR #16 proves the restricted $a = 3, b = 0$ sector for all its parameters and a fixed-column-alphabet obstruction [8]. The earlier global table fails at weight 41: two distinct cluster inputs share $U = (19, 12, 6, 3, 1)$ [7]. That failure is retained, not patched by adding entries. The newer `watch/openai-math` branch is a source-monitoring change and does not close KR1.

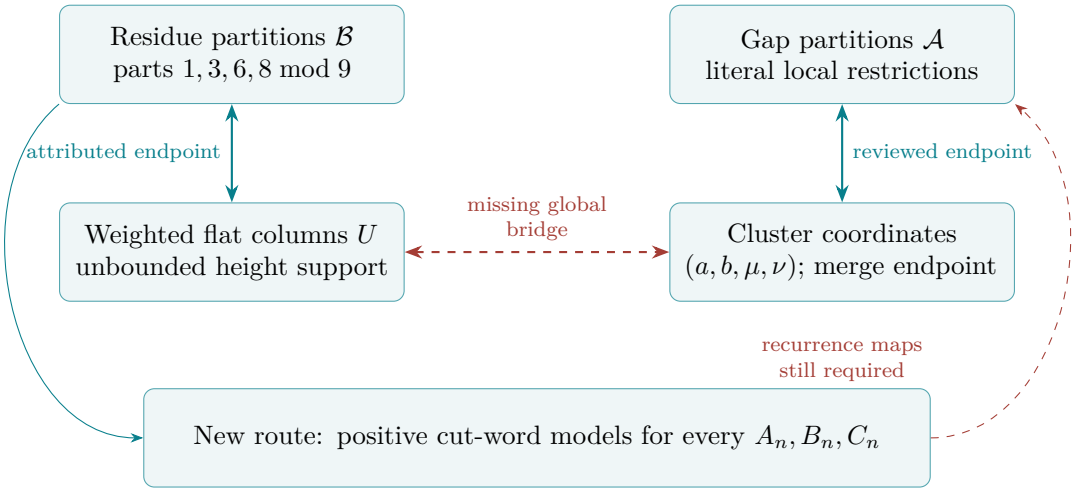


Figure 1: The full objective is unchanged. The new theorem removes signs from all auxiliary indices, but does not complete either dashed connection.

2 A general positive divided-difference theorem

The construction is more general than the specific KR1 parameters. Fix integers $d, c > 0$, $n \geq 0$, and $0 < r_i < c$ for $0 \leq i \leq n$. Assume the formal nodes $z_i = q^{r_i} + q^{c-r_i}$ are distinct. Put

$$\Phi_{d,c}(z) = \prod_{j \geq 0}(1 - zq^{dj} + q^{c+2dj})^{-1}, \qquad D = \Phi_{d,c}[z_0, \ldots, z_n].$$

Divided differences have their ordinary interpolation meaning; all evaluated expressions are interpreted coefficientwise as formal series.

2.1 Colored cut objects

For $n \geq 1$ choose $0 \leq k_1 \leq \cdots \leq k_n$, and set $k_0 = 0$. Layer $i < n$ contains particles at indices $k_i \leq j \leq k_{i+1}$; layer n allows every $j \geq k_n$. At each allowed cell there are independently repeatable particles

$$\alpha_i(j) : r_i + dj, \qquad \beta_i(j) : c - r_i + dj.$$

Each cut k_i contributes dk_i to the weight. A cell is identified by its layer, index, and color, even when two numeric weights happen to coincide. For $n = 0$ there is one layer and no cuts.

Lemma 2.1 (Cut expansion)**.** *The generating function of these colored cut objects is D.*

Proof. At z_i, the jth factor splits as $(1-q^{r_i+dj})^{-1}(1-q^{c-r_i+dj})^{-1}$. For a finite product, the divided-difference product rule assigns the n divided-difference steps, weakly in order, to factor indices $k_1 \le \cdots \le k_n$. A step at index j contributes q^{dj}, and each factor acquires exactly the node denominators for the layers meeting that index. Thus it gives the displayed cells, including both endpoints and repeated cuts. This is Xia's Lemma 2.2 [1]. Passing to the infinite product is legitimate coefficientwise: positive particle bases bound the number of particles, and the cut charge bounds the cut indices whenever they are nonzero. □

