Komorebi Resolution Semantics

An algebra of mathematical worlds and their transformations

Komorebi P.

Version 0.1.0 September 26, 2026

Abstract

We develop the proposal that mathematical worlds can be studied as related resolutions rather than as an unstructured collection of competing universes. A small category of worlds, a functor of Boolean algebras, and a family of semantic valuations define a transport defect. This defect obeys an additive law, a twisted product rule, and an exact composition law. Its kernel is an invariant Boolean subalgebra, while its quotient is naturally an additive object.

A finite resolution theorem classifies every Boolean retraction onto the invariant sector and identifies canonical lower and upper envelopes when no canonical retraction exists. A two-world obstruction shows that even the smallest symmetric system need not admit a symmetry-preserving Boolean projector. Graph integrability gives an exact reconstruction criterion, and a separate real-valued realization recovers orthogonal multiresolution projectors. We also prove that any set-indexed family of first-order models admits a common Boolean-valued presentation, while explaining why this fact alone does not identify a shared generative structure.

The algebraic results are proved in the note and supported by reproducible exact checks. Forcing supplies a conditional set-theoretic interpretation, not a proof of new independence, consistency, or universe-existence claims. The framework is a precise research language and a collection of proved structural statements, not a completed extension of ZFC.

Central distinction. A world may be a resolution of richer data without being merely a change of coordinates of another world. Invariant content, residual disagreement, and the existence of a common presentation are three different mathematical questions.

Included with this note: editable LaTeX, an exact Python checker, machine-readable verification results, a proof-boundary ledger, and an explicitly uncompiled Lean source scaffold.

Contents

1 The claim worth formalizing 3

2 Worlds, presentations, and the ambient theory 3

3 A common presentation exists, but need not explain anything 4

4 The resolution differential 5

5 What the invariant core actually is 6

6 Finite integrability and reconstruction 8

7 Projectors: existence, obstruction, and a canonical substitute 9

8 A real-valued realization of Komorebi multiresolution 11

9 Stabilization requires a path quantifier 12

10 Forcing as a conditional set-theoretic instance 13

11 Where nontrivial transport geometry can enter 14

12 A strengthened research program, with explicit obligations 15

13 Verification and proof boundary 16

14 Conclusion 17

A Reproduction and artifact contents 18

B Notation and distinctions at a glance 18

1 The claim worth formalizing

The motivating idea is not that different models of set theory are secretly identical. It is that a family of models may possess a meaningful transformation structure, and that this structure deserves its own calculus. The primitive question changes from "Which world is the correct one?" to the more local and testable question

Which observables are preserved, and what exactly changes along a specified transformation?

This does not replace the ordinary semantics of a model. It organizes comparisons between models.

The phrase *worlds as resolutions* has three strengths. The weakest says that a family of worlds can be encoded in one object. The middle claim says that there are explicit evaluation maps from richer data to individual worlds. The strongest says that a restricted transformation grammar explains the worlds, compresses their relationships, and reconstructs their differences. Theorem 3.1 establishes the middle claim for arbitrary set-indexed model families, but does not establish the strongest one.

The development below makes five substantive improvements to the original proposal. It separates Boolean and additive quotients, corrects the relationship between invariant observables and coherent sections, classifies finite projectors, replaces an unjustified demand for canonicity by canonical bounds, and distinguishes genuine transport information from a truth differential that is exact by construction.

The elementary algebra here is not claimed as a new alternative to the existing foundations of forcing, modal logic, Boolean-valued models, or rough sets. These are essential precedents [1, 2, 3]. The candidate Komorebi contribution is the explicit combination of typed transport, invariant extraction, residual recording, and reconstruction under auditable contracts. A claim of mathematical novelty for that combination requires a broader comparison than this note supplies.

2 Worlds, presentations, and the ambient theory

Fix a first-order language L and an L-theory T. The intended case is the membership language with $T = \mathsf{ZFC}$, but the algebraic results do not depend on that choice. Work in an external metatheory sufficient to discuss sets, categories, and satisfaction for *set-sized* structures.

A world is an L-structure $M \models T$. A chosen system of worlds is a small category C. Its arrows are specified transformations, not automatically elementary embeddings, isomorphisms, or inclusions. Their meaning must be supplied by an instance of the framework.

Forcing extensions, inner-model inclusions, elementary embeddings, and interpretations have different preservation properties. They cannot be put into one untyped class of "equivalences." A diagram must record which kind of arrow is present and which formulas and parameters may be compared.

Proposition 2.1 (Coordinate changes preserve truth)**.** *If $h : M \cong N$ is an isomorphism of L-structures, then for every formula $\varphi(\vec{x})$ and tuple $\vec{a}$ in M,*

$$M \models \varphi(\vec{a}) \quad \Longleftrightarrow \quad N \models \varphi(h\vec{a}).$$

Consequently, two worlds that disagree on an L-sentence are not related by a mere isomorphic change of coordinates in that language.

Proof. Induct on formulas. Atomic relations are preserved by the definition of an isomorphism. Boolean connectives commute with the induction. For existential quantification, transport a witness through h in one direction and through h^{-1} in the other. □

There is an equally important distinction between transformation and information refinement. In the information order on $\{\bot, 0, 1\}$, where $\bot \leq 0$ and $\bot \leq 1$ but 0 and 1 are incomparable, a monotone refinement cannot change a settled 0 into a settled 1. A truth-changing world transformation is therefore not a monotone refinement of a fixed partial truth assignment. It may instead change the quantifier domain, the admissible objects, or the semantic realization.

2.1 Existence is a separate obligation

All world theorems in this note are conditional on the models and transformations under discussion. ZFC does not acquire a proof that a set-sized model of ZFC exists merely because we introduce notation for one. Nor does a formal development of the algebra prove the existence of a countable transitive model, a generic extension with prescribed extra properties, or a universe containing all worlds.

Satisfaction for a set-sized structure is an external, stratified construction. We do not add a self-applicable truth predicate for the ambient universe. The limited algebraic construction can be expressed by ordinary definitions. An additional assertion that a nonempty diagram of ZFC worlds exists has a different logical status.