Let $\mathcal{W}_n$ be all finite words with exactly n separators $|$ and any number of letters α, β. Their length is L. The layer of a particle is the number of preceding separators. Define its intrinsic weight by

$$b(w) = \sum_{\alpha \text{ in layer } i} r_i + \sum_{\beta \text{ in layer } i} (c - r_i).$$

A descent is an adjacent $\beta\alpha$ in the word; a separator interrupts adjacency. Define

$$\operatorname{comaj}(w) = \sum_{\substack{1 \le s < L \\ w_s = \beta,\ w_{s+1} = \alpha}} (L - s).$$

Theorem 2.2 (Positive cut-word normal form)**.** *For every $0 \le m \le n$,*

$$(q^d; q^d)_m D = \sum_{w \in \mathcal{W}_n} \frac{q^{b(w) + d\operatorname{comaj}(w)}}{\prod_{h=m+1}^{L(w)} (1 - q^{dh})}. \tag{4}$$

Empty products have value one. The right side is the generating function of an explicit positive family with mutually inverse word coordinates. In particular, $(q^d; q^d)_m D \in \mathbb{N}[[q]]$.

Proof. **1. Flatten the object.** Read layer 0, cut 1, layer 1, and so on. Within each layer sort particles by index, breaking ties by $\alpha < \beta$. Record the resulting word w and the index j_t of every symbol, including cuts. Then $j_1 \le \cdots \le j_L$ and a strict rise occurs whenever $w_s w_{s+1} = \beta\alpha$.

Conversely, these are the only restrictions. A weakly increasing index list with those strict rises sorts correctly inside each layer; each particle lies between the surrounding cut indices, inclusively. Reading the cut symbols recovers the cuts and the layer labels. Thus flattening and decoding are mutually inverse, including repeated equal cuts and empty layers.

2. Subtract the forced staircase. If $L = 0$, there are no indices or gap coordinates and every sum below is empty. For $L > 0$, proceed as follows. Let $\epsilon_s = 1$ at a descent and 0 otherwise. Set

$$\delta_1 = j_1, \qquad \delta_t = j_t - j_{t-1} - \epsilon_{t-1} \quad (2 \le t \le L).$$

All δ_t are nonnegative. The inverse uses cumulative sums. The index charge is exactly

$$d \sum_{t=1}^{L} j_t = d \sum_{t=1}^{L} (L - t + 1)\delta_t + d\operatorname{comaj}(w). \tag{5}$$

For a fixed word, independent gap coordinates therefore give denominator $\prod_{h=1}^{L}(1 - q^{dh})$.

3. Cancel every marker jointly. Expand $\prod_{h=1}^{m}(1 - q^{dh})$ as signed marker sets. Marker h has weight dh and changes sign. Because $L \ge n \ge m$, the coordinate δ_{L-h+1} exists and has exactly the same weight. Select the smallest h for which either its marker is present or $\delta_{L-h+1} > 0$. If marked, remove the marker and add one to that coordinate. Otherwise insert the marker and

subtract one. Earlier coordinates stay inactive and the selected one stays active, so the same h is selected on reversal. This is a weight-preserving, sign-reversing involution.

Its fixed points have no markers and

$$\delta_{L-m+1} = \cdots = \delta_L = 0. \tag{6}$$

The other coordinates are free. Their generating function is (4). For each fixed weight the signed family is finite: positive particle bases bound $L - n$, and all coordinates have positive weight. Cancellation therefore requires no convergence assumption beyond formal coefficientwise finiteness. □

What is structural here? The word, layer labels, and all gap coordinates are recoverable from an individual cut object. The cancellation examines one coordinate at a time. It neither lists partitions of the same weight nor invokes the KR1 identity. The staircase principle is classical in character; priority for this particular general formulation and its KR1 specialization is not claimed.

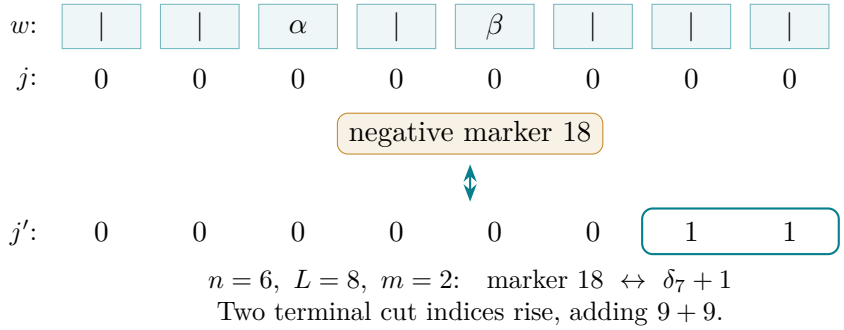


Figure 2: A concrete joint cancellation. The intrinsic particle weights are 28 and 32. Both signed partners have weight 78. The suffix-coordinate interpretation keeps earlier marker cancellations intact.

2.2 A forbidden initial particle

Corollary 2.3. *Multiplying the left side of* (4) *by* $(1 - q^{r_0})$ *also has a positive model. Replace the exponent on the right by* $b(w) + d\,\mathrm{comaj}(w) + dL\epsilon_0(w)$*, where* $\epsilon_0(w) = 1$ *exactly when the first symbol is* α*; set* $\epsilon_0(\varnothing) = 0$*.*

Proof. Layer 0, index 0 always exists. Its α factor can be cancelled directly, meaning that cell is forbidden. In a canonically ordered word, an α at index 0 in layer 0 would necessarily be the first symbol. Thus the exclusion is exactly $j_1 \geq 1$ when $w_1 = \alpha$. For a nonempty word, use $\delta_1 = j_1 - \epsilon_0$ instead. The empty word has no indices, no gap coordinates, and weight zero. Formula (5) acquires $dL\epsilon_0$ and the marker involution is unchanged. □

The restriction $m \leq n$ cannot be dropped universally: for $n = 0, m = 1, d = 1, c = 5, r_0 = 2$, the expression $(1 - q)/((q^2; q)_\infty (q^3; q)_\infty)$ has coefficient -1 at q.

3 Positive models for every KR auxiliary index

Put

$$\mathcal{E}_9 = \frac{1}{(q^3, q^6; q^9)_\infty}, \qquad \mathcal{D}_c(r_0, \ldots, r_n) = \Phi_{9,c}[q^{r_0} + q^{c-r_0}, \ldots, q^{r_n} + q^{c-r_n}].$$

The product sequences used by Xia [1, eqs. (3.5)–(3.7)] are

$$C_n = \mathcal{E}_9(q^3;q^3)_n \mathcal{D}_{9n+9}(3n+4, 3n+7, \dots, 6n+4), \tag{7}$$

$$A_n = \mathcal{E}_9(q^3;q^3)_n (1-q^{3n-1}) \mathcal{D}_{9n}(3n-1, 3n+2, \dots, 6n-1) \quad (n \geq 1), \tag{8}$$

$$B_n = C_n + q^{3n+2} A_{n+1}. \tag{9}$$

The separately defined A_0 is the KR1 product. Distinctness of the nodes follows because equality of unequal complementary exponents would require $3(i+j) = 3n+1$ or $3n+2$, respectively.

Let $m = \lfloor n/3 \rfloor$ and define the positive common tail

$$T_n = \prod_{\substack{j>n \\ 3\nmid j}} (1-q^{3j})^{-1}.$$

The essential cancellation of scalar factors is

$$\mathcal{E}_9(q^3;q^3)_n = (q^9;q^9)_m T_n. \tag{10}$$

This follows by cancelling the factors with $3 \nmid j$ and retaining exactly $j = 3, 6, \dots, 3m$ from the numerator.

Corollary 3.1 (All-index auxiliary families)**.** *For every $n \geq 0$, C_n has the positive model consisting of a common-tail partition, a word with n cuts, and nonnegative staircase coordinates with the last $\lfloor n/3 \rfloor$ coordinates zero. Its intrinsic particle bases are*

$$\alpha_i : 3n+4+3i, \qquad \beta_i : 6n+5-3i.$$

For every $n \geq 1$, A_n has the same form with bases

$$\alpha_i : 3n-1+3i, \qquad \beta_i : 6n+1-3i,$$

and the initial α exclusion in Corollary 2.3. B_n is the indicated disjoint union with a fixed weight shift.