3 A common presentation exists, but need not explain anything

Theorem 3.1 (Boolean product presentation)**.** *Assume the axiom of choice in the metatheory. Let I be a nonempty set and let $(M_i)_{i\in I}$ be a set-indexed family of nonempty L-structures. There is a full $\mathcal{P}(I)$-valued structure $\mathcal{M}$ such that for each $i \in I$, quotienting at the principal ultrafilter*

$$U_i = \{S \subseteq I : i \in S\}$$

recovers M_i up to isomorphism. If every $M_i \models T$, every axiom of T has Boolean value I in $\mathcal{M}$.

Proof. Take the domain $X = \prod_{i\in I} |M_i|$. Interpret function symbols coordinatewise. Set

$$[\![x = y]\!] = \{i : x_i = y_i\}, \qquad [\![R(\vec{x})]\!] = \{i : M_i \models R(\vec{x}_i)\}.$$

The equality and substitution axioms hold coordinatewise. Extend the semantics by complements, intersections, and joins in $\mathcal{P}(I)$. An induction on formulas gives

$$[\![\varphi(\vec{x})]\!] = \{i : M_i \models \varphi(\vec{x}_i)\}. \tag{1}$$

Only the quantifier step requires comment. A coordinate i belongs to $\bigcup_{y\in X} [\![\varphi(y, \vec{x})]\!]$ precisely when a witness exists in M_i. To see this from right to left, fill the other coordinates using a chosen family of default elements. To obtain fullness, choose witnesses simultaneously at every coordinate where one exists, and default elements elsewhere. The resulting $y \in X$ attains the Boolean value of the existential formula.

Modulo U_i, equality is exactly equality of the ith coordinate. The map $[x]_{U_i} \mapsto x_i$ is well-defined, injective, and surjective. It preserves functions and relations, so it is an isomorphism to M_i. Equation (1) also gives Boolean value I to each axiom satisfied by every coordinate structure. □

This is a standard Boolean-valued product construction, included here with its proof because it puts the central intuition on an exact footing. Boolean-valued structures and their quotients are treated systematically in [2].

Corollary 3.2. *Every supplied set-indexed family of models of T can be regarded as a family of two-valued resolutions of a richer Boolean-valued structure. This conclusion alone imposes no transformation relation on distinct models in the family.*

Proof. The construction requires only the family of structures. It does not use any maps between them. Consequently, its existence cannot by itself establish additional relationships among the coordinates. □

The theorem supplies a useful warning against overinterpretation. A product can contain unrelated structures. It is not a proof that all worlds share one economical source, that their disagreements are representational artifacts, or that they are forcing extensions of a common ground. The stronger Komorebi question is whether a *constrained* presentation and a *restricted* grammar explain relations that a raw product merely stores.

4 The resolution differential

We first give a fixed-observable version. This is sufficient for sentences in a common language and isolates the algebra from the existence of particular worlds.

Definition 4.1 (Resolution datum)**.** A resolution datum consists of a small category C, a unital Boolean algebra A, a functor

$$B : C \longrightarrow \mathbf{Bool},$$

and a unital Boolean homomorphism $v_M : A \to B_M$ for each object M. Write $B_f : B_M \to B_N$ for the map attached to $f : M \to N$.

The category **Bool** here uses ordinary Boolean homomorphisms. It does *not* require preservation of arbitrary joins. In Boolean-ring notation, + is symmetric difference, multiplication is conjunction, and $\neg a = 1 + a$. In particular, $a + a = 0$.

For sentences, A can be the Lindenbaum Boolean algebra of L-sentences modulo provable equivalence over T. An ordinary model gives $B_M = \mathbb{F}_2$ and $v_M([\varphi]) = 1$ precisely when $M \models \varphi$. The datum can also be studied without claiming that its valuations arise from models of ZFC.

Definition 4.2 (Transport defect)**.** For $f : M \to N$ define

$$D_f(a) = B_f(v_M(a)) + v_N(a) \in B_N. \tag{2}$$

The terms on the right live in the same algebra. A proposed residual without this typing condition is not a mathematical definition.

Theorem 4.3 (Defect calculus)**.** *Put $u = B_f \circ v_M$ and $w = v_N$. Then*

$$D_f(0) = D_f(1) = 0, \qquad D_f(\neg a) = D_f(a), \tag{3}$$

$$D_f(a + b) = D_f(a) + D_f(b), \tag{4}$$

$$D_f(ab) = D_f(a)u(b) + w(a)D_f(b). \tag{5}$$

For composable arrows $M \xrightarrow{f} N \xrightarrow{g} P$,

$$D_{gf}(a) = B_g(D_f(a)) + D_g(a), \qquad D_{\mathrm{id}_M}(a) = 0. \tag{6}$$

Proof. Both u and w preserve 0, 1, addition, and multiplication. The first identities follow immediately, including

$$D_f(1 + a) = (1 + u(a)) + (1 + w(a)) = D_f(a).$$

For the product rule, expand

$$(u(a) + w(a))u(b) + w(a)(u(b) + w(b)) = u(a)u(b) + w(a)w(b).$$

The mixed terms cancel in characteristic two. For composition, functoriality and additivity give

$$\begin{aligned} B_g(D_f(a)) + D_g(a) &= B_g B_f v_M(a) + B_g v_N(a) + B_g v_N(a) + v_P(a) \\ &= B_{gf} v_M(a) + v_P(a). \end{aligned}$$

The identity case follows from $B_{\mathrm{id}_M} = \mathrm{id}$. □

The product rule is a twisted derivation identity between two homomorphisms. It does not make D_f a ring homomorphism or an ordinary derivation relative to one fixed algebra action.