Proof. Apply Theorem 2.2, Corollary 2.3, and (10). All bases are positive. The assertion needs only the cut expansion and the displayed definitions; it does not assume the analytic KR1 evaluation or the three-state recurrence. □

3.1 Why treating each cut separately fails

At the all-zero cut tuple, the C_n summand has coefficient -1 at q^9 for every $n \geq 3$. Indeed, the residual marker product contains $(1-q^9)$, while every remaining common or node particle has weight strictly greater than 9. There is no positive weight-9 particle to cancel that term inside the fixed cut.

This rules out a positive construction that preserves every cut tuple. It does not rule out positivity of the sum. For $n = 3$, increasing the last empty cut introduces charge 9. With particles present, a simple cut shift can remove an occupied boundary cell. The staircase construction handles both cases: its suffix increments move whichever mixture of cuts and particles is actually present. It also avoids the failure of cancelling a q^{18} marker by shifting one particle twice after a q^9 cancellation. Such a shift can re-enter a previously cancelled image; the joint least-coordinate rule preserves the already inactive coordinates.

3.2 An exact finite check, with its limits

The independent program `positive_coefficients.py` builds the positive word series directly (with A_0 separately implemented as its defining product) and compares it with a literal last-two-parts gap recursion. For $0 \le n \le 8$, A_n agrees with minimum-part bound $3n+1$ and C_n with bound $3n+3$, through weight 45, except that the explicitly costlier A_1 word traversal is capped at 30. All comparisons agree. There are 813 checked coefficients across these eighteen comparisons. This validates the transcription, especially the initial descent condition; it does not supply a bijection between these families.

4 The missing recurrence bijections

The gap side has a direct three-state branching description. With $x = 3n+1$, $y = 3n+2$, $z = 3n+3$, its transfer matrix is

$$M_n = \begin{pmatrix} (1+q^x)(1+q^y) & q^z(1+q^x) & q^{2z} \\ 1+q^y & q^z & q^{2z} \\ 1 & q^z & q^{2z} \end{pmatrix}. \tag{11}$$

The entries correspond to literal choices of small gap parts, so these branches have actual inverse insertion rules. This is the Chern–Li linked-ideal structure [5]. Xia proves the same matrix recurrence for the product sequences analytically [1].

Our positive families give concrete objects at every product-side state. What remains is to construct, for those *particular* objects, the disjoint branch maps of (11), prove recoverability of the branch labels, and prove termination to the empty boundary. Equality of their generating functions does not choose those maps.

4.1 A symmetry that almost gives a branch

The C_n and A_{n+1} particle bases satisfy the exact identities

$$\alpha_i^{C_n}(j) = \beta_{n+1-i}^{A_{n+1}}(j), \tag{12}$$

$$\beta_i^{C_n}(j) = \alpha_{n+1-i}^{A_{n+1}}(j). \tag{13}$$

Thus reversing layers and exchanging colors preserves every intrinsic particle weight. Unfortunately, it reverses the index order. Keeping all indices and merely inserting one cut already fails at $n = 2$, cuts $(0, 0)$, and a terminal α at index 1: the proposed target layer is bounded above by a zero cut.

A stronger candidate uses $w' =|$ reverse(swap(w)). Its intrinsic bases match, but its forced staircase changes. For a source descent ending at position t in length L, the minimum-weight change is

$$9(2t - L - 2).$$

Summing this over descents gives the exact drift. At $n = 1$, $w =| \beta\alpha$ has minimum weight 27, whereas $w' =| \beta\alpha \,|$ has minimum weight 36. Nonnegative free coordinates cannot repair a positive minimum-weight excess.

4.2 The next precise global obligation

Construct the product-side realization of the C_n row:

$$\mathscr{C}_n \longleftrightarrow \mathscr{A}_{n+1} \;\sqcup\; (\{3n+3\} \times \mathscr{C}_{n+1}) \;\sqcup\; (\{6n+6\} \times \mathscr{C}_{n+1}) \;\sqcup\; (\{6n+8\} \times \mathscr{A}_{n+2}).$$



The node weights match, but the extra cut also pays for the strict rise.
A future rule must transport this drift, change the word, or choose another branch.

Figure 3: An exact obstruction to a specific recurrence recipe. It does not refute the recurrence or a more flexible bijection.

Here script letters denote the positive auxiliary objects, not the original gap class; the singleton notation is a fixed weight shift. A solution must change more than the cut positions while retaining a fixed color word: modulo three, the unshifted A_{n+1} branch already requires the difference between color counts to be divisible by three. It must also respect the phase change when $n + 1$ is divisible by three: the common tail loses no new part at that step, but one additional terminal gap becomes locked. These are exact requirements, not evidence that the needed map exists in any proposed form.

There is also an obstruction to preserving the number of colored particles, even if colors, indices, and the common tail may otherwise change arbitrarily. In the $n = 0$ row at weight 8, the C_0 object $\beta_0(0)$ together with one common part 3 has one colored particle. The other source object is two copies of $\alpha_0(0)$, each of weight 4. The target A_1 branch has only a two-particle object at weight 8; the shifted C_1 branches have none; and the $q^8 A_2$ branch has the zero-particle empty object. Thus no complete realization of this row can preserve colored-particle count. Cut rotations alone cannot suffice.

5 A global obstruction to bounded expansion

The seven possible patterns in a block $\{3j + 1, 3j + 2, 3j + 3\}$ are

$$E = \varnothing, \quad A = (1), \quad B = (1, 2), \quad C = (1, 3), \quad D = (2), \quad F = (3), \quad G = (3, 3).$$

Every displayed part is translated by $3j$. The successor sets are

Current state	Allowed next states
E, A, B, D	E, A, B, C, D, F, G
C, F	E, D, F, G
G	E, F, G

Checking the boundary pairs and triples proves that this is a complete encoding, not merely a generating recurrence. The model itself is established prior work [5].

Let H be its 7×7 zero-one adjacency matrix in the listed order. The positive integer vector

$$v = (10000, 10000, 10000, 5321, 10000, 5321, 3473)^T$$

satisfies

$$H^3 v = (1584707, 1584707, 1584707, 843205, 1584707, 843205, 550363)^T \leq 159v.$$

Monotonicity and induction give a bound for *all* powers. Set $\max \varnothing = 0$. For $M \geq 0$, padding with empty blocks consequently shows

$$\#\{\lambda \in \mathcal{A} : \max \lambda \leq M\} \leq \frac{10000}{3473}\, 159^{\lceil M/9 \rceil}. \tag{14}$$

Theorem 5.1 (Maximum-part expansion barrier)**.** *Let $\mathcal{B}^*$ be the residue partitions with multiplicities at most seven in the $1,8$ classes and at most one in the $3,6$ classes. If an injection $f : \mathcal{B}^* \to \mathcal{A}$ satisfies $\max f(\lambda) \le \gamma \max \lambda + K$ for all inputs, with fixed finite real γ, K, then*

$$\gamma \ge \frac{\log 256}{\log 159} = 1.093959803\ldots.$$

Weight preservation is not required for this obstruction.

Proof. The first s nine-blocks allow $8 \cdot 2 \cdot 2 \cdot 8 = 256$ choices each, so there are 256^s source objects with maximum at most $9s$. For $\gamma \ge 0$, injectivity and (14) imply $256^s \le (10000/3473)159^{\gamma s + K/9 + 1}$ whenever the maximum bound is nonnegative. If γ is smaller than the stated threshold, the quotient of the two exponential terms is unbounded. A negative maximum bound is already impossible. If $\gamma < 0$, single-part inputs with unbounded maximum give an immediate contradiction. □

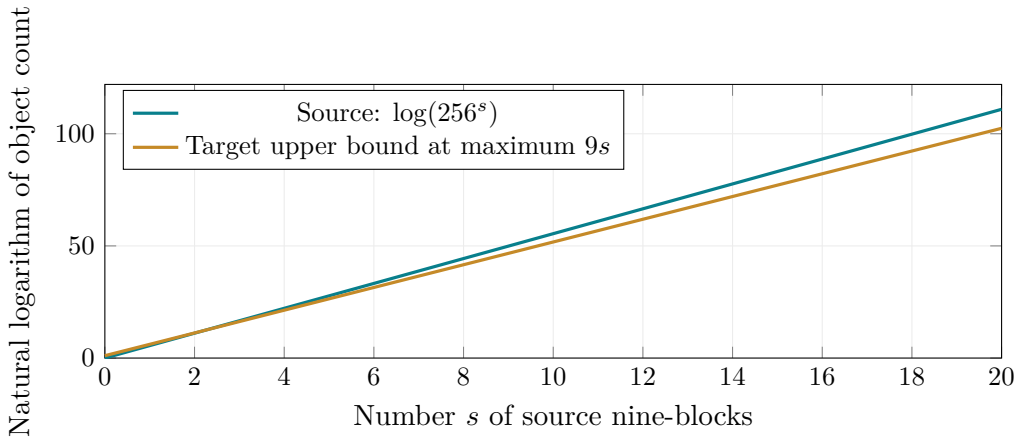


Figure 4: A slope mismatch survives every fixed additive offset. The theorem uses a bounded-multiplicity source subset; it is stronger than the immediate obstruction from arbitrarily many ones.