Proposition 4.4 (Product telescope)**.** *For $a_1, \dots, a_n \in A$,*

$$D_f\left(\prod_{i=1}^{n} a_i\right) = \sum_{k=1}^{n} \left(\prod_{i<k} w(a_i)\right) D_f(a_k) \left(\prod_{i>k} u(a_i)\right).$$

Proof. Expand $D_f(a_k) = u(a_k) + w(a_k)$. Every intermediate product occurs twice, once with transition index k and once with transition index $k+1$. They cancel, leaving the all-u product plus the all-w product. This also gives an induction on n, with the empty product case equal to $D_f(1) = 0$. □

4.1 Transported syntax and parameters

For parameters or changing languages, replace A by a functor A_M with translation maps $\alpha_f : A_M \to A_N$. The correct defect is

$$D_f(a) = B_f v_M(a) + v_N \alpha_f(a), \qquad a \in A_M. \tag{7}$$

Assuming both transport systems are functorial, the composition law becomes

$$D_{gf}(a) = B_g D_f(a) + D_g(\alpha_f(a)).$$

The same cancellation proves it. The fixed-A theory is the special case $\alpha_f = \mathrm{id}$. A single aggregate quotient of A will only be used in that special case. This prevents a silent comparison of formulas whose parameters have no specified transport.

5 What the invariant core actually is

For the fixed-A datum set

$$C^0 = \prod_{M \in \mathrm{Ob}(C)} B_M, \qquad C^1 = \prod_{f \in \mathrm{Mor}(C)} B_{\mathrm{target}(f)}.$$

As additive groups these are vector spaces over $\mathbb{F}_2$. Define

$$(d^0 c)_f = B_f(c_M) + c_N, \qquad \mathrm{Ev}(a) = (v_M(a))_M, \qquad D = d^0 \,\mathrm{Ev}\,.$$

Here Ev is a Boolean homomorphism, while d^0 is generally only additive.

Theorem 5.1 (Invariant algebra and residual quotient)**.** *The set*

$$\mathcal{I} = \ker D = \bigcap_f \{a : B_f v_M(a) = v_N(a)\}$$

is a unital Boolean subalgebra of A. There is a natural exact sequence of $\mathbb{F}_2$-vector spaces

$$0 \longrightarrow \mathcal{I} \longrightarrow A \xrightarrow{D} \operatorname{im} D \longrightarrow 0, \tag{8}$$

and therefore a natural additive isomorphism

$$A/\mathcal{I} \;\cong\; \operatorname{im} D.$$

Two observables have the same residual profile exactly when their symmetric difference belongs to $\mathcal{I}$.

Proof. Each defining condition for $\mathcal{I}$ is an equalizer of two Boolean homomorphisms. Equalizers are closed under 0, 1, addition, multiplication, and complement. Additivity of D gives the exact sequence and the first isomorphism theorem for vector spaces. Finally $D(a) = D(b)$ holds exactly when $D(a + b) = 0$. □

The quotient in (8) is not a Boolean-algebra quotient. The subalgebra $\mathcal{I}$ contains 1. Unless $\mathcal{I} = A$, it cannot be an ideal, since an ideal containing 1 is the whole ring. In particular, a nonzero residual profile does not inherit a well-defined conjunction from this quotient.

An immediate information loss is $D(a) = D(\neg a)$. A residual detects change, not its absolute orientation. It cannot replace the original observable without boundary data or invariant information.

Proposition 5.2 (Generating transformations suffice)**.** *If a family Γ generates the arrows of C under composition, then*

$$\mathcal{I} = \bigcap_{f \in \Gamma} \ker D_f.$$

Adding admissible transformations can only shrink the invariant core.

Proof. The composition law propagates vanishing along finite composites. Identity defects vanish. The monotonicity statement is an immediate property of intersections. □

5.1 Coherent sections and definability

The limit $\varprojlim_C B$ consists of families $(b_M)_M$ satisfying $B_f(b_M) = b_N$. Thus $a \in \mathcal{I}$ precisely when $\mathrm{Ev}(a)$ is a coherent family. But not every coherent family need be represented by a sentence, and distinct sentences can have the same evaluations.

Proposition 5.3 (The precise section statement)**.** *Let $K = \bigcap_M \ker v_M$. Then K is a Boolean ideal contained in $\mathcal{I}$, and*

$$\mathcal{I}/K \;\cong\; \operatorname{im}\big(\mathrm{Ev}\,|_{\mathcal{I}}\big) \;\subseteq\; \varprojlim_C B$$

as Boolean algebras.

Proof. The kernel of a Boolean-ring homomorphism is an ideal. An observable whose evaluations are all zero has zero defect, so $K \subseteq \mathcal{I}$. Restrict Ev to $\mathcal{I}$ and apply the Boolean-ring first isomorphism theorem. □

For example, a countable sentence algebra cannot cover all coherent families over a countably infinite *discrete* diagram with $B_M = \mathbb{F}_2$, since that limit is $\mathbb{F}_2^{\mathbb{N}}$ and is uncountable. The correct phrase is therefore *definable coherent observables*, not "all global sections." No sheaf has yet been defined. A sheaf additionally requires a site and a gluing condition, neither of which follows from a limit alone.

6 Finite integrability and reconstruction

Let $G = (W, E)$ be a finite simple undirected graph. Edges record which worlds are directly compared. Choose an orientation only to index them. For $a \in \mathbb{F}_2^W$ set

$$(da)_{\{u,v\}} = a_u + a_v.$$

Unlike the semantic algebra A, the space $\mathbb{F}_2^W$ contains every binary labeling. An actual family of models and a selected language may realize only a proper subalgebra of these labelings.

Theorem 6.1 (Finite residual integrability)**.** *Suppose G has n vertices, m edges, and c connected components.*

1. $\ker d$ *consists exactly of the componentwise constant labelings and has dimension c.*
2. *An edge field $r \in \mathbb{F}_2^E$ equals da for some a if and only if its sum around every cycle is zero.*
3. *Every integrable r has exactly 2^c reconstructions. Fixing one vertex value per component makes the reconstruction unique.*
4. $\dim \operatorname{im} d = n - c$ *and* $\dim(\mathbb{F}_2^E / \operatorname{im} d) = m - n + c$.

Proof. If $da = 0$, endpoints of every edge have equal values, so values agree along each path. Conversely componentwise constants have zero derivative. This proves the first statement.