$\mathcal{B}^*$ is a subset, not a canonical normalization of $\mathcal{B}$. The tempting carry $1^8 \mapsto 8$ collides with the original source (8). Thus ordinary uncolored Glaisher carrying cannot be invoked to identify these two source families.

6 Reversible local surgery and recorded carry failures

6.1 A bidegree-preserving switch

There is a direct local bijection realizing

$$\frac{1}{(1-x)(1-y)(1-x^2y)} = \frac{1}{(1-y)(1-x^2)(1-xy)} + \frac{x}{(1-x^2)(1-xy)(1-x^2y)}. \tag{15}$$

The input is $(a,b,c) \in \mathbb{N}^3$, of bidegree $X = a + 2c,\ Y = b + c$. Put

$$r = X \bmod 2, \qquad p = \max\{0, \lfloor (\min(X,Y) - r)/2 \rfloor + 1\}.$$

If $c < p$, output the unmarked triple

$$(d,e,f) = (Y - r - 2c,\ (X - r - 2c)/2,\ r + 2c).$$

Otherwise put $h = (X - 1 - Y) \bmod 2$ and $g = h + 2(c - p)$, and output the marked triple

$$(e, f, g) = \big((X - 1 - Y - g)/2,\ Y - g,\ g\big).$$

The marker supplies one x.

For the inverse, reconstruct (X, Y) from the output, then use $c = (f - r)/2$ in the unmarked branch or $c = p + (g - h)/2$ in the marked branch. Finally $a = X - 2c$ and $b = Y - c$. At fixed (X, Y), the admissible input coordinate is the interval $0 \le c \le \min(\lfloor X/2 \rfloor, Y)$. The two output arithmetic progressions partition this interval. Splitting into $Y \ge X$, $Y \le \lfloor X/2 \rfloor$, and the intermediate range proves nonnegativity, coverage, and both inverses, including zero boundaries. This is fixed integer arithmetic on one local multiplicity coordinate; it is not a ranking of whole partitions.

The specialization $x = q^{9j+2}$, $y = q^{9k+3}$ maps chosen residue-$2, 3, 7$ species to residue-$3, 4, 5$ species or a marked alternative. To compose switches globally one must still recover j, k, branch labels, and shared-species context. These are not supplied by (15). The full formulas and the separate interval-proof review are in the supplement.

6.2 A reversible carry that exposes the missing information

For a part $s \equiv \pm 1 \pmod 9$, write its multiplicity $m_s = 4q_s + 2d_s + e_s$, where $d_s, e_s \in \{0, 1\}$. Replace it by e_s copies of s, d_s copies of $2s$, and q_s copies of $4s$; leave the ± 3 classes unchanged. The four residue lanes $\pm 1, \pm 2, \pm 4, \pm 3$ are disjoint, so division by $1, 2, 4$ recovers the parent species. This is a total weight-preserving auxiliary bijection. It does not produce the gap class: 1^8 becomes $(4, 4)$.

Proposed repair	Two distinct source inputs	Shared output or failure
$(4s, 4s) \mapsto 8s$	1^8 and (8)	(8)
$(4s, 4s) \mapsto (s, 7s)$	10^8 and $(10, 35, 35)$	$(10, 70)$
$(4s, 4s) \mapsto (2s, 6s)$	1^8 and $(1, 1, 6)$	$(2, 6)$
$(4s, 4s) \mapsto (3s, 5s)$	8^8 and $(10, 10, 10, 10, 24)$	$(24, 40)$
Leave $(4s, 4s)$	1^8	Invalid close pair $(4, 4)$

The collisions assume the other lanes are unchanged and uncolored. They do not exclude context-dependent coordinated repairs.