For any cycle, summing da counts every vertex on the cycle twice, hence gives zero. Conversely choose a spanning forest and a root value in each tree. Define a_v by adding the edge values along the unique tree path from its root. This gives $da = r$ on tree edges. A non-tree edge and the tree path between its endpoints form a cycle. Its zero-sum condition gives the desired equation on that edge as well.

The difference of any two reconstructions lies in $\ker d$. There are 2^c such differences. The rank and quotient-dimension statements now follow from rank-nullity. □

The quotient $\mathbb{F}_2^E / \operatorname{im} d$ is the first graph cohomology group when the graph is regarded as a one-dimensional complex. A nonzero class obstructs integrating *proposed* edge data. This does not make an observed truth differential curved: an observed differential already lies in $\operatorname{im} d$ and has zero class.

6.1 A worked diamond

Take edges $01, 02, 13, 23$ and a labeling

$$a = (0, 1, 0, 1).$$

Then

$$da = (1, 0, 0, 1).$$

The path $0 \to 1 \to 3$ has accumulated defect $1 + 0 = 1$. The path $0 \to 2 \to 3$ has defect $0 + 1 = 1$. The unique cycle has sum $1 + 0 + 1 + 0 = 0$.

With root value $a_0 = 0$ the labeling above is recovered. With root value $a_0 = 1$ the reconstruction is $(1, 0, 1, 0)$. The derivative cannot distinguish these two complementary labelings. That ambiguity is exactly the one-dimensional invariant sector of this connected graph.

By contrast, assigning $r = (1, 0, 0)$ around a triangle is impossible. The cycle sum is 1. No collection of vertex truth values has that derivative.

6.2 Reconstruction with nontrivial transport

For a path $M_0 \xrightarrow{f_1} \cdots \xrightarrow{f_n} M_n$, repeated use of (2) gives

$$v_{M_n}(a) = B_{f_n \cdots f_1} v_{M_0}(a) + \sum_{k=1}^{n} B_{f_n \cdots f_{k+1}} D_{f_k}(a), \tag{9}$$

where the empty tail transport is the identity. This formula needs no inverse transport maps. It reconstructs a reachable target from a source value and the transported residuals on a directed path. It does not claim that every target has such a path or that a general arrow can be reversed.

7 Projectors: existence, obstruction, and a canonical substitute

The formula $a = P(a) + R(a)$ is meaningful only after the type and existence of P have been established.

Proposition 7.1 (Additive splitting)**.** *Over a metatheory with choice, the inclusion $\mathcal{I} \subseteq A$ always admits an $\mathbb{F}_2$-linear retraction $r : A \to \mathcal{I}$. For $P = \iota r$ and $R = \mathrm{id} + P$,*

$$P^2 = P, \quad R^2 = R, \quad PR = RP = 0, \quad a = P(a) + R(a).$$

This is a vector-space decomposition. The retraction need not preserve multiplication or be canonical.

Proof. Choose a vector-space basis of $\mathcal{I}$ and extend it to a basis of A. Let r fix the first basis and kill the added vectors. Then $P^2 = P$. In characteristic two, $(\mathrm{id} + P)^2 = \mathrm{id} + P$, and $P(\mathrm{id} + P) = (\mathrm{id} + P)P = 0$. □

If r is a Boolean retraction, its kernel is an ideal J and the additive decomposition is $A = \mathcal{I} \oplus J$. In these coordinates multiplication is

$$(i, x)(j, y) = (ij,\ iy + jx + xy).$$

It is not generally a direct-product multiplication. Writing a direct sum symbol must not silently erase these cross terms.

Example 7.2 (Completeness of the ambient algebra is not enough)**.** Let $A = \mathcal{P}(\mathbb{N})$ and let $\mathcal{I}$ be the finite-cofinite subalgebra. There is no Boolean retraction $r : A \to \mathcal{I}$ fixing $\mathcal{I}$. Indeed, for every $S \subseteq \mathbb{N}$ and $n \in \mathbb{N}$,

$$r(S) \cap \{n\} = r(S \cap \{n\}) = S \cap \{n\}.$$

Thus $r(S)$ has exactly the same members as S, so $r(S) = S$. An infinite coinfinite S gives a contradiction.

This example can occur as an invariant core of a resolution datum, rather than merely as an arbitrary subalgebra. Take worlds to be the nonprincipal ultrafilters on $\mathbb{N}$, join every ordered pair by an arrow, and use $v_U(S) = 1$ iff $S \in U$. All these worlds agree on finite and cofinite sets. Every infinite coinfinite set belongs to some nonprincipal ultrafilter and has its complement in another, by the ultrafilter extension argument. Hence the invariant core is exactly the finite-cofinite algebra. These worlds are propositional valuations, not asserted models of ZFC.

Completeness of the ambient algebra therefore does not solve the Boolean retraction problem. Neither completeness of the invariant core nor a canonical choice of retraction follows from the definition of a resolution datum. Even when an additional hypothesis supplies a Boolean retraction, canonicity remains a separate question.

7.1 A two-world obstruction

Theorem 7.3 (No symmetric Boolean projector)**.** *Let $A = \mathbb{F}_2^2$ and $\mathcal{I} = \{(0,0),(1,1)\}$. There is no Boolean retraction $P : A \to \mathcal{I}$ that commutes with swapping the two coordinates. There is not even a swap-equivariant additive retraction onto $\mathcal{I}$.*

Proof. Write $e_1 = (1,0)$ and $e_2 = (0,1)$. Swapping coordinates fixes every element of $\mathcal{I}$, so equivariance forces $P(e_1) = P(e_2)$. Additivity then gives

$$P(1,1) = P(e_1 + e_2) = P(e_1) + P(e_2) = 0,$$

contrary to the retraction condition $P(1,1) = (1,1)$. Every Boolean retraction is additive, so both claims follow. □

A demand for a projector determined only by a symmetric two-world structure is therefore inconsistent with these requirements. Adding an ordering, a distinguished representative, weights, or a different output type changes the problem. It does not retroactively make the original Boolean projector canonical.

7.2 The finite resolution envelope theorem

Let Π be a partition of a finite nonempty set W. Let $\mathcal{I}_\Pi \subseteq \mathbb{F}_2^W$ be the Boolean algebra of functions constant on each block. For $w \in C \in \Pi$ define

$$L_\Pi(a)_w = \bigwedge_{u \in C} a_u, \qquad U_\Pi(a)_w = \bigvee_{u \in C} a_u.$$

These lower and upper approximations are the classical equivalence-relation rough-set operators [3]. Here they give a precise replacement for a missing canonical projector.

Theorem 7.4 (Finite resolution envelope)**.** *Every Boolean retraction $P : \mathbb{F}_2^W \to \mathcal{I}_\Pi$ is obtained by choosing one representative $s(C) \in C$ per block and setting*

$$P_s(a)_w = a_{s(C)} \quad (w \in C). \tag{10}$$

For each fixed observable a, its possible projected values are exactly

$$\{P_s(a) : s \text{ chooses one point in each block}\} = \{b \in \mathcal{I}_\Pi : L_\Pi(a) \le b \le U_\Pi(a)\}. \tag{11}$$

Consequently,

$$L_\Pi(a) = \bigwedge_s P_s(a), \qquad U_\Pi(a) = \bigvee_s P_s(a).$$

If a is mixed on exactly k blocks, there are exactly 2^k possible invariant outputs for that a.

Proof. A unital Boolean homomorphism $h : \mathbb{F}_2^W \to \mathbb{F}_2$ must evaluate at exactly one point. To see this, the singleton indicators are pairwise disjoint and their join is 1. Their images are disjoint elements of $\mathbb{F}_2$ whose join is 1, so exactly one is 1. The image of every subset is then determined by whether it contains that point.

For a retraction P, evaluation of its output on a block C is such a homomorphism, hence selects a point $s(C) \in W$. Since P fixes the indicator of C, this point must belong to C. This proves the classification. Conversely every choice in (10) gives a Boolean homomorphism that fixes blockwise constants.

For fixed a, a selected value on C lies between its minimum and maximum there. Conversely, take $b \in \mathcal{I}_\Pi$ in that interval. When a is constant on C, b has that same value. When a is mixed, choose a point of C at which a has the value prescribed by b. These independent choices produce $P_s(a) = b$. Each mixed block allows precisely two outputs and every other block allows only one. □

The extrema are canonical even though an individual retraction need not be. They preserve order and satisfy

$$L_\Pi(a) \le a \le U_\Pi(a), \qquad L_\Pi(\neg a) = \neg U_\Pi(a).$$

The lower map preserves meets and the upper map preserves joins. Neither generally preserves both Boolean operations. Their role is a certified interval of invariant summaries, not a new binary truth valuation.

Proposition 7.5 (Adjoint characterization)**.** *For $b \in \mathcal{I}_\Pi$,*

$$U_\Pi(a) \le b \iff a \le b, \qquad b \le L_\Pi(a) \iff b \le a.$$

Thus U_Π is left adjoint to the inclusion and L_Π is right adjoint to it, in the order-theoretic sense.

Proof. $U_\Pi(a)$ is exactly the smallest union of blocks containing the support of a. Dually $L_\Pi(a)$ is the largest union of blocks contained in that support. □

Theorem 7.6 (Residual support determines ambiguity)**.** *Take Π to be the connected components of a finite graph G. Define*

$$\mathcal{A}_\Pi(a) = U_\Pi(a) + L_\Pi(a).$$

Then $\mathcal{A}_\Pi(a)$ is the indicator of the union of components containing a nonzero edge of da:

$$\operatorname{supp} \mathcal{A}_\Pi(a) = \operatorname{Sat}_\Pi(\operatorname{supp} da).$$

Here edge support is saturated by taking all vertices in components containing a supported edge.

Proof. A component is ambiguous exactly when it contains both truth values. If it contains an edge with unequal endpoint values, it is mixed. Conversely a path from a 0 vertex to a 1 vertex must contain an edge where the value changes. This proves both implications. □

The three objects now have distinct roles: da locates transitions, $(L_\Pi(a), U_\Pi(a))$ describes all invariant Boolean summaries, and a chosen base value or selector supplies reconstruction data. None of the three should be silently substituted for the others.

8 A real-valued realization of Komorebi multiresolution

The Boolean obstruction does not prohibit a canonical projector after extra structure and a different codomain are supplied. Let W be finite and assign positive weights μ_w. On $\mathbb{R}^W$ use

$$\langle f, g \rangle_\mu = \sum_{w \in W} \mu_w f_w g_w.$$

For a partition Π, define

$$(E_\Pi f)_w = \frac{\sum_{u \in C} \mu_u f_u}{\sum_{u \in C} \mu_u} \quad (w \in C \in \Pi).$$

Theorem 8.1 (Weighted invariant projection)**.** *E_Π is the unique orthogonal projector onto the real vector space of blockwise constant functions. It is positive, fixes $\mathbb{1}$, and is equivariant under weight-preserving permutations that preserve the partition. For binary a,*

$$L_\Pi(a) \le E_\Pi(a) \le U_\Pi(a).$$

Proof. The formula is blockwise constant and fixes such functions. For every blockwise constant g,

$$\langle f - E_\Pi f, g\rangle_\mu = \sum_{C\in\Pi} g_C \sum_{w\in C} \mu_w (f_w - (E_\Pi f)_w) = 0.$$

This characterizes the orthogonal projector uniquely. Positivity, equivariance, and the interval bounds follow directly from a positive weighted average. □

This average is not a Boolean truth assignment. With equal weights on two points, $(0, 1)$ projects to $(1/2, 1/2)$. Also, for the disjoint singleton indicators e_1, e_2, one has $E(e_1e_2) = 0$ while $(Ee_1)(Ee_2)$ is the constant $1/4$. Conjunction is not preserved.