One origin-preserving alternative is the involution

$$(m_s, m_{rs}) = (rq + d, b) \longmapsto (rb + d, q), \qquad 0 \le d < r.$$

It preserves $m_s + rm_{rs}$ and exchanges the quotient with the neighboring multiplicity. Applying it to disjoint 8-adic pairs in the ± 1 lanes and 2-adic pairs in the ± 3 lanes gives a total involution of the source. It retains information but does not normalize every lower digit. On two species alone, such a normalization is impossible: $(r, 0)$ and $(0, 1)$ would have only one normalized target of weight rs.

7 Formal assurance, reproduction, and preserved failures

7.1 What Lean actually checks

The four standalone modules use `import Std` with Lean `4.34.0`. Each was freshly compiled; the verification record contains the exact source hashes, commands, exit codes, and printed axiom

reports. No reported theorem uses `sorryAx` or an added axiom. Standard foundational axioms are reported explicitly rather than described as an axiom-free development.

Module	Universally checked content	Boundary of the certificate
`Cancellation.lean`	Least-eligible involution on arbitrary finite lists; weight, sign parity, exact fixed points	Does not formalize divided differences, cut flattening, or the full KR1 map
`Blocks.lean`	Seven-state multiplicity classification, transitions, existence and uniqueness of compatible sequences	Sorted partition lists, finite support and weight translation remain written arguments
`Entropy.lean`	Integer matrix witness and path bound for every power	Source counting, partition encoding and logarithmic limit remain written
`MergeBounds.lean`	Boundary identities and the nonadjacent selection inequalities for the cluster merge	Does not verify the entire merge algorithm

The central formal signatures include

```
cancel_involution : cancel (cancel xs) = xs
cancel_weight    : weight start (cancel xs) = weight start xs
cancel_parity    : hasEligible xs = true ->
                   parity (cancel xs) = !(parity xs)
fixed_iff_all_zero :
  cancel xs = xs <-> forall x in xs, x = (false, 0)
walks_bound : walks n s <= 159^n * v s
```

These are readable renderings; the exact checked Lean declarations are supplied.

7.2 Independent checks and reviews

The positive-word construction was checked against the primary cut expansion by a distinct agent. A second implementation tested 3,178,008 bounded cut-word states for $n = 3, 4, 6$; an independently written test checked 124,944 states for $n = 3, 6$. Both checked coordinate recovery, the involution, sign, weight, and survivors. The local surgery test covered 27,000 input triples and bounded boxes in both output branches. The lane and paired-carry checks covered 5,461 source objects through weight 40. These numbers describe finite checks, not proof coverage.

Separate reviews support the existing direct cluster merge and the new local constructions. The previous $a \leq 2$ formulas were briefly rediscovered, then recognized as already present in RR003 and removed from the contribution list. A contradictory draft sentence about the height-block mask $\{3\}$ was corrected. The failed uncolored Glaisher normalization, sequential marker cancellation, and layer-reversal proposals remain in the failure ledger.

7.3 Run the supplement

The accompanying package preserves the new mathematical notes, four Lean files, exact checks, source-bound reviews, and a reproduction script. Run

```
python3 verify.py --lean /path/to/lean
pdflatex -interaction=nonstopmode -halt-on-error kr1-positive-models.tex
pdflatex -interaction=nonstopmode -halt-on-error kr1-positive-models.tex
```

The complete research archive also contains the historical regression context. The article and its illustrations are generated from this LaTeX source.

8 Research judgement and next theorem

The main useful change is global in the auxiliary parameter: the signed product states now have explicit positive objects for every n. That resolves a real obstacle which first appears at $n = 3$, and it works simultaneously for all residual markers. It is broader than a new restricted KR1 subfamily, but it is still an intermediate construction.

The intuition behind the successful step was to change coordinates before adding more cases. A marker of weight $9h$ should act on an actual suffix of h objects, not on an arbitrarily chosen particle. Flattening identifies that suffix intrinsically, including cuts. The failed attempts identify what a future recurrence must retain: origin information in carries, layer labels at equal cuts, the forced-descent weight, and the phase of n modulo three.

The next decisive target is the explicit four-branch correspondence in Section 4, followed by the other two matrix rows and a terminating boundary argument. A proposed map should first meet the weight-27 versus 36 witness, the occupied shared-boundary cases, and the phase where an additional terminal gap is locked. Increasing the size of a finite apex table would not meet those obligations.