Theorem 8.2 (Nested resolution decomposition)**.** *Let $\Pi_0, \dots, \Pi_m$ be successively finer partitions, using the same weights. Write $E_k = E_{\Pi_k}$ and $Q_k = E_k - E_{k-1}$ for $1 \le k \le m$. Then*

$$E_iE_j = E_{\min(i,j)}, \qquad Q_k^2 = Q_k, \qquad Q_kQ_l = 0 \ (k \ne l),$$

and every $f \in \mathbb{R}^W$ has the orthogonal decomposition

$$f = E_0 f + \sum_{k=1}^{m} Q_k f + (\mathrm{id} - E_m) f. \tag{12}$$

In particular the squared norm is the sum of the squared norms of these terms.

Proof. The blockwise constant subspaces are nested. Orthogonal projection onto nested subspaces gives $E_iE_j = E_i$ for $i \le j$, and the reverse product agrees by self-adjointness. Expanding products proves the projector and orthogonality identities for the Q_k. Equation (12) telescopes. The norm identity is the Pythagorean theorem for these orthogonal summands. □

This is a genuine projector/detail calculus. Its hypotheses make the Komorebi analogy exact, but the result is a standard nested-projection construction, not a new theorem about ZFC. The weights and partition sequence are part of the input. Their existence does not establish a preferred measure or scale on a set-theoretic multiverse.

For the original Boolean bounds, finer partitions instead give monotone information:

$$L_{\Pi_0}(a) \le L_{\Pi_1}(a) \le \cdots \le a \le \cdots \le U_{\Pi_1}(a) \le U_{\Pi_0}(a).$$

The ambiguity support shrinks as blocks are refined. This is an order-theoretic resolution process, not an additive Boolean detail expansion.

9 Stabilization requires a path quantifier

For a path $M_0 \to M_1 \to \cdots$ with two-valued observations, write $a_n = v_{M_n}(a)$. Along that path, eventual stability means

$$\exists N\, \forall n \ge N \quad a_n = a_N,$$

equivalently $a_n + a_{n+1} = 0$ eventually. This differs from requiring stability on every continuation from some reachable world.

On a reflexive accessibility relation, $\Diamond\Box\varphi$ says that a world can be reached from which every later accessible world satisfies φ. A path that happens to stay at a φ-world does not establish this. A two-state switch system with arrows in both directions and identity loops already separates the notions.

Theorem 9.1 (Finite all-path stability criterion)**.** *Let G be a finite directed graph with a binary vertex labeling. Every infinite directed path is eventually constant in label if and only if every strongly connected component is monochromatic. When this holds, every path has at most $k-1$ label changes, where k is the number of strongly connected components.*

Proof. If one strongly connected component contains both labels, travel from a 0 vertex to a 1 vertex and back. Repeating the resulting closed walk gives an infinite path with infinitely many changes.

Conversely, a path can leave a strongly connected component only by moving forward in the finite acyclic condensation graph. It can never return to a component it has left. There are at most $k-1$ such transitions. Within each component the label is constant. Hence an infinite path eventually has constant label and has at most the asserted number of changes. □

There need not be a uniform *time index* for the last change. A path may remain at a self-looped 0 vertex for arbitrarily many steps before moving once to an absorbing 1 vertex. Thus a bound on the number of residual events is not a bound on when the last residual occurs.

A transformation grammar also needs an explicit cost before a "resolution depth" is meaningful. Subdividing one arrow into many arrows changes step counts without changing the underlying semantics. Minimum cost to change truth, stabilization time on one path, maximum number of changes, and robust modal stabilization are different invariants.

10 Forcing as a conditional set-theoretic instance

Fix externally a countable transitive $M \models \mathsf{ZFC}$. Let $\mathbb{B} \in M$ be a Boolean algebra that M regards as complete, and let G be an M-generic ultrafilter. For a sentence φ with names $\vec{\tau}$, the standard forcing truth lemma gives

$$M[G] \models \varphi(\vec{\tau}_G) \quad \Longleftrightarrow \quad [\![\varphi(\vec{\tau})]\!]^M_{\mathbb{B}} \in G. \tag{13}$$

Boolean-valued forcing and the role of generic quotient models are developed in [2]. The completeness assumption here is internal to M. The map $b \mapsto \mathbf{1}_{b \in G}$ is a Boolean homomorphism but need not preserve every externally given infinite join.

Equation (13) is a precise local instance of "richer data resolved into ordinary truth." It is not permission to replace G by an arbitrary ultrafilter and still claim an ordinary generic extension of the same ground model. Nor is the atomic product presentation of Section 3 the same construction as a forcing extension.

10.1 A changing observable and an invariant fragment

Taking two-valued fibers and an arrow $M \to M[G]$ gives

$$D_f(\varphi) = \mathbf{1}_{M \models \varphi} + \mathbf{1}_{M[G] \models \varphi}.$$

For suitable forcing, CH can change, so it is not invariant under unrestricted forcing. Both CH and its negation are forceable in the relevant forcing setting [6]. This is not a contradiction: the quantifiers of CH are interpreted in worlds with different available sets.

There is also a direct invariant fragment that requires no new theorem.

Proposition 10.1 (Bounded-formula invariance)**.** *If $M \subseteq N$ are transitive membership structures and all parameters belong to M, they agree on every bounded formula, namely a formula whose quantifiers are all of the form $\exists x \in y$ or $\forall x \in y$.*

Proof. Induct on the formula. Atomic membership and equality are the same actual relations in both structures. Boolean connectives preserve agreement. A bounded quantifier over $y \in M$ ranges over the same elements in both structures, because transitivity puts every member of y in M. □

This gives zero defect for the specified formulas and parameter transports. Stronger generic-absoluteness results require their own hypotheses and semantic fragments. They are not consequences of merely naming an invariant core.

10.2 Modal structure is established prior mathematics

The modal interpretation of forcing uses $\Diamond\varphi$ for forceability and $\Box\varphi$ for truth in every forcing extension. Assuming consistency, the ZFC-provably valid modal principles for unrestricted forcing are exactly S4.2 [1]. Restricted forcing classes have different structural behavior [4]. KRS does not reprove these results. It records individual transport defects and their reconstruction alongside the existing modal language.

10.3 Why universal gluing is not an admissible default

Specified generic extensions need not have a common further extension in the same generic multiverse. Nonamalgamation is a genuine structural phenomenon, not missing notation. The generic multiverse can encode finite patterns of absent upper bounds [5].

Thus the axiom "every pair of worlds has a common refinement" is not automatically valid for forcing worlds. Modal reachability of sentences is also weaker than amalgamating particular generic objects already supplied. A correct resolution atlas must record the domain on which amalgamation is available, rather than silently declaring the entire diagram directed.

11 Where nontrivial transport geometry can enter

For constant $\mathbb{F}_2$ fibers, a truth differential is a coboundary. On every closed comparison cycle its sum is zero. This is an exact identity, not evidence that every richer world system is geometrically trivial.

To retain information about structures rather than just truth values, give each world a fiber X_M and each arrow an actual transport map $T_f : X_M \to X_N$. Parallel paths $p, q : M \to N$ can then be compared only because their source and target types agree. For linear fibers define the path discrepancy

$$\Delta_{p,q} = T_p - T_q.$$

If the paths are declared to be the *same morphism* in the source category and T is a functor, this discrepancy must vanish. A nonzero value means that the paths are distinct morphisms, or that a claimed relation is not respected. It cannot be described as curvature of an already commuting strict functor square.

Example 11.1 (Nontrivial holonomy with an exact defect)**.** Let the one-object category be the group with generator s and relation $s^2 = 1$. Put $X = \mathbb{F}_2^2$ and let T_s swap coordinates. This is a valid nontrivial action. Its invariant vectors are $(0, 0)$ and $(1, 1)$. For $x = (1, 0)$,

$$T_s x + x = (1, 1) \neq 0,$$

but

$$T_s(T_s x + x) + (T_s x + x) = 0.$$

The second equation is the transported composition law for $s^2 = 1$. Thus nontrivial fiber transport is compatible with an exact defect. No contradiction or unexplained curvature is involved.

This suggests a useful stratification. Binary observations carry change profiles. Rich fibers carry witnesses, parameters, embeddings, and transport actions. A site or descent datum, when actually specified, carries gluing obligations. These layers should interact through explicit maps rather than by identifying them.

12 A strengthened research program, with explicit obligations

The defensible version of the central idea is now:

> A chosen family of mathematical worlds can be equipped with typed transformations and richer presentations. An invariant/residual calculus then classifies the observable effects of those transformations. Whether the presentation is explanatory, canonical, economical, or generative is an additional mathematical question.

A useful *resolution atlas* should specify the following data in one contract: a class of admissible worlds, local presentation objects, evaluation maps to their resolved worlds, a grammar of transformations, comparison maps on observables and witnesses, and the exact equations certified for those maps. It must also distinguish a failed comparison from a comparison that is undefined.

The next substantive target is not another name for the multiverse. It is a theorem for a specific nontrivial sector stating that a controlled set of generators determines the transport behavior of a rich observable family, with exact reconstruction and a recognizable obstruction to extension. The finite results here show the shape such a theorem could have.

12.1 Three concrete problems

Presentation sensitivity. Identify conditions under which equivalent presentations of the same typed resolution system induce the same invariant quotient and canonical bounds. For a finite partition system this follows under a bijection preserving blocks. The nontrivial case is a change of presentation that is not merely a relabeling.

Restricted forcing reconstruction. Choose a particular forcing class and a precise language fragment. Determine whether a structured family of residual certificates classifies transport on that fragment more economically than a list of all truth changes. This requires real preservation and nonamalgamation theorems, not just the abstract cocycle law.

Witness-sensitive invariants. Replace truth profiles by parameter and witness transport. Find an invariant that distinguishes diagrams with identical sentence-level defects. Example 11.1 shows why richer fibers can contain information that binary truth alone loses.

12.2 What would count as a new foundational axiom

The definitions and conditional algebra above do not assert a new set-theoretic existence principle. A new axiom would arise only when one requires that certain worlds, generics, coherent witnesses, or canonical structures exist. Its consistency strength and interaction with known obstructions must then be analyzed independently.

In particular, neither universal amalgamation nor a universal symmetry-preserving Boolean projector is available as an unrestricted default. The first conflicts with the forcing phenomena just cited. The second already fails in Theorem 7.3. A productive extension must respect these obstructions rather than declare them away.

13 Verification and proof boundary

The accompanying checker uses only Python's standard library and exact arithmetic. It does not use floating point, random sampling, or a satisfiability oracle. The precise run records and source hash are in `results/VERIFICATION.json`.

Check	Count	Range or meaning
Boolean polynomial identities	12	Exact normal forms
Monomial multiplication evaluations	4,096	Four variables
Simple undirected graphs	1,099	One to five vertices
Vertex truth assignments	33,866	Every graph above
Edge residual assignments	59,809	Every graph above
Partitions	75	One to five points
Observable inputs for partitions	1,954	All binary labelings
Selector evaluations	7,022	All block representatives
Independent coordinate maps	3,984	Up to four points
Directed labeled systems	4,164	Up to three vertices

The graph tests compare direct enumeration of all gradients with an independent propagation-based integrator and with Gaussian elimination over $\mathbb{F}_2$. They verify reconstruction counts, kernel dimension, and ambiguity support. The partition tests compare every selector output with the interval in Theorem 7.4, and independently enumerate coordinate homomorphisms through four points. The two-world test enumerates all functions $\mathbb{F}_2^2 \to \mathbb{F}_2$ and finds exactly two Boolean retractions, neither swap-equivariant.

13.1 Why the symbolic checks support universal identities

A Boolean polynomial is represented as a finite sum of squarefree monomials over $\mathbb{F}_2$. Addition cancels duplicate monomials. Multiplication unions their variable supports, using $x_i^2 = x_i$. These rules are valid in every Boolean ring. A polynomial identity whose difference reduces to the empty normal form therefore has an algebraic certificate applicable to every Boolean-ring substitution.

This statement uses the displayed algebraic reduction rules and their mathematical justification. The Python checker is not a replacement for a trusted proof-assistant kernel. The arbitrary-length product telescope is proved in the note by cancellation or induction, rather than inferred from the six product lengths tested by code.