Final claim boundary. Proved here: a general positive normalized divided-difference model, all-index positive KR auxiliary models, specified local bijections and obstructions, with scoped Lean certificates and independent reviews. Not proved here: an explicit structural bijection $\mathcal{B}_N \leftrightarrow \mathcal{A}_N$ for every N. No percentage of the open problem being solved is assigned.

A The audited gap-to-cluster endpoint

Every close pair is $P_{2r} = (3r + 1, 3r + 2)$ or $P_{2r+1} = (3r + 3, 3r + 3)$. Close pairs cannot overlap. Splitting a gap partition into these pair atoms and the remaining singles gives the boundary functions

	$L(\ell)$	$R(\ell)$
$\ell = 2r$	$3r - 1$	$3r + 4$
$\ell = 2r + 1$	$3r$	$3r + 6$

A single before P_ℓ must be at most $L(\ell)$; one after it must be at least $R(\ell)$. Successive pairs require a level difference of at least two. The key identity is $R(\ell) - 3 = L(\ell + 1) + 1$.

For single s_i with i earlier singles and q_i earlier pairs, and pair P_{ℓ_j} with j earlier pairs and z_j earlier singles, define

$$\mu_i = s_i - 3q_i - 2i - 1, \qquad \nu_j = \ell_j - 2j - z_j.$$

The boundary inequalities imply nonnegative increasing coordinate lists. Conversely, after emitting i singles and j pairs, compare $S = \mu_i + 3j + 2i + 1$ and $\ell = \nu_j + 2j + i$. Emit the single exactly when $S \le L(\ell + 1)$; otherwise emit P_ℓ. An exhausted list leaves the other type.

For nonadjacent competing atoms, the same threshold follows by bounding the intervening singles by steps of two, or the intervening pair levels by steps of two. These are the inequalities checked

in `MergeBounds.lean`. Extraction and merging recover each coordinate and each atom, hence are mutually inverse. Finally, $\sum_i q_i + \sum_j z_j = ab$, since each single-pair combination has one order. Summing the coordinates gives

$$|\lambda| = a^2 + 3ab + 3b^2 + |\mu| + 3|\nu|.$$

The existing repository proof of this endpoint received a fresh independent semantic review during this investigation. The source theorem and cluster idea are attributed to Kurşungöz; this endpoint does not supply the missing residue-to-cluster map.

References

[1] E. X. W. Xia, *The modulo 9 Kanade–Russell identities and their Nahm-sum duals*, arXiv:2609.08816v1 (2026). https://arxiv.org/html/2609.08816v1. Inspected: Lemma 2.2 and the auxiliary definitions and recurrences in Sections 3–4.

[2] Y. Mizuno, *The three Kanade–Russell identities modulo nine*, arXiv:2609.09653 (2026). https://arxiv.org/abs/2609.09653.

[3] Y. Mizuno, *KR-mod9-lean*, Palomar record PALOMAR-2026-09-10-000002, version 1. https://palomar-registry.org/entry?id=PALOMAR-2026-09-10-000002&version=1. Pinned source `d8b921339ecc`; registry-reported verification, not a local rerun.

[4] K. Kurşungöz, *Andrews–Gordon type series for Kanade–Russell conjectures*, arXiv:1808.01432, Theorem 7. https://arxiv.org/abs/1808.01432.

[5] S. Chern and Z. Li, *Linked partition ideals and Kanade–Russell conjectures*, Discrete Mathematics 343 (2020), 111876. https://arxiv.org/abs/1809.08655.

[6] G. E. Andrews and W. J. Keith, *Schmidt-type theorems for partitions with uncounted parts*, arXiv:2203.05202v2, Section 2 and Theorem 6. https://arxiv.org/abs/2203.05202.

[7] C. Yamba and contributors, *Phase 2: corrected tableau injections, beta recovery and bounded KR1 octant evidence*, Yamba AI Math PR #15. https://github.com/cyamba/yamba-ai-math/pull/15.

[8] C. Yamba and contributors, *Argon continuation: cone classification, all-parameter a3b0 proof and compiler obstruction*, Yamba AI Math PR #16. https://github.com/cyamba/yamba-ai-math/pull/16.