The graph, partition, and stability computations remain bounded checks. Their universal conclusions come from the proofs in Sections 6, 7, and 9. No finite graph in the test suite is claimed to be a verified model of ZFC or a realized forcing diagram.

13.2 Lean status and the actual failure

No Lean compilation was completed for this release. No Lean or Lake executable was found in the inspected runtime paths. A pinned Lean 4.19.0 download failed because the runtime could not resolve `github.com`. The external download route also did not provide the release asset. The failure was environmental, not a failed mathematical theorem check.

The bundle contains a small `formal/KRS_Core.lean` scaffold for two-valued identities and profile invariance. It has no intended placeholders, but it is explicitly *uncompiled*. It is not counted as verified proof evidence. No earlier Komorebi Lean checkpoint is imported, modified, or represented as verifying this note.

13.3 Claim ledger

Claim	Evidence in this release	Boundary
Common Boolean presentation	Theorem 3.1, full proof	Supplied set-indexed models and choice
Defect calculus	Theorem 4.3, symbolic checks	Boolean homomorphism hypotheses
Invariant algebra and quotient	Theorem 5.1	Quotient is additive
Coherent-observable description	Proposition 5.3	Only definable coherent sections
Finite integration	Theorem 6.1, exhaustive tests	Arbitrary finite graphs, not forcing realizability
Finite projector classification	Theorem 7.4, exhaustive tests	Full finite powerset algebra
Canonical Boolean projector	Refuted in Theorem 7.3	Extra structure or a different output type is needed
Canonical invariant interval	Theorem 7.4	Rough-set bounds, not a truth valuation
Weighted multiresolution	Theorems 8.1–8.2	Real fibers with fixed positive weights
All-path stability	Theorem 9.1, finite checks	Finite directed systems
Forcing interpretation	Conditional use of established forcing results	No new forcing theorem is claimed
Lean verification	Not run	Compiler unavailable and download failed
New axiom or consistency result	Not established	Separate foundational research obligation

14 Conclusion

The original intuition survives, but in a more discriminating form. Worlds can genuinely be presented as resolutions of richer Boolean-valued data. That fact is too general, by itself, to explain why the worlds are related. The mathematical content lies in the transformation grammar, typed comparison maps, preserved observables, and reconstruction constraints.

The resulting calculus has a proved additive core. Its finite realization has an exact integration theorem, a classification of all Boolean invariant projections, and canonical envelopes when symmetry forbids a preferred projection. After passing to weighted real observables, it has a genuine orthogonal multiresolution decomposition. These are compatible realizations with different contracts, not one flattened notion of truth.

The most useful central equation is therefore conditional rather than rhetorical:

$$\boxed{\text{transported source data} + \text{certified residual} = \text{target data.}}$$

When an invariant projector exists, it gives a stronger decomposition. When it does not, the theory retains canonical bounds and records the obstruction. This is the foundation of a workable resolution semantics, and a precise starting point for testing a broader Komorebi calculus of mathematical worlds.

A Reproduction and artifact contents

From the extracted release directory, run

```
python checks/verify_krs.py
pdflatex -interaction=nonstopmode -halt-on-error \
  Komorebi_Resolution_Semantics.tex
pdflatex -interaction=nonstopmode -halt-on-error \
  Komorebi_Resolution_Semantics.tex
```

The checker recreates its JSON evidence. A local Lean installation can separately be used to attempt

```
lean formal/KRS_Core.lean
```

That last command is a reproduction target, not a command that succeeded in this release. The source scaffold only covers a limited two-valued core, not the full categorical, graph, or forcing development.

The package includes the PDF, its LaTeX source, executable checks, results, a theorem ledger, a failure log, and checksums. Compiler binaries and font files are not distributed. This is the first versioned note produced from the discussion, not an update to an unseen earlier archive.

B Notation and distinctions at a glance

Symbol	Meaning
$M \models T$	A supplied model, not an assertion that the ambient theory proves such a model exists.
C	A small category of typed world transformations.
A	A Boolean algebra of observables, often sentences modulo provable equivalence.
B_M	The Boolean algebra of values at one world.
v_M	Semantic evaluation of observables.
D_f	Additive defect between transported source evaluation and target evaluation.
$\mathcal{I}$	Zero-defect observables. A subalgebra, generally not an ideal.
K	Observables vanishing at every world. An ideal.
$A/\mathcal{I}$	Additive residual quotient. Not a Boolean quotient in general.
$\mathcal{I}/K$	Definable coherent evaluations, a Boolean algebra.
L_Π, U_Π	Canonical lower and upper invariant approximations.
P_s	Boolean retraction determined by a representative in each finite block.
E_Π	Weighted real orthogonal projection. Not a Boolean homomorphism.
Q_k	Real multiresolution detail projector for nested partitions.
$\Delta_{p,q}$	Typed discrepancy between parallel transport paths.

References

[1] J. D. Hamkins and B. Löwe. *The modal logic of forcing*. Transactions of the American Mathematical Society 360 (2008), 1793–1817. arXiv:math/0509616.

[2] J. D. Hamkins and D. E. Seabold. *Well-founded Boolean ultrapowers as large cardinal embeddings*. 2012. arXiv:1206.6075. Sections 1–3 provide the Boolean-valued model and quotient background.

[3] Z. Pawlak. *Rough sets*. International Journal of Computer & Information Sciences 11 (1982), 341–356. doi:10.1007/BF01001956.

[4] J. D. Hamkins, G. Leibman, and B. Löwe. *Structural connections between a forcing class and its modal logic*. 2012 preprint. arXiv:1207.5841.

[5] M. E. Habič, J. D. Hamkins, L. D. Klausner, J. Verner, and K. J. Williams. *Set-Theoretic Blockchains*. Archive for Mathematical Logic 58 (2019), 965–997. arXiv:1808.01509.

[6] D. Asperó and M. Viale. *Incompatible bounded category forcing axioms*. 2021 preprint, version 1. arXiv:2101.03132. The introduction explicitly distinguishes generic absoluteness from unrestricted CH-changing forcing.